



A Comprehensive Real-World Assessment of Audio Watermarking Algorithms: Will They Survive Neural Codecs?

Yigitcan Özer^{1*}, Woosung Choi^{1*}, Joan Serra¹,
Mayank Kumar Singh¹, Wei-Hsiang Liao¹, Yuki Mitsufuji^{1,2}

¹Sony AI ²Sony Group Corporation
yiitozer@nii.ac.jp, woosung.choi@sony.com

Abstract

We present the Robust Audio Watermarking Benchmark (RAW-Bench) to foster the evaluation of deep learning-based audio watermarking algorithms, establishing a standardized benchmark and allowing systematic comparisons. To simulate real-world usage, we introduce a comprehensive audio attack pipeline featuring various distortions such as compression, background noise, and reverberation and propose a diverse test dataset, including speech, environmental sounds, and music recordings. By assessing the performance of four existing watermarking algorithms on our framework, two main insights stand out: (i) neural compression techniques pose the most significant challenge, even when algorithms are trained with such compressions; and (ii) training with audio attacks generally improves robustness, although it is insufficient in some cases. Furthermore, we find that specific distortions, such as polarity inversion, time stretching, or reverb, seriously affect certain algorithms. Our contributions strengthen the robustness and perceptual assessment of audio watermarking algorithms across a wide range of applications while ensuring a fair and consistent evaluation approach. The evaluation framework, including the attack pipeline, is accessible at github.com/SonyResearch/raw_bench.

Index Terms: robust audio watermarking, imperceptibility.

1. Introduction

Recent advances in audio-based applications have enabled seamless content sharing, improved creative workflows, and facilitated the widespread adoption of generative AI models [1–4]. However, such advancements have also introduced challenges in content authenticity and copyright protection [5, 6]. To address these challenges, audio watermarking has gained attention, embedding imperceptible but detectable information into signals. It embeds a hidden message within a carrier signal, ensuring inaudibility while enabling reliable detection and extraction [7]. Recent deep learning-based methods [8–13] have demonstrated remarkable improvements in robustness, imperceptibility, and efficiency over traditional approaches.

The effectiveness of an audio watermarking system is commonly evaluated using three criteria: *imperceptibility*, *robustness*, and *capacity*. Imperceptibility refers to the fidelity of the watermarked signal, which ensures that the embedded watermark remains inaudible. Robustness refers to the successful detection of the watermark, even under distortions or attacks that degrade the carrier signal and/or the watermark. Capacity represents the amount of information (that is, the number of message bits per unit time) that can be embedded in the carrier

signal. A key challenge in audio watermarking lies in the inherent trade-offs among these three properties, as optimizing one often comes at the expense of the others [14].

We propose the Robust Audio Watermarking Benchmark (RAW-Bench), focusing on imperceptibility and robustness, with comparable capacity across models. In this benchmark, we assume a threat model where adversaries have access only to the audio file, while the watermarking methods remain hidden. However, they can manipulate the audio file to prevent detection of the embedded watermark. For example, one might compress and then decompress the audio to damage the imperceptible watermark while maintaining the perceptible quality.

Our analysis evaluates four publicly available pre-trained baseline models (AudioSeal [8], SilentCipher [9], Timbre [10], and WavMark [11]) under various distortions, including mixing, background noise, filtering, reverberation, compression, and equalization. We also study the impact on the robustness of distortion-aware training, integrating a comprehensive audio attack pipeline into the training of AudioSeal and SilentCipher. To systematically evaluate performance, we introduce a novel test dataset covering multiple audio domains, including music, speech, and environmental sounds, with non-compressed raw recordings. Our findings provide two main insights: (i) neural compression (e.g., Encodec [15] and Descript Audio Codec [16]) poses the greatest challenge to audio watermarking systems, even when these systems are trained with such compressions; and (ii) training with audio attacks improves robustness, consistent with observations by Juvela and Wang [17], although it does not guarantee good performance. Additionally, we observe that specific distortions, such as polarity inversion, time stretching, and reverb, severely impact certain watermarking algorithms. We end with a discussion of future perspectives regarding the trade-off between audio watermarking and neural compression.

2. Related Work

To the best of our knowledge, the only study that compares deep learning-based audio watermarking models is AudioMarkBench [20]. AudioMarkBench is a benchmarking framework that evaluates the robustness of three audio watermarking models (AudioSeal, Timbre, and WavMark), using their publicly available pre-trained weights, on a subset of speech signals sampled at 16 kHz. Apart from additionally considering the imperceptibility criterion, our work diverges from AudioMarkBench in a number of important aspects. First, instead of relying on compressed recordings, we construct a high-fidelity test dataset containing raw, non-compressed audio at 44.1 kHz. This avoids any potential confounding factor introduced by low-bandwidth or compressed signals. Second, instead of focusing on speech,

*Equal contribution.

Table 1: Characteristics of baseline watermarking models. Capacity corresponds to the one we set/use for our evaluation.

Model	Domain	Sample rate (kHz)	Message (bits)	Capacity (bps)	Training size (h)	Training data
AS: AudioSeal [8]	Waveform	16	16	5.33	4500	Speech
SC: SilentCipher [9]	Spectral	16	23.8	5.33	372	Speech, music, TV shows
TI: Timbre [10]	Spectral	22.05	30	5.00	100	Speech
WM: WavMark [11]	Spectral	16	16	5.28	5000	Speech, music, environmental

Table 2: Audio attack pipeline for robustness analysis. The threshold is used to separate between loose (L) and strict (S) attacks.

Attack Category	Attack Type	Parameter	Range	L/S Threshold
Mixing	GN: Gaussian noise	SNR (dB)	[20, 60]	40
	BN: Background noise (from [18])	SNR (dB)	[20, 60]	35
	RV: Reverb (from [19])	SNR (dB)	[0, 12]	6
Dynamics	DC: Dynamic range compression	Threshold (dB)	[-36, -6]	-18
	DE: Dynamic range expansion	Threshold (dB)	[-16, -6]	-12
	LM: Limiter	Threshold (dB)	[-36, -6]	-18
Filtering	LP: Lowpass	Cutoff (Hz)	[3500, 8000]	6000
	HP: Highpass	Cutoff (Hz)	[10, 500]	250
	EQ: Equalization	Max gain (dB)	[-0.75, 0.75]	± 0.375
Low level	TS: Time stretch	Rate	[0.75, 1.25]	1.00 ± 0.05
	TJ: Time jittering	Scale	[0.10, 0.50]	0.20
	PI: Polarity inversion	N/A	N/A	N/A
	GA: Gain adjustment	Rate	[0.20, 5]	1.00 ± 0.50
	QN: Quantization	#Bits/sample	{8, 9, ... 16}	12
Neural compression	PS: Phase shift	Seconds	[-0.10, 0.10]	0 ± 0.05
	EN: Encoder [15] (at 24 kHz)	#Codebooks	{16, 32}	32
Conventional compression	DA: Descript Audio Codec [16] (at 44.1 kHz)	#Codebooks	{7, 8, 9}	9
	MP: MP3 codec	Bitrate (kbps)	{64, 128, 256}	64
	OG: OGG codec	Bitrate (kbps)	{48, 64, 128, 256}	48
	AA: AAC codec	Bitrate (kbps)	{64, 128, 256}	64

we base our results on a selection of speech, music, and environmental sounds. This ensures a more comprehensive evaluation under a variety of real-world signals. Third, we extend the analysis by including SilentCipher, another baseline approach with competitive performance. This broadens the scope of our work. Fourth, we study the effect of retraining watermarking algorithms using the proposed audio attack pipeline. This allows us to isolate the impact of training-time distortions on watermark robustness, providing deeper insights into the advantages of training with adversarial attacks. Finally, it is also worth mentioning that our attack/test pipeline is larger and more varied than the one of AudioMarkBench (20 vs. 12 distortions, respectively, see below).

All models considered in this paper encode a hidden message in a mono-carrier signal, but differ in architecture, sampling rate, training dataset, and operating domain, as detailed in Table 1. To ensure an unbiased evaluation, we verify that no training data from the considered models is included in our test set. This guarantees that our analysis and conclusions are based entirely on out-of-sample data. AudioSeal (AS) [8] is originally trained with various distortion augmentations, including time modifications, filtering, audio effects, and compression. SilentCipher (SC) [9] considers time jittering, additional noise, and non-differentiable compression techniques as augmentations. Additionally, it introduces a lower SDR bound on the watermarked signals to account for imperceptibility. Timbre (TI) [10] incorporates ISTFT, normalization, transformation, and wave reconstruction for robustness. WavMark (WM) [11] employs a curriculum learning strategy and applies various distortions during training, including noise, filtering, compression, echo, and time stretching.

3. Methodology

Test Dataset — To evaluate watermarking algorithms in various domains, we create a comprehensive test dataset using open-source collections from various sources. It includes classical and popular music, speech, and environmental sounds, which account for a wide range of real-world use cases. To maintain a high fidelity, all audio recordings in the dataset have sample rates equal to or exceeding 44.1 kHz, and they are provided as raw, non-compressed files. Our test dataset is formed by the union of the following publicly-available collections¹:

- Bach10 [21] – A dataset of ten classical ensemble recordings.
- Clotho [22] – A collection of diverse environmental sounds.
- Device and Produced Speech (DAPS) [23] – A dataset of studio-quality speech recordings alongside consumer-device recordings captured in real-world environments.
- FreiSchuetz [24] – A dataset of professional stereo mixes and raw multitrack recordings of three opera performances.
- GuitarSet [25] – A dataset of solo guitar recordings.
- jaCapella [26] – A corpus of 50 Japanese a cappella vocal ensemble recordings, including individual voice parts.
- MAESTRO [27] – A dataset of paired audio and MIDI recordings from the International Piano-e-Competition.
- MoisesDB [28] – A dataset of 240 musical tracks spanning twelve genres, performed by 45 artists (we use only the mixes).
- Piano Concerto Dataset (PCD) [29] – A collection of piano concerto excerpts (we use only the raw piano tracks).

¹github.com/SonyResearch/raw_bench

Attack Pipeline — For the robustness analysis, we develop a comprehensive audio attack pipeline that simulates 20 real-world distortions (Table 2). The distortions are organized into six categories and are designed to simulate real-world variability in playback and processing. As many attacks allow for parameter variation (e.g., different levels of noise, filtering, or compression strength), we first establish a range of suitable values that align with real-world conditions (Table 2, Range). Next, we consider two settings based on the strength of such parameters: *loose* and *strict*. The loose setting corresponds to imperceptible distortions, while the strict setting represents cases where distortions are audible but still acceptable. To define a threshold between these two settings (Table 2, Threshold), we conducted an internal listening test with five expert listeners who evaluated the perceptibility and acceptability of each attack. The threshold was then set based on whether they could perceive a notable subjective difference compared to the original audio. For example, the listeners unanimously agreed that Gaussian noise with Signal-to-Noise Ratio (SNR) values in the range of [40, 60] dB is almost imperceptible, placing it in the loose class. In contrast, noise in the range of [20, 40] dB is audible but still acceptable, categorizing it under the strict class. These thresholds allow us to systematically evaluate watermarking models under both realistic (loose) and challenging (strict) conditions, ensuring a meaningful robustness analysis.

Building on this attack framework, our robustness experiments follow a two-stage process. First, we evaluate pre-trained models by assessing their performance against the full set of distortions (loose and strict setups). Second, we retrain AS and SC using our audio attack pipeline under the strict parameter settings to examine the impact of adversarial training on watermark robustness. To ensure a balanced exposure to different distortions, we employ a uniform weighting scheme per attack category, along with spectrogram augmentation [30]. For re-training, we utilize a proprietary dataset consisting of approximately 1250 hours of musical mixes, along with 40 hours of VCTK [31] and a 40-hour subset of environmental sounds from BBC Sound Effects [32].

Evaluation Metrics — To evaluate the performance of the considered methods, we employ a set of metrics that measure either the robustness of watermark detection or the imperceptibility of the watermark. Our analysis focuses on these two aspects while keeping the capacity constant and comparable across all considered models (Table 1). For imperceptibility, we consider:

- Scale-invariant signal-to-noise ratio (SI-SNR) [33] – SI-SNR measures the distortion or noise in a processed signal relative to a reference, independent of scaling.
- Mel cepstral distance (MCD) [34] – MCD is a perceptually motivated measure that quantifies the spectral difference between the original and the watermarked audio.
- Virtual Speech Quality Objective Listener (MOS-LQO) [35] – An objective, full-reference metric for assessing perceived audio quality based on spectro-temporal similarity.

For robustness, we use:

- Bitwise accuracy – This metric measures the proportion of correctly decoded bits in the detected watermark message.
- Message accuracy – This metric assesses the overall success of the watermark extraction process (that is, if all bits in the extracted message match the original watermark).

Table 3: Comparison of models across different metrics on clean watermarked audio (no attacks). ACC shows average bitwise/full-message accuracy, and an asterisk (*) indicates that the model has been retrained with the strict attacks.

Model	SI-SNR \uparrow	MCD \downarrow	MOS-LQO \uparrow	ACC \uparrow
AS	22.73	0.53	4.93	0.997 / 0.962
SC	49.13	0.25	4.98	0.999 / 0.993
TI	21.91	1.74	4.59	1.000 / 1.000
WM	35.89	0.62	4.91	0.998 / 0.993
AS*	23.60	0.49	4.95	0.999 / 0.997
SC*	31.80	1.04	4.88	0.999 / 0.993

Table 4: Average bitwise/full-message accuracy across all strict attacks, for different audio domains.

Model	Environ.	Music	Speech
AS	.91 / .68	.91 / .68	.91 / .72
SC	.73 / .47	.75 / .49	.81 / .62
TI	.94 / .79	.93 / .78	.94 / .78
WM	.74 / .70	.77 / .72	.80 / .77
AS*	.95 / .81	.95 / .81	.94 / .80
SC*	.91 / .75	.90 / .79	.92 / .80

4. Results and Discussion

Imperceptibility — As a first step, we evaluate the considered models in clean (distortion-free) conditions, focusing on overall perceptual quality and detection accuracy (Table 3). Among all pre-trained models, SC consistently outperforms others in perceptual quality, achieving the highest SI-SNR and lowest MCD, indicating minimal perceptual impact from watermark insertion. Similarly, SC achieves the highest MOS-LQO score, closely followed by AS, while TI performs the worst across all perceptual metrics. The better performance of SC regarding the perceptual metrics can be attributed to its lower SDR bound constraint of the watermarked signals. In terms of overall robustness, all models achieve accuracies close to 1 in clean conditions (Table 3, ACC), and we additionally measure a true-positive rate between 0.97 and 1 at zero false-negative rate for all of them (not shown). Overall, this confirms a reliable watermark extraction in the absence of audio attacks. In this clean setup, results for the re-trained models AS* and SC* do not differ much from the pre-trained ones, except for the case of SC* with SI-SNR and MCD, which we on purpose re-train with a lower SDR bound to improve robustness (see below).

Audio Domains — Next, we analyze the robustness across the different audio domains found in our test dataset (Table 4). We find that all models exhibit similar performance across environmental sounds, music, and speech, with only minor variations in accuracy between domains. Interestingly, AS and TI were trained exclusively on speech data, yet they generalize well to the other two domains. This suggests that the considered models can generalize well across different audio domains. We also observe that training with adversarial attacks further improves robustness, especially for full-message accuracy.

Robustness — We now evaluate how different distortions impact the robustness of watermarking models. Table 5 presents the bitwise (top) and full-message (bottom) accuracy of both pre-trained and retrained models under strict and loose attack conditions. In general, TI demonstrates the highest robustness,

Table 5: Comparison of bitwise (top) and full-message (bottom) robustness for the considered models under various attacks (columns, see abbreviations in Table 2). For each model, we evaluate the strict (S) and loose (L) settings (Eval column). An asterisk (*) indicates that the model has been retrained with the strict attacks, and a checkmark (✓) indicates an accuracy of 0.99 or above.

Model	Eval	GN	BN	RV	DC	DE	LM	LP	HP	EQ	TS	TJ	PI	GA	QN	PS	EN	DA	MP	OG	AA
AS	S	✓	✓	.87	✓	✓	.98	✓	.96	.91	.97	✓	.18	✓	✓	.62	.96	.52	✓	✓	✓
SC	S	.63	.98	.80	.96	.91	.83	✓	.93	.92	.41	.86	✓	✓	✓	.63	.33	.32	.54	.61	.93
TI	S	.98	✓	.96	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	.64	.60	✓	✓	✓
WM	S	.83	✓	.89	.98	.95	.95	✓	✓	✓	.82	.98	✓	.98	✓	✓	.00	.00	.42	.89	✓
AS	L	✓	✓	.97	✓	✓	✓	✓	✓	✓	.97	✓	.18	✓	✓	.60	.97	.53	✓	✓	✓
SC	L	✓	✓	.89	✓	.97	✓	✓	✓	✓	.64	.88	✓	✓	✓	.65	.33	.32	.57	.65	.98
TI	L	✓	✓	.98	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	.65	.62	✓	✓	✓
WM	L	✓	✓	.98	✓	✓	✓	✓	✓	✓	.95	✓	✓	.98	✓	✓	.00	.00	.48	.83	✓
AS*	S	✓	✓	.91	✓	✓	✓	✓	✓	.96	✓	✓	.98	✓	✓	.62	.97	.60	✓	✓	✓
SC*	S	✓	✓	✓	✓	✓	✓	✓	✓	✓	.54	✓	✓	✓	✓	.91	.67	.42	.98	✓	✓
AS*	L	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	.98	✓	✓	.59	✓	.61	✓	✓	✓
SC*	L	✓	✓	✓	✓	✓	✓	✓	✓	✓	.84	✓	✓	✓	✓	.92	.75	.44	.98	✓	✓
AS	S	.93	.95	.22	.92	.89	.88	.93	.57	.39	.72	.96	.00	.91	.95	.06	.65	.00	.92	.95	.96
SC	S	.29	.93	.45	.78	.64	.49	✓	.71	.69	.00	.60	.98	✓	✓	.24	.00	.00	.14	.22	.75
TI	S	.80	.98	.57	.96	.94	.95	✓	✓	✓	.90	.97	✓	✓	✓	✓	.00	.00	.92	.97	✓
WM	S	.74	.98	.73	.95	.90	.90	✓	✓	.98	.60	.96	✓	.94	✓	✓	.00	.00	.29	.79	✓
AS	L	.95	.96	.72	.96	.94	.96	.96	.88	.90	.71	.95	.00	.91	.95	.02	.78	.00	.93	.95	.96
SC	L	✓	✓	.67	✓	.89	.97	✓	.96	.96	.20	.61	.98	✓	✓	.28	.00	.00	.17	.30	.94
TI	L	✓	✓	.77	✓	.98	✓	✓	✓	✓	.92	.97	✓	.98	✓	✓	.00	.00	.94	.94	✓
WM	L	✓	✓	.93	✓	.98	✓	✓	✓	✓	.88	.98	✓	.97	✓	✓	.00	.00	.34	.75	✓
AS*	S	.98	✓	.41	✓	✓	✓	✓	.98	.62	.95	✓	.72	✓	✓	.03	.79	.00	✓	✓	✓
SC*	S	.96	.98	.95	✓	.98	.98	.98	.97	.97	.00	✓	.98	✓	✓	.74	.06	.00	.85	.91	.98
AS*	L	✓	✓	.90	✓	✓	✓	✓	✓	✓	.93	✓	.72	✓	✓	.01	.90	.00	✓	✓	✓
SC*	L	✓	✓	.98	✓	✓	✓	.98	✓	.98	.37	✓	.98	✓	✓	.75	.12	.00	.87	.91	.98

which is expected as compensation for the low imperceptibility scores obtained above. In addition, different models show varying degrees of resilience to specific attacks. For example, AS struggles with PI and PS, SC is particularly vulnerable to TS, PS, MP, and OG, and WM is notably affected by MP compression. Among all distortions, neural compression methods (EN and DA) pose the greatest challenge, particularly for SC and WM under strict conditions, with bitwise accuracies dropping to 0.33 and 0, respectively. DA has a more severe effect on all models. Full-message accuracy is unacceptable for all models and settings, except for AS with EN compression, which can be explained by the fact that the architecture of AS is based on EN. However, importantly, the advantage of using the EN architecture does not extend to DA, suggesting that employing a neural codec architecture does not generalize to alternative neural codecs. Overall, this indicates that audio watermarking models fail under neural compression, highlighting a critical weakness.

Re-training — Now we assess the effect of incorporating adversarial attacks in SC and AS training (Table 5). While improving robustness against certain attacks (e.g., GN for SC and PI for AS), it does not fully resolve vulnerabilities in other cases: even after retraining, models continue to struggle with compression (EN, DA, MP, OG), RV, and PS. This suggests that some attacks introduce fundamental challenges that adversarial-attack training augmentations fail to overcome. Full-message accuracy is poor across all models, even after retraining, highlighting a fundamental limitation of existing methods.

Will Watermarks Survive Neural Codecs? — One of the main performance gaps we observe in our results is the robustness to neural codecs. Bitwise accuracies are generally below 0.5, and full-message accuracies are around 0 for almost all approaches in both EN and DA. While AS performs better than the rest for EN and training with neural-codec attacks can help

AS and SC, none of the considered watermarking algorithms survives the DA attack. Also, retraining does not bring robustness to neural codecs up to acceptable levels. This suggests that there is a fundamental issue underlying such poor generalization and lack of performance. In fact, watermarking algorithms and neural codecs compete for the same space. On the one hand, watermarking algorithms strive to insert imperceptible information into the audio signal, but on the other hand, neural codecs strive to remove the imperceptible information from the (possibly the same) audio signal. Deep learning methodologies have recently enhanced the capabilities of both types of algorithms. However, if we consider the limit situation where both algorithms successfully achieve their purpose, we believe that neural codecs will end up removing imperceptible watermarks. In addition, neural codecs are usually the final stage in the audio processing pipeline and thus have more chance/incentive to remove any imperceptible information, regardless of its origin.

5. Conclusion

We introduced a systematic evaluation framework for deep learning-based audio watermarking algorithms, addressing important gaps in robustness analysis and benchmarking. We designed a comprehensive audio attack pipeline that simulates real-world distortions, and introduced a diverse test dataset comprising multiple audio domains. By studying the performance of four existing watermarking algorithms within our framework, we were able to provide novel insights regarding imperceptibility and robustness to specific attacks. On the whole, our framework contributes to the development of more resilient and perceptually optimized audio watermarking systems. We believe future work should focus on the trade-off between audio watermarking and neural codecs, which our study discusses and shows to be a critical point.

6. References

- [1] Y. Ren, C. Hu, X. Tan, T. Qin, S. Zhao, Z. Zhao, and T.-Y. Liu, “Fastspeech 2: Fast and high-quality end-to-end text to speech,” in *Proc. of the Int. Conf. on Learn. Represent. (ICLR)*, Vienna, Austria, 2021.
- [2] M. Le, A. Vyas, B. Shi, B. Karrer, L. Sari, R. Moritz, M. Williamson, V. Manohar, Y. Adi, J. Mahadeokar, and W.-N. Hsu, “VoiceBox: Text-guided multilingual universal speech generation at scale,” in *Adv. in Neural Inf. Process. Syst. (NeurIPS)*, New Orleans, LA, 2023, pp. 14 005–14 034.
- [3] J. Copet, F. Kreuk, I. Gat, T. Remez, D. Kant, G. Synnaeve, Y. Adi, and A. Défossez, “Simple and controllable music generation,” in *Adv. in Neural Inf. Process. Syst. (NeurIPS)*, New Orleans, LA, USA, 2023.
- [4] S.-L. Wu, C. Donahue, S. Watanabe, and N. J. Bryan, “Music ControlNet: Multiple time-varying controls for music generation,” *IEEE/ACM Trans. on Audio, Speech, and Lang. Process.*, vol. 32, pp. 2692–2703, 2024.
- [5] X. Wang and J. Yamagishi, “A comparative study on recent neural spoofing countermeasures for synthetic speech detection,” in *Proc. of the Annual Conf. of the Int. Speech Commun. Assoc. (Interspeech)*, Brno, Czech Republic, 2021, pp. 4259–4263.
- [6] R. Batlle-Roca, W.-H. Liao, X. Serra, Y. Mitsufuji, and E. Gómez, “Towards assessing data replication in music generation with music similarity metrics on raw audio,” in *Proc. of the Int. Soc. for Music Inf. Retrieval Conf. (ISMIR)*, San Francisco, CA, USA, 2024, pp. 1004–1011.
- [7] G. Hua, J. Huang, Y. Q. Shi, J. Goh, and V. L. Thing, “Twenty years of digital audio watermarking – A comprehensive review,” *Signal Process.*, vol. 128, pp. 222–242, 2016.
- [8] R. S. Roman, P. Fernandez, H. Elshahar, A. Défossez, T. Furon, and T. Tran, “Proactive detection of voice cloning with localized watermarking,” in *Proc. of the Int. Conf. on Mach. Learn. (ICML)*, Vienna, Austria, 2024.
- [9] M. K. Singh, N. Takahashi, W. Liao, and Y. Mitsufuji, “SilentCipher: Deep audio watermarking,” in *Proc. of the Annual Conf. of the Int. Speech Commun. Assoc. (Interspeech)*, Kos Island, Greece, 2024, pp. 2235–2239.
- [10] C. Liu, J. Zhang, T. Zhang, X. Yang, W. Zhang, and N. Yu, “Detecting voice cloning attacks via Timbre Watermarking,” in *Netw. and Distrib. Syst. Secur. Symp.*, Vancouver, Canada, 2024.
- [11] G. Chen, Y. Wu, S. Liu, T. Liu, X. Du, and F. Wei, “Wavmark: Watermarking for audio generation,” 2023.
- [12] P. O’Reilly, Z. Jin, J. Su, and B. Pardo, “Maskmark: Robust neural watermarking for real and synthetic speech,” in *Proc. of the IEEE Int. Conf. on Acoust., Speech, and Signal Process. (ICASSP)*, Seoul, South Korea, 2024, pp. 4650–4654.
- [13] R. S. Roman, P. Fernandez, A. Deleforge, Y. Adi, and R. Serizel, “Latent watermarking of audio generative models,” in *Proc. of the IEEE Int. Conf. on Acoust., Speech, and Signal Process. (ICASSP)*, Hyderabad, India, 2025.
- [14] N. Agarwal, A. K. Singh, and P. K. Singh, “Survey of robust and imperceptible watermarking,” *Multimed. Tools Appl.*, vol. 78, no. 7, pp. 8603–8633, 2019.
- [15] A. Défossez, J. Copet, G. Synnaeve, and Y. Adi, “High fidelity neural audio compression,” *Trans. on Mach. Learn. Research (TMLR)*, 2023.
- [16] R. Kumar, P. Seetharaman, A. Luebs, I. Kumar, and K. Kumar, “High-fidelity audio compression with improved RVQGAN,” in *Adv. in Neural Inf. Process. Syst. (NeurIPS)*, New Orleans, LA, USA, 2023.
- [17] L. Juvela and X. Wang, “Audio codec augmentation for robust collaborative watermarking of speech synthesis,” in *Proc. of the IEEE Int. Conf. on Acoust., Speech, and Signal Process. (ICASSP)*, Hyderabad, India, 2025.
- [18] J. Thiemann, N. Ito, and E. Vincent, “The diverse environments multi-channel acoustic noise database (DEMAND): A database of multichannel environmental noise recordings,” *Proc. of Meet. on Acoust.*, vol. 19, no. 1, p. 035081, 2013.
- [19] M. Jeub, M. Schäfer, and P. Vary, “A binaural room impulse response database for the evaluation of dereverberation algorithms,” in *Proc. of Int. Conf. on Digital Signal Process. (DSP)*, Santorini, Greece, 2009, pp. 1–4.
- [20] H. Liu, M. Guo, Z. Jiang, L. Wang, and N. Z. Gong, “AudioMark-Bench: Benchmarking robustness of audio watermarking,” in *Adv. in Neural Inf. Process. Syst. (NeurIPS)*, Vancouver, Canada, 2024.
- [21] Z. Duan and B. Pardo, “Soundprism: An online system for score-informed source separation of music audio,” *IEEE Journal of Selected Topics in Signal Process.*, vol. 5, no. 6, pp. 1205–1215, 2011.
- [22] K. Drossos, S. Lipping, and T. Virtanen, “Clotho: An audio captioning dataset,” in *Proc. of the IEEE Int. Conf. on Acoust., Speech, and Signal Process. (ICASSP)*, 2020, pp. 736–740.
- [23] G. J. Mysore, “Can we automatically transform speech recorded on common consumer devices in real-world environments into professional production quality speech? – A dataset, insights, and challenges,” *IEEE Signal Process. Lett.*, vol. 22, no. 8, pp. 1006–1010, 2015.
- [24] T. Prätzlich, M. Müller, B. W. Bohl, and J. Veit, “Freischütz Digital: Demos of audio-related contributions,” in *Demos and Late Breaking News of the Int. Soc. for Music Inf. Retrieval Conf. (ISMIR)*, Málaga, Spain, 2015.
- [25] Q. Xi, R. Bittner, J. Pauwels, X. Ye, and J. P. Bello, “Guitarset: A dataset for guitar transcription,” in *Proc. of the Int. Soc. for Music Inf. Retrieval Conf. (ISMIR)*, Paris, France, 2018, pp. 453–460.
- [26] T. Nakamura, S. Takamichi, N. Tanji, S. Fukayama, and H. Saruwatari, “jaCappella corpus: A Japanese a cappella vocal ensemble corpus,” in *Proc. of the IEEE Int. Conf. on Acoust., Speech, and Signal Process. (ICASSP)*, 2023.
- [27] C. Hawthorne, A. Stasyuk, A. Roberts, I. Simon, C. A. Huang, S. Dieleman, E. Elsen, J. H. Engel, and D. Eck, “Enabling factorized piano music modeling and generation with the MAESTRO dataset,” in *Proc. of the Int. Conf. on Learn. Represent. (ICLR)*, New Orleans, LA, USA, 2019.
- [28] I. Pereira, F. Araújo, F. Korzeniowski, and R. Vogl, “MoisesDB: A dataset for source separation beyond 4-stems,” in *Proc. of the Int. Soc. for Music Inf. Retrieval Conf. (ISMIR)*, Milan, Italy, 2023, pp. 619–626.
- [29] Y. Özer, S. Schwär, V. Arifi-Müller, J. Lawrence, E. Sen, and M. Müller, “Piano Concerto Dataset (PCD): A multitrack dataset of piano concertos,” *Trans. of the Int. Soc. for Music Inf. Retrieval (TISMIR)*, vol. 6, no. 1, pp. 75–88, 2023.
- [30] D. S. Park, W. Chan, Y. Zhang, C.-C. Chiu, B. Zoph, E. D. Cubuk, and Q. V. Le, “SpecAugment: A simple data augmentation method for automatic speech recognition,” in *Proc. of the Annual Conf. of the Int. Speech Commun. Assoc. (Interspeech)*, Graz, Austria, 2019, pp. 2613–2617.
- [31] C. Veaux, J. Yamagishi, and K. MacDonald, “CSTR VCTK Corpus: English multi-speaker corpus for CSTR voice cloning toolkit (version 0.92),” 2019.
- [32] B. B. Corporation, “BBC sound effects library,” 1991, accessed: 2025-02-07. [Online]. Available: <https://sound-effects.bbcrewind.co.uk/>
- [33] Y. Luo and N. Mesgarani, “Conv-TasNet: Surpassing ideal time-frequency magnitude masking for speech separation,” *IEEE/ACM Trans. on Audio, Speech, and Lang. Process.*, vol. 27, no. 8, pp. 1256–1266, 2019.
- [34] R. F. Kubichek, “Mel-cepstral distance measure for objective speech quality assessment,” in *Proc. of IEEE Pacific Rim Conf. on Commun. Comp. and Signal Process.*, vol. 1, 1993, pp. 125–128.
- [35] A. Hines, J. Skoglund, A. C. Kokaram, and N. Harte, “ViSQOL: an objective speech quality model,” *EURASIP Journal on Audio, Speech, and Music Process.*, vol. 2015, pp. 1–18, 2015.