



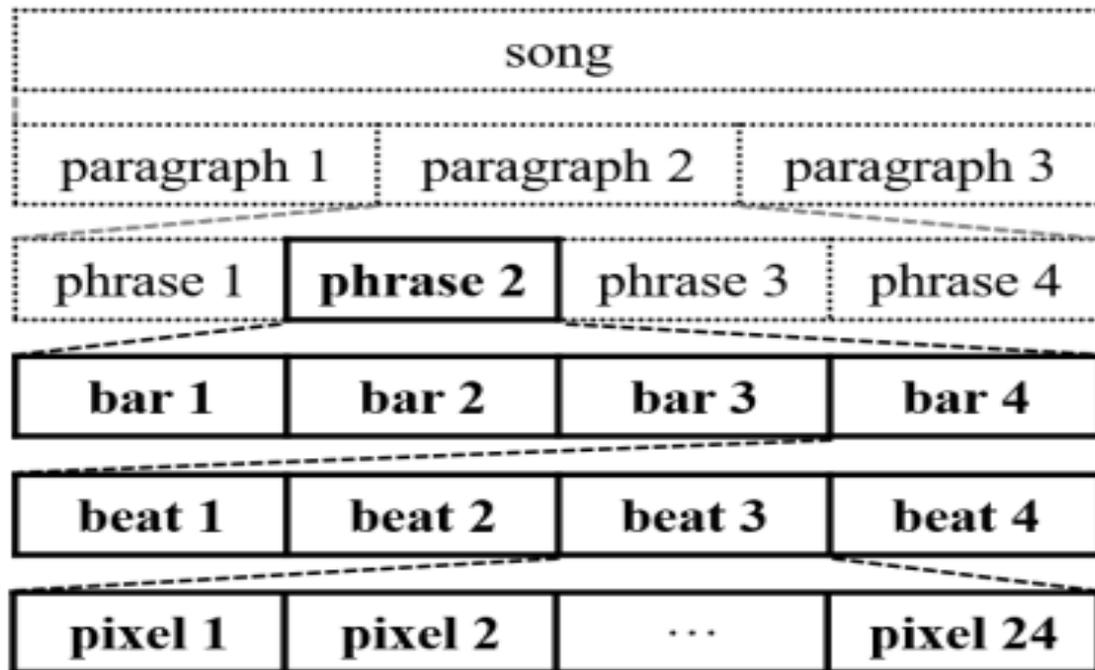
## AI BASED MUSIC TUNESMITH

**Shobhit sharma, Amritander Pratap Singh, Ramesh vaish, Aman singh, Anamika jadv**

**Abstraction:** Generating music has some distinct differences from generating images and videos. To begin with, it is an art of time, so a temporal model is essential. Additionally, music is often composed of multiple instruments/tracks with their own tempo, which all unfold over time in a mutually dependent manner. Musical notes are also often grouped into chords, arpeggios, or melodies in polyphonic music, thus introducing a hierarchical ordering of notes. In this paper, we propose three models for symbolic multi-track music generation based on a framework of generative adversarial networks (GANs). The three models differ in their underlying assumptions and corresponding network architectures, and are referred to as the jamming model, the composer model, and the hybrid model. We trained the proposed models on a dataset of more than one hundred thousand bars of rock music, and applied them to generate piano rolls.

**INTRODUCTION:** Generating realistic and aesthetic pieces is a highly sought-after challenge in the realm of Artificial Intelligence. In recent years, breakthroughs have been made in creating images, videos, and text using Generative Adversarial Networks (GANs) (Goodfellow et al. 2014; Radford et al. 2016; Vondrick et al. 2016; Saito et al. 2017; Yu et al. 2017). There have also been attempts to generate symbolic music, but the task remains difficult due to music's temporal nature. As seen in Figure 1, music has a hierarchical structure composed of higher-level components (e.g., a phrase) made up of repetitious segments (e.g., a bar). People pay attention to structural patterns relating to cohesion, rhythm, tension, and emotion progression.



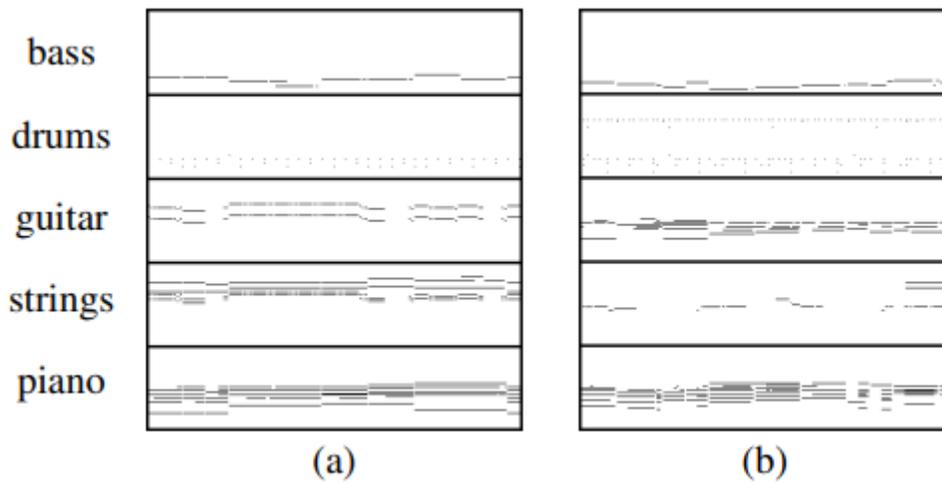


**Figure 1: Hierarchical structure of a music piece.**

Music is usually composed of multiple instruments and tracks. An orchestra typically consists of brass, strings, woodwinds, and percussion instruments, while a rock band often comprises bass, drum set, guitars, and sometimes vocals. These components interact with one another and evolve together. In music theory, there are considerable discussions around composition techniques, such as harmony and counterpoint. Moreover, musical notes are often organized into chords, arpeggios, or melodies. It is not feasible to impose a chronological ordering of notes in polyphonic music; therefore, success in natural language generation and monophonic music generation cannot be replicated in polyphonic music generation. Accordingly, most prior art (as discussed in the Related Work section) has chosen to simplify symbolic music generation in certain ways.

We propose two approaches to incorporate a temporal model: one that creates music from scratch without any human input, and another that learns to follow the given a priori temporal structure of a track. To address the interactions among tracks, we suggest three methods based on our understanding of pop music. The first generates tracks independently using separate generators for each. The second creates them collectively with one generator. The third uses separate generators for each, plus additional shared inputs to guide the tracks to sound harmonious and coordinated. To tackle the grouping of notes, we view bars rather than notes as the basic compositional unit and use transposed convolutional neural networks (CNNs) to generate music one bar at a time, which are known to be effective at locating point patterns.

**Generative Adversarial Networks:** The core concept of Generative Adversarial Networks (GANs) is to generate adversarial learning by constructing two networks: the generator and the discriminator (Goodfellow et al. 2014). The generator maps a random noise, sampled from a prior distribution, to the data space. On the other hand, the discriminator is trained to distinguish real data from data generated by the generator. The generator, on the other hand, is trained to fool the discriminator. The training process can be modeled as a two-player minimax game between the generator  $G$  and the discriminator  $D$ :  $\min_G \max_D E_x \text{sim}_p d[\log(D(x))] + E_z \text{sim}_p 2 [1 - \log(D(G(z)))]$ .



The core concept of Generative Adversarial Networks (GANs) is to generate adversarial learning by constructing two networks: the generator and the discriminator (Goodfellow et al. 2014). The generator maps a random noise, sampled from a prior distribution, to the data space. On the other hand, the discriminator is trained to distinguish real data from data generated by the generator. The generator, on the other hand, is trained to fool the discriminator. The training process can be modeled as a two-player minimax game between the generator  $G$  and the discriminator  $D$ :  $\min_G \max_D E_{x \sim p_d} [\log(D(x))] + E_{z \sim p_z} [1 - \log(D(G(z)))]$ .

**Proposed Model:** According to Yang, Chou, and Yang (2017), bar boundaries are often seen as the basic compositional unit for making harmonic changes (eg. chord changes). Furthermore, humans frequently employ bars as the fundamental components when writing songs.

**Implementation of Dataset:** The piano-roll dataset utilized in this work is derived from the Lakh MIDI dataset (LMD) (Raffel, 2016). We converted the MIDI files to piano-rolls with a height of 128 and width of 96 to capture common temporal patterns. The python library named pretty midi (Raffel and Ellis, 2014) was used to parse and process the MIDI files. Thus, the compiled dataset has been termed as the Lakh Pianoroll Dataset (LPD). Additionally, we present the subset LPD-matched, which is derived from the LMD-matched, an amalgamation of 45,129 MIDIS and entries from the Million Song Dataset (MSD) (Bertin-Mahieux et al., 2011). The conversion utilities, along with the datasets and the metadata, are accessible on the project website.

**Analysis of Training Data:** The piano-roll dataset utilized in this work is derived from the Lakh MIDI dataset (LMD) (Raffel, 2016). We converted the MIDI files to piano-rolls with a height of 128 and width of 96 to capture common temporal patterns. The python library named pretty midi (Raffel and Ellis, 2014) was used to parse and process the MIDI files. Thus, the compiled dataset has been termed as the Lakh Pianoroll Dataset (LPD). Additionally, we present the subset LPD-matched, which is derived from the LMD-matched, an amalgamation of 45,129 MIDIS and entries from the Million Song Dataset (MSD) (Bertin-Mahieux et al., 2011). The conversion utilities, along with the datasets and the metadata, are accessible on the project website.

**Training Process:** To gain insights into the training process, we examine the generator model used for creating music from scratch (other models have similar behaviors). Figure displays the training loss of  $D$  as a function of training steps, which is seen to decline rapidly in the beginning and reach a stable level afterward. However, there is also a slight increase in the trend after point B that appears on the graph, suggesting that  $G$  commences to learn something from that point onwards.

The generated piano-rolls at the five points marked in Figure are shown in Figure . We can observe how the generated piano-rolls develop over the training process. For instance,  $G$  grasps the pitch range of each track quite early and starts to produce some notes, though broken up but still within the right pitch range, at point B rather than the noise produced at point A. At point B, we can see clusters of points assembling at the lower part (all with lower pitches) of the bass. After point C,  $G$  begins to produce longer segments with a few tones.

**Related Work:** Our model design is based on prior work in video generation that utilize Generative Adversarial Networks (GANs). For example, VGAN (Vondrick, Pirsiavash, & Torralba, 2016) assumed that a video can be broken down into a dynamic

foreground and static background, and used 3D and 2D CNNs to generate each component in a two-stream architecture and then combine them with a mask generated by the foreground stream. TGAN (Saito, Matsumoto, & Saito, 2017) used a temporal generator (made of convolutions) to create a fixed-length series of latent variables, which were then fed one by one to an image generator to construct the video frame-by-frame. MOCOGAN (Tulyakov, et al., 2017) assumed that a video can be separated into content (objects) and motion (of objects) and used RNNs to capture the motion of objects.

Conclusion: In this paper, we present a new generative model using GANS for multi-track sequence generation. To demonstrate the effectiveness of the model, we have employed deep CNNs to generate multi-track piano-rolls. Additionally, we have designed several objective metrics to gain insights into the learning process. The objective metrics and subjective user study reveal that the proposed models can begin to learn about music. Although the musically and aesthetically generated tracks do not match the level of human musicians yet, the proposed model has certain desirable properties and we hope that future research will be able to further improve it.

#### References:

Arjovsky, M.; Chintala, S.; and Bottou, L. 2017. Wasserstein

GAN. arXiv preprint arXiv:1701.07875. Bertin-Mahieux, T.; Ellis, D. P.; Whitman, B.; and Lamere, P. 2011. The Million Song Dataset. In ISMIR.

Boulanger-Lewandowski, N.; Bengio, Y.; and Vincent, P. 2012. Modeling temporal dependencies in high-dimensional sequences: Application to polyphonic music generation and transcription. In ICML.

Briot, J.-P.; Hadjeres, G.; and Pachet, F. 2017. Deep learning techniques for music generation: A survey. arXiv preprint arXiv:1709.01620.

Harte, C.; Sandler, M.; and Gasser, M. 2006. Detecting harmonic change in musical audio. In ACM MM workshop on Audio and music computing multimedia.

Herremans, D., and Chew, E. 2017. Morpheus: generating structured music with constrained patterns and tension. IEEE Trans. Affective Computing.

Chu, H.; Urtasun, R.; and Fidler, S. 2017. Song from PI: A musically plausible network for pop music generation. In ICLR Workshop.

Goodfellow, I. J.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y.

2014. Generative adversarial nets. In NIPS.

Tulyakov, S.; Liu, M.; Yang, X.; and Kautz, J. 2017. MoCo- GAN: Decomposing motion and content for video genera-

tion. arXiv preprint arXiv:1707.04993. Vondrick, C.; Pirsiavash, H.; and Torralba, A. 2016. Gener- ating videos with scene dynamics. In NIPS.

Yang, L.-C.; Chou, S.-Y.; and Yang, Y.-H. 2017. MidiNet: A convolutional generative adversarial network for symbolic- domain music generation. In ISMIR.

Yu, L.; Zhang, W.; Wang, J.; and Yu, Y. 2017. SeqGAN: Sequence generative adversarial nets with policy gradient. In AAAI.