



NeuraScore: Music Generation using Deep Learning

Atharva Puranik, Abdul Rehman Tahami, Aanya Chourasiya, Ayushman Singh Chouhan, Amisha Prajapati

CSE Department, Acropolis Institute of Technology And Research, Bypass Road, Square, Mangalya Sadak, Indore, Madhya Pradesh 453771
atharvuranik20403@acropolis.in¹, abdulrehmantahmani20168@acropolis.in², aanyachourasiya20542@acropolis.in³,
ayushmanchouhan20089@acropolis.in⁴, amishaprajapati20792@acropolis.in⁵

ABSTRACT

Music is a form of expression that involves various methods and melodies to be created. Composing music is the process of creating music using melodies and rhythms.

With various improvements in technology music has grown and evolved and it continues to evolve. Leveraging deep neural networks, this approach enables a generation of various musical pieces without human intervention.

Music generation using deep learning offers a systematic and efficient way to generate music using analysis of vast musical datasets, deep learning models and various patterns and styles.

KEYWORDS: Music Generation, Deep Learning, Machine learning in music, AI Generated music, LSTM.

INTRODUCTION

Music generation using deep learning represents a cutting-edge intersection of artificial intelligence and the arts, where algorithms and neural networks are employed to compose and generate musical pieces autonomously.

This innovative approach to music composition leverages the capabilities of neural networks. Deep learning models, such as recurrent neural networks (RNNs) and generative adversarial networks (GANs), have been particularly influential in this domain. These models can capture complex musical structures and generate music that ranges from classical to contemporary, showcasing the adaptability of deep learning in the creative realm.

The process involves training these models on extensive datasets containing musical pieces, enabling them to learn the harmony, melody, rhythm, and other musical elements. Once trained, these models can generate entirely new compositions.

In this dynamic and evolving landscape, researchers and musicians are continuously exploring ways to enhance the capabilities of deep learning models, striving to create more emotionally resonant, diverse, and innovative compositions. As technology advances, the collaboration between artificial intelligence and human creativity in the realm of music generation is poised to redefine traditional notions of musical composition and offer new possibilities for the future of musical expression.

METHODOLOGY

The proposed solution for music generation utilizing Recurrent Neural Networks (RNNs) with LSTM (Long Short Term Memory) involves training the model on a dataset of musical sequences.

LSTMs, a type of RNN, excel in capturing long-term dependencies and are thus well-suited for sequential data like music. The model learns patterns and structures inherent in the music data during training.

During the generation phase, the trained model predicts the next note or sequence of notes based on the input provided. By sampling from the output probabilities, the model generates new music compositions that exhibit similar stylistic characteristics to the training data. This approach leverages the temporal dynamics of music to create harmonious and coherent compositions.

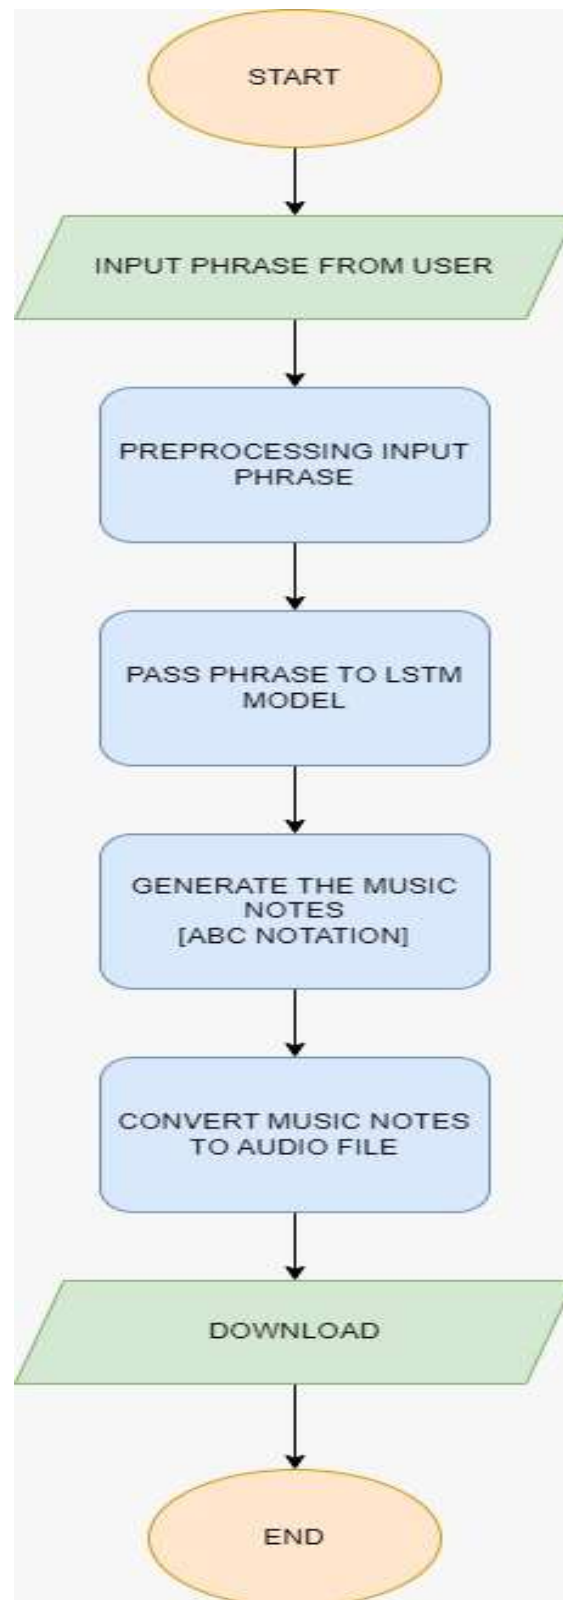


Fig 1 : Flowchart of Music Generation using Deep Learning

DEEP LEARNING

Deep Learning is a subfield of machine learning concerned with algorithms inspired by the structure and function of the brain called artificial neural networks. Deep learning (also known as deep structured learning or hierarchical learning) is part of a broader family of [machine learning](#) methods based on [learning data representations](#), as opposed to task-specific algorithms. Learning can be [supervised](#), [semi-supervised](#) or [unsupervised](#).

Deep learning models are loosely related to information processing and communication patterns in a biological [nervous system](#), such as [neural coding](#) that attempts to define a relationship between various stimuli and associated neuronal responses in the [brain](#).

Deep learning architectures such as [deep neural networks](#), [deep belief networks](#) and [recurrent neural networks](#) have been applied to fields including [computer vision](#), [speech recognition](#), [natural language processing](#), audio recognition, social network filtering, [machine translation](#), [bioinformatics](#) and [drug design](#), where they have produced results comparable to and in some cases superior to human experts.

RECURRENT NEURAL NETWORKS

Neural networks imitate the function of the human brain in the fields of AI, machine learning, and deep learning, allowing computer programs to recognize patterns and solve common issues.

RNNs are a type of neural network that can be used to model sequence data. RNNs, which are formed from feedforward networks, are similar to human brains in their behavior. Simply said, recurrent neural networks can anticipate sequential data in a way that other algorithms can't.

All of the inputs and outputs in standard neural networks are independent of one another, however in some circumstances, such as when predicting the next word of a phrase, the prior words are necessary, and so the previous words must be remembered. As a result, RNN was created, which used a Hidden Layer to overcome the problem. The most important component of RNN is the Hidden state, which remembers specific information about a sequence.

RNNs have a Memory that stores all information about the calculations. It employs the same settings for each input since it produces the same outcome by performing the same task on all inputs or hidden layers.




Fig 2: One - Many RNN [LSTM]

DEVELOPMENT

Being an important phase of software development, the design of the application's architecture was highly considered.

The application has a front end that acts as an interface to the users of the application and a backend that aids in the management of the application. The features of both are listed below.

Front-End

The frontend of the project likely involved a web-based interface or application where users could interact with the LSTM-generated music. This interface could have been built using HTML, CSS, and JavaScript, possibly with the assistance of frontend frameworks like React.js or Vue.js for dynamic user experiences. The frontend would display generated music samples, provide options for user input (such as seed phrases or style preferences), and offer controls for playback and evaluation.

Back-End

The backend of the project would handle the heavy lifting of LSTM model training, inference, and serving generated music samples to the frontend. Python is a common choice for backend development due to its popularity in machine learning and deep learning tasks. Frameworks like TensorFlow or PyTorch would likely be used to implement the LSTM model architecture and train it on large datasets of musical compositions. Additionally, Flask or Django could be utilized to build a RESTful API for communication between the frontend and backend, allowing for seamless integration and interaction.

The programming languages used in the development of the software were as follows:

Html was used in making the front end of the software. Javascript was used in making the logic of the software. CSS was used to design and beautify the external look of the software. Flask was used due to its various features. Python was used in the backend and server of the system. These Languages were used due to their popularity, how easy they were to learn and use and their compatibility with each other. Microsoft Visual Studio Code was used as the code editor to program these scripts. During implementation, most of the functionalities were developed as small chunks of code.

We also did intensive testing and debugging to make sure our program was working correctly.

RELATED WORK

Long Short Term memory neural networks (LSTMNNs)

□ This paper researches how to apply the RNN using Long Short Term Memory. This system has proved the efficiency of an LSTM RNN model at producing a musical sequence that matches the dataset in terms of its the sequence's grammatical coherence. The final output, once converted into MIDI, is surreal to most casual listeners' ears. Most of the audience couldn't identify any discrepancies with the sequence generated.

Automatic Music Generation Using Deep Learning

In this system starting and ending music can be added in every new generated tune to give a tune a better start and better ending. By doing this, generated music will become melodious. Model can also be trained with multi-instrument tunes. By filtering unknown notes and replacing them with known notes, model can generate more robust quality music.

Deep Learning Music Generation

In this research paper the results show that a batch size of 128 with training under 100 epochs produces the best accuracy. The accuracy for the train set produces a higher accuracy than the test and validation as expected, but is still significantly underfitting. Accuracy is still low as the model requires more epochs, additional tuning, and possibly a deeper network

DeepClassic

DeepClassic is used to design a network that could automatically generate piano music. The network is trained based on a predictive task: the developers fit the network to predict the representation of the note to come based on the previous notes that have been played. They then use this predictive model to generate notes that are likely to appear in a human-composed piece.

EXISTING SYSTEM DRAWBACKS

One of the significant drawbacks of implementing a system for music generation using deep learning is the extensive requirement for high-quality, labeled data. Training RNN models effectively demands large, diverse, and well-annotated musical datasets. In many cases, these datasets can be challenging to obtain due to ethical concerns such as copyright. Moreover, there can be a scarcity of originality as the model mimics whatever the musical data that is feeded into them. This data limitation not only hampers the development and training of robust RNN models but also introduces the risk of model biases, as training data might not adequately represent the full spectrum of diversity and uniqueness in the music. Overcoming these data-related challenges is an ongoing issue that developers and musicians in this field must address to maximize the potential of RNNs in music generation.

RESULT AND DISCUSSION

The results and discussion section of the research paper elucidates findings gleaned from LSTM-based music generation experiments. It delves into the efficacy of LSTM networks in capturing intricate musical patterns and generating coherent compositions. Through detailed analysis, the section examines model performance metrics and discusses observed limitations. Furthermore, it explores potential avenues for improvement and proposes future research directions, contributing to a deeper understanding of LSTM-based music generation and its implications for creative expression and AI-assisted composition.

CONCLUSION

In summary, this research paper delves into the use of LSTM networks for music generation, showcasing their ability to capture intricate temporal patterns in musical data. Despite challenges like model complexity and subjective evaluation criteria, LSTM models offer exciting prospects for creative exploration in music composition. Future endeavors may explore hybrid models, additional feature integration, and refined evaluation metrics. Overall, LSTM networks present a promising avenue for pushing the boundaries of musical expression and fostering innovative collaborations between humans and AI.

REFERENCE

- Music Generation Using Deep Learning, Vaishali Ingale, Anush Mohan, Divit Adlakha, Krishna Kumar, Mohit Gupta
- Ratika Jadhav Aarati Mohite Debashish Chakravarty and Sanjay Nalbalwar, Automatic Music Generation Using Deep Learning
- Shaily Malik, Gaurav Arora, Anvaya Ahlawat, Mandeep Payal, MUSIC GENERATION USING DEEP LEARNING – Char RNN
- Raphael Abbou, DeepClassic: Music Generation with Neural Neural Networks
- Kinbert Chou, Ryan Peng, Deep Learning Music Generation

-
- Pedro Ferreira, Ricardo Limongi, and Luiz Paulo Fávero, Generating Music with Data: Application of Deep Learning Models for Symbolic Music Composition
 - Jean-Pierre Briot, From Artificial Neural Networks to Deep Learning for Music Generation – History, Concepts and Trends

ACKNOWLEDGEMENT

We extend our heartfelt gratitude to the students and faculty members of Acropolis Institute of Technology and Research for their unwavering support and invaluable contributions to this research project. Their dedication and collaborative efforts have significantly enriched the outcome of this work. We are deeply thankful for their guidance, expertise, and commitment, which have been instrumental in the successful completion of this endeavor.