

# Lyrics Generation using Recurrent Neural Networks

Thendral Pyualnithi<sup>1</sup>, Murugeswari Kandavel<sup>2</sup>

Dept. of Computer Science and Engineering

Kalasalangam Academy of research and Education, Srivilliputtur, India

<sup>1</sup>thendral.p@klu.ac.in, <sup>2</sup>murugeswari.k@klu.ac.in

Rajvardhan Dixit

School of Computer Science and Engineering, Vellore Institute of Technology, Vellore, India

**Abstract**— It's a neural network that has been trained on Kanye West's (very famous American singer and music producer) discography, and can use any lyrics you feed it and write a new song word by word that rhymes and has a flow (to an extent). It is basically a markov chain will look at the lyrics you entered and generate new lines. Then, it feeds this to a recurrent neural net that will generate a sequence of tuples in the format of: (Desired rhyme, desired count of syllables). The program will then sift through the lines the markov chain generated, and match the lines to their corresponding tuples. The result is the rap. (Abstract)

**Keywords**- Recurrent Neural Network, Machine Learning, Rhyme generation

## I. INTRODUCTION

Neural Network is an information processing paradigm that is inspired by the way biological nervous systems, such as the brain, process information. The key element of this paradigm is the novel structure of the information processing system. It is composed of a large number of highly interconnected processing elements (neurones) working in unison to solve specific problems. It learns by example.

The history of neural networks that was described above can be divided into several periods. There were some initial simulations using formal logic. McCulloch and Pitts (1943) developed models of neural networks based on their understanding of neurology. These models made several assumptions about how neurons worked. Their networks were based on simple neurons which were considered to be binary devices with fixed thresholds. The results of their model were simple logic functions such as "a or b" and "a and b". Significant progress has been made in the field of neural networks-enough to attract a great deal of attention and fund further research. Advancement beyond current commercial applications appears to be possible, and research is advancing the field on many fronts. Neurally based chips are emerging and applications to complex problems developing. Clearly, today is a period of transition for neural network technology.

To design a neural network that has been trained on by songs of any artist by feeding lyrics and produce a new song word by word that rhymes and has a flow. This is a new way of applying machine learning attributes to a new region of people's interest. Having a system that comes up with new songs which is similar to your liking boosts up a person's mood. Having machines serve as composers actually opens the music world up to a "new creativity" that isn't possible otherwise. Latest researches are going on to write and publish new albums which give tough competitions to old hits by learning the lyrics of those songs. One of them is carried out by Sony CSL known as Flow Machines.

## II. LITERATURE REVIEW

Paper	Methodology	Conclusion
Text based LSTM networks for Automatic Music Composition [1]	<ul style="list-style-type: none"> <li>The proposed method includes two LSTM layers and the Keras deep learning framework.</li> <li>The proposed network is designed to learn relationships within text documents that represent chord progressions and drum tracks in two case studies.</li> <li>The proposed system can be used for fully automatic composition or as semi-automatic systems that help humans to compose music by controlling a diversity parameter of the model.</li> </ul>	They introduced an algorithm of text-based LSTM networks for automatic composition and reported results for generating chord progressions and rock drum tracks. Word-RNNs showed good results in both cases while char-RNNs only successfully learned chord progressions. The experiments show LSTM provides a way to learn the sequence of musical events even when the data is given as text.

<p>Lyric Generation using artificial intelligence [2]</p>	<ul style="list-style-type: none"> <li>• Lyric Generator is application that can be used to generate lyrics automatically i.e. all by itself.</li> <li>• Here the user only needs to describe the scenario of the lyric that is the mood, relationship, character description, rhyming, similarities, and inspirations for the lyrics.</li> <li>• These are certain criteria by which the computer or the application understands about what exactly we want it to write for us.</li> <li>• It is very essential that we provide lot of knowledge to the system prior to the creation of the application.</li> <li>• Updating the database of lyrics is also important. The database may vary from one language to the other but the concept remains the same.</li> </ul>	<p>In this paper, a creation technique has been proposed which can be used to create a lyric by using the ten prominent features of the lyrics. These Features have a pattern and we can create various patterns using appropriate algorithms. It is important to specify the meaning and usage of a word in a computer database. This application can be used to generate lyrics to any language provided the database has been loaded correctly. Future work includes techniques for finding similes and metaphors.</p>
<p>A survey on intelligent poetry Generation – Languages, Features, Techniques, Reutilization and Evaluation [3]</p>	<ul style="list-style-type: none"> <li>• Poetry generation is becoming popular among researchers of Natural Language Generation, Computational Creativity and, broadly, Artificial Intelligence.</li> <li>• To produce text that may be regarded as poetry, computational systems are typically knowledge- intensive and deal with several levels of language.</li> <li>• This paper surveys intelligent poetry generators around a set of relevant axis – target language, form and content features, applied techniques, reutilisation of material, and evaluation – and aims to organise work developed on this topic so far.</li> </ul>	<p>Intelligent poetry generators were surveyed in this paper, around a set of relevant axis where alternative approaches have been explored. Poetry has been automatically generated in different languages and forms, considering different sets of features, and through significantly different approaches. Poetry generators have been developed with different goals and intents, each with their stronger and weaker points, which adds to the subjectivity involved in the evaluation of poetry, even for humans.</p>
<p>Algorithmic Song Writing with ALYSIA [4]</p>	<ul style="list-style-type: none"> <li>• This paper introduces ALYSIA: Automated Lyrical Song- writing Application.</li> <li>• ALYSIA is based on a machine learning model using Random Forests, and we discuss its success at pitch and rhythm prediction.</li> <li>• Next, it is shown how ALYSIA was used to create original pop songs that were subsequently recorded and produced.</li> <li>• Finally, vision for the future of Automated Song writing for both co-creative and autonomous systems has been discussed.</li> </ul>	<p>Algorithmic song writing offers intriguing challenges as both an autonomous and a co-creative system. An autonomous song writing system producing works on par with those of expert human songwriters would mark a significant achievement.</p>

<p>Automatically Generating Rhythmic Verse with Neural Networks [5]</p>	<ul style="list-style-type: none"> <li>• Two novel methodologies for the automatic generation of rhythmic poetry in a variety of forms have been proposed.</li> <li>• The first approach uses a neural language model trained on a phonetic encoding to learn an implicit representation of both the form and content of English poetry.</li> <li>• This model can effectively learn common poetic de-vices such as rhyme, rhythm and alliteration.</li> <li>• The second approach considers poetry generation as a constraint satisfaction problem where a generative neural language model is tasked with learning a representation of content, and a discriminative weighted finite state machine con- strains it on the basis of form.</li> </ul>	<p>First, they developed a neural language model trained on a phonetic transliteration of poetic form and content. Although example output looked promising, this model was limited by its inability to generalise to novel forms of verse. They then proposed a more robust model trained on unformed poetic text, whose output form is constrained at sample time. This approach offers greater control over the style of the generated poetry than the earlier method, and facilitates themes and poetic devices.</p>
<p>Automated Composition of Lyrical Songs [6]</p>	<ul style="list-style-type: none"> <li>• They address the challenging task of automatically com- posing lyrical songs with matching musical and lyrical features, and present the first prototype to accomplish the task.</li> <li>• The focus of this paper is especially on generation of art songs.</li> <li>• The proposed approach writes lyrics first and then composes music to match the lyrics.</li> <li>• Some example songs composed by M.U. Sicus have been proposed, first steps towards a general system combining both music composition and writing of lyrics have been outlined.</li> </ul>	<p>They have proposed the task of generating lyrical songs as a research topic of computational creativity. This topic has received only little attention in the past although both music composition and poetry/lyrics generation have been studied on their own. An automatic generation procedure of lyrical and musical content also offers interesting possibilities for medicalization of data</p>
<p>Automatic Generation of Poetry inspired by Twitter Trends [7]</p>	<ul style="list-style-type: none"> <li>• This paper revisits PoeTryMe, a poetry generation platform, and presents it's most recent instantiation for producing poetry inspired by trends in the Twitter social network.</li> <li>• The presented system searches for tweets that mention a given topic, extracts the most frequent words in those tweets, and uses them as seeds for the generation of new poems.</li> <li>• Generation is performed by the classic PoeTryMe system, based on a semantic network and a grammar, with a previously used generate &amp;test strategy.</li> <li>• Illustrative results are presented using different seed expansion settings.</li> </ul>	<p>A new instantiation of PoeTryMe, a poetry generation platform, was presented. The singular feature of the presented system is that its poetry is inspired by Twitter trends, more precisely, words that are associated with those.</p>

<p>GhostWriter: Using an LSTM for Automatic Rap Lyric Generation [8]</p>	<ul style="list-style-type: none"> <li>• This paper demonstrates the effectiveness of a Long Short-Term Memory language model in our initial efforts to generate unconstrained rap lyrics.</li> <li>• The goal of this model is to generate lyrics that are similar in style to that of a given rapper, but not identical to existing lyrics: this is the task of ghostwriting.</li> <li>• Unlike previous work, which defines explicit templates for lyric generation, this model defines its own rhyme scheme, line length, and verse length.</li> <li>• The experiments show that a Long Short-Term Memory language model produces better “ghost-written” lyrics than a baseline model.</li> </ul>	<p>The performance of the LSTM model has been compared to a much simpler system: an n-gram model. The results of our experiments show that, as an unsupervised, non-template model, the LSTM model is better able to produce novel lyrics that also reflect the rhyming style of the target artist. In future work, we plan to use more data to train our model, making it easier for our system to actually identify rhyming pairs and use them in new contexts.</p>
<p>Automatic Generation of Lyrics in Bob Dylan’s Style [8]</p>	<ul style="list-style-type: none"> <li>• It is interesting to see if machine could learn Bob Dylan’s poetic style by looking at his lyrics.</li> <li>• In this Paper, N-grams and Recurrent Neural Network (RNN) with Long Short Term Memory (LSTM) have been used to model Dylan’s lyrics, and eventually use the algorithms to generate samples of lyrics in Bob Dylan’s style.</li> </ul>	<p>For N-gram training perplexity decreases as N increases because larger N leads to fewer choices of the next word, hence higher probability to select the right word. Character-level RNN seems good at capturing the grammar of the sentences, but may be weak in generating text that makes sense in the context.</p>
<p>SMUG: Scientific Music Generator [9]</p>	<ul style="list-style-type: none"> <li>• Music is based on the real world. Composers use their day-to-day lives as inspiration to create rhythm and lyrics.</li> <li>• Procedural music generators are capable of creating good quality pieces, and while some already use the world as inspiration, there is still much to be explored in this.</li> <li>• A system has been described to generate lyrics and melodies from real-world data, in particular from academic papers.</li> <li>• Through this they want to create a playful experience and establish a novel way of generating content.</li> <li>• For melody generation, they present an approach to Markov chains evolution and briefly discuss the advantages and disadvantages of this approach.</li> </ul>	<p>They have presented a method for creating melody and lyrics using real-world data. To do so, they developed a musical generator that evolves Markov chains to create melodies, and a lyric generator, that extracts content from academic papers and transforms them into songs. Hence there is fully functional system that completes tasks, taking an academic paper in PDF format and outputting a melody and the according lyrics. The generator seems to produce interesting music/lyrics combinations, but they still have to conduct further studies to prove their interestingness.</p>

### III. PROPOSED ARCHITECTURE AND FRAMEWORK

#### A. Recurrent Neural Networks

The basic feed forward network, there is a single direction in which the information flows: from input to output. But in a recurrent neural network, this direction constraint does not exist. There are a lot of possible networks that can be classified as recurrent, but we will focus on one of the simplest and most practical.

Basically, what we can do is take the output of each hidden layer, and feed it back to itself as an additional input. Each node of the hidden layer receives both the list of inputs from the previous layer and the list of outputs of the current layer in the last time step as shown in Fig. 1. This, after unwrapping along the time axis looks like as shown in Fig. 2.

In this representation, each horizontal line of layers is the network running at a single time step. Each hidden layer receives both input from the previous layer and input from itself one time step in the past. The power of this is that it enables the network to have a simple version of memory, with very minimal overhead. This opens up the possibility of variable length input and output: we can feed in inputs one- at-a-time, and let the network combine them using the state passed from each time step.

One problem with this is that the memory is very short-term. Any value that is output in one time step becomes input in the next, but unless that same value is output again, it is lost at the next tick. To solve this, we can use a Long Short-Term Memory (LSTM) node instead of a normal node. This introduces a “memory cell” value that is passed down for multiple time steps, and which can be added to or subtracted from at each tick. To create multileveled equations, it may be necessary to treat the equation as a graphic and insert it into the text after your paper is styled.

**B. Proposed System**

First of all the dataset is cleaned using Natural Language Processing Tool Kit. We have two datasets of positive words and negative words, which are compared with each line in the dataset. According to the calculated score of each word, negative sentences are removed from the dataset. This new filtered dataset is given to the neural network for computation. Stemmer function is a part of nltk library suit and it is import as the LensesterStreamer in the code, this function removes all the grammatical changes of the words and convert all the words into a root word. Tokenize function is used to tokenize the words in the sentences that is making an array of the words from the sentence.




Figure 1. Recurrent Neural Network




Figure 2. Recurrent Neural Network after unwrapping along time axis

A Markov chain is a stochastic model describing a sequence of possible events in which the probability of each event depends only on the state attained in the previous event. Thus the working of the model through Markov chain is through the states of lyrics obtained from numerous songs collected. This being said, the probability of every word occurring in a specified pattern is calculated, the rhyming words are played after certain time with huge probability and the lyrics are composed. The example is shown in Fig. 3.

Therefore, we see that according to data collected from numerous lyrics of the genre “rap” of the same singer, there is a high probability of getting two rhyming words separated by sentences, rather than having both words said one after the other.

The probability of “Bore” being joined with “adore” =  $P(\text{Adore} \rightarrow \text{Bore}) = 0.1$ . The probability of “Bore” being joined with a sentence first =  $P(\text{Adore} \rightarrow \text{“some sentence”} \rightarrow \text{Bore}) = 0.81$ .

Clearly the Markov chain helps to follow the specific rules of rap music, following the sequence of pattern regular rap music follows and the way a composer would write down the lyrics for it to rhyme.

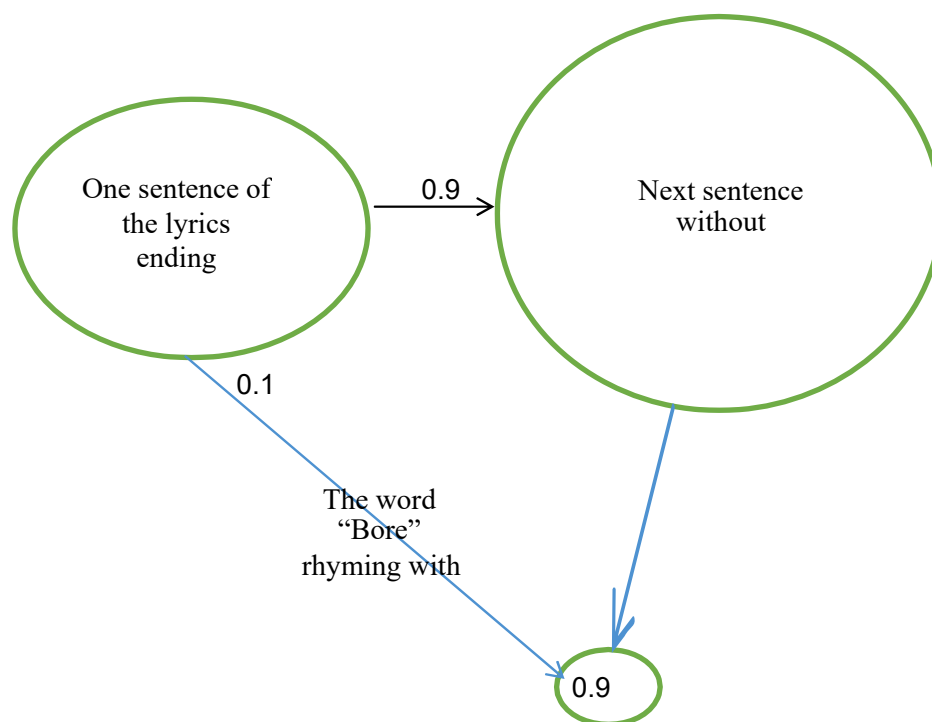


Figure 3. Example showing the conversion of word to rhyme using probability

#### IV. IMPLEMENTATION

##### A. Steps involved in the Proposed System

- **Data Preparation** – This stage involves collecting, compiling and cleaning the lyrics we want to feed into the neural networks.
- **Training** – Train the network with the fed lyrics.
- **Generating Lyrics - from** the already trained network.
- **Making Music** – Producing songs by using a text-to-speech conversion over a beat.

We used Python Anaconda Navigator interface and various python libraries for implementing the proposed system such as

- Keras library, is a deep learning and high level neural networks library in Python with various models.
- Markovify library which uses Markov chains to generate random semi-plausible sentences based on an existing text.
- pyttsx is a Python package supporting common text-to-speech engines.

Implementation of the proposed system begins with training a recurrent neural network on a very tough task of language modeling. We implement this on a set of lyrics by a rap artist, namely, Kanye West. We take his discography as a text file which becomes the input of the program.

The final aim is to fit a probabilistic model whose task would be to assign probabilities to the sentence and we use a markov chain model for that. Now this will be carried out by predicting the next words of the lyrics given the whole input of the discography of lyrics which will be used to learn i.e. use as a history of previous words.

The core of this model will be made of five LSTM cells. LSMT being the most widely used RNN and they are way better at capturing the dependencies, in this case from the set of lyrics inputted as the training set. These LSTM cells will invite words and then help predict the outcome of the next word by using the probabilities.

Now, by the method of design the outputs from the LSTM cells or in RNN in general are dependent on the set of random inputs. Although this makes back propagation difficult but in our case five epochs are done. One epoch will occur after one forward pass of the words followed by a step of back propagation to take the error into account and reduce it. That means, one epoch will occur after one forward and one backward pass as in Fig.4.

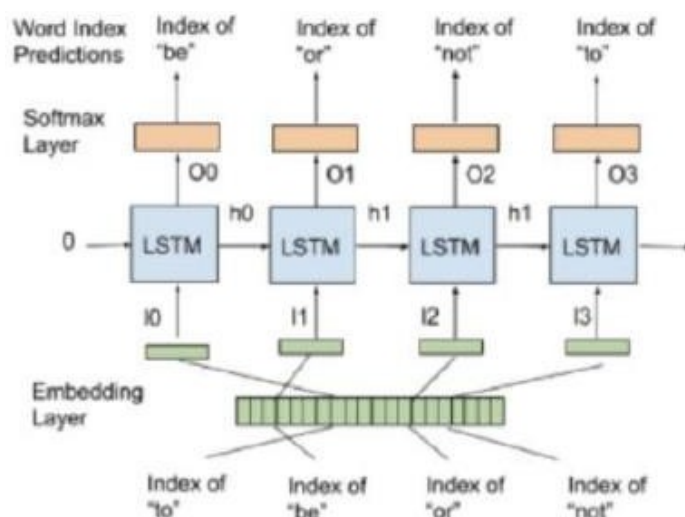


Figure 4. Architecture of RNN using LSTM

After a sequence of tuples is generated this will be in the form of (desired rhyme, desired count of syllable). The program will then sift through the lines the markov chain generated, and match the lines to their corresponding tuples.

#### B. Implementation Explanation

We take a set of lyrics or a set of text as an input which we eventually feed to the Recursive Net.

The code is designed in a way that it takes the lyrics that we have entered and a Markov Chain looks at it and generates new lines.

Further, this will be fed to a Recursive Neural Network which will create tuples in the form of (desired rhyme, desired count of syllable)

For desired syllable

The neural net generates all the sets of syllables in the input set. To perform this task it takes the following steps:

1. It strips out all the punctuations from the input set to avoid any discrepancy.
2. It then takes the set of all the vowels and picks out words with different vowel-vowel or a vowel-consonant or a consonant-consonant combination in them.
3. It ignores those rhyming syllables which are not a combination of a pair-pair syllable rhyme. For example: It will not match the syllable sound of “e” with the syllable sound of “le” even though they have the same sound because they are not in accordance with a pair-pair syllable match scenario.

For desired rhyme

A rhyme scheme is the pattern of rhymes that is being used, whenever we are constructing a musical flow.

We need to understand how the genre of rap music is based on rhyming lyrics, be it words or syllables. The better the rhymes the better the lyrics.

So we basically desire a combination of rhymes in a sentence to be able to create a rap song. To do just that, we take the following steps:

1. We decide how many maximum syllables we want in a single line. This is important as this will in turn determine how many rhyming syllables we will have in our lines.
2. It will then find out all the rhyming syllables in the set of generated syllables and identify them to be used in the, to be generated rap song.
3. The Recursive Neural Network will go through five Epochs to minimize the loss in each cycle and in turn find out the best combinations for the, to be generated rap song.

All these rhymes will be then saved in a .rhyme type file and be later used in the program. To understand this better let us take a verse and try and match all the pair-pair syllable sounds in it; that is, a two pair syllable sound.

Penitentiary chances, the devil dances/ 12,—2/  
 and eventually answers, to the call of Autumn/ -12,—3/  
 all of them fallin', for the love of ballin'/ 33—3/  
 got caught with 30 rocks the cop look like Alec Baldwin/ —4, 4-3/

We can see here that this piece of lyrics contain 2 consecutive multi-syllable rhymes in the beginning and then ending them with a rhyme that contains the same number of syllable as the previous of the same rhyme.

#### V. CONCLUSIONS

The following limitations were encountered.

1. **If the dataset is small, the machine will not be able to learn anything from it**, because the idea of a data set is to train the machine, so a large data set is needed. Even if the machine does learn anything, it will over fit or there will be a certain outlier condition and the answer won't be in accordance to our design.
2. **Our model's predictions are accurate, but often recycled.** It's important to note that many of our predicted lines turned out to be nearly identical to lines Logic has actually written, i.e. half of a line from one verse/song combined with half of a line from another verse/song. This is to be expected, as using bigrams yields less variability in predicted words due to basing predictions off the previous two words instead of the one most recent, resulting in sequences of three or more words coming from the same Logic lyric. To put it simply, using bigrams instead of single words increases readability and similarity to Logic's style, but decreases creativity.
3. **Our model is slower and generates less output.** The unigram model runs faster because the dictionary object representing its Markov Chain has far fewer keys. Our model has so many more keys because it has to process tuples of two words. Furthermore, as mentioned before, there were times when we received very little to no output, and generally we received less than we did from the unigram implementation. This can be attributed to the smaller number of possibilities for the next word when we're basing it off the previous two words.

The range of algorithms available for ML problem solving is astounding. Random forests, support vector machines, neural networks, and Bayesian estimation methods – the list goes on (and on, and on). The question of what algorithm is best for a given ML problem, however, is often less impactful than we might think.

It's true that some approaches, on some questions, will work better than others. In some cases, this difference can even be quite distinct. However, in my experience it's been quite rare that one modelling approach will strictly dominate all other options in answering a given ML question.

Therefore we also have faced cons of applying the markov model here.

#### REFERENCES

#### AUTHORS PROFILE

Dr.Thendral Puyalnithi, Associate Professor in Kalasalingam Academy of Research and Education has completed Bachelor of Engineering in PSG Tech, Coimbatore in 2005 in Electrical and Electronics Engineering , then completed his Master of Engineering(ME) in BITS Pilani in Software Systems and he has completed PhD in Computer Science in VIT University Vellore. He has previous worked in software companies such as Infosys and Samsung India with designation of Senior Software Engineer for 3 years and he was an Assistant Professor Senior in VIT University, Vellore for 9 years. His area of interest includes Computational Intelligence and Algorithm Design analysis and Programming languages and he has publications in these domains.

Murugeswari Kandavel received her Bachelor of Engineering in Computer Science and Engineering, Master of Technology in Information Technology and Ph.D in Information Communication Engineering in 1996, 2003 and 2018, respectively, from Madurai Kamaraj University, University of Punjab and Anna University vice versa. She started her carrier as a Lecturer in 1997 and continues her service in teaching field. She worked in various Engineering Institutions for the past 22 years and had several publications in National and International level journals. Currently, she is working as an Associate Professor in Kalasalingam Academy of Research and Education. Her research interests include Information Security, Steganalysis and Soft Computing techniques.