

DIFFUSION-BASED VOICE CONVERSION WITH FAST MAXIMUM LIKELIHOOD SAMPLING SCHEME

Vadim Popov, Ivan Vovk, Vladimir Gogoryan,

Huawei Noah's Ark Lab, Moscow, Russia

{vadim.popov, vovk.ivan, gogoryan.vladimir}@huawei.com

Taslima Sadekova, Mikhail Kudinov & Jiansheng Wei

Huawei Noah's Ark Lab, Moscow, Russia

{sadekova.taslima, kudinov.mikhail, weijiansheng}@huawei.com

ABSTRACT

Voice conversion is a common speech synthesis task which can be solved in different ways depending on a particular real-world scenario. The most challenging one often referred to as one-shot many-to-many voice conversion consists in copying target voice from only one reference utterance in the most general case when both source and target speakers do not belong to the training dataset. We present a scalable high-quality solution based on diffusion probabilistic modeling and demonstrate its superior quality compared to state-of-the-art one-shot voice conversion approaches. Moreover, focusing on real-time applications, we investigate general principles which can make diffusion models faster while keeping synthesis quality at a high level. As a result, we develop a novel Stochastic Differential Equations solver suitable for various diffusion model types and generative tasks as shown through empirical studies and justify it by theoretical analysis.

1 INTRODUCTION

Voice conversion (VC) is the task of copying the target speaker's voice while preserving the linguistic content of the utterance pronounced by the source speaker. Practical VC applications often require a model which is able to operate in one-shot mode (i.e. when only one reference utterance is provided to copy the target speaker's voice) for any source and target speakers. Such models are usually referred to as one-shot many-to-many models (or sometimes zero-shot many-to-many models, or just any-to-any VC models). It is challenging to build such a model since it should be able to adapt to a new unseen voice having only one spoken utterance pronounced with it, so it was not until recently that successful one-shot VC solutions started to appear.

Conventional one-shot VC models are designed as autoencoders whose latent space ideally contains only the linguistic content of the encoded utterance while target voice identity information (usually taking shape of speaker embedding) is fed to the decoder as conditioning. Whereas in the pioneering AutoVC model (Qian et al., 2019) only speaker embedding from the pre-trained speaker verification network was used as conditioning, several other models improved on AutoVC enriching conditioning with phonetic features such as pitch and loudness (Qian et al., 2020; Nercessian, 2020), or training voice conversion and speaker embedding networks jointly (Chou & Lee, 2019). Also, several papers (Lin et al., 2021; Ishihara & Saito, 2020; Liu et al., 2021b) made use of attention mechanism to better fuse specific features of the reference utterance into the source utterance thus improving the decoder performance. Apart from providing the decoder with sufficiently rich information, one of the main problems autoencoder VC models face is to disentangle source speaker identity from speech content in the encoder. Some models (Qian et al., 2019; 2020; Nercessian, 2020) solve this problem by introducing an information bottleneck. Among other popular solutions of the disentanglement problem one can mention applying vector quantization technique to the content information (Wu et al., 2020; Wang et al., 2021), utilizing features of Variational AutoEncoders (Luong & Tran, 2021; Saito et al., 2018; Chou & Lee, 2019), introducing instance normalization layers (Chou & Lee, 2019; Chen et al., 2021b), and using Phonetic Posteriorgrams (PPGs) (Nercessian, 2020; Liu et al., 2021b).

The model we propose in this paper solves the disentanglement problem by employing the encoder predicting “average voice”: it is trained to transform mel features corresponding to each phoneme into mel features corresponding to this phoneme averaged across a large multi-speaker dataset. As for decoder, in our VC model, it is designed as a part of a Diffusion Probabilistic Model (DPM) since this class of generative models has shown very good results in speech-related tasks like raw waveform generation (Chen et al., 2021a; Kong et al., 2021) and mel feature generation (Popov et al., 2021; Jeong et al., 2021). However, this decoder choice poses a problem of slow inference because DPM forward pass scheme is iterative and to obtain high-quality results it is typically necessary to run it for hundreds of iterations (Ho et al., 2020; Nichol & Dhariwal, 2021). Addressing this issue, we develop a novel inference scheme that significantly reduces the number of iterations sufficient to produce samples of decent quality and does not require model re-training. Although several attempts have been recently made to reduce the number of DPM inference steps (Song et al., 2021a; San-Roman et al., 2021; Watson et al., 2021; Kong & Ping, 2021; Chen et al., 2021a), most of them apply to some particular types of DPMs. In contrast, our approach generalizes to all popular kinds of DPMs and has a strong connection with likelihood maximization.

This paper has the following structure: in Section 2 we present a one-shot many-to-many VC model and describe DPM it relies on; Section 3 introduces a novel DPM sampling scheme and establishes its connection with likelihood maximization; the experiments regarding voice conversion task as well as those demonstrating the benefits of the proposed sampling scheme are described in Section 4; we conclude in Section 5.

2 VOICE CONVERSION DIFFUSION MODEL

As with many other VC models, the one we propose belongs to the family of autoencoders. In fact, any conditional DPM with data-dependent prior (i.e. terminal distribution of forward diffusion) can be seen as such: forward diffusion gradually adding Gaussian noise to data can be regarded as encoder while reverse diffusion trying to remove this noise acts as a decoder. DPMs are trained to minimize the distance (expressed in different terms for different model types) between the trajectories of forward and reverse diffusion processes thus, speaking from the perspective of autoencoders, minimizing reconstruction error. Data-dependent priors have been proposed by Popov et al. (2021) and Lee et al. (2021), and we follow the former paper due to the flexibility of the continuous DPM framework used there. Our approach is summarized in Figure 1.

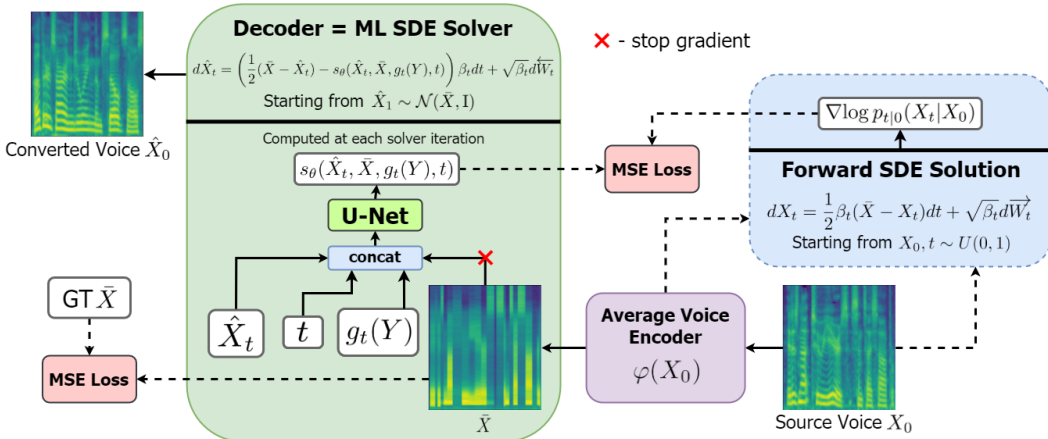


Figure 1: VC model training and inference. Y stands for the training mel-spectrogram at training and the target mel-spectrogram at inference. Speaker conditioning in the decoder is enabled by the speaker conditioning network $g_t(Y)$ where $Y = \{Y_t\}_{t \in [0,1]}$ is the whole forward diffusion trajectory starting at Y_0 . Dotted arrows denote operations performed only at training.

2.1 ENCODER

We choose average phoneme-level mel features as speaker-independent speech representation. To train the encoder to convert input mel-spectrograms into those of “average voice”, we take three steps: (i) first, we apply Montreal Forced Aligner (McAuliffe et al., 2017) to a large-scale multi-speaker LibriTTS dataset (Zen et al., 2019) to align speech frames with phonemes; (ii) next, we obtain average mel features for each particular phoneme by aggregating its mel features across the whole LibriTTS dataset; (iii) the encoder is then trained to minimize mean square error between output mel-spectrograms and ground truth “average voice” mel-spectrograms (i.e. input mel-spectrograms where each phoneme mel feature is replaced with the average one calculated on the previous step).

The encoder has exactly the same Transformer-based architecture used in Grad-TTS (Popov et al., 2021) except that its inputs are mel features rather than character or phoneme embeddings. Note that unlike Grad-TTS the encoder is trained separately from the decoder described in the next section.

2.2 DECODER

Whereas the encoder parameterizes the terminal distribution of the forward diffusion (i.e. the prior), the reverse diffusion is parameterized with the decoder. Following Song et al. (2021c) we use Itô calculus and define diffusions in terms of stochastic processes rather than discrete-time Markov chains.

The general DPM framework we utilize consists of forward and reverse diffusions given by the following Stochastic Differential Equations (SDEs):

$$dX_t = \frac{1}{2}\beta_t(\bar{X} - X_t)dt + \sqrt{\beta_t}d\vec{W}_t, \quad (1)$$

$$d\hat{X}_t = \left(\frac{1}{2}(\bar{X} - \hat{X}_t) - s_\theta(\hat{X}_t, \bar{X}, t) \right) \beta_t dt + \sqrt{\beta_t}d\overleftarrow{W}_t, \quad (2)$$

where $t \in [0, 1]$, \vec{W} and \overleftarrow{W} are two independent Wiener processes in \mathbb{R}^n , β_t is non-negative function referred to as *noise schedule*, s_θ is the score function with parameters θ and \bar{X} is n -dimensional vector. It can be shown (Popov et al., 2021) that the forward SDE (1) allows for explicit solution:

$$\text{Law}(X_t|X_0) = \mathcal{N}\left(e^{-\frac{1}{2}\int_0^t \beta_s ds} X_0 + \left(1 - e^{-\frac{1}{2}\int_0^t \beta_s ds}\right) \bar{X}, \left(1 - e^{-\int_0^t \beta_s ds}\right) \mathbf{I}\right), \quad (3)$$

where \mathbf{I} is $n \times n$ identity matrix. Thus, if noise follows linear schedule $\beta_t = \beta_0 + t(\beta_1 - \beta_0)$ for β_0 and β_1 such that $e^{-\int_0^1 \beta_s ds}$ is close to zero, then $\text{Law}(X_1)$ is close to $\mathcal{N}(\bar{X}, \mathbf{I})$ which is the prior in this DPM. The reverse diffusion (2) is trained by minimizing weighted L_2 loss:

$$\theta^* = \arg \min_{\theta} \mathcal{L}(\theta) = \arg \min_{\theta} \int_0^1 \lambda_t \mathbb{E}_{X_0, X_t} \|s_\theta(X_t, \bar{X}, t) - \nabla \log p_{t|0}(X_t|X_0)\|_2^2 dt, \quad (4)$$

where $p_{t|0}(X_t|X_0)$ is the probability density function (pdf) of the conditional distribution (3) and $\lambda_t = 1 - e^{-\int_0^t \beta_s ds}$. The distribution (3) is Gaussian, so we have

$$\nabla \log p_{t|0}(X_t|X_0) = -\frac{X_t - X_0 e^{-\frac{1}{2}\int_0^t \beta_s ds} - \bar{X}(1 - e^{-\frac{1}{2}\int_0^t \beta_s ds})}{1 - e^{-\int_0^t \beta_s ds}}. \quad (5)$$

At training, time variable t is sampled uniformly from $[0, 1]$, noisy samples X_t are generated according to the formula (3) and the formula (5) is used to calculate loss function \mathcal{L} on these samples. Note that X_t can be sampled without the necessity to calculate intermediate values $\{X_s\}_{0 < s < t}$ which makes optimization task (4) time and memory efficient. A well-trained reverse diffusion (2) has trajectories that are close to those of the forward diffusion (1), so generating data with this DPM can be performed by sampling \hat{X}_1 from the prior $\mathcal{N}(\bar{X}, \mathbf{I})$ and solving SDE (2) backwards in time.

The described above DPM was introduced by Popov et al. (2021) for text-to-speech task and we adapt it for our purposes. We put $\bar{X} = \varphi(X_0)$ where φ is the encoder, i.e. \bar{X} is the ‘‘average voice’’ mel-spectrogram which we want to transform into that of the target voice. We condition the decoder $s_\theta = s_\theta(\hat{X}_t, \bar{X}, g_t(Y), t)$ on some trainable function $g_t(Y)$ to provide it with information about the target speaker (Y stands for forward trajectories of the target mel-spectrogram at inference and the ones of the training mel-spectrogram at training). This function is a neural network trained jointly with the decoder. We experimented with three input types for this network:

- *d-only* – the input is the speaker embedding extracted from the target mel-spectrogram Y_0 with the pre-trained speaker verification network employed in (Jia et al., 2018);
- *wodyn* – in addition, the noisy target mel-spectrogram Y_t is used as input;
- *whole* – in addition, the whole dynamics of the target mel-spectrogram under forward diffusion $\{Y_s | s = 0.5/15, 1.5/15, \dots, 14.5/15\}$ is used as input.

The decoder architecture is based on U-Net (Ronneberger et al., 2015) and is the same as in Grad-TTS but with four times more channels to better capture the whole range of human voices. The speaker conditioning network $g_t(Y)$ is composed of 2D convolutions and MLPs and described in detail in Appendix H. Its output is 128-dimensional vector which is broadcast-concatenated to the concatenation of \hat{X}_t and \bar{X} as additional 128 channels.

2.3 RELATED VC MODELS

To the best of our knowledge, there exist two diffusion-based voice conversion models: VoiceGrad (Kameoka et al., 2020) and DiffSVC (Liu et al., 2021a). The one we propose differs from them in several important aspects. First, neither of the mentioned papers considers a one-shot many-to-many voice conversion scenario. Next, these models take no less than 100 reverse diffusion steps at inference while we pay special attention to reducing the number of iterations (see Section 3) achieving good quality with as few as 6 iterations. Furthermore, VoiceGrad performs voice conversion by running Langevin dynamics starting from the source mel-spectrogram, thus implicitly assuming that forward diffusion trajectories starting from the mel-spectrogram we want to synthesize are likely to pass through the neighborhood of the source mel-spectrogram on their way to Gaussian noise. Such an assumption allowing to have only one network instead of encoder-decoder architecture is too strong and hardly holds for real voices. Finally, DiffSVC performs singing voice conversion and relies on PPGs as speaker-independent speech representation.

3 MAXIMUM LIKELIHOOD SDE SOLVER

In this section, we develop a fixed-step first-order reverse SDE solver that maximizes the log-likelihood of sample paths of the forward diffusion. This solver differs from general-purpose Euler-Maruyama SDE solver (Kloeden & Platen, 1992) by infinitesimally small values which can however become significant when we sample from diffusion model using a few iterations.

Consider the following forward and reverse SDEs defined in Euclidean space \mathbb{R}^n for $t \in [0, 1]$:

$$dX_t = -\frac{1}{2}\beta_t X_t dt + \sqrt{\beta_t} d\vec{W}_t \quad (F), \quad d\hat{X}_t = \left(-\frac{1}{2}\beta_t \hat{X}_t - \beta_t s_\theta(\hat{X}_t, t) \right) dt + \sqrt{\beta_t} d\overleftarrow{W}_t \quad (R), \quad (6)$$

where \vec{W} is a forward Wiener process (i.e. its forward increments $\vec{W}_t - \vec{W}_s$ are independent of \vec{W}_s for $t > s$) and \overleftarrow{W} is a backward Wiener process (i.e. backward increments $\overleftarrow{W}_s - \overleftarrow{W}_t$ are independent of \overleftarrow{W}_t for $s < t$). Following Song et al. (2021c) we will call DPM (6) Variance Preserving (VP). For simplicity we will derive maximum likelihood solver for this particular type of diffusion models. The equation (1) underlying VC diffusion model described in Section 2 can be transformed into the equation (6-F) by a constant shift and we will call such diffusion models Mean Reverting Variance Preserving (MR-VP). VP model analysis carried out in this section can be easily extended (see Appendices D, E and F) to MR-VP model as well as to other common diffusion model types such as sub-VP and VE described by Song et al. (2021c).

The forward SDE (6-F) allows for explicit solution:

$$\text{Law}(X_t|X_s) = \mathcal{N}(\gamma_{s,t}X_s, (1 - \gamma_{s,t}^2)\mathbf{I}), \quad \gamma_{s,t} = \exp\left(-\frac{1}{2}\int_s^t \beta_u du\right), \quad (7)$$

for all $0 \leq s < t \leq 1$. This formula is derived by means of Itô calculus in Appendix A. The reverse SDE (6-R) parameterized with a neural network s_θ is trained to approximate gradient of the log-density of noisy data X_t :

$$\theta^* = \arg \min_{\theta} \int_0^1 \lambda_t \mathbb{E}_{X_t} \|s_\theta(X_t, t) - \nabla \log p_t(X_t)\|_2^2 dt, \quad (8)$$

where the expectation is taken with respect to noisy data distribution $\text{Law}(X_t)$ with pdf $p_t(\cdot)$ and λ_t is some positive weighting function. Note that certain Lipschitz constraints should be satisfied by coefficients of SDEs (6) to guarantee existence of strong solutions (Liptser & Shiryaev, 1978), and throughout this section we assume these conditions are satisfied as well as those from (Anderson, 1982) which guarantee that paths \hat{X} generated by the reverse SDE (6-R) for the optimal θ^* equal forward SDE (6-F) paths X in distribution.

The generative procedure of a VP DPM consists in solving the reverse SDE (6-R) backwards in time starting from $\hat{X}_1 \sim \mathcal{N}(0, \mathbf{I})$. Common Euler-Maruyama solver introduces discretization error (Kloeden & Platen, 1992) which may harm sample quality when the number of iterations is small. At the same time, it is possible to design unbiased (Henry-Labordère et al., 2017) or even exact (Beskos & Roberts, 2005) numerical solvers for some particular SDE types. The Theorem 1 shows that in the case of diffusion models we can make use of the forward diffusion (6-F) and propose a reverse SDE solver which is better than the general-purpose Euler-Maruyama one in terms of likelihood.

The solver proposed in the Theorem 1 is expressed in terms of the values defined as follows:

$$\mu_{s,t} = \gamma_{s,t} \frac{1 - \gamma_{0,s}^2}{1 - \gamma_{0,t}^2}, \quad \nu_{s,t} = \gamma_{0,s} \frac{1 - \gamma_{s,t}^2}{1 - \gamma_{0,t}^2}, \quad \sigma_{s,t}^2 = \frac{(1 - \gamma_{0,s}^2)(1 - \gamma_{s,t}^2)}{1 - \gamma_{0,t}^2}, \quad (9)$$

$$\begin{aligned} \kappa_{t,h}^* &= \frac{\nu_{t-h,t}(1 - \gamma_{0,t}^2)}{\gamma_{0,t}\beta_t h} - 1, & \omega_{t,h}^* &= \frac{\mu_{t-h,t} - 1}{\beta_t h} + \frac{1 + \kappa_{t,h}^*}{1 - \gamma_{0,t}^2} - \frac{1}{2}, \\ (\sigma_{t,h}^*)^2 &= \sigma_{t-h,t}^2 + \frac{1}{n} \nu_{t-h,t}^2 \mathbb{E}_{X_t} [\text{Tr}(\text{Var}(X_0|X_t))], \end{aligned} \quad (10)$$

where n is data dimensionality, $\text{Var}(X_0|X_t)$ is the covariance matrix of the conditional data distribution $\text{Law}(X_0|X_t)$ (so, $\text{Tr}(\text{Var}(X_0|X_t))$ is the overall variance across all n dimensions) and the expectation $\mathbb{E}_{X_t}[\cdot]$ is taken with respect to the unconditional noisy data distribution $\text{Law}(X_t)$.

Theorem 1. Consider a DPM characterized by SDEs (6) with reverse diffusion trained till optimality. Let $N \in \mathbb{N}$ be any natural number and $h = 1/N$. Consider the following class of fixed step size h reverse SDE solvers parameterized with triplets of real numbers $\{(\hat{\kappa}_{t,h}, \hat{\omega}_{t,h}, \hat{\sigma}_{t,h}) | t = h, 2h, \dots, 1\}$:

$$\hat{X}_{t-h} = \hat{X}_t + \beta_t h \left(\left(\frac{1}{2} + \hat{\omega}_{t,h} \right) \hat{X}_t + (1 + \hat{\kappa}_{t,h}) s_{\theta^*}(\hat{X}_t, t) \right) + \hat{\sigma}_{t,h} \xi_t, \quad (11)$$

where θ^* is given by (8), $t = 1, 1-h, \dots, h$ and ξ_t are i.i.d. samples from $\mathcal{N}(0, \mathbf{I})$. Then:

- (i) Log-likelihood of sample paths $X = \{X_{kh}\}_{k=0}^N$ under generative model \hat{X} is maximized for $\hat{\kappa}_{t,h} = \kappa_{t,h}^*$, $\hat{\omega}_{t,h} = \omega_{t,h}^*$ and $\hat{\sigma}_{t,h} = \sigma_{t,h}^*$.
- (ii) Assume that the SDE solver (11) starts from random variable $\hat{X}_1 \sim \text{Law}(X_1)$. If X_0 is a constant or a Gaussian random variable with diagonal isotropic covariance matrix (i.e. $\delta^2 \mathbf{I}$ for $\delta > 0$), then generative model \hat{X} is exact for $\hat{\kappa}_{t,h} = \kappa_{t,h}^*$, $\hat{\omega}_{t,h} = \omega_{t,h}^*$ and $\hat{\sigma}_{t,h} = \sigma_{t,h}^*$.

Table 1: Input types for speaker conditioning $g_t(Y)$ compared in terms of speaker similarity.

	Diff-LibriTTS			Diff-VCTK		
	<i>d-only</i>	<i>wodyn</i>	<i>whole</i>	<i>d-only</i>	<i>wodyn</i>	<i>whole</i>
Most similar	27.0%	38.0%	34.1%	27.2%	46.7%	23.6%
Least similar	28.9%	29.3%	38.5%	25.3%	23.9%	48.6%

The Theorem 1 provides an improved DPM sampling scheme which comes at no additional computational cost compared to standard methods (except for data-dependent term in σ^* as discussed in Section 4.3) and requires neither model re-training nor extensive search on noise schedule space. The proof of this theorem is given in Appendix C. Note that it establishes optimality of the reverse SDE solver (11) with the parameters (10) in terms of likelihood of *discrete* paths $X = \{X_{kh}\}_{k=0}^N$ while the optimality of *continuous* model (6-R) on *continuous* paths $\{X_t\}_{t \in [0,1]}$ is guaranteed for a model with parameters $\theta = \theta^*$ as shown in (Song et al., 2021c).

The class of reverse SDE solvers considered in the Theorem 1 is rather broad: it is the class of all fixed-step solvers whose increments at time t are linear combination of \hat{X}_t , $s_\theta(\hat{X}_t, t)$ and Gaussian noise with zero mean and diagonal isotropic covariance matrix. As a particular case it includes Euler-Maruyama solver ($\hat{\kappa}_{t,h} \equiv 0$, $\hat{\omega}_{t,h} \equiv 0$, $\hat{\sigma}_{t,h} \equiv \sqrt{\beta_t h}$) and for fixed t and $h \rightarrow 0$ we have $\kappa_{t,h}^* = \bar{o}(1)$, $\omega_{t,h}^* = \bar{o}(1)$ and $\sigma_{t,h}^* = \sqrt{\beta_t h}(1 + \bar{o}(1))$ (the proof is given in Appendix B), so the optimal SDE solver significantly differs from general-purpose Euler-Maruyama solver only when N is rather small or t has the same order as h , i.e. on the final steps of DPM inference. Appendix G contains toy examples demonstrating the difference of the proposed optimal SDE solver and Euler-Maruyama one depending on step size.

The result (ii) from the Theorem 1 strengthens the result (i) for some particular data distributions, but it may seem useless since in practice data distribution is far from being constant or Gaussian. However, in case of generation with strong conditioning (e.g. mel-spectrogram inversion) the assumptions on the data distribution may become viable: in the limiting case when our model is conditioned on $c = \psi(X_0)$ for an injective function ψ , random variable $X_0|c$ becomes a constant $\psi^{-1}(c)$.

4 EXPERIMENTS

We trained two groups of models: *Diff-VCTK* models on VCTK (Yamagishi et al., 2019) dataset containing 109 speakers (9 speakers were held out for testing purposes) and *Diff-LibriTTS* models on LibriTTS (Zen et al., 2019) containing approximately 1100 speakers (10 speakers were held out). For every model both encoder and decoder were trained on the same dataset. Training hyperparameters, implementation and data processing details can be found in Appendix I. For mel-spectrogram inversion, we used the pre-trained universal HiFi-GAN vocoder (Kong et al., 2020) operating at 22.05kHz. All subjective human evaluation was carried out on Amazon Mechanical Turk (AMT) with Master assessors to ensure the reliability of the obtained Mean Opinion Scores (MOS). In all AMT tests we considered unseen-to-unseen conversion with 25 unseen (for both *Diff-VCTK* and *Diff-LibriTTS*) speakers: 9 VCTK speakers, 10 LibriTTS speakers and 6 internal speakers. For VCTK source speakers we also ensured that source phrases were unseen during training. We place other details of listening AMT tests in Appendix J. A small subset of speech samples used in them is available at our demo page <https://diffvc-fast-ml-solver.github.io> which we encourage to visit. The code will soon be published at <https://github.com/huawei-noah/Speech-Backbones>.

As for sampling, we considered the following class of reverse SDE solvers:

$$\hat{X}_{t-h} = \hat{X}_t + \beta_t h \left(\left(\frac{1}{2} + \hat{\omega}_{t,h} \right) (\hat{X}_t - \bar{X}) + (1 + \hat{\kappa}_{t,h}) s_\theta(\hat{X}_t, \bar{X}, g_t(Y), t) \right) + \hat{\sigma}_{t,h} \xi_t, \quad (12)$$

where $t = 1, 1-h, \dots, h$ and ξ_t are i.i.d. samples from $\mathcal{N}(0, I)$. For $\hat{\kappa}_{t,h} = \kappa_{t,h}^*$, $\hat{\omega}_{t,h} = \omega_{t,h}^*$ and $\hat{\sigma}_{t,h} = \sigma_{t,h}^*$ (where $\kappa_{t,h}^*$, $\omega_{t,h}^*$ and $\sigma_{t,h}^*$ are given by (10)) it becomes maximum likelihood reverse SDE solver for MR-VP DPM (1-2) as shown in Appendix D. In practice it is not trivial to estimate variance of the conditional distribution $\text{Law}(X_0|X_t)$, so we skipped this term in $\sigma_{t,h}^*$,

Table 2: Subjective evaluation (MOS) of one-shot VC models trained on VCTK. Ground truth recordings were evaluated only for VCTK speakers.

	VCTK test (9 speakers, 54 pairs)		Whole test (25 speakers, 350 pairs)	
	Naturalness	Similarity	Naturalness	Similarity
<i>AGAIN-VC</i>	1.98 ± 0.05	1.97 ± 0.08	1.87 ± 0.03	1.75 ± 0.04
<i>FragmentVC</i>	2.20 ± 0.06	2.45 ± 0.09	1.91 ± 0.03	1.93 ± 0.04
<i>VQMVC</i>	2.89 ± 0.06	2.60 ± 0.10	2.48 ± 0.04	1.95 ± 0.04
<i>Diff-VCTK-ML-6</i>	3.73 ± 0.06	3.47 ± 0.09	3.39 ± 0.04	2.69 ± 0.05
<i>Diff-VCTK-ML-30</i>	3.73 ± 0.06	3.57 ± 0.09	3.44 ± 0.04	2.71 ± 0.05
<i>Ground truth</i>	4.55 ± 0.05	4.52 ± 0.07	4.55 ± 0.05	4.52 ± 0.07

Table 3: Subjective evaluation (MOS) of one-shot VC models trained on large-scale datasets.

	VCTK test (9 speakers, 54 pairs)		Whole test (25 speakers, 350 pairs)	
	Naturalness	Similarity	Naturalness	Similarity
<i>Diff-LibriTTS-EM-6</i>	1.68 ± 0.06	1.53 ± 0.07	1.57 ± 0.02	1.47 ± 0.03
<i>Diff-LibriTTS-PF-6</i>	3.11 ± 0.07	2.58 ± 0.11	2.99 ± 0.03	2.50 ± 0.04
<i>Diff-LibriTTS-ML-6</i>	3.84 ± 0.08	3.08 ± 0.11	3.80 ± 0.03	3.27 ± 0.05
<i>Diff-LibriTTS-ML-30</i>	3.96 ± 0.08	3.23 ± 0.11	4.02 ± 0.03	3.39 ± 0.05
<i>BNE-PPG-VC</i>	3.95 ± 0.08	3.27 ± 0.12	3.83 ± 0.03	3.03 ± 0.05

assuming this variance to be rather small because of strong conditioning on $g_t(Y)$ and just used $\hat{\sigma}_{t,h} = \sigma_{t-h,t}$ calling this sampling method *ML-N* ($N = 1/h$ is the number of SDE solver steps). We also experimented with Euler-Maruyama solver *EM-N* (i.e. $\hat{\kappa}_{t,h} = 0, \hat{\omega}_{t,h} = 0, \hat{\sigma}_{t,h} = \sqrt{\beta_t h}$) and “probability flow sampling” from (Song et al., 2021c) which we denote by *PF-N* ($\hat{\kappa}_{t,h} = -0.5, \hat{\omega}_{t,h} = 0, \hat{\sigma}_{t,h} = 0$).

4.1 SPEAKER CONDITIONING ANALYSIS

For each dataset we trained three models – one for each input type for the speaker conditioning network $g_t(Y)$ (see Section 2.2). Although these input types had much influence neither on speaker similarity nor on speech naturalness, we did two experiments to choose the best models (one for each training dataset) in terms of speaker similarity for further comparison with baseline systems. We compared voice conversion results (produced by *ML-30* sampling scheme) on 92 source-target pairs. AMT workers were asked which of three models (if any) sounded most similar to the target speaker and which of them (if any) sounded least similar. For *Diff-VCTK* and *Diff-LibriTTS* models each conversion pair was evaluated 4 and 5 times respectively. Table 1 demonstrates that for both *Diff-VCTK* and *Diff-LibriTTS* the best option is *wodyn*, i.e. to condition the decoder at time t on the speaker embedding together with the noisy target mel-spectrogram Y_t . Conditioning on Y_t allows making use of diffusion-specific information of how the noisy target sounds whereas embedding from the pre-trained speaker verification network contains information only about the clean target. Taking these results into consideration, we used *Diff-VCTK-wodyn* and *Diff-LibriTTS-wodyn* in the remaining experiments.

4.2 ANY-TO-ANY VOICE CONVERSION

We chose four recently proposed VC models capable of one-shot many-to-many synthesis as the baselines:

- *AGAIN-VC* (Chen et al., 2021b), an improved version of a conventional autoencoder AdaIN-VC solving the disentanglement problem by means of instance normalization;
- *FragmentVC* (Lin et al., 2021), an attention-based model relying on wav2vec 2.0 (Baevski et al., 2020) to obtain speech content from the source utterance;

Table 4: Reverse SDE solvers compared in terms of FID. N is the number of SDE solver steps.

	VP DPM		sub-VP DPM		VE DPM	
	$N=10$	$N=100$	$N=10$	$N=100$	$N=10$	$N=100$
Euler-Maruyama	229.6	19.68	312.3	19.83	462.1	24.77
Reverse Diffusion	679.8	65.95	312.2	19.74	461.1	303.2
Probability Flow	88.92	5.70	64.22	4.42	495.3	214.5
Ancestral Sampling	679.8	68.35	—	—	454.7	17.83
Maximum Likelihood ($\tau = 0.1$)	260.3	4.34	317.0	6.63	461.9	23.63
Maximum Likelihood ($\tau = 0.5$)	24.45	7.82	30.90	6.43	462.0	10.07
Maximum Likelihood ($\tau = 1.0$)	41.78	7.94	48.02	6.51	48.51	12.37

- *VQMIVC* (Wang et al., 2021), state-of-the-art approach among those employing vector quantization techniques;
- *BNE-PPG-VC* (Liu et al., 2021b), an improved variant of PPG-based VC models combining a bottleneck feature extractor obtained from a phoneme recognizer with a seq2seq-based synthesis module.

As shown in (Kim et al., 2021), PPG-based VC models provide high voice conversion quality competitive even with that of the state-of-the-art VC models taking text transcription corresponding to the source utterance as input. Therefore, we can consider *BNE-PPG-VC* a state-of-the-art model in our setting.

Baseline voice conversion results were produced by the pre-trained VC models provided in official GitHub repositories. Since only *BNE-PPG-VC* has the model pre-trained on a large-scale dataset (namely, LibriTTS + VCTK), we did two subjective human evaluation tests: the first one comparing *Diff-VCTK* with *AGAIN-VC*, *FragmentVC* and *VQMIVC* trained on VCTK and the second one comparing *Diff-LibriTTS* with *BNE-PPG-VC*. The results of these tests are given in Tables 2 and 3 respectively. Speech naturalness and speaker similarity were assessed separately. AMT workers evaluated voice conversion quality on 350 source-target pairs on 5-point scale. In the first test, each pair was assessed 6 times on average both in speech naturalness and speaker similarity evaluation; as for the second one, each pair was assessed 8 and 9 times on average in speech naturalness and speaker similarity evaluation correspondingly. No less than 41 unique assessors took part in each test.

Table 2 demonstrates that our model performs significantly better than the baselines both in terms of naturalness and speaker similarity even when 6 reverse diffusion iterations are used. Despite working almost equally well on VCTK speakers, the best baseline *VQMIVC* shows poor performance on other speakers perhaps because of not being able to generalize to different domains with lower recording quality. Although *Diff-VCTK* performance also degrades on non-VCTK speakers, it achieves good speaker similarity of MOS 3.6 on VCTK ones when *ML-30* sampling scheme is used and only slightly worse MOS 3.5 when 5x less iterations are used at inference.

Table 3 contains human evaluation results of *Diff-LibriTTS* for four sampling schemes: *ML-30* with 30 reverse SDE solver steps and *ML-6*, *EM-6* and *PF-6* with 6 steps of reverse diffusion. The three schemes taking 6 steps achieved real-time factor (RTF) around 0.1 on GPU (i.e. inference was 10 times faster than real time) while the one taking 30 steps had RTF around 0.5. The proposed model *Diff-LibriTTS-ML-30* and the baseline *BNE-PPG-VC* show the same performance on the VCTK test set in terms of speech naturalness the latter being slightly better in terms of speaker similarity which can perhaps be explained by the fact that *BNE-PPG-VC* was trained on the union of VCTK and LibriTTS whereas our model was trained only on LibriTTS. As for the whole test set containing unseen LibriTTS and internal speakers also, *Diff-LibriTTS-ML-30* outperforms *BNE-PPG-VC* model achieving MOS 4.0 and 3.4 in terms of speech naturalness and speaker similarity respectively. Due to employing PPG extractor trained on a large-scale ASR dataset LibriSpeech (Panayotov et al., 2015), *BNE-PPG-VC* has fewer mispronunciation issues than our model but synthesized speech suffers from more sonic artifacts. This observation makes us believe that incorporating PPG features in the proposed diffusion VC framework is a promising direction for future research.

Table 3 also demonstrates the benefits of the proposed maximum likelihood sampling scheme over other sampling methods for a small number of inference steps: only *ML-N* scheme allows us to

use as few as $N = 6$ iterations with acceptable quality degradation of MOS 0.2 and 0.1 in terms of naturalness and speaker similarity respectively while two other competing methods lead to much more significant quality degradation.

4.3 MAXIMUM LIKELIHOOD SAMPLING



Figure 2: CIFAR-10 images randomly sampled from VP DPM by running 10 reverse diffusion steps with the following schemes (from left to right): “euler-maruyama”, “probability flow”, “maximum likelihood ($\tau = 0.5$)”, “maximum likelihood ($\tau = 1.0$)”.

To show that the maximum likelihood sampling scheme proposed in Section 3 generalizes to different tasks and DPM types, we took the models trained by Song et al. (2021c) on CIFAR-10 image generation task and compared our method with other sampling schemes described in that paper in terms of Fréchet Inception Distance (FID).

The main difficulty in applying maximum likelihood SDE solver is estimating data-dependent term $\mathbb{E}[\text{Tr}(\text{Var}(X_0|X_t))]$ in $\sigma_{t,h}^*$. Although in the current experiments we just set this term to zero, we can think of two possible ways to estimate it: (i) approximate $\text{Var}(X_0|X_t)$ with $\text{Var}(\hat{X}_0|\hat{X}_t = X_t)$: sample noisy data X_t , solve reverse SDE with sufficiently small step size starting from terminal condition $\hat{X}_t = X_t$ several times, and calculate sample variance of the resulting solutions at initial points \hat{X}_0 ; (ii) use the formula (58) from Appendix C to calculate $\text{Var}(X_0|X_t)$ assuming that X_0 is distributed normally with mean and variance equal to sample mean and sample variance computed on the training dataset. Experimenting with these techniques and exploring new ones seems to be an interesting future research direction.

Another important practical consideration is that the proposed scheme is proven to be optimal only for score matching networks trained till optimality. Therefore, in the experiments whose results are reported in Table 4 we apply maximum likelihood sampling scheme only when $t \leq \tau$ while using standard Euler-Maruyama solver for $t > \tau$ for some hyperparameter $\tau \in [0, 1]$. Such a modification relies on the assumption that score matching network is closer to being optimal for smaller noise.

Table 4 shows that despite likelihood and FID are two metrics that do not perfectly correlate (Song et al., 2021b), in most cases our maximum likelihood SDE solver performs best in terms of FID. Also, it is worth mentioning that although $\tau = 1$ is always rather a good choice, tuning this hyperparameter can lead to even better performance. One can find randomly chosen generated images for various sampling methods in Figure 2.

5 CONCLUSION

In this paper, the novel one-shot many-to-many voice conversion model has been presented. Its encoder design and powerful diffusion-based decoder made it possible to achieve good results both in terms of speaker similarity and speech naturalness even on out-of-domain unseen speakers. Subjective human evaluation verified that the proposed model delivers scalable VC solution with competitive performance. Furthermore, aiming at fast synthesis, we have developed and theoretically justified the novel sampling scheme. The main idea behind it is to modify the general-purpose Euler-Maruyama SDE solver so as to maximize the likelihood of discrete sample paths of the forward diffusion. Due to the proposed sampling scheme, our VC model is capable of high-quality voice conversion with as few as 6 reverse diffusion steps. Moreover, experiments on the image generation task show that all known diffusion model types can benefit from the proposed SDE solver.

REFERENCES

- Brian D.O. Anderson. Reverse-time Diffusion Equation Models. *Stochastic Processes and their Applications*, 12(3):313 – 326, 1982. ISSN 0304-4149.
- Alexei Baeovski, Yuhao Zhou, Abdelrahman Mohamed, and Michael Auli. wav2vec 2.0: A Framework for Self-Supervised Learning of Speech Representations. In *Advances in Neural Information Processing Systems*, volume 33, pp. 12449–12460. Curran Associates, Inc., 2020.
- Alexandros Beskos and Gareth O. Roberts. Exact Simulation of Diffusions. *The Annals of Applied Probability*, 15(4):2422–2444, 2005. ISSN 10505164.
- Nanxin Chen, Yu Zhang, Heiga Zen, Ron J Weiss, Mohammad Norouzi, and William Chan. WaveGrad: Estimating Gradients for Waveform Generation. In *International Conference on Learning Representations*, 2021a.
- Yen-Hao Chen, D. Wu, Tsung-Han Wu, and Hung yi Lee. Again-VC: A One-Shot Voice Conversion Using Activation Guidance and Adaptive Instance Normalization. In *ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 5954–5958, 2021b.
- Ju-Chieh Chou and Hung-yi Lee. One-Shot Voice Conversion by Separating Speaker and Content Representations with Instance Normalization. In *Interspeech 2019, 20th Annual Conference of the International Speech Communication Association, Graz, Austria, 15-19 September 2019*, pp. 664–668. ISCA, 2019.
- Pierre Henry-Labordère, Xiaolu Tan, and Nizar Touzi. Unbiased Simulation of Stochastic Differential Equations. *The Annals of Applied Probability*, 27(6):3305–3341, 2017. ISSN 10505164.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising Diffusion Probabilistic Models. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, volume 33. Curran Associates, Inc., 2020.
- Aapo Hyvärinen. Estimation of Non-Normalized Statistical Models by Score Matching. *Journal of Machine Learning Research*, 6(24):695–709, 2005.
- Tatsuma Ishihara and Daisuke Saito. Attention-Based Speaker Embeddings for One-Shot Voice Conversion. In *Interspeech 2020, 21st Annual Conference of the International Speech Communication Association, Virtual Event, Shanghai, China, 25-29 October 2020*, pp. 806–810. ISCA, 2020.
- Myeonghun Jeong, Hyeongju Kim, Sung Jun Cheon, Byoung Jin Choi, and Nam Soo Kim. Diff-TTS: A Denoising Diffusion Model for Text-to-Speech. In *Proc. Interspeech 2021*, pp. 3605–3609, 2021.
- Ye Jia, Yu Zhang, Ron Weiss, Quan Wang, Jonathan Shen, Fei Ren, Zhifeng Chen, Patrick Nguyen, Ruoming Pang, Ignacio Lopez Moreno, and Yonghui Wu. Transfer Learning from Speaker Verification to Multispeaker Text-To-Speech Synthesis. In *Advances in Neural Information Processing Systems 31*, pp. 4480–4490. Curran Associates, Inc., 2018.
- Hirokazu Kameoka, Takuhiro Kaneko, Kou Tanaka, Nobukatsu Hojo, and Shogo Seki. VoiceGrad: Non-Parallel Any-to-Many Voice Conversion with Annealed Langevin Dynamics, 2020.
- Kang-wook Kim, Seung-won Park, and Myun-chul Joe. Assem-VC: Realistic Voice Conversion by Assembling Modern Speech Synthesis Techniques, 2021.
- Diederik P. Kingma, Tim Salimans, Ben Poole, and Jonathan Ho. Variational Diffusion Models, 2021.
- Peter E. Kloeden and Eckhard Platen. *Numerical Solution of Stochastic Differential Equations*, volume 23 of *Stochastic Modelling and Applied Probability*. Springer-Verlag Berlin Heidelberg, 1992.

- Jungil Kong, Jaehyeon Kim, and Jaekyoung Bae. HiFi-GAN: Generative Adversarial Networks for Efficient and High Fidelity Speech Synthesis. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, virtual, 2020*.
- Zhifeng Kong and Wei Ping. On Fast Sampling of Diffusion Probabilistic Models. In *ICML Workshop on Invertible Neural Networks, Normalizing Flows, and Explicit Likelihood Models, 2021*.
- Zhifeng Kong, Wei Ping, Jiaji Huang, Kexin Zhao, and Bryan Catanzaro. DiffWave: A Versatile Diffusion Model for Audio Synthesis. In *International Conference on Learning Representations, 2021*.
- Sang-gil Lee, Heeseung Kim, Chaehun Shin, Xu Tan, Chang Liu, Qi Meng, Tao Qin, Wei Chen, Sungroh Yoon, and Tie-Yan Liu. PriorGrad: Improving Conditional Denoising Diffusion Models with Data-Driven Adaptive Prior, 2021.
- Yist Y. Lin, Chung-Ming Chien, Jheng-Hao Lin, Hung-yi Lee, and Lin-Shan Lee. FragmentVC: Any-To-Any Voice Conversion by End-To-End Extracting and Fusing Fine-Grained Voice Fragments with Attention. In *ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 5939–5943, 2021.
- Robert S. Liptser and Albert N. Shiryaev. *Statistics of Random Processes*, volume 5 of *Stochastic Modelling and Applied Probability*. Springer-Verlag, 1978.
- Songxiang Liu, Yuewen Cao, Dan Su, and Helen Meng. DiffSVC: A Diffusion Probabilistic Model for Singing Voice Conversion, 2021a.
- Songxiang Liu, Yuewen Cao, Disong Wang, Xixin Wu, Xunying Liu, and Helen Meng. Any-to-Many Voice Conversion With Location-Relative Sequence-to-Sequence Modeling. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 29:1717–1728, 2021b.
- Manh Luong and Viet Anh Tran. Many-to-Many Voice Conversion Based Feature Disentanglement Using Variational Autoencoder. In *Proc. Interspeech 2021*, pp. 851–855, 2021.
- Michael McAuliffe, Michaela Socolof, Sarah Mihuc, Michael Wagner, and Morgan Sonderegger. Montreal Forced Aligner: Trainable Text-Speech Alignment Using Kaldi. In *Proc. Interspeech 2017*, pp. 498–502, 2017.
- Shahan Nercessian. Improved Zero-Shot Voice Conversion Using Explicit Conditioning Signals. In *Interspeech 2020, 21st Annual Conference of the International Speech Communication Association, Virtual Event, Shanghai, China, 25-29 October 2020*, pp. 4711–4715. ISCA, 2020.
- Alexander Quinn Nichol and Prafulla Dhariwal. Improved Denoising Diffusion Probabilistic Models. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 8162–8171. PMLR, 18–24 Jul 2021.
- Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur. Librispeech: An ASR Corpus Based on Public Domain Audio Books. In *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 5206–5210, 2015.
- Vadim Popov, Ivan Vovk, Vladimir Gogoryan, Tasnima Sadekova, and Mikhail Kudinov. Grad-TTS: A Diffusion Probabilistic Model for Text-to-Speech. In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, pp. 8599–8608. PMLR, 2021.
- Kaizhi Qian, Yang Zhang, Shiyu Chang, Xuesong Yang, and Mark Hasegawa-Johnson. AutoVC: Zero-Shot Voice Style Transfer with Only Autoencoder Loss. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 5210–5219. PMLR, 09–15 Jun 2019.
- Kaizhi Qian, Zeyu Jin, Mark Hasegawa-Johnson, and Gautham J. Mysore. F0-Consistent Many-To-Many Non-Parallel Voice Conversion Via Conditional Autoencoder. In *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 6284–6288, 2020.

- Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional Networks for Biomedical Image Segmentation. In *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, pp. 234–241. Springer International Publishing, 2015.
- Yuki Saito, Yusuke Ijima, Kyosuke Nishida, and Shinnosuke Takamichi. Non-Parallel Voice Conversion Using Variational Autoencoders Conditioned by Phonetic Posteriorgrams and D-Vectors. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 5274–5278, 2018.
- Robin San-Roman, Eliya Nachmani, and Lior Wolf. Noise Estimation for Generative Diffusion Models, 2021.
- Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising Diffusion Implicit Models. In *International Conference on Learning Representations*, 2021a.
- Yang Song, Conor Durkan, Iain Murray, and Stefano Ermon. Maximum Likelihood Training of Score-Based Diffusion Models, 2021b.
- Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-Based Generative Modeling through Stochastic Differential Equations. In *International Conference on Learning Representations*, 2021c.
- Pascal Vincent. A Connection Between Score Matching and Denoising Autoencoders. *Neural Computation*, 23(7):1661–1674, 2011.
- Disong Wang, Liqun Deng, Yu Ting Yeung, Xiao Chen, Xunying Liu, and Helen Meng. VQMIVC: Vector Quantization and Mutual Information-Based Unsupervised Speech Representation Disentanglement for One-Shot Voice Conversion. In *Proc. Interspeech 2021*, pp. 1344–1348, 2021.
- Daniel Watson, Jonathan Ho, Mohammad Norouzi, and William Chan. Learning to Efficiently Sample from Diffusion Probabilistic Models, 2021.
- Da-Yi Wu, Yen-Hao Chen, and Hung-yi Lee. VQVC+: One-Shot Voice Conversion by Vector Quantization and U-Net Architecture. In Helen Meng, Bo Xu, and Thomas Fang Zheng (eds.), *Interspeech 2020, 21st Annual Conference of the International Speech Communication Association, Virtual Event, Shanghai, China, 25-29 October 2020*, pp. 4691–4695. ISCA, 2020.
- Junichi Yamagishi, Christophe Veaux, and Kirsten MacDonald. CSTR VCTK Corpus: English Multi-speaker Corpus for CSTR Voice Cloning Toolkit (version 0.92), 2019.
- Heiga Zen, Rob Clark, Ron J. Weiss, Viet Dang, Ye Jia, Yonghui Wu, Yu Zhang, and Zhifeng Chen. LibriTTS: A Corpus Derived from LibriSpeech for Text-to-Speech. In *Interspeech*, 2019.

A FORWARD VP SDE SOLUTION

Since function $\gamma_{0,t}^{-1}X_t$ is linear in X_t , taking its differential does not require second order derivative term in Itô's formula:

$$\begin{aligned} d(\gamma_{0,t}^{-1}X_t) &= d\left(e^{\frac{1}{2}\int_0^t\beta_u du}X_t\right) \\ &= e^{\frac{1}{2}\int_0^t\beta_u du} \cdot \frac{1}{2}\beta_t X_t dt + e^{\frac{1}{2}\int_0^t\beta_u du} \cdot \left(-\frac{1}{2}\beta_t X_t dt + \sqrt{\beta_t}d\vec{W}_t\right) \\ &= \sqrt{\beta_t}e^{\frac{1}{2}\int_0^t\beta_u du}d\vec{W}_t. \end{aligned} \quad (13)$$

Integrating this expression from s to t results in an Itô's integral:

$$e^{\frac{1}{2}\int_0^t\beta_u du}X_t - e^{\frac{1}{2}\int_0^s\beta_u du}X_s = \int_s^t \sqrt{\beta_\tau}e^{\frac{1}{2}\int_0^\tau\beta_u du}d\vec{W}_\tau, \quad (14)$$

or

$$X_t = e^{-\frac{1}{2}\int_s^t\beta_u du}X_s + \int_s^t \sqrt{\beta_\tau}e^{-\frac{1}{2}\int_s^\tau\beta_u du}d\vec{W}_\tau. \quad (15)$$

The integrand on the right-hand side is deterministic and belongs to $L_2[0, 1]$ (for practical noise schedule choices), so its Itô's integral is a normal random variable, a martingale (meaning it has zero mean) and satisfies Itô's isometry which allows to calculate its variance:

$$\text{Var}(X_t|X_s) = \int_s^t \beta_\tau e^{-\int_s^\tau\beta_u du} I d\tau = \left(1 - e^{-\int_s^t\beta_u du}\right) I. \quad (16)$$

Thus

$$\text{Law}(X_t|X_s) = \mathcal{N}\left(e^{-\frac{1}{2}\int_s^t\beta_u du}X_s, \left(1 - e^{-\int_s^t\beta_u du}\right)I\right) = \mathcal{N}(\gamma_{s,t}X_s, (1 - \gamma_{s,t}^2)I) \quad (17)$$

B THE OPTIMAL COEFFICIENTS ASYMPTOTICS

First derive asymptotics for γ :

$$\gamma_{t-h,t} = e^{-\frac{1}{2}\int_{t-h}^t\beta_u du} = 1 - \frac{1}{2}\beta_t h + \bar{o}(h), \quad (18)$$

$$\gamma_{0,t-h}^2 = e^{-\int_0^{t-h}\beta_u du} = e^{-\int_0^t\beta_u du}e^{\int_{t-h}^t\beta_u du} = \gamma_{0,t}^2(1 + \beta_t h) + \bar{o}(h), \quad (19)$$

$$\gamma_{0,t-h} = e^{-\frac{1}{2}\int_0^{t-h}\beta_u du} = e^{-\frac{1}{2}\int_0^t\beta_u du}e^{\frac{1}{2}\int_{t-h}^t\beta_u du} = \gamma_{0,t}\left(1 + \frac{1}{2}\beta_t h\right) + \bar{o}(h), \quad (20)$$

$$\gamma_{t-h,t}^2 = e^{-\int_{t-h}^t\beta_u du} = 1 - \beta_t h + \bar{o}(h). \quad (21)$$

Then find asymptotics for μ , ν and σ^2 :

$$\mu_{t-h,t} = \left(1 - \frac{1}{2}\beta_t h + \bar{o}(h)\right) \frac{1 - \gamma_{0,t}^2 - \gamma_{0,t}^2\beta_t h + \bar{o}(h)}{1 - \gamma_{0,t}^2} = 1 - \frac{1}{2}\beta_t h - \frac{\gamma_{0,t}^2}{1 - \gamma_{0,t}^2}\beta_t h + \bar{o}(h), \quad (22)$$

$$\nu_{t-h,t} = (\gamma_{0,t}(1 + \frac{1}{2}\beta_t h) + \bar{o}(h)) \frac{\beta_t h + \bar{o}(h)}{1 - \gamma_{0,t}^2} = \frac{\gamma_{0,t}}{1 - \gamma_{0,t}^2}\beta_t h + \bar{o}(h), \quad (23)$$

$$\sigma_{t-h,t}^2 = \frac{1}{1 - \gamma_{0,t}^2}(\beta_t h + \bar{o}(h))(1 - \gamma_{0,t}^2(1 + \beta_t h) + \bar{o}(h)) = \beta_t h + \bar{o}(h). \quad (24)$$

Finally we get asymptotics for κ^* , ω^* and σ^* :

$$\begin{aligned}\kappa_{t,h}^* &= \frac{\nu_{t-h,t}(1-\gamma_{0,t}^2)}{\gamma_{0,t}\beta_t h} - 1 = \frac{\gamma_{0,t-h}(1-\gamma_{t-h,t}^2)}{\gamma_{0,t}\beta_t h} - 1 \\ &= \frac{(\beta_t h + \bar{o}(h))((1 + \frac{1}{2}\beta_t h)\gamma_{0,t} + \bar{o}(h))}{\gamma_{0,t}\beta_t h} - 1 = \bar{o}(1),\end{aligned}\quad (25)$$

$$\begin{aligned}\beta_t h \omega_{t,h}^* &= \mu_{t-h,t} - 1 + \frac{\nu_{t-h,t}}{\gamma_{0,t}} - \frac{1}{2}\beta_t h = 1 - \frac{1}{2}\beta_t h - \frac{\gamma_{0,t}^2}{1-\gamma_{0,t}^2}\beta_t h - 1 - \frac{1}{2}\beta_t h + \frac{1}{\gamma_{0,t}} \times \\ &\times \left(\frac{\gamma_{0,t}}{1-\gamma_{0,t}^2}\beta_t h + \bar{o}(h) \right) + \bar{o}(h) = \beta_t h \left(-1 - \frac{\gamma_{0,t}^2}{1-\gamma_{0,t}^2} + \frac{1}{1-\gamma_{0,t}^2} \right) + \bar{o}(h) = \bar{o}(h),\end{aligned}\quad (26)$$

$$\begin{aligned}(\sigma_{t,h}^*)^2 &= \sigma_{t-h,t}^2 + \nu_{t-h,t}^2 \mathbb{E}_{X_t} [\text{Tr}(\text{Var}(X_0|X_t))] / n = \beta_t h + \bar{o}(h) \\ &+ \frac{\gamma_{0,t}^2}{(1-\gamma_{0,t}^2)^2} \beta_t^2 h^2 \mathbb{E}_{X_t} [\text{Tr}(\text{Var}(X_0|X_t))] / n = \beta_t h(1 + \bar{o}(1)).\end{aligned}\quad (27)$$

C PROOF OF THE THEOREM 1

The key fact necessary to prove the Theorem 1 is established in the following

Lemma 1. *Let $p_{0|t}(\cdot|x)$ be pdf of conditional distribution $\text{Law}(X_0|X_t = x)$. Then for any $t \in [0, 1]$ and $x \in \mathbb{R}^n$*

$$s_{\theta^*}(x, t) = -\frac{1}{1-\gamma_{0,t}^2} \left(x - \gamma_{0,t} \mathbb{E}_{p_{0|t}(\cdot|x)} X_0 \right). \quad (28)$$

Proof of the Lemma 1. As mentioned in (Song et al., 2021c), an expression alternative to (8) can be derived for θ^* under mild assumptions on the data density (Hyvärinen, 2005; Vincent, 2011):

$$\theta^* = \arg \min_{\theta} \int_0^1 \lambda_t \mathbb{E}_{X_0 \sim p_0(\cdot)} \mathbb{E}_{X_t \sim p_{t|0}(\cdot|X_0)} \|s_{\theta}(X_t, t) - \nabla \log p_{t|0}(X_t|X_0)\|_2^2 dt, \quad (29)$$

where $\text{Law}(X_0)$ is data distribution with pdf $p_0(\cdot)$ and $\text{Law}(X_t|X_0 = x_0)$ has pdf $p_{t|0}(\cdot|x_0)$. By Bayes formula we can rewrite this in terms of pdfs $p_t(\cdot)$ and $p_{0|t}(\cdot|x_t)$ of distributions $\text{Law}(X_t)$ and $\text{Law}(X_0|X_t = x_t)$ correspondingly:

$$\theta^* = \arg \min_{\theta} \int_0^1 \lambda_t \mathbb{E}_{X_t \sim p_t(\cdot)} \mathbb{E}_{X_0 \sim p_{0|t}(\cdot|X_t)} \|s_{\theta}(X_t, t) - \nabla \log p_{t|0}(X_t|X_0)\|_2^2 dt. \quad (30)$$

For any n -dimensional random variable ξ with finite second moment and deterministic vector a we have

$$\begin{aligned}\mathbb{E}\|\xi - a\|_2^2 &= \mathbb{E}\|\xi - \mathbb{E}\xi + \mathbb{E}\xi - a\|_2^2 = \mathbb{E}\|\xi - \mathbb{E}\xi\|_2^2 + 2\langle \mathbb{E}[\xi - \mathbb{E}\xi], \mathbb{E}\xi - a \rangle \\ &+ \mathbb{E}\|\mathbb{E}\xi - a\|_2^2 = \mathbb{E}\|\xi - \mathbb{E}\xi\|_2^2 + \|\mathbb{E}\xi - a\|_2^2.\end{aligned}\quad (31)$$

In our case $\xi = \nabla \log p_{t|0}(X_t|X_0)$ and $a = s_{\theta}(X_t, t)$, so $\mathbb{E}\|\xi - \mathbb{E}\xi\|_2^2$ is independent of θ . Thus

$$\theta^* = \arg \min_{\theta} \int_0^1 \lambda_t \mathbb{E}_{X_t \sim p_t(\cdot)} \|s_{\theta}(X_t, t) - \mathbb{E}_{X_0 \sim p_{0|t}(\cdot|X_t)} [\nabla \log p_{t|0}(X_t|X_0)]\|_2^2 dt. \quad (32)$$

Therefore, the optimal score estimation network s_{θ^*} can be expressed as

$$s_{\theta^*}(x, t) = \mathbb{E}_{p_{0|t}(\cdot|x)} [\nabla \log p_{t|0}(x|X_0)] \quad (33)$$

for all $t \in [0, 1]$ and $x \in \text{supp} \{p_t\} = \mathbb{R}^n$.

As proven in Appendix A, Law $(X_t|X_0)$ is Gaussian with mean vector $\gamma_{0,t}X_0$ and covariance matrix $(1 - \gamma_{0,t}^2)\mathbf{I}$, so finally we obtain

$$s_{\theta^*}(x, t) = \mathbb{E}_{p_{0|t}(\cdot|x)} \left[-\frac{1}{1 - \gamma_{0,t}^2} (x - \gamma_{0,t}X_0) \right] = -\frac{1}{1 - \gamma_{0,t}^2} \left(x - \gamma_{0,t}\mathbb{E}_{p_{0|t}(\cdot|x)}X_0 \right). \quad (34)$$

□

Now let us prove the Theorem 1.

Proof of the Theorem 1. The sampling scheme (11) consists in adding Gaussian noise to a linear combination of \hat{X}_t and $s_{\theta^*}(\hat{X}_t, t)$. Combining (11) and the Lemma 1 we get

$$\begin{aligned} \hat{X}_{t-h} &= \hat{\sigma}_{t,h}\xi_t + \hat{X}_t + \beta_t h \left(\left(\frac{1}{2} + \hat{\omega}_{t,h} \right) \hat{X}_t + (1 + \hat{\kappa}_{t,h})s_{\theta^*}(\hat{X}_t, t) \right) = \hat{\sigma}_{t,h}\xi_t \\ &+ \left(1 + \beta_t h \left(\frac{1}{2} + \hat{\omega}_{t,h} \right) \right) \hat{X}_t + \beta_t h (1 + \hat{\kappa}_{t,h}) \left(-\frac{1}{1 - \gamma_{0,t}^2} \left(\hat{X}_t - \gamma_{0,t}\mathbb{E}_{p_{0|t}(\cdot|\hat{X}_t)}X_0 \right) \right) \\ &= \hat{\sigma}_{t,h}\xi_t + \left(1 + \beta_t h \left(\frac{1}{2} + \hat{\omega}_{t,h} - \frac{1 + \hat{\kappa}_{t,h}}{1 - \gamma_{0,t}^2} \right) \right) \hat{X}_t + \frac{\gamma_{0,t}\beta_t h (1 + \hat{\kappa}_{t,h})}{1 - \gamma_{0,t}^2} \mathbb{E}_{p_{0|t}(\cdot|\hat{X}_t)}X_0, \end{aligned} \quad (35)$$

where ξ_t are i.i.d. random variables from standard normal distribution $\mathcal{N}(0, \mathbf{I})$ for $t = 1, 1-h, \dots, h$. Thus, the distribution $\hat{X}_{t-h}|\hat{X}_t$ is also Gaussian:

$$\text{Law}(\hat{X}_{t-h}|\hat{X}_t) = \mathcal{N}\left(\hat{\mu}_{t,h}(\hat{\kappa}_{t,h}, \hat{\omega}_{t,h})\hat{X}_t + \hat{\nu}_{t,h}(\hat{\kappa}_{t,h})\mathbb{E}_{p_{0|t}(\cdot|\hat{X}_t)}X_0, \hat{\sigma}_{t,h}^2\mathbf{I}\right), \quad (36)$$

$$\hat{\mu}_{t,h}(\hat{\kappa}_{t,h}, \hat{\omega}_{t,h}) = 1 + \beta_t h \left(\frac{1}{2} + \hat{\omega}_{t,h} - \frac{1 + \hat{\kappa}_{t,h}}{1 - \gamma_{0,t}^2} \right), \quad (37)$$

$$\hat{\nu}_{t,h}(\hat{\kappa}_{t,h}) = \frac{\gamma_{0,t}\beta_t h (1 + \hat{\kappa}_{t,h})}{1 - \gamma_{0,t}^2}, \quad (38)$$

which leads to the following formula for the transition densities of the reverse diffusion:

$$\hat{p}_{t-h|t}(x_{t-h}|x_t) = \frac{1}{\sqrt{2\pi}\hat{\sigma}_{t,h}^n} \exp\left(-\frac{\|x_{t-h} - \hat{\mu}_{t,h}x_t - \hat{\nu}_{t,h}\mathbb{E}_{p_{0|t}(\cdot|x_t)}X_0\|_2^2}{2\hat{\sigma}_{t,h}^2}\right). \quad (39)$$

Moreover, comparing $\hat{\mu}_{t,h}$ and $\hat{\nu}_{t,h}$ with $\mu_{t-h,t}$ and $\nu_{t-h,t}$ defined in (9) we deduce that

$$\hat{\nu}_{t,h} = \nu_{t-h,t} \Leftrightarrow \frac{\gamma_{0,t}\beta_t h (1 + \hat{\kappa}_{t,h})}{1 - \gamma_{0,t}^2} = \nu_{t-h,t} \Leftrightarrow \hat{\kappa}_{t,h} = \kappa_{t,h}^*. \quad (40)$$

If we also want $\hat{\mu}_{t,h} = \mu_{t-h,t}$ to be satisfied, then we should have

$$1 + \beta_t h \left(\frac{1}{2} + \hat{\omega}_{t,h} - \frac{1 + \kappa_{t,h}^*}{1 - \gamma_{0,t}^2} \right) = \mu_{t-h,t} \Leftrightarrow \left(\frac{\mu_{t-h,t} - 1}{\beta_t h} - \omega_{t,h}^* + \hat{\omega}_{t,h} \right) \beta_t h + 1 = \mu_{t-h,t}, \quad (41)$$

i.e. $\hat{\nu}_{t,h} = \nu_{t-h,t}$ and $\hat{\mu}_{t,h} = \mu_{t-h,t}$ iff $\hat{\kappa}_{t,h} = \kappa_{t,h}^*$ and $\hat{\omega}_{t,h} = \omega_{t,h}^*$ for the parameters $\kappa_{t,h}^*$ and $\omega_{t,h}^*$ defined in (10).

As for the corresponding densities of the forward process X , they are Gaussian when conditioned on the initial data point X_0 :

$$\text{Law}(X_{t-h}|X_t, X_0) = \mathcal{N}(\mu_{t-h,t}X_t + \nu_{t-h,t}X_0, \sigma_{t-h,t}^2 \mathbf{I}), \quad (42)$$

where coefficients $\mu_{t-h,t}$, $\nu_{t-h,t}$ and $\sigma_{t-h,t}$ are defined in (9). This formula for $\text{Law}(X_{t-h}|X_t, X_0)$ follows from the general fact about Gaussian distributions appearing in many recent works on diffusion probabilistic modeling (Kingma et al., 2021): if $Z_t|Z_s \sim \mathcal{N}(\alpha_{t|s}Z_s, \sigma_{t|s}^2 \mathbf{I})$ and $Z_t|Z_0 \sim \mathcal{N}(\alpha_{t|0}Z_0, \sigma_{t|0}^2 \mathbf{I})$ for $0 < s < t$, then

$$\text{Law}(Z_s|Z_t, Z_0) = \mathcal{N}\left(\frac{\sigma_{s|0}^2}{\sigma_{t|0}^2}\alpha_{t|s}Z_t + \frac{\sigma_{t|s}^2}{\sigma_{t|0}^2}\alpha_{s|0}Z_0, \frac{\sigma_{s|0}^2\sigma_{t|s}^2}{\sigma_{t|0}^2} \mathbf{I}\right). \quad (43)$$

This fact is a result of applying Bayes formula to normal distributions. In our case $\alpha_{t|s} = \gamma_{s,t}$ and $\sigma_{t|s}^2 = 1 - \gamma_{s,t}^2$.

To get an expression for the densities $p_{t-h|t}(x_{t-h}|x_t)$ similar to (39), we need to integrate out the dependency on data X_0 from Gaussian distribution $\text{Law}(X_{t-h}|X_t, X_0)$:

$$\begin{aligned} p_{t-h|t}(x_{t-h}|x_t) &= \int p_{t-h,0|t}(x_{t-h}, x_0|x_t) dx_0 = \int p_{t-h|t,0}(x_{t-h}|x_t, x_0) p_{0|t}(x_0|x_t) dx_0 \\ &= \mathbb{E}_{X_0 \sim p_{0|t}(\cdot|x_t)} [p_{t-h|t,0}(x_{t-h}|x_t, X_0)], \end{aligned} \quad (44)$$

which implies the following formula:

$$p_{t-h|t}(x_{t-h}|x_t) = \frac{1}{\sqrt{2\pi}\sigma_{t-h,t}^n} \mathbb{E}_{p_{0|t}(\cdot|x_t)} \left[\exp\left(-\frac{\|x_{t-h} - \mu_{t-h,t}x_t - \nu_{t-h,t}X_0\|_2^2}{2\sigma_{t-h,t}^2}\right) \right]. \quad (45)$$

Note that in contrast with the transition densities (39) of the reverse process \hat{X} , the corresponding densities (45) of the forward process X are not normal in general.

Our goal is to find parameters $\hat{\kappa}$, $\hat{\omega}$ and $\hat{\sigma}$ that maximize log-likelihood of sample paths X under probability measure with transition densities \hat{p} . Put $t_k = kh$ for $k = 0, 1, \dots, N$ and write down this log-likelihood:

$$\begin{aligned} &\int p(x_1, x_{1-h}, \dots, x_0) \left(\sum_{k=0}^{N-1} \log \hat{p}_{t_k|t_{k+1}}(x_{t_k}|x_{t_{k+1}}) + \log \hat{p}_1(x_1) \right) dx_1 dx_{1-h} \dots dx_0 \\ &= \sum_{k=0}^{N-1} \int p(x_{t_k}, x_{t_{k+1}}) \log \hat{p}_{t_k|t_{k+1}}(x_{t_k}|x_{t_{k+1}}) dx_{t_{k+1}} dx_{t_k} + \int p(x_1) \log \hat{p}_1(x_1) dx_1. \end{aligned} \quad (46)$$

The last term does not depend on $\hat{\kappa}$, $\hat{\omega}$ and $\hat{\sigma}$, so we can ignore it. Let R_k be the k -th term in the sum above. Since we are free to have different coefficients $\hat{\kappa}_{t,h}$, $\hat{\omega}_{t,h}$ and $\hat{\sigma}_{t,h}$ for different steps, we can maximize each R_k separately. Terms R_k can be expressed as

$$\begin{aligned}
R_k &= \int p(x_{t_k}, x_{t_{k+1}}) \log \hat{p}_{t_k|t_{k+1}}(x_{t_k}|x_{t_{k+1}}) dx_{t_{k+1}} dx_{t_k} \\
&= \int p(x_{t_{k+1}}) p_{t_k|t_{k+1}}(x_{t_k}|x_{t_{k+1}}) \log \hat{p}_{t_k|t_{k+1}}(x_{t_k}|x_{t_{k+1}}) dx_{t_{k+1}} dx_{t_k} \\
&= \mathbb{E}_{X_{t_{k+1}}} \left[\int p_{t_k|t_{k+1}}(x_{t_k}|X_{t_{k+1}}) \log \hat{p}_{t_k|t_{k+1}}(x_{t_k}|X_{t_{k+1}}) dx_{t_k} \right].
\end{aligned} \tag{47}$$

From now on we will skip subscripts of $\mu, \nu, \sigma, \hat{\mu}, \hat{\nu}, \hat{\sigma}, \hat{\kappa}, \hat{\omega}, \kappa^*$ and ω^* for brevity. Denote

$$Q(x_{t_k}, X_{t_{k+1}}, X_0) = \frac{1}{\sqrt{2\pi}\sigma^n} \exp\left(-\frac{\|x_{t_k} - \mu X_{t_{k+1}} - \nu X_0\|_2^2}{2\sigma^2}\right) \log \hat{p}_{t_k|t_{k+1}}(x_{t_k}|X_{t_{k+1}}). \tag{48}$$

Using the formula (44) for the densities of X together with the explicit expression for the Gaussian density $p_{t_k|t_{k+1},0}(x_{t_k}|X_{t_{k+1}}, X_0)$ and applying Fubini's theorem to change the order of integration, we rewrite R_k as

$$\begin{aligned}
R_k &= \mathbb{E}_{X_{t_{k+1}}} \left[\int p_{t_k|t_{k+1}}(x_{t_k}|X_{t_{k+1}}) \log \hat{p}_{t_k|t_{k+1}}(x_{t_k}|X_{t_{k+1}}) dx_{t_k} \right] \\
&= \mathbb{E}_{X_{t_{k+1}}} \left[\int \mathbb{E}_{X_0 \sim p_{0|t_{k+1}}(\cdot|X_{t_{k+1}})} [p_{t_k|t_{k+1},0}(x_{t_k}|X_{t_{k+1}}, X_0) \log \hat{p}_{t_k|t_{k+1}}(x_{t_k}|X_{t_{k+1}})] dx_{t_k} \right] \\
&= \mathbb{E}_{X_{t_{k+1}}} \left[\int \mathbb{E}_{X_0 \sim p_{0|t_{k+1}}(\cdot|X_{t_{k+1}})} [Q(x_{t_k}, X_{t_{k+1}}, X_0)] dx_{t_k} \right] \\
&= \mathbb{E}_{X_{t_{k+1}}} \mathbb{E}_{X_0 \sim p_{0|t_{k+1}}(\cdot|X_{t_{k+1}})} \left[\int Q(x_{t_k}, X_{t_{k+1}}, X_0) dx_{t_k} \right].
\end{aligned} \tag{49}$$

The formula (48) implies that the integral of $Q(x_{t_k}, X_{t_{k+1}}, X_0)$ with respect to x_{t_k} can be seen as expectation of $\log \hat{p}_{t_k|t_{k+1}}(\xi|X_{t_{k+1}})$ with respect to normal random variable ξ with mean $\mu X_{t_{k+1}} + \nu X_0$ and covariance matrix $\sigma^2 \mathbf{I}$. Plugging in the expression (39) into (48), we can calculate this integral:

$$\begin{aligned}
&\mathbb{E}_\xi \left[-\log \sqrt{2\pi} - n \log \hat{\sigma} - \frac{\|\xi - \hat{\mu} X_{t_{k+1}} - \hat{\nu} \mathbb{E}_{X_0 \sim p_{0|t_{k+1}}(\cdot|X_{t_{k+1}})} X_0'\|_2^2}{2\hat{\sigma}^2} \right] \\
&= -\log \sqrt{2\pi} - n \log \hat{\sigma} - \frac{\mathbb{E}_\xi \|\xi - \hat{\mu} X_{t_{k+1}} - \hat{\nu} \mathbb{E}_{p_{0|t_{k+1}}(\cdot|X_{t_{k+1}})} X_0'\|_2^2}{2\hat{\sigma}^2}.
\end{aligned} \tag{50}$$

Thus, terms R_k we want to maximize equal

$$R_k = -\log \sqrt{2\pi} - n \log \hat{\sigma} - \mathbb{E}_{X_{t_{k+1}}} \mathbb{E}_{X_0 \sim p_{0|t_{k+1}}(\cdot|X_{t_{k+1}})} \frac{\mathbb{E}_\xi \|\xi - \hat{\mu} X_{t_{k+1}} - \hat{\nu} \mathbb{E}_{p_{0|t_{k+1}}(\cdot|X_{t_{k+1}})} X_0'\|_2^2}{2\hat{\sigma}^2} \tag{51}$$

Maximizing R_k with respect to $(\hat{\kappa}, \hat{\omega}, \hat{\sigma})$ is equivalent to minimizing $\mathbb{E}_{X_{t_{k+1}}} S_k$ where S_k is given by

$$S_k = n \log \hat{\sigma} + \frac{1}{2\hat{\sigma}^2} \mathbb{E}_{X_0 \sim p_{0|t_{k+1}}(\cdot|X_{t_{k+1}})} \mathbb{E}_\xi \|\xi - \hat{\mu} X_{t_{k+1}} - \hat{\nu} \mathbb{E}_{p_{0|t_{k+1}}(\cdot|X_{t_{k+1}})} X_0'\|_2^2, \tag{52}$$

where the expectation with respect to $\xi \sim \mathcal{N}(\mu X_{t_{k+1}} + \nu X_0, \sigma^2 \mathbf{I})$ can be calculated using the fact that for every vector \hat{a} we can express $\mathbb{E}_\xi \|\xi - \hat{a}\|_2^2$ as

$$\mathbb{E}\|\xi - \mathbb{E}\xi + \mathbb{E}\xi - \hat{a}\|_2^2 = \mathbb{E}\|\xi - \mathbb{E}\xi\|_2^2 + 2\langle \mathbb{E}[\xi - \mathbb{E}\xi], \mathbb{E}\xi - \hat{a} \rangle + \mathbb{E}\|\mathbb{E}\xi - \hat{a}\|_2^2 = n\sigma^2 + \|\mathbb{E}\xi - \hat{a}\|_2^2. \quad (53)$$

So, the outer expectation with respect to $\text{Law}(X_0|X_{t_{k+1}})$ in (52) can be simplified:

$$\begin{aligned} & \mathbb{E}_{X_0 \sim p_0|\cdot(\cdot|X_{t_{k+1}})} \left[n\sigma^2 + \|(\mu - \hat{\mu})X_{t_{k+1}} + \nu X_0 - \hat{\nu} \mathbb{E}_{X'_0 \sim p_0|\cdot(\cdot|X_{t_{k+1}})} X'_0\|_2^2 \right] \\ &= n\sigma^2 + \mathbb{E}_{X_0} \|((\mu - \hat{\mu})X_{t_{k+1}} + \nu X_0 - \hat{\nu} \mathbb{E}_{X'_0} X'_0)\|_2^2 = n\sigma^2 + (\mu - \hat{\mu})^2 \|X_{t_{k+1}}\|_2^2 \\ &+ 2\langle (\mu - \hat{\mu})X_{t_{k+1}}, (\nu - \hat{\nu}) \mathbb{E}_{X_0} X_0 \rangle + \mathbb{E}_{X_0} \|\nu X_0 - \hat{\nu} \mathbb{E}_{X'_0} X'_0\|_2^2 = (\mu - \hat{\mu})^2 \|X_{t_{k+1}}\|_2^2 \\ &+ 2\langle (\mu - \hat{\mu})X_{t_{k+1}}, (\nu - \hat{\nu}) \mathbb{E}_{X_0} X_0 \rangle + \nu^2 \mathbb{E}_{X_0} \|X_0\|_2^2 + \hat{\nu}^2 \|\mathbb{E}_{X_0} X_0\|_2^2 + n\sigma^2 \\ &- 2\nu\hat{\nu} \langle \mathbb{E}_{X_0} X_0, \mathbb{E}_{X'_0} X'_0 \rangle = \|(\mu - \hat{\mu})X_{t_{k+1}} + (\nu - \hat{\nu}) \mathbb{E}_{X_0} X_0\|_2^2 + \nu^2 \mathbb{E}_{X_0} \|X_0\|_2^2 \\ &- \nu^2 \|\mathbb{E}_{X_0} X_0\|_2^2 + n\sigma^2, \end{aligned} \quad (54)$$

where all the expectations in the formula above are taken with respect to the conditional data distribution $\text{Law}(X_0|X_{t_{k+1}})$. So, the resulting expression for the terms S_k whose expectation with respect to $\text{Law}(X_{t_{k+1}})$ we want to minimize is

$$\begin{aligned} S_k &= n \log \hat{\sigma} + \frac{1}{2\hat{\sigma}^2} \left(n\sigma^2 + \|(\mu - \hat{\mu})X_{t_{k+1}} + (\nu - \hat{\nu}) \mathbb{E}[X_0|X_{t_{k+1}}]\|_2^2 \right. \\ &\quad \left. + \nu^2 (\mathbb{E}[\|X_0\|_2^2|X_{t_{k+1}}] - \|\mathbb{E}[X_0|X_{t_{k+1}}]\|_2^2) \right). \end{aligned} \quad (55)$$

Now it is clear that $\kappa_{t_{k+1},h}^*$ and $\omega_{t_{k+1},h}^*$ are optimal because $\hat{\mu}_{t_{k+1},h}(\kappa_{t_{k+1},h}^*, \omega_{t_{k+1},h}^*) = \mu_{t_k, t_{k+1}}$ and $\hat{\nu}_{t_{k+1},h}(\kappa_{t_{k+1},h}^*) = \nu_{t_k, t_{k+1}}$. For this choice of parameters we have

$$\mathbb{E}_{X_{t_{k+1}}} S_k = n \log \hat{\sigma} + \frac{1}{2\hat{\sigma}^2} \left(n\sigma^2 + \nu^2 \mathbb{E}_{X_{t_{k+1}}} [\mathbb{E}[\|X_0\|_2^2|X_{t_{k+1}}] - \|\mathbb{E}[X_0|X_{t_{k+1}}]\|_2^2] \right). \quad (56)$$

Note that $\mathbb{E}[\|X_0\|_2^2|X_{t_{k+1}}] - \|\mathbb{E}[X_0|X_{t_{k+1}}]\|_2^2 = \text{Tr}(\text{Var}(X_0|X_{t_{k+1}}))$ is the overall variance of $\text{Law}(X_0|X_{t_{k+1}})$ along all n dimensions. Differentiating $\mathbb{E}_{X_{t_{k+1}}} S_k$ with respect to $\hat{\sigma}$ shows that the optimal $\sigma_{t_{k+1},h}^*$ should satisfy

$$\frac{n}{\sigma_{t_{k+1},h}^*} - \frac{1}{(\sigma_{t_{k+1},h}^*)^3} \left(n\sigma_{t_k, t_{k+1}}^2 + \nu_{t_k, t_{k+1}}^2 \mathbb{E}_{X_{t_{k+1}}} [\text{Tr}(\text{Var}(X_0|X_{t_{k+1}}))] \right) = 0, \quad (57)$$

which is indeed satisfied by the parameters $\sigma_{t,h}^*$ defined in (10). Thus, the statement (i) is proven.

When it comes to proving that \hat{X} is exact, we have to show that $\text{Law}(\hat{X}_{t_k}) = \text{Law}(X_{t_k})$ for every $k = 0, 1, \dots, N$. By the assumption that $\text{Law}(\hat{X}_1) = \text{Law}(X_1)$ it is sufficient to prove that $\hat{p}_{t_k|t_{k+1}}(x_{t_k}|x_{t_{k+1}}) \equiv p_{t_k|t_{k+1}}(x_{t_k}|x_{t_{k+1}})$ since the exactness will follow from this fact by mathematical induction. If X_0 is a constant random variable, $\text{Law}(X_0|X_t) = \text{Law}(X_0)$ also corresponds to the same constant, so $\text{Var}(X_0|X_t) = 0$ meaning that $\sigma_{t,h}^* = \sigma_{t-h,t}$, and the formulae (39) and (45) imply the desired result.

Let us now consider the second case when $X_0 \sim \mathcal{N}(\bar{\mu}, \delta^2 \mathbf{I})$. It is a matter of simple but lengthy computations to prove another property of Gaussian distributions similar to (43): if $Z_0 \sim \mathcal{N}(\bar{\mu}, \delta^2 \mathbf{I})$ and $Z_t|Z_0 \sim \mathcal{N}(a_t Z_0, b_t^2 \mathbf{I})$, then $Z_0|Z_t \sim \mathcal{N}(\frac{b_t^2}{b_t^2 + \delta^2 a_t^2} \bar{\mu} + \frac{\delta^2 a_t}{b_t^2 + \delta^2 a_t^2} Z_t, \frac{\delta^2 b_t^2}{b_t^2 + \delta^2 a_t^2} \mathbf{I})$ and $Z_t \sim \mathcal{N}(\bar{\mu} a_t, (b_t^2 + \delta^2 a_t^2) \mathbf{I})$. In our case $a_t = \gamma_{0,t}$ and $b_t^2 = 1 - \gamma_{0,t}^2$, therefore

$$\text{Law}(X_0|X_t) = \mathcal{N} \left(\frac{1 - \gamma_{0,t}^2}{1 - \gamma_{0,t}^2 + \delta^2 \gamma_{0,t}^2} \bar{\mu} + \frac{\delta^2 \gamma_{0,t}}{1 - \gamma_{0,t}^2 + \delta^2 \gamma_{0,t}^2} X_t, \frac{\delta^2 (1 - \gamma_{0,t}^2)}{1 - \gamma_{0,t}^2 + \delta^2 \gamma_{0,t}^2} \mathbf{I} \right). \quad (58)$$

So, $\text{Var}(X_0|X_t)$ does not depend on X_t and

$$(\sigma_{t,h}^*)^2 = \sigma_{t-h,t}^2 + \frac{\nu_{t-h,t}^2}{n} \mathbb{E}_{X_t} [\text{Tr}(\text{Var}(X_0|X_t))] = \sigma_{t-h,t}^2 + \nu_{t-h,t}^2 \frac{\delta^2(1-\gamma_{0,t}^2)}{1-\gamma_{0,t}^2 + \delta^2\gamma_{0,t}^2}. \quad (59)$$

Since $\text{Law}(X_t|X_{t-h})$, $\text{Law}(X_{t-h})$ and $\text{Law}(X_t)$ are Gaussian, Bayes formula implies that $\text{Law}(X_{t-h}|X_t)$ is Gaussian as well with the following mean and covariance matrix:

$$\mathbb{E}[X_{t-h}|X_t] = \frac{\gamma_{0,t-h}(1-\gamma_{t-h,t}^2)}{1-\gamma_{0,t}^2 + \delta^2\gamma_{0,t}^2} \bar{\mu} + \frac{\gamma_{t-h,t}(1-\gamma_{0,t-h}^2 + \delta^2\gamma_{0,t-h}^2)}{1-\gamma_{0,t}^2 + \delta^2\gamma_{0,t}^2} X_t, \quad (60)$$

$$\text{Var}(X_{t-h}|X_t) = \frac{(1-\gamma_{t-h,t}^2)(1-\gamma_{0,t-h}^2 + \delta^2\gamma_{0,t-h}^2)}{1-\gamma_{0,t}^2 + \delta^2\gamma_{0,t}^2} \mathbf{I}. \quad (61)$$

The distribution $\text{Law}(\hat{X}_{t-h}|\hat{X}_t)$ is also Gaussian by the formula (39), so to conclude the proof we just need to show that $\mathbb{E}[\hat{X}_{t-h}|\hat{X}_t = x] = \mathbb{E}[X_{t-h}|X_t = x]$ and $\text{Var}(\hat{X}_{t-h}|\hat{X}_t = x) = \text{Var}(X_{t-h}|X_t = x)$ for every $x \in \mathbb{R}^n$ for the optimal parameters (10). Recall that for $\kappa_{t,h}^*$ and $\omega_{t,h}^*$ we have $\hat{\mu}_{t,h}(\kappa_{t,h}^*, \omega_{t,h}^*) = \mu_{t-h,t}$ and $\hat{\nu}_{t,h}(\kappa_{t,h}^*) = \nu_{t-h,t}$. Utilizing the formulae (9), (39), (58) and the fact that $\gamma_{0,t-h} \cdot \gamma_{t-h,t} = \gamma_{0,t}$ (following from the definition of γ in (7)) we conclude that

$$\begin{aligned} \mathbb{E}[\hat{X}_{t-h}|\hat{X}_t = x] &= \hat{\mu}_{t,h}(\kappa_{t,h}^*, \omega_{t,h}^*)x + \hat{\nu}_{t,h}(\kappa_{t,h}^*)\mathbb{E}_{p_{0|t}(\cdot|x)}X_0 \\ &= \gamma_{t-h,t} \frac{1-\gamma_{0,t-h}^2}{1-\gamma_{0,t}^2}x + \gamma_{0,t-h} \frac{1-\gamma_{t-h,t}^2}{1-\gamma_{0,t}^2} \left[\frac{1-\gamma_{0,t}^2}{1-\gamma_{0,t}^2 + \delta^2\gamma_{0,t}^2} \bar{\mu} + \frac{\delta^2\gamma_{0,t}}{1-\gamma_{0,t}^2 + \delta^2\gamma_{0,t}^2} x \right] \\ &= \gamma_{0,t-h} \frac{1-\gamma_{t-h,t}^2}{1-\gamma_{0,t}^2 + \delta^2\gamma_{0,t}^2} \bar{\mu} + \gamma_{t-h,t} \frac{(1-\gamma_{0,t-h}^2)(1-\gamma_{0,t}^2) + \delta^2\gamma_{0,t}^2(1-\gamma_{0,t-h}^2)}{(1-\gamma_{0,t}^2)(1-\gamma_{0,t}^2 + \delta^2\gamma_{0,t}^2)} x \\ &\quad + \gamma_{t-h,t} \frac{\delta^2\gamma_{0,t-h}^2(1-\gamma_{t-h,t}^2)}{(1-\gamma_{0,t}^2)(1-\gamma_{0,t}^2 + \delta^2\gamma_{0,t}^2)} x = \gamma_{0,t-h} \frac{1-\gamma_{t-h,t}^2}{1-\gamma_{0,t}^2 + \delta^2\gamma_{0,t}^2} \bar{\mu} \\ &\quad + \gamma_{t-h,t} \frac{(1-\gamma_{0,t-h}^2)(1-\gamma_{0,t}^2) + \delta^2\gamma_{0,t-h}^2(1-\gamma_{0,t}^2)}{(1-\gamma_{0,t}^2)(1-\gamma_{0,t}^2 + \delta^2\gamma_{0,t}^2)} x = \gamma_{0,t-h} \frac{1-\gamma_{t-h,t}^2}{1-\gamma_{0,t}^2 + \delta^2\gamma_{0,t}^2} \bar{\mu} \\ &\quad + \gamma_{t-h,t} \frac{1-\gamma_{0,t-h}^2 + \delta^2\gamma_{0,t-h}^2}{1-\gamma_{0,t}^2 + \delta^2\gamma_{0,t}^2} x = \mathbb{E}[X_{t-h}|X_t = x], \end{aligned} \quad (62)$$

$$\begin{aligned} \text{Var}(\hat{X}_{t-h}|\hat{X}_t = x) &= (\sigma_{t,h}^*)^2 \mathbf{I} = \left(\sigma_{t-h,t}^2 + \nu_{t-h,t}^2 \frac{\delta^2(1-\gamma_{0,t}^2)}{1-\gamma_{0,t}^2 + \delta^2\gamma_{0,t}^2} \right) \mathbf{I} \\ &= \left(\frac{(1-\gamma_{0,t-h}^2)(1-\gamma_{t-h,t}^2)}{1-\gamma_{0,t}^2} + \gamma_{0,t-h}^2 \frac{\delta^2(1-\gamma_{t-h,t}^2)^2}{(1-\gamma_{0,t}^2)(1-\gamma_{0,t}^2 + \delta^2\gamma_{0,t}^2)} \right) \mathbf{I} \\ &= \frac{1-\gamma_{t-h,t}^2}{(1-\gamma_{0,t}^2)(1-\gamma_{0,t}^2 + \delta^2\gamma_{0,t}^2)} \left((1-\gamma_{0,t-h}^2)(1-\gamma_{0,t}^2) + \delta^2\gamma_{0,t}^2(1-\gamma_{0,t-h}^2) \right) \mathbf{I} \\ &\quad + \frac{1-\gamma_{t-h,t}^2}{(1-\gamma_{0,t}^2)(1-\gamma_{0,t}^2 + \delta^2\gamma_{0,t}^2)} (\delta^2\gamma_{0,t-h}^2(1-\gamma_{t-h,t}^2)) \mathbf{I} \\ &= \frac{1-\gamma_{t-h,t}^2}{(1-\gamma_{0,t}^2)(1-\gamma_{0,t}^2 + \delta^2\gamma_{0,t}^2)} \left((1-\gamma_{0,t-h}^2)(1-\gamma_{0,t}^2) + \delta^2\gamma_{0,t-h}^2(1-\gamma_{0,t}^2) \right) \mathbf{I} \\ &= \frac{(1-\gamma_{t-h,t}^2)(1-\gamma_{0,t-h}^2 + \delta^2\gamma_{0,t-h}^2)}{1-\gamma_{0,t}^2 + \delta^2\gamma_{0,t}^2} \mathbf{I} = \text{Var}(X_{t-h}|X_t = x). \end{aligned} \quad (63)$$

□

D REVERSE MR-VP SDE SOLVER

MR-VP DPM is characterized by the following forward and reverse diffusions:

$$dX_t = \frac{1}{2}\beta_t(\bar{X} - X_t)dt + \sqrt{\beta_t}d\bar{W}_t, \quad (64)$$

$$d\hat{X}_t = \left(\frac{1}{2}\beta_t(\bar{X} - \hat{X}_t) - \beta_t s_\theta(\hat{X}_t, \bar{X}, t) \right) dt + \sqrt{\beta_t}d\hat{W}_t. \quad (65)$$

Using the same method as in Appendix A, we can show that for $s < t$

$$\text{Law}(X_t|X_s) = \mathcal{N}(\gamma_{s,t}X_s + (1 - \gamma_{s,t})\bar{X}, (1 - \gamma_{s,t}^2)\mathbf{I}), \quad \gamma_{s,t} = e^{-\frac{1}{2}\int_s^t \beta_u du}. \quad (66)$$

With the following notation:

$$\mu_{s,t} = \gamma_{s,t} \frac{1 - \gamma_{0,s}^2}{1 - \gamma_{0,t}^2}, \quad \nu_{s,t} = \gamma_{0,s} \frac{1 - \gamma_{s,t}^2}{1 - \gamma_{0,t}^2}, \quad \sigma_{s,t}^2 = \frac{(1 - \gamma_{0,s}^2)(1 - \gamma_{s,t}^2)}{1 - \gamma_{0,t}^2} \quad (67)$$

we can write down the parameters of Gaussian distribution $X_s|X_t, X_0$:

$$\mathbb{E}[X_s|X_t, X_0] = \bar{X} + \mu_{s,t}(X_t - \bar{X}) + \nu_{s,t}(X_0 - \bar{X}), \quad \text{Var}(X_s|X_t, X_0) = \sigma_{s,t}^2 \mathbf{I}. \quad (68)$$

The Lemma 1 for MR-VP DPMs takes the following shape:

$$\begin{aligned} s_{\theta^*}(x, \bar{X}, t) &= -\frac{1}{1 - \gamma_{0,t}^2} \left(x - (1 - \gamma_{0,t})\bar{X} - \gamma_{0,t}\mathbb{E}_{p_{0|t}(\cdot|x)}X_0 \right) \\ &= -\frac{1}{1 - \gamma_{0,t}^2} \left((x - \bar{X}) - \gamma_{0,t} \left(\mathbb{E}_{p_{0|t}(\cdot|x)}X_0 - \bar{X} \right) \right). \end{aligned} \quad (69)$$

The class of reverse SDE solvers we consider is

$$\hat{X}_{t-h} = \hat{X}_t + \beta_t h \left(\left(\frac{1}{2} + \hat{\omega}_{t,h} \right) (\hat{X}_t - \bar{X}) + (1 + \hat{\kappa}_{t,h}) s_\theta(\hat{X}_t, \bar{X}, t) \right) + \hat{\sigma}_{t,h} \xi_t, \quad (70)$$

where $t = 1, 1-h, \dots, h$ and ξ_t are i.i.d. samples from $\mathcal{N}(0, \mathbf{I})$. Repeating the argument of the Theorem 1 leads to the following optimal (in terms of likelihood of the forward diffusion sample paths) parameters:

$$\begin{aligned} \kappa_{t,h}^* &= \frac{\nu_{t-h,t}(1 - \gamma_{0,t}^2)}{\gamma_{0,t}\beta_t h} - 1, \quad \omega_{t,h}^* = \frac{\mu_{t-h,t} - 1}{\beta_t h} + \frac{1 + \kappa_{t,h}^*}{1 - \gamma_{0,t}^2} - \frac{1}{2}, \\ (\sigma_{t,h}^*)^2 &= \sigma_{t-h,t}^2 + \frac{1}{n} \nu_{t-h,t}^2 \mathbb{E}_{X_t} [\text{Tr}(\text{Var}(X_0|X_t))], \end{aligned} \quad (71)$$

which are actually the same as the optimal parameters (10) for VP DPM. It is of no surprise since MR-VP DPM and VP-DPM differ only by a constant shift.

E REVERSE SUB-VP SDE SOLVER

Sub-VP DPM is characterized by the following forward and reverse diffusions:

$$dX_t = -\frac{1}{2}\beta_t X_t dt + \sqrt{\beta_t(1 - e^{-2\int_0^t \beta_u du})}d\bar{W}_t, \quad (72)$$

$$d\hat{X}_t = \left(-\frac{1}{2}\beta_t\hat{X}_t - \beta_t \left(1 - e^{-2\int_0^t \beta_u du} \right) s_\theta(\hat{X}_t, t) \right) dt + \sqrt{\beta_t(1 - e^{-2\int_0^t \beta_u du})} d\overleftarrow{W}_t. \quad (73)$$

Using the same method as in Appendix A, we can show that for $s < t$

$$\text{Law}(X_t|X_s) = \mathcal{N}(\gamma_{s,t}X_s, (1 + \gamma_{0,t}^4 - \gamma_{s,t}^2(1 + \gamma_{0,s}^4))\mathbf{I}), \quad \gamma_{s,t} = e^{-\frac{1}{2}\int_s^t \beta_u du}. \quad (74)$$

Note that for $s = 0$ this expression simplifies to

$$\text{Law}(X_t|X_0) = \mathcal{N}(\gamma_{0,t}X_0, (1 - \gamma_{0,t}^2)^2\mathbf{I}). \quad (75)$$

With the following notation:

$$\begin{aligned} \mu_{s,t} &= \gamma_{s,t} \left(\frac{1 - \gamma_{0,s}^2}{1 - \gamma_{0,t}^2} \right)^2, & \nu_{s,t} &= \gamma_{0,s} \frac{1 + \gamma_{0,t}^4 - \gamma_{s,t}^2(1 + \gamma_{0,s}^4)}{(1 - \gamma_{0,t}^2)^2}, \\ \sigma_{s,t}^2 &= \frac{(1 - \gamma_{0,s}^2)^2(1 + \gamma_{0,t}^4 - \gamma_{s,t}^2(1 + \gamma_{0,s}^4))}{(1 - \gamma_{0,t}^2)^2} \end{aligned} \quad (76)$$

we can write down the parameters of Gaussian distribution $X_s|X_t, X_0$:

$$\mathbb{E}[X_s|X_t, X_0] = \mu_{s,t}X_t + \nu_{s,t}X_0, \quad \text{Var}(X_s|X_t, X_0) = \sigma_{s,t}^2\mathbf{I}. \quad (77)$$

The Lemma 1 for sub-VP DPMs takes the following shape:

$$s_{\theta^*}(x, t) = -\frac{1}{(1 - \gamma_{0,t}^2)^2} \left(x - \gamma_{0,t}\mathbb{E}_{p_{0,t}(\cdot|x)}X_0 \right). \quad (78)$$

The class of reverse SDE solvers we consider is

$$\hat{X}_{t-h} = \hat{X}_t + \beta_t h \left(\left(\frac{1}{2} + \hat{\omega}_{t,h} \right) \hat{X}_t + (1 + \hat{\kappa}_{t,h}) \left(1 - e^{-2\int_0^t \beta_u du} \right) s_\theta(\hat{X}_t, t) \right) + \hat{\sigma}_{t,h}\xi_t, \quad (79)$$

where $t = 1, 1-h, \dots, h$ and ξ_t are i.i.d. samples from $\mathcal{N}(0, \mathbf{I})$. Repeating the argument of the Theorem 1 leads to the following optimal (in terms of likelihood of the forward diffusion sample paths) parameters:

$$\begin{aligned} \kappa_{t,h}^* &= \frac{\nu_{t-h,t}(1 - \gamma_{0,t}^2)}{\gamma_{0,t}\beta_t h(1 + \gamma_{0,t}^2)} - 1, & \omega_{t,h}^* &= \frac{\mu_{t-h,t} - 1}{\beta_t h} + \frac{(1 + \kappa_{t,h}^*)(1 + \gamma_{0,t}^2)}{1 - \gamma_{0,t}^2} - \frac{1}{2}, \\ (\sigma_{t,h}^*)^2 &= \sigma_{t-h,t}^2 + \frac{1}{n}\nu_{t-h,t}^2\mathbb{E}_{X_t}[\text{Tr}(\text{Var}(X_0|X_t))]. \end{aligned} \quad (80)$$

F REVERSE VE SDE SOLVER

VE DPM is characterized by the following forward and reverse diffusions:

$$dX_t = \sqrt{(\sigma_t^2)'} d\overrightarrow{W}_t, \quad (81)$$

$$d\hat{X}_t = -(\sigma_t^2)' s_\theta(\hat{X}_t, t) dt + \sqrt{(\sigma_t^2)'} d\overleftarrow{W}_t. \quad (82)$$

Since for $s < t$

$$X_t = X_s + \int_s^t \sqrt{(\sigma_u^2)'} d\vec{W}_u, \quad (83)$$

similar argument as in Appendix A allows showing that

$$\text{Law}(X_t|X_s) = \mathcal{N}\left(X_s, \mathbf{I} \cdot \int_s^t (\sigma_u^2)' du\right) = \mathcal{N}(X_s, (\sigma_t^2 - \sigma_s^2) \mathbf{I}). \quad (84)$$

With the following notation:

$$\mu_{s,t} = \frac{\sigma_s^2 - \sigma_0^2}{\sigma_t^2 - \sigma_0^2}, \quad \nu_{s,t} = \frac{\sigma_t^2 - \sigma_s^2}{\sigma_t^2 - \sigma_0^2}, \quad \sigma_{s,t}^2 = \frac{(\sigma_t^2 - \sigma_s^2)(\sigma_s^2 - \sigma_0^2)}{\sigma_t^2 - \sigma_0^2}, \quad (85)$$

we can write down the parameters of Gaussian distribution $X_s|X_t, X_0$:

$$\mathbb{E}[X_s|X_t, X_0] = \mu_{s,t}X_t + \nu_{s,t}X_0, \quad \text{Var}(X_s|X_t, X_0) = \sigma_{s,t}^2 \mathbf{I}. \quad (86)$$

The Lemma 1 for VE DPMs takes the following shape:

$$s_{\theta^*}(x, t) = -\frac{1}{\sigma_t^2 - \sigma_0^2} \left(x - \mathbb{E}_{p_{0|t}(\cdot|x)} X_0 \right). \quad (87)$$

Repeating the argument of the Theorem 1 leads to the following optimal (in terms of likelihood of the forward diffusion sample paths) reverse SDE solver:

$$\hat{X}_{t-h} = \hat{X}_t + (\sigma_t^2 - \sigma_{t-h}^2) s_{\theta}(\hat{X}_t, t) + \sigma_{t,h}^* \xi_t, \quad (88)$$

where

$$(\sigma_{t,h}^*)^2 = \frac{(\sigma_t^2 - \sigma_{t-h}^2)(\sigma_{t-h}^2 - \sigma_0^2)}{\sigma_t^2 - \sigma_0^2} + \frac{1}{n} \nu_{t-h,t}^2 \mathbb{E}_{X_t} [\text{Tr}(\text{Var}(X_0|X_t))], \quad (89)$$

$t = 1, 1-h, \dots, h$ and ξ_t are i.i.d. samples from $\mathcal{N}(0, \mathbf{I})$.

G TOY EXAMPLES

In this section we consider toy examples where data distribution X_0 is represented by a single point (corresponding to the case (ii) of the Theorem 1) and by two points (corresponding to more general case (i)). In the first case the point is unit vector $i = (1, 1, \dots, 1)$ of dimensionality 100, in the second one two points i and $-2i$ have the same probability. We compare performance of two solvers, Euler-Maruyama and the proposed Maximum Likelihood, depending on the number $N \in \{1, 2, 5, 10, 100, 1000\}$ of solver steps. The output of the perfectly trained score matching network s_{θ^*} is computed analytically and Gaussian noise with variance $\varepsilon \in \{0.0, 0.1, 0.5\}$ is added to approximate the realistic case when the network s_{θ} we use is not trained till optimality. We considered VP diffusion model (6) with $\beta_0 = 0.05$ and $\beta_1 = 20.0$.

The results of the comparison are given in Table 5 and can be summarized in the following:

- for both methods, larger N means better quality;
- for both methods, more accurate score matching networks (smaller ε) means better quality;
- for large number of steps, both methods perform the same;
- it takes less number of steps for the proposed Maximum Likelihood solver to converge with a good accuracy to data distribution than it does for Euler-Maruyama solver;

Table 5: Maximum Likelihood (ML) and Euler-Maruyama (EM) solvers comparison in terms of Mean Square Error (MSE). MSE < 0.001 is denoted by *conv*, MSE > 1.0 is denoted by *div*. N is the number of SDE solver steps, ε is variance of Gaussian noise added to perfect scores s_{θ^*} .

MSE ML / EM	$X_0 = \{i\}$			$X_0 = \{i, -2i\}$		
	$\varepsilon = 0.0$	$\varepsilon = 0.1$	$\varepsilon = 0.5$	$\varepsilon = 0.0$	$\varepsilon = 0.1$	$\varepsilon = 0.5$
$N = 1$	<i>conv / div</i>	<i>div / div</i>	<i>div / div</i>	<i>div / div</i>	<i>div / div</i>	<i>div / div</i>
$N = 2$	<i>conv / div</i>	<i>div / div</i>	<i>div / div</i>	0.15 / <i>div</i>	<i>div / div</i>	<i>div / div</i>
$N = 5$	<i>conv / div</i>	0.017 / <i>div</i>	0.085 / <i>div</i>	<i>conv / div</i>	0.017 / <i>div</i>	0.085 / <i>div</i>
$N = 10$	<i>conv / 0.57</i>	0.001/0.59	0.005/0.67	<i>conv / 0.57</i>	0.001/0.59	0.006/0.67
$N = 100$	<i>conv / 0.01</i>	<i>conv / 0.01</i>	<i>conv / 0.01</i>	<i>conv / 0.01</i>	<i>conv / 0.01</i>	<i>conv / 0.01</i>
$N = 1000$	<i>conv / conv</i>	<i>conv / conv</i>	<i>conv / conv</i>	<i>conv / conv</i>	<i>conv / conv</i>	<i>conv / conv</i>

- in accordance with the statement (ii) of the Theorem 1, the optimal Maximum Likelihood solver leads to exact data reconstruction in the case when data distribution is constant and score matching network is trained till optimality (i.e. $\varepsilon = 0.0$) irrespective of the number of steps N .

Also, in the second example where $X_0 \in \{i, -2i\}$ the Maximum Likelihood SDE solver reconstructs the probabilities of these two points better than Euler-Maruyama which tends to output “ i -samples” (which are closer to the origin) more frequently than “ $-2i$ -samples”. E.g. for $\varepsilon = 0.0$ and $N = 10$ the frequency of “ i -samples” generated by Euler-Maruyama scheme is 54% while this frequency for Maximum Likelihood scheme is 50% (500k independent runs were used to calculate these frequencies).

H SPEAKER CONDITIONING NETWORK

The function $x \cdot \tanh(\text{softplus}(x))$ is used as a non-linearity in the speaker conditioning network $g_t(Y)$. First, time embedding t_e is obtained by the following procedure: time $t \in [0, 1]$ is encoded with positional encoding (Song et al., 2021c), then resulting 256-dimensional vector t' is passed through the first linear module with 1024 units, then a non-linearity is applied to it and then it is passed through the second linear module with 256 units. Next, noisy mel-spectrogram Y_t for *wodyn* input type or Y_t concatenated with $\{Y_s | s = 0.5/15, 1.5/15, \dots, 14.5/15\}$ for *whole* is passed through 6 blocks consisting of 2D convolutional layers each followed by instance normalization and Gated Linear Unit. The number of input and output channels of these convolutions is (1, 64), (32, 64), (32, 128), (64, 128), (64, 256), (128, 256) for *wodyn* input type and the same but with 16 input channels in the first convolution for *whole* input type. After the 2nd and 4th blocks $MLP_1(t_e)$ and $MLP_2(t_e)$ are broadcast-added where MLP_1 (MLP_2) are composed of a non-linearity followed by a linear module with 32 (64) units. After the last 6th block the result is passed through the final convolution with 128 output channels and average pooling along both time and frequency axes is applied resulting in 128-dimensional vector. All convolutions except for the final one have (kernel, stride, zero padding) = (3, 1, 1) while for the final one the corresponding parameters are (1, 0, 0). Denote the result of such processing of Y by c for *wodyn* and *whole* input types.

Clean target mel-spectrogram Y_0 is used to obtain 256-dimensional speaker embedding d with the pre-trained speaker verification network (Jia et al., 2018) which is *not* trained. Vectors d , c and t' are concatenated (except for *d-only* input type where we concatenate only d and t'), passed through a linear module with 512 units followed by a non-linearity and a linear module with 128 units. The resulting 128-dimensional vector is the output of the speaker conditioning network $g_t(Y)$.

I TRAINING HYPERPARAMETERS AND OTHER DETAILS

Encoders and decoders were trained with batch sizes 128 and 32 and Adam optimizer with initial learning rates 0.0005 and 0.0001 correspondingly. Encoders and decoders in VCTK models were trained for 500 and 200 epochs respectively; as for LibriTTS models, they were trained for 300 and

110 epochs. The datasets were downsampled to 22.05kHz which was the operating rate of our VC models. VCTK recordings were preprocessed by removing silence in the beginning and in the end of utterances. To fit GPU memory, decoders were trained on random speech segments of approximately 1.5 seconds rather than on the whole utterances. Training segments for reconstruction and the ones used as input to the speaker conditioning network $g_t(Y)$ were different random segments extracted from the same training utterances. Noise schedule parameters β_0 and β_1 were set to 0.05 and 20.0.

Our VC models operated on mel-spectrograms with 80 mel features and sampling rate 22.05kHz. Short-Time Fourier Transform was used to calculate spectra with 1024 frequency bins. Hann window of length 1024 was applied with hop size 256.

For *Diff-LibriTTS* models we used simple spectral subtraction algorithm in mel domain with spectral floor parameter $\beta = 0.02$ as post-processing to reduce background noise sometimes produced by these models. Noise spectrum was estimated on speech fragments automatically detected as the ones corresponding to silence in source mel-spectrogram.

J DETAILS OF AMT TESTS

For fair comparison with the baselines all the recordings were downsampled to 16kHz; we also normalized their loudness. In speech naturalness tests workers chosen by geographic criterion were asked to assess the overall quality of the synthesized speech, i.e to estimate how clean and natural (human-sounding) it was. Five-point Likert scale was used: 1 - “Bad”, 2 - “Poor”, 3 - “Fair”, 4 - “Good”, 5 - “Excellent”. Assessors were asked to wear headphones and work in a quiet environment. As for speaker similarity tests, workers were asked to assess how similar synthesized samples sounded to target speech samples in terms of speaker similarity. Assessors were asked not to pay attention to the overall quality of the synthesized speech (e.g. background noise or incorrect pronunciation). Five-point scale was used: 1 - “Different: absolutely sure”, 2 - “Different: moderately sure”, 3 - “Cannot decide more same or more different”, 4 - “Same: moderately sure”, 5 - “Same: absolutely sure”.