

Generative Modelling for Controllable Audio Synthesis of Expressive Piano Performance

Hao Hao Tan^{*1} Yin-Jyun Luo^{*1,2} Dorien Herremans^{1,2}

Abstract

We present a controllable neural audio synthesizer based on Gaussian Mixture Variational Autoencoders (GM-VAE), which can generate realistic piano performances in the audio domain that closely follows temporal conditions of two essential style features for piano performances: *articulation* and *dynamics*. We demonstrate how the model is able to apply fine-grained style morphing over the course of synthesizing the audio. This is based on conditions which are latent variables that can be sampled from the prior or inferred from other pieces. One of the envisioned use cases is to inspire creative and brand new interpretations for existing pieces of piano music.

1. Introduction

Synthesizing audio of piano performances from MIDI requires either a huge collection of recordings of individual notes, or a model that simulates an actual piano. These approaches, however, have two main limitations: (i) the “stitching” of individual notes might not optimally capture the various interactions between notes (Hawthorne et al., 2019); and (ii) the quality of the synthesized audio is restricted by the recordings in the sound library or the piano simulator. This further motivates the effort to build realistic piano synthesizers. In this work, we propose a neural network based synthesizer which takes the onset roll as input, and generates realistic piano performances in the audio domain. Our model takes the onset roll instead of the complete piano roll as input, thereby waiving the necessity of fine-grained frame and velocity information, which are often unavailable, to synthesize expressive piano performances. Without the constraint of frame and velocity information, performance style transfer could also be achieved by interpreting each onset note with the style inferred from a given audio performance. On top of the onset roll, the generation

^{*}Equal contribution ¹Singapore University of Technology and Design ²Institute of High Performance Computing, A*STAR, Singapore. Correspondence to: Hao Hao Tan <hao-hao.tan@sutd.edu.sg>.

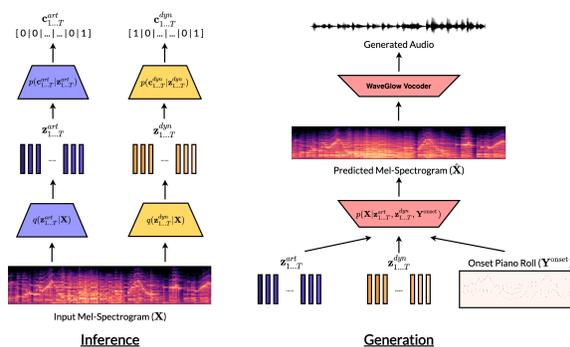


Figure 1. Model architecture.

is further conditioned on variables corresponding to style features, which enable the rendition of expressive piano performances. In particular, we consider *articulation* and *dynamics*, which are two significant features related to expressiveness in piano performances, and base our framework on variational autoencoders (VAEs) (Kingma & Welling, 2014) with a Gaussian mixture prior (Jiang et al., 2016).

The underlying task of this work is to build a neural audio synthesizer that maps piano performances in MIDI to audio. The current literature either focuses on such direct mapping without additional controllability with regard to performance styles (Hawthorne et al., 2019; Thakkar Vijay Manzelli, 2018), or concerns only the global attributes such as synthesizing different instruments (Kim et al., 2019). We distinguish ourselves from the previous works by incorporating GM-VAE (Jiang et al., 2016), which has been applied to speech (Hsu et al., 2019) and instrument modelling (Luo et al., 2019), to disentangle two significant expressive performance factors. This allows us to achieve creative applications such as gradual style morphing over time. The sequence of conditions can be sampled from the prior or inferred from other pieces, as shown in Section 3.

2. Experimental Details

Data Representation: We use the MAESTRO v2.0.0 dataset (Hawthorne et al., 2019) which consists of 1,282 performances with aligned audio-MIDI pairs, and we split the data as annotated. We train on 20-second random crops from audio clips, which are converted into a log-scale Mel-spectrogram with 80 Mel-filters (X). The MIDI note se-

quences are represented as piano rolls, from which the onset roll ($\mathbf{Y}^{\text{onset}}$) is extracted. Both articulation (*staccato*, *legato*) and dynamics (loud, soft) are represented as binary sequences $\mathbf{c}_{1..T}^{\text{art}}, \mathbf{c}_{1..T}^{\text{dyn}} \in \{0, 1\}$. $c_t^{\text{art}} = 1$ if at least one note is held at t ; and $c_t^{\text{dyn}} = 1$ if the average velocity across all notes is greater than 70 at t , which the threshold is determined by our preliminary data analysis.

Model Formulation: Figure 1 shows our model architecture, which is adapted from (Hsu et al., 2019). The joint distribution $p_\theta(\mathbf{X}, \mathbf{z}^{\text{art}}, \mathbf{z}^{\text{dyn}} | \mathbf{Y}^{\text{onset}}, \mathbf{c}^{\text{art}}, \mathbf{c}^{\text{dyn}}) = p_\theta(\mathbf{X} | \mathbf{Y}^{\text{onset}}, \mathbf{z}^{\text{art}}, \mathbf{z}^{\text{dyn}}) p_\theta(\mathbf{z}^{\text{art}} | \mathbf{c}^{\text{art}}) p_\theta(\mathbf{z}^{\text{dyn}} | \mathbf{c}^{\text{dyn}})$, where θ refers to parameters of the generation network. Both the prior distributions $p_\theta(\mathbf{z}^{\text{art}})$ and $p_\theta(\mathbf{z}^{\text{dyn}})$ are Gaussian mixtures of two components. A variational distribution $q_\phi(\mathbf{z} | \mathbf{X})$ is introduced to approximate the true posterior, where ϕ refers to parameters of the inference network. The model is trained to optimize the evidence lower bound:

$$\begin{aligned} \mathcal{L}(\theta, \phi; \mathbf{X}) = & \mathbb{E}_{q_\phi(\mathbf{z}^{\text{art}} | \mathbf{X}) q_\phi(\mathbf{z}^{\text{dyn}} | \mathbf{X})} [\log p_\theta(\mathbf{X} | \mathbf{Y}^{\text{onset}}, \mathbf{z}^{\text{art}}, \mathbf{z}^{\text{dyn}})] \\ & - \mathcal{D}_{\text{KL}}(q_\phi(\mathbf{z}^{\text{art}} | \mathbf{X}) \| p(\mathbf{z}^{\text{art}} | \mathbf{c}^{\text{art}})) \\ & - \mathcal{D}_{\text{KL}}(q_\phi(\mathbf{z}^{\text{dyn}} | \mathbf{X}) \| p(\mathbf{z}^{\text{dyn}} | \mathbf{c}^{\text{dyn}})). \end{aligned} \quad (1)$$

Both the generation and inference network are implemented with two-layer bidirectional LSTMs.¹ Note that we simplify $\mathbf{c}_{1..T}^{\text{art}}$ and $\mathbf{c}_{1..T}^{\text{dyn}}$ as \mathbf{c}^{art} and \mathbf{c}^{dyn} (similarly for the latent variables \mathbf{z}^{art} and \mathbf{z}^{dyn}). For each of \mathbf{c}^{art} and \mathbf{c}^{dyn} , an additional cross-entropy loss is introduced such that the posterior $p(\mathbf{c} | \mathbf{z})$ can also learn from the labelled ground-truth \mathbf{c} .

Audio Synthesis: We leverage WaveGlow (Prenger et al., 2019) to invert the Mel-spectrogram to audio, due to its fast inference and superior performance (Govalkar et al.; Zhao et al., 2020). We adopt the implementation from (Yu, 2019).

3. Results and Discussion

Gradual Style Morphing Over Time: Utilizing the Gaussian mixture prior distribution enables style morphing by linear interpolation between mixture components. In particular, given μ_0^{art} and μ_1^{art} representing the mean vectors of mixture components corresponding to *staccato* and *legato*, we can set $\mathbf{z}_t^{\text{art}} = \mu_0^{\text{art}} + (\mu_1^{\text{art}} - \mu_0^{\text{art}}) \times \frac{t}{T}$ in the latent space of articulation (similarly for $\mathbf{z}_t^{\text{dyn}}$). In other words, the sequence of conditioning vectors travel from μ_0^{art} to μ_1^{art} , whereby we can expect the articulation of the synthesized piano performance to morph gradually from *staccato* to *legato*. Figure 2 demonstrates the generated Mel-spectrograms of four different scenarios. From *staccato* to *legato*, one can observe that notes are gradually sustained longer; and from *soft* to *loud*, it is clear that the amplitude increases over time, and the frequency coverage expands. This matches with the human auditory perception that mid-high frequencies are easier

¹More implementation details at: (source code url suppressed)

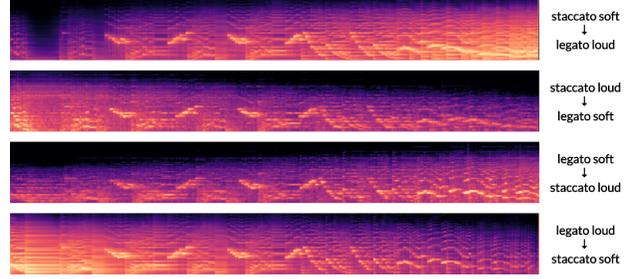


Figure 2. Generated Mel-spectrograms of a given piece, with combinations of gradual morphing of articulation and dynamics.

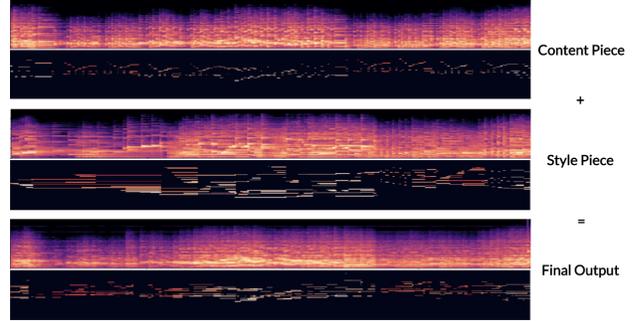


Figure 3. An example of performance style transfer. The piano roll for the final output is estimated using the state-of-the-art model for piano transcription (Hawthorne et al., 2018).

to be perceived as “loud” as compared to low frequencies under the same amplitude (Fletcher & Munson, 1933).

Performance Style Transfer: In addition to sampling from the prior distribution, we can also infer the sequence of conditioning vectors from another piece of music. Specifically, we let $\mathbf{z}^{\text{art}} \sim q_\phi(\mathbf{z}^{\text{art}} | \mathbf{X}^{\text{style}})$ (similarly for \mathbf{z}^{dyn} , where $\mathbf{X}^{\text{style}}$ is the *style piece* that would determine the fine-grained style over time of the synthesized piano performance. Figure 3 shows an example that renders a piece from the Baroque era (which is more detached and constant in terms of dynamics) in the style of a piece from the Romantic era (which is more *legato* and expressive in terms of dynamics). By observing both the Mel-spectrograms and piano rolls, one can see that the final output closely follows the style features of the style piece in terms of note duration (articulation) and amplitude (dynamics), while preserving the musical content. Audio examples can be found online.²

We envision that this framework could learn to achieve fine-grained control on multiple performance style factors, which allows us to explore new performance directions for any given piece, even by taking inspirations from other pieces via style transfer. Future work will involve extending the set of performance features (e.g. onset deviation, pedalling) in order to generate more realistic piano performance.

²<https://piano-audio-synthesis.github.io>

References

- Fletcher, H. and Munson, W. A. Loudness, its definition, measurement and calculation. *Bell System Technical Journal*, 12(4):377–430, 1933.
- Govalkar, P., Fischer, J., Zalkow, F., and Dittmar, C. A comparison of recent neural vocoders for speech signal reconstruction.
- Hawthorne, C., Elsen, E., Song, J., Roberts, A., Simon, I., Raffel, C., Engel, J., Oore, S., and Eck, D. Onsets and frames: Dual-objective piano transcription. In *18th International Society for Music Information Retrieval Conference, ISMIR*, 2018.
- Hawthorne, C., Stasyuk, A., Roberts, A., Simon, I., Huang, C.-Z. A., Dieleman, S., Elsen, E., Engel, J., and Eck, D. Enabling factorized piano music modeling and generation with the maestro dataset. In *International Conference of Learning Representations, ICLR*, 2019.
- Hsu, W.-N., Zhang, Y., Weiss, R. J., Zen, H., Wu, Y., Wang, Y., Cao, Y., Jia, Y., Chen, Z., Shen, J., et al. Hierarchical generative modeling for controllable speech synthesis. In *International Conference of Learning Representations, ICLR*, 2019.
- Jiang, Z., Zheng, Y., Tan, H., Tang, B., and Zhou, H. Variational deep embedding: An unsupervised and generative approach to clustering. *arXiv preprint arXiv:1611.05148*, 2016.
- Kim, J. W., Bittner, R., Kumar, A., and Bello, J. P. Neural music synthesis for flexible timbre control. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 176–180. IEEE, 2019.
- Kingma, D. P. and Welling, M. Auto-encoding variational bayes. In *International Conference of Learning Representations, ICLR*, 2014.
- Luo, Y.-J., Agres, K., and Herremans, D. Learning disentangled representations of timbre and pitch for musical instrument sounds using gaussian mixture variational autoencoders. In *International Society for Music Information Retrieval Conference, ISMIR*, 2019.
- Prenger, R., Valle, R., and Catanzaro, B. Waveglow: A flow-based generative network for speech synthesis. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 3617–3621. IEEE, 2019.
- Thakkar Vijay Manzelli, R. Combining deep generative raw audio models for structured automatic music. In *19th International Society for Music Information Retrieval Conference, ISMIR*, 2018.
- Yu, C. Y. Constant memory waveglow: A pytorch implementation of waveglow with constant memory cost. <https://github.com/yoyololicon/constant-memory-waveglow>, 2019.
- Zhao, Y., Wang, X., Juvela, L., and Yamagishi, J. Transferring neural speech waveform synthesizers to musical instrument sounds generation. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 6269–6273. IEEE, 2020.