



# SpecGrad: Diffusion Probabilistic Model based Neural Vocoder with Adaptive Noise Spectral Shaping

Yuma Koizumi<sup>1</sup>, Heiga Zen<sup>1</sup>, Kohei Yatabe<sup>2</sup>, Nanxin Chen<sup>1</sup>, Michiel Bacchiani<sup>1</sup>

<sup>1</sup>Google Research, <sup>2</sup>Tokyo University of Agriculture and Technology

{koizumiyuma, heigazen, nanxinchen, michiel}@google.com

## Abstract

Neural vocoder using denoising diffusion probabilistic model (DDPM) has been improved by adaptation of the diffusion noise distribution to given acoustic features. In this study, we propose *SpecGrad* that adapts the diffusion noise so that its time-varying spectral envelope becomes close to the conditioning log-mel spectrogram. This adaptation by time-varying filtering improves the sound quality especially in the high-frequency bands. It is processed in the time-frequency domain to keep the computational cost almost the same as the conventional DDPM-based neural vocoders. Experimental results showed that *SpecGrad* generates higher-fidelity speech waveform than conventional DDPM-based neural vocoders in both analysis-synthesis and speech enhancement scenarios. Audio demos are available at [wavegrad.github.io/specgrad/](http://wavegrad.github.io/specgrad/).

**Index Terms:** Denoising diffusion probabilistic model, neural vocoder, spectral envelope, and speech enhancement.

## 1. Introduction

Neural vocoders [1–4] are neural networks that generate a speech waveform given acoustic features. They are indispensable building blocks of recent speech applications. For example, they are used as a backend model of text-to-speech (TTS) [1, 5] and speech enhancement (SE) [6–10]. A challenge in neural vocoder research is to generate a high-fidelity speech waveform with low computational costs. Autoregressive models [1, 11, 12] have revolutionized the quality of output signals, yet the nature of the models requires a large number of sequential operations for generation. To speed up the inference, various non-autoregressive models have been proposed such as generative adversarial networks [13–17], flow-based models [3, 4], and signal processing-inspired models [18–20].

Among non-autoregressive models, denoising diffusion probabilistic models (DDPMs) [21] have recently gained increased attention due to its quality and the controllable computational cost [22–28]. DDPMs convert a random signal into a speech waveform by the iterative sampling process called *denoising process* as illustrated in Fig. 1 (a). Since it iteratively refines a waveform, DDPMs have a tradeoff between the output quality and computational costs [22], i.e., many iterations are necessary for obtaining a high-fidelity waveform. To reduce the number of iterations while maintaining the quality, the conventional studies have proposed a proper inference noise schedule [25, 26] and/or network architecture [27, 28].

PriorGrad [24] provided a new approach to DDPM-based neural vocoders by considering the prior distribution on the acoustic model [29]. It adapts diffusion noises based on the conditioning log-mel spectrogram as illustrated in Fig. 1 (b). More specifically, the diffusion distribution is Gaussian with the diagonal covariance matrix whose diagonal entries are frame-wise energies of the mel-spectrogram. This proposal can be regarded

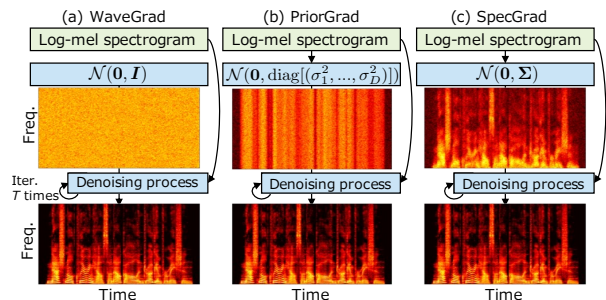


Figure 1: (a) *WaveGrad* [22], (b) *PriorGrad* with a diagonal covariance matrix [24], and (c) *Proposed SpecGrad*. While (a) samples noise from the standard Gaussian, (b) and (c) sample noise from Gaussian with covariance matrix calculated from the conditioning log-mel spectrogram. *Proposed SpecGrad* in (c) adapts noise to both signal power and spectral envelope.

as scaling of noise schedule for each sample point because a diagonal covariance matrix represents the power of a waveform in the time domain. Introduction of this well-known relation in signal processing to DDPMs is one of the important contributions of *PriorGrad*. The success of *PriorGrad* suggests that DDPM-based neural vocoders can be improved further by incorporating more knowledge from signal processing.

In this study, we propose *SpecGrad* that adapts spectral envelope of diffusion noise to the conditioning log-mel spectrogram as illustrated in Fig. 1 (c). We begin our discussion with the decomposed covariance matrix that is used in both diffusion noise generation and cost function of DDPMs. Then, we design a covariance matrix to manipulate the spectral envelope of diffusion noise, which is realized by time-varying filtering in the time-frequency (T-F) domain. We conducted objective and subjective experiments to show that *SpecGrad* generates a waveform whose quality is better than *WaveGrad* [22] and *PriorGrad* [24] on both analysis-synthesis and SE tasks.

## 2. Conventional methods

### 2.1. DDPM-based neural vocoder

Let a  $D$ -point speech waveform  $\mathbf{x}_0 \in \mathbb{R}^D$  be generated from a conditioning log-mel spectrogram  $\mathbf{c} = (\mathbf{c}_1, \dots, \mathbf{c}_K) \in \mathbb{R}^{FK}$ , where  $\mathbf{c}_k \in \mathbb{R}^F$  is an  $F$ -point log-mel spectrum at  $k$ th time frame, and  $K$  is the number of time frames. Our goal is to find a probability density function (PDF) of  $\mathbf{x}_0$  as  $q(\mathbf{x}_0 | \mathbf{c})$ .

A DDPM-based neural vocoder is a latent variable model based on a Markov chain of  $\mathbf{x}_t \in \mathbb{R}^D$  with learned Gaussian transitions, starting from  $q(\mathbf{x}_T) = \mathcal{N}(\mathbf{0}, \mathbf{I})$ , defined as

$$q(\mathbf{x}_0 | \mathbf{c}) = \int_{\mathbb{R}^{DT}} q(\mathbf{x}_T) \prod_{t=1}^T q(\mathbf{x}_{t-1} | \mathbf{x}_t, \mathbf{c}) d\mathbf{x}_1 \cdots d\mathbf{x}_T. \quad (1)$$

By modeling  $q(\mathbf{x}_{t-1} | \mathbf{x}_t, \mathbf{c})$ ,  $\mathbf{x}_0$  can be obtained from  $\mathbf{x}_T$  via recursive sampling of  $\mathbf{x}_{t-1}$  from  $\mathbf{x}_t$ .

In DDPMs,  $\mathbf{x}_{t-1}$  is generated by the *diffusion process* that gradually adds Gaussian noise to the waveform according to a noise schedule  $\{\beta_1, \dots, \beta_T\}$  given by  $p(\mathbf{x}_t | \mathbf{x}_{t-1}) = \mathcal{N}(\sqrt{1 - \beta_t} \mathbf{x}_{t-1}, \beta_t \mathbf{I})$ . This formulation enables us to sample  $\mathbf{x}_t$  at an arbitrary timestep  $t$  in a closed form as

$$\mathbf{x}_t = \sqrt{\bar{\alpha}_t} \mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t} \boldsymbol{\epsilon}, \quad (2)$$

where  $\alpha_t = 1 - \beta_t$ ,  $\bar{\alpha}_t = \prod_{s=1}^t \alpha_s$ , and  $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ . As proposed by Ho *et al.* [21], DDPM-based neural vocoders use a deep neural network (DNN)  $\mathcal{F}$  with parameter  $\theta$  for predicting  $\boldsymbol{\epsilon}$  from  $\mathbf{x}_t$  as  $\boldsymbol{\epsilon} = \mathcal{F}_\theta(\mathbf{x}_t, \mathbf{c}, \beta_t)$ . Then, if  $\beta_t$  is small enough,  $q(\mathbf{x}_{t-1} | \mathbf{x}_t, \mathbf{c})$  can be given by  $\mathcal{N}(\boldsymbol{\mu}_t, \gamma_t \mathbf{I})$ , where

$$\boldsymbol{\mu}_t = \frac{1}{\sqrt{\alpha_t}} \left( \mathbf{x}_t - \frac{1 - \alpha_t}{\sqrt{1 - \bar{\alpha}_t}} \mathcal{F}_\theta(\mathbf{x}_t, \mathbf{c}, \beta_t) \right), \quad (3)$$

and  $\gamma_t = \frac{1 - \bar{\alpha}_{t-1}}{1 - \bar{\alpha}_t} \beta_t$ . The DNN  $\mathcal{F}$  can be trained by maximizing the evidence lower bound (ELBO), though most of DDPM-based neural vocoders use a simplified loss function for the training; for example, WaveGrad [22] used the  $\ell_1$  norm as

$$\mathcal{L}^{\text{WG}} = \|\boldsymbol{\epsilon} - \mathcal{F}_\theta(\mathbf{x}_t, \mathbf{c}, \beta_t)\|_1, \quad (4)$$

where  $\|\cdot\|_p$  denotes the  $\ell_p$  norm.

## 2.2. PriorGrad

Lee *et al.* proposed PriorGrad [24] by introducing an adaptive prior  $\mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma})$ , where  $\boldsymbol{\Sigma}$  is computed from  $\mathbf{c}$ . Compared to the conventional DDPM-based neural vocoders, PriorGrad is different in two points: (i)  $\boldsymbol{\epsilon}$  is sampled from  $\mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma})$  for all diffusion steps, and (ii) the Mahalanobis distance according to  $\boldsymbol{\Sigma}$ ,

$$\mathcal{L}^{\text{PG}} = (\boldsymbol{\epsilon} - \mathcal{F}_\theta(\mathbf{x}_t, \mathbf{c}, \beta_t))^\top \boldsymbol{\Sigma}^{-1} (\boldsymbol{\epsilon} - \mathcal{F}_\theta(\mathbf{x}_t, \mathbf{c}, \beta_t)), \quad (5)$$

is used for the loss function, where  $^\top$  is the transpose. In their experiments, a specific form of the covariance matrix was given by  $\boldsymbol{\Sigma} = \text{diag}[(\sigma_1^2, \sigma_2^2, \dots, \sigma_D^2)]$ , where  $\sigma_d^2$  denotes the signal power at  $d$ th sample calculated by interpolating the normalized frame-wise energy of  $c_k$  [24], and  $\text{diag}$  constructs the diagonal matrix whose diagonal entries are those of the input vector.

## 3. Proposed method

### 3.1. SpecGrad

The performance of PriorGrad is determined by the covariance matrix  $\boldsymbol{\Sigma}$ . As Lee *et al.* showed, ELBO of PriorGrad becomes small when  $\boldsymbol{\Sigma}$  is close to the covariance of  $\mathbf{x}_0$  [24]. One way to obtain such  $\boldsymbol{\Sigma}$  is to make the amplitude spectrum of  $\boldsymbol{\epsilon}$  similar to that of  $\mathbf{x}_0$ . In this paper, we propose *SpecGrad* by incorporating the information of spectral envelope into  $\boldsymbol{\Sigma}$ .

Since  $\boldsymbol{\Sigma}$  is positive semi-definite, it can be decomposed as  $\boldsymbol{\Sigma} = \mathbf{L}\mathbf{L}^\top$ . Then, sampling from  $\mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma})$  can be written as  $\boldsymbol{\epsilon} = \mathbf{L}\tilde{\boldsymbol{\epsilon}}$  using  $\tilde{\boldsymbol{\epsilon}} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ , and Eq. (5) can be rewritten as

$$\mathcal{L}^{\text{SG}} = \|\mathbf{L}^{-1}(\boldsymbol{\epsilon} - \mathcal{F}_\theta(\mathbf{x}_t, \mathbf{c}, \beta_t))\|_2^2. \quad (6)$$

Thus, our interest is to design  $\mathbf{L} \in \mathbb{R}^{D \times D}$  so that a high-fidelity waveform can be generated with low computational costs.

Some desired properties of  $\mathbf{L}$  are as follows. First, the amplitude spectrum of  $\boldsymbol{\epsilon}$  ( $= \mathbf{L}\tilde{\boldsymbol{\epsilon}}$ ) approximates that of  $\mathbf{x}_0$ . This property is required to lower ELBO. Second, multiplication of

---

### Algorithm 1: Training of SpecGrad.

---

```

1 Function TrainOneStep( $\mathbf{x}_0, \mathbf{c}, \mathbf{M}, \beta_t$ ):
2    $\boldsymbol{\epsilon} \leftarrow \text{SampleNoise}(\mathbf{M})$ 
3    $\mathbf{x}_t \leftarrow \sqrt{\bar{\alpha}_t} \mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t} \boldsymbol{\epsilon}$ 
4    $\mathcal{L} \leftarrow \|\mathbf{G}^+ \mathbf{M}^{-1} \mathbf{G}(\boldsymbol{\epsilon} - \mathcal{F}_\theta(\mathbf{x}_t, \mathbf{c}, \beta_t))\|_2^2$ 
5   Update the model parameter  $\theta$  based on  $\nabla_\theta \mathcal{L}$ 
6 Function SampleNoise( $\mathbf{M}$ ):
7   Sample  $\tilde{\boldsymbol{\epsilon}} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ 
8   return  $\mathbf{G}^+ \mathbf{M} \mathbf{G} \tilde{\boldsymbol{\epsilon}}$ 

```

---



---

### Algorithm 2: Inference of SpecGrad.

---

```

1 Function Sampling( $\mathbf{c}, \mathbf{M}, \beta_1, \dots, \beta_T$ ):
2    $\mathbf{x}_T \leftarrow \text{SampleNoise}(\mathbf{M})$ 
3   for  $t = T$  to 1 do
4      $\hat{\boldsymbol{\epsilon}} \leftarrow \mathcal{F}_\theta(\mathbf{x}_t, \mathbf{c}, \beta_t)$ 
5      $\mathbf{x}_{t-1} \leftarrow \frac{1}{\sqrt{\alpha_t}} (\mathbf{x}_t - \frac{\beta_t}{\sqrt{1 - \bar{\alpha}_t}} \hat{\boldsymbol{\epsilon}})$ 
6     if  $t > 1$  then
7        $\mathbf{x}_{t-1} \leftarrow \mathbf{x}_{t-1} + \gamma_t \cdot \text{SampleNoise}(\mathbf{M})$ 
8   return  $\mathbf{x}_0$ 

```

---

$\mathbf{L}$  and  $\mathbf{L}^{-1}$  should be efficiently computed. Their computation directly impacts the total cost of training because  $\mathbf{L}$  and  $\mathbf{L}^{-1}$  are repeatedly applied in the training and depend on the training sample  $\mathbf{x}_0$  due to the first property.

To meet the requirements given in the previous paragraph, we propose to apply a time-varying filter in the T-F domain. Let the short-time Fourier transform (STFT) be represented by an  $NK \times D$  matrix  $\mathbf{G}$ , where  $N$  is the window size. We consider the following time-varying filter in the T-F domain:

$$\mathbf{L} = \mathbf{G}^+ \mathbf{M} \mathbf{G}, \quad (7)$$

where  $\mathbf{M} = \text{diag}[(m_{1,1}, \dots, m_{N,K})] \in \mathbb{C}^{NK \times NK}$  is the diagonal matrix that defines the filter,  $m_{n,k} \neq 0$  is a coefficient multiplied to the  $(n, k)$ th T-F bin, and  $\mathbf{G}^+$  is the matrix representation of the inverse STFT (iSTFT) using a dual window.<sup>1</sup> This representation allows us to recast the problem of designing  $\mathbf{L}$  to a filter design problem. We propose to design  $\mathbf{M}$  so that  $\boldsymbol{\epsilon}$  ( $= \mathbf{L}\tilde{\boldsymbol{\epsilon}}$ ) approximates the spectral envelope of  $\mathbf{x}_0$ . We also propose to approximate  $\mathbf{L}^{-1}$  as  $\mathbf{L}^{-1} \approx \mathbf{G}^+ \mathbf{M}^{-1} \mathbf{G}$ .

Since  $\mathbf{M}$  and  $\mathbf{M}^{-1}$  are diagonal, both  $\mathbf{L}$  and the approximate  $\mathbf{L}^{-1}$  can be applied with  $O(KN \log N)$  operations using a fast Fourier transform (FFT) algorithm. In addition,  $K$  FFTs can be computed in parallel. Therefore, Eq. (7) provides a good compromise between the flexibility and computational cost.

### 3.2. Implementation

Pseudocodes of the training and inference of the proposed method are shown in **Algorithm 1** and **2**, respectively. The differences from the vanilla DDPM-based neural vocoders [22,23] and PriorGrad are (i) the diffusion noise sampling  $\mathbf{G}^+ \mathbf{M} \mathbf{G} \tilde{\boldsymbol{\epsilon}}$  and (ii) the loss function in Eq. (6). Thus, the proposed method can coexist with the noise schedules and/or network architectures in the literature [25–28]. Although SpecGrad is slightly

<sup>1</sup>We define STFT and iSTFT so that they satisfy  $\mathbf{G}^+ \mathbf{G} = \mathbf{I}$  and  $\mathbf{G}^+ \mathbf{M} \mathbf{G} \in \mathbb{R}^{D \times D}$ . This is realized by (i) using a pair of windows that fulfills the perfect reconstruction condition and (ii) preserving the conjugate symmetry of spectra within  $\mathbf{G}^+ \mathbf{M}$ .

Table 1: Inference noise schedules used in experiments.

Name	Schedule
WG-3	[3E-4, 6E-2, 9E-1]
WG-6	[7E-6, 1.4E-4, 2.1E-3, 2.8E-2, 3.5E-1, 7E-1]
PG-6	[1E-4, 1E-3, 1E-2, 5E-2, 2E-1, 5E-1]
WG-50	<code>linspace(1e-4, 0.05, 50)</code>

more costly than WaveGrad, the inference speed of SpecGrad and WaveGrad is almost the same because the forward propagation of a typical DNN used in the neural vocoders is significantly more expensive than STFT.

For computation of the T-F domain filter  $M$ , we estimate the spectral envelope via cepstrum as follows. First, pseudoinverse of the mel-compression matrix is applied to  $c$  for computing the corresponding power spectrogram. Then, the  $r$ th order lifter is applied to compute the spectral envelope for each time frame. As with PriorGrad, to ensure numerical stability during training [24], we add 0.01 to the estimated envelope. The coefficients  $m_{1,k}, \dots, m_{N,k}$  for the  $k$ th time frame are computed from the obtained envelope with minimum phase response.

Note that any other method can be a choice for constructing the filter  $M$ . The above choice was due to our preliminary investigation. Directly using the spectrogram did not provide a satisfactory result, and thus we apply envelope estimation. We chose the cepstrum-based spectral envelope estimation method because its implementation is simpler than the other methods. These arguments are based on our informal experiments, and detailed study of the filter design is left as a future work.

## 4. Experiment

In this paper, we compared SpecGrad with WaveGrad [22] and PriorGrad [24] by both objective and subjective experiments. We also evaluated their performance as a backend module of speech enhancement. Since WaveGrad has been compared with an autoregressive model, WaveRNN [12], and several non-autoregressive models [15–17] in the literature, we focus on the DDPM-based neural vocoders with different diffusion PDFs. Audio demos are available in our demo page.<sup>2</sup>

### 4.1. Experimental setup

**Dataset:** We trained the models using a proprietary speech dataset consisted of 184 hours of high quality US English speech spoken by 11 female and 10 male speakers. For evaluation, we used 1,000 holdout samples of US English speech spoken by the same 21 speakers as the training dataset. The signals were downsampled to 24 kHz, and then 128-dimensional log-mel spectrograms (50 ms Hann window, 12.5 ms frame shift, 2048-point FFT, and 20 Hz and 12 kHz lower and upper frequency cutoffs, respectively) were extracted as  $c$ .

**Model and training setup:** To evaluate the performance difference due to the difference in diffusion PDF, we used the same network architecture and noise schedule for all three methods. We used the “WaveGrad Base model [22]” having 15M parameters. We trained all models using 128 Google TPU v3 cores with a global batch size of 512. To accelerate training, we randomly picked 120 frames (1.5 seconds,  $D = 36,000$  samples) as input. We trained all models for 1M steps (around 3 days) with the optimizer setting same as that of WaveGrad [22].

<sup>2</sup>[wavegrad.github.io/specgrad/](https://wavegrad.github.io/specgrad/)

Table 2: Mean opinion scores (MOS) and WARP-Q scores with their 95% confidence intervals. All models were trained by using training noise schedule of WaveGrad [22]. GT means ground-truth.

Method	Schedule	MOS ( $\uparrow$ )	WARP-Q ( $\downarrow$ )
WaveGrad	WG-3	$3.56 \pm 0.08$	$1.54 \pm 0.010$
PriorGrad		$3.45 \pm 0.08$	$1.33 \pm 0.008$
SpecGrad		<b><math>3.88 \pm 0.07</math></b>	<b><math>1.22 \pm 0.007</math></b>
WaveGrad	WG-6	$4.10 \pm 0.06$	$1.38 \pm 0.009$
PriorGrad		$4.01 \pm 0.07$	<b><math>1.12 \pm 0.007</math></b>
SpecGrad		<b><math>4.25 \pm 0.06</math></b>	<b><math>1.12 \pm 0.007</math></b>
WaveGrad	WG-50	$4.30 \pm 0.06$	$1.28 \pm 0.008$
PriorGrad		$4.32 \pm 0.06$	$1.11 \pm 0.007$
SpecGrad		<b><math>4.39 \pm 0.06</math></b>	<b><math>1.09 \pm 0.006</math></b>
GT	—	$4.47 \pm 0.05$	-

We tested two training noise schedules and several inference noise schedules listed in Table 1, which were used in the papers of WaveGrad [22] and PriorGrad [24]. One is the same as WaveGrad; training schedule was `linspace(1e-4, 0.005, 1000)`,<sup>3</sup> and inference schedules were WG-3, WG-6, and WG-50. The other one is the same as PriorGrad; training schedule was `linspace(1e-4, 0.05, 50)` and inference schedules were PG-6 and WG-50.

For PriorGrad and SpecGrad, we used the generalized energy distance (GED) [30], which was used in the first version of the PriorGrad paper,<sup>4</sup> as an auxiliary loss function. The weight for GED was 0.01 which is the same as that in the PriorGrad paper. The lifter order was  $r = 24$ . For STFT, we used the same settings as the log-mel spectrogram calculation.

**Metrics:** To evaluate subjective quality, we rated speech naturalness on a 5-point mean opinion score (MOS) scale (1: Bad, 2: Poor, 3: Fair, 4: Good, 5: Excellent) with rating increments of 0.5. Test stimuli were randomly chosen and presented to subjects in isolation, i.e., each stimulus was evaluated by one subject. Each subject was allowed to evaluate up to six stimuli. The subjects were paid native English speakers living in United States. They were requested to use headphones in a quiet room.

For objective evaluation, we used quality prediction for generative neural speech codecs (WAPR-Q) [31] which correlates with MOS of output sounds of neural vocoders. Since the default parameters of WARP-Q were for signals sampled at 16 kHz, we changed the cut-off frequency to 10 kHz and the number of mel-frequency cepstrum coefficients (MFCCs) to 24.

### 4.2. Results

The results for the noise schedules of WaveGrad and PriorGrad are shown in Tables 2 and 3, respectively. For all training and inference noise schedules, SpecGrad achieved the best scores for both subjective and objective metrics; SpecGrad provided better quality than WaveGrad and PriorGrad. Since it worked well regardless of the choice of noise schedule, SpecGrad should be able to enhance the performance of other DDPM-based neural vocoders that use extended noise schedules [25, 26].

We additionally conducted a 7-scale (-3 to 3) side-by-side (SxS) preference test for PG-6 inference schedule. The results are shown in Table 4. This test also indicated that SpecGrad

<sup>3</sup>`linspace` is the function that returns evenly spaced real numbers over the interval specified by the first two arguments.

<sup>4</sup><https://arxiv.org/pdf/2106.06406v1.pdf>

Table 3: Mean opinion scores (MOS) and WARP-Q scores with their 95% confidence intervals. All models were trained by using training noise schedule of PriorGrad [24].

Method	Schedule	MOS ( $\uparrow$ )	WARP-Q ( $\downarrow$ )
WaveGrad	PG-6	$4.14 \pm 0.06$	$1.32 \pm 0.009$
PriorGrad		$4.02 \pm 0.06$	$1.10 \pm 0.007$
SpecGrad		<b><math>4.31 \pm 0.05</math></b>	<b><math>1.05 \pm 0.006</math></b>
WaveGrad	WG-50	$4.29 \pm 0.05$	$1.31 \pm 0.008$
PriorGrad		$4.29 \pm 0.05$	$1.08 \pm 0.007$
SpecGrad		<b><math>4.40 \pm 0.05</math></b>	<b><math>1.03 \pm 0.006</math></b>

Table 4: Results of side-by-side (SxS) test with 95% confidence intervals. Positive score means Method-A is preferred.

Method-A	Method-B	SxS
WaveGrad	PriorGrad	$0.165 \pm 0.061$
SpecGrad	WaveGrad	$0.161 \pm 0.065$
SpecGrad	PriorGrad	$0.360 \pm 0.075$

is better than WaveGrad and PriorGrad. Note that PriorGrad was rated worse than WaveGrad in both MOS and SxS scores. We found that artifacts generated by phase distortion of high frequency components were noticeable in PriorGrad. Our experiments set the cutoff frequency of log-mel spectrogram to 12 kHz because the sampling frequency was 24 kHz, whereas that of the PriorGrad paper was 7.6 kHz [24]. This difference can be a reason for the inconsistency between our and their experiments. The unique property of SpecGrad is reduction of high frequency components according to the conditioning log-mel spectrogram (see Fig. 1). This feature can contribute to prevention of the artifacts because the amount of subtraction of the estimated noise ( $\hat{\epsilon}$  in Algorithm 2) becomes smaller for higher frequency.

### 4.3. Evaluation as speech enhancement backend

Since the diffusion PDF of SpecGrad depends on the conditioning log-mel spectrogram  $c$ , its error might degrade the quality of output signals. Therefore, we investigated the robustness of SpecGrad to the error of  $c$ . As a realistic scenario, we applied the DDPM-based neural vocoders as backends of an SE system. WARP-Q and extended short-time objective intelligibility measure (ESTOI) [32] were used for objectively evaluating the speech quality and intelligibility, respectively. Note that we used WARP-Q instead of the common metrics in speech enhancement [6, 7] such as the perceptual evaluation of speech quality (PESQ) [33], CSIG, CBLK, and COVL [34] because they are designed for waveforms sampled at 16 kHz.

**Model and training setup:** We followed the SE scheme of parametric resynthesis [6], i.e., the frontend predicts the clean log-mel spectrogram from an observed noisy log-mel spectrogram, and then the backend generates a waveform given the predicted log-mel spectrogram. For the SE frontend, we used a combination of a part of DF-Conformer [35] and Post-Net of Tacotron2 [5]; an observed noisy log-mel spectrogram was enhanced by the mask predictor of DF-Conformer, and then its output was further cleaned up by the Post-Net.

We used the ‘‘DF-Conformer-8’’ model [35], whose bottleneck feature dimension was adapted to that of the input log-mel spectrogram (i.e., 128), and Post-Net in Tacotron2 [5] without modification. The total number of their parameters was 7.6M.

Table 5: Results for speech enhancement experiment. ESTOI and WARP-Q scores with their 95% confidence intervals.

Method	ESTOI [%] ( $\uparrow$ )	WARP-Q ( $\downarrow$ )
WaveGrad	$82.7 \pm 0.7$	$1.97 \pm 0.019$
PriorGrad	$81.9 \pm 0.7$	$1.92 \pm 0.019$
SpecGrad	<b><math>83.6 \pm 0.7</math></b>	<b><math>1.89 \pm 0.019</math></b>

This frontend was pretrained to minimize the sum of the mean-absolute- and mean-squared-error between the clean and predicted log-mel spectrograms at before and after Post-Net [5]. The optimizer and batch size settings were the same as those of the training of SpecGrad, and we pretrained all models for 500k steps. Then, the joint network was built by concatenating the pretrained frontend and a backend in Section 4.2 that was trained using WaveGrad noise schedule. This joint network was finetuned for 500k steps to minimize the sum of the frontend loss and backend loss. In the inference stage, WG-50 noise schedule was used for all models.

**Dataset:** Training and test datasets were generated by contaminating the clean data used in Section 4.2 with reverberation and noise. A room impulse response (RIR) for each sample was generated by a stochastic RIR generator using the image method [36]. Its parameters were drawn from the following uniform distributions  $\mathcal{U}$ : the distance between the source and microphone was  $\mathcal{U}(0.5, 3.0)$  [cm], the length of one side of the square room was  $\mathcal{U}(2.0, 10.0)$  [m], and the reflection ratio was  $\mathcal{U}(0.5, 0.95)$ . For the noise dataset, the TAU Urban Audio-Visual Scenes 2021 dataset [37] were used. The average ratio of the energy of noise to the reverberated clean speech was 0 dB. The average ESTOI and WARP-Q scores of the generated noisy samples were  $51.9 \pm 1.1$  % and  $2.78 \pm 0.024$ , respectively.

**Results:** The results are shown in Table 5. For both metrics, the proposed method outperformed WaveGrad and PriorGrad. The differences of the scores between the proposed and conventional methods were smaller compared to those of the previous experiments, which should be because the frontend limited the maximum possible quality of the output signals (note that both ESTOI and WARP-Q measure the quality relative to the clean signals). Even so, these results suggested that SpecGrad is robust to the error of the conditioning log-mel spectrogram, at least on the similar level with the conventional methods.

## 5. Conclusion

We proposed *SpecGrad* that adapts the spectral envelope of diffusion noise based on the conditioning log-mel spectrogram. We designed the decomposed covariance matrix  $L$  and its approximate inverse using the idea from T-F domain filtering. This design allows us to use an FFT algorithm for computation of the matrix multiplication, which only adds a negligible extra computational cost compared to the forward computation of the conventional DDPM-based neural vocoders. The experimental results showed that SpecGrad generated waveforms of higher quality than the conventional methods. As SpecGrad performed well for different noise schedules, its combination with recently proposed extended noise schedules [25, 26] is promising.

We believe SpecGrad could be further improved by incorporating more ideas from signal processing. Some time-domain filters can be considered for the covariance matrix [38, 39]. The ideas from classic vocoders should be useful [19, 20] because SpecGrad can be viewed as a classic vocoder with noise excitation. These can be possible future work.

## 6. References

- [1] S. Mehri, K. Kumar, I. Gulrajani, R. Kumar, S. Jain, and J. Sotelo, "SampleRNN: An unconditional end-to-end neural audio," in *Proc. Int. Conf. Learn. Represent. (ICLR)*, 2018.
- [2] A. Tamamori, T. Hayashi, K. Kobayashi, K. Takeda, and T. Toda, "Speaker-dependent WaveNet vocoder," in *Proc. Interspeech*, 2017.
- [3] R. Prenger, R. Valle, and B. Catanzaro, "WaveGlow: A flow-based generative network for speech synthesis," in *Proc. Int. Conf. on Acoust., Speech, and Signal Process. (ICASSP)*, 2019.
- [4] W. Ping, K. Peng, K. Zhao, and Z. Song, "WaveFlow: A compact flow-based model for raw audio," in *Proc. Int. Conf. Mach. Learn. (ICML)*, 2020.
- [5] J. Shen, R. Pang, R. J. Weiss, M. Schuster, N. Jaitly, Z. Yang, Z. Chen, Y. Zhang, Y. Wang, R. Skerrv-Ryan, R. A. Saurous, Y. Agiomvrgiannakis, and Y. Wu, "Natural TTS synthesis by conditioning WaveNet on mel spectrogram predictions," in *Proc. Int. Conf. on Acoust., Speech, and Signal Process. (ICASSP)*, 2018.
- [6] S. Maiti and M. I. Mandel, "Parametric resynthesis with neural vocoders," in *Proc. IEEE Workshop Appl. Signal Process. Audio Acoust. (WASPAA)*, 2019.
- [7] —, "Speaker independence of neural vocoders and their effect on parametric resynthesis speech enhancement," in *Proc. Int. Conf. on Acoust., Speech, and Signal Process. (ICASSP)*, 2020.
- [8] J. Su, Z. Jin, and A. Finkelstein, "HiFi-GAN: High-fidelity denoising and dereverberation based on speech deep features in adversarial networks," in *Proc. Interspeech*, 2020.
- [9] —, "HiFi-GAN-2: Studio-quality speech enhancement via generative adversarial networks conditioned on acoustic features," in *Proc. IEEE Workshop Appl. Signal Process. Audio Acoust. (WASPAA)*, 2021.
- [10] H. Liu, Q. Kong, Q. Tian, Y. Zhao, D. L. Wang, C. Huang, and Y. Wang, "VoiceFixer: Toward general speech restoration with neural vocoder," *arXiv:2109.13731*, 2021.
- [11] A. van den Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. Senior, and K. Kavukcuoglu, "WaveNet: A generative model for raw audio," *arXiv:1609.03499*, 2016.
- [12] N. Kalchbrenner, W. Elsen, K. Simonyan, S. Noury, N. Casagrande, W. Lockhart, F. Stimberg, A. van den Oord, S. Dieleman, and K. Kavukcuoglu, "Efficient neural audio synthesis," in *Proc. Int. Conf. Mach. Learn. (ICML)*, 2018.
- [13] C. Donahue, J. McAuley, and M. Puckette, "Adversarial audio synthesis," in *Proc. Int. Conf. Learn. Represent. (ICLR)*, 2019.
- [14] J. Kong, J. Kim, and J. Bae, "HiFi-GAN: Generative adversarial networks for efficient and high fidelity speech synthesis," in *Proc. Adv. Neural Inf. Process. Syst. (NeurIPS)*, 2020.
- [15] K. Kumar, R. Kumar, T. de Boissiere, L. Gestin, W. Z. Teoh, J. Sotelo, A. de Brébisson, Y. Bengio, and A. C. Courville, "MelGAN: Generative adversarial networks for conditional waveform synthesis," in *Proc. Adv. Neural Inf. Process. Syst. (NeurIPS)*, 2019.
- [16] R. Yamamoto, E. Song, and J.-M. Kim, "Parallel WaveGAN: A fast waveform generation model based on generative adversarial networks with multi-resolution spectrogram," in *Proc. Int. Conf. on Acoust., Speech, and Signal Process. (ICASSP)*, 2020.
- [17] G. Yang, S. Yang, K. Liu, P. Fang, W. Chen, and L. Xie, "Multi-band MelGAN: Faster waveform generation for high-quality text-to-speech," in *Proc. IEEE Spoken Language Technology Workshop (SLT)*, 2021.
- [18] L. Juvela, B. Bollepalli, V. Tsias, and P. Alku, "GlottNet—a raw waveform model for the glottal excitation in statistical parametric speech synthesis," *IEEE/ACM Trans. Audio Speech Lang. Process.*, 2019.
- [19] X. Wang, S. Takaki, and J. Yamagishi, "Neural source-filter waveform models for statistical parametric speech synthesis," *IEEE/ACM Trans. Audio Speech Lang. Process.*, 2020.
- [20] Y. Hono, S. Takaki, K. Hashimoto, K. Oura, Y. Nankaku, and K. Tokuda, "PeriodNet: A non-autoregressive raw waveform generative model with a structure separating periodic and aperiodic components," *IEEE Access*, 2021.
- [21] J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models," in *Proc. Adv. Neural Inf. Process. Syst. (NeurIPS)*, 2020.
- [22] N. Chen, Y. Zhang, H. Zen, R. J. Weiss, M. Norouzi, and W. Chan, "WaveGrad: Estimating gradients for waveform generation," in *Proc. Int. Conf. Learn. Represent. (ICLR)*, 2021.
- [23] Z. Kong, W. Ping, J. Huang, K. Zhao, and B. Catanzaro, "DiffWave: A versatile diffusion model for audio synthesis," in *Proc. Int. Conf. Learn. Represent. (ICLR)*, 2021.
- [24] S. Lee, H. Kim, C. Shin, X. Tan, C. Liu, Q. Meng, T. Qin, W. Chen, S. Yoon, and T.-Y. Liu, "PriorGrad: Improving conditional denoising diffusion models with data-dependent adaptive prior," in *Proc. Int. Conf. Learn. Represent. (ICLR)*, 2022.
- [25] M. W. Y. Lam, J. Wang, D. Su, and D. Yu, "BDDM: Bilateral denoising diffusion models for fast and high-quality speech synthesis," in *Proc. Int. Conf. Learn. Represent. (ICLR)*, 2022.
- [26] Z. Chen, X. Tan, K. Wang, S. Pan, D. Mandic, L. He, and S. Zhao, "InferGrad: Improving diffusion models for vocoder by considering inference in training," in *Proc. Int. Conf. on Acoust., Speech, and Signal Process. (ICASSP)*, 2022.
- [27] T. Okamoto, T. Toda, Y. Shiga, and H. Kawai, "Noise level limited sub-modeling for diffusion probabilistic vocoders," in *Proc. Int. Conf. on Acoust., Speech, and Signal Process. (ICASSP)*, 2021.
- [28] K. Goel, A. Gu, C. Donahue, and C. Ré, "It's Raw! audio generation with state-space models," *arXiv:2202.09729*, 2022.
- [29] V. Popov, I. Vovk, V. G. T. Sadekova, and M. A. Kudinov, "GradTTS: A diffusion probabilistic model for text-to-speech," in *Proc. Int. Conf. Mach. Learn. (ICML)*, 2021.
- [30] A. A. Gritsenko, T. Salimans, R. van den Berg, J. Snoek, and N. Kalchbrenner, "A spectral energy distance for parallel speech synthesis," in *Proc. Adv. Neural Inf. Process. Syst. (NeurIPS)*, 2020.
- [31] W. A. Jassim, J. Skoglund, M. Chinen, and A. Hines, "WARP-Q: Quality prediction for generative neural speech codecs," in *Proc. Int. Conf. on Acoust., Speech, and Signal Process. (ICASSP)*, 2021.
- [32] J. Jensen and C. H. Taal, "An algorithm for predicting the intelligibility of speech masked by modulated noise maskers," *IEEE Trans. Audio Speech Lang. Process.*, 2016.
- [33] "P.862.2: Wideband extension to recommendation P.862 for the assessment of wideband telephone networks and speech codecs," *ITU-T Std.*, 2007.
- [34] Y. Hu and P. C. Loizou, "Evaluation of objective quality measures for speech enhancement," *IEEE/ACM Trans. Audio Speech Lang. Process.*, 2008.
- [35] Y. Koizumi, S. Karita, S. Wisdom, H. Erdogan, J. R. Hershey, L. Jones, and M. Bacchiani, "DF-Conformer: Integrated architecture of Conv-TasNet and Conformer using linear complexity self-attention for speech enhancement," in *Proc. IEEE Workshop Appl. Signal Process. Audio Acoust. (WASPAA)*, 2021.
- [36] J. B. Allen and D. A. Berkley, "Image method for efficiently simulating small-room acoustics," *J. Acoust. Soc. Am.*, 1979.
- [37] S. Wang, A. Mesaros, T. Heittola, and T. Virtanen, "A curated dataset of urban scenes for audio-visual scene analysis," in *Proc. Int. Conf. on Acoust., Speech, and Signal Process. (ICASSP)*, 2021.
- [38] K. Tokuda and H. Zen, "Directly modeling speech waveforms by neural networks for statistical parametric speech synthesis," in *Proc. Int. Conf. on Acoust., Speech, and Signal Process. (ICASSP)*, 2015.
- [39] —, "Directly modeling voiced and unvoiced components in speech waveforms by neural networks," in *Proc. Int. Conf. on Acoust., Speech, and Signal Process. (ICASSP)*, 2016.