

# MMT-BERT: CHORD-AWARE SYMBOLIC MUSIC GENERATION BASED ON MULTITRACK MUSIC TRANSFORMER AND MUSICBERT

Jinlong Zhu<sup>1</sup>      Keigo Sakurai<sup>1</sup>      Ren Togo<sup>2</sup>  
Takahiro Ogawa<sup>2</sup>      Miki Haseyama<sup>2</sup>

<sup>1</sup> Graduate School of Information Science and Technology, Hokkaido University, Japan

<sup>2</sup> Faculty of Information Science and Technology, Hokkaido University, Japan

jinlong@lmd.ist.hokudai.ac.jp, mhaseyama@lmd.ist.hokudai.ac.jp

## ABSTRACT

We propose a novel symbolic music representation and Generative Adversarial Network (GAN) framework specially designed for symbolic multitrack music generation. The main theme of symbolic music generation primarily encompasses the preprocessing of music data and the implementation of a deep learning framework. Current techniques dedicated to symbolic music generation generally encounter two significant challenges: training data’s lack of information about chords and scales and the requirement of specially designed model architecture adapted to the unique format of symbolic music representation. In this paper, we solve the above problems by introducing new symbolic music representation with MusicLang chord analysis model. We propose our MMT-BERT architecture adapting to the representation. To build a robust multitrack music generator, we fine-tune a pre-trained MusicBERT model to serve as the discriminator, and incorporate relativistic standard loss. This approach, supported by the in-depth understanding of symbolic music encoded within MusicBERT, fortifies the consonance and humanity of music generated by our method. Experimental results demonstrate the effectiveness of our approach which strictly follows the state-of-the-art methods.

## 1. INTRODUCTION

Music plays an indispensable role in our daily lives, and there is a significant demand for creating new musical contents. Automatic music generation is one of the most intriguing tasks in bringing new music experiences to consumers [1]. The earliest studies in the 1950s focused on a combination of music theory and Markov-chains-based probabilistic models, and realized randomly creating music parts and combining them into a synthesis [2]. Contemporary studies have achieved higher quality and faster music generation by utilizing advanced neural networks

such as Generative Adversarial Networks (GANs), Transformer, and diffusion models [3–5]. Despite significant advancements, previous methods continue to suffer from challenges such as insufficient data extraction and unstable training trajectories. Consequently, there is room for new approaches for more effective music representations and more robust deep learning architectures.

In particular, chords are crucial for conveying emotional and humanistic expressions in music, yet few methods take chords into account in symbolic music representation. Consequently, previous methods are deprived of indispensable information about chords and scales. This lack results in the generation of music that exhibits a diminished degree of humanity. Therefore, previous methods for music generation face limitations in their ability to produce human-like and high quality expressions [6]. A feasible solution to overcome this difficulty is the integration of a chord analysis model [7]. Chord analysis model aids in the extraction of chord data from raw audio, fostering a novel representation method that encompasses chord information [8–11]. With the aid of state-of-the-art chord analysis models, we can generate more harmonious and structured music with more regular chord progressions by automatically extracting and encoding chords from raw audio files. Therefore, it is expected that adopting chord analysis models in creating new symbolic representations of music will enable the generation of music that is closer to human composition.

Another problem arises from the ever-changing format of symbolic music representation, which makes designing the model’s architecture that fits symbolic music generation to be another challenge. GANs are widely applied in the symbolic music generation field because the addition of a discriminator obviously strengthens the fidelity of the overall generative model [12–17]. The performance of GANs is deeply influenced by the architecture of the generator and discriminator. Previous studies have demonstrated the effectiveness of transformer-based generators [6, 7, 18–23]. Whereas, the architecture of the discriminator has been extensively discussed in recent years. Some methods [12, 14, 16, 17] involve constructing a discriminator based on CNN or Transformer, while others [15] utilize pre-trained models adapted to their tasks. Compared to hand-crafted discriminators, using pre-trained models often achieves a fairly good result



© J. Zhu, K. Sakurai, R. Togo, T. Ogawa and M. Haseyama. Licensed under a Creative Commons Attribution 4.0 International License (CC BY 4.0). **Attribution:** J. Zhu, K. Sakurai, R. Togo, T. Ogawa and M. Haseyama, “MMT-BERT: Chord-aware Symbolic Music Generation Based on Multitrack Music Transformer and MusicBERT”, in *Proc. of the 25th Int. Society for Music Information Retrieval Conf.*, San Francisco, United States, 2024.

because pre-trained models are already trained on large and diverse datasets. Therefore, the application of pre-trained models will allow the GAN to leverage their learning and knowledge, ensuring the training efficiency and stability.

However, there are few choices of pre-trained models designed for symbolic music representation that can be used as the discriminator considering the input format and length limitation. We solve this problem by applying BERT-based scores, which are well correlated with human ranking and can jointly measure quality and diversity [24,25]. Since BERT is trained using a self-supervised loss on bidirectional contexts of all attention layers, it can effectively extract representations [26–28]. Muhamed *et al.* employ a pre-trained Span-BERT model and achieve considerable results on harmonic choices and overall music quality, showing that pre-trained BERT-based models outperform CNN-based discriminator [15]. Hence, employing a pre-trained model as a discriminator can amplify the overall performance of the GAN model.

In this paper, we propose a novel symbolic music generation method using the chord-aware symbolic music representation and MusicBERT-based discriminator. In terms of symbolic musical expression, we introduce the novel symbolic music representation with chord information derived from MusicLang<sup>1</sup>, one of the state-of-the-art chord analysis models. By employing symbolic music representation with chord information, our model can achieve the generation of more human-like music that considers chord progressions. For the model architecture, we employ the Multitrack Music Transformer (MMT) [6] as the generator and fine-tune the MusicBERT [29], a symbolic music understanding model pre-trained in large-scale dataset, as the discriminator. Leveraging the superior comprehension capabilities of MusicBERT, we can improve GAN’s performance, thereby facilitating the creation of higher-quality music. Furthermore, we introduce relativistic standard loss to further optimize the stability and consistency of the training process [30]. The use of Relativistic Standard GAN (RS-GAN) has realized great results in the field of image generation. It enables models to account for the fact that half of the data in a mini-batch is fake, leading to more accurate estimations of data realism [14,31]. Building upon the innovations mentioned above, our model is capable of retrieving substantial information about chords and scales, acquiring knowledge in music theory, and autonomously generating multitrack music of superior quality and enriched with human-like characteristics.

The contributions of this paper are summarized as follows.

- We propose a modified MMT style symbolic music representation including chord and scale information.
- We develop MMT-BERT, an optimized GAN architecture utilizing MMT and MusicBERT, with relativistic standard loss to enhance the stability of the training process and achieve better results.

<sup>1</sup> <https://musiclang.github.io/tokenizer/>

Representation	Multitrack	Instrument control	Compound tokens	Chord awareness
REMI [7]				✓
MMM [21]	✓			
CP [18]			✓	✓
FIGARO [23]	✓			✓
MMT [6]	✓	✓	✓	
MMT-BERT (ours)	✓	✓	✓	✓

**Table 1.** Comparisons of related representations.

## 2. RELATED WORKS

### 2.1 Symbolic Music Representation

To enable computers to properly understand music, research on symbolic music representation has been conducted for many years [32]. Musical Instrument Digital Interface (MIDI) is the most commonly used format for symbolic music representation, containing performance data and control information for musical notes. In the music processing community, many researchers symbolize music with MIDI-like events [33].

Huang *et al.* have proposed REvamped MIDI-derived events (REMI), which adds note duration and bar events, enabling models to generate music with subtle rhythmic repetition [7]. However, the REMI representation often encounters a challenge that the sequence is too long. Building upon the REMI framework, Hsiao *et al.* have proposed Compound Word Transformer (CP) [18]. CP modifies REMI’s approach by transforming one-dimensional sequence tokens into compound words sequence using specific rules. Although this modification significantly shortens the average token sequence length and simplifies the model’s ability to capture musical nuances, CP is hard to generate multitrack music [6]. Dong *et al.* have proposed their multitrack music representation, which represents music with a sequence of sextuple tokens, along with a Transformer-XL-based generation method Multitrack Music Transformer (MMT). This approach utilizes a decoder-only Transformer architecture, adept at processing multi-dimensional inputs and outputs. MMT leverages the advantages of the Transformer to enable the generation of longer multitrack music compositions than previous music generation methods. However, MMT’s representation scheme lacks chord event inclusion, an essential element in musical compositions. In contrast, our symbolic music representation technique builds on the foundation laid by MMT by integrating chord information, enabling our model to produce more harmonically rich compositions.

### 2.2 Generative Adversarial Network-based Music Generation

Previous studies have employed various GANs to realize symbolic music generation [12, 13, 17]. In early states, Dong *et al.* have proposed MuseGAN, a CNN-based GAN architecture, managing to generate multitrack music pieces [16]. However, CNN-based GANs often suffer from problems such as limited local perception, fixed-size inputs, etc. Muhamed *et al.* solved this problem by introducing their Transformer-GANs model, using a Transformer-

Event type $t_j$	Quintuple token $\mathbf{x}_i^{t_j}$
start-of-song	(0, 0, 0, 0, 0)
instrument	(0, 0, 0, 0, instrument)
start-of-score	(0, 0, 0, 0, 0)
note	(beat, position, pitch, duration, instrument)
chord	(beat, degree, root, mode, extension)
end-of-song	(0, 0, 0, 0, 0)

**Table 2.** The elements of the quintuple token  $\mathbf{x}_i^{t_j}$  for each event type  $t_j$ .

XL-based generator and pre-trained Span-BERT as the discriminator [15]. Transformer-XL introduces the notion of recurrence into the deep self-attention network, enabling the reuse of hidden states from previous segments as memory for the current segment, allowing for the modeling of long-range dependencies. For the discriminator, Span-BERT is utilized to extract sequence embeddings followed by a pooling and linear layer. The bidirectional transformer has a comparable capacity to the transformer-based generator and uses the self-attention mechanism to capture meaningful aspects of the input music sequence. Their research validates the efficacy of employing a Transformer-XL-based generator in conjunction with a BERT-based discriminator [34]. Building on this concept, we developed the MMT-BERT model by utilizing MMT as the generator and Music-BERT as the discriminator.

### 3. METHODOLOGY

#### 3.1 Proposed Symbolic Music Representation

In our approach, we introduce a novel symbolic music representation that incorporates chords. Table 1 shows differences between the conventional approaches and the proposed symbolic music representation. While most conventional representations omit details about chords, we focus on chord information-aware representation to facilitate the process of generating music that more closely resembles humans. First, we extract music data including chords and notes from MIDI files based on MuspyToolkit [35] and MusicLang. During the extraction process, we recognize a chord once per bar, i.e., every four beats. We exclude songs with a time signature other than 4/4, limit the number of chords in a bar to one, and ignore chord changes within a bar since MusicLang only detects chord changes once per bar. Each time MusicLang detects a chord change, it extracts the scale degree, tonality root, tonality mode, chord octave and extension note of the chord. After extracting chord and note information, we encode a piece of music into a sequence of quintuple tokens  $\mathbf{X} = (\mathbf{x}_0, \dots, \mathbf{x}_{N-1})$ , where  $\mathbf{x}_i$  and  $N$  denote the  $i$ -th quintuple token and the total number of quintuple tokens, respectively. Here,  $t$  represents the following event type: {start-of-song, instrument, start-of-score, note, chord, end-of-song}. The meanings of each event type are shown as follows:

- **Start-of-song:** The beginning of the music piece
- **Instrument:** An instrument used in the music piece

Quintuple tokens	Meaning
$\mathbf{x}_0^{\text{start-of-song}}$ : ( 0, 0, 0, 0, 0 )	Start of song
$\mathbf{x}_1^{\text{instrument}}$ : ( 0, 0, 0, 0, 6 )	Instrument : harpsichord
$\mathbf{x}_2^{\text{instrument}}$ : ( 0, 0, 0, 0, 40 )	Instrument : violin
$\mathbf{x}_3^{\text{instrument}}$ : ( 0, 0, 0, 0, 42 )	Instrument : cello
$\mathbf{x}_4^{\text{start-of-score}}$ : ( 0, 0, 0, 0, 0 )	Start of score
$\mathbf{x}_5^{\text{chord}}$ : ( 1, 1, 7, 2, 9 )	Chord: beat=1, degree=1, root=G, mode=minor, extension=9th
$\mathbf{x}_6^{\text{note}}$ : ( 1, 1, 46, 6, 6 )	Note : beat=1, position=1, pitch=A2, duration=6, instrument=harpsichord bar 1
$\mathbf{x}_{12}^{\text{chord}}$ : ( 5, 1, 7, 2, 6 )	Chord: beat=5, degree=1, root=G, mode=minor, extension=6th
$\mathbf{x}_{13}^{\text{note}}$ : ( 5, 7, 48, 5, 42 )	Note : beat=5, position=7, pitch=G2, duration=5, instrument=cello bar 2
$\mathbf{x}_{N-1}^{\text{end-of-song}}$ : ( 0, 0, 0, 0, 0 )	End of song

**Figure 1.** An example of the proposed representation. Compared to the conventional representation, the proposed representation incorporates an additional chord event (highlighted by red blocks) per bar, thereby aiding the model in understanding the relationship between the notes and chords.

- **Start-of-score:** The beginning of a sequence of musical events, including notes and chords
- **Note:** A note characterized by beat, position, pitch, duration, and instrument
- **Chord:** A chord characterized by beat, scale degree, root note, mode, and extension note
- **End-of-song:** The end of the music piece

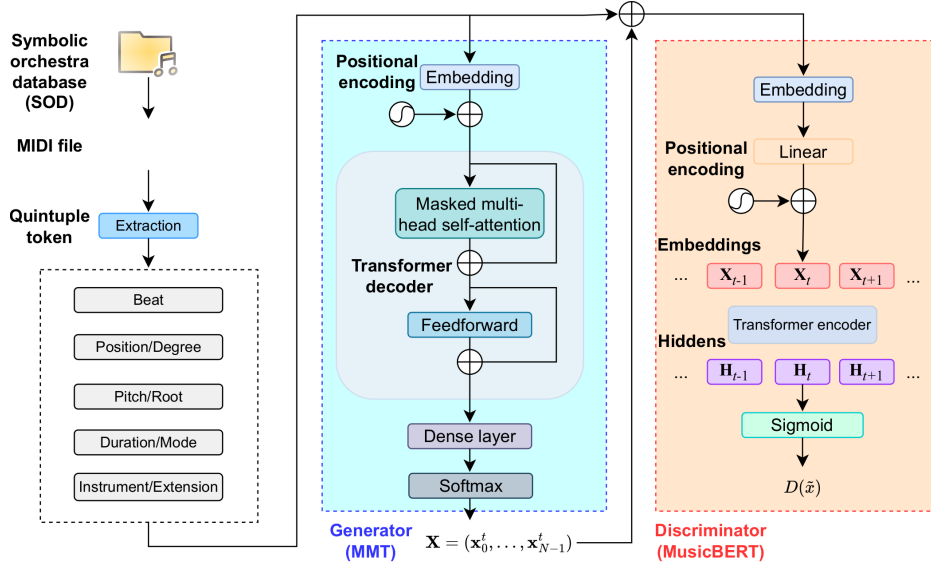
The meaning of each element in the quintuple token  $\mathbf{x}_i^t$  varies depending on the event type  $t$ . The correspondence between the event type  $t$  and the meanings of the quintuple token  $\mathbf{x}_i^t$  is shown in Table 2. Additionally, it is noted that we apply different embeddings for the different features sharing the same axis. A schematic diagram of the proposed representation is illustrated in Figure 1. In this way, we can obtain a symbolic music representation that incorporates chords that is suitable for input into the aforementioned MMT-BERT architecture.

#### 3.2 MMT-BERT Architecture

The fundamental structure of our MMT-BERT architecture is based on a GAN architecture, employing MMT as the generator and MusicBERT as the discriminator. The overview diagram of MMT-BERT is illustrated in Figure 2. The primary concept of GAN is minimizing the loss to enhance the generator’s ability to deceive the discriminator by producing fake music indistinguishable from real music, while simultaneously maximizing the discriminator’s accuracy in distinguishing between real and fake music. Details of the generator and discriminator will be discussed later.

##### 3.2.1 Generator

As the generator, we employ MMT [6], a Transformer-XL-based model that consists solely of decoders. In MMT, elements in the quintuple token  $\mathbf{x}_i$  are individually embedded first, and then concatenated, followed by the addition of



**Figure 2.** The diagram of the MMT-BERT. The generator is built upon Transformer-XL architecture, and the discriminator is built upon MusicBERT. A MIDI file from the dataset is firstly encoded into a sequence of quintuple tokens before fed into the model. Embeddings of the 6 elements are concatenated by a linear layer and converted into a single vector. Then they are fed into the encoder and decoder layers with the addition of position embeddings.

positional embeddings. Subsequently, this combined input is passed through transformer decoder blocks, which are composed of a masked multi-head self-attention layer and a feedforward layer. The output from the decoder blocks is then processed by a dense layer and a softmax layer, resulting in the generation of new music samples. MMT proves advantageous due to its capability to handle multi-dimensional input and output spaces, aligning perfectly with the requirements of symbolic music representation. Significantly, MMT can retain hidden states from previous segments, thereby eliminating the need for recalculating from scratch with each new segment. These retained states function as a memory aid for the current segment, establishing a recurrent connection between segments.

The application of MMT as the generator allows for instrument-controllable multitrack music generation with extended duration and higher training speed. Such a key feature facilitates the modeling of extensive long-range dependencies.

### 3.2.2 Discriminator

As the discriminator, we adopted MusicBERT [29], a large-scale Transformer model developed for symbolic music understanding. MusicBERT consists of a Transformer encoder and utilizes a masked language modeling approach where certain tokens in the input music sequence are masked and then predicted by the model output. The original proposed encoding method, called OctupleMIDI process transforms a symbolic music piece into a sequence of octuple tokens, each containing eight basic elements related to a musical note. In order to make MusicBERT act as a discriminator adapted to the proposed representation mentioned in Sec. 3.1, we refine the input and output format of MusicBERT. Quintuple tokens are converted into a single vector through the concatenation of embeddings and

a linear layer. The resulting vector is combined with position embeddings and provided as input to the Transformer encoder. To predict each of the five tokens within the quintuple, separate softmax layers are added to map the hidden states of the Transformer encoder to the vocabulary sizes of the different element types. MusicBERT’s proficiency in comprehending symbolic music as the discriminator integrates with MMT’s generation process, thereby aiding in the stability of the training process and faster convergence.

### 3.2.3 Relativistic Standard Loss

Inspired by RS-GAN [31], one of the state-of-the-art methods in GANs, we adopt the relativistic standard loss as our objective function. Applying relativistic standard loss prevents the network from becoming overconfident, leading to slower and more careful decisions, allowing the generator more room to adjust its weights and improve the training process [30]. The probability that the given fake data is more realistic than a randomly sampled real data is defined as follows:

$$D(\tilde{x}) = \text{sigmoid}(C(f) - C(r)), \quad (1)$$

where  $C(\cdot)$  denotes a non-transformed layer, and  $\tilde{x}$  denotes real/fake data pairs  $\tilde{x} = (r, f)$ . Hence, the loss function of the generator  $G$  and the discriminator  $D$  are defined as follows:

$$L_G = \mathbb{E}_{(r,f) \sim p(r,f)} [\log(\text{sigmoid}(C(r) - C(f)))] - \sum_i r_i \log f_i, \quad (2)$$

$$L_D = \mathbb{E}_{(r,f) \sim p(r,f)} [\log(\text{sigmoid}(C(f) - C(r)))] \quad (3)$$

where  $r_i$  and  $f_i$  denote ground truth logits and generated music logits, respectively. It is noted that we add cross entropy to the loss function of generator in order to accelerate

the convergence process of the loss function. By training both the generator and discriminator with the relativistic standard loss to emulate human-like musical compositions, our MMT-BERT model can generate high quality music pieces that incorporate sophisticated chord information.

## 4. EXPERIMENT

### 4.1 Experiment Setup

In the experiment, we utilize the Symbolic Orchestral Database (SOD) [36], which comprises 5,864 music pieces encoded as MIDI files along with associated metadata. The dataset is partitioned into training, testing, and validation sets, receiving 80%, 10%, and 10% of the data, respectively. We set a temporal resolution of 12 time steps per quarter note for detailed timing accuracy. The Transformer-XL generator is composed of six decoder layers with 512 dimensions and eight self-attention heads, and the MusicBERT discriminator consists of one encoder layer with two self-attention heads. The maximum length for symbolic music sequences is set at 1024, with a maximum of 256 beats. To optimize the models, we employ the Adagrad optimizer to mitigate issues of gradient explosion and vanishing [37]. Additionally, to enhance the robustness of the data, we augment it by randomly transposing all pitches by  $s \sim U(-5, 6)$  ( $s \in \mathbb{Z}$ ) semitones and assign a starting beat. Here,  $U$  denotes a uniform distribution.

As comparative methods, we employ three state-of-the-art music generation models: MMM [21], FIGARO [23], and MMT [6]. We validate the performance of our MMT-BERT model by conducting quantitative evaluations using existing metrics and subjective experiments to assess the human-like qualities of the generated music pieces.

### 4.2 Quantitative Evaluation

Following [6], we evaluate the generated music pieces using four metrics: pitch class entropy similarity (PCES), scale consistency similarity (SCS), groove consistency similarity (GCS), and average length (AL). We consider higher values of PCES, SCS, and GCS as indicators of superior quality, while a higher AL denotes a greater capability to produce long-duration music pieces.

In preparation for calculating PCES, the pitch class entropy (PCE) is defined as follows:

$$\text{PCE} = - \sum_{i=0}^{11} h_i \log_2(h_i), \quad (4)$$

where  $h_i$  denotes the number of occurrences of each note name in the 12-dimensional pitch class histogram. As the PCE values increase, the tonality of the generated music pieces exhibits greater instability. However, it is important to recognize that more stable tonality does not necessarily imply higher quality. Subsequently, we calculate PCES between generated music samples and human compositions as follows:

$$\text{PCES} = 1 - \frac{|\text{PCE}_{\text{gen}} - \text{PCE}_{\text{tr}}|}{\text{PCE}_{\text{tr}}}, \quad (5)$$

where  $\text{PCE}_{\text{gen}}$  and  $\text{PCE}_{\text{tr}}$  denotes the PCE value of generated music samples and human compositions, respectively. Moreover, noticing that PCE is intrinsically linked to the volume of data, we truncate the generated musical pieces to the preceding  $k$  seconds and calculate their PCES.

The scale consistency (SC) is derived by calculating the proportion of tones that conform to a conventional scale and presenting the value for the most closely aligned scale [38]. SC serves as an indicator of the model’s proficiency in generating musical segments that demonstrate cognizance of chords and scales within the current bar. The SCS between generated music samples and human compositions is defined as follows:

$$\text{SCS} = 1 - \frac{|\text{SC}_{\text{gen}} - \text{SC}_{\text{tr}}|}{\text{SC}_{\text{tr}}}, \quad (6)$$

where  $\text{SC}_{\text{gen}}$  and  $\text{SC}_{\text{tr}}$  denote the SC values of generated music samples and human compositions, respectively.

To calculate GCS, we first define a groove pattern  $\mathbf{g}$  as a 64-dimensional binary vector. The groove consistency (GC) between two grooving patterns ( $\mathbf{g}^a, \mathbf{g}^b$ ) is defined as follows:

$$\text{GC} = 1 - \frac{1}{Q} \sum_{i=0}^{Q-1} \text{XOR}(g_i^a, g_i^b), \quad (7)$$

where  $\text{XOR}(\cdot, \cdot)$  denotes the exclusive OR operation, and  $g_i$  denotes a position in a bar at which there is at least a note onset.  $Q$  is the dimensionality of  $\mathbf{g}^a$  and  $\mathbf{g}^b$ . GC is a measure of music’s rhythmicity. The value of GC stands for the steadiness in rhythm of the generated music pieces. The GCS between generated music samples and human compositions is defined as follows:

$$\text{GCS} = 1 - \frac{|\text{GC}_{\text{gen}} - \text{GC}_{\text{tr}}|}{\text{GC}_{\text{tr}}}, \quad (8)$$

where  $\text{GC}_{\text{gen}}$  and  $\text{GC}_{\text{tr}}$  denote the GC values of generated music samples and human compositions, respectively.

AL denotes the mean duration of the generated music pieces, which collectively illustrates the model’s ability to generate musical sequences with significant length.

The results of the quantitative evaluation are shown in Table 3. To facilitate a fair comparison by standardizing the lengths of music pieces, PCES is assessed over a 15-second span due to the limitations of MMM and FIGARO in producing extended compositions. Experimental results show that MMT-BERT achieves higher performance in PCES, SCS, and GCS compared to the other methods, demonstrating its effectiveness in generating high quality music pieces. This achievement is attributed to its chord awareness and the symbolic music understanding facilitated by MusicBERT. MMT-BERT’s AL is marginally less than that of MMT, and this results from integrating chord events that are not converted to audio during the decoding phase. However, MMT-BERT’s AL significantly surpasses that of MMM and FIGARO, confirming its capability to generate longer compositions. Additionally, the AL of all the music pieces in the SOD we used, which also serve as

	PCES (%)	SCS (%)	GCS (%)	AL (sec)
MMM [21]	92.93±1.22	98.64±0.92	98.28±0.29	38.69
FIGARO [23]	94.33±0.31	98.70±0.22	<u>98.84±0.67</u>	28.69
MMT [6]	95.19±0.45	98.94±0.77	98.44±0.55	<b>100.42</b>
MMT-BERT w/o Chord event	95.57±1.32	98.81±0.23	99.56±0.32	<u>100.25</u>
MMT-BERT w/o MusicBERT	<u>96.22±0.44</u>	<u>99.14±0.29</u>	98.61±0.44	97.43
MMT-BERT (ours)	<b>99.73±0.21</b>	<b>99.64±0.31</b>	<b>99.66±0.25</b>	99.87

**Table 3.** Quantitative evaluation results. The **boldface** denotes the highest value, and the underlined denotes the second highest value, respectively.

	R	H	C	S	O
MMM [21]	3.83±0.92	3.78±0.87	3.78±0.73	3.67±0.84	3.83±0.79
FIGARO [23]	3.78±1.11	3.78±1.11	3.89±1.02	3.89±1.13	3.83±0.92
MMT [6]	3.22±0.70	3.17±0.98	<b>3.44±1.03</b>	3.33±1.09	3.22±0.78
MMT-BERT (ours)	<b>3.55±0.94</b>	<b>3.55±0.92</b>	3.33±0.98	<b>3.39±0.90</b>	<b>3.44±0.80</b>

**Table 4.** Subjective evaluation results. Each metric is rated on a five-point scale, with the average score being calculated.

the ground truth, is 99.88 seconds. Evaluation results show that MMT-BERT can produce music of higher quality than MMT, and of longer duration than MMM and FIGARO.

### 4.3 Impacts of Chord Event and Discriminator

MMT-BERT aims to generate more harmonious, more human-like music pieces through the addition of chord events and adversarial generative learning by employing MusicBERT as its discriminator. To evaluate aspects related to richness and humanness, we have conducted subjective experiment and ablation study.

In the subjective experiment, we asked 18 music amateurs as the following five questions and requested that they rated each on a five-point scale.

- **Richness (R):** Does the music piece have diversity and interestingness?
- **Humanness (H):** Does the music piece sound like it was composed by an expressive human musician?
- **Correctness (C):** Does the music piece contain perceived mistakes in composition or performance?
- **Structureness (S):** Does the music piece exhibit structural patterns such as repeating themes or the development of musical ideas?
- **Overall (O):** What is the general score of the music piece?

As mentioned in Sec. 4.2, FIGARO and MMM employ a music representation that considers percussive sounds and typically generates much shorter pieces. Therefore, the nature of the music pieces generated by these models, FIGARO and MMM, differs significantly from that of MMT-BERT and MMT due to their use of percussive sounds and shorter compositions. To fairly evaluate the human-like quality of the generated music pieces, we compared MMT-BERT with MMT, a state-of-the-art approach whose generated compositions have lengths and musical styles that are

relatively similar to those of MMT-BERT. Additionally, to ensure clarity in subjective evaluation, we included the results for MMM and FIGARO. The results of the subjective evaluation are shown in Table 4. Table 4 indicates that MMT-BERT scores are particularly high in both richness and humanness compared to MMT. This suggests that the application of chord events and MusicBERT contribute to the generation of music pieces that more closely resemble human compositions. On the other hand, regarding correctness, our method did not specifically aim to enhance this metric, which may cause the gap in this value. For the same reason, our method exceeds MMT by a small margin in structureness mainly because of uncertainty. Although there is no clear advantage between MMT and MMT-BERT in correctness and structureness, our method still outperforms MMT in richness and humanness. The overall score also proves the superiority of MMT-BERT, which indicates that chord events and MusicBERT enhance the ability to create music similar to that produced by humans.

The results of the ablation study are shown in Table 3 along with the quantitative evaluation results. It is evident that the addition of chord events improves PCES and SCS. MusicBERT also contributes to the enhancement of PCES and GCS.

## 5. CONCLUSION

In this paper, we have proposed the chord-aware symbolic music generation approach, named MMT-BERT. By extracting chord information from raw audio files, we have devised a chord-aware symbolic music representation. We also developed a novel RS-GAN architecture based on MMT and MusicBERT. Both experimental evaluations validate the efficacy of our method in producing music pieces of superior quality, enhanced human likeness, and considerable length. In future works, we plan to explore methods that refine musical structure and incorporate information from various musical modalities.

## 6. ACKNOWLEDGMENTS

This work was partly supported by JSPS KAKENHI Grant Numbers JP21H03456, JP23K11141 and JP23KJ0044.

## 7. REFERENCES

- [1] W. Liu, "Literature survey of multi-track music generation model based on generative confrontation network in intelligent composition," *The Journal of Supercomputing*, vol. 79, no. 6, pp. 6560–6582, 2023.
- [2] F. P. Brooks, A. L. Hopkins, P. G. Neumann, and W. V. Wright, "An experiment in musical composition," *IRE Transactions on Electronics Computers*, vol. 6, pp. 175–182, 1957.
- [3] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," *Advances in Neural Information Processing Systems*, vol. 27, pp. 53–65, 2014.
- [4] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in Neural Information Processing Systems*, vol. 30, pp. 5999–6009, 2017.
- [5] J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models," *Advances in Neural Information Processing Systems*, vol. 33, pp. 6840–6851, 2020.
- [6] H.-W. Dong, K. Chen, S. Dubnov, J. McAuley, and T. Berg-Kirkpatrick, "Multitrack music transformer," in *Proc. of the International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023, pp. 1–5.
- [7] Y.-S. Huang and Y.-H. Yang, "Pop music transformer: Beat-based modeling and generation of expressive pop piano compositions," in *Proc. of the ACM International Conference on Multimedia (ACMMM)*, 2020, pp. 1180–1188.
- [8] F. Takuya, "Realtime chord recognition of musical sound: A system using common lisp music," in *Proc. of the International Computer Music Conference (ICMC)*, 1999, pp. 464–467.
- [9] J. Park, K. Choi, S. Jeon, D. Kim, and J. Park, "A bi-directional transformer for musical chord recognition," in *Proc. of the International Society for Music Information Retrieval Conference (ISMIR)*, 2019, pp. 620–627.
- [10] R. E. Scholz and G. L. Ramalho, "Cochonut: Recognizing complex chords from midi guitar sequences," in *Proc. of the International Society for Music Information Retrieval Conference (ISMIR)*, 2008, pp. 27–32.
- [11] E. Demirel, B. Bozkurt, and X. Serra, "Automatic chord-scale recognition using harmonic pitch class profiles," in *Proc. of the Sound and Music Computing Conference (SMC)*, 2019, pp. 72–79.
- [12] H. Jhamtani and T. Berg-Kirkpatrick, "Modeling self-repetition in music generation using generative adversarial networks," in *Proc. of the Machine Learning for Music Discovery Workshop, ICML*, 2019.
- [13] H. Zhang, L. Xie, and K. Qi, "Implement music generation with GAN: A systematic review," in *Proc. of the International Conference on Computer Engineering and Application (ICCEA)*, 2021, pp. 352–355.
- [14] S. Walter, G. Mougeot, Y. Sun, L. Jiang, K.-M. Chao, and H. Cai, "MidiPGAN: A progressive gan approach to midi generation," in *Proc. of the IEEE International Conference on Computer Supported Cooperative Work in Design (CSCWD)*, 2021, pp. 1166–1171.
- [15] A. Muhamed, L. Li, X. Shi, S. Yaddanapudi, W. Chi, D. Jackson, R. Suresh, Z. C. Lipton, and A. J. Smola, "Symbolic music generation with transformer-gans," in *Proc. of the AAAI Conference on Artificial Intelligence*, 2021, pp. 408–417.
- [16] H.-W. Dong, W.-Y. Hsiao, L.-C. Yang, and Y.-H. Yang, "MuseGAN: Multi-track sequential generative adversarial networks for symbolic music generation and accompaniment," in *Proc. of the AAAI Conference on Artificial Intelligence*, 2018.
- [17] H.-W. Dong and Y.-H. Yang, "Convolutional generative adversarial networks with binary neurons for polyphonic music generation," in *Proc. of the International Society for Music Information Retrieval Conference (ISMIR)*, 2018, pp. 190–196.
- [18] W.-Y. Hsiao, J.-Y. Liu, Y.-C. Yeh, and Y.-H. Yang, "Compound word transformer: Learning to compose full-song music over dynamic directed hypergraphs," in *Proc. of the AAAI Conference on Artificial Intelligence*, 2021, pp. 178–186.
- [19] J. Copet, F. Kreuk, I. Gat, T. Remez, D. Kant, G. Synnaeve, Y. Adi, and A. Défossez, "Simple and controllable music generation," *Advances in Neural Information Processing Systems*, vol. 36, 2024.
- [20] S.-L. Wu and Y.-H. Yang, "The jazz transformer on the front line: Exploring the shortcomings of ai-composed music through quantitative measures," in *Proc. of the International Society for Music Information Retrieval Conference (ISMIR)*, 2020, pp. 142–149.
- [21] J. Ens and P. Pasquier, "MMM: Exploring conditional multi-track music generation with the transformer," *arXiv preprint arXiv:2008.06048*, 2020.
- [22] B. Yu, P. Lu, R. Wang, W. Hu, X. Tan, W. Ye, S. Zhang, T. Qin, and T.-Y. Liu, "Museformer: Transformer with fine-and coarse-grained attention for music generation," *Advances in Neural Information Processing Systems*, vol. 35, pp. 1376–1388, 2022.

- [23] D. von Rütte, L. Biggio, Y. Kilcher, and T. Hofmann, “FIGARO: Generating symbolic music with fine-grained artistic control,” in *Proc. of the International Conference on Learning Representations (ICLR)*, 2023, pp. 1–18.
- [24] J. Salazar, D. Liang, T. Q. Nguyen, and K. Kirchoff, “Masked language model scoring,” in *Proc. of the Annual Meeting of the Association for Computational Linguistics (ACL)*, Jul. 2020, pp. 2699–2712.
- [25] E. Montahaei, D. Alihosseini, and M. S. Baghshah, “Jointly measuring diversity and quality in text generation models,” in *Proc. of the Workshop on Methods for Optimizing and Evaluating Neural Language Generation (NeuralGen)*, 2019, pp. 90–98.
- [26] Z. Dai, Z. Yang, Y. Yang, J. Carbonell, Q. V. Le, and R. Salakhutdinov, “Transformer-XL: Attentive language models beyond a fixed-length context,” in *Proc. of the Annual Meeting of the Association for Computational Linguistics (ACL)*, 2019, pp. 2978–2988.
- [27] T. Zhang\*, V. Kishore\*, F. Wu\*, K. Q. Weinberger, and Y. Artzi, “Bertscore: Evaluating text generation with bert,” in *Proc. of the International Conference on Learning Representations (ICLR)*, 2020.
- [28] E. Montahaei, D. Alihosseini, and M. S. Baghshah, “Jointly measuring diversity and quality in text generation models,” in *Proc. of the Workshop on Methods for Optimizing and Evaluating Neural Language Generation (NeuralGen)*, 2019, pp. 90–98.
- [29] M. Zeng, X. Tan, R. Wang, Z. Ju, T. Qin, and T.-Y. Liu, “MusicBERT: Symbolic music understanding with large-scale pre-training,” in *Proc. of the Association for Computational Linguistics: ACL-IJCNLP*, 2021, pp. 791–800.
- [30] A. Jolicoeur-Martineau, “The relativistic discriminator: a key element missing from standard gan,” in *Proc. of the International Conference on Learning Representations (ICLR)*, 2018.
- [31] C. Li, R. Fu, and Y. Liu, “Algorithm for generating tire defect images based on RS-GAN,” in *Proc. of the International Conference on Neural Information Processing (ICNIP)*, 2023, pp. 388–399.
- [32] J.-P. Briot, G. Hadjeres, and F.-D. Pachet, “Deep learning techniques for music generation—a survey,” *arXiv preprint arXiv:1709.01620*, 2017.
- [33] S. Ji, J. Luo, and X. Yang, “A comprehensive survey on deep music generation: Multi-level representations, algorithms, evaluations, and future directions,” *arXiv preprint arXiv:2011.06801*, 2020.
- [34] P. Neves, J. Fornari, and J. Florindo, “Generating music with sentiment using transformer-gans,” in *Proc. of the International Society for Music Information Retrieval (ISMIR)*, 2022, pp. 717–725.
- [35] H.-W. Dong, K. Chen, J. McAuley, and T. Berg-Kirkpatrick, “MusPy: A toolkit for symbolic music generation,” in *Proc. of the International Society for Music Information Retrieval Conference (ISMIR)*, 2020, pp. 101–108.
- [36] L. Crestel, P. Esling, L. Heng, and S. McAdams, “A database linking piano and orchestral midi scores with application to automatic projective orchestration,” in *Proc. of the International Society for Music Information Retrieval Conference (ISMIR)*, 2018, pp. 592–598.
- [37] J. Duchi, E. Hazan, and Y. Singer, “Adaptive sub-gradient methods for online learning and stochastic optimization.” *Journal of machine learning research*, vol. 12, no. 7, 2011.
- [38] O. Mogren, “C-RNN-GAN: A continuous recurrent neural network with adversarial training,” in *Proc. of the Constructive Machine Learning Workshop (CML) at NIPS*, 2016, pp. 1–6.