
Incorporating Music Knowledge in Continual Dataset Augmentation for Music Generation

Alisa Liu¹ Alexander Fang^{1,2} Gaëtan Hadjeres³ Prem Seetharaman¹ Bryan Pardo¹

1. Introduction

Deep learning has rapidly become the state-of-the-art approach for music generation (Briot et al., 2017). However, training a deep model typically requires a large training set, and it is well known that the performance of deep learning systems scales up with more data, even when the data is noisy. However, when training a model to emulate the style of a particular composer, the size of the dataset is inherently limited to the number of compositions by that musician.

In this paper, we present **augmentative generation** (AUG-GEN), a method of dataset augmentation for any music generation system trained on a resource-constrained domain. The key intuition of this method is that the training data for a generative system can be augmented by examples the system produces during the course of training, provided these examples are of sufficiently high quality and variety. To the best of our knowledge, our paper is the first to introduce a framework in which a generative system continuously generates output and adds it to its training dataset.

We apply AUG-GEN to chorale generation in the style of J.S. Bach. We perform experiments using a Transformer model (Vaswani et al., 2017) as the music generation system. To select model output to include in the training data, we use the grading function in (Fang et al., 2020), which evaluates a given chorale along nine musical features important for Bach chorales. This allows us to incorporate non-differentiable domain knowledge (e.g. 18th century counterpoint rules) into the training procedure.

There has been recent work to use generative models to synthesize training examples for image classification (Kong et al., 2019) and text classification (Anaby-Tavor et al.,

2020). However, these works aim to train a classifier rather than a generative system like ours. Moreover, the informativeness of candidate training examples is measured relative to the model, whereas in our work, selection of training examples is done by an external critic that does not depend on the model’s current performance.

2. Augmentative Generation

Algorithm 1 AUG-GEN training algorithm

```
1: input: true training dataset
2: initialize epoch  $i = 0$ 
3: while stopping criterion is false do
4:   for  $j = 1, \dots, N$  do
5:     generate chorale  $c_{ij}$ 
6:     grade  $c_{ij}$  to obtain  $g(c_{ij})$ 
7:     if  $g(c_{ij}) \leq t$  and  $c_{ij}$  is sufficiently diverse then
8:       add  $c_{ij}$  to the train dataset
9:     end if
10:  end for
11:  train on  $m$  batches of size  $k$  selected from the dataset
12:   $i \leftarrow i + 1$ 
13: end while
```

In AUG-GEN, the true dataset is first split into training and validation data. During model training, the training dataset is continuously augmented by model output, while the validation data is fixed and is used to determine when to terminate training. Each epoch of training includes a **generation step** and a **training step**. In the generation step, N examples are generated and graded by the grading function g . We add all generated output that (1) passes a pre-determined quality threshold t , and (2) is sufficiently diverse, to the training dataset. In this work, we use a simple uniqueness criterion (i.e. the chorale must not have been seen before) as the diversity requirement. Next, in the training step of the epoch, we train the model on the augmented dataset. This consists of training on m batches of size k , so that the amount of training in each epoch is independent of the size of the entire training dataset. These steps continue until the validation loss ceases to improve. In Algorithm 1, the AUG-GEN algorithm is shown in pseudocode. Note that the grading function used returns a measure of distance from

¹Department of Computer Science ²Bienen School of Music, Northwestern University, Evanston, IL, USA ³Sony CSL, Paris. Correspondence to: Alisa Liu <alisa@u.northwestern.edu>, Alexander Fang <alexanderfang2019@u.northwestern.edu>, Gaëtan Hadjeres <gaetan.hadjeres@etu.upmc.fr>, Prem Seetharaman <prem@u.northwestern.edu>, Bryan Pardo <pardo@northwestern.edu>.

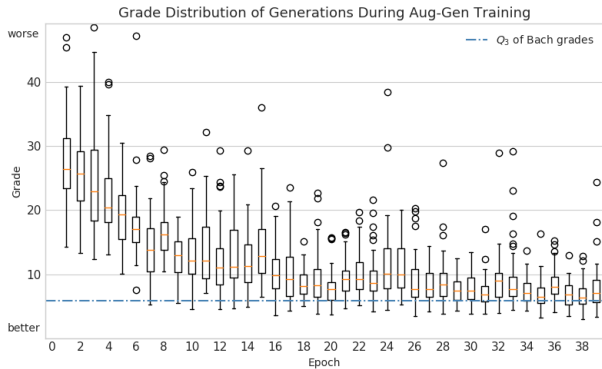


Figure 1. A series of boxplots representing the grade distribution of 50 chorales generated at each epoch of training in AUG-GEN. Note the lowest validation loss was achieved at epoch 19. Lower grades represent better chorales.

Bach chorales, so lower grades represent better chorales.

3. Experiments

In our experiments¹, we evaluated the effectiveness of AUG-GEN in improving the output quality of a Transformer model trained to generate Bach-style chorales. We encoded Bach chorales in XML notation using music21 (Cuthbert & Ariza, 2010), and used the same data representation as in (Hadjeres et al., 2017). Our generative model consists of a Transformer network with relative attention (Huang et al., 2018).

3.1. Comparison of Training Methods

We compared three training methods on the same model architecture. In each method, we initialized the dataset to the set of 351 Bach chorales, and split the dataset into 80% training and 20% validation. In the generation step of each training epoch, we generated $N = 50$ chorales, and include a generated example in the training dataset if it passes a threshold t and is unique. In the training step, we trained on 2048 randomly selected batches of size 8 from the full training dataset. We trained each model for 40 epochs.

The three models have equivalent hyperparameters and differ only in the threshold t used for including generated chorales in the training dataset: (1) AUG-GEN ($t = 5.87$, i.e. the third quartile of Bach grades), which includes only generated chorales that receive a better grade than 25% of Bach chorales; (2) **Baseline-none** ($t = -\infty$), which includes no generated chorales, equivalent to training a model on only Bach chorales; (3) **Baseline-all** ($t = \infty$), which includes all generated chorales, regardless of quality.

¹<https://github.com/asdfang/constraint-transformer-bach>




Figure 2. The grade distribution of the 351 Bach chorales, and 351 output generated from each model. The distributions are truncated at 50 due to a long tail from the Baseline-all model. AUG-GEN tends to produce chorales with better grades than either baseline.

3.2. Analysis of Training

Figure 1 shows the grade distribution of output generated by AUG-GEN during each epoch of training. We see that model improvement is clearly reflected in the grading function, suggesting that the grading function can be used to assess model quality independently from the loss function. The first chorale is added to the dataset at epoch 7; by epoch 19, generated examples are 12% of the training dataset. The lowest validation loss achieved by Baseline-none is 0.841 at epoch 16, and 0.846 for AUG-GEN at epoch 19. This suggests AUG-GEN allows a model to be trained for longer without overfitting.

3.3. Grade Distribution

To compare the three methods, we used each model’s epoch that achieved the lowest validation loss. Figure 2 shows the grade distribution of the 351 Bach chorales and 351 generations from each method. We see that a threshold that selects only high-quality generations results in a tighter grade distribution that more closely resembles Bach’s, compared to thresholds that select all or no generations for training.

4. Conclusion

Our experimental evidence suggests AUG-GEN, paired with an effective grading function, allows for longer training and results in better generative output. Listening to generated chorales indicates that remaining errors tend to be along dimensions not measured by the grading function, including excessive modulation, weak metric structure, and unmusical repetition. Future work includes improving the grading function to account for these issues, exploring richer measures of diversity within a dataset, applying AUG-GEN to different models and musical domains, and devising other training methods that utilize generated music data.

References

- Anaby-Tavor, A., Carmeli, B., Goldbraich, E., Kantor, A., Kour, G., Shlomov, S., Tepper, N., and Zwerdling, N. Not enough data? deep learning to the rescue! *Association for the Advancement of Artificial Intelligence*, 2020.
- Briot, J.-P., Hadjeres, G., and Pachet, F. Deep learning techniques for music generation - a survey. *ArXiv*, abs/1709.01620, 2017.
- Cuthbert, M. S. and Ariza, C. music21: A toolkit for computer-aided musicology and symbolic music data. In *Conference of the International Society of Music Information Retrieval (ISMIR)*, pp. 637–642, 2010.
- Fang, A., Liu, A., Seetharaman, P., and Pardo, B. A grading function for four-part chorales in the style of j.s. bach. *Under review*, 2020.
- Hadjeres, G. and Crestel, L. Vector quantized contrastive predictive coding for template-based music generation. *arXiv preprint arXiv:2004.10120*, 2020.
- Hadjeres, G., Pachet, F., and Nielsen, F. DeepBach: a steerable model for Bach chorales generation. In *Proceedings of the 34th International Conference on Machine Learning*, pp. 1362–1371, 2017.
- Huang, C. A., Vaswani, A., Uszkoreit, J., Shazeer, N., Hawthorne, C., Dai, A. M., Hoffman, M. D., and Eck, D. An improved relative self-attention mechanism for transformer with application to music generation. *CoRR*, abs/1809.04281, 2018. URL <http://arxiv.org/abs/1809.04281>.
- Kong, Q., Tong, B., Klinkigt, M., Watanabe, Y., Akira, N., and Murakami, T. Active generative adversarial network for image classification. In *Association for the Advancement of Artificial Intelligence*, 2019.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L. u., and Polosukhin, I. Attention is all you need. In Guyon, I., Luxburg, U. V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, 2017.