FULLY-HIERARCHICAL FINE-GRAINED PROSODY MODELING FOR INTERPRETABLE SPEECH SYNTHESIS

Guangzhi Sun[1], Yu Zhang[2], Ron J. Weiss[2], Yuan Cao[2], Heiga Zen[2], Yonghui Wu[2]*

[1]University of Cambridge   [2]Google

ABSTRACT

This paper proposes a hierarchical, fine-grained and interpretable latent variable model for prosody based on the Tacotron 2 text-to-speech model. It achieves multi-resolution modeling of prosody by conditioning finer level representations on coarser level ones. Additionally, it imposes hierarchical conditioning across all latent dimensions using a conditional variational auto-encoder (VAE) with an auto-regressive structure. Evaluation of reconstruction performance illustrates that the new structure does not degrade the model while allowing better interpretability. Interpretations of prosody attributes are provided together with the comparison between word-level and phone-level prosody representations. Moreover, both qualitative and quantitative evaluations are used to demonstrate the improvement in the disentanglement of the latent dimensions.

Index Terms— text-to-speech, Tacotron 2, fine-grained VAE, hierarchical

1. INTRODUCTION

Significant developments have taken place in the neural end-to-end text-to-speech (TTS) synthesis models for generating high fidelity speech with a simplified pipeline [1–4]. Such systems usually incorporate an encoder-decoder neural network architecture [5] that maps a given text sequence to a sequence of acoustic features. More recent advancement in such models enables the use of crowd-sourced data by disentangling and controlling different attributes such as speaker identity, noise, recording channels as well as prosody [6–8]. The focus of this paper, *prosody*, is a collection of attributes including fundamental frequency (F_0), energy and duration [9]. Efforts have been made to model and control these attributes by factorizing the latent attributes (e.g. prosody) from observed attributes (e.g. speaker). Although most of these works use latent representations at utterance level which captures the salient features of the utterance [6, 10–12], fine-grained prosody that are aligned with the phone sequence can be captured using techniques recently proposed in [13]. This model provides a localized prosody control that achieves more variability and higher robustness to speaker perturbations.

Even though prosody attributes such as F_0 and energy can be treated as latent features, interpreting the phone-level latent space is still difficult since latent dimensions can be entangled with each other. Moreover, coherence of the prosody within a word (e.g. accented syllables), noise level and channel properties are important attributes not captured at the phone-level alone. Respecting the hierarchical nature of spoken language and aiming at interpretation of prosody at fine-scale such as F_0 for a vowel, this paper aims to achieve disentangled control of each prosody attribute at different levels.

This paper proposes a multilevel model based on Tacotron 2 [14] integrated with a hierarchical latent variable model. In addition to the

prosody representation at utterance level, the representation is also extracted at word and phone levels. Apart from utterance-level characteristics such as noise and channel properties, phone-level prosodic features are expected to capture fine-grained information associated with each phone, and word-level features are expected to capture the prosody at each word while maintaining a natural prosody structure within the word. To better interpret the representation of each latent dimension, the original VAE is replaced by a conditional VAE driven by the information contained in the previous latent dimensions. This setup gives a hierarchy where finer level features are conditioned on coarser, and latent variables at each level are hierarchically factorized. The proposed model is thus referred to as a fully-hierarchical VAE. Furthermore, imposing a training schedule for each latent dimension results in a phone-level representation which reflects a consistent ordering of prosody attributes. Finally, we assess the disentanglement property of our model on three most significant attributes.

2. PRIOR WORK

Abundant research has been performed on learning latent representations for styles and prosody [15, 16] such as the use of an utterance-level VAE in [17]. Our multilevel model is based on the fine-grained VAE structure which extends the idea in [13], and is closely related to the hierarchical VQVAE model in [18]. While the latter uses down sampling to extract coarser features for image processing, our proposed model takes advantage of the hierarchical structure of spoken language. The multilevel alignment is also similar to the multilevel information extraction model proposed by [19].

Meanwhile, exhaustive exploration has been made in the unsupervised learning of disentangled latent representations these years in various scenarios including speech recognition [20, 21]. Progress have been made mostly in the direction of learning independently distributed latent variables without associating them with the actual latent factors [22–24]. However, [25] demonstrated the impossibility of unsupervised learning of disentangled representations without any inductive bias. They pointed out that there exists an infinite number of bijective mappings from the learned latent space to another space with the same marginal distribution, but the two spaces are fully entangled. Other works try to learn disentangled representations via semi-supervised learning [26–28] that guides a subset of latent variables to learn some labelled features, or via adversarial training [29, 30]. The most similar hierarchical decomposition to our approach is proposed by [31], but the intention of this decomposition is to facilitate learning statistically independent random variables.

3. MULTILEVEL PROSODY MODELING STRUCTURE

Different from utterance-level VAE [6] where a single latent feature is extracted for each utterance, the fine-grained VAE [13] aligns the

---

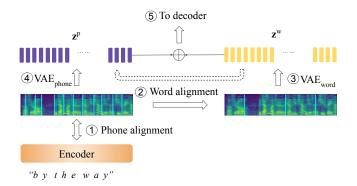*Work performed while interning at Google Brain.

Fig. 1. Multilevel fine-grained VAE structure on the encoder side. Phone-level alignment implemented with a location-sensitive attention. Utterance-level latent features can be extracted separately and used as input to steps ③, ④ and ⑤.

Fig. 2. **Left**: graphical models for a 3-dimensional latent space. ① shows the generation process, while ② and ③ show the inference process. **Right**: conditional VAE structure where Proj. is a projection layer mapping each latent dimension to a higher dimension.

target spectrogram with the phone sequence and extract a sequence of phone-level latent prosody features. These latent prosody features are concatenated with their corresponding phone encodings before sending to the decoder. Extending the fine-grained VAE, an illustration of our proposed multilevel model is shown in Fig. 1.

This structure is integrated with the encoder of the Tacotron-2 [14]. Location-sensitive attention [32] is used to align the target spectrogram with encodings of each phone at step 1. After the aligned target spectrogram is obtained, the average of spectrograms associated with phones in each word is calculated using the known phone-word alignment. The word-level latent prosody features are then extracted from these averaged spectrograms. Phone-level latent features are extracted conditioning on word-level latent features, and both features are concatenated using the phone-word alignment again. These features are used by the decoder for reconstruction. The system is optimized with the multilevel evidence lower bound (ELBO):

$$
\begin{aligned}
\mathcal{L}(p, q) = {} & \mathbb{E}_{q(\mathbf{z}|\mathbf{X})}\left[\log p(\mathbf{X} \mid \mathbf{Y}, \mathbf{z})\right] \\
& - \beta_1 \sum_{n=1}^{N} \mathbb{E}_{q(\mathbf{z}_{f(n)}^w|\mathbf{X})}\left[D_{\mathrm{KL}}\big(q(\mathbf{z}_n^p \mid \mathbf{X}, \mathbf{z}_{f(n)}^w) \,\|\, p(\mathbf{z}_n^p)\big)\right] \\
& - \beta_2 \sum_{m=1}^{M} D_{\mathrm{KL}}(q(\mathbf{z}_m^w \mid \mathbf{X}) \,\|\, p(\mathbf{z}_m^w)), \qquad (1)
\end{aligned}
$$

where M is the number of words, N is the number of phones and βs have the same function as [22]. $\mathbf{z}$ refers to the sequence of concatenated phone and word-level latent features. $f(n)$ maps the phone index n to the corresponding word index. $\mathbf{z}^w$ and $\mathbf{z}^p$ represent latent features associated with each word and each phone respectively. The model incorporates the utterance-level feature $\mathbf{z}^u$ by conditioning other fine-grained latent features on the utterance level latent one, and subtracting $\beta_3 D_{\mathrm{KL}}(q(\mathbf{z}^u \mid \mathbf{X}) \,\|\, p(\mathbf{z}^u))$ in Eq.(1) accordingly. For a more complete model one could also introduce the dependency on text and speaker information for the posterior, e.g. $q(\mathbf{z}_n^p \mid \mathbf{X}, \mathbf{Y}_n^p, \mathbf{S}, \mathbf{z}_n^w)$ and $q(\mathbf{z}_m^w \mid \mathbf{X}, \mathbf{Y}_m^w, \mathbf{S})$ where $\mathbf{S}$ is the speaker embedding and $\mathbf{Y}^w$ and $\mathbf{Y}^p$ for phone and word encodings.

4. INTERPRETABLE CONDITIONAL VAE STRUCTURE

In the previous multilevel structure, the VAE layer models the data distribution in the latent space as a multi-dimensional Gaussian distribution with diagonal covariance matrices, which is based on the assumption that different latent dimensions have independent effects to the prosody attributes. A corresponding graphical model is illustrated in Fig. 2 where the graph 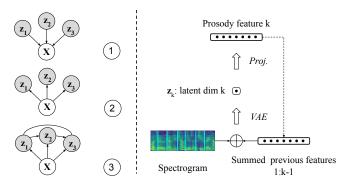with label 1 shows the generation sub-graph and label 2 shows the corresponding inference sub-graph. When inferring the posterior of z_2, the model in label 2 where z are independent of each other indicates when the observation X is given, knowing what z_1 represents does not provide any further information. However, if latent dimensions control disentangled factors and z_1 is known to represent the energy of the phone, it adds extra information to z_2 indicating that z_2 captures attributes other than energy. Therefore, this conditional dependency as shown by the graph with label 3 should be modeled when inferring the posterior distribution.

To incorporate the dependency into inference, we extend the hierarchical structure described in the previous section further to include a conditional VAE according to an auto-regressive decomposition of the posterior. The model structure is shown in the right part of Fig. 2. Unlike auto-regressive density estimators [33, 34] which uses recurrent structures directly on top of the latent variables, our recurrent model conditions on the projection of latent variables and extract latent dimensions one at a time. Specifically, when extracting the k-th latent dimension, the input to the VAE is the aligned spectrogram concatenated with the summation of all previously extracted latent features. Because these latent dimensions after projection are directly used by the decoder, they are effectively representing the prosody attributes being captured. Using these features with the aligned spectrogram as inputs to the VAE implicitly encourages the current latent dimension k to extract information about the prosody attribute other than what has already been represented. The training objective for this VAE model can still be written in the ELBO form as shown in Eq. (2), where expectation is estimated by single sample for the KL-divergence between the auto-regressive posterior $q(\mathbf{z} \mid X)$ and the prior $p(\mathbf{z})$.

$$
\begin{aligned}
\mathcal{L}(p, q) = {} & \mathbb{E}_{q(\mathbf{z}|\mathbf{X})}\left[\log p(\mathbf{X} \mid \mathbf{Y}, \mathbf{z})\right] \\
& - \beta \sum_{k=1}^{d} D_{\mathrm{KL}}(q(z_k \mid \mathbf{Z}_{1:k-1}, \mathbf{X}) \,\|\, p(z_k)), \qquad (2)
\end{aligned}
$$

where $\mathbf{Z}_{1:k-1}$ are samples of dimensions 1 to $k-1$ from their posterior distributions. The prior uses the isometric standard normal distribution for each latent dimension. When applied to the multilevel framework, this loss function is minimized at each time step for each level and the subscription n and m in Eq. (1) for phone and word indices can be directly added to each latent variable. Combining the two approaches where auto-regressive decomposition of the posterior is applied across different levels and different latent dimensions, the model thus covers the full hierarchy from phone to utterance level.

2

Finally, disentangled prosody features are observed to be extracted following an energy-duration-F_0 order guided by scheduled training across latent dimensions. Scheduled training refers to the process where the first latent dimension is trained for a certain number of steps before the second dimension starts to train. The rest of the dimensions are started consecutively in the same way. Therefore, when scheduling is imposed together with the conditional VAE, the first latent dimension will capture the energy information, and the second dimension being aware that energy has been represented by the first dimension, will seek for representations of the duration.

5. EXPERIMENTS

The proposed models are evaluated on the LibriTTS multi-speaker audiobook dataset [35] and the Blizzard Challenge 2013 single-speaker audiobook dataset [36]. LibriTTS includes approximately 585 hours of read English audiobooks at 24kHz sampling rate. It covers a wide range of speakers, recording conditions and speaking styles. The latent space is expected to control the prosody without affecting speaker characteristics. On the other hand, the Blizzard Challenge 2013 dataset contains 147 hours US English speech with highly varying prosody, recorded by a female professional speaker.

Three attributes are considered in this paper for fine-grained prosody interpretation including F_0, energy and duration. To quantitatively evaluate each attributes, we leverage the decoder alignment attention weights to obtain the duration of the phone by counting the frames which have a peak value at that specific phone in their attention weights. After obtaining the duration frames n_1 to n_2 and also converting to signal sample indices t_1 to t_2, the energy can be estimated using the average signal magnitude in $[t_1, t_2-1]$ divided by the average signal magnitude of the whole utterance. F_0 can be similarly measured using the average F_0 estimated from an F_0 tracker [37] among the frames in $[n_1, n_2 - 1]$. To decrease the variance due to bad alignments, we exclude 50 samples at both margins.

Finally, the mel-cepstral distortion (MCD), the F_0 Frame Error (FFE) [38], which is a combination of the Gross Pitch Error (GPE) and the Voicing Decision Error (VDE), are used to quantify the reconstruction performance. FFE evaluates the reconstruction of the F_0 track, and MCD evaluates the timbral distortion. We strongly recommend readers to listen to the samples on the demo page [39].

5.1. Reconstruction Performance

Table 1 shows reconstruction performance measured using FFE and MCD_{13} for the first 13 MFCCs. Lower is better for both metrics. Both the fine-grained VAE with 3-dimensional latent features and the fully-hierarchical VAE achieve similar reconstruction performance. Furthermore, progressive improvements from global to the system with 3-dimensional latent space can be interpreted: By introducing fine-grained VAE, there is a significant drop in VDE as the fine-grained VAE could capture phone-level energy and duration information. Conditioning the posterior distribution on the speaker embedding at the encoder side significantly reduced GPE because the speaker identity is closely related to the average F_0. Increasing the latent space size to 3 again significantly reduces the F_0 error, which confirms the fact that F_0 information is the last to be captured.

5.2. Multilevel Controllability

The model selected for the demonstration is trained on the LibriTTS dataset, and both phone-level and word-level latent spaces are 3-dimensional. To illustrate the effect of controlling a single attribute at different levels clearly, we traverse one dimension of the latent features to control a certain attribute while keeping other dimensions

Model	GPE	VDE	FFE	MCD
Global VAE	0.39	0.34	0.52	16.0
Phone-level VAE 2d (no spk.)	0.33	0.18	0.35	10.5
Phone-level VAE 2d	0.25	0.15	0.28	9.0
Phone-level VAE 3d	0.10	0.12	0.18	8.6
Phone-level conditional VAE	0.10	0.13	0.20	9.2
Fully-hierarchical VAE	0.10	0.12	0.19	8.8

Table 1. Reconstruction performance results. 2d and 3d refers to the dimension of the latent space. 3-dimensional latent space is used for the conditional and the fully-hierarchical VAE. If not specified, the posterior is conditioned on the speaker embedding.

constant. We demonstrate the control of a single vowel or a word using phone-level or word-level latent features in Fig 4.

The effect on a single vowel is clear: Increasing F_0 raised harmonic frequencies within that phone. Increasing energy brightens the area in the box while darkens the rest as signals are normalized. Increasing duration stretched the corresponding area. Because the influence at the word level contains a mixture of effects on vowels and consonants, when changing the dimension controlling energy, the duration of the phone n also changes which affects the duration of the word. However, the most significant changes can still be interpreted as the same three attributes exerted on all the vowels in the word. Meanwhile, the traversing of each latent dimension for a phone at different word prosody level is shown in Fig. 3. The control of each latent attribute is linear when traversing each dimension from -1 to $+1$. The word-level control shifts the phone-level curves up and down as the phone-level latent features are conditioned on the word-level.

Furthermore, adjusting word-level prosody features also retains the prosody distribution within a word while phone-level adjustment required a manual assignment to keep it natural. This effect can be evaluated by generating utterances with phone-level or word-level independent sampling of one latent feature. Subjective mean opinion score (MOS) test results are shown in Table 2 where random samples of the latent dimension controlling the F_0 were used. Consequently, the word-level independent sampling sounds more natural, as the prosody structure within each word is retained as neutral prosody.

Sampling level	MOS
Neutral-prosody	4.04 ± 0.06
Phone-level	3.75 ± 0.09
Word-level	3.91 ± 0.10

Table 2. MOS evaluation of speech generated with the phone/word level independent F_0 sampling. When sampling at one level, the other is set to all zero to give neutral prosody.

5.3. Improved phone-level interpretability

The interpretability is improved by the disentanglement for the three prosody attributes with the conditional VAE model. To illustrate this improvement, a vowel was selected and its F_0, energy and duration were measured with the method in Sec. 5. For each model, 100 samples were generated by drawing from a standard Gaussian distribution for one latent dimension while keeping other dimensions constant. Then, the standard deviations for measured attributes were
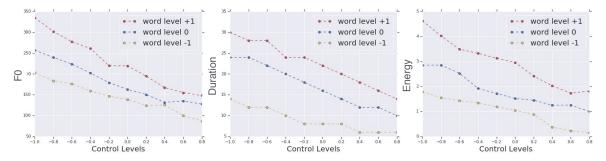
Fig. 3. Prosody variations when traversing the nominated controlling dimension in the latent space under different word latent values.
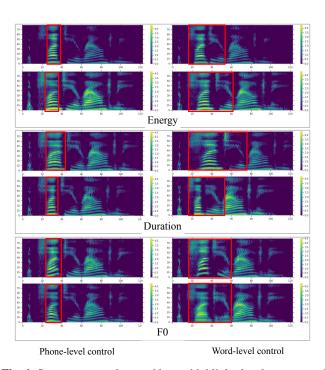

Fig. 4. Spectrograms where red boxes highlight the phone or word being controlled. For the two spectrograms in each sub-figure, the corresponding latent dimensions are taking $-2, +2$ respectively. the word of interest is `plenty` and the phone is `e`.

Model	Average variance ratio
Baseline fine-grained VAE	7.5 ± 1.0
DIP-VAE I	9.9 ± 1.9
Fully-hierarchical VAE	$\mathbf{11.5 \pm 2.3}$

Table 3. Variance ratio for different influencing control factors associated with a vowel on LibriTTS. DIP-VAE-I refers the model proposed in [24] which essentially enforces the covariance matrix of the marginal posterior $q(\mathbf{z})$ to be diagonal.

Model	Average variance ratio
Baseline fine-grained VAE	5.8 ± 0.8
Fully-hierarchical VAE	$\mathbf{8.0 \pm 2.9}$

Table 4. Variance ratio for different influencing control factors associated with a vowel on the single-speaker audio book dataset.

latent dimension, each attribute is adjusted by traversing the corresponding latent dimension. Additionally, when training schedule is imposed on latent dimensions, the order of prosody attributes being captured is always found to be energy, duration, and F_0 on both datasets. Energy is the amplitude of the signal which is directly related to the reconstruction loss and is easier to be captured, and F_0 is the last which coincides with the findings from the reconstruction evaluation. Moreover, the first dimension captures the duration of silence since that is the most significant attribute. The effect of latent features for a consonant in spectrograms can also be categorized into these three attributes but is hard to interpret directly from the audio.

6. CONCLUSIONS

A fully-hierarchical model to achieve multilevel control of prosody attributes is proposed in this paper. The model consists of a hierarchical structure across different levels covering phone, word and utterance. Besides, a conditional VAE is applied at the phone and word-level which also adopts a hierarchical structure across all latent dimensions. Experimental results demonstrate improved interpretability by showing improved disentanglement, and the order of prosody attributes to be extracted is explained. Furthermore, the difference in phone and word level control effects is also analyzed.

7. ACKNOWLEDGEMENTS

calculated and scaled to lie in a similar range. Next, the ratio of standard deviations between the attribute under control and the attribute with the second-largest deviation was obtained to represent the disentanglement in that dimension. This was repeated for each latent dimension and the sum of the ratios for each repetition is used to generally represent the degree of disentanglement. Though the ratio is not capturing the exact entanglement, it still reflects the degree of disentanglement because disentangled systems should have a much larger variation in one factor than the other two when only one latent dimension is varying. The experiment with each model was repeated 5 times with different random seeds to show a consistent improvement in disentanglement. Results are displayed in the form of $\mu \pm \sigma$ in Table 3 and Table 4 respectively, where μ is the average of summed ratios and σ is the standard deviation across repetitions.

Even though the standard deviation increases, using the conditional VAE model on average improves the degree of disentanglement. As the variation in prosody is found to be linearly correlated with the

8. REFERENCES

[1] J. Sotelo, S. Mehri, K. Kumar, et al., "Char2wav: End-to-end speech synthesis.," in *Proc. Int. Conf. on Learning Representations (ICLR)*, 2017.

[2] Y. Wang, R. Skerry-Ryan, D. Stanton, et al., "Tacotron: Towards end-to-end speech synthesis.," in *Proc. Interspeech*, 2017, pp. 4006–4010.

[3] W. Ping, K. Peng, A. Gibiansky, et al., "Deep Voice 3: 2000-speaker neural text-to-speech.," in *Proc. Int. Conf. on Learning Representations (ICLR)*, 2018.

[4] J. Shen, R. Pang, R. J. Weiss, et al., "Natural TTS synthesis by conditioning WaveNet on mel spectrogram predictions.," in *Proc. ICASSP*, 2018, pp. 4779–4783.

[5] I. Sutskever, O. Vinyals, and Q. V. Le, "Sequence to sequence learning with neural networks," in *Advances in Neural Information Processing Systems*, 2014.

[6] W.-N. Hsu, Y. Zhang, R. J. Weiss, et al., "Hierarchical generative modeling for controllable speech synthesis," in *Proc. Int. Conf. on Learning Representations (ICLR)*, 2019.

[7] Y. Wang, D. Stanton, Y. Zhang, et al., "Style tokens: Unsupervised style modeling, control and transfer in end-to-end speech synthesis," in *Proc. Int. Conf. on Machine Learning (ICML)*, 2018, pp. 5167–5176.

[8] S. Ma, D. Mcduff, and Y. Song., "A generative adversarial network for style modeling in a text-to-speech system," in *Proc. Int. Conf. on Learning Representations (ICLR)*, 2019.

[9] M. Wagner and D. G. Watson, "Experimental and theoretical advances in prosody: A review," in *Language and Cognitive Processes*, 2010, pp. 905–945.

[10] W.-N. Hsu, Y. Zhang, R. J. Weiss, et al., "Disentangling correlated speaker and noise for speech synthesis via data augmentation and adversarial factorization," in *Proc. ICASSP*, 2019.

[11] E. Battenberg, S. Mariooryad, D. Stanton, et al., "Effective use of variational embedding capacity in expressive end-to-end speech synthesis," *arXiv: 1906.03402*, 2019.

[12] R. Skerry-Ryan, E. Battenberg, Y. Xiao, et al., "Towards end-to-end prosody transfer for expressive speech synthesis with Tacotron," in *Proc. Int. Conf. on Machine Learning (ICML)*, 2018.

[13] Y. Lee and T. Kim, "Robust and fine-grained prosody control of end-to-end speech synthesis," in *Proc. ICASSP*, 2019.

[14] J. Shen, R. Pang, R. J. Weiss, et al., "Natural TTS synthesis by conditioning wavenet on mel spectrogram predictions," in *Proc. ICASSP*, 2018.

[15] Y.-J. Zhang, S. Pan, L. He, and Z.-H. Ling, "Learning latent representations for style control and transfer in end-to-end speech synthesis," in *Proc. ICASSP*, 2019, pp. 6945–6949.

[16] G. E. Henter, J. Lorenzo-Trueba, X. Wang, and J. Yamagishi, "Deep encoder-decoder models for unsupervised learning of controllable speech synthesis," *arXiv:1807.11470*, 2018.

[17] K. Akuzawa, Y. Iwasawa, and Y. Matsuo, "Expressive speech synthesis via modeling expressions with variational autoencoder," in *Proc. Interspeech*, 2018, pp. 3067–3071.

[18] A. Razavi, A. van den Oord, and O. Vinyals, "Generating diverse high-fidelity images with VQ-VAE-2," *arXiv: 1904.02882*, 2019.

[19] Y.-A. Chung, Y. Wang, W.-N. Hsu, Y. Zhang, and R. Skerry-Ryan, "Semi-supervised training for improving data efficiency in end-to-end speech synthesis," in *Proc. ICASSP*, 2019.

[20] W.-N. Hsu, H. Tang, and J. Glass, "Unsupervised adaptation with interpretable disentangled representations for distant conversational speech recognition," in *Proc. Interspeech*, 2018.

[21] W.-N. Hsu, Y. Zhang, and J. Glass, "Unsupervised learning of disentangled and interpretable representations from sequential data.," in *Advances in Neural Information Processing Systems*, 2017.

[22] I. Higgins, L. Matthey, A. Pal, et al., "Beta-VAE: Learning basic visual concepts with a constrained variational framework," in *Proc. Int. Conf. on Learning Representations (ICLR)*, 2017.

[23] H. Kim and A. Mnih, "Disentangling by factorising," in *Proc. Int. Conf. on Machine Learning (ICML)*, 2018, pp. 2649–2658.

[24] A. Kumar, P. Sattigeri, and A. Balakrishnan, "Variational inference of disentangled latent concepts from unlabeled observations," in *Proc. Int. Conf. on Learning Representations (ICLR)*, 2017.

[25] F. Locatello, S. Bauer, M. Lucic, et al., "Challenging common assumptions in the unsupervised learning of disentangled representations," in *Proc. Int. Conf. on Machine Learning (ICML)*, 2019.

[26] S. Narayanaswamy, T. B. Paige, J.-W. van de Meent, et al., "Learning disentangled representations with semi-supervised deep generative models," in *Advances in Neural Information Processing Systems*, 2017, pp. 5925–5935.

[27] P. K. Gyawali, Z. Li, S. Ghimire, and L. Wang, "Semi-supervised learning by disentangling and self-ensembling over stochastic latent space," *arXiv: 1907.09607*, 2019.

[28] R. Habib, S. Mariooryad, M. Shannon, et al., "Semi-supervised generative modeling for controllable speech synthesis," *arXiv preprint arXiv:1910.01709*, 2019.

[29] X. Chen, Y. Duan, R. Houthooft, et al., "Infogan: Interpretable representation learning by information maximizing generative adversarial nets," in *Advances in Neural Information Processing Systems*, 2016, pp. 2172–2180.

[30] M. Mathieu, J. Zhao, P. Sprechmann, A. Ramesh, and Y. LeCun, "Disentangling factors of variation in deep representations using adversarial training," in *Advances in Neural Information Processing Systems*, 2016.

[31] B. Esmaeili, H. Wu, S. Jain, et al., "Structured disentangled representations," in *Proc. Int. Conf. on Artificial Intelligence and Statistics*, 2019, pp. 2525–2534.

[32] A. Vaswani, N. Shazeer, N. Parmar, et al., "Attention is all you need," in *Advances in Neural Information Processing Systems*, 2017.

[33] B. Uria, I. Murray, and H. Larochelle, "RNADE: The real-valued neural autoregressive density-estimator," in *Advances in Neural Information Processing Systems*, 2013.

[34] G. Papamakarios, T. Pavlakou, and I. Murray, "Masked autoregressive flow for density estimation," in *Advances in Neural Information Processing Systems*, 2017, pp. 2338–2347.

[35] H. Zen, V. Dang, R. Clark, et al., "LibriTTS: A corpus derived from LibriSpeech for text-to-speech," in *Proc. Interspeech*, 2019.

[36] S. King and V. Karaiskos, "The blizzard challenge 2013," in *Blizzard Challenge Workshop*, 2013.

[37] A. de Cheveigné and H. Kawahara, "YIN, a fundamental frequency estimator for speech and music," *Journal of the Acoustical Society of America*, vol. 111, no. 4, pp. 1917–1930, 2002.

[38] W. Chu and A. Alwan, "Reducing f0 frame error of f0 tracking algorithms under noisy conditions with an unvoiced/voiced classification frontend," in *Proc. ICASSP*, 2009.

[39] "Audio samples from "Fully-hierarchical Fine-grained Prosody Modelling for Interpretable Speech Synthesis"," https://google.github.io/tacotron/publications/hierarchical_prosody.