

MIDI-GPT: A Controllable Generative Model for Computer-Assisted Multitrack Music Composition

Philippe Pasquier¹, Jeff Ens¹, Nathan Fradet¹, Paul Triana¹, Davide Rizzotti¹,
Jean-Baptiste Rolland², Maryam Safi²

¹Metacreation Lab - Simon Fraser University, Vancouver, Canada

²Steinberg Media Technologies GmbH, Hamburg, Germany
pasquier@sfu.ca

Abstract

We present and release MIDI-GPT, a generative system based on the Transformer architecture that is designed for computer-assisted music composition workflows. MIDI-GPT supports the infilling of musical material at the track and bar level, and can condition generation on attributes including: instrument type, musical style, note density, polyphony level, and note duration. In order to integrate these features, we employ an alternative representation for musical material, creating a time-ordered sequence of musical events for each track and concatenating several tracks into a single sequence, rather than using a single time-ordered sequence where the musical events corresponding to different tracks are interleaved. We also propose a variation of our representation allowing for expressiveness. We present experimental results that demonstrate that MIDI-GPT is able to consistently avoid duplicating the musical material it was trained on, generate music that is stylistically similar to the training dataset, and that attribute controls allow enforcing various constraints on the generated material. We also outline several real-world applications of MIDI-GPT, including collaborations with industry partners that explore the integration and evaluation of MIDI-GPT into commercial products, as well as several artistic works produced using it.

1 Introduction

Recent research on generative music systems (Huang et al. 2019; Huang and Yang 2020; Briot, Hadjeres, and Pachet 2019; Fradet et al. 2023a,b) has mainly focused on modeling musical material as an end-goal, rather than on their affordance in practical scenarios (Sturm et al. 2019). Although these works pave the way for efficient generative methods for music, their usability in real-world co-creative conditions remains limited. As a result, while we have seen a wide adoption of generative models for language and vision tasks, this has not occurred to the same extent for symbolic music composition. In contrast, artists recently expressed their concerns about the misuse of artificial intelligence in the music field (Artist Rights Alliance 2024). If we want musicians to adopt generative systems, we must work on making models controllable, able to generate content that the user will appropriate as theirs, and integrated into their existing workflows.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Motivated by these considerations, we introduce MIDI-GPT, a style-agnostic generative model that builds on an alternative representation for multi-track musical material (Ens and Pasquier 2020), resulting in an expressive and steerable generative system. We outline the ongoing real-world usage of MIDI-GPT and provide quantitative evidence to support our claim that MIDI-GPT rarely duplicates the training data as the length of the generated material increases; generates musical material that retains the stylistic characteristics of the training data; and that attribute control methods are an effective way to steer generation.

2 Background

Considering our interest in developing a system that is well suited to practical and interactive computer-assisted composition applications, we must identify the factors that enhance the real-world usability of generative music systems. Our first design decision is to use the General MIDI format as input and output given that it is the most supported symbolic music encoding standard. We consider two main categories: I/O specifications, which place restrictions on the musical material that can be processed and generated by the system; and generation methods, favoring using existing musical content as prompt vs. unconditioned generation.

2.1 Input/Output Specification

We first define a track (t_{inst}), which is a distinct set of musical material (i.e. notes) played by an instrument ($inst$). In some cases, a track may be distinguished by its musical purpose (ex. t_{melody}) rather than its instrument. For example, MusicVAE (Roberts et al. 2018) aggregates all melodies into a single track type, rather than distinguishing between melodies based on their instrumentation (piano, synth, saxophone, etc.). In what follows, t_{inst}^m , t_{inst}^p and t_{drum} denote a monophonic track, a polyphonic track, and a drum track, respectively. For example, t_{bass}^m denotes a monophonic bass track.

A generated excerpt can be described by a list of track types (ex. $[t_{bass}^m, t_{piano}^p]$), and thus, we can define the output specification (O^*) for an arbitrary generative music system as a set of track lists (ex. $O^* = \{[t_{bass}^m, t_{piano}^p], [t_{piano}^p, t_{synth}^p]\}$). We consider a system to have a fixed schema when O^* contains a single tracklist. For

example, CoCoNet (Roberts et al. 2018) has a fixed schema, as $O^* = \{[t_{\text{soprano}}^m, t_{\text{alto}}^m, t_{\text{tenor}}^m, t_{\text{bass}}^m]\}$, meaning that the system is only capable of generating 4 track music containing soprano, alto, tenor and bass tracks.

In Table 1, we describe the following features of symbolic music generation systems: the number of tracks, the number of instruments, whether a fixed schema of instruments is assumed, support for drum tracks, and support for polyphony at the track-level. Clearly, reducing the restrictions on a system’s output increases the usability of a system, as it can accommodate a greater number of practices and user workflows. We leave aside style-specific generative systems (Ren et al. 2020; Collins and Barthet 2023; Wu and Yang 2020; Liu et al. 2022; Huang and Yang 2020; Hsiao et al. 2021), as we aim for a style-agnostic system that can accommodate as many users as possible.

As shown in Table 1, most systems either support a single track or require a fixed schema of instruments. One exception is MuseNet (Payne 2019), which supports up to 10 tracks and any subset of the 10 available instruments. However, there are significant differences between MuseNet and MIDI-GPT. MuseNet uses separate `NOTE_ON` and `NOTE_OFF` tokens for each pitch on each track, placing inherent limitations on the number of tracks that can be represented, as the token vocabulary size cannot grow unbounded (Fradet et al. 2023a). Considering that MuseNet is currently the largest model in terms of the number of weights, the number of tracks is unlikely to be increased without altering the representation. Instead, we decouple track information from `NOTE_ON`, `NOTE_DUR`, and `NOTE_POS` tokens, allowing the use of the same tokens for each track. Although this is a relatively small change, it enables us to accommodate all 128 General MIDI instruments. Furthermore, there is no inherent limit on the number of tracks, as long as the entire n -bar multi-track sequence can be encoded using less than 2048 tokens. Practically, this means more than 10 tracks can be generated at once depending on their content. Note that the upper limit of 2048 tokens is not a limitation of the representation itself, but rather the size of the model, and this limitation could be addressed with larger and more memory-intensive models. Both MuseNet and MIDI-GPT do not require a fixed instrument schema, however, MuseNet treats instrument selections as a suggestion, while MIDI-GPT guarantees a particular instrument will be used.

2.2 Generation Tasks

We consider four different generation tasks: unconditional generation, continuation, infilling, and attribute control. Unconditional generation produces music from scratch. Besides changing the data that the model is trained on, the user has limited control over the output of the model. Continuation involves conditioning the model with musical material temporally preceding the music that is to be generated. Since both unconditional generation and continuation come for free with any auto-regressive model trained on a temporally ordered sequence of musical events, all systems are capable of generating musical material in this manner. Infilling conditions generation on a subset of musical material,

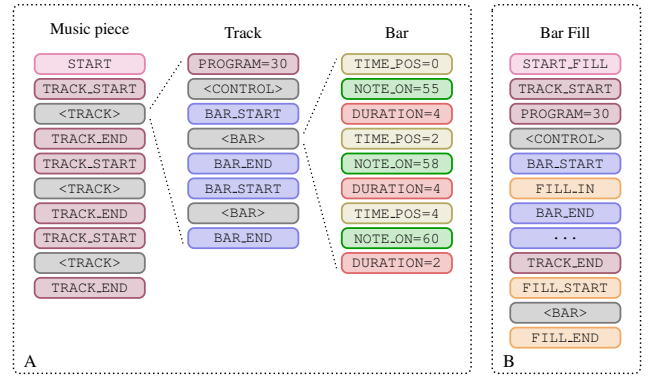


Figure 1: The Multitrack (A) and Bar-Fill (B) tokenizations. The grey `<BAR>`, `<TRACK>` and `<CONTROL>` placeholders correspond token subsequences of complete bars, complete tracks, and attribute controls, respectively.

asking the model to fill in the blanks, so to speak. Although the terms infilling and inpainting are often used interchangeably, some important distinctions must be made in our context. In contrast to inpainting a section of an image, where the exact location and number of pixels to be inpainted are defined before generation, when infilling a section of music the number of tokens to be generated is unknown. Furthermore, in the context of multi-track music that is represented using a single time-ordered sequence where tracks are interleaved, the location of tokens to be added is unknown. This makes bar-level and track-level infilling quite complex, directly motivating the representation we describe in Section 3. With tracks ordered sequentially, the location of the tokens to be infilled is then known. Infilling can occur at different levels (i.e. note-level, bar-level, track-level). Track-level infilling is the most coarse and allows a set of n -tracks to be generated that are conditioned on a set of k existing tracks, resulting in a composition with $k + n$ tracks. Bar-level infilling allows for n -bars selected across one or more tracks to be re-generated, and conditioned on the remaining content - past, current, and future - both on the track(s), and all other tracks.

3 Proposed Music Tokenization

In this section, we introduce two tokenizations to interpret musical compositions: the Multi-Track representation and the Bar-Fill representation. In contrast to other systems (Oore et al. 2020; Huang et al. 2019), which use `NOTE_ON`, `NOTE_OFF` and `TIME_DELTA` tokens, we represent musical material using an approach which was previously employed for the Pop Music Transformer (Huang and Yang 2020). In our Multi-Track representation, each bar of music is represented by a sequence of tokens, which include:

- 128 `NOTE_ON` tokens: These represent the pitch of each note in the bar.
- 96 `TIME_POSITION` tokens: These represent the absolute start time (the time elapsed since the beginning of the bar, as opposed to the time elapsed since the last event) of each note within the bar.

System	I/O Specifications					Generation Tasks	
	Number of Tracks	Number of Instruments	Fixed Schema	Drums	Track-Level Polyphony	Infilling	Attribute Control
MIDI-GPT	Any	128	no	yes	yes	yes	yes
FIGARO (von Rütte et al. 2023)	128	128	no	yes	yes	no	no
MMT (Dong et al. 2023)	64	64	yes	yes	yes	no	no
MuseNet (Payne 2019)	10	10	no	yes	yes	no	yes
MuseGAN (Hong et al. 2019)	4	4	yes	yes	yes	no	no
LahkNES (Donahue et al. 2019)	4	4	yes	yes	no	no	no
CoCoNet (Huang et al. 2018)	4	4	yes	no	no	yes	no
MusicVAE (Roberts et al. 2018)	3	3	yes	yes	no	no	no
MusIAC (Guo et al. 2022)	3	3	yes	no	yes	yes	yes
SketchNet (Chen et al. 2020)	1	no	yes	no	no	yes	yes
(Pati, Lerch, and Hadjeres 2019),(Mittal et al. 2021)	1	no	yes	no	no	yes	no
(Chang, Lee, and Yang 2021),(Chi et al. 2020)	1	no	yes	no	yes	yes	no
(TAN and Herremans 2020),(Wang and Xia 2021), (Wang et al. 2020)	1	no	yes	no	yes	no	yes
(Haki et al. 2022), (Nuttall, Haki, and Jorda 2021)	1	no	yes	yes	yes	no	no
(Huang et al. 2019)	1	no	yes	no	yes	no	no

Table 1: A summary of the I/O specifications and generation tasks of recently published generative music systems.

- **96 DURATION tokens:** These represent the duration of each note. Both the `DURATION` and `TIME_POSITION` tokens range from a sixteenth-note triplet to a double whole note in sixteenth-note triplet increments.

We delimit a bar with `BAR_START` and `BAR_END` tokens. A sequence of bars makes up a track, which is delimited by `TRACK_START` and `TRACK_END` tokens. At the beginning of each track, one of 128 `INSTRUMENT` token specifies its MIDI program. Tokens that condition the generation of each track on various musical attributes follow the `INSTRUMENT` token, and will be discussed in Section 4. The tracks are then nested within a multi-track piece, which begins with a `START` token. Note that all tracks are played simultaneously, not sequentially. This process of nesting bars within a track and tracks within a piece is illustrated in Figure 1A. Notably, we do not use an `END` token, as we can simply sample until we reach the n^{th} `TRACK_END` token if we wish to generate n tracks. This tokenization is implemented in `Mid-iTok` (Fradet et al. 2021) for ease of use.

The Multi-Track representation allows the model to condition the generation of each track on the tracks that precede it, which allows for a subset of the musical material to be fixed while generating additional tracks. However, this representation doesn’t provide control at the bar level, except in cases where the model is asked to complete the remaining bars of a track. In other words, the model cannot fill in bars that are in the middle of a track. To generate a specific bar in a track conditioned on the other bars, we introduce the Bar-Fill representation. In this representation, bars to be predicted are replaced by a `FILL_IN` token. These bars are then placed/generated at the end of the piece after the last track token, and each bar is delimited by `FILL_START` and `FILL_END` tokens (instead of `BAR_START` and `BAR_END` tokens).

Note that during training, the bars with `FILL_IN` tokens appear in the same order as they appeared in the original Multi-Track representation, shown in Figure 1B. By order-

ing the bars consistently, the model learns to always output tokens in the same order as the bars that are marked for generation. The Bar-Fill representation begins with a `START_FILL` token instead of a `START` token. The Multi-Track representation is simply a special case of the Bar-Fill representation, where no bars are selected for infilling.

3.1 Adding Interpretation Expressiveness

Multiple attempts at generating expressive symbolic music have been made either as an independent process (Gillick et al. 2019; Cancino-Chacón and Grachten 2016; Malik and Ek 2017; Maezawa, Yamamoto, and Fujishima 2019) or a simultaneous process with the musical content generation (Oore et al. 2020; Huang et al. 2019; Hawthorne et al. 2018; Huang and Yang 2020; Wu et al. 2022). None, however, allow for expressive multitrack generation. Here, we focus on velocity, as a proxy for dynamics, and microtiming, as the two main aspects of expressive music interpretation. We implement two extensions to our current tokenisation allowing the simultaneous generation of expressive MIDI. This allows us to leverage the 31% of MIDI files in GigaMIDI that have been marked as expressive (varying velocity, and non-quantized micro-timing).

Firstly, we include 128 `VELOCITY` tokens that encode every possible velocity level of a MIDI note, as velocity is a proxy for dynamics and an important aspect of expressiveness in musical performances.

Secondly, we include new tokens to represent microtiming. Our current tokenization allows for 96 different `TIME_POSITION` tokens within a bar. Therefore, this level of quantization occurs in the model which does not capture microtiming. Intuitively, a solution to this problem would be to increase the vocabulary and time resolution of the `TIME_POSITION` tokens. However, to maintain the possibility of using the current downsampled and non-expressive tokenization while allowing the possibility to add expressiveness, we introduce a new token `DELTA` which encodes

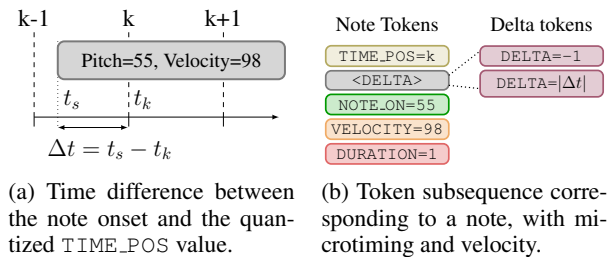


Figure 2: Adding expressivity to the token sequence

the time difference of the original MIDI note onset t_s from the quantized token onset t_k , illustrated in Figure 2a. The DELTA tokens encode the offset in increments of 1/160th of a sixteenth-note triplet. We consider 80 additional tokens because the maximal absolute time difference is half of a sixteenth-note triplet, and we use a DELTA_-1 token when this time difference is negative. This resolution allows for a small addition to the vocabulary, yet is enough to encode 99% of the expressive tracks of GigaMIDI. The use of expressive tokens is illustrated in Figure 2b.

4 Controlling Music Generation

The premise behind attribute controls is that given a musical excerpt x , and a measurable musical attribute a for which we can compute a categorical or ordinal value from x (i.e. $a(x)$), the model will learn the conditional relationship between tokens representing $a(x)$ and the musical material on a track, provided these tokens precede the musical material. Practically, this is accomplished by inserting one or more CONTROL tokens which specify the level of a particular musical attribute $a(x)$ immediately after the INSTRUMENT token (see Figure 1), and before the tokens which specify the musical material. As a result, our approach is most certainly not limited to the specific musical attributes we discuss below, and can be applied to control any musical feature that can be measured. We employ three approaches to control musical attributes of the generated material: categorical controls, which condition generation on one of n different categories; value controls, which condition generation on one of n different ordinal values; and range controls, which condition the system to generate music wherein a particular musical attribute has values that fall within a specified range.

Instrument control is an example of a categorical control, as one of 128 different instrument types can be selected. We use a value control for note density, however, the density categories are determined relative to the instrument type, as average note density varies significantly between instruments. For each of the 128 general MIDI instruments, we calculate the number of note onsets for each bar in the dataset. We divide the distribution for each instrument σ into 10 regions with the range $[P_{10i}(\sigma), P_{10(i+1)}(\sigma)]$ for $0 \leq i < 10$, where $P_n(\sigma)$ denotes the n^{th} percentile of the distribution σ . Each region corresponds to a different note density level for a particular instrument.

We choose to apply range controls to note duration and polyphony. Each note duration (d) is quantized as $\lfloor \log_2(d) \rfloor$.

The quantization process groups note durations into 5 different bins $[\frac{1}{32}, \frac{1}{16})$, $[\frac{1}{16}, \frac{1}{8})$, $[\frac{1}{8}, \frac{1}{4})$, $[\frac{1}{4}, \frac{1}{2})$, and $[\frac{1}{2}, 1)$, which we will refer to as note duration levels. Then the 15th and 85th percentiles of a distribution containing all note duration levels within a track are used to condition generation. Polyphony levels follow a similar approach. The number of notes simultaneously sounding (i.e. polyphony level) at each timestep is calculated (a timestep is one 16th note triplet). Then we use the 15th and 85th percentiles of a distribution containing all polyphony levels within a track. For both these controls, we use two tokens, one to specify the lower bound and another for the upper bound. Admittedly, this is fuzzy range control, as strict range control would typically use the smallest and largest values in the distribution (0th and 100th percentiles respectively). We elected to use the 15th and 85th percentiles in order to mitigate the effect of outliers within the distribution, decreasing the probability of exposing the model to ranges in which values are heavily skewed to one side of the range.

5 Training MIDI-GPT

We use the new GigaMIDI (Lee et al. 2024) dataset, which builds on the MetaMIDI dataset (Ens and Pasquier 2021), to train with a split of: $p_{train} = 80\%$, $p_{valid} = 10\%$, and $p_{test} = 10\%$. Our model is built on the GPT2 architecture (Radford et al. 2019), implemented using the HuggingFace Transformers library (Wolf et al. 2020). The configuration of this model includes 8 attention heads and 6 layers, utilizing an embedding size of 512 and an attention window encompassing 2048 tokens. This results in approximately 20 million parameters.

For each batch, we pick 32 random MIDI files (batch size) from the respective split of the dataset (train, test, valid) and pick random 4 or 8-bar multi-track segments from each MIDI file. For a segment with n tracks, we pick k tracks randomly selecting a value for k on the range $[2, \min(n, 12)]$. With 75% probability, we do bar infilling on a segment where we mask up to 75% of the bars. The number of bars is selected uniformly from values on the range $[0, \lfloor n_{tracks} * n_{bars} * 0.75 \rfloor]$. Then, we randomly transpose the musical pitches (except for the drum track, of course) with a value for the range $[-6, 5]$. Each time we select a n -bar segment during training, we randomly order the tracks so that the model learns each possible conditional ordering between different types of tracks. The model is trained to predict bar, track, and instrument tokens. As a result, when generating a new track, the model can select a sensible instrument to accompany the pre-existing tracks, thus learning instrumentation.

We train with the Adam optimizer, a learning rate of 10^{-4} , without dropout. Training to convergence typically takes 2-3 days using 4 V100 GPUs. We pick the model with the best validation loss.

6 Sampling with MIDI-GPT

To achieve syntactically valid outputs from the system (with respect to the tokenization used), we incorporate specific masking constraints. More precisely, we mask select tokens

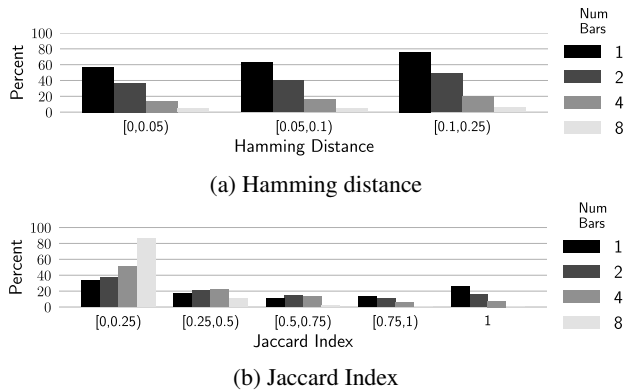


Figure 3: The percentage of generated excerpts (g_i) with a Hamming distance (resp. Jaccard Index $\mathcal{J}(o_i, g_i)$) between any excerpt o_i from the training dataset and g_i on the range [a,b]. A Hamming distance (resp. Jaccard Index) of distance 0 (resp. of 1) indicates two excerpts are identical, while 1 (resp. 0) indicates they are very different.

during various stages of the model’s inference process to preserve the sequence necessary for encoding, decoding, and prediction tasks. For example, a scenario in which a `BAR_END` token appears after a `NOTE_ON` token is not feasible, given that the `DURATION` and `NOTE_POSITION` have yet to be determined. In such instances, we mask the `BAR_END` to prevent its sampling. Subsequently, we sample among the remaining unmasked tokens. This rule-based sampling approach ensures that the model maintains a logical structure throughout its operations.

7 Release, Evaluation, and Applications

MIDI-GPT has been released¹ and is seeing real-world usage in several contexts, which directly supports our assertion that MIDI-GPT is a practical model for computer-assisted composition. There are ongoing collaborations for the integration of MIDI-GPT into synthesizers, like the OP-Z by Teenage Engineering², and game music composition software by Elias. MIDI-GPT has been integrated into the Cubase³ digital audio workstation, the Calliope⁴ web application, and an Ableton⁵ plugin has been developed. MIDI-GPT has been used to compose music⁶, including two entries to the 2022 and 2023 AI Song contests, four albums (two by Philip Tremble, one by Monobor, and one by a selection of American and Canadian composers). It has been used to compose adaptive music for games (Plut et al. 2022), and a yearly artistic residency⁷ with French composers is

¹<https://www.metacreation.net/projects/mmm> links to models and various examples of generations.

²<https://www.metacreation.net/projects/opz-mmm>

³www.metacreation.net/projects/mmm-cubase

⁴<https://www.metacreation.net/projects/calliope>

⁵<https://www.metacreation.net/projects/mmm4live>

⁶<https://www.metacreation.net/projects/mmm-music>

⁷<https://vancouver.consulfrance.org/Artificial-Muse-residency-Vancouver>

ongoing.

A user study (Bougueng Tchemeube et al. 2023) was conducted to evaluate the integration of a previous version of MIDI-GPT into a popular digital audio workstation. The study measured usability, user experience, and technology acceptance for two groups of experienced composers: hobbyists and professionals with convincing results. Since we have already conducted a comprehensive user study, we do not repeat a listening study here. Instead, our experiments are designed to address other aspects that impact the usability of the system in real-world settings.

In the following, our evaluation of MIDI-GPT gauges the performance of the system by addressing the following research questions:

1. Originality: Does MIDI-GPT generate original variations or simply duplicate material from the dataset?
2. Stylistic Similarity: Does MIDI-GPT generate musical material that is stylistically similar to the dataset (i.e., well-formed music)?
3. Attribute Controls: How effective are density level, polyphony range, and note duration range controls?

7.1 Evaluating the Originality of Generated Material

Intra-Dataset Originality It is increasingly important to quantify the frequency with which a generative system is producing musical material that is nearly identical to the training dataset, given potential legal issues that may arise when these systems are deployed into the real world, and the difficulty of guaranteeing that a generative system does not engage in this type of behavior (Papadopoulos, Roy, and Pachet 2014). To accomplish this, we represent musical material in each track as a piano roll and use hamming distance to calculate the distance between two piano rolls. A piano roll is a $T \times 128$ boolean matrix specifying when particular pitches are sounding, where T is the number of time-steps. Note that when we calculate Hamming distance between two piano rolls, we normalize the distance by the size of the piano roll. Therefore, the distance between maximally different piano rolls would be 1.

Since the dataset contains hundreds of thousands of unique MIDI files, we are faced with a time complexity issue, and must employ some heuristics to speed this process up. First, rather than searching nearly identical n -bar piano rolls, we search using single-bar piano rolls, and aggregate the results of n search processes. Note that this means that if $n - 1$ of the bars have a match in the dataset, but one of the bars does not, the n -bar excerpt will not be considered to have a match in the dataset. However, since we are interested in identifying nearly identical matches, this is unlikely to cause much of an issue. To filter out highly dissimilar candidate matches efficiently, we compute the Hamming distance between compressed piano rolls first. Given a 48×128 piano roll x that represents a single 4/4 bar of musical material, we discard notes outside the range [21, 109) and take the maximum value over each consecutive set of 6 time-steps (equivalent to one 1/8 note) on the first axis, producing a 8×88 matrix x . We calculate the Hamming distance between the

compressed piano rolls, discarding any candidate matches that have a distance greater than 0.25, and then compute Hamming distance on the full-sized piano rolls for the remaining candidate matches. Even with these optimizations, the search is executed in parallel on a 32-core machine and takes an average of 83 seconds to complete a search for a single 4-bar excerpt. In the worst case, it can take up to an hour for a single query. Although we would have preferred to use Jaccard index rather than Hamming distance, as we do in the next sub-section, the nature of the heuristics employed to increase computation speed prohibited this option.

We compute 100 trials where we randomly select a 4-track 8-bar musical segment from the test split of the dataset, blank out n consecutive bars on a single track, and generate (i.e. infill) a new set of bars (g_i). Given a Hamming distance threshold, we determine if g_i is nearly identical to any n -bar excerpt in the training split of the dataset, using the method described above. In Figure 3a we present the percentage of trials for which the Hamming distance between any excerpts in the training dataset and g_i is on the specified range. Unsurprisingly, as the number of bars increases, the percentage of instances where MIDI-GPT duplicates the training data decreases significantly. This correlation was expected as shorter generations are more constrained by the surrounding musical content. There are not that many one-track, one-bar musical excerpts that are not already in the data set.

Infilling Originality We ought to also measure the data reproduction for the infilling task, which may occur when the model predicts the exact segment that a user wants to infill, resulting in no change and inevitable frustration from the user. To measure the frequency with which this occurs, we randomly select a 4 track 8 bar musical segment from the test split of the dataset, blank out n consecutive bars on a single track (o_i), and generate a new set of n bars (g_i) to replace (o_i). Then, we measure the Jaccard index between piano roll representations of o_i and g_i . We repeat this process 250 times for each number of bars ($n = 1, 2, 4, 8$) and report the results in Figure 3b. On the whole, as the number of bars increases, the frequency with which the original material is duplicated decreases. Taken collectively, the results in this section indicate that MIDI-GPT can reliably produce original variations when generating 4 or more bars. In practice, when deployed in products, it always does as we actually test that the material generated is different from what is being replaced, and regenerate otherwise. We also re-generate tracks or bar infilling resulting in silence, which the user never intends.

7.2 Quantifying Stylistic Similarity

It is also important that the variations generated by the system are stylistically similar to the dataset. To be clear, we define musical style as the stylistic characteristics delineated by a set of musical data. As a result, when we claim to measure stylistic similarity to the training data, we are measuring similarity to the style that is delineated by this set of data. Consequently, we avoid having to make subjective decisions about what constitutes a particular musical style, while maintaining an evaluation framework that is generic




Figure 4: The percentage of trials where $S_{O_{25}, G_{25}}^{C_{50}}(\hat{O}_{25}^*, C_{50}^*) \leq S_{O_{25}, G_{25}}^{C_{50}}(\hat{G}_{25}^*, C_{50}^*)$. Hatching indicates that the binomial test was insignificant, indicating that \hat{O}^* is not more similar to C than \hat{G}^* .




Figure 5: The percentage for each absolute difference between requested and actual note density.

enough to handle any arbitrary set of data.

We use StyleRank (Ens and Pasquier 2019) to measure the stylistic similarity of generated material. StyleRank is designed to measure the similarity of two or more groups of musical excerpts ($\mathcal{G}_1, \dots, \mathcal{G}_k$) relative to a style delineated by a collection of ground truth musical excerpts (\mathcal{C}). Each musical excerpt is represented using a set of features, described in detail in the original paper, and a Random Forest classifier is trained to discriminate between $\mathcal{G}_1, \dots, \mathcal{G}_k$ and \mathcal{C} . Using an embedding space constructed from the trained Random Forest classifier, the average similarity between \mathcal{G}_i and \mathcal{C} can be computed for each i . In what follows, let $S_{\mathcal{G}_1, \dots, \mathcal{G}_k}^{\mathcal{C}}(a, b)$ denote the median similarity between a and b , calculated using a StyleRank instance trained on $\mathcal{G}_1, \dots, \mathcal{G}_k$ and \mathcal{C} .

For this experiment, we use the same musical excerpts from Section 7.1 ($\mathcal{O} = \{o_1, \dots, o_{250}\}, \mathcal{G} = \{g_1, \dots, g_{250}\}$), however, we remove each pair (o_i, g_i) where $\mathcal{J}(o_i, g_i) \geq 0.75$, producing $\hat{\mathcal{O}}$ and $\hat{\mathcal{G}}$. This ensures that we do not bias our measurements by including generated material that is nearly identical to the original preexisting material (o_i) from the dataset. We also assemble a set of 1000 n -bar segments (\mathcal{C}) from the dataset. For each trial, we compute $S_{\hat{\mathcal{O}}_{25}, \hat{\mathcal{G}}_{25}}^{C_{50}}(\hat{\mathcal{O}}_{25}^*, C_{50}^*) \leq S_{\hat{\mathcal{O}}_{25}, \hat{\mathcal{G}}_{25}}^{C_{50}}(\hat{\mathcal{G}}_{25}^*, C_{50}^*)$, where X_n^* denotes a subset of X containing n elements, which are selected randomly for each trial. In other words, for each trial, we determine if the median similarity between two subsets of the corpus is less than or equal to the median similarity between a subset of the corpus and a set of generated excerpts. We collect the results for 100 trials and compute a binomial test. If there is no significant difference between the count of trials for which the condition is true and the count of trials for which the condition is false, we can conclude that there is not a significant difference between the generated material and the corpus with respect to the similarity metric we are using. We report the results of this test using different numbers of bars and temperatures in Figure 4.

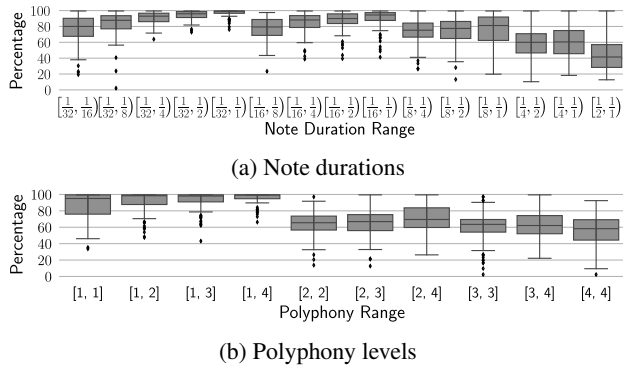


Figure 6: The percentage of note durations (a) and polyphony levels (b) within the range shown for 100 trials.

The results indicate that when generating with a temperature of 1.0, infilled generations are equivalent to the original preexisting material in terms of musical style (as quantified by StyleRank). Our results also show that when the temperature is greater than 1.0, the generated material is more frequently considered less similar to the dataset (\mathcal{C}) than $\hat{\mathcal{O}}$, an effect that increases along with the number of generated bars. This demonstrates that our measurement instrument is sensitive enough to detect small increases in the entropy of the music generated by the model, which are the byproduct of slightly changing the temperature.

This result does not only serve as an analytical evaluation of the model (as human evaluations are covered elsewhere (Bougueng Tchemeube et al. 2023)), but future work is to replicate these when conditioning on a set of files representing a given musical style. MIDI-GPT includes a categorical style control based on the style metadata extracted from the MetaMIDI dataset (a subset of GigaMIDI) and aligned to the MusicMap (Crauwels 2022) ontology (Disco, Rock, Jazz, ...).

7.3 Evaluating the Effectiveness of Attribute Controls

MIDI-GPT allows the user to condition generation not only on the existing musical content at the time of generation, but also on various control attributes such as the instrument, musical style, note density, polyphony level, and note duration. To evaluate how effective these control mechanisms are, we focus on the three last ones for brevity. We conduct 100 trials where we generate 8-bar segments from scratch using a particular attribute control method, and measure the difference between the anticipated outcome and the actual outcome. For note density control, we measure the absolute difference between the density level the generation was conditioned on and the density level of the generated material. For polyphony level and note duration, we compute the distribution of values (either polyphony level or note duration level) from the generated material, and count the percentage of values that fall within the specified range. If attribute control is successful, we would expect at least 70% of the values to be within this range, as we used the 15th and 85th

percentiles while training.

The results for note density, shown in Figure 5, demonstrate that the majority of times the absolute difference between the anticipated and actual density level is most of ten 0, and rarely exceeds 1. This indicates that this control method is effective. The results for Note Duration, shown in Figure 6a demonstrate that this attribute control method is quite effective, as the median outcome (in terms of percentage of note durations within the specified range) is at or above 70% in all cases except for $[\frac{1}{4}, \frac{1}{2})$, $[\frac{1}{4}, \frac{1}{1})$ and $[\frac{1}{2}, \frac{1}{1})$. In contrast, the Polyphony Level control is less effective, with the median outcome lying below the 70% threshold in many cases. Calculating polyphony level at a single time-step is inherently more difficult than note duration, as the former requires knowledge of where multiple notes start and end, while the latter only requires knowing where one note starts and ends. This difference in difficulty seems to be reflected in the results, as MIDI-GPT is better at controlling note duration than polyphony.

In addition to these “soft controls”, we can also implement hard controls through rule-based sampling. This involves some bookkeeping. For example, for practical reasons, we also provide a hard polyphony limit l_{poly} . In places where it would otherwise be valid for a note to be inserted in the token sequence, we first check that the size of the set of currently sounding pitches n_{pitch} satisfies $n_{pitch} < l_{poly}$. If $n_{pitch} = l_{poly}$, we mask note tokens.

8 Conclusion

We present MIDI-GPT, a style-agnostic generative system released as an Open RAIL-M licenced MMM model (Ens and Pasquier 2020). MIDI-GPT builds on an alternative approach to representing musical material, resulting in increased control over the generated output. We provided experimental evidence demonstrating the effectiveness of the system and outlined several ongoing real-world applications. The system runs on most personal computer with an attention window of 2048 tokens, corresponding to 8-16 bars depending on the number of tracks and their density. However, using an auto-regressive approach, and sliding windows, longer parts can be generated by repeatedly conditioning on portions of preexisting along with newly generated material. Future work involves: optimizing the model for real-time generation in musical agents, training larger models to expand the model attention window and attend to larger musical structures, expanding the set of attribute controls, and continuing integration of MIDI-GPT into real-world products and practices.

Ethical Statement

While MIDI-GPT is style-agnostic, and does include more musical styles and content than any human can know, it still inherits from the dataset’s bias, which only encompasses the musical styles afforded by the MIDI notation and available online. Arguably, there are still many musical styles under-represented (mostly non-western styles) or even absent from the dataset (i.e. not representable or available in MIDI) and therefore the model (although some suspecting users have

noticed its generalization capabilities). Conversely, some musical styles are over-represented (e.g., pop). Further work is needed to qualify and quantify these biases.

Regarding copyright and intellectual property, MetaMIDI and GigaMIDI were acquired under the Fair Dealing law of Canada, which limits their use to research and non-commercial use. We abide by these limitations.

Regarding MIDI-GPT: at the time of writing, it is unclear what the legal status of such a model is, as it does not contain the data itself, and produces original content for any non-over-constrained request (as shown above). So far, we restricted its use to non-commercial use. It is either released as-free-for use (Calliope), used for research purposes with collaborating companies, or used for research-creation purposes with selected artists. The ethical or legal implications of current creative use by other artists, and the potential use of the released model (under an Open RAIL-M licence) for commercial purposes by other parties does not rest with the authors of this research. We thus decline any responsibility for misuse.

Acknowledgments

We would like to acknowledge the support of the National Science and Engineering Research Council (NSERC), the Social Sciences and Humanities Research Council (SSHRC), the Canada Council for the Arts (CCA), and Steinberg Media Technology. We also thank Griffin Page and the anonymous reviewers for their edits and feedback.

References

- Artist Rights Alliance. 2024. 200+ Artists Urge Tech Platforms: Stop Devaluing Music. <https://t.ly/-LbHx>. Accessed: 2024-12-19.
- Bougueng Tchameube, R.; Ens, J.; Plut, C.; Pasquier, P.; Safi, M.; Grabit, Y.; and Rolland, J.-B. 2023. Evaluating Human-AI Interaction via Usability, User Experience and Acceptance Measures for MMM-C: A Creative AI System for Music Composition. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, IJCAI-23*, 5769–5778.
- Briot, J.-P.; Hadjeres, G.; and Pachet, F.-D. 2019. *Deep Learning Techniques for Music Generation*. Computational Synthesis and Creative Systems. Springer International Publishing.
- Cancino-Chacón, C. E.; and Grachten, M. 2016. The Basis Mixer: A Computational Romantic Pianist. In *Proceedings of the Late Breaking/ Demo Session, 17th International Society for Music Information Retrieval Conference (ISMIR)*.
- Chang, C.-J.; Lee, C.-Y.; and Yang, Y.-H. 2021. Variable-Length Music Score Infilling via XLNet and Musically Specialized Positional Encoding. In *Proceedings of the 22nd ISMIR Conference*, 97–104.
- Chen, K.; Wang, C.; Berg-Kirkpatrick, T.; and Dubnov, S. 2020. Music SketchNet: Controllable music generation via factorized representations of pitch and rhythm. In *Proceedings of the 21st ISMIR Conference*, 77–84.
- Chi, W.; Kumar, P.; Yaddanapudi, S.; Rahul, S.; and Isik, U. 2020. Generating music with a self-correcting non-chronological autoregressive model. In *Proceedings of the 21st ISMIR Conference*, 893–900.
- Collins, T.; and Barthet, M. 2023. Expressor: A Transformer Model for Expressive MIDI Performance. In *Proceedings of the 16th International Symposium on Computer Music Multidisciplinary Research*, 740–743. Zenodo.
- Crauwels, K. 2022. MusicMap. <https://musicmap.info/>. Accessed: 2024-12-19.
- Donahue, C.; Mao, H. H.; Li, Y. E.; Cottrell, G.; and McAuley, J. 2019. LakhNES: Improving Multi-instrumental Music Generation with Cross-domain Pre-training. In *Proceedings of the 20th ISMIR Conference*, 685–692.
- Dong, H.-W.; Chen, K.; Dubnov, S.; McAuley, J.; and Berg-Kirkpatrick, T. 2023. Multitrack Music Transformer. In *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 1–5.
- Ens, J.; and Pasquier, P. 2019. Quantifying Musical Style: Ranking Symbolic Music based on Similarity to a Style. In *Proceedings of the 20th ISMIR Conference*, 870–877.
- Ens, J.; and Pasquier, P. 2020. MMM : Exploring Conditional Multi-Track Music Generation with the Transformer. *CoRR*, abs/2008.06048.
- Ens, J.; and Pasquier, P. 2021. Building the MetaMIDI Dataset: Linking Symbolic and Audio Musical Data. In *Proceedings of the 22nd ISMIR Conference*, 182–188.
- Fradet, N.; Briot, J.-P.; Chhel, F.; El Fallah Seghrouchni, A.; and Gutowski, N. 2021. MidiTok: A Python package for MIDI file tokenization. In *Extended Abstracts for the Late-Breaking Demo Session of the 22nd ISMIR Conference*.
- Fradet, N.; Gutowski, N.; Chhel, F.; and Briot, J.-P. 2023a. Byte Pair Encoding for Symbolic Music. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, 2001–2020.
- Fradet, N.; Gutowski, N.; Chhel, F.; and Briot, J.-P. 2023b. Impact of time and note duration tokenizations on deep learning symbolic music modeling. In *Proceedings of the 24th ISMIR Conference*.
- Gillick, J.; Roberts, A.; Engel, J.; Eck, D.; and Bamman, D. 2019. Learning to Groove with Inverse Sequence Transformations. In Chaudhuri, K.; and Salakhutdinov, R., eds., *Proceedings of the 36th International Conference on Machine Learning*, 2269–2279.
- Guo, R.; Simpson, I.; Kiefer, C.; Magnusson, T.; and Herremans, D. 2022. MusIAC: An extensible generative framework for Music Infilling Applications with multi-level Control. In *Proceeding of the Artificial Intelligence in Music, Sound, Art and Design: 11th International Conference, EvoMUSART 2022, Held as Part of EvoStar 2022, Madrid, Spain, April 20–22, 2022*, 341–356. Springer.
- Haki, B.; Nieto, M.; Pelinski, T.; and Jordà, S. 2022. Real-Time Drum Accompaniment Using Transformer Architecture. In *Proceedings of the 3rd Conference on AI Music Creativity. AIMC*.

- Hawthorne, C.; Huang, A.; Ippolito, D.; and DouglasEck. 2018. Transformer-NADE for Piano Performances. In *NeurIPS 2018 Workshop on Machine Learning for Creativity and Design*.
- Hong, Y.; Hwang, U.; Yoo, J.; and Yoon, S. 2019. How Generative Adversarial Networks and Their Variants Work: An Overview. *ACM Computing Surveys*, 52(1).
- Hsiao, W.-Y.; Liu, J.-Y.; Yeh, Y.-C.; and Yang, Y.-H. 2021. Compound Word Transformer: Learning to Compose Full-Song Music over Dynamic Directed Hypergraphs. *Proceedings of the 21st AAAI Conference on Artificial Intelligence*, 178–186.
- Huang, C. A.; Vaswani, A.; Uszkoreit, J.; Simon, I.; Hawthorne, C.; Shazeer, N.; Dai, A. M.; Hoffman, M. D.; Dinculescu, M.; and Eck, D. 2019. Music Transformer: Generating Music with Long-Term Structure. In *7th International Conference on Learning Representations (ICLR)*.
- Huang, C.-Z. A.; Cooijmans, T.; Roberts, A.; Courville, A. C.; and Eck, D. 2018. Counterpoint by Convolution. In *Proceedings of the 18th ISMIR Conference*, 211–218.
- Huang, Y.-S.; and Yang, Y.-H. 2020. Pop Music Transformer: Beat-Based Modeling and Generation of Expressive Pop Piano Compositions. In *Proceedings of the 28th ACM International Conference on Multimedia*, MM 2020, 1180–1188.
- Lee, K. J. M.; Ens, J.; Adkins, S.; Sarmiento, P.; Barthet, M.; and Pasquier, P. 2024. The GigaMIDI Dataset with Features for Expressive Music Performance Detection. Forthcoming.
- Liu, J.; Dong, Y.; Cheng, Z.; Zhang, X.; Li, X.; Yu, F.; and Sun, M. 2022. Symphony Generation with Permutation Invariant Language Model. In *Proceedings of the 23rd ISMIR Conference*, 551–558.
- Maezawa, A.; Yamamoto, K.; and Fujishima, T. 2019. Rendering Music Performance With Interpretation Variations Using Conditional Variational RNN. In *Proceedings of the 20th ISMIR Conference*, 855–861.
- Malik, I.; and Ek, C. H. 2017. Neural Translation of Musical Style. In *NeurIPS 2017 Workshop on Machine Learning for Creativity and Design*.
- Mittal, G.; Engel, J.; Hawthorne, C.; and Simon, I. 2021. Symbolic Music Generation with Diffusion Models. In *Proceedings of the 22nd ISMIR Conference*, 468–475.
- Nuttall, T.; Haki, B.; and Jorda, S. 2021. Transformer Neural Networks for Automated Rhythm Generation. In *Proceedings of the International Conference on New Interfaces for Musical Expression*.
- Oore, S.; Simon, I.; Dieleman, S.; Eck, D.; and Simonyan, K. 2020. This time with feeling: learning expressive musical performance. *Nueral Computing and Applications*, 32(4): 955–967.
- Papadopoulos, A.; Roy, P.; and Pachet, F. 2014. Avoiding Plagiarism in Markov Sequence Generation. In *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence*, 2731–2737.
- Pati, A.; Lerch, A.; and Hadjeres, G. 2019. Learning to Traverse Latent Spaces for Musical Score Inpainting. In *Proc. of the 20th ISMIR Conference*, 343–351.
- Payne, C. 2019. MuseNet. openai.com/blog/musenet. Accessed: 2024-12-19.
- Plut, C.; Pasquier, P.; Ens, J.; and Tchemeube, R. 2022. PreGLAM-MMM: Application and evaluation of affective adaptive generative music in video games. In *Proceedings of the 17th International Conference on the Foundations of Digital Games*, FDG 2022.
- Radford, A.; Wu, J.; Child, R.; Luan, D.; Amodei, D.; and Sutskever, I. 2019. Language models are unsupervised multitask learners. https://cdn.openai.com/better-language-models/language_models_are_unsupervised_multitask_learners.pdf. Accessed: 2024-12-19.
- Ren, Y.; He, J.; Tan, X.; Qin, T.; Zhao, Z.; and Liu, T.-Y. 2020. PopMAG: Pop Music Accompaniment Generation. In *Proceedings of the 28th ACM International Conference on Multimedia*, 1198–1206.
- Roberts, A.; Engel, J.; Raffel, C.; Hawthorne, C.; and Eck, D. 2018. A Hierarchical Latent Vector Model for Learning Long-Term Structure in Music. In Dy, J.; and Krause, A., eds., *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of PMLR, 4364–4373.
- Sturm, B. L.; Ben-Tal, O.; Monaghan, Ú.; Collins, N.; Herremans, D.; Chew, E.; Hadjeres, G.; Deruty, E.; and Pachet, F. 2019. Machine learning research that matters for music creation: A case study. *Journal of New Music Research*, 48(1): 36–55.
- TAN, H. H.; and Herremans, D. 2020. Music FaderNets: Controllable music generation based on high-Level features via low-level feature modelling. In *Proceedings of the 21st ISMIR Conference*, 109–116.
- von Rütte, D.; Biggio, L.; Kilcher, Y.; and Hofmann, T. 2023. FIGARO: Controllable Music Generation using Learned and Expert Features. In *11th ICLR*.
- Wang, Z.; Wang, D.; Zhang, Y.; and Xia, G. 2020. Learning interpretable representation for controllable polyphonic music generation. In *Proceedings of the 21st ISMIR Conference*, 662–669.
- Wang, Z.; and Xia, G. 2021. MuseBERT: Pre-training Music Representation for Music Understanding and Controllable Generation. In *Proceedings of the 22nd ISMIR Conference*, 722–729.
- Wolf, T.; Debut, L.; Sanh, V.; Chaumond, J.; Delangue, C.; Moi, A.; Cistac, P.; Rault, T.; Louf, R.; Funtowicz, M.; Davison, J.; Shleifer, S.; von Platen, P.; Ma, C.; Jernite, Y.; Plu, J.; Xu, C.; Scao, T. L.; Gugger, S.; Drame, M.; Lhoest, Q.; and Rush, A. M. 2020. Transformers: State-of-the-Art Natural Language Processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, 38–45.
- Wu, S.-L.; and Yang, Y.-H. 2020. The jazz transformer on the front line: Exploring the shortcomings of AI-composed music through quantitative measures. In *Proceedings of the 21st ISMIR Conference*, 142–149.
- Wu, Y.; Manilow, E.; Deng, Y.; Swavely, R.; Kastner, K.; Cooijmans, T.; Courville, A.; Huang, C.-Z. A.; and Engel, J. 2022. MIDI-DDSP: Detailed Control of Musical Performance via Hierarchical Modeling. In *10th ICLR*.