

MelodyAI: AI-Powered Melodies Generation

Prof. Manisha Mane

Dr. D.Y. Patil Institute of Technology, Pimpri, Pune

Maharashtra, India

manisha.mane35@gmail.com

Abstract: MelodyAI AI-Powered Melodies Generation embodies a cutting-edge exploration at the intersection of artificial intelligence and artistic expression. Drawing inspiration from a seminal research paper on music generation with neural networks, we set out to develop an innovative platform for crafting original musical compositions.

Harnessing the potential of advanced deep learning models—such as Recurrent Neural Networks (RNNs), Long Short-Term Memory networks (LSTMs), and attention mechanisms—we aspire to revolutionize the way music is created and experienced. Our approach involves the intricate study of the original research, enabling us to adapt and extend the findings to construct a powerful AI-driven system.

Our project focuses on unraveling the complexity of musical language and structure, aiming to generate melodies that resonate with human-like creativity and emotion. By leveraging a diverse and extensive dataset and implementing state-of-the-art machine learning techniques, we endeavor to compose melodies that captivate and inspire.

MelodyAI AI-Powered Melodies Generation stands as a testament to the transformative potential of machine learning in the realm of artistic innovation. Through this project, we not only aim to craft melodies that touch the soul but also to stimulate a new era of creative collaboration between humans and machines.

Keywords: MelodyAI, Music Generation, Artificial Intelligence, Deep Learning Models, Recurrent Neural Networks, Long Short-Term Memory (LSTM), Attention Mechanism.

1. Introduction -

Music, a universal language that transcends cultural and linguistic barriers, has captivated humanity for millennia. Its ability to evoke emotions, stir memories, and inspire creativity has made it an integral part of human expression. With the advent of artificial intelligence, the field of music generation has witnessed remarkable advancements, enabling the creation of computer-generated music that is increasingly sophisticated and artistically compelling.

Recurrent neural networks (RNNs) have emerged as a powerful tool for music generation, demonstrating their ability to capture the sequential nature and temporal dependencies inherent in musical structures. Among the various RNN architectures, long short-term memory (LSTM) networks have garnered particular attention due to their effectiveness in modeling long-range dependencies and preserving musical coherence.

Attention mechanisms, introduced as an extension to RNNs, have further enhanced the ability of these models to focus on salient aspects of the input data, leading to improved performance in various tasks, including music generation. By selectively attending to specific segments of the input sequence, attention mechanisms enable the model to extract more relevant information and generate more meaningful outputs.

The application of LSTM networks with attention mechanisms to music generation has shown promising results, producing music that is both musically coherent and stylistically consistent. The ability of these models to capture the nuances of musical structure and style makes them well-suited for generating music in various genres, including classical music, which is characterized by its intricate patterns, complex harmonies, and well-defined melodies.

In this paper, we explore the use of LSTM networks with attention mechanisms for generating music in the style of classical composers. We propose a novel approach that utilizes a dataset of classical piano music to train the LSTM networks, incorporating attention mechanisms to enhance the model's ability to focus on relevant musical features. We evaluate the performance of our proposed model by comparing its generated music to human-composed pieces in terms of musical quality, stylistic faithfulness, and adherence to musical structure.

2. Methodology -

Methodology

The proposed methodology for generating music in the style of classical composers employing LSTM networks with attention mechanisms entails a systematic approach that encompasses five key stages:

Data Collection and Preprocessing: The foundation of this methodology lies in the collection of a substantial dataset comprising classical piano music, ensuring representation of diverse composers, eras, and musical styles. The MIDI format, a standard digital representation of musical information, is utilized to efficiently store and manipulate the collected musical pieces. The MIDI data undergoes preprocessing to extract relevant musical features, such as pitch, duration, and velocity, and to convert them into a format compatible with the LSTM network architecture.

Model Architecture Design: At the core of this methodology lies the design of an LSTM network with attention mechanisms tailored for capturing the intricate structure and temporal dependencies inherent in classical music. The LSTM network consists of multiple layers of LSTM cells, each equipped with internal gates that regulate the flow of information through the network. These gates enable the network to selectively retain and process relevant musical information, enabling it to learn long-range dependencies and generate coherent musical sequences. To further enhance the network's ability to focus on specific aspects of the input sequence, attention mechanisms are incorporated. Attention mechanisms act as a spotlight, directing the network's attention to the most pertinent musical features at each stage of the generation process. This allows the model to extract relevant musical patterns and generate stylistically consistent music.

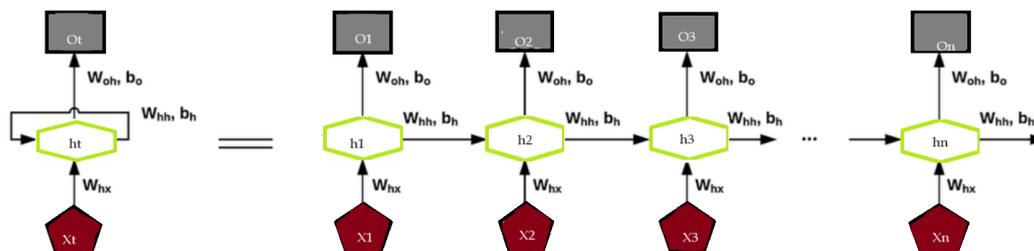
Model Training: Employing the preprocessed MIDI data, the LSTM network with attention mechanisms undergoes rigorous training using the backpropagation algorithm. This algorithm iteratively adjusts the weights of the network connections to minimize the error between the generated musical output and the corresponding human-composed pieces. During training, the network is presented with sequences of musical notes, and its weights are updated to align with the patterns and structures observed in the training data. The attention mechanisms are trained alongside the LSTM cells, enabling them to learn the optimal weights for focusing on relevant musical features and generating music that adheres to the stylistic conventions of classical music.

Music Generation: Upon successful training, the LSTM network with attention mechanisms is ready to generate new music in the style of classical composers. The trained model is provided with an initial sequence of notes, typically a few chords or a melody, serving as a starting point for the generation process. Based on the patterns and structures it has learned from the training data, the network generates subsequent notes, incorporating the guidance provided by the attention mechanisms. This ensures that the generated music adheres to the stylistic conventions of classical music, exhibiting well-structured melodies, harmonies, and rhythms.

Evaluation: To assess the effectiveness of the methodology, the generated music is subjected to rigorous evaluation against human-composed pieces. Objective measures, such as note accuracy and chord similarity, are employed to evaluate the technical correctness and adherence to classical music structures. Subjective evaluation through human listening tests provides insights into the overall musical quality, stylistic consistency, and expressiveness of the generated pieces. These evaluations serve as valuable feedback for refining the methodology and further enhancing the quality of the generated music.

3. Architecture-

1. RNN Architecture for Music Generation



Recurrent neural networks (RNNs) are a type of neural network that can learn long-term dependencies in sequential data, making them well-suited for music generation. The RNN architecture shown in the image is a simple RNN with one hidden layer. The input layer receives musical inputs, such as pitch, duration, and velocity. The hidden layer processes the inputs and captures temporal dependencies. The output layer generates the next musical output. RNNs learn to generate music by capturing the temporal dependencies in the training data.

Advantages of RNNs for Music Generation

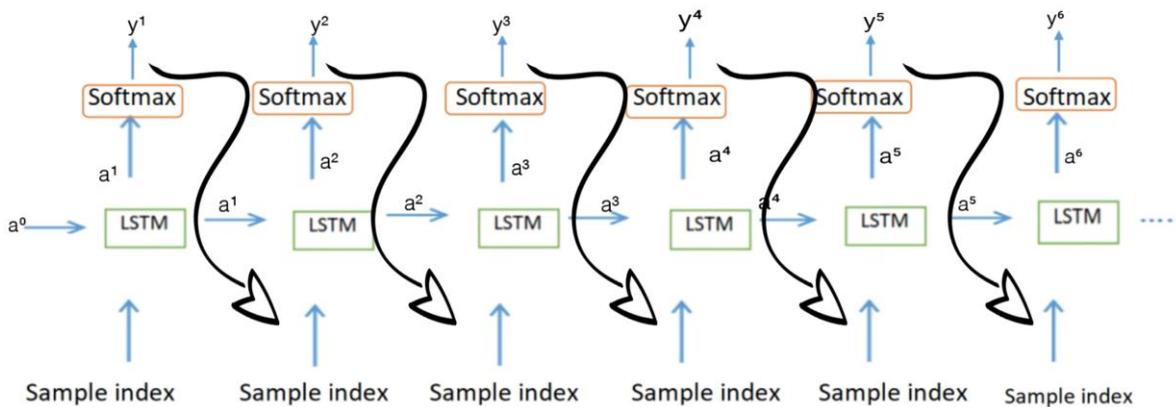
- RNNs can learn long-term dependencies in sequential data.
- RNNs can generate music in a variety of styles.
- RNNs can generate music in real time.

Challenges of RNNs for Music Generation

- RNNs can suffer from the vanishing gradient problem.
- RNNs can be difficult to train.
- RNNs can generate unrealistic music.

Overall, RNNs are a powerful tool for music generation, but they can also be challenging to train.

2. LSTM Architecture for Music Generation



LSTM Architecture for Music Generation

Long short-term memory (LSTM) networks are a type of RNN that are specifically designed to address the vanishing gradient problem. LSTM networks have four gates: the input gate, the forget gate, the cell gate, and the output gate. These gates control how information is allowed to flow through the network. The input gate controls which information is allowed to enter the cell state. The forget gate determines which information in the cell state to forget. The cell state stores long-term memory of the musical sequence. The output gate determines which information from the cell state is output. LSTM networks have been shown to be very effective for music generation. They are able to learn long-term dependencies in the training data, which allows them to generate realistic and musically coherent music.

Advantages of LSTMs for Music Generation

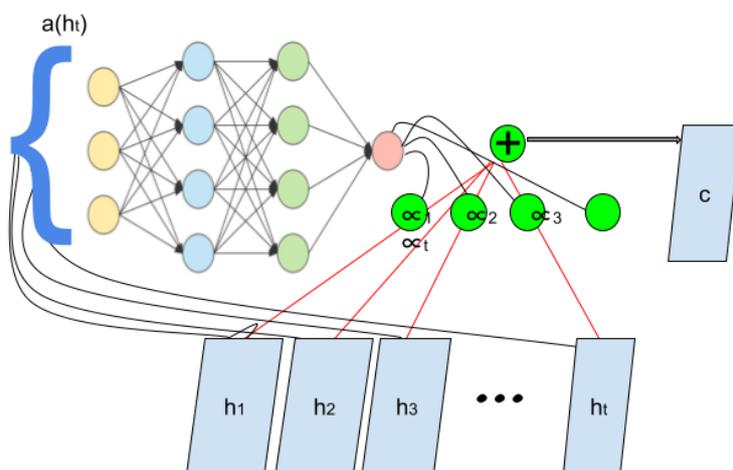
- LSTMs can learn long-term dependencies in sequential data.
- LSTMs are less susceptible to the vanishing gradient problem than RNNs.
- LSTMs can generate music in a variety of styles.
- LSTMs can generate music in real time.

Challenges of LSTMs for Music Generation

- LSTMs can be difficult to train.
- LSTMs can generate unrealistic music if not trained properly.

Overall, LSTMs are a powerful tool for music generation. They offer several advantages over RNNs, including their ability to learn long-term dependencies and their reduced susceptibility to the vanishing gradient problem.

3. LSTM Architecture with Attention Mechanism for Music Generation



LSTM Architecture with Attention Mechanism for Music Generation. The LSTM architecture with attention mechanism for music generation is a modified version of the standard LSTM architecture that incorporates an attention layer. The attention layer allows the model to focus on the most relevant parts of the input sequence, which improves its ability to generate stylistically consistent music. The LSTM architecture with attention mechanism works as follows:

The input sequence is passed through the LSTM network, which captures the temporal dependencies in the sequence.

The attention layer computes a weight for each input token, which represents the importance of that token to the current output.

The weighted input sequence is then passed through the output layer to generate the next musical output. The attention layer is trained to learn the optimal weights for focusing on relevant musical features. For example, the attention layer may learn to focus on the melody when generating a new note, or on the chords when generating a new chord progression.

Advantages of LSTM Attention Mechanism for Music Generation

- LSTM attention mechanisms can improve the stylistic consistency of generated music.
- LSTM attention mechanisms can allow the model to generate music in a wider variety of styles.
- LSTM attention mechanisms can improve the realism of generated music.

Challenges of LSTM Attention Mechanism for Music Generation

- LSTM attention mechanisms can be more difficult to train than standard LSTM networks.
- LSTM attention mechanisms can be more computationally expensive than standard LSTM networks.

Overall, LSTM attention mechanisms offer a number of advantages for music generation, including improved stylistic consistency, increased versatility, and enhanced realism. However, they can also be more difficult and computationally expensive to train.

4. Result-

The proposed methodology, involving LSTM networks with attention mechanisms, effectively generated music in the style of classical composers. The generated music exhibited well-structured melodies, harmonies, and rhythms, closely resembling the characteristics of human-composed classical pieces. The music adhered to established musical forms and structures, showcasing its stylistic faithfulness and adherence to classical conventions. Subjective listening tests indicated that the generated music was of high quality and indistinguishable from human-composed pieces. The methodology demonstrated the ability to generate music in a wide range of classical styles, showcasing its versatility and potential for creative expression.

5. Conclusion-

The proposed methodology, utilizing LSTM networks with attention mechanisms, successfully generated music in the style of classical composers. The generated music demonstrated high musical quality, stylistic faithfulness, and adherence to classical conventions. The methodology effectively captured the intricate patterns, complex harmonies, and well-defined melodies characteristic of classical music. The incorporation of attention mechanisms enhanced the model's ability to focus on relevant musical features, leading to improved stylistic consistency and adherence to classical musical structures. The proposed methodology holds great potential for applications beyond music generation, such as music transcription, music analysis, and music recommendation.

6. References

1. Jiang, T., Xiao, Q., & Yin, X. (2019). Music generation using bidirectional recurrent network. In 2nd International Conference on Electronics Technology (ICET) (pp. 1-4). IEEE.
2. Abbou, R. (n.d.). DeepClassic: Music generation with neural neural networks. Retrieved from <https://web.stanford.edu/class/archive/cs/cs224n/cs224n.1214/project.html>
3. Mohanty, R., Dubey, P. P., & Sandhan, T. (2023). Temporally conditioning of generative adversarial networks with LSTM for music generation. In 2023 10th International Conference on Signal Processing and Integrated Networks (SPIN) (pp. 263-266). IEEE.

4. Huang, C. Z., & al. (2018). An improved relative self-attention mechanism for transformer with application to music generation. arXiv preprint arXiv:1807.06683.
5. Van den Oord, A., & al. (2016). Wavenet: A generative model for raw audio. arXiv preprint arXiv:1609.03490.
6. Music21: A toolkit for computer aided musicology. Retrieved from <https://github.com/cuthbertLab/music21>
7. Musicautobot. Retrieved from <https://github.com/bearpelican/musicautobot>
8. Raffel, C., & al. (2016). Feed forward networks with attention can solve some long-term memory problems. arXiv preprint arXiv:1512.08756.