A Cold Diffusion Approach for Percussive Dereverberation

Dimos Makris
Department of Music Technology and Acoustics
Hellenic Mediterranean University
Rethymno, Greece
dimakr@hmu.gr

András Barják
XLN Audio
Stockholm, Sweden
andras.barjak@xlnaudio.com

Maximos Kaliakatsos-Papakostas
Department of Music Technology and Acoustics
Hellenic Mediterranean University
Rethymno, Greece
maximoskp@hmu.gr

arXiv:2605.10256v1 [cs.SD] 11 May 2026

***Abstract*—Most recent advances in audio dereverberation focus almost exclusively on speech, leaving percussive and drum signals largely unexplored despite their importance in music production. Percussive dereverberation poses distinct challenges due to sharp transients and dense temporal structure. In this work, we propose a cold diffusion framework for dereverberating stereo drum stems (downmixes), modeling reverberation as a deterministic degradation process that progressively transforms anechoic signals into reverberant ones. We investigate two reverse-process parameterizations—Direct (next-state) and a Δ-normalized residual (velocity-style) prediction—and implement the framework using both a UNet and a diffusion Transformer backbone. The models are trained and evaluated on curated datasets comprising both acoustic and electronic drum recordings, with reverberation generated using a combination of synthetic and real room impulse responses. Extensive experiments on in-domain and fully out-of-domain test sets demonstrate that the proposed method consistently outperforms strong score-based and conditional diffusion baselines, evaluated using signal-based and perceptual metrics tailored to percussive audio. Code and audio examples are available online[1].**

***Index Terms*—cold diffusion, drums dereverberation, audio enhancement, generative models, music production.**

I. Introduction

AUDIO enhancement algorithms aim to improve the quality of recorded sound by suppressing unwanted distortions such as background noise or reverberation [1]. In particular, reverberation – the persistence of sound due to reflections in an acoustic environment – can smear temporal and spectral details of audio, thereby degrading clarity and intelligibility. Therefore, *dereverberation* is the process of removing or reducing reverberant effects and has long been recognised as a crucial task for improving the overall perceived quality, specifically on speech [2].

Indeed, most prior research in audio enhancement has focused on speech signals, motivated by applications in robust telecommunication, automatic speech recognition, and speaker separation [3]. Early speech dereverberation methods ranged from classical signal processing techniques (e.g., estimating and subtracting late reflections via a weighted prediction error algorithm [4]) to modern deep learning models that directly learn to map reverberant speech to its anechoic counterpart (e.g., [5], [6]). In parallel, generative approaches, such as diffusion-based, have recently gained traction in this domain with prominent results [7]–[10]. Within diffusion-based speech dereverberation, score-based formulations in the complex STFT domain have demonstrated strong performance and flexibility [11], [12]. Despite this ongoing rapid progress [13], the methodological and evaluation conventions of the field remain strongly speech-centric.

In contrast, dereverberation for *music* signals has received comparatively less attention and is complicated by the diversity of sources, production effects, and perceptual objectives [14]. Recent music-oriented dereverberation studies have primarily focused on vocals (often under artificial reverb) and highlight the difficulty of collecting sufficiently diverse paired dry/wet data [15]. For *percussive* material, the problem is especially delicate: Drum hits have sharp onsets and fast-decaying transients, and reverberation can blur onset timing and redistribute energy into the tail, which impacts rhythmic clarity and downstream analysis/production operations [16].

Reverberation is commonly added to drums in music production for aesthetic effect, but there is growing interest in tools that can remove or reduce reverb from recorded percussion – for instance, to tighten up a drum mix or adapt a recording to a different acoustic context. Moreover, many of the commercial "de-reverb" plugins are typically optimised for dialogue/vocals (e.g., Waves Clarity[2]) and may not generalise well to percussive inputs. Therefore, this work targets a specific and practically motivated setting: **Drums Dereverberation**. Unlike isolated single-hit drum samples, stereo downmixed drums exhibit dense transient activity and overlapping resonances across kit elements, which complicates blind dereverberation [17]. Additionally, percussion lacks continuous harmonic content or phonetic structure to guide a model; small errors in a drum's transient or timbre can be perceptually obvious. These factors make percussive dereverberation a distinct and challenging task.

To the best of our knowledge, this is the first attempt to investigate a learning-based solution specifically for blind dereverberation of percussive signals. We adopt and extend *cold diffusion*, a family of diffusion-like iterative restoration methods that replace additive Gaussian corruption with a *de-*

[1]https://github.com/dimakr169/drums_dereverb

[2]https://www.waves.com/plugins/clarity-vx-dereverb

terministic degradation operator and learn to invert it through iterative reconstruction [18]. Cold diffusion has recently been instantiated for speech enhancement [19], demonstrating that deterministic degradation schedules can yield strong restoration performance and robustness. It has also been used in music source separation settings where a structured forward process enables learning from imperfect targets [20].

Inspired by these developments, we define a deterministic forward process that interpolates between anechoic and reverberant complex stereo spectrograms. We study two reverse prediction modes: (i) *Direct (next-state) prediction*, which predicts the next less-reverberant intermediate spectrogram at each step, and (ii) Δ*-normalized residual (velocity-style) prediction*, which predicts the step-size-normalized difference between consecutive intermediate spectrograms. We instantiate both with a UNet [21] and a diffusion Transformer [22] (DiT), and compare them against two strong diffusion baselines originally developed for speech enhancement/dereverberation: SGMSE+ [11], [12] and CDiffuSE [10].

Since most related diffusion-based pipelines are designed and evaluated in the speech domain, this is reflected not only in the training data and degradations but also in common evaluation protocols. Widely used perceptual/intelligibility metrics such as PESQ, POLQA, and ESTOI are explicitly designed for speech quality or intelligibility assessment [23]–[25]. More broadly, objective quality measures can be strongly *application-domain dependent* and may not transfer reliably across domains such as speech, music and source separation [26]. Consequently, adapting such evaluation practice directly to percussive inputs is not well-motivated. Thus, we propose a standard signal-oriented evaluation scheme with a dedicated set of *percussive-oriented* low-level and high-level objective metrics tailored to transient-rich audio. Experimental results demonstrate that our cold diffusion models consistently outperform these baselines, supporting cold diffusion as an effective and practical paradigm for transient-rich dereverberation.

The remainder of this paper is organised as follows: Sec. II presents the proposed cold diffusion framework and model variants, Sec. III describes the dataset and experimental setup including baselines, Sec. IV reports results and examples, while Sec. V concludes the paper.

II. Method

Following recent enhancement work in the complex STFT domain [11], [12], we operate on the *stereo* complex spectrogram of the drum stem. Audio is segmented into 2 s excerpts and transformed using an STFT with FFT size 1024 and hop size 384, yielding F frequency bins and K time frames. For each excerpt, we compute complex STFTs for the left and right channels, $X^{(L)}, X^{(R)} \in \mathbb{C}^{F\times K}$, and represent them using real and imaginary (RI) components. Concretely, we stack the channels as $\mathbf{x} \in \mathbb{R}^{4\times F\times K}$ with $\mathbf{x} = [\Re\{X^{(L)}\}, \Im\{X^{(L)}\}, \Re\{X^{(R)}\}, \Im\{X^{(R)}\}]$. This representation preserves phase information and allows the model to directly learn stereo-consistent dereverberation in the complex domain.

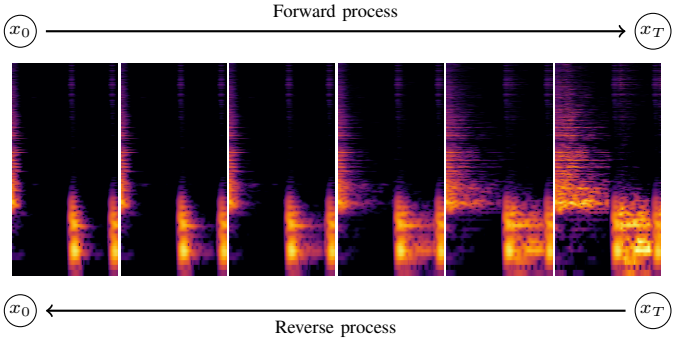

Fig. 1. Illustration of the cold diffusion process on a spectrogram. In the forward process, the anechoic signal $\mathbf{x}_0$ is progressively mixed with its reverberant counterpart, while the reverse process learns to iteratively remove reverberation starting from the reverberant signal $\mathbf{x}_T$.

A. Forward Process

We adopt a *cold diffusion* formulation, where the forward "process" is defined via a deterministic degradation operator rather than additive Gaussian noise [18]. Let $\mathbf{x}_0$ denote the anechoic (clean) stereo RI spectrogram and $\mathbf{y}$ the corresponding reverberant spectrogram. We define a sequence $\{\mathbf{x}_t\}_{t=0}^{T}$ that deterministically interpolates between $\mathbf{x}_0$ and $\mathbf{y}$ using a time-dependent mixing schedule:

$$\mathbf{x}_t \;=\; a_t\,\mathbf{x}_0 \;+\; (1-a_t)\,\mathbf{y}, \qquad a_0 = 1,\; a_T = 0, \tag{1}$$

where $\{a_t\}_{t=0}^{T}$ is a monotone schedule. In our implementation, we use a cosine-squared schedule

$$a_t \;=\; \cos^2\!\Big(\frac{\pi}{2}\frac{t}{T}\Big), \qquad t = 0,\ldots,T, \tag{2}$$

which yields a smooth progression from $a_0 = 1$ (clean) to $a_T = 0$ (fully reverberant). For the Δ-normalized parameterization we further define the schedule step size

$$g_t \;=\; a_{t-1} - a_t, \qquad t = 1,\ldots,T, \tag{3}$$

which is used both to update the state in the reverse process (Eq. (6)) and to form the normalized update target in training (Eq. (9)).

B. Reverse Process

Given $\mathbf{x}_T = \mathbf{y}$, we aim to iteratively recover $\mathbf{x}_0$ by learning a reverse transition model $f_\theta(\cdot, t)$ conditioned on the step index t. We explore two reverse parameterizations:

(i) Direct prediction: In this mode, the network predicts the next (less reverberant) state:

$$\widehat{\mathbf{x}}_{t-1} \;=\; f_\theta(\mathbf{x}_t, t), \qquad t = T,\ldots,1, \tag{4}$$

and the reverse trajectory is obtained by iterating Eq. (4) from $\mathbf{x}_T$ down to $\widehat{\mathbf{x}}_0$.

(ii) Δ-normalized residual prediction: In this case, the network predicts a step-size-normalized residual increment between consecutive intermediate states:

$$\widehat{\mathbf{v}}_t \;=\; f_\theta(\mathbf{x}_t, t), \tag{5}$$

which is converted into the next state using the step size g_t from Eq. (3):

$$\widehat{\mathbf{x}}_{t-1} \;=\; \mathbf{x}_t \;+\; g_t\, \widehat{\mathbf{v}}_t. \tag{6}$$

This parameterization predicts a schedule-normalized step update, similar in spirit to step-wise deterministic inversion strategies in diffusion restoration [18], [19]. Figure 1 illustrates the proposed deterministic forward degradation and the learned reverse process featuring representative percussive signals (i.e. snare and kick).

C. Training objective

Training uses paired anechoic/reverberant data $\{(\mathbf{x}_0, \mathbf{y})\}$. For a randomly sampled step $t \in \{1, \ldots, T\}$ we construct $\mathbf{x}_t$ and $\mathbf{x}_{t-1}$ using Eq. (1) and train the network with a weighted sum of a frequency-domain loss and a time-domain loss:

$$\mathcal{L} \;=\; \mathcal{L}_{\text{spec}} \;+\; \lambda_{\text{aud}}\, \mathcal{L}_{\text{aud}}. \tag{7}$$

In *Direct* mode, we supervise the next-state prediction with an ℓ_1 loss,

$$\mathcal{L}_{\text{spec}}^{\text{direct}} \;=\; \left\|\widehat{\mathbf{x}}_{t-1} - \mathbf{x}_{t-1}\right\|_1. \tag{8}$$

Regarding Δ-*normalized* mode, we supervise both (i) the reconstructed next state and (ii) a normalized update target. Using Eq. (3), we define

$$\mathbf{v}_t \;=\; \frac{\mathbf{x}_{t-1} - \mathbf{x}_t}{g_t}, \tag{9}$$

and use a weighted combination of a "delta" term and a "next-state" term:

$$\mathcal{L}_{\text{spec}}^{\Delta} \;=\; 0.7\, \left\|\widehat{\mathbf{v}}_t - \mathbf{v}_t\right\|_1 \;+\; 0.3\, \left\|\widehat{\mathbf{x}}_{t-1} - \mathbf{x}_{t-1}\right\|_1. \tag{10}$$

$\mathcal{L}_{\text{aud}}$ is computed after inverse STFT between the estimated and reference waveforms for step $(t-1)$ (ℓ_1 loss). In practice, we set a larger weight $\lambda_{\text{aud}} = 8$ to emphasize waveform fidelity. Finally, in all experiments we use $T = 16$ diffusion steps, starting inference from $\mathbf{x}_T = \mathbf{y}$ and iterating the learned reverse transitions to obtain $\widehat{\mathbf{x}}_0$. Preliminary experiments showed no consistent benefit from using more than 16 reverse steps, so we retained the smallest effective setting.

D. Backbones

We consider two different backbone network architectures.

UNet: It is based on the NCSN++ architecture commonly used in score-based diffusion models [21]. We only modify the *input projection* convolution layer with a $(9, 1)$ kernel to better match the transient-dominated nature of drum spectra. The network follows a standard encoder–decoder design with four resolution levels and two residual blocks per level, starting from a base width of 64 channels. We include an attention block at the bottleneck to improve global time–frequency context aggregation. This results in a network capacity of roughly 54.6M parameters.

Transformer Diffuser (DiT): We additionally consider a diffusion transformer that operates on patch-tokenized spectrograms [22]. We employ rotary positional embeddings for sequence modeling [27] and inject diffusion-step information through an auxiliary timestep embedding that modulates token representations [28]. We use a stack of 5 layer blocks with 8 attention heads and embedding dimension of 768, resulting in a 57.1M parameters model, and therefore enables a fair comparison with the UNet backbone.

III. Experimental Setup

A. Dataset

We construct a paired anechoic–reverberant dataset using a combination of the MUSDB18-HQ dataset [29] and the Groove MIDI Dataset (GMD) [30]. MUSDB18-HQ provides high-quality (44.1 kHz, stereo) music stems from real recordings, from which we extract downmixed drum stems corresponding to physically recorded drum kits. To complement these with electronically generated material, we additionally include drum performances rendered from GMD, yielding a mixed content with diverse rhythmic and spectral characteristics.

Because both source datasets may contain printed reverberation, room ambience, or production effects, all drum stems were manually curated by an experienced audio engineer. Only stems judged to be perceptually *dry*—i.e., lacking audible room tails, algorithmic reverb, or heavy ambience—were retained. This manual filtering is crucial for dereverberation, where the boundary between the "clean" signal and reverberant components is inherently ambiguous, particularly for percussive instruments. Drum sounds include resonances from the shell and membrane, sympathetic vibrations across kit elements, and early reflections that may be perceived as part of the source rather than the room. Furthermore, common algorithmic and analog reverbs (e.g., plate or spring simulations) introduce modulation and non-linearities (e.g., gating artifacts) that do not correspond to physical room acoustics and are outside the scope of room dereverberation [31]. Our dataset, therefore, targets physically plausible room reverberation, including both early reflections and late reverberant decay, rather than production-style effects.

The selected dry drum mixtures are randomly segmented into 2 s excerpts. To increase variability in timbre and production style, we apply data augmentation using the `audiomentations` library[3], including pitch shifting, time stretching, and random equalization. Reverberant counterparts are generated by convolving the dry segments with room impulse responses (RIRs). Following prior work in speech enhancement [12], we use `pyroomacoustics` [32] to synthesize artificial RIRs with randomized room geometries and reverberation times. To further improve realism and reduce bias toward simulated acoustics, we additionally incorporate

[3]https://github.com/iver56/audiomentations

measured RIRs from the OpenAIR database [33]. For each segment, a synthetic or real RIR is randomly selected. Dry excerpts are normalized before RIR rendering; reverberant counterparts then undergo wet/dry energy control and peak safety. Finally, the resulting dataset comprises approximately 38 hours of stereo audio and is available upon request.

B. Baselines

We compare the proposed cold-diffusion models against two representative diffusion-based baselines originally developed for speech enhancement and dereverberation: **SGMSE+** [12] and **CDiffuSE** [10]. To isolate the effect of the diffusion formulation in a controlled comparison, we retrain both baselines using the same data, input representation (stereo RI spectrograms), backbone architecture (UNet), and training configuration as our models. The only difference lies in the diffusion formulation and training objective prescribed by each method. This reduces confounds from backbone capacity, input features, and optimization settings. Inference step counts, however, follow each method's original configuration and are therefore not matched for computational cost or inference latency: our method uses $T = 16$, SGMSE+ is evaluated with 30 reverse steps using the original predictor-corrector sampler with one-step correction, and CDiffuSE uses 50 reverse steps. The comparison should therefore be interpreted as a standard-configuration comparison rather than a matched-compute benchmark.

C. Training configuration

The dataset is split into 80% / 10% / 10% train, validation, and test sets. All models are trained using the Adam optimizer with a learning rate of 10^{-4}. During training, we maintain an exponential moving average (EMA) of the model parameters with decay 0.995; at inference, we use the EMA weights rather than the raw training weights, as this commonly improves stability and perceptual quality in diffusion models [12], [34].

For the transformer-based diffuser (DiT), we restrict our experiments to the Δ-*normalized* mode. In our preliminary inference experiments, *direct* per-step spectrogram prediction led to unstable reverse diffusion, consistent with the general exposure-bias / error-accumulation phenomenon in diffusion sampling chains [35]. The Δ-normalized formulation mitigates this effect by constraining each reverse step to a normalized, schedule-aware update, resulting in stable and reliable inference for the DiT backbone.

D. Evaluation Metrics

Standard objective measures in dereverberation are largely *speech-oriented* (e.g., PESQ [23], POLQA [24], (E)STOI [25]) and are not well matched to drum downmix dereverberation. We therefore report two complementary groups of metrics:

1) Low-level (signal-based) metrics: We report low-level metrics that quantify waveform fidelity, time–frequency magnitude/phase consistency, and improvement over the reverberant input. Let $\mathbf{y}, \mathbf{x}, \hat{\mathbf{x}} \in \mathbb{R}^{2\times T}$ denote the stereo reverberant input, anechoic reference, and dereverberated estimate, respectively.

(a) **Multi-resolution STFT magnitude MAE ($\mathbf{mSTFT_{mag}}$):** We measure magnitude fidelity using a multi-resolution STFT log-magnitude distance [36]. Multiple resolutions are important for drums because short windows capture sharp transients, while long windows capture decay tails and low-frequency resonances. We average the MAE over $n_{\text{fft}} \in \{256, 1024, 4096, 8192\}$ with hop sizes $\{64, 256, 1024, 2048\}$.

(b) **Multi-resolution STFT phase MAE ($\mathbf{mSTFT_{phase}}$):** We compute phase error on complex STFTs as the mean absolute wrapped phase difference between $\hat{\mathbf{x}}$ and $\mathbf{x}$ (wrapping to $(-\pi, \pi]$). Phase-aware modeling and evaluation has been shown to impact perceived quality in enhancement tasks [37].

(c) **Error-to-Signal Ratio (ESR):** ESR [36] measures the *residual error energy* relative to the target energy, and is sensitive to gain mismatches:

$$\mathrm{ESR}(\hat{\mathbf{x}}, \mathbf{x}) = \frac{\|\hat{\mathbf{x}} - \mathbf{x}\|_2^2}{\|\mathbf{x}\|_2^2 + \epsilon}. \tag{11}$$

The normalization by $\|\mathbf{x}\|_2^2$ enables fair comparison across excerpts with different dynamic ranges.

(d) **SI-SDR improvement (SI-SDRi):** We report *improvement* (SI-SDRi) rather than absolute SI-SDR [38] because dereverberation is a hard condition where absolute SI-SDR can remain low even for perceptually improved outputs (e.g., [11]). We do not report SI-SIR/SI-SAR, whose interference/artifact decomposition is less meaningful for single-source dereverberation.

(e) **Normalized Mutual Information (NMI):** We compute NMI [39] between target and estimated magnitude spectrograms:

$$\mathrm{NMI}(U, V) = \frac{I(U; V)}{\sqrt{H(U)\, H(V)}}, \tag{12}$$

where $I(\cdot;\cdot)$ is mutual information and $H(\cdot)$ entropy. In our setting, NMI captures non-linear similarity and is less sensitive to exact local time-frequency alignment.

2) High-level (perceptual) metrics: Low-level signal-fidelity metrics do not directly quantify perceptually salient aspects of *drums* dereverberation, such as transient clarity, envelope restoration, and preservation of rhythmic event timing. We therefore define the following metrics:

(a) **Modulation Spectrum Distance (MSD):** Reverberation alters temporal-envelope modulations by smearing energy and changing the distribution of modulation frequencies [40]. We compute an auditory-inspired modulation representation from subband Hilbert envelopes and measure the distance between target and estimate. Lower values indicate modulation characteristics closer to the anechoic target.

(b) **Envelope Correlation (ENV):** We compute the Pearson correlation between target and estimated RMS envelopes. This metric is sensitive to transient distortion and tail overhang; higher correlation indicates better preservation of macro-dynamics and event shapes.

TABLE I
IN-DOMAIN TEST SET RESULTS (MEAN ± STD) ON SIGNAL-BASED AND PERCEPTUAL METRICS.

Model	$\text{mSTFT}_{\text{mag}}$ ↓	$\text{mSTFT}_{\text{phase}}$ ↓	ESR ↓	SI-SDRi ↑	NMI ↑	MSD ↓	ENV ↑	TTER ↓	ONFi ↑
SGMSE+	0.12 ± 0.52	1.32 ± 0.29	1.35 ± 0.75	4.06 ± 6.32	0.36 ± 0.18	0.29 ± 0.09	0.62 ± 0.31	5.90 ± 4.51	0.08 ± 0.17
CDiffuSE	0.12 ± 0.49	1.36 ± 0.24	1.37 ± 0.73	2.77 ± 4.01	0.34 ± 0.18	0.30 ± 0.10	0.59 ± 0.31	6.03 ± 4.08	0.04 ± 0.17
Cold UNet Δ-*norm*	$\mathbf{0.08 \pm 0.52}$	$\mathbf{1.21 \pm 0.36}$	$\mathbf{0.79 \pm 0.74}$	$\mathbf{11.09 \pm 10.25}$	$\mathbf{0.55 \pm 0.16}$	$\mathbf{0.22 \pm 0.09}$	$\mathbf{0.92 \pm 0.12}$	$\mathbf{2.07 \pm 2.13}$	$\mathbf{0.16 \pm 0.21}$
Cold DiT Δ-*norm*	0.10 ± 0.52	1.28 ± 0.31	1.05 ± 0.83	7.36 ± 9.10	0.45 ± 0.17	0.25 ± 0.09	0.84 ± 0.19	3.57 ± 3.24	0.07 ± 0.21
Cold UNet *Direct*	0.09 ± 0.52	1.23 ± 0.35	0.88 ± 0.79	9.91 ± 10.13	0.52 ± 0.17	0.23 ± 0.09	0.89 ± 0.16	2.72 ± 3.05	0.14 ± 0.21

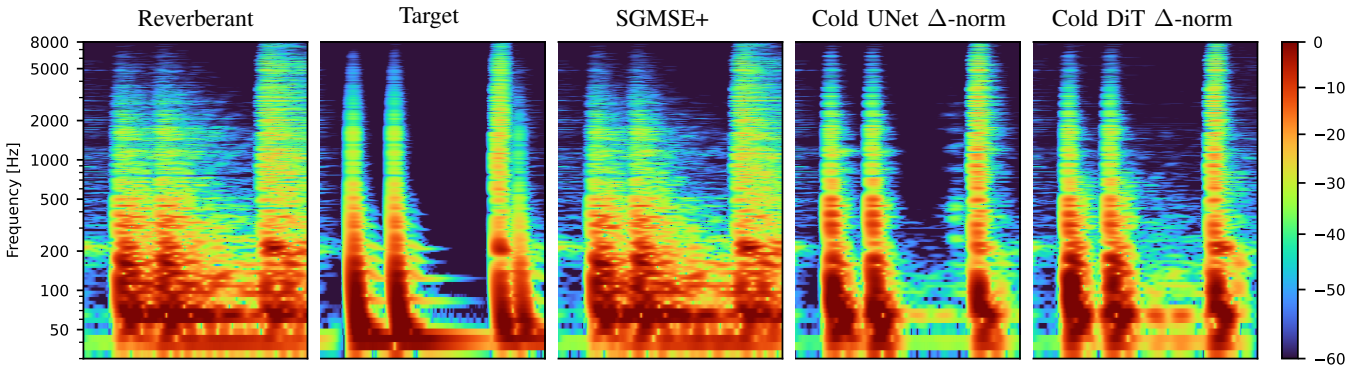

Fig. 2. Qualitative spectrogram comparison on an out-of-domain drum excerpt processed with a highly reverberant unseen impulse response.

(c) **Per-hit Transient-to-Tail Energy Ratio deviation (TTER):** Early-to-late energy balance is related to perceived reverberance [41]. For drums, we propose a local onset-conditioned analogue that measures whether dereverberation restores energy concentration near each hit. We detect onset times $\{t_k\}$ on the clean target using `librosa` [42]. For each onset, we compute transient and tail energies from fixed windows of length T_{tr} and T_{tail}:

$$R(\mathbf{s}; t_k) = 10 \log_{10} \frac{\sum_{t \in [t_k,\, t_k + T_{\text{tr}})} s[t]^2 + \epsilon}{\sum_{t \in [t_k + T_{\text{tr}},\, t_k + T_{\text{tr}} + T_{\text{tail}})} s[t]^2 + \epsilon}. \tag{13}$$

We report the mean absolute deviation of $R(\hat{\mathbf{x}}; t_k)$ from $R(\mathbf{x}; t_k)$ across hits.

(d) **Onset F-measure improvement (ONFi):** Dereverberation should sharpen hit boundaries and reduce spurious detections caused by smeared reverberant tails. We therefore compute an onset-detection F-measure by comparing detected onsets against reference onsets obtained from the clean target, using a fixed temporal tolerance window [43]. We report an *improvement* score as the difference between the estimate and the reverberant input with respect to the same clean reference.

IV. RESULTS

A. In-domain Evaluation

Table I reports in-domain dereverberation performance on the held-out test split. Both speech-oriented diffusion baselines (SGMSE+ and CDiffuse) achieve limited improvement, with similar spectral distortion, high residual energy (ESR), and weak transient–tail separation (TTER), indicating substantial remaining reverberant energy. In contrast, all proposed cold-diffusion models consistently improve both signal-based and perceptual metrics. The best variant, Cold UNet (Δ-norm) achieves the best overall performance, yielding the largest SI-SDR improvement (11.09 dB), the lowest ESR, and the highest envelope correlation, while substantially reducing TTER. These results indicate effective suppression of late reverberation while preserving transient structure. Although the onset F-measure improvement (ONFi) values are numerically modest (best 0.16), this is expected for an improvement in a bounded onset F-measure and still indicates more reliable hit boundary recovery compared to the reverberant input. Finally, Δ-normalized residual inference is consistently stronger than Direct prediction—especially on ESR/TTER/ONFi—supporting the claim that step-normalized updates reduce error accumulation over the iterative reverse process; the DiT backbone benefits similarly but remains below the UNet, particularly on transient-sensitive measures.

B. Out-of-domain Evaluation

To evaluate generalization beyond the training distribution, we construct a fully out-of-domain test set using a curated subset of MoisesDB [44]. Following the same procedure as in the training data, an experienced audio engineer manually selected perceptually dry drum stems, ensuring the absence of existing reverberation. Reverberant counterparts were generated using *only physical impulse responses*, drawn from real-room recordings (ACE Challenge Workshop[4]). Importantly, this setup introduces a complete domain shift: none of the

[4]https://zenodo.org/records/6257551

TABLE II
OUT-OF-DOMAIN DATASET RESULTS (MEAN ± STD) ON SIGNAL-BASED AND PERCEPTUAL METRICS.

Model	$\text{mSTFT}_{\text{mag}}$ ↓	$\text{mSTFT}_{\text{phase}}$ ↓	ESR ↓	SI-SDRi ↑	NMI ↑	MSD ↓	ENV ↑	TTER ↓	ONFi ↑
SGMSE+	0.22 ± 0.10	1.37 ± 0.24	1.42 ± 0.75	2.01 ± 5.08	0.31 ± 0.15	0.32 ± 0.08	0.58 ± 0.30	6.70 ± 4.39	0.05 ± 0.15
CDiffuse	0.23 ± 0.10	1.39 ± 0.21	1.44 ± 0.69	0.17 ± 3.81	0.29 ± 0.16	0.32 ± 0.09	0.55 ± 0.30	6.85 ± 4.22	0.03 ± 0.14
Cold UNet Δ-*norm*	$\mathbf{0.16 \pm 0.08}$	$\mathbf{1.25 \pm 0.33}$	$\mathbf{1.09 \pm 0.88}$	$\mathbf{7.52 \pm 8.61}$	$\mathbf{0.45 \pm 0.17}$	$\mathbf{0.25 \pm 0.08}$	$\mathbf{0.84 \pm 0.17}$	$\mathbf{3.60 \pm 3.35}$	$\mathbf{0.13 \pm 0.20}$
Cold DiT Δ-*norm*	0.17 ± 0.09	1.32 ± 0.27	1.19 ± 0.86	5.59 ± 7.52	0.41 ± 0.16	0.27 ± 0.08	0.79 ± 0.21	4.58 ± 3.90	0.05 ± 0.20
Cold UNet *Direct*	$\mathbf{0.16 \pm 0.08}$	1.26 ± 0.32	1.11 ± 0.89	7.20 ± 8.48	0.44 ± 0.17	$\mathbf{0.25 \pm 0.09}$	0.83 ± 0.19	4.00 ± 3.76	0.11 ± 0.20

drum performances, recording characteristics, or reverberation conditions overlap with the training data.

Results are reported in Table II. As expected, all models exhibit degraded performance compared to the in-domain evaluation, reflecting the difficulty of generalizing dereverberation across unseen drum sources and room acoustics. This degradation is most pronounced for the speech-oriented diffusion baselines, whose SI-SDRi drops to 2.01 dB (SGMSE+) and near zero for CDiffuSE, alongside increased spectral distortion, residual energy (ESR), and transient–tail imbalance (TTER). These results indicate that baseline diffusion objectives fail to generalize when both the source characteristics and reverberation statistics change.

In contrast, the proposed cold-diffusion models retain substantially better performance across all metric groups. The Cold UNet (Δ-norm) remains the strongest model, achieving more than 7.5 dB SI-SDR improvement and consistently lower ESR and TTER compared to all baselines. Perceptual proxies confirm this trend: envelope correlation remains high (ENV 0.84), and modulation distortion (MSD) is significantly reduced relative to baseline methods. These results suggest that the deterministic cold-diffusion formulation provides increased robustness to unseen reverberation conditions, likely because the model learns to invert a structured degradation process.

Compared to the in-domain setting, the performance gap between Δ-normalized and direct prediction for the UNet narrows considerably. While the Δ-normalized variant still performs best overall, both modes yield similar scores across most metrics. This behaviour may suggest that under strong domain shift, modeling capacity and inductive bias dominate over fine-grained reverse-process parameterization, and that the UNet backbone itself is the primary driver of robustness. The DiT again benefits from the Δ-normalized formulation but underperforms the UNet, particularly on transient-sensitive metrics (TTER, ONFi).

Figure 2 shows a qualitative comparison on an out-of-domain electronic drum excerpt processed with a highly reverberant impulse response ($T_{60} > 2\,\text{s}$). The reverberant input exhibits strong energy smearing across time and a pronounced low-frequency bed that masks transient structure. SGMSE+ shows almost no improvement, leaving substantial residual energy and blurred transients. In contrast, both proposed cold-diffusion models substantially reduce late reverberation while preserving sharp onsets; UNet yields the cleanest decay and most compact transient structure. Still, this example also highlights a representative failure case under unseen conditions: UNet partially attenuates low-frequency content, while DiT leaves more residual tail energy. These artifacts suggest that highly reverberant unseen conditions remain challenging despite the overall gains.

V. CONCLUSIONS

We introduced a cold-diffusion framework for stereo drum-stem dereverberation, addressing a largely unexplored percussive enhancement task beyond the speech-centric focus of prior work. By modeling reverberation as a deterministic degradation and evaluating both Direct and Δ-normalized residual reverse parameterizations, our approach consistently outperformed strong diffusion-based baselines on both in-domain and out-of-domain datasets, with improved transient preservation and reduced late reverberation. Future work will focus on expanding the availability of clean percussive data and extending the framework to better handle production-style reverbs, whose non-physical characteristics differ substantially from impulse-response-based room acoustics. We also plan to explore more flexible reverse-diffusion settings, including alternative schedules and variable numbers of reverse steps, to improve robustness and reduce dependence on a fixed inference configuration. Additionally, we will explore using selected metrics from our evaluation suite as potential candidates for auxiliary training losses, aiming to better align optimization with perceptual and transient-sensitive quality.

ACKNOWLEDGMENT

This publication is financed by the Project "Strengthening and optimizing the operation of MODY services and academic and research units of the Hellenic Mediterranean University", funded by the Public Investment Program of the Greek Ministry of Education and Religious Affairs.

The authors gratefully acknowledge XLN Audio for early support and for discussions that helped shape the initial development and motivation of this work.

REFERENCES

[1] R. C. Hendriks, T. Gerkmann, and J. Jensen, *DFT-domain Based Single-microphone Noise Reduction for Speech Enhancement: A Survey of the State-of-the-art*. Morgan & Claypool Publishers, 2013, vol. 11.

[2] K. Kinoshita, M. Delcroix, T. Yoshioka, T. Nakatani, E. Habets, R. Haeb-Umbach, V. Leutnant, A. Sehr, W. Kellermann, R. Maas *et al.*, "The reverb challenge: A common evaluation framework for dereverberation and recognition of reverberant speech," in *2013 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics*. IEEE, 2013, pp. 1–4.

[3] D. Wang and J. Chen, "Supervised speech separation based on deep learning: An overview," *IEEE/ACM transactions on audio, speech, and language processing*, vol. 26, no. 10, pp. 1702–1726, 2018.

[4] T. Nakatani, T. Yoshioka, K. Kinoshita, M. Miyoshi, and B.-H. Juang, "Speech dereverberation based on variance-normalized delayed linear prediction," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 18, no. 7, pp. 1717–1731, 2010.

[5] O. Ernst, S. E. Chazan, S. Gannot, and J. Goldberger, "Speech dereverberation using fully convolutional networks," in *2018 26th European Signal Processing Conference (EUSIPCO)*. IEEE, 2018, pp. 390–394.

[6] Z.-Q. Wang and D. Wang, "Deep learning based target cancellation for speech dereverberation," *IEEE/ACM transactions on audio, speech, and language processing*, vol. 28, pp. 941–950, 2020.

[7] S.-W. Fu, C.-F. Liao, Y. Tsao, and S.-D. Lin, "Metricgan: Generative adversarial networks based black-box metric scores optimization for speech enhancement," in *International Conference on Machine Learning*. PmLR, 2019, pp. 2031–2041.

[8] J. Su, Z. Jin, and A. Finkelstein, "Hifi-gan: High-fidelity denoising and dereverberation based on speech deep features in adversarial networks," *arXiv preprint arXiv:2006.05694*, 2020.

[9] Y.-J. Lu, Y. Tsao, and S. Watanabe, "A study on speech enhancement based on diffusion probabilistic model," in *2021 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*. IEEE, 2021, pp. 659–666.

[10] Y.-J. Lu, Z.-Q. Wang, S. Watanabe, A. Richard, C. Yu, and Y. Tsao, "Conditional diffusion probabilistic model for speech enhancement," in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. Ieee, 2022, pp. 7402–7406.

[11] S. Welker, J. Richter, and T. Gerkmann, "Speech enhancement with score-based generative models in the complex stft domain," *arXiv preprint arXiv:2203.17004*, 2022.

[12] J. Richter, S. Welker, J.-M. Lemercier, B. Lay, and T. Gerkmann, "Speech enhancement and dereverberation with diffusion-based generative models," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 31, pp. 2351–2364, 2023.

[13] J.-M. Lemercier, J. Richter, S. Welker, E. Moliner, V. Välimäki, and T. Gerkmann, "Diffusion models for audio restoration," *arXiv preprint arXiv:2402.09821*, 2024.

[14] N. Yasuraoka, T. Yoshioka, T. Nakatani, A. Nakamura, and H. G. Okuno, "Music dereverberation using harmonic structure source model and wiener filter," in *2010 IEEE International Conference on Acoustics, Speech and Signal Processing*. IEEE, 2010, pp. 53–56.

[15] K. Saito, N. Murata, T. Uesaka, C.-H. Lai, Y. Takida, T. Fukui, and Y. Mitsufuji, "Unsupervised vocal dereverberation with diffusion-based generative models," in *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2023, pp. 1–5.

[16] T. Wilmering, G. Fazekas, and M. Sandler, "The effects of reverberation on onset detection tasks," in *Audio Engineering Society Convention 128*. Audio Engineering Society, 2010.

[17] G. Grindlay, "Blind dereverberation of audio signals," *E4810 Final Project, University of Columbia*, 2008.

[18] A. Bansal, E. Borgnia, H.-M. Chu, J. Li, H. Kazemi, F. Huang, M. Goldblum, J. Geiping, and T. Goldstein, "Cold diffusion: Inverting arbitrary image transforms without noise," *Advances in Neural Information Processing Systems*, vol. 36, pp. 41 259–41 282, 2023.

[19] H. Yen, F. G. Germain, G. Wichern, and J. Le Roux, "Cold diffusion for speech enhancement," in *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2023, pp. 1–5.

[20] G. Plaja-Roglans, M. Miron, A. Shankar, and X. Serra, "Carnatic singing voice separation using cold diffusion on training data with bleeding," 2023.

[21] Y. Song, J. Sohl-Dickstein, D. P. Kingma, A. Kumar, S. Ermon, and B. Poole, "Score-based generative modeling through stochastic differential equations," *arXiv preprint arXiv:2011.13456*, 2020.

[22] W. Peebles and S. Xie, "Scalable diffusion models with transformers," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2023, pp. 4195–4205.

[23] A. W. Rix, J. G. Beerends, M. P. Hollier, and A. P. Hekstra, "Perceptual evaluation of speech quality (pesq)-a new method for speech quality assessment of telephone networks and codecs," in *2001 IEEE international conference on acoustics, speech, and signal processing. Proceedings (Cat. No. 01CH37221)*, vol. 2. IEEE, 2001, pp. 749–752.

[24] ITU-T, "Itu-t rec. p.863, "perceptual objective listening quality prediction,"," Int. Telecom. Union (ITU), Tech. Rep., 2018, [Online]. Available: https://www.itu.int/rec/T-REC-P.863-201803-I/en. [Online]. Available: https://www.itu.int/rec/T-REC-P.863-201803-I/en

[25] J. Jensen and C. H. Taal, "An algorithm for predicting the intelligibility of speech masked by modulated noise maskers," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 24, no. 11, pp. 2009–2022, 2016.

[26] M. Torcoli, T. Kastner, and J. Herre, "Objective measures of perceptual audio quality reviewed: An evaluation of their application domain dependence," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 1530–1541, 2021.

[27] J. Su, M. Ahmed, Y. Lu, S. Pan, W. Bo, and Y. Liu, "Roformer: Enhanced transformer with rotary position embedding," *Neurocomputing*, vol. 568, p. 127063, 2024.

[28] H. R. Guimarães, J. Su, R. Kumar, T. H. Falk, and Z. Jin, "Ditse: High-fidelity generative speech enhancement via latent diffusion transformers," *arXiv preprint arXiv:2504.09381*, 2025.

[29] Z. Rafii, A. Liutkus, F.-R. Stöter, S. I. Mimilakis, and R. Bittner, "Musdb18-hq-an uncompressed version of musdb18," *(No Title)*, 2019.

[30] J. Gillick, A. Roberts, J. Engel, D. Eck, and D. Bamman, "Learning to groove with inverse sequence transformations," in *International conference on machine learning*. PMLR, 2019, pp. 2269–2279.

[31] N. H. Fletcher and T. D. Rossing, *The physics of musical instruments*. Springer Science & Business Media, 2012.

[32] R. Scheibler, E. Bezzam, and I. Dokmanić, "Pyroomacoustics: A python package for audio room simulation and array processing algorithms," in *2018 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2018, pp. 351–355.

[33] S. Shelley and D. T. Murphy, "Openair: An interactive auralization web resource and database," in *129th Audio Engineering Society Convention 2010*, 2010, pp. 1270–1278.

[34] J.-M. Lemercier, J. Richter, S. Welker, and T. Gerkmann, "Storm: A diffusion-based stochastic regeneration model for speech enhancement and dereverberation," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 31, pp. 2724–2737, 2023.

[35] M. Ning, E. Sangineto, A. Porrello, S. Calderara, and R. Cucchiara, "Input perturbation reduces exposure bias in diffusion models," *arXiv preprint arXiv:2301.11706*, 2023.

[36] C. J. Steinmetz and J. D. Reiss, "auraloss: Audio focused loss functions in pytorch," in *Digital music research network one-day workshop (DMRN+ 15)*, 2020, p. 124.

[37] K. Tan and D. Wang, "Learning complex spectral mapping with gated convolutional recurrent networks for monaural speech enhancement," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 28, pp. 380–390, 2019.

[38] J. Le Roux, S. Wisdom, H. Erdogan, and J. R. Hershey, "Sdr–half-baked or well done?" in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 626–630.

[39] J. Foote, "A similarity measure for automatic audio classification," in *Proc. AAAI 1997 Spring Symposium on Intelligent Integration and Use of Text, Image, Video, and Audio Corpora*, vol. 3, 1997.

[40] T. H. Falk, C. Zheng, and W.-Y. Chan, "A non-intrusive quality and intelligibility measure of reverberant and dereverberated speech," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 18, no. 7, pp. 1766–1774, 2010.

[41] E. Larsen, N. Iyer, C. R. Lansing, and A. S. Feng, "On the minimum audible difference in direct-to-reverberant energy ratio," *The Journal of the Acoustical Society of America*, vol. 124, no. 1, pp. 450–461, 2008.

[42] B. McFee, C. Raffel, D. Liang, D. P. Ellis, M. McVicar, E. Battenberg, and O. Nieto, "librosa: Audio and music signal analysis in python." *SciPy*, vol. 2015, pp. 18–24, 2015.

[43] C. Raffel, B. McFee, E. J. Humphrey, J. Salamon, O. Nieto, D. Liang, D. P. Ellis, and C. C. Raffel, "Mir_eval: A transparent implementation of common mir metrics." in *ISMIR*, vol. 10, 2014, p. 2014.

[44] I. Pereira, F. Araújo, F. Korzeniowski, and R. Vogl, "Moisesdb: A dataset for source separation beyond 4-stems," *arXiv preprint arXiv:2307.15913*, 2023.