



Few-step Adversarial Schrödinger Bridge for Generative Speech Enhancement

Seungu Han^{†1}, Sungho Lee¹, Juheon Lee², Kyogu Lee^{1,3}

¹Department of Intelligence and Information, Seoul National University, Republic of Korea

²Supertone Inc., Republic of Korea

³Artificial Intelligence Institute, Seoul National University, Republic of Korea

hansw0326@snu.ac.kr, sh-lee@snu.ac.kr, juheon@supertone.ai, kglee@snu.ac.kr

Abstract

Deep generative models have recently been employed for speech enhancement to generate perceptually valid clean speech on large-scale datasets. Several diffusion models have been proposed, and more recently, a tractable Schrödinger Bridge has been introduced to transport between the clean and noisy speech distributions. However, these models often suffer from an iterative reverse process and require a large number of sampling steps—more than 50. Our investigation reveals that the performance of baseline models significantly degrades when the number of sampling steps is reduced, particularly under low-SNR conditions. We propose integrating Schrödinger Bridge with GANs to effectively mitigate this issue, achieving high-quality outputs on full-band datasets while substantially reducing the required sampling steps. Experimental results demonstrate that our proposed model outperforms existing baselines, even with a single inference step, in both denoising and dereverberation tasks.

Index Terms: speech enhancement, speech denoising, speech dereverberation, Schrödinger bridge, Diffusion-GAN Hybrids

1. Introduction

Deep generative models have been employed for various speech enhancement and restoration tasks [1–6]. This approach is motivated by the inherent nature of the problem: under severe conditions, such as low signal-to-noise ratio (SNR) scenarios, some information is lost, making it necessary to generate the missing parts [7]. Furthermore, unlike predictive models, which map inputs to a single deterministic output, generative models learn the underlying data distribution, allowing them to generate multiple valid estimates for a given input. This ability enhances their performance on non-intrusive metrics [8], making them valuable for real-world applications. Following the recent success in various domains, diffusion models or score-based generative models (SGMs) [9], have been proposed for speech enhancement [1–3]. Diffusion models aim to learn the reverse process that transports the tractable prior distribution, typically Gaussian, into the data distribution. To better address the speech enhancement task, [1, 3] proposed conditional diffusion models, in which the reverse process begins with a mixture of noisy speech and Gaussian noise. However, these models suffer from a mismatch in the mean of the prior. To resolve this, the Schrödinger Bridge (SB) framework [10–13], whose goal is to find the path measure between one probability distribution and another, has been introduced as a natural solution. It enables the generative process to start directly from the noisy input [4–6]. However, solving the general SB requires executing costly it-

erative procedures for simulating stochastic processes [11–13], which limits its scalability and applicability. On the other hand, in the special case where paired data is available, SB can be solved using a simulation-free approach [14]. Several studies have built on this, exploring its incorporation into speech-related tasks, such as text-to-speech [15] and audio super-resolution [16]. Furthermore, recent studies have demonstrated that SB outperforms diffusion models in speech enhancement for both wide-band [4–6] and full-band [17] scenarios.

Diffusion models and the SB approach achieve high-quality results but suffer from slow inference speeds. This is due to the need for discretizing time into small step size to formulate each step of the reverse process as a Gaussian form. On the other hand, [4] reported that the SB model is robust to the number of steps used in the reverse process for speech enhancement on wide-band data. However, our experiments reveal that this robustness does not always hold for full-band data, where a large number of iterations is still required, particularly in low-SNR cases. Meanwhile, studies such as [18–20] have explored combining Generative Adversarial Networks (GANs) with diffusion models, achieving both high-quality outputs and reduced sampling steps through adversarial training to match the denoising distribution. Building on these advancements, we propose a novel approach that integrates the distribution matching GANs into the SB framework for speech enhancement. Our model not only avoids the prior mismatch, just like the original SB models, but also enables inference with as few as one to four steps. We evaluated our approach on full-band datasets, specifically using EARS-WHAM for denoising and EARS-Reverb for dereverberation [21]. Experimental results demonstrate that our model outperforms 50-step SB model with only a single-step inference of our model. Notably, our model achieves a large margin of improvement over the SB baselines under low-SNR conditions. Audio samples are available at <https://bit.ly/4i06Pgm>.

2. Background

2.1. Schrödinger Bridge for Speech Enhancement

The Schrödinger Bridge (SB) problem [10] considers the optimization of path measures with boundary conditions:

$$\min_{\mathcal{Q} \in \mathcal{P}_{[p_0, p_T]}} D_{\text{KL}}(\mathcal{Q} || \mathcal{P}). \quad (1)$$

Here, \mathcal{P} denotes a reference measure, which is generally set to the path measure of the forward stochastic differential equation (SDE): $d\mathbf{x}_t = f(t, \mathbf{x}_t)dt + g(t)d\mathbf{w}_t$, $\mathbf{x}_0 \sim p_0$ as SGMs [9]. The drift term is denoted by f , the diffusion term by g , and $\mathcal{P}_{[p_0, p_T]}$ is a set of path measures with marginal densities p_0 and p_T at $t = 0$ and $t = T$, respectively. Then, the solution

[†] Work partially done at Supertone Inc.

to the Eq. 1 can be expressed by the path measure of a pair of forward and reverse SDEs [12]:

$$d\mathbf{x}_t = [f(t, \mathbf{x}_t) + g^2(t)\nabla_{\mathbf{x}_t} \log \Psi(t, \mathbf{x}_t)]dt + g(t) d\mathbf{w}_t \quad (2a)$$

$$d\bar{\mathbf{x}}_t = [f(t, \mathbf{x}_t) - g^2(t)\nabla_{\mathbf{x}_t} \log \bar{\Psi}(t, \mathbf{x}_t)]dt + g(t) d\bar{\mathbf{w}}_t \quad (2b)$$

where $\mathbf{x}_0 \sim p_0$ and $\mathbf{x}_T \sim p_T$. While f and g are the same as in the reference SDEs, the additional nonlinear drift terms, $\nabla_{\mathbf{x}_t} \log \Psi(t, \mathbf{x}_t)$ and $\nabla_{\mathbf{x}_t} \log \bar{\Psi}(t, \mathbf{x}_t)$, are introduced. The functions Ψ and $\bar{\Psi}$ are time-varying energy potentials governed by coupled partial differential equations (PDEs):

$$\frac{\partial \Psi}{\partial t} = -\nabla_{\mathbf{x}} \Psi^\top f - \frac{1}{2} \text{Tr} (g^2 \nabla_{\mathbf{x}}^2 \Psi), \quad (3a)$$

$$\frac{\partial \bar{\Psi}}{\partial t} = -\nabla_{\mathbf{x}} \cdot (\bar{\Psi} f) + \frac{1}{2} \text{Tr} (g^2 \nabla_{\mathbf{x}}^2 \bar{\Psi}), \quad (3b)$$

subject to the conditions $\Psi_0 \bar{\Psi}_0 = p_0$ and $\Psi_T \bar{\Psi}_T = p_T$. The marginal of the SB state \mathbf{x}_t is then given as $p_t = \Psi_t \bar{\Psi}_t$ [12]. SB allows transportation between two arbitrary distributions beyond Gaussian priors. Thus, it is natural to formulate the speech enhancement task as an SB problem, where the goal is to transport between noisy and clean data distributions.

While the SB problem typically has no analytic solution, closed-form solutions exist for specific cases, such as when Gaussian boundary conditions are employed [14, 15]. This formulation can be adapted to avoid the prior mismatch encountered in SGMs for speech enhancement. For that, [4, 6] assume a linear drift $f(\mathbf{x}_t) = f(t)\mathbf{x}_t$ and Gaussian boundary conditions $p_0 = \mathcal{N}(\mathbf{x}; \mathbf{x}_0, \epsilon_0^2 \mathbf{I})$ and $p_T = \mathcal{N}(\mathbf{x}; \mathbf{y}, \epsilon_T^2 \mathbf{I})$ with $\epsilon_T = \exp(\int_0^T f(\tau) d\tau) \cdot \epsilon_0$. For $\epsilon_0 \rightarrow 0$, $\bar{\Psi}_t, \Psi_t$ converge to the tractable solution between clean data \mathbf{x}_0 and noisy data \mathbf{y} :

$$\bar{\Psi}_t = \mathcal{N}(\alpha_t \mathbf{x}_0, \alpha_t^2 \sigma_t^2 \mathbf{I}), \quad \Psi_t = \mathcal{N}(\bar{\alpha}_t \mathbf{y}, \bar{\alpha}_t^2 \bar{\sigma}_t^2 \mathbf{I}) \quad (4)$$

where the parameters are given as $\alpha_t = \exp(\int_0^t f(\tau) d\tau)$, $\sigma_t = \int_0^t g^2(\tau) / \alpha_\tau^2 d\tau$, $\bar{\alpha}_t = \alpha_t \alpha_T^{-1}$ and $\bar{\sigma}_t^2 = \sigma_T^2 - \sigma_t^2$. Therefore, the marginal distributions of \mathbf{x}_t can be derived, and they follow Gaussian distributions, given as:

$$q(\mathbf{x}_t | \mathbf{x}_0, \mathbf{y}) = \mathcal{N}(\mathbf{x}_t; w_{\mathbf{x}}(t)\mathbf{x}_0 + w_{\mathbf{y}}(t)\mathbf{y}, \sigma_x^2(t)\mathbf{I}), \quad (5)$$

where the parameters are defined as

$$\sigma_x^2(t) = \frac{\alpha_t^2 \bar{\sigma}_t^2 \sigma_t^2}{\sigma_T^2}, \quad w_{\mathbf{x}}(t) = \frac{\alpha_t \bar{\sigma}_t^2}{\sigma_T^2}, \quad w_{\mathbf{y}}(t) = \frac{\bar{\alpha}_t \sigma_t^2}{\sigma_T^2}. \quad (6)$$

Here, $w_{\mathbf{x}}(t)$ and $w_{\mathbf{y}}(t)$ are the weights for \mathbf{x}_0 and \mathbf{y} , respectively, while $\sigma_x^2(t)$ is the variance of the Gaussian distribution at time t . Based on this design, the intermediate step can be sampled directly, and the same denoising objective as in the diffusion model can be applied. Moreover, the inference process starts with a noisy sample and iteratively follows the reverse SDE or the probability flow ordinary differential equation (ODE) to produce a clean speech estimate [4, 6].

2.2. Diffusion-GAN Hybrids

Diffusion models require small step sizes, resulting in a large number of iterations to maintain the Gaussian assumption on the denoising distribution. In contrast, when the denoising step is large and the data distribution is non-Gaussian, the Gaussian assumption may not hold, resulting in a more complex distribution. To address this, DDGAN [18] was the first to propose combining adversarial training with diffusion models, enabling

denoising with larger discretization sizes and reducing the number of sampling steps. Continuous processes are discretized by selecting intermediate time points as $t_n = n/N$ with fixed N , $t_0 = \epsilon$ and $t_N = T = 1$. They leverage conditional GANs to optimize the divergence between the complex denoising distribution $q(\mathbf{x}_{t_{n-1}} | \mathbf{x}_{t_n})$ with $p_\theta(\mathbf{x}'_{t_{n-1}} | \mathbf{x}_{t_n})$:

$$\min_{\theta} \mathbb{E}_{q(\mathbf{x}_{t_n})} [D_{\text{adv}}(q(\mathbf{x}_{t_{n-1}} | \mathbf{x}_{t_n}) || p_\theta(\mathbf{x}'_{t_{n-1}} | \mathbf{x}_{t_n}))]. \quad (7)$$

On the other hand, SIDDM [19] has been introduced to address scalability issues observed with DDGAN on large-scale datasets. The objective in Eq. 7 can be reformulated as matching the joint distributions $q(\mathbf{x}_{t_{n-1}}, \mathbf{x}_{t_n})$ and $p_\theta(\mathbf{x}'_{t_{n-1}}, \mathbf{x}_{t_n})$, since DDGAN indirectly aligns the conditional distributions $q(\mathbf{x}_{t_{n-1}} | \mathbf{x}_{t_n})$ and $p_\theta(\mathbf{x}'_{t_{n-1}} | \mathbf{x}_{t_n})$ through simple concatenation. Furthermore, the joint distributions can be factorized in the reverse direction as $q(\mathbf{x}_{t_{n-1}}, \mathbf{x}_{t_n}) = q(\mathbf{x}_{t_n} | \mathbf{x}_{t_{n-1}}) q(\mathbf{x}_{t_{n-1}})$ and $p_\theta(\mathbf{x}'_{t_{n-1}}, \mathbf{x}_{t_n}) = p_\theta(\mathbf{x}_{t_n} | \mathbf{x}'_{t_{n-1}}) p_\theta(\mathbf{x}'_{t_{n-1}})$. Thus, joint distribution matching can be formulated as a combination of matching marginal distributions using adversarial divergence and matching conditional distributions using KL divergence:

$$\min_{\theta} \mathbb{E}_{q(\mathbf{x}_{t_n})} [D_{\text{adv}}(q(\mathbf{x}_{t_{n-1}}) | p_\theta(\mathbf{x}'_{t_{n-1}})) + \lambda_{\text{KL}} D_{\text{KL}}(p_\theta(\mathbf{x}_{t_n} | \mathbf{x}'_{t_{n-1}}) || q(\mathbf{x}_{t_n} | \mathbf{x}_{t_{n-1}}))]. \quad (8)$$

This factorization enables improved performance on complex data distributions. Unlike the loss of DDGAN in Eq. 7, where the discriminator incorporates \mathbf{x}_{t_n} , this approach excludes \mathbf{x}_{t_n} for the discriminator's input.

In Diffusion-GAN hybrids, the generator produces a sample of $\mathbf{x}_{t_{n-1}}$. Instead of directly producing $\mathbf{x}_{t_{n-1}}$ from the generator G_θ , DDGAN and SIDDM parameterize $p_\theta(\mathbf{x}'_{t_{n-1}} | \mathbf{x}_{t_n})$ with posterior sampling of $q(\mathbf{x}_{t_{n-1}} | \mathbf{x}_{t_n}, \mathbf{x}'_0 = G_\theta(\mathbf{x}_{t_n}, t_n))$. Moreover, a modified parameterization is proposed by UFOGen [20], where $p_\theta(\mathbf{x}'_{t_{n-1}}) = q(\mathbf{x}_{t_{n-1}} | \mathbf{x}'_0 = G_\theta(\mathbf{x}_{t_n}, t_n))$, based on the forward process instead of the posterior. Starting with the objective of SIDDM in Eq. 8, this parameterization enables the model to be optimized, incorporating a simple reconstruction loss at \mathbf{x}_0 :

$$\min_{\theta} \mathbb{E}_{q(\mathbf{x}_{t_n})} [D_{\text{adv}}(q(\mathbf{x}_{t_{n-1}}) | p_\theta(\mathbf{x}'_{t_{n-1}})) + \lambda_{\text{KL}} \|\mathbf{x}'_0 - \mathbf{x}_0\|^2]. \quad (9)$$

They hypothesize that SIDDM struggles with single-step generation due to the noise variance when sampling \mathbf{x}_{t-1} with \mathbf{x}_0 , whereas UFOGen enables successful one-step sampling by directly matching \mathbf{x}_0 in the reconstruction term. Note that the final loss can also be interpreted as training a diffusion model with adversarial refinement.

3. Methods

3.1. Adversarial Schrödinger Bridge

We combine the Schrödinger Bridge (SB) for speech enhancement with GANs. Since the SB is defined by a pair of SDE, it inherently satisfies the Markov property. Thus, the discrete version is obtained by discretizing the continuous process and then adapting it into the Diffusion-GAN Hybrids, in the same way as the usual SGMs. Furthermore, since noisy data \mathbf{y} can serve as a condition in the speech enhancement task, the reverse process can be formulated with this condition at each step. In this case, the true denoising distribution $q(\mathbf{x}_{t_{n-1}} | \mathbf{x}_{t_n}, \mathbf{y})$ is matched with $p_\theta(\mathbf{x}'_{t_{n-1}} | \mathbf{x}_{t_n}, \mathbf{y})$, where each distribution is conditioned on \mathbf{y} . UFOGen can be naturally extended to incorporate condition,

Algorithm 1 UFOGen Training**Require:** Generator G_θ , Discriminator D_ϕ

- 1: **repeat**
- 2: $\mathbf{x}_0 \sim q(\mathbf{x}_0), n \sim \text{Uniform}(1, \dots, N), t_n = n/N$.
- 3: $\mathbf{x}_{t_{n-1}} \sim q(\mathbf{x}_{t_{n-1}} | \mathbf{x}_0, \mathbf{y}), \mathbf{x}_{t_n} \sim q(\mathbf{x}_{t_n} | \mathbf{x}_{t_{n-1}}, \mathbf{y})$.
- 4: $\mathbf{x}'_{t_{n-1}} \sim q(\mathbf{x}_{t_{n-1}} | \mathbf{x}'_0, \mathbf{y})$ where $\mathbf{x}'_0 = G_\theta(\mathbf{x}_{t_n}, \mathbf{y}, t_n)$.
- 5: Update D_ϕ with gradient:

$$\nabla_\phi \left(-\log(D_\phi(\mathbf{x}_{t_{n-1}}, \mathbf{y}, t_{n-1})) \right. \\ \left. - \log(1 - D_\phi(\mathbf{x}'_{t_{n-1}}, \mathbf{y}, t_{n-1})) \right).$$
- 6: Update G_θ with gradient:

$$\nabla_\theta \left(-\log(D_\phi(\mathbf{x}'_{t_{n-1}}, \mathbf{y}, t_{n-1})) + \lambda_{\text{recon}} L_{\text{recon}} \right).$$
- 7: **until** converged

ensuring that the noisy input is consistently integrated into the model for speech enhancement. The conditional transition probability is also Gaussian and is expressed as:

$$q(\mathbf{x}_{t_n} | \mathbf{x}_{t_{n-1}}, \mathbf{y}) = \mathcal{N} \left(\mathbf{x}_{t_n}; \boldsymbol{\mu}_{t_n | t_{n-1}}, \sigma_{t_n | t_{n-1}}^2 \mathbf{I} \right), \quad (10)$$

where the mean $\boldsymbol{\mu}_{t_n | t_{n-1}}$ and variance $\sigma_{t_n | t_{n-1}}^2$ are given by:

$$\boldsymbol{\mu}_{t_n | t_{n-1}} = \frac{w_{\mathbf{x}}(t_n) \mathbf{x}_{t_{n-1}}}{w_{\mathbf{x}}(t_{n-1})} + \left(w_{\mathbf{y}}(t_n) - \frac{w_{\mathbf{x}}(t_n) w_{\mathbf{y}}(t_{n-1})}{w_{\mathbf{x}}(t_{n-1})} \right) \mathbf{y}, \quad (11a)$$

$$\sigma_{t_n | t_{n-1}}^2 = \sigma_{t_n}^2 - \left(\frac{w_{\mathbf{x}}(t_n)}{w_{\mathbf{x}}(t_{n-1})} \right)^2 \sigma_{t_{n-1}}^2 \quad (11b)$$

based on the mean and variance defined at each step in Eq. 6. The training process of our model is outlined in Algorithm 1 following the UFOGen scheme. The model includes a generator G_θ and a discriminator D_ϕ , conditioned on noisy data \mathbf{y} , and trained within an adversarial framework. First, at a given time t_n , sample $\mathbf{x}_{t_{n-1}}$ using the marginal distribution in Eq. 5 and then sample \mathbf{x}_{t_n} using the conditional transition probability in Eq. 10. Then, the generator produces a denoised estimate \mathbf{x}'_0 and sample $\mathbf{x}'_{t_{n-1}}$ through $q(\mathbf{x}_{t_{n-1}} | \mathbf{x}'_0, \mathbf{y})$. The discriminator and generator are updated using an adversarial loss to match intermediate data at time t_{n-1} and a reconstruction loss. Instead of relying solely on the l_2 -norm, we incorporate an auxiliary loss for the reconstruction loss, as done in [4, 6]. Note that this framework can be applied to other diffusion processes that can be defined as a Markov chain, such as the Ornstein-Uhlenbeck Variance Exploding (OUVE) process [3]. While Adversarial Schrödinger Bridge Matching (ASBM) [22] requires an explicit procedure for solving the Schrödinger Bridge problem, our model does not, as it relies on a simulation-free framework.

3.2. Implementation Details

For the noise schedule, the Variance Exploding (VE) schedule is chosen, following [4, 6], given by $g^2(t) = ck^{2t}$ and $f(t) = 0$. This results in $\alpha_t = 1$ and $\sigma_t^2 = c(k^{2t} - 1)/2 \log k$. The recommended hyperparameters from [4] are employed for the parameterization of the diffusion process. We select $N = 4$ for the discretization, as [18–20] reported it to yield the best results.

The generator is based on the Noise Conditional Score Network (NCSN++) architecture. It utilizes power-compressed complex spectrograms as the input representation and uses the same model hyperparameters and STFT parameterization as in [21], consistent with the baseline models.

We employ a multi-scale short-time Fourier transform (MS-STFT) discriminator with 5 different scales [23]. The FFT sizes

Table 1: *Denoising results on EARS-WHAM.*

Model	NFEs	PESQ	SISDR	ESTOI	DNSMOS	SIGMOS
Unprocessed		1.23	5.88	0.49	2.76	1.94
Conv-TasNet [33]	1	2.31	16.9	0.70	3.47	2.69
CDiffuSE [1]	6	1.60	8.35	0.53	2.87	2.08
Demucs [34]	1	2.37	16.9	0.71	3.66	2.87
SGMSE+ [3]	60	2.49	16.8	0.73	3.91	3.42
SB [17]	50	2.32	17.9	0.73	3.87	3.44
	4	2.32	16.3	0.72	3.84	3.34
	1	2.20	13.4	0.67	3.68	3.15
SB-UFOGen	4	2.52	17.7	0.74	3.88	3.49
	1	2.56	17.9	0.74	3.88	3.48

and window lengths are set to [4096, 2048, 1024, 512, 256], and their hop lengths are given as [1024, 512, 256, 128, 64]. We follow [24] in applying the diffusion process in the time domain, instead of the non-linear spectrogram domain used in [3–6], so our discriminator operates on white Gaussian noise.

For denoising, we use the same reconstruction loss as in [4, 6, 17], given by $L_{\text{recon}} = \|\mathbf{X}'_0 - \mathbf{X}_0\|_2^2 + \alpha \|\mathbf{x}'_0 - \mathbf{x}_0\|_1$, where \mathbf{X}'_0 is the estimated power-compressed complex spectrogram output by the generator, and \mathbf{x}'_0 is the corresponding time-domain signal obtained through inverse STFT. \mathbf{X}_0 and \mathbf{x}_0 are the corresponding targets. For dereverberation, since the model does not train effectively with simple losses alone, we additionally incorporate the multi-scale mel loss, following the setup in [25]. The weighting coefficients for reconstruction loss are set as follows: $\alpha = 10^{-3}$, 0.01 for the additional multi-scale mel loss for dereverberation, and $\lambda_{\text{recon}} = 100$.

4. Evaluation**4.1. Experimental Setup**

We trained our model with a batch size of 8, and the optimizer was AdamW [26] with a learning rate of 10^{-4} . We use the exponential moving average (EMA) of the weights with a decay of 0.999 [27]. The average SI-SDR of 20 randomly selected validation examples is logged during training, and the best-performing model is chosen for evaluation. All models are trained on 4 NVIDIA RTX 4090 or 2 A6000 GPUs.

4.2. Denoising Results

We trained and tested the models on the full-band EARS-WHAM dataset [21] for the denoising task. We used the ODE sampler as it has demonstrated better performance for the denoising task [4]. Results were evaluated using intrusive metrics: Perceptual Evaluation of Speech Quality (PESQ) [28], Scale-Invariant Signal-to-Distortion Ratio (SI-SDR) [29], and Extended Short-Time Objective Intelligibility (ESTOI) [30]. Additionally, we report non-intrusive metrics: DNSMOS [31] and SIGMOS [32], where SIGMOS is included because it supports full-band signals, unlike DNSMOS. We compared our proposed model with the following baselines: Conv-TasNet [33] and Demucs [34] for the predictive models and CDiffuSE [1], SGMSE+ [3], and SB [17] for the generative models.

The results on the EARS-WHAM test set are presented in Table 1. Our model, SB-UFOGen, significantly outperforms competing models across all evaluation metrics when using the same sampling steps, 4 or 1. Moreover, SB-UFOGen, with only a single step, even outperforms the baselines that require 50 steps or more to achieve their best performance. We also calculate the Real-Time Factor (RTF) by measuring the processing time per second, averaged over 20 test set utterances, using an

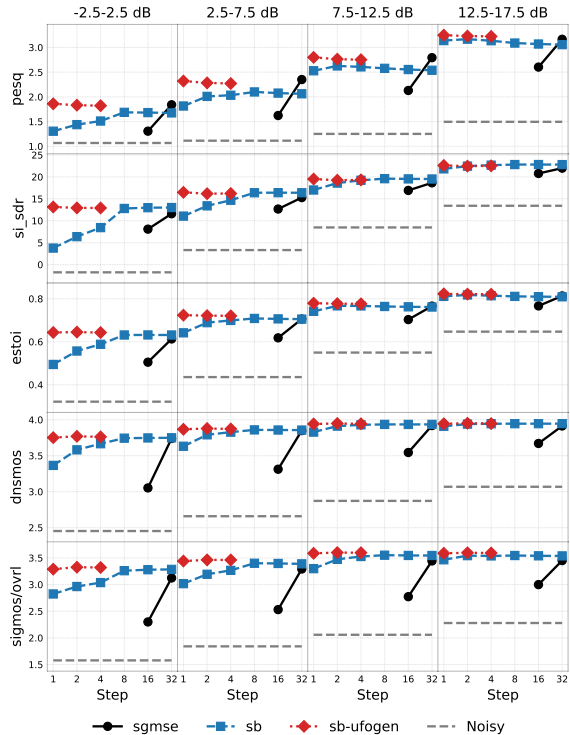


Figure 1: Results on EARS-WHAM test set across different sampling steps and SNR ranges.

Table 2: Denoising results on EARS-WHAM (with PESQ loss).

Model	NFEs	PESQ	SISDR	ESTOI	DNSMOS	SIGMOS
SB [4]	50	3.08	16.2	0.73	3.74	3.17
	4	3.03	15.9	0.72	3.72	3.10
	1	3.09	16.1	0.73	3.71	3.10
SB-UFOGen	4	3.13	17.1	0.74	3.76	3.16
	1	3.14	17.2	0.74	3.76	3.15

NVIDIA RTX A6000 GPU. While the SB model [17] with 50 iterations achieves an RTF of 2.112 proc/s, our SB-UFOGen model with a single iteration achieves 0.046 proc/s, making it competitive with predictive models.

Figure 1 shows the influence of the number of steps in the reverse process on the EARS-WHAM test dataset, segmented by input SNR. While SGMSE+ suffers significant performance degradation as the number of steps decreases, the SB model with the ODE sampler effectively maintains denoising performance with fewer iterations, especially in high-SNR scenarios. However, for low SNR, the performance of SB still declines sharply with fewer iterations. Comparing 32-step and 1-step inference of SB, SI-SDR decreases by only 0.97 dB and SIGMOS by 0.07 in the SNR range of [12.5, 17.5] dB, whereas in [-2.5, 2.5] dB range, the decreases are drastic, with a drop of 9.23 dB and 0.46, respectively. Therefore, diffusion or SB-based speech enhancement models still require increasing the number of sampling steps, which is not the case for our model.

Since PESQ is regarded as a standard evaluation metric for speech enhancement, a differentiable version, PESQ loss [35], was introduced in [6, 17]. Given its state-of-the-art performance on the PESQ metric, we also evaluate another version of our model with PESQ loss, using the same weighting coefficient as in [17]. The results are presented in Table 2. Unlike the formal results, the SB model is more robust to the number of sampling steps when incorporating the PESQ loss. Nevertheless, the SB-UFOGen model outperforms the SB model. On the other hand,

Table 3: Dereverberation results on EARS-Reverb.

Model	NFEs	PESQ	SISDR	ESTOI	SIGMOS	MOS Reverb
Unprocessed		1.47	-16.22	0.52	2.77	2.97
SGMSE+ [3]	60	3.04	6.35	0.85	3.49	4.73
SB [4]	50	3.41	6.65	0.88	3.37	4.73
	4	3.24	0.94	0.86	3.32	4.71
	1	2.12	-12.31	0.77	3.16	4.47
SB-UFOGen	4	3.33	6.71	0.87	3.37	4.72
	1	3.36	8.73	0.88	3.33	4.71




Figure 2: EARS-Reverb results across different sampling steps.

incorporating the PESQ loss leads to a decline in other metrics, particularly non-intrusive metrics such as DNSMOS and SIGMOS. This is expected, as PESQ loss only prioritizes PESQ, increasing the risk of overfitting to this specific metric [36].

4.3. Dereverberation Results

We also trained and tested the models using the full-band EARS-Reverb dataset [21] for the dereverberation task. We use the SDE sampler as it has demonstrated superior performance for the dereverberation task [4]. Results are evaluated using the same metrics as for denoising, except that SIGMOS Reverb is reported instead of the DNSMOS.

The results on the EARS-Reverb test set are presented in Table 3. SB-UFOGen also significantly outperforms baselines on the dereverberation task when using the same sampling steps, 4 or 1, especially in terms of SI-SDR. Moreover, single-step SB-UFOGen achieves performance comparable to baselines that require 50 steps or more to reach their optimal results.

Figure 2 illustrates the impact of the number of steps in the reverse process on the EARS-Reverb test dataset. While the SGMSE+ model experiences significant performance degradation as the number of steps decreases, the SB model shows some robustness to the number of steps. However, the SB model yet shows a decline in performance with fewer iterations. Therefore, for dereverberation tasks based on the SB model, reducing the number of sampling steps remains a challenge, which our model effectively addresses.

5. Conclusion

Diffusion models have gained attention in speech enhancement for their generative capabilities, while SB models have recently emerged as a promising alternative. However, both approaches are still constrained by the iterative reverse process. Our investigation reveals that SB models suffer notable performance decline, particularly under low-SNR conditions. To address this, we combine SB and GAN models, achieving better performance than the baselines with single-step inference for both denoising and dereverberation. Notably, single-step inference, as a direct predictive mapping, yields higher scores on intrusive metrics than 4-step inference. In future work, similar studies can be conducted under more severe conditions and a wider range of degradations. Moreover, our model can be extended to handle a broader range of degradations, including super-resolution [37] and universal speech enhancement [2].

6. Acknowledgements

This work was supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (No. RS-2024-00461617, 50%), and the Institute of Information & Communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) [No. RS-2022-II220641, 40%], [No. RS-2021-II211343, Artificial Intelligence Graduate School Program (Seoul National University), 5%], and [No. RS-2021-II212068, Artificial Intelligence Innovation Hub, 5%].

7. References

- [1] Y.-J. Lu, Z.-Q. Wang, S. Watanabe, A. Richard, C. Yu, and Y. Tsao, "Conditional diffusion probabilistic model for speech enhancement," in *IEEE ICASSP*, 2022, pp. 7402–7406.
- [2] J. Serrà, S. Pascual, J. Pons, R. O. Araz, and D. Scaini, "Universal speech enhancement with score-based diffusion," *arXiv:2206.03065*, 2022.
- [3] J. Richter, S. Welker, J.-M. Lemercier, B. Lay, and T. Gerkmann, "Speech enhancement and dereverberation with diffusion-based generative models," *IEEE/ACM TASLP*, vol. 31, pp. 2351–2364, 2023.
- [4] A. Jukić, R. Korostik, J. Balam, and B. Ginsburg, "Schrödinger bridge for generative speech enhancement," in *Interspeech*, 2024, pp. 1175–1179.
- [5] S. Wang, S. Liu, A. Harper, P. Kendrick, M. Salzmann, and M. Cernak, "Diffusion-based speech enhancement with schrödinger bridge and symmetric noise schedule," *arXiv:2409.05116*, 2024.
- [6] J. Richter, D. De Oliveira, and T. Gerkmann, "Investigating training objectives for generative speech enhancement," in *IEEE ICASSP*, 2025, pp. 1–5.
- [7] S. S. Shetu, E. A. P. Habets, and A. Brendel, "Gan-based speech enhancement for low snr using latent feature conditioning," in *IEEE ICASSP*, 2025, pp. 1–5.
- [8] D. de Oliveira, J. Richter, J.-M. Lemercier, T. Peer, and T. Gerkmann, "On the behavior of intrusive and non-intrusive speech enhancement metrics in predictive and generative settings," in *Speech Communication; 15th ITG Conference*, 2023, pp. 260–264.
- [9] Y. Song, J. Sohl-Dickstein, D. P. Kingma, A. Kumar, S. Ermon, and B. Poole, "Score-based generative modeling through stochastic differential equations," in *ICLR*, 2021.
- [10] E. Schrödinger, "Sur la théorie relativiste de l'électron et l'interprétation de la mécanique quantique," in *Annales de l'institut Henri Poincaré*, vol. 2, no. 4, 1932, pp. 269–310.
- [11] V. De Bortoli, J. Thornton, J. Heng, and A. Doucet, "Diffusion schrödinger bridge with applications to score-based generative modeling," in *NeurIPS*, vol. 34, 2021, pp. 17 695–17 709.
- [12] T. Chen, G.-H. Liu, and E. A. Theodorou, "Likelihood training of schrödinger bridge using forward-backward sdes theory," in *ICLR*, 2022.
- [13] Y. Shi, V. De Bortoli, A. Campbell, and A. Doucet, "Diffusion schrödinger bridge matching," in *NeurIPS*, vol. 36, 2023, pp. 62 183–62 223.
- [14] G.-H. Liu, A. Vahdat, D.-A. Huang, E. A. Theodorou, W. Nie, and A. Anandkumar, "I² sb: Image-to-image schrödinger bridge," in *ICML*, 2023.
- [15] Z. Chen, G. He, K. Zheng, X. Tan, and J. Zhu, "Schrödinger bridges beat diffusion models on text-to-speech synthesis," *arXiv:2312.03491*, 2023.
- [16] C. Li, Z. Chen, F. Bao, and J. Zhu, "Bridge-sr: Schrödinger bridge for efficient sr," in *IEEE ICASSP*, 2025, pp. 1–5.
- [17] J. Richter and T. Gerkmann, "Diffusion-based speech enhancement: Demonstration of performance and generalization," *Audio Imagination: NeurIPS 2024 Workshop AI-Driven Speech, Music, and Sound Generation*.
- [18] Z. Xiao, K. Kreis, and A. Vahdat, "Tackling the generative learning trilemma with denoising diffusion GANs," in *ICLR*, 2022.
- [19] y. xu, M. Gong, S. Xie, W. Wei, M. Grundmann, K. Batmanghelich, and T. Hou, "Semi-implicit denoising diffusion models (siddms)," in *NeurIPS*, vol. 36, 2023, pp. 17 383–17 394.
- [20] Y. Xu, Y. Zhao, Z. Xiao, and T. Hou, "Ufogen: You forward once large scale text-to-image generation via diffusion gans," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024, pp. 8196–8206.
- [21] J. Richter, Y.-C. Wu, S. Krenn, S. Welker, B. Lay, S. Watanabe, A. Richard, and T. Gerkmann, "Ears: An anechoic fullband speech dataset benchmarked for speech enhancement and dereverberation," in *Interspeech*, 2024, pp. 4873–4877.
- [22] N. Gushchin, D. Selikhanovych, S. Kholkin, E. Burnaev, and A. Korotin, "Adversarial schrödinger bridge matching," in *NeurIPS*, vol. 37, 2024, pp. 89 612–89 651.
- [23] A. Défossez, J. Copet, G. Synnaeve, and Y. Adi, "High fidelity neural audio compression," *TMLR*, 2023.
- [24] R. Scheibler, Y. Ji, S.-W. Chung, J. Byun, S. Choe, and M.-S. Choi, "Diffusion-based generative speech source separation," in *IEEE ICASSP*, 2023, pp. 1–5.
- [25] R. Kumar, P. Seetharaman, A. Luebs, I. Kumar, and K. Kumar, "High-fidelity audio compression with improved rvqgan," in *NeurIPS*, vol. 36, 2023, pp. 27 980–27 993.
- [26] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," in *ICLR*, 2019.
- [27] Y. Song and S. Ermon, "Improved techniques for training score-based generative models," in *NeurIPS*, vol. 33, 2020, pp. 12 438–12 448.
- [28] A. Rix, J. Beerends, M. Hollier, and A. Hekstra, "Perceptual evaluation of speech quality (pesq)-a new method for speech quality assessment of telephone networks and codecs," in *IEEE International Conference on Acoustics, Speech, and Signal Processing*, vol. 2, 2001, pp. 749–752 vol.2.
- [29] J. L. Roux, S. Wisdom, H. Erdogan, and J. R. Hershey, "Sdr – half-baked or well done?" in *IEEE ICASSP*, 2019, pp. 626–630.
- [30] J. Jensen and C. H. Taal, "An algorithm for predicting the intelligibility of speech masked by modulated noise maskers," *IEEE/ACM TASLP*, vol. 24, no. 11, pp. 2009–2022, 2016.
- [31] C. K. A. Reddy, V. Gopal, and R. Cutler, "Dnsmos: A non-intrusive perceptual objective speech quality metric to evaluate noise suppressors," in *IEEE ICASSP*, 2021, pp. 6493–6497.
- [32] N.-C. Ristea, A. Saabas, R. Cutler, B. Naderi, S. Braun, and S. Branets, "Icassp 2024 speech signal improvement challenge," in *IEEE International Conference on Acoustics, Speech, and Signal Processing Workshops*, 2024, pp. 15–16.
- [33] Y. Luo and N. Mesgarani, "Conv-tasnet: Surpassing ideal time-frequency magnitude masking for speech separation," *IEEE/ACM TASLP*, vol. 27, no. 8, pp. 1256–1266, 2019.
- [34] S. Rouard, F. Massa, and A. Défossez, "Hybrid transformers for music source separation," in *IEEE ICASSP*, 2023, pp. 1–5.
- [35] I. A. L. Erlangen, "torch-pesq: Loss function inspired by the pesq score," <https://github.com/audiolabs/torch-pesq>, 2023, accessed: 2025-01-03.
- [36] D. de Oliveira, S. Welker, J. Richter, and T. Gerkmann, "The pesqetarian: On the relevance of goodhart's law for speech enhancement," in *Interspeech*, 2024, pp. 3854–3858.
- [37] J. Lee and S. Han, "Nu-wave: A diffusion probabilistic model for neural audio upsampling," in *Interspeech*, 2021, pp. 1634–1638.