



# Complex-valued neural networks for voice anti-spoofing

Nicolas M. Müller, Philip Sperl, Konstantin Böttinger

Fraunhofer AISEC

nicolas.mueller@aisec.fraunhofer.de

## Abstract

Current anti-spoofing and audio deepfake detection systems use either magnitude spectrogram-based features (such as CQT or Melspectrograms) or raw audio processed through convolution or sinc-layers. Both methods have drawbacks: magnitude spectrograms discard phase information, which affects audio naturalness, and raw-feature-based models cannot use traditional explainable AI methods. This paper proposes a new approach that combines the benefits of both methods by using complex-valued neural networks to process the complex-valued, CQT frequency-domain representation of the input audio. This method retains phase information and allows for explainable AI methods. Results show that this approach outperforms previous methods on the "In-the-Wild" anti-spoofing dataset and enables interpretation of the results through explainable AI. Ablation studies confirm that the model has learned to use phase information to detect voice spoofing.

**Index Terms:** voice anti-spoofing, audio deepfake detection, complex neural network

## 1. Introduction

Artificial Intelligence (AI) is advancing at a rapid pace, and has enabled the development of numerous good and helpful applications that are changing the way we live, work, and interact with each other. However, along with these positive advancements, new threats have emerged, particularly in the form of deepfakes and spoofs. Deepfakes are computer-generated images, videos, or audio that are created using AI algorithms to manipulate and alter original content. Although deepfakes can be used for harmless fun, they can also be misused for malicious purposes such as creating fake news, information campaigns, slander, or fraud. This is concerning as deepfakes can be highly convincing, making it challenging to distinguish between genuine and fake content. Spoofing, on the other hand, involves creating a fake or manipulated version of oneself to circumvent biometric identification systems such as facial recognition or fingerprint scanners. This has significant implications, especially in the realm of security, as it can compromise the integrity of identity verification systems. As a result, the detection of spoofed or faked content is becoming increasingly important. It is crucial to develop effective methods for detecting such misuse of AI in order to ensure that these technologies are used ethically and responsibly.

Using Machine Learning (ML) for voice anti-spoofing requires adequate preprocessing because human speech has long-range temporal dependencies, with one second of audio usually represented by 16,000 single data points or more. Previous approaches to voice anti-spoofing have taken one of two approaches to audio preprocessing. Some have used spectral

preprocessing, transforming the time-domain waveform into a magnitude spectrogram using the short-time Fourier transform (STFT) or similar techniques. This approach has shown good performance but has one key disadvantage: it discards phase information from the spectrogram. This phase information is crucial for audio quality and voice naturalness [1], and in the domain of speech-to-text (STT), it is laboriously recreated using Griffin-Lim or neural vocoders [1, 2, 3]. As bad phase is very audible<sup>1</sup>, we argue that it is a useful feature that should not be discarded.

The second approach to voice anti-spoofing is to design models that process raw audio directly, which have been shown to be more effective [4, 5, 6], but lack transparency because explainable AI methods (XAI) such as saliency maps require input data with spatial dimensions (i.e. at least two-dimensional input).

In this paper, we propose a new approach to voice anti-spoofing that combines the benefits of both previous approaches. We transform input audio into a complex-valued spectrogram using the constant-Q transform, a technique closely related to the STFT, which yields a complex-valued frequency representation. This representation is mathematically equivalent to the time-domain representation of the input audio, but more suited to machine-learning algorithms. We then process this frequency representation using a complex-valued convolutional neural network, which allows us to process all information present in the audio file, without the need to discard the phase. We demonstrate that our proposed approach outperforms both magnitude spectrogram-based models and raw-feature models. Our approach is conceptually and architecturally simple, which sets it apart from some of the recent state-of-the-art raw models. Moreover, it enables the use of explainable AI techniques.

In summary, our contributions to the field of voice anti-spoofing can be outlined as follows:

- We emphasize the significance of phase information in the complex Short-Time Fourier Transform (STFT) output, which has been neglected in previous research.
- Based on this insight, we propose a corresponding complex-valued input feature and neural architecture for jointly processing both magnitude and phase information.
- We evaluate our approach and demonstrate its superiority over both magnitude spectrogram-based models and raw-based models.

We provide an online web interface<sup>2</sup> where one can test arbitrary input, including YouTube videos, against our model.

<sup>1</sup>Audio examples available at [google.github.io/phase-prediction](https://github.com/google.github.io/phase-prediction).

<sup>2</sup><https://deepfake-total.com>

## 2. Previous Work

As previously mentioned, voice anti-spoofing traditionally employs one of two feature preprocessing approaches. The first approach involves using magnitude-spectral features like the constant-Q transform (CQT) [7], log or mel-scaled spectrograms [8]. These techniques use the complex-valued STFT of a time-domain audio signal as a foundation, extracting the magnitude and discarding phase information. The result is then squared, scaled, and possibly binned using the mel-scale [8] to mimic the human ear’s perception. Previously suggested voice anti-spoofing architectures have then employed neural components such as convolutions, attention mechanisms, recurrence, or transformer blocks [9, 10, 11, 12, 13] to process the resulting spectrogram. These models have shown good performance, but no longer achieve state of the art.

The second approach involves directly processing the time-domain “raw” audio using either neural convolutions on the audio file itself [14] or stacks of sinc-layers [15] corresponding to rectangular band-pass filters. Max-pooling is often used [5] to avoid phase mismatch between the audio waveform and the sinc wavelets. The result is then processed using stacks of convolutions and gated recurrence (RawNet2) [5], more advanced architectures like spectro-temporal graph attention neural networks (RawGat-ST) [4], or differentiable architecture search [6].

Although these “raw” models have been shown to deliver outstanding performance [4, 5, 6], they lack explainability. This is because existing XAI techniques such as saliency maps [16] or Smooth Grad [17] are designed for spatially-dimensioned input, while “raw” audio is a one-dimensional vector. As a result, it becomes challenging to comprehend and trust the models’ judgments. Among other disadvantages, this impedes the detection of “learning shortcuts” [18], i.e. highly predictive, but artificial correlations between input and target which do not reflect real-world causality. Detecting such shortcuts is particularly crucial in voice spoof detection, where they have already been identified [19] and may result in a lack of generalizability to real-world scenarios [20]. Thus, there is an urgent demand for explainable anti-spoofing models.

## 3. Proposed Approach

### 3.1. Complex-valued CQT spectrograms

In digital signal processing, an input audio signal  $X = [X[0], X[1], \dots, X[n]]$  is represented in the time domain as a series of scalar values  $X[t] \in [-1, 1] \in \mathbb{R}$ , which correspond to measurements of air pressure at time  $t/s$ , where  $s$  is the sampling rate (e.g., 16 kHz or 16,000 samples per second). To convert this signal into a frequency representation  $Z$ , one may use the Short Time Fourier Transform, which is a bijective transform:

$$Z[t, k] = \sum_{n=0}^{N-1} W[n]X[n+t]e^{-2\pi i k n N^{-1}} \quad (1)$$

Here,  $X$  is the time-domain signal, and  $W$  is a window such as the Gaussian, Hann, or Hamming window of size  $N$ . The frequency-domain representation  $Z$  is a complex-valued, two-dimensional vector, where  $t$  is the temporal and  $k$  the frequency index. Elements in  $Z$  are complex-valued scalars, i.e.  $Z[t, k] = \alpha e^{i\theta} \in \mathbb{C}$ , where  $\alpha \in \mathbb{R}$  is the magnitude and  $\theta \in [0, 2\pi[$  the phase.

Although  $Z$  is better suited for neural processing than the corresponding time-domain representation  $X$ , it also has its

own drawbacks: the uniform time and frequency resolution for all frequency bins  $k$  does not align with the non-linearities of the human auditory system [21, 22]. To address this issue, we employ the constant Q transform (CQT) [7, 23] which processes the same number of cycles  $Q$  for each frequency and spaces the frequency bins  $k$  geometrically (i.e., higher frequencies are farther apart). Therefore, for high frequencies  $k$ , CQT employs smaller window sizes  $N_k$  (and thus higher time resolution), while for lower frequencies, it employs larger window sizes and thus larger frequency resolution. Thus, instead of using the STFT, we convert time-domain audio  $X$  into a complex-valued CQT spectrogram:

$$Z_Q[t, k] = \frac{1}{N_k} \sum_{n=0}^{N_k-1} W_k[n]X[n+t]e^{-2\pi i Q n N_k^{-1}} \quad (2)$$

Previous work has already shown that magnitude-based CQT is better suited to voice antispoofing than magnitude-based STFT [24], which motivates our approach.

Finally, we convert the complex-valued CQT-spectrogram into a log representation by means of the following transformation:

$$\mathbb{C} \rightarrow \mathbb{C} \quad (3)$$

$$|z|e^{i\theta} \mapsto \max(\epsilon, -\log_e |z| + c) \cdot \alpha \cdot e^{i\theta} \quad (4)$$

To account for the logarithmic perception of sound level by the human ear, we log-scale the magnitude of the complex number while keeping the phase intact. However,  $\log_e |z|$  can be negative, because  $\log_e |z|e^{i\theta} = -\log_e |z|e^{i\theta+i\pi}$ , leading to ambiguity. To address this, we apply a lower threshold of  $\epsilon > 0$  to ensure that the magnitude remains positive. Additionally, to ease model training, we introduce trainable parameters  $\alpha, c \in \mathbb{R}^{>0}$  to scale the magnitude. In our experiments, these converge to approximately  $\alpha = 0.15$  and  $c = -0.3$ . To summarize, we utilize CQT followed by eq. (4) to generate complex-valued CQT log-spectrograms (C-CQT).

### 3.2. Complex-valued neural networks

Complex-valued neural networks (CVNN) are networks where both inputs  $X$  and weights  $\theta$  are complex numbers  $a + ib \in \mathbb{C}$ . Apart from that, CVNNs rely on many of the same fundamental components as traditional real-valued ones. For instance, linear layers, convolutions, and simple recurrent networks like RNN operate identically in both types of networks. However, more complex components such as attention and recurrence, such as GRU [25], need to be adapted, since typical activation functions like *sigmoid*, *tanh*, and *softmax* cannot be applied directly to complex-valued scalars [26]. For example, the complex *ReLU* function can be implemented as

$$CReLU(x) = \max(0, \Re(x)) + i\max(0, \Im(x)) \quad (5)$$

where  $\Re$  and  $\Im$  denote the real and imaginary parts of a complex number. In supervised classification, class scores  $y$  for  $n$  classes can be derived from complex-valued logits  $z$  as follows:

$$\text{softmax}(|z|)_i = \frac{e^{|z_i|}}{\sum_{j=1}^n e^{|z_j|}} \quad (6)$$

Classes with high probability are represented by logits with a large magnitude. These models, like their real counterparts, are trained using gradient descent [26] through the use of Wirtinger Calculus [27].




Figure 1: *Our proposed complex-valued architecture.*

### 3.3. Proposed Architecture

In order to process complex-valued, frequency-domain audio input, we use a complex convolutional network. It consists of four blocks of convolutions, where each block is composed of two-dimensional complex convolutions (kernel size 3, stride 2 and padding 1), complex ReLU and complex Batch-Normalisation. This is followed by linear projection, which consists of three complex-valued linear layers with CReLU activation and Dropout of 40%. This is followed by two-dimensional complex time pooling. The output is a two-class complex vector of shape  $(B, 2)$ , to which we apply eq. (6) in order to obtain real-valued class scores. There are two target classes: class 0, which corresponds to bonafide audio, and class 1, which corresponds to spoofed audio. The input frequency-domain audio has shape  $B, F, T$  where  $B$  is the batch-size,  $F$  the number of frequency bins, and  $T$  the number of time bins. The convolutional stack outputs a tensor  $B, C, L$  where  $C$  is the number of output channels and  $L$  is the aggregated time. Figure 1 provides an overview.

## 4. Evaluation

### 4.1. Datasets

We are utilizing two datasets for our experiments. The first dataset is the Logical Access (LA) section of the ASVspoof 2019 dataset [28]. This dataset is widely used in related research [4, 5, 6, 9, 10, 11, 12, 13] and is considered one of the most established anti-spoofing datasets. It includes audio files that are either genuine human speech recordings (bona-fide) or manipulated/synthesized audio (spoofed). The spoofed audio files are generated using 19 different Text-To-Speech (TTS) synthesis algorithms that are considered a threat to the authenticity of human voice in terms of spoofing detection. Therefore, the spoofed audio files are labeled as “attacks” by ASVspoof. The dataset consists of 19 attackers, labeled as A1 to A19, and for each attacker, there are 4914 synthetic audio recordings and 7355 genuine samples.

Furthermore, we utilize the “In-the-Wild” (ITW)

dataset [20], which comprises 37.9 hours of audio recordings, either spoofed (17.2 hours) or authentic (20.7 hours), obtained from 58 English-speaking celebrities and politicians via online video sharing platforms, such as YouTube. It tries to match the speaking style for each pair of spoofed and authentic instance to prevent the model from distinguishing spoofs based on extraneous characteristics (also called “machine learning shortcuts”), such as speaking style, background noise, or setting. For instance, for a spoofed speech by Barack Obama, an authentic speech was also included. Our objective is for the model to distinguish spoofed audio based on either disfluencies that indicate poor naturalness or artifacts resulting from the text-to-speech (TTS) process.

Since our primary concern is the real-world performance of the models on unseen audio recordings, we train them on the entire ASVspoof dataset (all splits) and assess their generalization abilities using the “In-the-Wild” dataset.

### 4.2. Data-augmentation and Hyperparameters

Previous research has identified a significant learning shortcut in the ASVspoof 2019 dataset, whereby the duration of silence in an audio file is highly correlated with its corresponding label. This correlation has led to an overestimation of anti-spoofing system performance [19]. To obtain a realistic estimate of the anti-spoofing model performance, it is crucial to eliminate this shortcut. We address this issue by following the suggestions made by related work [19, 30] and trim the beginning and end silences of an audio file. Moreover, we randomly select a 2-second segment from the remaining audio to ensure that all audio files have equal length. This prevents our models from using the presence or position of the cut-off as a new learning shortcut. We apply this technique to all subsequent analyses.

In addition, inspired by the widespread use of data augmentation in the image domain, we employ the following audio data augmentations in all subsequent analyses. By doing so, we further aim to avoid shortcuts, model overfitting, and promote model generalization.

First, we use Adversarial Retraining [31]. Namely, we employ the Fast Gradient Sign Method (FGSM) to craft adversarial examples with a probability of 20% and a magnitude of 0.25% of the input magnitude and supply these as additional training examples. This method has been shown to improve model generalization [31] and regularization [32]. Second, we use audio augmentations [33] such as room impulse response, frequency masking, pitch- and timeshift, gain, high, low- and bandpass filters, polarity inversion, and clipping to diversify the input data with a probability of 20%. Third, again with a probability of 20%, we partially encode the input using MP3, ulaw, and alaw codecs to perform compression. Lastly, we add random clips of noise and music to the audio files with a probability of 5% by utilizing the MusDB18 [34] and Noise ESC 50 [35] datasets.

We utilize the Adam optimizer with a learning rate of  $5 \cdot 10^{-3}$  to train all models. The models are trained for 25 epochs with a weight decay of  $10^{-6}$  and a batch-size of 32. To prevent overfitting, we employ early stopping with a patience of 3 epochs and a minimum delta of  $5 \cdot 10^{-4}$ . We implement the models using Python and PyTorch. For the Constant-Q Transform (CQT), we use a hop size of 32 samples with the “Hann” windowing function. We train all models using an Nvidia DGX-A100 server, which is equipped with eight A100 GPUs, each having a memory of 40GB, and a total server memory of 1024GB. It features two AMD Rome 7742 processors, each with 64 cores.

Table 1: *Model evaluation results on the “In-the-Wild” dataset. The results are averaged over three independent runs, with the standard deviation displayed.*

Model Name	Feature Type	EER Test ITW	EER Train	Epoch Time (s)	Num. Parameters
Proposed Model	C-CQT	<b>26.95±3.12</b>	12.382±6.58	276.6±6.5	2,459,492
CRNN Spoof [14]	Raw	29.90±4.33	8.103±0.59	177.0±5.0	3,330,562
MesoNet [12]	CQT	37.05±3.93	11.496±0.19	187.0±21.2	11,698
RawGat-ST [5]	Raw	39.29±2.41	11.802±0.12	402.5±1.1	440,810
RawNet2 [4]	Raw	40.00±0.25	11.680±0.02	113.4±5.9	17,648,770
LCNN [14]	Mel-Spec	55.07±7.09	21.575±1.82	170.8±12.2	178,306
Deep ResNet [29]	Mel-Spec	58.80±0.23	36.21±0.65	192.7±6.8	111,874

Table 2: *Ablation study results for the C-CQT feature: full phase retention, random phase selection, or zero assignment.*

Model Name	Phase	EER Test ITW	EER Train
Proposed Model	full	<b>26.95±3.12</b>	12.38±6.58
	zero	31.29±1.31	5.94±0.37
	random	35.79±0.13	8.09±0.18

### 4.3. Evaluation and Results

We adopt the Equal Error Rate (EER) as the performance metric for our model evaluation, consistent with previous studies [4, 5, 6, 28, 36]. EER is the point on the Receiver Operating Characteristic (ROC) curve where the false acceptance rate and false rejection rate are equal. For each experimental configuration, we conduct three independent trials and report the mean and standard deviation of the EER. Table 1 summarizes the results of our proposed model compared to models from related work. We observe the following: Firstly, our complex-valued model demonstrates an absolute improvement over related work by about 3% EER when compared to “raw” features and 10% EER when using magnitude-based CQT features. Secondly, we observe the least amount of overfitting, as indicated by the small gap between the training and testing performance.

### 4.4. Ablation Study

In order to evaluate the impact of the phase information, we perform an ablation study where we keep all parameters for our proposed model, but discard or randomize the phase information from the complex spectrogram. Table 2 presents the results. We observe that the removal of phase information (“zero phase”) results in a 4% degradation in model performance with respect to EER. This suggests that the proposed model has indeed learned to detect spoofing attacks by utilizing the mismatch between phase and magnitude information to some extent. Furthermore, randomizing the phase information causes an additional 4% decrease in performance with respect to EER. We presume that this is due to the fact that the model not only lacks phase information, but also has to cope with a significant amount of noise when dealing with randomly generated phase information.

### 4.5. Explainable AI

Our model allows the application of explainable AI techniques such as saliency maps [16] and Smooth Grad [17]. These techniques allow us to visualize how our model processes audio in the frequency domain, thereby enabling us to gain insight into




Figure 2: *SmoothGrad [17] applied on an instance from the ASVspoof dataset.*

the decision-making process of our model. To demonstrate the effectiveness of these techniques, we apply Smooth Grad and present an exemplary output in Figure 2. Our analysis shows that the model does not rely on any obvious shortcuts, such as the duration of silence [19] or other artifacts such as sampling or ringing issues in the upper frequency bands. Instead, the model focuses on the speech signal itself, as evidenced by the red XAI-overlay coinciding with the spectro-temporal bins containing descriptive features.

## 5. Conclusion and Future Work

This paper introduces a novel approach to voice anti-spoofing, which involves utilizing complex-valued features and complex-valued neural networks. This approach is motivated by the fact that magnitude-based feature extraction discards phase information, which is difficult to spoof yet crucial for producing natural-sounding synthesized speech, thereby constituting a useful input feature to voice anti-spoofing systems.

Building on previous research, we thus propose using the complex-valued Constant-Q Transform (C-CQT) as a new input feature and suggest a corresponding complex-valued neural architecture. Our proposed approach not only outperforms previous related work but also allows for explainability through established XAI-methods, which is lacking in “raw”-feature-based models. Through ablation studies, we demonstrate that our model has effectively learned to use phase information to distinguish between spoofed and bona-fide audio. These results indicate the potential of our approach in improving the accuracy of voice anti-spoofing and advancing the understanding of the underlying decision-making process.

In summary, we suggest an effective architecture that can be useful for anti-spoofing purposes. Despite the simplicity of our current model, it already surpasses the performance of related works that employ more sophisticated components. We are confident that incorporating further model refinements would be straightforward and could potentially enhance the ability to detect spoofed voice signals.

## 6. References

- [1] D. Griffin and J. Lim, "Signal estimation from modified short-time fourier transform," *IEEE Transactions on acoustics, speech, and signal processing*, vol. 32, no. 2, pp. 236–243, 1984.
- [2] A. van den Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. W. Senior, and K. Kavukcuoglu, "Wavenet: A generative model for raw audio," *CoRR*, vol. abs/1609.03499, 2016. [Online]. Available: <http://arxiv.org/abs/1609.03499>
- [3] N. Chen, Y. Zhang, H. Zen, R. J. Weiss, M. Norouzi, and W. Chan, "Wavegrad: Estimating gradients for waveform generation," in *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021.
- [4] H. Tak, J. Jung, J. Patino, M. R. Kamble, M. Todisco, and N. W. D. Evans, "End-to-end spectro-temporal graph attention networks for speaker verification anti-spoofing and speech deepfake detection," *CoRR*, vol. abs/2107.12710, 2021.
- [5] H. Tak, J. Patino, M. Todisco, A. Nautsch, N. Evans, and A. Larcher, "End-to-End anti-spoofing with RawNet2," in *ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021, pp. 6369–6373.
- [6] W. Ge, J. Patino, M. Todisco, and N. Evans, "Raw Differentiable Architecture Search for Speech Deepfake and Spoofing Detection," in *Proc. 2021 Edition of the Automatic Speaker Verification and Spoofing Countermeasures Challenge*, 2021, pp. 22–28.
- [7] J. C. Brown, "Calculation of a constant q spectral transform," *The Journal of the Acoustical Society of America*, vol. 89, no. 1, pp. 425–434, 1991.
- [8] S. S. Stevens, J. Volkman, and E. B. Newman, "A scale for the measurement of the psychological magnitude pitch," *The journal of the acoustical society of america*, vol. 8, no. 3, pp. 185–190, 1937.
- [9] X. Wang and J. Yamagishi, "A comparative study on recent neural spoofing countermeasures for synthetic speech detection," in *Interspeech*, H. Hermansky, H. Cernocký, L. Burget, L. Lamel, O. Scharenborg, and P. Motlíček, Eds., 2021.
- [10] G. Lavrentyeva, S. Novoselov, E. Malykh, A. Kozlov, O. Kudashev, and V. Shchemelinin, "Audio replay attack detection with deep learning frameworks," in *Interspeech 2017*. ISCA, 2017, pp. 82–86. [Online]. Available: [http://www.isca-speech.org/archive/Interspeech\\_2017/abstracts/0360.html](http://www.isca-speech.org/archive/Interspeech_2017/abstracts/0360.html)
- [11] G. Lavrentyeva, S. Novoselov, A. Tseren, M. Volkova, A. Gorlanov, and A. Kozlov, "STC antispoofing systems for the ASVspoof2019 challenge," in *Interspeech*. ISCA, 2019, pp. 1033–1037. [Online]. Available: [http://www.isca-speech.org/archive/Interspeech\\_2019/abstracts/1768.html](http://www.isca-speech.org/archive/Interspeech_2019/abstracts/1768.html)
- [12] D. Afchar, V. Nozick, J. Yamagishi, and I. Echizen, "MesoNet: A Compact Facial Video Forgery Detection Network," in *IEEE International Workshop on Information Forensics and Security (WIFS)*, 2018, pp. 1–7.
- [13] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in neural information processing systems*, vol. 30, 2017.
- [14] A. Chintla, B. Thai, S. J. Sohrawardi, K. M. Bhatt, A. Hickerson, M. Wright, and R. Ptucha, "Recurrent Convolutional Structures for Audio Spoof and Video Deepfake Detection," *IEEE Journal of Selected Topics in Signal Processing*, pp. 1–1, 2020.
- [15] M. Ravanelli and Y. Bengio, "Speaker recognition from raw waveform with sincnet," in *2018 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2018, pp. 1021–1028.
- [16] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, Y. Bengio and Y. LeCun, Eds., 2015.
- [17] D. Smilkov, N. Thorat, B. Kim, F. Viégas, and M. Wattenberg, "Smoothgrad: removing noise by adding noise," *arXiv preprint arXiv:1706.03825*, 2017.
- [18] R. Geirhos, J.-H. Jacobsen, C. Michaelis, R. Zemel, W. Brendel, M. Bethge, and F. A. Wichmann, "Shortcut learning in deep neural networks," *Nature Machine Intelligence*, vol. 2, no. 11, pp. 665–673, 2020.
- [19] N. M. Müller, F. Dieckmann, P. Czempin, R. Canals, J. Williams, and K. Böttinger. Speech is Silver, Silence is Golden: What do ASVspoof-trained Models Really Learn?
- [20] N. M. Müller, P. Czempin, F. Dieckmann, A. Frogghar, and K. Böttinger, "Does audio deepfake detection generalize?" *Interspeech*, 2022.
- [21] R. McEachern, "How the ear really works," in *[1992] Proceedings of the IEEE-SP International Symposium on Time-Frequency and Time-Scale Analysis*. IEEE, 1992, pp. 437–440.
- [22] B. C. Moore, *An introduction to the psychology of hearing*. Brill, 2012.
- [23] B. Blankertz, "The constant q transform," *URL http://doc.ml.tu-berlin.de/bbci/material/publications/Bla\_constQ.pdf*, 2001.
- [24] M. Todisco, H. Delgado, and N. Evans, "Constant q cepstral coefficients: A spoofing countermeasure for automatic speaker verification," *Computer Speech & Language*, vol. 45, pp. 516–535, 2017.
- [25] K. Cho, B. van Merriënboer, D. Bahdanau, and Y. Bengio, "On the properties of neural machine translation: Encoder-decoder approaches," in *Proceedings of SSST@EMNLP 2014*, D. Wu, M. Carpuat, X. Carreras, and E. M. Vecchi, Eds., 2014.
- [26] J. Bassey, L. Qian, and X. Li, "A survey of complex-valued neural networks," *arXiv preprint arXiv:2101.12249*, 2021.
- [27] W. Wirtinger, "Zur formalen theorie der funktionen von mehr komplexen veränderlichen," *Mathematische Annalen*, vol. 97, no. 1, pp. 357–375, 1927.
- [28] M. Todisco, X. Wang, V. Vestman, M. Sahidullah, H. Delgado, A. Nautsch, J. Yamagishi, N. W. D. Evans, T. H. Kinnunen, and K. A. Lee, "Asvspoof 2019: Future horizons in spoofed and fake audio detection," in *Interspeech 2019*, G. Kubin and Z. Kacic, Eds. ISCA, 2019.
- [29] M. Alzantot, Z. Wang, and M. B. Srivastava, "Deep Residual Neural Networks for Audio Spoofing Detection," in *Interspeech*. ISCA, 2019, pp. 1078–1082.
- [30] Y. Zhang<sup>12</sup>, W. Wang<sup>12</sup>, and P. Zhang<sup>12</sup>, "The effect of silence and dual-band fusion in anti-spoofing system," *Interspeech*, 2021.
- [31] I. J. Goodfellow, J. Shlens, and C. Szegedy, "Explaining and harnessing adversarial examples," in *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, Y. Bengio and Y. LeCun, Eds., 2015.
- [32] "Regularization effect of adversarial examples learning," <http://chandlerzuo.github.io/blog/2018/10/fgsmreg>, (Accessed on 03/07/2023).
- [33] "iver56/audiomentations: A python library for audio data augmentation. inspired by albumations. useful for machine learning." <https://github.com/iver56/audiomentations>, (Accessed on 03/07/2023).
- [34] Z. Rafii, A. Liutkus, F.-R. Stöter, S. I. Mimitakis, and R. Bittner, "The MUSDB18 corpus for music separation," Dec. 2017. [Online]. Available: <https://doi.org/10.5281/zenodo.1117372>
- [35] K. J. Piczak, "ESC: Dataset for Environmental Sound Classification," in *Proceedings of the 23rd Annual ACM Conference on Multimedia*. ACM Press, 2015, pp. 1015–1018. [Online]. Available: <http://dl.acm.org/citation.cfm?doi=2733373.2806390>
- [36] A. Nautsch, X. Wang, N. Evans, T. H. Kinnunen, V. Vestman, M. Todisco, H. Delgado, M. Sahidullah, J. Yamagishi, and K. A. Lee, "ASVspoof 2019: Spoofing Countermeasures for the Detection of Synthesized, Converted and Replayed Speech," vol. 3, no. 2, pp. 252–265, 2021.