
Learning to Generate Music with BachProp

Florian Colombo

School of Computer Science and School of Life Sciences
École Polytechnique Fédérale de Lausanne, Lausanne, Switzerland
florian.colombo@epfl.ch

Johanni Brea

School of Computer Science and School of Life Sciences
École Polytechnique Fédérale de Lausanne, Lausanne, Switzerland
johanni.brea@epfl.ch

Wulfram Gerstner

School of Computer Science and School of Life Sciences
École Polytechnique Fédérale de Lausanne, Lausanne, Switzerland
wulfram.gerstner@epfl.ch

Abstract

As deep learning advances, algorithms of music composition increase in performance. However, most of the successful models are designed for specific musical structures. Here, we present BachProp, an algorithmic composer that can generate music scores in many styles given sufficient training data. To adapt BachProp to a broad range of musical styles, we propose a novel representation of music and train a deep network to predict the note transition probabilities of a given music corpus. In this paper, new music scores generated by BachProp are compared with the original corpora as well as with different network architectures and other related models. We show that BachProp captures important features of the original datasets better than other models and invite the reader to a qualitative comparison on a large collection of generated songs.

1 Introduction

In search of the computational creativity frontier [1], machine learning algorithms are more and more present in creative domains such as painting [2, 3] and music [4, 5, 6]. Already in 1847, Ada Lovelace predicted the potential of analytical engines for algorithmic music composition [7]. Current models of music generation include rule based approaches, genetic algorithms, Markov models or more recently artificial neural networks [8].

One of the first artificial neural networks applied to music composition was a recurrent neural network trained to generate monophonic melodies [9]. In 2002, networks of long short-term memory (LSTM) [10] were applied for the first time to music composition, so as to generate Blues monophonic melodies constrained on chord progressions [11]. Since then, music composition algorithms employing LSTM units, have been used to generate monophonic [4, 5] and polyphonic music [12, 13, 14, 6] or to harmonize chorales in the style of Bach [14, 6]. However, most of these algorithms make strong assumptions about the structure of the music they model.

Here, we present a neural composer algorithm named *BachProp* designed to generate new music scores in an arbitrary style implicitly defined by the corpus of training data. To this end, we do not

assume any specific musical structure of the data except that it is composed of sequences of notes that are characterized by pitch, duration and time-shift relative to the previous note. This time-shift can be zero to represent chords, i.e. notes played at the same time. We indicate why our novel representation of music is superior to previous propositions [12, 14, 6, 15] for the purpose of training style-agnostic generative models of music. We compare BachProp with other models on a standard datasets of chorales written by Johann Sebastian Bach [16] and establish new benchmarks on the musically complex datasets of MIDI recordings by John Sankey [17] and string quartets by Haydn and Mozart [18]. As the evaluation and comparison of generative models is not trivial [19], we invite the reader, first, to a subjective comparison on a large collection of samples generated from the different models on the accompanying media webpage[20] and, second, we propose a new set of metrics to quantify differences between the models.

2 Related work

Unlike approaches to image generation, where the standard data consists of rows and columns of pixel values for multiple color channels, approaches to music generation lack a standard representation of music data. This is reflected by the zoo of music notation file formats (ABC, LilyPond, MusicXML, NIFF, MIDI) and the fact that lossless conversion from one to the other is usually not possible. The MIDI file format captures most features of music, like polyphony, dynamics, micro tuning, expressive timing and tempo changes. But its representational richness and the possibility to represent the exact same song in multiple ways, make it challenging to work directly with MIDI. Therefore, all approaches discussed in the following use a first preprocessing step to transform all songs into a simpler representation. The subsequent design choices of the generative model are heavily influenced by this first preprocessing step.

DeepBach [6] is designed exclusively for songs with a constant number of voices (e.g. four voices for a typical Bach chorale) and a discretization of the rhythm into multiples of a base unit, e.g. 16th notes. The model achieves good results not only in generating novel songs but allows also in reharmonizing given melodies while respecting user-provided meta-information like the temporal position of fermatas. The model works with a Gibbs-sampling-like procedure, where, for each voice and time step, one note is sampled from conditional distributions parameterized by deep neural networks. The conditioning is on the other voices in a time window surrounding the current time-step. Additionally a “temporal backbone” signals the position of the current 16th note relative to quarter notes and other meta-information. A special hold symbol can also be sampled instead of a note, to represent notes with a duration longer than one time-step.

BachBot [14] and its Magenta implementation Polyphony-RNN [15] contain no assumption about the number of voices; they can be fit to any corpus of polyphonic music, if the rhythm can be discretized into multiples of a base unit, e.g. 16th notes. Songs are represented as sequences of `NEW_NOTE(PITCH)`, `CONTINUED_NOTE(PITCH)` and `STEP_END` events, where the `STEP_END` event indicates the end of the current time-step. Between two `STEP_END` events, typically several `NEW_NOTE(PITCH)` and `CONTINUED_NOTE(PITCH)` events can be found sorted by `PITCH`. A generative model parametrized by a recurrent neural network model is fit to these sequences of events, in the same way as recurrent neural network models are used for language modeling on a character- or word-level [21, 22, 23].

Common to the models discussed above is a discretization of the rhythm into multiples of a base unit like the 16th note. This limits the representable rhythms considerably; e.g. triplets, grace notes or expressive variations in timing cannot be represented in this way. To overcome this limitation, [24] replace the repertoire of symbols employed by the Polyphony-RNN by `NOTE_ON`, `NOTE_OFF`, `TIME_SHIFT` and `SET_VELOCITY` events, where the `TIME_SHIFT` events allows the model to move forward in time by multiples of 8 ms up to 1 second and the `SET_VELOCITY` events allow to model the loudness of a note (which depends on the piano on the velocity with which a key is pressed).

3 Method

In written music, the n^{th} note $\mathbf{note}[n]$ of a piece of music $\mathbf{song} = (\mathbf{note}[1], \dots, \mathbf{note}[N])$ can be characterized by its pitch $P[n]$, duration $T[n]$ and the time-shift $dT[n]$ of its onset relative to the previous note, i.e. $\mathbf{note}[n] = (dT[n], T[n], P[n])$. The time-shift $dT[n]$ is zero for notes played at

Table 1: **Duration and time-shift dictionary**. The values on the right for the dotted, double dotted and triplet notes should be multiplied with 2^{-4} to 2^3 to get the full set of $4 \cdot 8 = 32$ possible durations $T[n]$ and $32 + 1$ time-shifts $dT[n]$, including a time-shift of zero.

music notation												
in our dictionaries	2^{-4}	2^{-3}	2^{-2}	2^{-1}	2^0	2^1	2^2	2^3		$3/2$	$7/4$	$2/3$

the same time as the previous note. In contrast to most other approaches that discretize the rhythm into multiples of a base unit (except e.g. [24]), we round all durations into a set of common musical durations which allows a more faithful representation of timing that is limited only by the number of possible values considered for $T[n]$ and $dT[n]$. For example, our representation allows to easily and without any distortion represent 32nd notes, triplets and dotted notes in the same dataset (see Table 1). As well as any other more complex note durations that can be needed for specific corpora.

Our approach is to approximate probability distributions over note sequences in music scores $\text{song}_1, \dots, \text{song}_S$ with distributions parameterized by recurrent neural networks and move its weights θ towards the maximum likelihood estimate

$$\theta^* = \arg \max_{\theta} Pr(\text{song}_1, \dots, \text{song}_S | \theta), \quad (1)$$

Since each note in each song consists of the triplet $(dT[n], T[n], P[n])$ we can parametrize the distributions in a similar way as the pixel-RNN [25] that was developed for the (red, green, blue) triplets of pixels in images. Importantly, our model takes into account that pitch and duration of a note are generally not independent. For example in classical music, the fundamental, e.g. the note C in a piece written in C major, tends to be longer than other notes.

In the following we describe in more details our representation of music, the structure of the model and our approach to comparing different models that use different representations of music.

3.1 Conversion of MIDI files into our representation of music




Figure 1: **From MIDI to our representation of music**. An illustration of the steps involved in the proposed conversion of MIDI sequences. See text for details.

A MIDI file contains a header (meta parameters) and possibly multiple tracks that contain a sequence of MIDI messages. For BachProp, we merge all tracks and consider only the MIDI messages defining when a note starts (ON events) or ends (OFF events). For each ON event we look forward at the next OFF event with the same pitch P to convert sequences of MIDI messages into a sequences of notes (Figure 1A). We then translate timings from the internal MIDI TICK representation to quarter note lengths (Figure 1B).

We round all durations such that they are in a set of 32 possible note lengths (duration dictionary; see Table 1) expressed in units of a quarter note, similar to durations in standard music notation software. Similarly, we round the time-shifts to the 0 or one of the 32 possible note lengths. Mapping to the closest value in the set removes temporal jitter around the standard note duration that may have been introduced accidentally at the moment of recording the MIDI file (Figure 1C). While this standardization may be desired when expressive timing is not taken into account, it is straightforward to extend the duration dictionary to include also values that allow to model expressive timing.

In order for BachProp to learn tonality and transposition invariance of music, we transpose each song within the available bounds of the pitch set. For each song we compute the possible shifts of

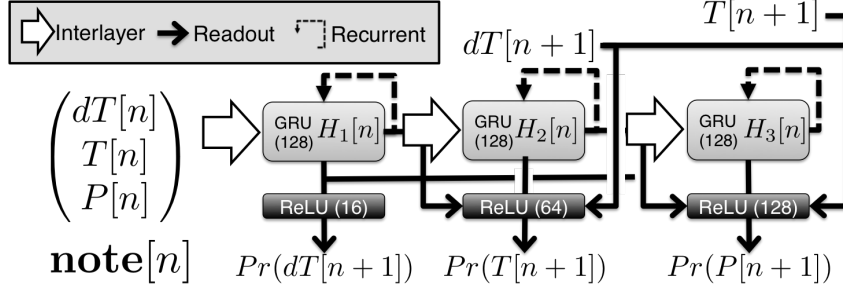


Figure 2: **BachProp neural architecture.** See text for details.

semitones and apply them as an offset to all pitches in the song. Because a single MIDI sequence will be transposed with up to 20 offsets, this augmentation method allows BachProp to learn the temporal structure of music on more examples.

Finally, we add an artificial note at the beginning and end of each score. After training, the inaudible ‘end note’ is generated by the model to seed and end the generation of songs.

3.2 The BachProp neural network

We used a deep GRU [26] network with three consecutive layers as schematized in Figure 2. The network’s task is to infer the probability distribution over the next possible notes from the representation of the current note and the network’s internal state (the network representation of the history of notes).

The probability of a sequence of N notes $\mathbf{note}[1 : N] = (\mathbf{note}[1], \dots, \mathbf{note}[N])$ is given by

$$Pr(\mathbf{note}[1 : N]) = Pr(\mathbf{note}[1]) \prod_{n=1}^{N-1} Pr(\mathbf{note}[n+1] | \mathbf{note}[1 : n]). \quad (2)$$

Each term on the right hand side can be further split into

$$Pr(\mathbf{note}[n+1] | \mathbf{note}[1 : n]) = Pr(dT[n+1] | \mathbf{note}[1 : n]) \times Pr(T[n+1] | \mathbf{note}[1 : n], dT[n+1]) \times Pr(P[n+1] | \mathbf{note}[1 : n], dT[n+1], T[n+1]). \quad (3)$$

The goal of training the Bachprop network with parameters θ is to approximate the conditional probability distributions on the right hand side of Equation 3.

In the BachProp network (Figure 2), the conditioning on the history $\mathbf{note}[1 : n]$ (see Equation 3) is implemented by the values of the shared hidden states. The hidden state is composed of 3 recurrent layers with 128 gated-recurrent units (GRU). The state $H_1[n]$ of the first hidden layer is updated with input $\mathbf{note}[n]$ and previous state $H_1[n-1]$. The state of the upper layers $H_i[n]$ for $i = 2, 3$ is updated with input $H_{i-1}[n]$ and $H_i[n-1]$. To generate $\mathbf{note}[n+1]$, one third ($H_1[n]$ in Figure 2) of the full hidden state is fed into a feedforward network with one layer of 16 Relu units and one output softmax-layer that represents $Pr(dT[n+1] | H_1[n]) \approx Pr(dT[n+1] | \mathbf{note}[1 : n])$. The chosen $dT[n+1]$ together with $H_1[n]$ and $H_2[n]$ is fed into a second feedforward network with one layer of 64 Relu units and an output softmax-layer that represents $Pr(T[n+1] | H_1[n], H_2[n], dT[n+1]) \approx Pr(T[n+1] | \mathbf{note}[1 : n], dT[n+1])$. In a similarly way, the pitch is sampled from $Pr(P[n+1] | H_1[n], H_2[n], H_3[n], dT[n+1], T[n+1]) \approx Pr(P[n+1] | \mathbf{note}[1 : n], dT[n+1], T[n+1])$. These three small steps of sampling $dT[n+1]$, $T[n+1]$ and $P[n+1]$ form together one big step from note n to note $n+1$.

The resulting sequence of notes is a newly generated score sampled from BachProp. Note that, the temperature of sampling can be adapted to the confidence we give to the model predictions [27, 5]. In particular, any model trained with a corpus that exhibits many repetition of patterns, will generate scores with more examples of these repetitions for lower sampling temperatures. Indeed, a lower temperature will reduce the probability to select an undesired note that is not part of the pattern to be

repeated. Finally, the generated sequence of notes in our representation can easily be translated back to a MIDI sequence by reversing the method schematized in Figure 1.

BachProp has been implemented in Python using the Keras API [28]. Code is available on GitHub¹.

3.3 Comparison against plagiarism and other models

Even in well-established domains such as computer vision and image generation, it is not clear how to compare generative models [19]. But in order to turn generative models of music eventually into useful tools for composers, they should be able to generate (1) plagiarism-free music of (2) a predefined style or mood that is (3) pleasant to listen to.

A way of measuring plagiarism is to control overfitting by comparing the loss on training and validation data. While this is a simple method it is rather coarse since it works on songs as a whole. Instead we propose *novelty profiles* that compare the co-occurrence of short note sequences across different data sets. A crucial parameter of novelty profiles is the length of a note sequence on which the comparison takes place. We adapted the novelty profile, a measure of similarity between any given score and a reference corpus, from [5]. For a pattern size of 6 notes, a novelty score of 1 indicates that all patterns of 6 consecutive notes are not present in the reference corpus. On the other hand, a note sequence that contains only patterns found in the reference corpus would exhibit a novelty score of 0. We define the binary novelty of a single pattern by checking if all three features ($dT[n - m : n]$, $T[n - m : n]$, $P[n - m : n]$) of the notes included in the pattern are found in the same order anywhere in the reference corpus. The novelty score of an entire song is the average binary novelty over all possible patterns.

Models that are trained on the same representation of music can be compared by their likelihood to assess how well they generate pieces of a predefined type. But if the models represent probability distributions over different spaces, which is quickly the case when different representations are used, they are unfortunately not comparable in terms of likelihood. For example, the event based representation from [24] can in principle produce all possible note sequences. But it could also generate nonsensical sequences of multiple consecutive NOTE_OFF events, without corresponding previous NOTE_ON events. To nevertheless compare models that build on different representations of music we propose simple statistics like interval distributions that can be applied to the samples of each generative model of music.

Finally, to compare the pleasantness of the generated music, one can ask people to rate different pieces; an approach that is followed in previous works (e.g. [6]). We also invite the reader to listen to the large collections of non-cherry-picked generated examples [20].

4 Results and discussion

4.1 Datasets

We consider four MIDI corpora with different musical structures and styles (see Table 3). The Nottingham database [29] contains British and American folk tunes. The musical structure of all songs is very similar with a melody on top of simple chords. The Chorales corpus [16] includes hundreds of four-part chorales harmonized by Bach. All chorales share some common structures, such as the number of voices and rhythmical patterns. For comparison we used the same filtering of songs as DeepBach [30] to exclude chorales with number of voices unequal four. We consider both Nottingham and Chorales corpora as homogeneous data sets. The John Sankey data set [17] is a collection of MIDI sequences recorded by John Sankey on a digital keyboard. Even though all songs were composed by Bach, the pieces are rather different. In addition, this data set was recorded live from the digital keyboard and thus we applied the temporal normalization described above. At last, the string quartets data set [18] includes string quartets from Haydn and Mozart. Here again, there is a large heterogeneity of pieces across the corpus.

Renderings of scores generated by BachProp are available for listening on the webpage containing media for this paper². They are the result of five BachProp Networks. All networks had the same

¹<https://github.com/FlorianColombo/BachProp>

²Media webpage: <https://goo.gl/Z4AfPg>

architecture, number of neurons, and learning parameters, but each of the network was trained on a different corpus.

4.2 Alternative models

We trained six alternatives to BachProp. PolyDAC and IndepBP are direct BachProp variants. MidiBP is a version of BachProp that utilizes a different representation of MIDI note sequences inspired by [24]. Along with two state-of-the-art artificial composers, DeepBach [6] and PolyRNN [15], it allows us to compare our representation of music scores with five score generating models of our design. The 6th model is a multi-layer perceptron model (MLP) and serves as a baseline control.

PolyDAC is a polyphonic version of [5]. It models the same conditional distribution as BachProp but instead of reading out the probabilities from shared hidden layer states, it models each note feature with three independent neural networks. The time-shift, duration, and pitch networks are composed of three recurrent layers with 16, 128, and 256 GRUs respectively. **IndepBP** assumes that all note features are independent from each others. As such, $Pr(dT[n+1])$, $Pr(T[n+1])$, and $Pr(P[n+1])$ are read out by three softmax output layers directly from the hidden state of three hidden layers composed of 128 GRUs that takes as input the one-hot encoding of the n^{th} note. **MidiBP** neural architecture consists of three recurrent layers composed of 128 GRUs. Here, the MIDI note sequences are represented differently. While the normalization and preprocessing is done as described above (Figure 1), we then convert the normalized music score back to the MIDI-like format proposed in [24] where in each time step a single on-hot vector defines either a NOTE_ON event and its corresponding pitch, a NOTE_OFF event and its corresponding pitch, or a time-shift and its corresponding duration (defined by our duration representation). Therefore, a single softmax read out layer is used to sample the upcoming MIDI event. **MLP** has no recurrent layers but 3 feedforward hidden layers of 124 ReLUs each that gets as input the 5 most recent notes $\mathbf{note}[n-4:n]$ together with the current time-shift $dT[n+1]$ and duration $T[n+1]$ to sample the pitch $P[n+1]$. To sample the duration $T[n+1]$ and the time-shift $dT[n+1]$, appropriate parts of the input are masked with zeros.

Models BachProp, PolyDAC, MidiBP, IndepBP were trained with truncated back propagation through time and the Adam optimizer [31]. The MLP model was trained with standard back propagation and the Adam optimizer. The mini-batch size is 32 scores, the validation set a 0.1 fraction of the augmented original corpora, and one training epoch consists of updating the network parameters with all training examples and evaluating the performances on the entire validation set. Training is stopped when the performances on the validation set saturates and the model leading to the highest accuracy is used for generating new music scores. DeepBach was trained for 15 epochs with the standard settings of the current master branch [30]. PolyRNN was trained for 26000 steps with the standard settings of the current master branch [15].

Table 2: **Comparison of architectures on our representation of music.** NLL stands for negative log-likelihood on the validation set. Columns dT , T and P indicate the accuracy (fraction of correct predictions) for time-shifts, durations and pitches, respectively.

MODEL	NLL	dT	T	P
BACHPROP	0.419	0.97	0.91	0.77
POLYDAC	0.647	0.97	0.94	0.69
INDEPBP	0.647	0.97	0.75	0.63
MLP	0.796	0.95	0.76	0.49

4.3 BachProp performs better than alternative models with same representation

On the Bach Chorales we find that the BachProp architecture performs considerably better than the alternative architectures using the same representation of music (see Table 2). As expected, the standard feedforward MLP with ReLUs yields the worst performance. It lacks the ability to model long range dependencies, which the other models can do through their recurrent connections. When we remove the conditioning on each of probability terms on the right side of Equation 3, as done for the IndepBP model, we get poorer performances. We further observe that sharing a common hidden state allowed BachProp to outperform PolyDAC on the pitch predictions.

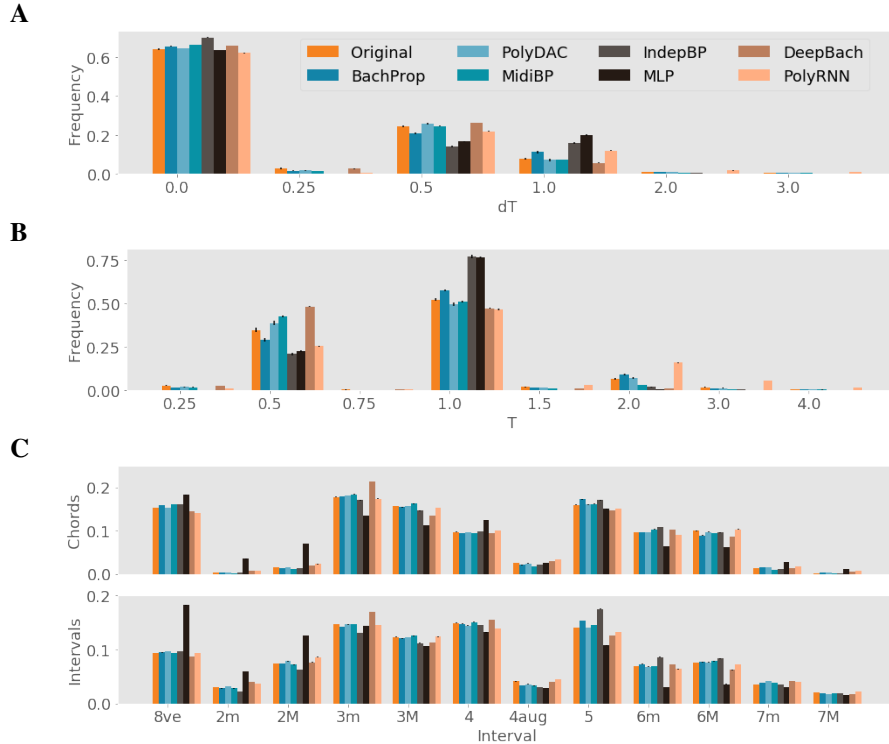


Figure 3: **Local statistics.** **A** Distribution of dT . **B** Distribution of T . **C** Distribution of intervals in chords (top) and between each note (bottom). For all figures, we show the mean and standard deviation (in black) obtained with bootstrapping (50% of the entire corpus resampled 10 times). All models were trained on the Bach Chorales corpus.

4.4 BachProp performs at least as good as alternatives with different representation

To compare models that use a different representation of music, we look at a set of metrics that includes local statistics, song-length statistics and novelty profiles. To evaluate these metrics for each model, we generated from each model a set containing as many scores as the original corpus (400 songs). We include the baseline models from the last section for comparison reasons.

4.4.1 Local statistics

A model that has captured the underlying structure of the sequences of notes present in a corpus, should be able to generate new scores matching the local statistics of what they modeled. As such, we suggest to compute the distributions of generated dT and T and compare them to the original corpus distributions as a first metric to evaluate generative models of music. Note that for such direct local statistics, a simple n-gram model would match the original distributions perfectly. Figure 3A and B shows that BachProp and PolyDAC match the original distributions best, followed by MidiBP, DeepBach and PolyRNN, while IndepBP and MLP match the least.

Next, we look at interval distributions. An interval is the number of half-tone separating two notes. Here, BachProp, PolyDAC, MidiBP and PolyRNN match the distribution quite well. DeepBach seems to generate minor thirds considerably more often than present in the training data (Figure 3C).

4.4.2 Distribution of song lengths

The distribution of song lengths can indicate whether a model captured really long-range dependencies in the training set. On this measure MidiBP matches the distribution slightly better than BachProp, PolyDAC, IndepBP and MLP (see Figure 4A). Since DeepBach and PolyRNN do not model score endings, we manually set their duration.




Figure 4: **Song lengths and novelty profiles.** **A** Distribution of the duration of scores in quarter note length. **B** Novelty profile of all corpora with respect to the auto-novelty of the original corpus. **C** The auto-novelty profiles of all corpora. See text for details.

Table 3: **BachProp on other datasets.** See Table 2 for description of labels.

DATASET	NLL	dT	T	P	SIZE [SCORE]	SIZE [NOTE]
CHORALES	0.419	0.97	0.91	0.77	357	95'337
NOTTINGHAM	0.587	0.98	0.89	0.70	1037	313'975
JOHN SANKEY	1.002	0.89	0.77	0.45	135	358'211
STRING QUARTETS	0.936	0.88	0.83	0.49	215	738'739

4.4.3 Novelty profiles

In Figure 4B, we compare the novelty profiles for all models with respect to the original Chorales corpus with which each model was trained. We compare the different profiles with the auto-novelty of the reference corpus. The auto-novelty is the novelty profile for each song in the reference corpus with respect to the same corpus without the song for which the novelty score is computed. It reflects, how similar is the music within the original corpus and is consequently the distribution to match for an ideal generative model of music. Here, the only model that is clearly outside the target distribution is the MLP model. While the IndepBP and MidiBP models match the target distributions, their novelty distributions for bigger pattern sizes is lower than the original corpus auto-novelty. This is an indicator that these models are generating music examples that are too similar to the original data. In other words, these models adopted a strategy closer to reproducing or recombining observed patterns rather than inferring the actual temporal dependencies between music notes. DeepBach, BachProp and PolyDAC have their medians close and above the original distributions. However, DeepBach and PolyRNN have a surprisingly low variance for each of the pattern sizes.

In Figure 4C we compare the auto-novelty of all generated corpora with the original corpus. An auto-novelty profile exhibiting distributions with lower novelty scores than the original data set, is suspected to generate new music scores of little diversity. The auto-novelty profile of BachProp and PolyDAC match the one of the original corpus best.

4.5 BachProp generates pleasant examples on more complex datasets

As a reference for future comparisons, we report here the results of BachProp trained on more complex datasets. In Table 3, we observe that for homogeneous corpora with many examples of similar structures (Chorales, Nottingham), BachProp can predict notes with higher accuracies than for more heterogeneous data sets (John Sankey, String Quartets).

We encourage readers to listen to the examples provided on the accompanying webpage to convince themselves of the ability of BachProp and its variants to generate unique and heterogeneous new music scores.

5 Conclusion

In this paper, we presented BachProp, an algorithm for general automated music composition. Our main contributions are (1) a note-sequence based representation of music with minimal distortion of the rhythm for training neural network models, (2) a network architecture that learns to generate pleasant music in this representation and (3) a set of metrics to compare generative models that operate on different representations of music.

References

- [1] Simon Colton, Geraint A Wiggins, et al. Computational creativity: The final frontier? In *ECAI*, volume 12, pages 21–26, 2012.
- [2] Alexander Mordvintsev, Christopher Olah, and Mike Tyka. Inceptionism: Going deeper into neural networks. *Google Research Blog*. Retrieved June, 20(14):5, 2015.
- [3] Leon A Gatys, Alexander S Ecker, and Matthias Bethge. Image style transfer using convolutional neural networks. In *Computer Vision and Pattern Recognition (CVPR), 2016 IEEE Conference on*, pages 2414–2423. IEEE, 2016.
- [4] Bob L Sturm, Joao Felipe Santos, Oded Ben-Tal, and Iryna Korshunova. Music transcription modelling and composition using deep learning. In *1st Conference on Computer Simulation of Musical Creativity*, 2016.
- [5] Florian Colombo, Alexander Seeholzer, and Wulfram Gerstner. Deep artificial composer: A creative neural network model for automated melody generation. In *International Conference on Evolutionary and Biologically Inspired Music and Art*, pages 81–96. Springer, 2017.
- [6] Gaëtan Hadjeres, François Pachet, and Frank Nielsen. DeepBach: a steerable model for Bach chorales generation. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 1362–1371. PMLR, 2017.
- [7] Ada Lovelace. Notes on l. menabrea’s ”sketch of the analytical engine invented by charles babbage, esq.”. *Taylor’s Scientific Memoirs*, 3:1843, 1843.
- [8] Jose D Fernández and Francisco Vico. Ai methods in algorithmic composition: A comprehensive survey. *Journal of Artificial Intelligence Research*, 48:513–582, 2013.
- [9] Peter M Todd. A connectionist approach to algorithmic composition. *Computer Music Journal*, 13(4):27–43, 1989.
- [10] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [11] Douglas Eck and Juergen Schmidhuber. Finding temporal structure in music: Blues improvisation with lstm recurrent networks. In *Proceedings of the 12th IEEE Workshop on Neural Networks for Signal Processing*, pages 747–756. IEEE, 2002.
- [12] Nicolas Boulanger-Lewandowski, Yoshua Bengio, and Pascal Vincent. Modeling Temporal Dependencies in High-Dimensional Sequences: Application to Polyphonic Music Generation and Transcription. *ArXiv:1206.6392*, 2012.
- [13] Stefan Lattner, Maarten Grachten, and Gerhard Widmer. Imposing higher-level structure in polyphonic music generation using convolutional restricted boltzmann machines and constraints. *Journal of Creative Music Systems*, 2(1), 2018.
- [14] Feynman Liang, Mark Gotham, Matthew Johnson, and Jamie Shotton. Automatic stylistic composition of bach chorales with deep lstm. October 2017.
- [15] Magenta Team Google Brain. Polyphony RNN, revision ca73164. https://github.com/tensorflow/magenta/tree/master/magenta/models/polyphony_rnn, 2016.
- [16] J.S. Bach Chorales. <http://web.mit.edu/music21/>.
- [17] Bach MIDI sequences by John Sankey. http://www.jsbach.net/midi/midi_johnsankey.html. Accessed: 2018-02-04.
- [18] String Quartets by Mozart and Haydn. <http://www.stringquartets.org>.

- [19] Lucas Theis, Aäron van den Oord, and Matthias Bethge. A note on the evaluation of generative models. *ArXiv:1511.01844*, page arXiv:1511.01844, 2015.
- [20] <https://sites.google.com/view/bachprop>.
- [21] Ilya Sutskever, James Martens, and Geoffrey Hinton. Generating text with recurrent neural networks. In *Proceedings of the 28th International Conference on International Conference on Machine Learning, ICML'11*, pages 1017–1024, USA, 2011. Omnipress.
- [22] Alex Graves. Generating Sequences With Recurrent Neural Networks. *ArXiv:1308.0850*, 2013.
- [23] Tomáš Mikolov. *Statistical Language Models Based on Neural Networks*. PhD thesis, 2012.
- [24] Saageev Oore, Ian Simon, Sander Dieleman, and Douglas Eck. Learning to create piano performances. *NIPS 2017 Workshop on Machine Learning for Creativity and Design*, 2017.
- [25] Aaron van den Oord, Nal Kalchbrenner, and Koray Kavukcuoglu. Pixel Recurrent Neural Networks. *ArXiv:1601.06759*.
- [26] Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555*, 2014.
- [27] Andrej Karpathy. The unreasonable effectiveness of recurrent neural networks, 2015. URL <http://karpathy.github.io/2015/05/21/rnn-effectiveness>, 2016.
- [28] François Chollet. keras. <https://github.com/fchollet/keras>, 2015.
- [29] Nottingham data set of folk songs. <http://www-etud.iro.umontreal.ca/~boulanni/icml2012>.
- [30] DeepBach, revision f069695. <https://github.com/Ghadjeres/DeepBach>.
- [31] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.