



Expressive Speech Synthesis via Modeling Expressions with Variational Autoencoder

Kei Akuzawa¹, Yusuke Iwasawa¹, Yutaka Matsuo¹

¹Graduate School of Engineering, The University of Tokyo, Japan

{akuzawa-kei, iwasawa, matsuo}@weblab.t.u-tokyo.ac.jp

Abstract

Recent advances in neural autoregressive models have improved the performance of speech synthesis (SS). However, as they lack the ability to model global characteristics of speech (such as speaker individualities or speaking styles), particularly when these characteristics have not been labeled, making neural autoregressive SS systems more expressive is still an open issue. In this paper, we propose to combine VoiceLoop, an autoregressive SS model, with Variational Autoencoder (VAE). This approach, unlike traditional autoregressive SS systems, uses VAE to model the global characteristics explicitly, enabling the expressiveness of the synthesized speech to be controlled in an unsupervised manner. Experiments using the VCTK and Blizzard2012 datasets show the VAE helps VoiceLoop to generate higher quality speech and to control the expressions in its synthesized speech by incorporating global characteristics into the speech generating process.

Index Terms: autoregressive model, variational autoencoder, expressive speech synthesis

1. Introduction

Natural human speech is very expressive, and varies based on the speaker individualities (such as age and gender), emotions and speaking styles (see, e.g., [1, 2]). Many studies have suggested properly incorporating such expressiveness makes speech synthesis (SS) systems more pleasant to listen to and interact with, and have investigated the problem of expressive speech synthesis (see, e.g., [2, 3, 4]).

This paper addresses the problem of synthesizing expressive speech without relying on speech expression labels, which we refer to as *unsupervised expressive speech synthesis* (UESS). Many studies have reported that such labels are helpful for modeling complex audio data [5, 6, 4, 7]. Unsupervised methods, however, are more desirable because expressive speech is easy to obtain from video hosting websites (e.g., Youtube) or audiobooks but annotating such sources is costly. Moreover, manually-annotated labels are not always reliable: for example, not all emotions in a given category have the same strength [4].

Another important aspect of this paper is that it focuses on the *neural autoregressive models*, which have been shown to offer significant performance improvements to SS systems. For example, WaveNet [8] generates more natural speech than traditional parametric or unit-selection based SS methods. In addition, autoregressive-based sequence-to-sequence (seq2seq) speech synthesis models have simple structure, and can be trained on <text, audio> pairs with minimal human annotation. Such end-to-end systems have many advantages, for example, they alleviate the need for laborious feature engineering, which may involve heuristics, and are likely to be more robust than multi-stage models where each component's errors can compound [9].

However, finding ways to add the expressiveness to autoregressive SS models is still an open issue. One of the difficulties here is such models are typically unable to model the global characteristics of data because they model data densities autoregressively, i.e., point-by-point [10, 11]. Given that certain sources of speech expressiveness (e.g., gender or emotions) characterize speech in a global manner (sentence level), autoregressive SS models may suffer from the difficulty: it reduces the quality of the synthesized speech, and they have no structured way to control the expressions in the synthesized speech.

In this paper, we propose a model called VAE-Loop, which combines VoiceLoop [12], an autoregressive SS model, with Variational Autoencoder (VAE) [13]. Several recent studies (e.g., [14]) have shown VAEs can model global characteristics of speech such as speaker individualities, but to our knowledge no study has yet suggested using VAEs for SS or UESS. We use VAE to deal with the problem incorporating global characteristics into the speech generating process when using an autoregressive model for UESS. Specifically, VAE encodes such global characteristics as a tractable probability distribution, which is used to give hints about them to VoiceLoop, allowing it to generate higher quality speech and to control the expressions in the synthesized speech. The proposed VAE-based method is both effective and simple, trained in end-to-end as well as seq2seq SS models.

Our experiments show, by incorporating global characteristics in this way, the VAE can help VoiceLoop to attain lower test errors and higher mean opinion scores (MOSs) when no labels are available. Also, latent variables yielded by the VAE show it has the ability to control speaker individualities and speaking styles, and interpolate between them.

2. Related Work

Seq2seq SS systems have simple structures that directly predict acoustic features from text. In addition, they have demonstrated the ability to generate natural and intelligible speech [9, 15, 16] and to robustly handle different prosodies [12, 17]. [12] has shown even when VoiceLoop is trained on data obtained from YouTube containing various speaking styles, it can generate high quality speech. We have incorporated VoiceLoop into the proposed model, expecting that it will be effective for UESS.

Several studies tackled the problem of UESS on both neural autoregressive and non-neural paradigms. On the non-neural paradigm, representation of emotions, acquired by unsupervised learning methods such as clustering and principal component analysis, have been used to generate speech [2, 3, 18]. The most relevant research related to ours might be [19, 12], which proposed the seq2seq SS models that can learn and control speaking styles in an unsupervised manner. However, the proposed VAE-based method is different in that it learns speech expressions as a tractable distribution, which can be useful for

downstream tasks such as interpolation and semi-supervised learning.

[2, 18] pointed out that UESS can be divided into two parts: predicting expressive information from text; and synthesizing the speech with a particular expression. In this paper only the latter stage is considered for simplicity.

Several recent studies have proposed using VAEs for modeling speech [20, 21, 14, 22, 23]. The most relevant works might be [21, 23], which conditioned the VAE on speaker labels and perform voice conversion. In contrast, we perform SS by conditioning the model on text, and verify that the VAE as a SS model is also able to learn and control speech expressions.

Many other studies in areas outside the SS field have also proposed combining VAEs and autoregressive models. For example, it has been shown a recurrent neural network language model combined with VAE can generate sentences with consistent global characteristics (e.g., style, topics) [24]. That study also pointed out the problem that autoregressive models often ignore the latent variables obtained from the VAE. Several authors have proposed countermeasures for dealing with this problem [11, 25, 22]. Based on these studies, we employ the *KL cost annealing* approach used in [24] to alleviate this problem.

3. Models

In this section, we first introduce conditional VAE and VoiceLoop, which form the basis of the proposed method, and then present our VAE-Loop method.

3.1. Conditional Variational Autoencoder

Here we present the variant of VAE used in VAE-Loop, which is simply conditioned on the auxiliary features \mathbf{c} . In this study, \mathbf{x} and \mathbf{c} correspond to the acoustic features and phonemes, respectively.

Using the latent variable vector \mathbf{z} and the approximate distribution $q_\phi(\mathbf{z}|\mathbf{x}, \mathbf{c})$ (with parameter ϕ) of the true posterior $p_\theta(\mathbf{z}|\mathbf{x}, \mathbf{c})$ (with parameter θ), we can obtain the following lower bound $\mathcal{L}(\theta, \phi; \mathbf{x}, \mathbf{c})$ on the marginal likelihood of $p_\theta(\mathbf{x}|\mathbf{c})$:

$$\mathcal{L}(\theta, \phi; \mathbf{x}, \mathbf{c}) = \int q_\phi(\mathbf{z}|\mathbf{x}, \mathbf{c}) \log \frac{p_\theta(\mathbf{x}, \mathbf{z}|\mathbf{c})}{q_\phi(\mathbf{z}|\mathbf{x}, \mathbf{c})} dz \quad (1)$$

$$= \int q_\phi(\mathbf{z}|\mathbf{x}) \log \frac{p_\theta(\mathbf{x}, \mathbf{z}, \mathbf{c})p(\mathbf{z})}{q_\phi(\mathbf{z}|\mathbf{x})} dz \quad (2)$$

$$= -D_{KL}(q_\phi(\mathbf{z}|\mathbf{x})||p(\mathbf{z})) + \mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{x})}[\log p_\theta(\mathbf{x}|\mathbf{z}, \mathbf{c})] \quad (3)$$

where we have assumed $q_\phi(\mathbf{z}|\mathbf{x}, \mathbf{c}) = q_\phi(\mathbf{z}|\mathbf{x})$ and $p(\mathbf{z}|\mathbf{c}) = p(\mathbf{z})$ for simplicity. The prior $p(\mathbf{z})$ and approximate posterior $q_\phi(\mathbf{z}|\mathbf{x})$ are modeled by Gaussian distributions, namely $p(\mathbf{z}) = \mathcal{N}(\mathbf{z}|\mathbf{0}, \mathbf{I})$ and $q_\phi(\mathbf{z}|\mathbf{x}) = \mathcal{N}(\mathbf{z}|\boldsymbol{\mu}_\phi(\mathbf{x}), \boldsymbol{\sigma}_\phi^2(\mathbf{x})\mathbf{I})$.

During training, we update the parameters θ and ϕ to maximize \mathcal{L} . We call $q_\phi(\mathbf{z}|\mathbf{x})$ an encoder and $p_\theta(\mathbf{x}|\mathbf{z}, \mathbf{c})$ a decoder.

3.2. VoiceLoop

Let $\mathbf{x} = [\mathbf{x}_1, \dots, \mathbf{x}_T]$ be a variable length sequence of audio features that we want to predict. VoiceLoop can be regarded as a conditional autoregressive model with a parameter $\boldsymbol{\xi}$ as follows:

$$p_\xi(\mathbf{x}|\mathbf{c}) = \prod_{t=1}^T p_\xi(\mathbf{x}_t|\mathbf{x}_{1:t-1}, \mathbf{c}) \quad (4)$$

$$p_\xi(\mathbf{x}_t|\mathbf{x}_{1:t-1}, \mathbf{c}) = \mathcal{N}(\mathbf{x}_t|\boldsymbol{\mu}_\xi(\mathbf{x}_{1:t-1}, \mathbf{c}), \mathbf{I}) \quad (5)$$

where $\mathbf{x}_{1:t-1}$ is the vector of audio features between time steps 1 and $t-1$ and we estimate $p_\xi(\mathbf{x}_t|\mathbf{x}_{1:t-1}, \mathbf{c})$ in order for each time step $t \in T$. Eq.(5) assumes $p_\xi(\mathbf{x}_t|\mathbf{x}_{1:t-1}, \mathbf{c})$ is modeled by a Gaussian distribution with mean $\boldsymbol{\mu}_\xi$ and variance \mathbf{I} (identity matrix).

Next, we describe the procedure for estimating $p_\xi(\mathbf{x}_t|\mathbf{x}_{1:t-1}, \mathbf{c})$. VoiceLoop has a shifting buffer, which can be seen as a matrix $S \in \mathbb{R}^{d \times k}$ with columns $S[1] \dots S[k]$. At each time step, all the columns shift to the right as follows:

$$S_t[i+1] = S_{t-1}[i] \text{ for } 1 \leq i < k \quad (6)$$

$$S_t[1] = u \quad (7)$$

Here, u is a function of four parameters, namely the current attention-mediated context \mathbf{c}_t , buffer S_{t-1} itself, latest ‘‘spoken’’ output \mathbf{x}_{t-1} and speaker embedding s , as follows:

$$C_t = [\mathbf{c}_t + \tanh(F_u(s)), \mathbf{x}_{t-1}] \quad (8)$$

$$u = N_u([S_{t-1}, C_t]) \quad (9)$$

where $[a, b]$ is the concatenation of the two column vectors a and b to one column vector. VoiceLoop then estimates \mathbf{x}_t using the buffer S_t and embedding s , as follows:

$$\hat{\mathbf{x}}_t = N_o(S_t + F_o(s)) \quad (10)$$

where F_u, N_u, F_o and N_o are the respective neural networks, and $\hat{\mathbf{x}}_t$ is equivalent to $\boldsymbol{\mu}_\xi$ in Eq.(5).

3.3. Proposed model: VAE-Loop

VoiceLoop has no structured way to model the complex global characteristics in an unsupervised manner since it relies on the point-by-point estimation. In contrast, VAE-Loop explicitly incorporates them into the speech generating process in the VAE framework. Specifically, VAE-Loop regards VoiceLoop as a decoder for the conditional VAE, i.e., VoiceLoop is conditioned on the global latent variable \mathbf{z} .

3.3.1. Modeling various expressions using VAE

We first change VoiceLoop’s probability distribution (Eq.(4) and (5)) so that it is conditioned on the latent variable \mathbf{z} , as follows:

$$p_\theta(\mathbf{x}|\mathbf{z}, \mathbf{c}) = \prod_{t=1}^T p_\theta(\mathbf{x}_t|\mathbf{x}_{1:t-1}, \mathbf{z}, \mathbf{c}) \quad (11)$$

$$p_\theta(\mathbf{x}_t|\mathbf{x}_{1:t-1}, \mathbf{z}, \mathbf{c}) = \mathcal{N}(\mathbf{x}_t|\boldsymbol{\mu}_\theta(\mathbf{x}_{1:t-1}, \mathbf{z}, \mathbf{c}), \mathbf{I}) \quad (12)$$

where we have set $\boldsymbol{\xi} = \theta$ because VoiceLoop is regarded as the decoder in Eq.(3).

In the VAE framework, information which is useful to estimate \mathbf{x} but is not contained in the text \mathbf{c} is encoded into \mathbf{z} . Since certain types of expressions are difficult to predict from the spoken text alone, \mathbf{z} is expected to acquire latent representations of such expressions (i.e., the global characteristics).

3.3.2. Generating speech using the global latent variable

The $p_\theta(\mathbf{x}_t|\mathbf{x}_{1:t-1}, \mathbf{z}, \mathbf{c})$ in Eq.(12) is estimated by changing Eq.(9) to incorporate the latent variable \mathbf{z} into the speech generating process of VoiceLoop, as follows:

$$u = N_u([S_{t-1}, C_t, \mathbf{z}]) \quad (13)$$

As previously mentioned, \mathbf{z} is expected to acquire the expression information. In addition, since \mathbf{z} does not depend on the time step t unlike S_{t-1} and C_t , it conditions the speech generating process in a global manner.

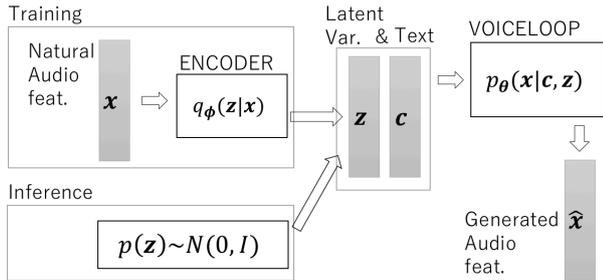


Figure 1: Speech generating process of VAE-Loop

3.3.3. Training and inference

By combining Eq.(3) and (11), we can obtain the objective function of VAE-Loop:

$$\mathcal{L}(\theta, \phi; \mathbf{x}, \mathbf{c}) = -D_{KL}(q_{\phi}(z|\mathbf{x})||p(z)) + \mathbb{E}_{q_{\phi}(z|\mathbf{x})} \left[\sum_{t=1}^T \log p_{\theta}(\mathbf{x}_t | \mathbf{x}_{1:t-1}, z, \mathbf{c}) \right] \quad (14)$$

where the first and second terms are the regularizer and reconstruction error, respectively. We can estimate the reconstruction error by taking the mean squared error between the estimators $\hat{\mathbf{x}}_t$ and true audio features \mathbf{x}_t . The second term is thus equivalent to the objective function of the original VoiceLoop, except that z is used in the generating process.

At training, z is sampled from the encoder. Here, as with a conventional VAE, the encoder $q_{\phi}(z|\mathbf{x})$ is parameterized as a deep neural network (DNN). At inference, z is sampled from the prior $p(z)$. Figure 1 illustrates the speech generating process of VAE-Loop, showing its training and inference procedures are simple; and do not require any additional training stage or data preprocessing compared with VoiceLoop alone. In addition, in spite of these simple procedures, it offers higher performance as we will demonstrate in Section 4.

3.3.4. KL cost annealing

We exploit the ideas outside the SS field and employ the simple *KL cost annealing* technique to alleviate the problem that autoregressive models often ignore the latent variables [24]. They argued that the latent variables were ignored because the regularizer in Eq.(14), which we call Kullback-Leibler divergence (KLD) term, acted too strongly at the start of training; therefore, Eq.(14) is adjusted to include the weight λ , as follows:

$$\mathcal{L}(\theta, \phi; \mathbf{x}, \mathbf{c}) = -D_{KL}(q_{\phi}(z|\mathbf{x})||p(z)) + \lambda \mathbb{E}_{q_{\phi}(z|\mathbf{x})} \left[\sum_{t=1}^T \log p_{\theta}(\mathbf{x}_t | \mathbf{x}_{1:t-1}, z, \mathbf{c}) \right] \quad (15)$$

We set λ to 0 at the start of training so that the model learns to encode as much information as possible, and then increase it lineally to 1 over the course of the annealing process.

4. Experiments

4.1. Datasets

We used two datasets: one featuring multiple speakers and another containing a variety of emotions and speaking styles. The first dataset was VCTK Corpus[26] (VCTK), which contains

Table 1: Test errors for different numbers of annealing epochs and z dimensions, on the VCTK dataset.

Model	annealing epochs	z-dim	Rec. error	KLD term	Total
VoiceLoop (w/o)	N/A	N/A	15.946	N/A	15.946
VoiceLoop (w/)	N/A	N/A	15.759	N/A	15.759
VAE-Loop	0	64	15.832	0.073	15.905
VAE-Loop	15(10%)	64	15.684	0.090	15.774
VAE-Loop	30(20%)	64	15.749	0.086	15.835
VAE-Loop	15	32	15.839	0.082	15.921
VAE-Loop	15	128	15.724	0.084	15.808

speech samples from 109 English speakers. We used the version of the dataset from VoiceLoop’s source code page¹ instead of the complete VCTK, in order to replicate the conditions of [12]. This contained about 5 hours of speech by 21 North American speakers (4 males and 17 females), and each utterance lasted less than 5 seconds. Our second dataset was from the Blizzard Challenge 2012 (Blizzard2012), and consisting of four audiobooks [27, 28]. Audiobooks are often used by expressive speech synthesis studies because they include a variety of emotions and speaking styles. Unlike those in the VCTK, all the utterances in this dataset were read by the same male speaker. To match the conditions with VCTK and avoid exploding gradients, we used utterances of less than 5 seconds only, resulting in a total of about 10 hours of speech. Both of the datasets were divided into three parts, with 90% used for training and the remaining 10% used for validation, with 50 samples set aside as test data.

4.2. Experimental setup

We used DNN based on time-domain convolution for the encoder of VAE-Loop. Specifically, the first half of the encoder consisted of five repeated convolutional layers with a stride size 2, with dropout, batch normalization and ReLU, while the rest consisted of time-domain global max-pooling and fully connected layers. The model hyperparameters used for the baseline VoiceLoop and VAE-Loop decoder were the same as in the authors’ implementation¹.

During training, we employ a variant of teacher forcing technique as well as the original VoiceLoop paper [12], which aims to stabilize both training and inference. We refer to this as *semi-teacher-forcing*. Specifically, \mathbf{x}_{t-1} used as input to the network N_u in Eq.(8) is replaced by $\tilde{\mathbf{x}}_{t-1}$ below:

$$\tilde{\mathbf{x}}_{t-1} = \frac{\mathbf{x}_{t-1} + \hat{\mathbf{x}}_{t-1}}{2} + \eta \quad (16)$$

where we assumed $\eta \sim \mathcal{N}(0, \mathbf{I})$.

Unless otherwise noted, we used the Adam [29] optimizer and 150 training epochs. The learning rate was chosen to achieve the minimum possible validation error from the set [1e-3, 1e-4, 5e-5, 1e-5], resulting in 1e-4 and 5e-5 for the VCTK and Blizzard2012, respectively.

4.3. Effect of latent variables on the test error

We compared the test errors for VAE-Loop with those for VoiceLoop alone, to demonstrate how adding latent variables to VoiceLoop enables it to estimate audio features more accurately. For this experiment, the models were trained on the VCTK, using the setup described in Section 4.2. However, to stabilize training on various hyperparameters, we set the learning rate to 5e-5. In addition, since the baseline VoiceLoop had not converged sufficiently after 150 epochs at that learning rate,

¹<https://github.com/facebookresearch/loop/>

Table 2: Mean opinion scores (mean \pm CI) for both datasets.

Method	VCTK	Blizzard2012
Ground Truth	4.07 \pm 0.23	3.94 \pm 0.30
VoiceLoop(w/o)	2.51 \pm 0.34	2.23 \pm 0.24
VoiceLoop(w/)	3.24 \pm 0.27	N/A
VoiceLoop(orig, w/)	3.57	N/A
VAE-Loop($\sigma = 1$)	3.25 \pm 0.29	2.47 \pm 0.32
VAE-Loop($\sigma = 0.7$)	N/A	2.89 \pm 0.32
VAE-Loop($\sigma = 0$)	N/A	3.03 \pm 0.32

we extended the training period to 200 epochs. We tested with annealing for 10 or 20% of the training period, and without annealing. The test errors were calculated using semi-teacher-forcing in order to use the same objective function as during training.

Table 1 presents the test errors, calculated using Eq.(14) and then divided by the sequence length. Here, (w) and (w/o) mean “with speaker labels” and “without speaker labels” respectively. These show proper use of KL cost annealing leads to a higher KLD term and a lower test error, suggesting it allows the decoder (VoiceLoop) to receive more useful information from the latent variables. Moreover, the test errors of VAE-loop are smaller than that of VoiceLoop without speaker labels, suggesting incorporating latent variables into the speech generating process enables VoiceLoop to estimate audio features more accurately.

4.4. Mean opinion score tests

To demonstrate that incorporating global characteristics enables VAE-Loop to generate higher quality speech, we conducted a mean opinion score (MOS) study, using the crowdMOS toolkit [30] and Amazon Mechanical Turk. The MOS is a popular subjective audio quality measure, obtained by asking people to rate the audio’s naturalness on a scale of 1 to 5. More than 15 people living in the US rated each of the two datasets. Table 2 shows MOSs for the two models, together with their 95% confidence intervals(CIs). Here, “Ground Truth” recordings were the audio reconstructed using the WORLD vocoder [31].

For the VCTK, the MOS achieved by VAE-Loop was higher than that by VoiceLoop without speaker labels, matching even VoiceLoop with labels, despite not using labels. In addition, in our informal listening tests, we observed VAE-Loop was less likely than the baseline to generate unintelligible speech (e.g., several seconds of just breath or a certain phoneme). Therefore, these results could indicate that where the original VoiceLoop struggled to model the various speaker individualities, adding VAE stabilized its speech generating process by giving hints about them. Here, we acknowledge that VoiceLoop’s MOS by our implementation is lower than that reported in [12] (“VoiceLoop(orig, w/)” in Table 2), probably because there might be a different choice of hyperparameters, including the use of pre-training.

For the Blizzard2012 data, we observed that the high variance of the $p(z)$ used for generating test samples meant VAE-Loop often generated unintelligible speech in much the same way as VoiceLoop. To investigate this issue, we instead assumed that $p(z) = \mathcal{N}(z|\mathbf{0}, \sigma^2 \mathbf{I})$ at inference time, and sampled z using different parameters σ , where $\sigma = 0$ means we always sample $z = \mathbf{0}$. When the variance of $p(z)$ was suppressed, in this way, VAE-Loop’s MOS improved, exceeding that of the baseline. Here, note that using small σ values means always sampling similar z values; therefore, VAE-Loop can make a tradeoff between stable inference and latent variable va-

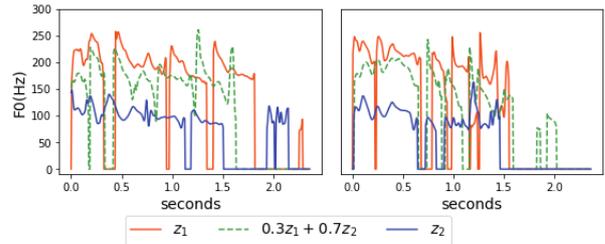


Figure 2: F0 Trajectories for two utterances generated by VAE-Loop, trained on VCTK. Here, z_1 and z_2 correspond to high-pitched (female as we heard) and low-pitched (male) voices, respectively. Averaged F0 trajectories are also shown, generated by interpolating between z_1 and z_2 .

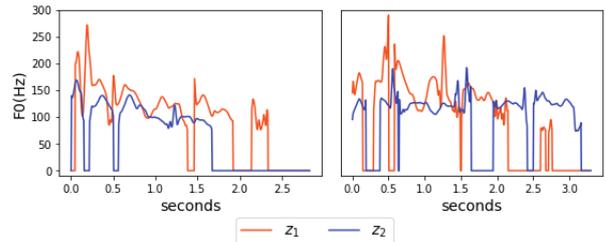


Figure 3: F0 Trajectories for two utterances generated by VAE-Loop, trained on Blizzard2012. Here, z_1 and z_2 correspond to voices with large (dramatic as we heard) and small (calm) pitch fluctuations, respectively.

riety.

4.5. Controlling speech expressions using latent variables

To demonstrate that VAE-Loop is able to control the expressions in its synthesized speech, we presented the trajectories of the fundamental frequency (F0). Figure 2 shows F0 trajectories generated by VAE-Loop, trained on the VCTK. The left and right figures correspond to different texts; however, both were generated using the same z values. Here, different latent variable values, z_1 and z_2 , lead to different F0 characteristics, indicating our model can control speaker individualities expressed in the synthesized speech using latent variables. Moreover, when the speech is synthesized using a latent variable value that interpolated between previous two, the F0 trajectories were also averaged. Likewise, Figure 3 shows F0 trajectories generated by VAE-Loop, trained on the Blizzard2012. Here, the latent variables characterize the pitch fluctuations of the F0 trajectories. Some audio samples can be found at:

<https://akuzeee.github.io/VAELoopDemo/>.

5. Conclusions

In this paper, we have proposed to combine VoiceLoop with VAE, in order to enable this autoregressive SS model to be more expressive by using VAE to help model a range of expressions. Even though autoregressive SS models have shown promising results, they typically lack the ability to model the global characteristics of speech. However, the proposed method can incorporate such expressions explicitly into the speech generating process in an unsupervised manner. Our experiments have shown taking advantage of these global characteristics could enable our method to generate higher quality speech than VoiceLoop without labels and to control speech expressions.

In future studies, we plan to extend this approach to semi-supervised learning with a small amount of labeled data, and to infer the latent variables from text.

6. References

- [1] D. Erickson, “Expressive speech: Production, perception and application to speech synthesis,” *Acoustical Science and Technology*, vol. 26, no. 4, pp. 317–325, 2005.
- [2] F. Eyben, S. Buchholz, and N. Braunschweiler, “Unsupervised clustering of emotion and voice styles for expressive TTS,” in *2012 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, March 2012, pp. 4009–4012.
- [3] M. Charfuelan and I. Steiner, “Expressive speech synthesis in MARY TTS using audiobook data and emotionML,” in *Proc. interspeech 2013*, 2013, pp. 1564–1568.
- [4] G. E. Henter, J. Lorenzo-Trueba, X. Wang, and J. Yamagishi, “Principles for Learning Controllable TTS from Annotated and Latent Variation,” in *Proc. Interspeech 2017*, 2017, pp. 3956–3960.
- [5] Y. Fan, Y. Qian, F. K. Soong, and L. He, “Multi-speaker modeling and speaker adaptation for DNN-based TTS synthesis,” in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, April 2015, pp. 4475–4479.
- [6] H. T. Luong, S. Takaki, G. E. Henter, and J. Yamagishi, “Adapting and controlling DNN-based speech synthesis using input codes,” in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, March 2017, pp. 4905–4909.
- [7] Y. Lee, A. Rabiee, and S. Lee, “Emotional End-to-End Neural Speech Synthesizer,” *CoRR*, vol. abs/1711.05447, 2017. [Online]. Available: <http://arxiv.org/abs/1711.05447>
- [8] A. van den Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. W. Senior, and K. Kavukcuoglu, “WaveNet: A Generative Model for Raw Audio,” *CoRR*, vol. abs/1609.03499, 2016. [Online]. Available: <http://arxiv.org/abs/1609.03499>
- [9] Y. Wang, R. Skerry-Ryan, D. Stanton, Y. Wu, R. J. Weiss, N. Jaitly, Z. Yang, Y. Xiao, Z. Chen, S. Bengio, Q. Le, Y. Agiomyrgiannakis, R. Clark, and R. A. Saurous, “Tacotron: Towards End-to-End Speech Synthesis,” in *Proc. Interspeech 2017*, 2017, pp. 4006–4010.
- [10] A. Kolesnikov and C. H. Lampert, “PixelCNN models with auxiliary variables for natural image modeling,” in *Proc. 34th International Conference on Machine Learning*, vol. 70, 2017, pp. 1905–1914.
- [11] I. Gulrajani, K. Kumar, F. Ahmed, A. A. Taiga, F. Visin, D. Vazquez, and A. Courville, “PixelVAE: A Latent Variable Model for Natural Images,” in *Proc. 5th International Conference on Learning Representations*, 2017.
- [12] Y. Taigman, L. Wolf, A. Polyak, and E. Nachmani, “VoiceLoop: Voice Fitting and Synthesis via a Phonological Loop,” in *Proc. 6th International Conference on Learning Representations*, 2018.
- [13] D. P. Kingma and M. Welling, “Auto-Encoding Variational Bayes,” in *Proc. 2nd International Conference on Learning Representations*, 2014.
- [14] W.-N. Hsu, Y. Zhang, and J. Glass, “Learning latent representations for speech generation and transformation,” in *Proc. Interspeech 2017*, 2017, pp. 1273–1277.
- [15] J. Sotelo, S. Mehri, K. Kumar, J. F. Santos, K. Kastner, A. Courville, and Y. Bengio, “Char2Wav: End-to-end speech synthesis,” in *International Conference on Learning Representations (Workshop Track)*, April 2017.
- [16] W. Ping, K. Peng, A. Gibiansky, S. O. Arik, A. Kannan, S. Narang, J. Raiman, and J. Miller, “Deep Voice 3: Scaling Text-to-Speech with Convolutional Sequence Learning,” in *Proc. 6th International Conference on Learning Representations*, 2018.
- [17] S. Ronanki, O. Watts, and S. King, “A Hierarchical Encoder-Decoder Model for Statistical Parametric Speech Synthesis,” in *Proc. Interspeech 2017*, 2017, pp. 1133–1137.
- [18] L. Chen, M. J. F. Gales, N. Braunschweiler, M. Akamine, and K. Knill, “Integrated Expression Prediction and Speech Synthesis From Text,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 8, no. 2, pp. 323–335, April 2014.
- [19] Y. Wang, R. J. Skerry-Ryan, Y. Xiao, D. Stanton, J. Shor, E. Battenberg, R. Clark, and R. A. Saurous, “Uncovering Latent Style Factors for Expressive Speech Synthesis,” *CoRR*, vol. abs/1711.00520, 2017. [Online]. Available: <http://arxiv.org/abs/1711.00520>
- [20] M. Blaauw and J. Bonada, “Modeling and Transforming Speech Using Variational Autoencoders,” in *Proc. Interspeech 2016*, 2016, pp. 1770–1774.
- [21] C. C. Hsu, H. T. Hwang, Y. C. Wu, Y. Tsao, and H. M. Wang, “Voice conversion from non-parallel corpora using variational auto-encoder,” in *2016 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA)*, Dec 2016, pp. 1–6.
- [22] A. van den Oord, O. Vinyals, and k. kavukcuoglu, “Neural Discrete Representation Learning,” in *Advances in Neural Information Processing Systems 30*, 2017, pp. 6306–6315.
- [23] C.-C. Hsu, H.-T. Hwang, Y.-C. Wu, Y. Tsao, and H.-M. Wang, “Voice conversion from unaligned corpora using variational autoencoding wasserstein generative adversarial networks,” in *Proc. Interspeech 2017*, 2017, pp. 3364–3368.
- [24] S. R. Bowman, L. Vilnis, O. Vinyals, A. M. Dai, R. Józefowicz, and S. Bengio, “Generating Sentences from a Continuous Space,” in *Proceedings of the 20th SIGNLL Conference on Computational Natural Language Learning*, 2016.
- [25] X. Chen, D. P. Kingma, T. Salimans, Y. Duan, P. Dhariwal, J. Schulman, I. Sutskever, and P. Abbeel, “Variational Lossy Autoencoder,” in *Proc. 5th International Conference on Learning Representations*, 2017.
- [26] C. Veaux, J. Yamagishi, and K. MacDonald, “CSTR VCTK Corpus: English Multi-speaker Corpus for CSTR Voice Cloning Toolkit,” 2017. [Online]. Available: <http://dx.doi.org/10.7488/ds/1994>
- [27] S. King and V. Karaiskos, “The blizzard challenge 2012,” in *Proc. Blizzard Challenge workshop*, 2012.
- [28] N. Braunschweiler, M. J. F. Gales, and S. Buchholz, “Lightly supervised recognition for automatic alignment of large coherent speech recordings,” in *Proc. interspeech 2010*, 2010, pp. 2222–2225.
- [29] D. P. Kingma and J. Ba, “Adam: A Method for Stochastic Optimization,” in *Proc. 3rd International Conference on Learning Representations*, 2015.
- [30] F. Ribeiro, D. Florêncio, C. Zhang, and M. Seltzer, “CROWD-MOS: An approach for crowdsourcing mean opinion score studies,” in *2011 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, May 2011, pp. 2416–2419.
- [31] M. Morise, F. YOKOMORI, and K. Ozawa, “WORLD: A Vocoder-Based High-Quality Speech Synthesis System for Real-Time Applications,” in *IEICE Transactions on Information and Systems*, vol. E99.D, 07 2016, pp. 1877–1884.