

Spectrogram Inpainting for Interactive Generation of Instrument Sounds

Théis Bazin^{1,2*}, Gaëtan Hadjeres¹, Philippe Esling², and Mikhail Malt²

¹ Sony CSL Paris

² IRCAM - CNRS UMR 9912 STMS - Sorbonne Université

Abstract. Modern approaches to sound synthesis using deep neural networks are hard to control, especially when fine-grained conditioning information is not available, hindering their adoption by musicians.

In this paper, we cast the generation of individual instrumental notes as an inpainting-based task, introducing novel and unique ways to iteratively shape sounds. To this end, we propose a two-step approach: first, we adapt the VQ-VAE-2 image generation architecture to spectrograms in order to convert real-valued spectrograms into compact discrete codemaps, we then implement token-masked Transformers for the inpainting-based generation of these codemaps.

We apply the proposed architecture on the NSynth dataset on masked resampling tasks. Most crucially, we open-source an interactive web interface to transform sounds by inpainting, for artists and practitioners alike, opening up to new, creative uses.

Keywords: VQ-VAE-2 · Transformers · inpainting · Vector-Quantization · interactive interfaces · sound synthesis · NSynth · web interfaces

1 Introduction

Generative modeling of audio has seen a surge in fidelity in recent years, thanks to advances in generative neural network-based architectures (van den Oord et al., 2016, 2018; J. H. Engel et al., 2019; Aouameur, Esling, & Hadjeres, 2019; Esling, Masuda, Bardet, Despres, & Chemla-Romeu-Santos, 2019). These new models are progressively bridging the gap with dedicated synthesis software in terms of the versatility of sounds produced without requiring the involved domain-specific knowledge of audio synthesis. Yet these approaches still fall short of providing convenient, interpretable control over the generated sounds, only giving access either to *global* parameters that change the texture of the whole sounds or to *dense*, high-frequency conditioning, impractical for interactive uses.

Global conditioning labels such as pitch and instrument family set aside, introducing *local* control in such models is made difficult by the typically high sampling frequency of sounds in waveform representation. Indeed, this high sampling frequency makes autoregressive modeling of sounds highly challenging. This is

* This author was supported by the ANRT via the CIFRE contract number 2019.0094.

the case for models of the WAVENET family (van den Oord et al., 2016, 2018), and such models are bound to operate with low receptive fields and display highly unstationary behaviour. Controlling these models therefore typically requires very dense conditioning sequences, sampled at a rate close to the audio rate, which is impractical for end-users. Conversely, models such as GANSYNTH (J. H. Engel et al., 2019), that use spectrogram representations and generate full-scale sounds of several seconds in parallel through transposed convolutions conditioned on a single vector, lose the ability to control the small-scale structure of the sounds.

Promising solutions have been proposed to introduce control at a mid-level range in autoregressive models through the use of Style Tokens and similar techniques (J. Engel et al., 2017; Wang et al., 2018). Here, an intermediary, subsampled latent space is learned along with the autoregressive decoder and used to condition the decoder. Nevertheless, the resulting conditioning vectors are real-valued vectors with unidentified dimensions, making them hard to interpret. Existing control schemes are therefore limited to extracting the sequence of conditioning vectors of an example sound and transferring it to another sound, but further control is difficult.

We turn to recent advances in *inpainting*-based interactive generation, which have brought new approaches to controllable sampling for creative applications. Inpainting models are prediction models trained to reconstruct hidden or missing data in an input. These models were previously proposed for the interactive generation of symbolic music (Hadjeres, Pachet, & Nielsen, 2017; Hadjeres & Nielsen, 2018) and are furthermore readily amenable to interactive applications (Bazin & Hadjeres, 2019).

In this work, we propose to tackle the problem of controlled sound generation through a two-steps approach. First, an encoder-decoder network is trained to compute a highly compressed representation of the sounds through down-sampling and discretization, with efficient encoding as well as decoding. This mid-level representation, used as input for our subsequent generative models, bridges the gap between the dense conditioning sequences of WAVENET and the very sparse, label-based conditioning of GANSYNTH. We then train powerful Transformer-based sequence generation models on these discrete codemaps for interactive generation by inpainting.

2 Contributions

We adapt the VQ-VAE-2 image generation model to allow for the interactive generation of sounds via inpainting. Our key contributions are:

1. We cast the generation of musical instrument sounds as an interactive, inpainting-based generative-modeling task,
2. We adapt the hierarchical VQ-VAE-2 model to operate on spectrograms, using a custom modification of the IF-Mel spectrogram-based representation originally proposed for the GANSYNTH model,
3. In order to sample new sounds in this compact latent space, we model a factorization of the joint probability of hierarchical codemaps from the

VQ-VAE-2. To this end we introduce hierarchically conditioned sequence-masked Spectrogram Transformers, compatible with interactive generation by inpainting. These models furthermore support conditioning through global labels for additional control,

4. We open-source the code for this model and release an open-source interactive web interface for the proposed model³, which allows to edit sounds via inpainting whilst viewing and listening to the results, enabling novel and unique schemes of sound transformation.

3 Proposed approach

In order to provide control on the generated sounds, a sweet spot must be found between full autoregressive sampling of full-scale spectrograms, which is computationally demanding when using powerful Transformer architectures, and one-shot parallel generation *à-la*-GANSYNTH, which discards local control in favor of computational efficiency. We therefore follow an autoregressive approach but introduce models with highly downsampled and compact latent representations, so that these latent sequences can in turn be conveniently generated and edited using autoregressive models.

3.1 Compact Representations of Sounds

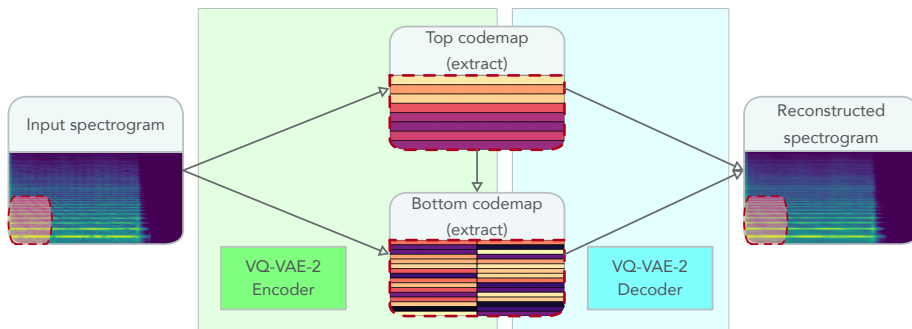


Fig. 1. Proposed VQ-VAE-2 architecture for spectrograms. The encoder converts 1024×128 -large spectrograms into two compact integer maps, c_{top} and c_{bottom} of size 32×4 and 64×8 . The decoder is trained to reconstruct the original spectrogram from these discrete codemaps.

The VQ-VAE-2 model, initially proposed for the generation of images (Razavi, van den Oord, & Vinyals, 2019), defines an autoencoding architecture using a hierarchy of discrete latent variables. In the VQ-VAE-2 model, the VAE’s latent

³ Code for the models and interface available at: <https://github.com/SonyCSLParis/interactive-spectrogram-inpainting>

space is made discrete through *vector-quantization*, a classic dictionary-learning algorithm (Gray, 1984): the latent representation computed by the VQ-VAE-2 encoder is projected onto a learned *codebook* (a dictionary of code vectors), effectively replacing the input data by a discrete grid of codebook indexes. The decoder in turn receives only the projection and reconstructs the original input by upsampling (via transposed convolutions). In the hierarchical approach, a top codemap is obtained from the bottom map by further downsampling and vector-quantization. In this setting, the highly downsampled top codes can be interpreted as depicting the overall, large-scale structure of the images. The bottom codemap decoder is subsequently conditioned on this global structure and refines it with fine-grained detail. This is shown on Figure 1 via the arrow depicting the conditional dependency of the bottom codemap on the top one.

We note that VQ has previously been applied to the generation of audio signals with impressive results (Dieleman, van den Oord, & Simonyan, 2018; van den Oord, Vinyals, & Kavukcuoglu, 2017), most notably in the recent Jukebox model by Open AI (Dhariwal et al., 2020). Nonetheless, VQ is used in these approaches as a powerful information bottleneck allowing to scale autoregressive models of audio to very large temporal scales of several tens of seconds. Controlling the huge models proposed in these papers remains highly challenging, if only due to the very high computational load they incur, making these approaches somewhat orthogonal to our goal of intuitive and interactive control.

3.2 Spectrogram Representations for Intuitive Control

Since the VQ-VAE-2 model was initially introduced for the generation of images, we borrow the invertible Hi-Res Mel-IF (Mel-Amplitude and Instantaneous Frequency) spectrogram representation that was shown by J. H. Engel et al. (2019) to provide the best results for generation with convolutional neural networks. This representation, which performs a temporal unrolling of the phase channel of the spectrograms, introduces smoothness and stability as opposed to directly learning to model the phase. This representation can furthermore be inverted back to the original spectrogram at no cost and therefore allows for efficient conversion back to audio using the inverse Fourier transform.

The auto-encoding training of the VQ-VAE-2 furthermore requires introducing a reconstruction loss, as opposed to the adversarial training of GANSYNTH. We propose to use a slightly modified $L2$ -loss, using a phase thresholding process: we ignore IF reconstruction errors (by setting those gradients to zero) in areas of low sound amplitude, where the IF starts becoming highly noisy (J. H. Engel et al., 2019), effectively shifting gradient signals to areas with significant amplitude.

Working on spectrograms has the additional advantage of separating time and frequencies onto two separate axis in the input representation. This provides an intuitive interpretation of the latent codes learned by the hierarchical VQ-VAE-2 as time-frequency atoms. Editing the resulting 2D codemaps by inpainting can then be seen as performing localized transformations in the time-frequency plane over multiple frequency bands, akin to the multi-band equalizers musicians are familiar with.

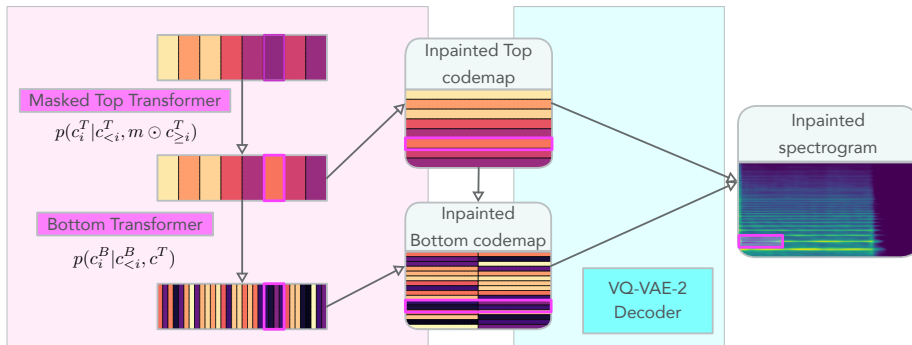


Fig. 2. Diagram of the proposed spectrogram inpainting process. After a linearization-procedure informed by harmonic structure, sequence-masked Transformers are trained to perform inpainting on the codemaps of the VQ-VAE-2, which allows to sample new sounds by first autoregressively sampling from the factorized distribution $p(c_{top})p(c_{bottom}|c_{top})$ then decoding these sequences.

3.3 Spectrogram Transformers

After training the VQ-VAE, the continuous-valued spectrograms can be replaced with the hierarchical latent codemaps. Generating new sounds then amounts to modeling and sampling from the joint probability $p(c_{top}, c_{bottom})$. We follow the approach of (Razavi et al., 2019) and factorize this probability into $p(c_{top})p(c_{bottom}|c_{top})$, which is also in line with the conditional structure of the VQ-VAE-2 encoder. We therefore train an autoregressive model for the top codes and another one – conditioned on the top codes – for the bottom codes.

To this end we adapt the VQ-CPC architecture (Hadjeres & Crestel, 2020) for hierarchical discrete latent variables. For the top codemap $c_{i \dots N}^T$, we use a causally masked, self-attentive Transformer, therefore modeling the probability $p(c_i^T | c_{<i}^T, m \odot c_{\geq i}^T)$, where m denotes the randomly sampled *inpainting mask*, used to stochastically hide data during training. Typically, a simple Bernoulli can be used here for each token. This mask then allows, at inference time, to select areas to regenerate, and provide the model with information about the fixed, future tokens, ensuring coherent generation. For the upsampled, bottom codemaps we use another Transformer which can attend via attention to the whole top codemap in a patch-based fashion: indeed, to help the attention procedure, we align patches of the bottom map to the top code which they were upsampled from. We obtain the following formulation: $p(c^B | c^T) = \prod_i p(c_i^B | c_{<i}^B, c^T)$. The resulting architecture is depicted on Figure 2.

4 Experiments

The proposed hierarchical model allows to perform various generative modeling tasks, we provide some usage examples here. We refer the reader to the compan-

ion website⁴ where many uncurated audio samples are provided for each task described here, using all of the instrument types in the NSynth dataset and across a large range of pitches, along with a more thorough discussion of the performances of the model. Note that in this paper, we focus on the general framework of interactive sound inpainting and its applications and we reserve a quantitative evaluation of architectural choices for future work.

4.1 Training

We train the proposed architecture on the NSynth dataset (J. Engel et al., 2017). As is done in the GANSynth paper (J. H. Engel et al., 2019), we restrict the dataset to sounds with pitch from MIDI 24-84, avoiding extreme pitches. We compute the input spectrograms using windows of 2048 samples with $N_{fft} = 2048$ with a 75% overlap factor, resulting in an effective dimension of 1024×128 for the spectrograms. As opposed to the representation used in GANSynth, we lower the Mel-scale break frequency from 700Hz to 240Hz, in order to increase resolution in the low frequencies, bringing more precision in the interactive editing of the lower harmonics.

The VQ-VAE is trained for reconstruction of the Mel-IF representation using the phase-masked L2 loss described above with the ADAM optimizer. After training the VQ-VAE, we use its encoder to extract the latent codemaps for all of the sounds in the training dataset. The bottom and top Transformers can then be independently trained on these codemaps. Here, we train in an autoregressive fashion with the RAdam optimizer (Liu et al., 2020) and a label-smoothed prediction loss (M  ller, Kornblith, & Hinton, 2019). More details are given in the repository accompanying the paper.

4.2 Results

Inpainting operations On Figure 3 we display uncurated samples for three possible uses of the model. First, a spectrogram for a 4 second organ sound at pitch 60 generated by our model is shown. Three successive inpainting operations are then presented. The first two examples operate on both codemaps, by first autoregressively resampling the displayed portion of the top codemap, then resampling the aligned tokens in the bottom codemap. The third one only regenerates the bottom codemap, keeping the top one fixed. In the first example, the five lowest frequency bands of the top map are regenerated over the whole duration of the sound with constraint "**Bass + pitch C1**", resulting in a visible increase in lower-frequency energy. Then, the first second of the spectrogram is regenerated over all frequency bands with constraint "**Mallet + pitch C3**", resulting in a more marked attack. In the last step, the top codemap is kept fixed but the full bottom codemap is resampled conditioned on it, with a constraint "**Synth Lead + pitch C5**". We can see that the global structure of the spectrogram remains similar, but the local timbral texture indeed changes. We again refer the reader to the accompanying webpage for audio examples.

⁴ <https://sonycslparis.github.io/interactive-spectrogram-inpainting/>




Fig. 3. Successive inpainting operations over a generated sample. (1) Generated sample ; (2) inpainting of the lower frequencies over the whole duration of the sound ; (3) inpainting of the full first second ; (4) inpainting of the full bottom codemap with the top codemap fixed.

Interactive interface We believe that generative models for music only reveal their potential through actual usage in music production contexts, which requires providing musicians with intuitive interfaces to these complex models. In line with this philosophy, we propose a TypeScript-based web interface that allows to generate and edit sounds using the models presented in this paper⁵, using a PyTorch-based computing back-end. This interface is represented on Figure 4.

The *modus operandi* is a simple one, akin to classic image editing software, which we can expect our users to be familiar with. An initial sound is either generated from scratch using the autoregressive models, sampled from a dataset or provided by the user through drag and drop. The VQ-VAE-reconstructed spectrogram is displayed on the interface, along with a grid overlay depicting the cells of the VQ-VAE’s latent codemaps. The user can then perform inpainting operations as described in 4.2 by selecting zones on this grid, either on the coarse top codemap or on the more detailed bottom map.

We note that the possibility of using different instrument and pitch constraints for successive inpainting operations opens the way to the generation of heterogeneous spectrograms, unseen in the training dataset – such as inpainting the attack of an organ sound with a "Guitar" instrument constraint to obtain a plucked organ. The successive inpainting operations performed in Figure 3 are an example of this use case. We finally observe that, due to the inherently linearly-scaling nature of the autoregressive generation, each token-wise predic-

⁵ Available within the code’s repository at <https://github.com/SonyCSLParis/interactive-spectrogram-inpainting>




Fig. 4. Screenshot of the proposed interactive web interface. The highlighted zone is currently being regenerated by the back-end.

tion of the spectrogram transformers should be fast if responsive interaction is desired. Currently, typical local operations (over a region of one second) take around one second to complete, and work is undergoing to reduce this latency for a seamless interactive process.

5 Conclusion and perspectives

We have presented an approach to interactively generating instrument sounds by inpainting. To this end, we adapted the VQ-VAE-2 image generation model to spectrograms via a representation inspired by the Mel-IF representation from GANSYNTH. This allows to encode sounds into compact discrete sequences, greatly compressing the informational complexity of the data. We then introduced efficient hierarchical Transformers with sequence-masking, in order to model a factorization of the joint probability of the top and bottom VQ-VAE-2 codemaps. These models are trained to perform autoregressive prediction, but incorporate information both from the past and the future of the sequences, so that they can be used to perform inpainting on the codemaps at inference time. We introduced a web interface that allows to use the trained models to generate and edit sounds, which we distribute in an open-source fashion, for researchers and musicians.

Future work includes speeding up and scaling the model for better reconstruction and faster generation, via, for instance, recent advancements to Transformer-like architectures (Katharopoulos, Vyas, Pappas, & Fleuret, 2020), along with a more in-depth, quantitative assessment of the performances of the proposed framework and architectures.

References

- Aouameur, C., Esling, P., & Hadjeres, G. (2019, June). Neural Drum Machine : An Interactive System for Real-time Synthesis of Drum Sounds. In *International Conference on Computational Creativity*. Charlotte, United States. Retrieved from <https://hal.archives-ouvertes.fr/hal-02479211> (8 pages, accepted at the International Conference on Computational Creativity 2019)
- Bazin, T., & Hadjeres, G. (2019). NONOTO: A Model-agnostic Web Interface for Interactive Music Composition by Inpainting. In A. for Computational Creativity (Ed.), *Proceedings of the 10th International Conference on Computational Creativity, ICCO 2019, Charlotte, North Carolina, USA, June 17-21* (Vol. abs/1907.10380). Retrieved from <http://arxiv.org/abs/1907.10380>
- Dhariwal, P., Jun, H., Payne, C., Kim, J. W., Radford, A., & Sutskever, I. (2020). Jukebox: A Generative Model for Music.. Retrieved from <http://arxiv.org/abs/2005.00341>
- Dieleman, S., van den Oord, A., & Simonyan, K. (2018). The Challenge of Realistic Music Generation: Modelling Raw Audio at Scale. In S. Bengio, H. M. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, & R. Garnett (Eds.), *Advances in Neural Information Processing Systems 31, NeurIPS 2018, 3-8 December 2018, Montréal, Canada* (pp. 8000–8010). Retrieved from <http://papers.nips.cc/paper/8023-the-challenge-of-realistic-music-generation-modelling-raw-audio-at-scale>
- Engel, J., Resnick, C., Roberts, A., Dieleman, S., Norouzi, M., Eck, D., & Simonyan, K. (2017). Neural Audio Synthesis of Musical Notes with WaveNet Autoencoders. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70* (p. 1068–1077). JMLR.org.
- Engel, J. H., Agrawal, K. K., Chen, S., Gulrajani, I., Donahue, C., & Roberts, A. (2019). GANSynth: Adversarial Neural Audio Synthesis. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net. Retrieved from <https://openreview.net/forum?id=H1xQVn09FX>
- Esling, P., Masuda, N., Bardet, A., Despres, R., & Chemla-Romeu-Santos, A. (2019). Universal Audio Synthesizer Control with Normalizing Flows. *CoRR*, abs/1907.00971. Retrieved from <http://arxiv.org/abs/1907.00971>
- Gray, R. (1984). Vector Quantization. *IEEE ASSP Magazine*, 1(2), 4-29.
- Hadjeres, G., & Crestel, L. (2020). Vector-Quantized Contrastive Predictive Coding for Template-based Music Generation.. Retrieved from <https://arxiv.org/abs/2004.10120>
- Hadjeres, G., & Nielsen, F. (2018). Anticipation-RNN: Enforcing Unary Constraints in Sequence Generation, with Application to Interactive Music Generation. *Neural Computing and Applications*, 32(4), 995–1005. Retrieved from <https://doi.org/10.1007/s00521-018-3868-4> doi: 10.1007/s00521-018-3868-4

- Hadjeres, G., Pachet, F., & Nielsen, F. (2017). DeepBach: a Steerable Model for Bach Chorales Generation. In D. Precup & Y. W. Teh (Eds.), *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017* (Vol. 70, pp. 1362–1371). PMLR. Retrieved from <http://proceedings.mlr.press/v70/hadjeres17a.html>
- Katharopoulos, A., Vyas, A., Pappas, N., & Fleuret, F. (2020). Transformers are RNNs: Fast Autoregressive Transformers with Linear Attention. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020*. Retrieved from <https://arxiv.org/pdf/2006.16236.pdf>
- Liu, L., Jiang, H., He, P., Chen, W., Liu, X., Gao, J., & Han, J. (2020, April). On the Variance of the Adaptive Learning Rate and Beyond. In *Proceedings of the Eighth International Conference on Learning Representations (ICLR 2020)*.
- M  ller, R., Kornblith, S., & Hinton, G. E. (2019). When Does Label Smoothing Help? In H. M. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alch  -Buc, E. B. Fox, & R. Garnett (Eds.), *Advances in Neural Information Processing Systems 32: NeurIPS 2019, 8-14 December 2019, Vancouver, BC, Canada* (pp. 4696–4705). Retrieved from <http://papers.nips.cc/paper/8717-when-does-label-smoothing-help>
- Razavi, A., van den Oord, A., & Vinyals, O. (2019). Generating Diverse High-fidelity Images with VQ-VAE-2. In H. M. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alch  -Buc, E. B. Fox, & R. Garnett (Eds.), *Advances in Neural Information Processing Systems 32, NeurIPS 2019, 8-14 December 2019, Vancouver, BC, Canada* (pp. 14837–14847). Retrieved from <http://papers.nips.cc/paper/9625-generating-diverse-high-fidelity-images-with-vq-vae-2>
- van den Oord, A., Dieleman, S., Zen, H., Simonyan, K., Vinyals, O., Graves, A., ... Kavukcuoglu, K. (2016). WaveNet: A Generative Model for Raw Audio. In *Arxiv*. Retrieved from <https://arxiv.org/abs/1609.03499>
- van den Oord, A., Li, Y., Babuschkin, I., Simonyan, K., Vinyals, O., Kavukcuoglu, K., ... Hassabis, D. (2018, 10–15 Jul). Parallel WaveNet: Fast High-fidelity Speech Synthesis. In J. Dy & A. Krause (Eds.), *Proceedings of the 35th International Conference on Machine Learning* (Vol. 80, pp. 3918–3926). Stockholm, Sweden: PMLR. Retrieved from <http://proceedings.mlr.press/v80/oord18a.html>
- van den Oord, A., Vinyals, O., & Kavukcuoglu, K. (2017). Neural Discrete Representation Learning. In I. Guyon et al. (Eds.), *Advances in Neural Information Processing Systems 30* (pp. 6306–6315). Curran Associates, Inc. Retrieved from <http://papers.nips.cc/paper/7210-neural-discrete-representation-learning.pdf>
- Wang, Y., Stanton, D., Zhang, Y., Skerry-Ryan, R., Battenberg, E., Shor, J., ... Saurous, R. A. (2018). Style Tokens: Unsupervised Style Modeling, Control and Transfer in End-to-End Speech Synthesis. Retrieved from <https://arxiv.org/abs/1803.09017>