



HAL
open science

GLA-Grad: A Griffin-Lim Extended Waveform Generation Diffusion Model

Haocheng Liu, Teysir Baoueb, Mathieu Fontaine, Jonathan Le Roux, Gael
Richard

► **To cite this version:**

Haocheng Liu, Teysir Baoueb, Mathieu Fontaine, Jonathan Le Roux, Gael Richard. GLA-Grad: A Griffin-Lim Extended Waveform Generation Diffusion Model. IEEE International Conference on Acoustics, Speech and Signal Processing, Apr 2024, Seoul (Korea), South Korea. hal-04424100

HAL Id: hal-04424100

<https://hal.science/hal-04424100>

Submitted on 9 Feb 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

GLA-GRAD: A GRIFFIN-LIM EXTENDED WAVEFORM GENERATION DIFFUSION MODEL

Haocheng Liu¹, Teysir Baoueb¹, Mathieu Fontaine¹, Jonathan Le Roux², Gaël Richard¹

¹LTCI, Télécom Paris, Institut polytechnique de Paris, France

²Mitsubishi Electric Research Laboratories (MERL), Cambridge, MA, USA

ABSTRACT

Diffusion models are receiving a growing interest for a variety of signal generation tasks such as speech or music synthesis. WaveGrad, for example, is a successful diffusion model that conditionally uses the mel spectrogram to guide a diffusion process for the generation of high-fidelity audio. However, such models face important challenges concerning the noise diffusion process for training and inference, and they have difficulty generating high-quality speech for speakers that were not seen during training. With the aim of minimizing the conditioning error and increasing the efficiency of the noise diffusion process, we propose in this paper a new scheme called GLA-Grad, which consists in introducing a phase recovery algorithm such as the Griffin-Lim algorithm (GLA) at each step of the regular diffusion process. Furthermore, it can be directly applied to an already-trained waveform generation model, without additional training or fine-tuning. We show that our algorithm outperforms state-of-the-art diffusion models for speech generation, especially when generating speech for a previously unseen target speaker.

Index Terms— Diffusion models, speech generation, Griffin-Lim algorithm, domain adaptation

1. INTRODUCTION

In the era of deep learning, generating high-fidelity audio data is a major research topic for improving the performance of machine listening algorithms. Apart from audio synthesis for listening applications, it can also be seen as an important asset for data augmentation in unbalanced settings, such as in sound event detection, speech processing, or music information retrieval-related tasks.

Speech generation has witnessed the utilization of various architectures and input methods to produce high-quality audio speech. Notable algorithms in speech generation encompass text-to-speech (TTS) variational autoencoder [1], as well as convolutional neural network (CNN) architectures operating directly on raw audio such as Wavenet [2]. Mel spectrograms serve as an alternative input for speech generation, with such examples as Waveglow [3], based on normalizing flow [4], and HiFi-GAN [5], based on a generative adversarial network (GAN) [6].

Diffusion models [7, 8] have recently emerged as highly advanced deep generative models where a stochastic flow to solve differential equations is governed by a fixed noise schedule. They have surpassed the long-standing dominance of GANs in the field

This work was funded by the ANR Project SAROUMANE (ANR-22-CE23-0011) and by the European Union (ERC, HI-Audio, 101052978). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them.

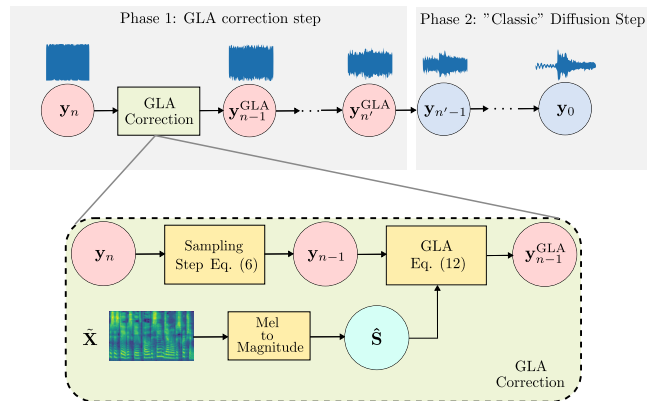


Fig. 1: Griffin-Lim corrected diffusion model

of image synthesis [9–11]. Diffusion models have also been considered for various speech-related tasks such as speech enhancement [12–14] or speech separation [15, 16]. Diffusion-based speech generation methods that use conditioning on mel spectrograms have also been proposed, notably Wavegrad [17] and DiffWave [18]. Building upon DiffWave, PriorGrad [19] replaces the standard Gaussian prior noise assumption by an adaptive prior, while SpecGrad [20] applies constraints to ensure the time-varying spectral envelope aligns with the log-mel spectrogram.

Nonetheless, diffusion models face an ongoing challenge in efficient training. Their stability relies heavily on large datasets to ensure the stability of the diffusion process through the iterations [21]. An essential component is the noise schedule governing the reverse process for generating speech. Recent advances such as FastDiff [22] propose to predict the noise schedule dynamically throughout iterations to reduce their number and enhance the signal generation. Furthermore, the versatility of speech signals presents domain adaptation limitations with a diffusion model framework [23]. These limitations result in poor generalization to unseen speakers in speech generation scenarios.

Our paper addresses the latest drawback which, to the best of our knowledge, has not been explored within the diffusion model framework for speech generation. In this article, we introduce a straightforward extension to WaveGrad, aiming to rectify conditioning errors on the mel spectrogram during the reverse process iterations. As illustrated in Fig. 1, we use the Griffin-Lim algorithm (GLA) [24] to partially reconcile the phase of the short-time Fourier transform of the current sample in the diffusion process with the magnitude spectrogram obtained as pseudo-inverse of the conditioning mel spectrogram. Our objective is to ensure consistency in the reverse process, even when the user-provided signal significantly differs from those in the training set.

The achieved scores demonstrate GLA-Grad’s competitiveness

with WaveGrad and SpecGrad when using a single-speaker dataset. However, GLA-Grad significantly outperforms them as the dataset’s speaker count increases, highlighting its superior generalization capability across multiple speakers. Additionally, experiments involving unseen speakers during training demonstrate GLA-Grad’s effectiveness in addressing domain adaptation challenges.

2. RELATED WORK

DDPM: Let $\mathbf{y}_0 \sim q(\mathbf{y}_0)$ be a data sample. The denoising diffusion probabilistic model (DDPM) [7] involves the gradual addition of noise to \mathbf{y}_0 in the forward process and then learning how to recover it in the reverse process. Specifically, in the forward process, \mathbf{y}_n at each time step $n \in \llbracket 1, N \rrbracket$ is derived via $q(\mathbf{y}_n | \mathbf{y}_{n-1}) = \mathcal{N}(\mathbf{y}_n; \sqrt{1 - \beta_n} \mathbf{y}_{n-1})$, where $(\beta_n)_n$ is a variance schedule and N is the maximum number of time steps. Therefore we have:

$$\mathbf{y}_0 \sim q(\mathbf{y}_0), \quad q(\mathbf{y}_{1:N} | \mathbf{y}_0) \triangleq \prod_{n=1}^N q(\mathbf{y}_n | \mathbf{y}_{n-1}). \quad (1)$$

Denoting $\alpha_n \triangleq 1 - \beta_n$, $\bar{\alpha}_n \triangleq \prod_{s=1}^n \alpha_s$ and $\boldsymbol{\epsilon} \sim \mathcal{N}(\boldsymbol{\epsilon}; \mathbf{0}, \mathbf{I})$, \mathbf{y}_n can be obtained in closed form for any given time n as follows:

$$\mathbf{y}_n = \sqrt{\bar{\alpha}_n} \mathbf{y}_0 + \sqrt{1 - \bar{\alpha}_n} \boldsymbol{\epsilon}. \quad (2)$$

In the reverse process, we want to reconstruct \mathbf{y}_0 . Since $q(\mathbf{y}_{n-1} | \mathbf{y}_n)$ is intractable, a neural network model is trained to match the true denoising distribution. The equation of the reverse process is given by:

$$p_\theta(\mathbf{y}_{0:N}) \triangleq p(\mathbf{y}_N) \prod_{n=1}^N p_\theta(\mathbf{y}_{n-1} | \mathbf{y}_n), \quad (3)$$

where, assuming N is large enough, $p_\theta(\mathbf{y}_{n-1} | \mathbf{y}_n)$ is modeled as a Gaussian distribution with mean and variance determined by the neural network, and θ represents the parameters of the model, and $p(\mathbf{y}_N) = \mathcal{N}(\mathbf{y}_N; \mathbf{0}, \mathbf{I})$.

In practice, the neural network, denoted by $\boldsymbol{\epsilon}_\theta(\mathbf{y}_n, n)$, is parameterized to estimate the noise added to \mathbf{y}_0 in Equation 2, and thus minimizes the following training objective, which bears resemblance to the concept of denoising score-matching:

$$\mathbb{E}_{n, \boldsymbol{\epsilon}} \left[\left\| \boldsymbol{\epsilon}_\theta(\sqrt{\bar{\alpha}_n} \mathbf{y}_0 + \sqrt{1 - \bar{\alpha}_n} \boldsymbol{\epsilon}, n) - \boldsymbol{\epsilon} \right\|_2^2 \right]. \quad (4)$$

WaveGrad: WaveGrad [17] is based on DDPM and proposes to perform waveform generation with a conditional diffusion process $p_\theta(\mathbf{y}_{0:n} | \tilde{\mathbf{X}})$ where $\tilde{\mathbf{X}}$ is a mel spectrogram. Starting from a Gaussian white noise signal, it progressively enhances the signal through an iterative refinement process using a gradient-based sampler that takes into account $\tilde{\mathbf{X}}$ as a conditioning factor. WaveGrad provides a flexible mechanism for balancing inference speed against sample quality by controlling the number of refinement steps. In contrast with DDPM [7], WaveGrad [17] reparameterized the model to condition on a continuous noise level $\bar{\alpha}$ instead of the discrete iteration index n . The loss function is then modified as follows:

$$\mathbb{E}_{\bar{\alpha}, \boldsymbol{\epsilon}} \left[\left\| \boldsymbol{\epsilon}_\theta \left(\sqrt{\bar{\alpha}} \mathbf{y}_0 + \sqrt{1 - \bar{\alpha}} \boldsymbol{\epsilon}, \tilde{\mathbf{X}}, \sqrt{\bar{\alpha}} \right) - \boldsymbol{\epsilon} \right\|_1 \right]. \quad (5)$$

Considering $\mathbf{z} \sim \mathcal{N}(\mathbf{z}; \mathbf{0}, \mathbf{I})$ for $n > 1$, $\mathbf{z} = \mathbf{0}$ for $n = 1$ and $\sigma_n = \frac{1 - \bar{\alpha}_n - 1}{1 - \bar{\alpha}_n} \beta_n$, Wavegrad applies the following iterative procedure to generate \mathbf{y}_{n-1} :

$$\mathbf{y}_{n-1} = \frac{\left(\mathbf{y}_n - \frac{1 - \bar{\alpha}_n}{\sqrt{1 - \bar{\alpha}_n}} \boldsymbol{\epsilon}_\theta \left(\mathbf{y}_n, \tilde{\mathbf{X}}, \sqrt{\bar{\alpha}_n} \right) \right)}{\sqrt{\bar{\alpha}_n}} + \sigma_n \mathbf{z}. \quad (6)$$

SpecGrad: Based on WaveGrad, SpecGrad [20] proposed to adjust the diffusion noise in a way that aligns its dynamic spectral characteristics with those of the conditioning mel spectrogram. The adaptation performed by SpecGrad through time-varying filtering results in notable enhancements in audio fidelity, particularly in the higher frequency regions. SpecGrad uses the implementation of a time-varying filter within the time-frequency domain. We denote by \mathbf{T} a matrix representation of the short-time Fourier transform (STFT) expressed from the time-domain to a flattened version of the time-frequency domain where all frames are concatenated, and by \mathbf{T}^\dagger the similar matrix representation of the inverse STFT (iSTFT). The time-varying filter can be expressed as:

$$\mathbf{L} = \mathbf{T}^\dagger \mathbf{D} \mathbf{T}, \quad (7)$$

where \mathbf{D} is a diagonal matrix determining the filter in the time-frequency domain, here obtained from the spectral envelope. From \mathbf{L} , we can obtain the covariance matrix $\boldsymbol{\Sigma} = \mathbf{L} \mathbf{L}^\dagger$ of the Gaussian noise $\mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma})$ in the diffusion process. The loss function is then obtained as:

$$\mathbb{E}_{\bar{\alpha}, \boldsymbol{\epsilon}} \left[\left\| \mathbf{L}^{-1} \left(\boldsymbol{\epsilon}_\theta \left(\sqrt{\bar{\alpha}} \mathbf{y}_0 + \sqrt{1 - \bar{\alpha}} \boldsymbol{\epsilon}, \tilde{\mathbf{X}}, \sqrt{\bar{\alpha}} \right) - \boldsymbol{\epsilon} \right) \right\|_2^2 \right]. \quad (8)$$

3. PROPOSED GRIFFIN-LIM DIFFUSION EXTENSION

3.1. Griffin-Lim Algorithm

The Griffin-Lim Algorithm (GLA) [24] iteratively reconstructs a time-domain signal $s \in \mathbb{R}^L$ of length L from a given magnitude spectrogram $\hat{\mathbf{S}} \in \mathbb{R}^{T \times F}$ by estimating a phase that is most consistent with that magnitude [25]. T and F denote the time and frequency dimensions of the time-frequency (TF) representation. For simplicity, we use the same notations \mathbf{T} and \mathbf{T}^\dagger for the STFT and iSTFT operators as for their matrix representations introduced earlier. GLA relies on two projection operations in the complex time-frequency domain $\mathbb{C}^{T \times F}$, the first onto the subset $\mathcal{C} \subset \mathbb{C}^{T \times F}$ of consistent spectrograms (the elements of $\mathbb{C}^{T \times F}$ that can be obtained as STFT of a signal in \mathbb{R}^L):

$$P_{\mathcal{C}}(\mathbf{C}) = \mathbf{T} \mathbf{T}^\dagger \mathbf{C}, \quad (9)$$

and the second onto the subset $\{\mathbf{C} \in \mathbb{C}^{T \times F} \mid |\mathbf{C}| = \hat{\mathbf{S}}\}$ of spectrograms with magnitude equal to $\hat{\mathbf{S}}$:

$$P_{|\cdot|=\hat{\mathbf{S}}}(\mathbf{C}) = \hat{\mathbf{S}} \odot \frac{\mathbf{C}}{|\mathbf{C}|}, \quad (10)$$

where \odot, \div are element-wise product and division respectively. The algorithm is initialized with $\mathbf{C}_0 = \hat{\mathbf{S}} \odot e^{i\Phi_0}$ where Φ_0 is typically random. The k -th iteration can be summarized as:

$$\mathbf{C}_k = P_{\mathcal{C}} \circ P_{|\cdot|=\hat{\mathbf{S}}}(\mathbf{C}_{k-1}). \quad (11)$$

Note that in practice, we use the Fast Griffin-Lim Algorithm (FGLA) [26], a slightly modified version of GLA which accelerates the optimization and can be readily integrated in the modifications of the diffusion model based on GLA described below.

3.2. GLA-Grad: Griffin-Lim corrected diffusion model

We propose to introduce an inference-time improvement to a trained WaveGrad model. In current diffusion models such as WaveGrad, the iteration process (cf. Eq. (6)) may lead to a signal \mathbf{y}_{n-1} that is

out of distribution of the training data if not carefully considered, as introduced in [27]. We believe that only conditioning on $\tilde{\mathbf{X}}$ may not be sufficient to solve the problem and a stronger constraint could be beneficial. Our goal is to use GLA to fix the bias between the generated signal and our expected signal (see Fig. 1). We propose to introduce GLA into each step of the diffusion process, in which an estimate of the desired magnitude spectrogram is used to guide the generation of the signal estimate \mathbf{y}_{n-1} at step $n-1$ by minimizing the disparity with its magnitude $|\mathbf{T}\mathbf{y}_{n-1}|$, enabling the inference to have a better convergence. We do not have access to the desired magnitude, but can consider the magnitude spectrogram estimate $\hat{\mathbf{S}}$ obtained from the desired mel spectrogram $\tilde{\mathbf{X}}$ through the pseudo-inverse \mathbf{M}^\dagger of the mel transform as $\hat{\mathbf{S}} = \mathbf{M}^\dagger \tilde{\mathbf{X}}$. At diffusion time step $n-1$, we can then use GLA with the magnitude spectrogram $\hat{\mathbf{S}}$ as desired magnitude.

Starting from the current diffusion estimate \mathbf{y}_{n-1} obtained using WaveGrad with Eq. (6), we can introduce a Griffin-Lim correction leading to a new estimate $\mathbf{y}_{n-1}^{\text{GLA}}$ defined as:

$$\mathbf{y}_{n-1}^{\text{GLA}} = \mathbf{T}^\dagger (P_C \circ P_{|\cdot|=\hat{\mathbf{S}}})^K (\mathbf{T}\mathbf{y}_{n-1}), \quad (12)$$

where $(P_C \circ P_{|\cdot|=\hat{\mathbf{S}}})^K$ indicates the application of K iterations of GLA, starting from the spectrogram $\mathbf{T}\mathbf{y}_{n-1}$ of \mathbf{y}_{n-1} as initial value.

In practical implementation, our algorithm comprises two distinct phases, as depicted in Fig. 1, showcasing the entire workflow. The initial phase employs the GLA-corrected approach outlined in this subsection at each sampling step of the diffusion process. More precisely, we obtain \mathbf{y}_{n-1} using the original WaveGrad update in Eq. (6), then correct it with GLA to obtain $\mathbf{y}_{n-1}^{\text{GLA}}$ using Eq. (12). We posit that commencing the diffusion process with a projection method like GLA aids in maintaining signal consistency. Subsequently, the second phase involves a conventional diffusion process, exclusively considering the diffusion update. The proposed overall scheme is referred to as GLA-Grad.

4. EVALUATION

4.1. Datasets and Metrics

For our experiments, we selected two widely used speech datasets, namely:

- *LJ Speech dataset* [28], a public collection of 13,100 short audio clips featuring a single speaker reading passages from 7 non-fiction books. The duration of the clips ranges from 1 to 10 seconds, resulting in a total length of approximately 24 hours. The sampling frequency is 22050 Hz. We employ 150 files for testing and allocate the remainder for training.
- *Centre for Speech Technology Voice Cloning Toolkit (VCTK)* dataset [29], a speech dataset spoken by 109 native English speakers representing various accents. Each speaker was prompted to read approximately 400 sentences. We resample the VCTK dataset at the same sample rate as LJ Speech.

For all datasets, we use an STFT with 2048 FFT points, a hop size of 300, and a Hann window of length 1200. For the mel filterbank, 128 filterbanks were considered with a lower frequency cutoff at 20 Hz. The mel spectrograms were computed from the ground-truth audio.

The three metrics used for evaluation are:

- *Perceptual Evaluation of Speech Quality (PESQ)* [30], a family of standards comprising a test methodology for automated assessment of speech quality as experienced by a user of a telephony system.

- *Short Term Objective Intelligibility measure (STOI)* [31], a metric to predict intelligibility of (quite) noisy speech.
- *WARP-Q* [32], an objective speech quality metric based on a subsequence dynamic time warping (SDTW) algorithm which yields a raw quality score of the similarity between the ground truth and the generated speech signal.

4.2. Methods

The four models considered for this evaluation include our GLA-Grad model and three baseline algorithms, namely WaveGrad, SpecGrad, and the original Griffin-Lim Algorithm (GLA) applied to the magnitude spectrogram obtained as pseudo-inverse of the conditioning mel spectrogram. We provide further implementation details on all four algorithms below.

WaveGrad: We train and evaluate a WaveGrad Base model [17]. The noise schedule is set to the WaveGrad 6-step schedule, which is denoted as WG-6 and specified as $[7 \times 10^{-6}, 1.4 \times 10^{-4}, 2.1 \times 10^{-3}, 2.8 \times 10^{-2}, 3.5 \times 10^{-1}, 7 \times 10^{-1}]$. This schedule is also employed and evaluated by SpecGrad [20]. The rest of the parameters (batch size, sample rate) are all set as given in [17] for the original WaveGrad model. We also report results with the much longer 50-step inference schedule WG-50 [17], to investigate whether a longer reverse diffusion process can lead to better results, at the expense of speed. This latter model is referred to as WaveGrad-50 in the results.

SpecGrad: For fair comparison, the same parameters as WaveGrad are used for training SpecGrad, and we keep the same noise-shaping parameters as in the original SpecGrad algorithm. SpecGrad actually introduced conditioning information into the sampling process since the noise-shaping could be regarded as a reconstruction. Its goal is thus somehow aligned with our approach for minimizing the conditioning error. The same WG-6 inference schedule is used, as in [20].

Griffin-Lim Algorithm: We employ the Fast Griffin-Lim Algorithm [26] for phase reconstruction, whose implementation is available in torchaudio. A total of 1000 iterations is used to estimate the phase and get the reconstructed signal.

GLA-Grad: GLA-Grad itself does not require training, and is built on top of a WaveGrad model trained as described above. The same WG-6 noise schedule is used in this experiment for comparing with WaveGrad and SpecGrad. We find that the alternative projection cannot keep correcting the signal in the entire reverse process. On the contrary, the projection may be harmful when the quality of the generated wave is high enough. Thus, we only apply the correction at the first 3 reverse steps, then using the unmodified diffusion update for the remaining steps. A total of $K = 32$ alternative projections are used in Eq. 12 for the correction during each reverse step in the first phase. We use the pseudo-inverse of the mel matrix to convert mel spectrogram to magnitude STFT.

Audio examples as well as the full code of our model released as open-source software are available on our demo page¹.

4.3. Experiments and discussion

We propose several experiments with the objective to assess the performance and effectiveness of the proposed approach in addressing generation efficiency, in particular for speech generation of unseen speakers at training time.

We conducted experiments under three setups, training WaveGrad and SpecGrad and estimating all methods in each scenario:

¹<https://GLA-Grad.github.io/>

Table 1: Results when training and evaluating on LJ Speech

Model	PESQ (\uparrow)	STOI (\uparrow)	WARP-Q (\downarrow)
GLA-Grad	3.46 ± 0.11	0.963 ± 0.005	1.677 ± 0.076
WaveGrad	3.59 ± 0.13	0.970 ± 0.004	1.654 ± 0.075
WaveGrad-50	3.72 ± 0.11	0.978 ± 0.004	1.363 ± 0.054
SpecGrad	3.62 ± 0.14	0.963 ± 0.005	1.408 ± 0.054
Griffin-lim	1.02 ± 0.00	0.565 ± 0.042	3.234 ± 0.118

Table 2: Results when training on LJ Speech and evaluating on 19 speakers of VCTK

Model	PESQ (\uparrow)	STOI (\uparrow)	WARP-Q (\downarrow)
GLA-Grad	2.73 ± 0.28	0.944 ± 0.017	1.722 ± 0.132
WaveGrad	2.08 ± 0.31	0.873 ± 0.035	1.913 ± 0.128
WaveGrad-50	2.00 ± 0.29	0.670 ± 0.111	2.122 ± 0.411
SpecGrad	2.48 ± 0.38	0.812 ± 0.066	1.593 ± 0.103
Griffin-lim	1.04 ± 0.01	0.522 ± 0.098	3.411 ± 0.164

- *Closed single speaker:* We train and evaluate on the LJ Speech Dataset, using the isolated validation set for test.
- *Generalization from a single speaker:* The models trained on LJ Speech as above are evaluated on the VCTK multi-speaker dataset for testing the generalization capability of the model from a single speaker’s data. We select 19 speakers from VCTK and use 150 utterances for evaluation.
- *Adaptation to new speakers:* We train and evaluate the models on VCTK, for further evaluating their domain adaption ability when given a larger variety of training data. The speech recordings of 19 Speakers (7400 utterances in total) are used as the training set, while we pick 150 utterances from the remaining 90 speakers for evaluation.

WaveGrad and SpecGrad slightly outperform GLA-Grad on LJ Speech (see Table 1), which confirms their performance on speakers known at training time. In our experiments, we observed that the performance of our model varies with the number of GLA steps in inference, especially for the steps near y_0 . An optimisation of this amount of GLA steps should help the model be more stable. Note though that GLA-Grad has the lowest standard deviation on PESQ, indicating a superior stability in output quality.

When considering generalization from a single speaker, GLA-Grad outperforms both WaveGrad and SpecGrad (see Table 2). SpecGrad outperforms other models on WARP-Q metric, while GLA-Grad has far higher score on PESQ and STOI. As in the closed speaker setup, GLA-Grad has the lowest standard deviation on both PESQ and STOI. It can be observed that the performance of WaveGrad and SpecGrad drops dramatically due to the external evaluation data set which includes more speakers, unseen at training, whereas GLA-Grad sees a much smaller decrease in performance. The results for adaptation to new speakers on VCTK, presented in Table 3, show that GLA-Grad also outperforms other models for a larger unseen speaker set, and that its generalization capability is not limited to the extreme scenario where only data from a single speaker is used for training.

Looking at the results of the more computationally intensive WaveGrad-50 across the various settings, we see that it does improve performance in the closed single-speaker setting, but actually resulted in a significantly decreased performance in the open speaker settings, the longer inference process harming the generalization ability of the WaveGrad model.

Table 3: Results when training on 19 speakers of VCTK and evaluating on 90 other speakers of VCTK

Model	PESQ (\uparrow)	STOI (\uparrow)	WARP-Q (\downarrow)
GLA-Grad	2.88 ± 0.44	0.856 ± 0.081	1.520 ± 1.102
WaveGrad	2.09 ± 0.48	0.803 ± 0.076	1.801 ± 0.133
WaveGrad-50	1.99 ± 0.38	0.706 ± 0.093	2.024 ± 0.157
SpecGrad	2.56 ± 0.35	0.814 ± 0.080	1.492 ± 0.127
Griffin-lim	1.04 ± 0.02	0.542 ± 0.112	3.410 ± 0.169

Table 4: Inference speed compared to real-time (\uparrow)

Model	LJ→LJ	LJ→VCTK	VCTK→VCTK
GLA-Grad	39.0	25.1	24.1
WaveGrad	54.3	55.7	56.1
WaveGrad-50	7.3	7.5	7.5
SpecGrad	40.5	42.7	43.3
Griffin-lim	26.4	11.6	10.8

In summary, it can be noticed that, while WaveGrad and SpecGrad perform best in the closed single-speaker setting, GLA-Grad clearly outperforms all baselines in cross-domain situations where there exists important speaker differences between training and evaluation sets. This indicates that our Griffin-Lim corrected WaveGrad model exhibits stronger generalization performance to speakers who are unseen at training, without needing fine-tuning or modification to the training.

4.4. Complexity

Table 4 shows the inference speed compared to real-time evaluated with 150 utterances in our 3 setups on an NVIDIA A100 GPU. In terms of complexity, the proposed model is comparable in speed with SpecGrad on LJ Speech and somewhat slower on VCTK, while the difference with the baseline WaveGrad model is higher. The computation of the pseudo-inverse of the mel spectrogram accounts for a significant amount of the speed difference between the methods, while the extra Griffin-Lim steps performed at a given diffusion step are not that computationally intensive, and are in fact faster than the diffusion model sampling iteration itself. Adding these Griffin-Lim iterations in GLA-Grad does reduce speed, but it is both an effective and efficient way to improve performance, as adding more diffusion steps has a higher cost in terms of speed, while potentially leading to worse performance as shown in the results of WaveGrad-50. The proposed GLA-Grad thus provides a good trade-off between quality and inference speed.

5. CONCLUSION & FUTURE WORKS

In this paper, we proposed GLA-Grad, a novel scheme for diffusion-based generation of speech from mel spectrogram. Our experiments show in particular that our model clearly outperforms several baselines for unseen speakers at training. These results are obtained by refining the spectrogram phase for the first steps of the diffusion process using a Griffin-Lim algorithm. This shows that it is beneficial to incorporate a phase retrieval module within the diffusion process, which opens new perspectives for diffusion-based speech generation. Future work will be dedicated to the extension of our model with recent and more accurate phase retrieval methods which should further help for speech generation of unseen speakers.

6. REFERENCES

- [1] H. Lu, Z. Wu, X. Wu, X. Li, S. Kang, X. Liu *et al.*, “VAENAR-TTS: Variational Auto-Encoder Based Non-AutoRegressive Text-to-Speech Synthesis,” in *Proc. Interspeech*, 2021, pp. 3775–3779.
- [2] A. van den Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves *et al.*, “Wavenet: A generative model for raw audio,” in *Proc. ISCA SSW*, 2016, pp. 125–125.
- [3] R. Prenger, R. Valle, and B. Catanzaro, “Waveglow: A flow-based generative network for speech synthesis,” in *Proc. ICASSP*, 2019, pp. 3617–3621.
- [4] D. Rezende and S. Mohamed, “Variational inference with normalizing flows,” in *Proc. ICML*, 2015, pp. 1530–1538.
- [5] J. Kong, J. Kim, and J. Bae, “HiFi-GAN: Generative adversarial networks for efficient and high fidelity speech synthesis,” *Proc. NeurIPS*, vol. 33, pp. 17 022–17 033, 2020.
- [6] I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair *et al.*, “Generative adversarial networks,” *arXiv preprint arXiv:1406.2661*, 2014.
- [7] J. Ho, A. Jain, and P. Abbeel, “Denosing diffusion probabilistic models,” *Proc. NeurIPS*, vol. 33, pp. 6840–6851, 2020.
- [8] Y. Song, J. Sohl-Dickstein, D. P. Kingma, A. Kumar, S. Ermon, and B. Poole, “Score-based generative modeling through stochastic differential equations,” in *Proc. ICLR*, 2021.
- [9] H. Huang, P. S. Yu, and C. Wang, “An introduction to image synthesis with generative adversarial nets,” *arXiv preprint arXiv:1803.04469*, 2018.
- [10] M.-Y. Liu, X. Huang, J. Yu, T.-C. Wang, and A. Mallya, “Generative adversarial networks for image and video synthesis: Algorithms and applications,” *Proc. IEEE*, 2021.
- [11] P. Dhariwal and A. Nichol, “Diffusion models beat gans on image synthesis,” *Proc. NeurIPS*, vol. 34, pp. 8780–8794, 2021.
- [12] J.-M. Lemercier, J. Richter, S. Welker, and T. Gerkmann, “Storm: A diffusion-based stochastic regeneration model for speech enhancement and dereverberation,” *IEEE Trans. Audio, Speech, Lang. Process.*, 2023.
- [13] J. Richter, S. Welker, J.-M. Lemercier, B. Lay, and T. Gerkmann, “Speech enhancement and dereverberation with diffusion-based generative models,” *IEEE Trans. Audio, Speech, Lang. Process.*, 2023.
- [14] Y.-J. Lu, Z.-Q. Wang, S. Watanabe, A. Richard, C. Yu, and Y. Tsao, “Conditional diffusion probabilistic model for speech enhancement,” in *Proc. ICASSP*, 2022, pp. 7402–7406.
- [15] R. Scheibler, Y. Ji, S.-W. Chung, J. Byun, S. Choe, and M.-S. Choi, “Diffusion-based generative speech source separation,” in *Proc. ICASSP*, 2023, pp. 1–5.
- [16] B. Chen, C. Wu, and W. Zhao, “Sepdiff: Speech separation based on denosing diffusion model,” in *Proc. ICASSP*, 2023, pp. 1–5.
- [17] N. Chen, Y. Zhang, H. Zen, R. J. Weiss, M. Norouzi, and W. Chan, “Wavegrad: Estimating gradients for waveform generation,” in *Proc. ICLR*, 2020.
- [18] Z. Kong, W. Ping, J. Huang, K. Zhao, and B. Catanzaro, “Diffwave: A versatile diffusion model for audio synthesis,” *Proc. ICLR*, 2021.
- [19] S.-g. Lee, H. Kim, C. Shin, X. Tan, C. Liu, Q. Meng *et al.*, “Priorgrad: Improving conditional denosing diffusion models with data-dependent adaptive prior,” *Proc. ICLR*, 2021.
- [20] Y. Koizumi, H. Zen, K. Yatabe, N. Chen, and M. Bacchiani, “SpecGrad: Diffusion probabilistic model based neural vocoder with adaptive noise spectral shaping,” in *Proc. Interspeech*, 2022.
- [21] Z. Wang, Y. Jiang, H. Zheng, P. Wang, P. He, Z. Wang *et al.*, “Patch diffusion: Faster and more data-efficient training of diffusion models,” *arXiv preprint arXiv:2304.12526*, 2023.
- [22] R. Huang, M. W. Lam, J. Wang, D. Su, D. Yu, Y. Ren *et al.*, “Fastdiff: A fast conditional diffusion model for high-quality speech synthesis,” in *Proc. IJCAI*, 2022.
- [23] A. Farahani, S. Voghoei, K. Rasheed, and H. R. Arabnia, “A brief review of domain adaptation,” *Proc. ICDATA and IKE*, pp. 877–894, 2021.
- [24] D. Griffin and J. Lim, “Signal estimation from modified short-time fourier transform,” *IEEE Transactions on acoustics, speech, and signal processing*, vol. 32, no. 2, pp. 236–243, 1984.
- [25] J. Le Roux, N. Ono, and S. Sagayama, “Explicit consistency constraints for STFT spectrograms and their application to phase reconstruction,” in *Proc. SAPA*, Sep. 2008, pp. 23–28.
- [26] N. Perraudin, P. Balazs, and P. L. Søndergaard, “A fast griffinlim algorithm,” in *Proc. WASPAA*, 2013, pp. 1–4.
- [27] M. S. Graham, W. H. Pinaya, P.-D. Tudosiu, P. Nachev, S. Ourselin, and J. Cardoso, “Denosing diffusion models for out-of-distribution detection,” in *Proc. CVPR*, 2023, pp. 2947–2956.
- [28] K. Ito and L. Johnson, “The lj speech dataset,” <https://keithito.com/LJ-Speech-Dataset/>, 2017.
- [29] J. Yamagishi, C. Veaux, and K. MacDonald, “CSTR VCTK Corpus: English multi-speaker corpus for CSTR voice cloning toolkit (version 0.92),” University of Edinburgh. The Centre for Speech Technology Research (CSTR), Tech. Rep., 2019.
- [30] 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 *Proc. ICASSP*, vol. 2, 2001, pp. 749–752 vol.2.
- [31] C. H. Taal, R. C. Hendriks, R. Heusdens, and J. Jensen, “A short-time objective intelligibility measure for time-frequency weighted noisy speech,” in *Proc. ICASSP*, 2010, pp. 4214–4217.
- [32] W. Jassim, J. Skoglund, M. Chinen, and A. Hines, “Speech quality assessment with WARP-Q: From similarity to sub-sequence dynamic time warp cost,” *IET Signal Processing*, vol. 16, 08 2022.