



# Learning to Automatically Generating Genre-Specific Song Lyrics: A Comparative Study

Tze Huat Tee, Belicia Qiao Bei Yeap, Keng Hoon Gan<sup>(✉)</sup>, and Tien Ping Tan

School of Computer Sciences, Universiti Sains Malaysia, USM, 11800 Pulau Pinang, Malaysia  
{tee.tze.huat.ucom20,beliciayqb}@student.usm.my, {khgan,  
tienping}@usm.my

**Abstract.** The impact of music on the many dimensions of human life can be partly attributed to its linguistic component – the lyrics. In hopes of helping song-writers reach their full potential, researchers have implemented advanced artificial intelligence (AI) technology to automatically generate song lyrics. These efforts, however, were met with challenges that accompany the distinctive qualities of song lyrics, such as word repetition, structural patterns, and line breaks; all of which are dependent on the music genre. Seeing as most previous research either focuses on a given approach or genre, or performs the task without consideration of lyric variation among genres, this study attempts to address the gap by exploring and comparing the capabilities of three promising methods, specifically Markov chains, long short-term memory (LSTM), and gated recurrent units (GRU), in algorithmically generating lyrics for six selected music genres, namely rock, pop, country, hip-hop, electronic dance music (EDM), and rhythm and blues (R&B). Our findings show that LSTM scored better in the average readability index in overall, however, GRU produced the overall highest Rhyme Density score.

**Keywords:** Lyrics · Text generation · Natural Language Processing (NLP) · Markov chains · Long Short-Term Memory (LSTM) · Gated Recurrent Units (GRU)

## 1 Introduction

It is generally believed that music has been deeply ingrained in our societies since the dawn of humanity, with a significant amount of ancient musical instruments dating back as far as the Middle and Upper Palaeolithic [1]. Indeed, the tremendous influence music has on people of all ages from pre-schoolers [2], to adolescents [3], and to seniors [4]; is undeniable.

One of the fundamental elements of music is its linguistic content, i.e., the lyrics. In addition to intensifying emotions such as sadness, nostalgia, and astonishment, song lyrics have been observed to activate certain psychological mechanisms, including episodic memory, evaluative conditioning, contagion, and visual imagery [5].

Moreover, despite being initially centered on a limited number of themes, lyrics have, since the 1960s, evolved into a vessel for writers and performers to convey a broad spectrum of symbolic messages [6]. In particular, a considerable number of artists have leveraged the capacity of song lyrics to raise awareness in important issues, such as mental health, gender equality, and racial harmony [7].

## 2 Problem Statement

With the purpose of helping songwriters overcome the many challenges of lyric writing, notable efforts in the automatic generation of song lyrics have been made. Nevertheless, the application of artificial intelligence (AI) to the writing of lyrics has been proven to be no easy feat. Due to its unique features, an in-depth understanding of songwriting techniques, on top of sound knowledge in natural language processing (NLP), is crucial [8]. The necessity of modelling the line breaks, stylistic elements (e.g., flow, rhyming, and repetition), and structural layout and components (e.g., verse, refrain, chorus, and bridge) observed in lyrics adds another layer of complexity to the already difficult task [9]. Furthermore, these linguistic attributes may vary among different music genres. For instance, it has been demonstrated that rap songs incorporate significantly more word repetition as compared to country songs [9].

Regardless of the intricacies, several methods, including Markov chains [10], long short-term memory (LSTM) [9], and gated recurrent units (GRU) [11], have been shown to produce promising results on separate occasions. Therefore, it would be interesting to expand on previous research, such as that of Gill et al., and explore and compare the performance of these three approaches in the algorithmic generation of song lyrics. In this case, sub-genres are classified into their parent genres (e.g., categorizing metal as part of rock) due to computational and time constraints. This study thus focuses on six popular music genres of the English language, namely rock, pop, country, hip-hop, electronic dance music (EDM), and rhythm and blues (R&B) [12, 13].

## 3 Literature Review

### 3.1 Generating Non-genre-specific Lyrics

In 2010, Settles presented two interactive computational creativity tools designed to aid the song-writing process – Titular, a text synthesis algorithm capable of generating song titles semi-automatically, and LyriCloud, which displays a cloud of suggested lyrics based on a word input [14]. These intelligent tools were developed based on the criteria that their recommendations should be both unlikely and meaningful. Although the results were semantically satisfactory, they failed to exhibit any notion of stylistic qualities such as lyrical wordplay (e.g., rhyme) and other devices of creative writing (e.g., repetition).

On the other hand, Pudaruth et al. attempted to generate the lyrics of an entire song using context-free grammars (CFGs) [8]. By imposing grammatical rules and statistical constraints, they successfully produced lyrics that were grammatically correct and rather convincing, with more than half (52%) of their respondents evaluating one of their generated lyrics as an existing song. However, their output often lacked semantic meaning

due to the impossibility of defining all grammatical rules which exist in the English language.

The studies above approached the task at hand without taking into account the influence of the genre on a song's lyrics, though Pudaruth et al. examined a few themes (i.e., love, pain, and cause) commonly found in popular songs [8]. Since writing is usually performed with an audience in mind [9], capturing the differences among genres, be it semantically or stylistically, could be an essential matter.

### 3.2 Generating Lyrics for a Specific Genre

An article published by Barbieri et al. in 2012 describes a framework of Constrained Markov Processes which generates lyrics in the style of a particular writer while maintaining the structural properties (in terms of rhyme and meter) of a provided text [10]. Apart from these features, their demonstration of mapping Bob Dylan's songwriting style onto the structure of the Beatles' "Yesterday" showed syntactic correctness and semantic relatedness. Nevertheless, additional cases should be investigated to ensure that this technique can be generalized to different writers and styles.

A more recent study by Fernandez et al. compared the performance of three character-level deep learning models, namely plain recurrent neural network (PRNN), long short-term memory (LSTM), and gated recurrent units (GRU), in the composition of rap lyrics [11]. The resulting lyrics achieved positive overall evaluation, convincing 67% of the participants who are familiar with rap lyrics in one of the instances, in spite of low rhyme density. Consequently, they suggested incorporating rhymes and intelligibility in the algorithm to improve rhythmic flow and coherency.

Despite promising results, these methods were formulated to address the issue for a specific genre (e.g., rap). In view of the broad spectrum of music preferences, it would perhaps be useful to explore the application of these approaches to other genres to appeal to a wider audience.

### 3.3 Generating Lyrics for Multiple Genres

In 2020, Gill et al. proposed a method which uses state-of-the-art long short-term memory (LSTM) to automatically generate lyrics for a specified music genre [9]. Upon evaluating their output using linguistic metrics, it was found that their model performed better in capturing the characteristics of pop and rap lyrics, in comparison to other genres such as rock, metal, country, and jazz. Seeing as only a single technique, i.e., LSTM, was employed, further research should be conducted to explore and compare the potential of other algorithms in computationally composing lyrics of various genres.

## 4 Methodology

The following section consists of descriptions of the dataset used in this study as well as details regarding data pre-processing, exploration, and cleaning.

## 4.1 Dataset Description

The dataset is self-collected by using Geniuslyrics API (Genius 2020) and Spotify Web API.

At the beginning, an account is required in Spotify to request access to Spotify Web API. After Spotify verified and approved the application, the client key and client secret are granted for access to Spotify Web API. By using Spotify Web API, the categories provided in Spotify playlist are retrieved and the genre of each playlist (rock, pop, country, hip-hop, EDM, or R&B) is identified. Following that, the track details are extracted from the identified playlist.

On the other hand, setting up an account in the Genius Lyrics Website authorized the access to apply for API Clients. A new API Client can be created with the application name and application website URL information. Upon confirmation of the API Client, the page generated a Client ID and Client Secret that authorize the usage of Geniuslyrics API.

Once the Client ID and Client Secret are provided, the lyricsgenius package in Python called the API and scraped the lyrics based on the track details retrieved from Spotify Web API. To avoid duplication of songs, a filter is added to skip live, demo, and remix versions in the scraping process. The relevant attributes of the collected dataset are as described in Table 1.

**Table 1.** Description of attributes.

Attribute	Description	Data type
id	The unique identifier for the track generated by Spotify	Nominal categorical data
name	The title of the track	Nominal categorical data
artist	The artist(s) who performed the track	Nominal categorical data
album	The album of the track	Nominal categorical data
lyric	The lyric of the track	Nominal categorical data
genre	The music genre of the artist(s)	Nominal categorical data
popularity	The popularity of the track based on the total number of plays and the recency of plays	Continuous numerical data

## 4.2 Data Pre-processing

As mentioned above, the song lyrics are collected by using web scraping API, Genius-lyrics. Within the scraped data, there are unwanted strings such as “EmbedShare”, “URL-CopyEmbedCopy”, and new line “\n” etc. All the unwanted string are replaced with a space. Other than that, the null data for the lyric column is removed and only the top 100 rows being selected as our dataset in this experiment.

Next, the lyrics strings are converted to lower case and punctuation is removed. Finally, tokenization breaks the lyrics strings into tokens.

## 4.3 Data Analysis

Text analysis of the song lyrics is carried out to further understand the six different music genres in terms of their linguistic content. The most common words in the song lyrics are identified and visualized in a word cloud for each genre. Apart from that, bar charts are also created to visualize the frequency distribution of the number of words in the song lyrics for each genre.

As shown in Table 2, the highest average word length in song lyrics can be seen in hip-hop. On top of that, hip-hop also has the highest average unique word counts. This indicates that hip-hop has the highest complexity among all genres and could possibly impact the model performance.

Besides that, the genres having the highest and second highest noun term frequencies can be determined as hip-hop and EDM respectively. These two genres are also the highest and second highest in terms of verb term frequencies in song lyrics. Thus, it can be deduced that the noun term frequencies and verb terms frequencies in song lyrics are correlated to each other.

Since the usage of adverbs in song lyrics are relatively close for every genre, this characteristic plays an insignificant role in analytics.

Interestingly, EDM has the greatest maximum number of words (3980) as well as the lowest minimum number of words (37) in song lyrics. In contrast, pop and R&B seem to have rather short lyrics in general as shown by their maximum number of words.

**Table 2.** Text analysis of lyrics

Characteristic	Genres & counts
Average word length in song lyrics	Rock: 282.44 Pop: 368.69 Country: 323.69 Hip-Hop: 581.32 EDM: 332.37 R&B: 444.04

*(continued)*

**Table 2.** (continued)

Characteristic	Genres & counts
Average unique word counts in song lyrics	Rock: 110.06 Pop: 132.25 Country: 139.77 Hip-Hop: 254.13 EDM: 107.42 R&B: 159.61
Noun term frequencies in song lyrics	Rock: 1188 Pop: 1209 Country: 1350 Hip-Hop: 2423 EDM: 1388 R&B: 1066
Verb term frequencies in song lyrics	Rock: 630 Pop: 637 Country: 662 Hip-Hop: 1124 EDM: 775 R&B: 678
Adverb term frequencies in song lyrics	Rock: 214 Pop: 231 Country: 232 Hip-Hop: 271 EDM: 232 R&B: 233
Maximum number of words in song lyrics	Rock: 1390 Pop: 748 Country: 1019 Hip-Hop: 1125 EDM: 3980 R&B: 747
Minimum number of words in song lyrics	Rock: 136 Pop: 167 Country: 140 Hip-Hop: 281 EDM: 37 R&B: 196

Figure 1 illustrates the word cloud generated from the lyrics of the collected country songs. From the diagram, the outliers and most common terms, such as “got”, “yeah”, “oh”, and “know”, are identified; all of which will introduce bias to the model.

The bar chart in Fig. 2 depicts the frequency distribution of the number of words in lyrics of the selected country songs. Based on Fig. 2, most of the number of words are




Fig. 1. Word cloud of country song lyrics.

scattered between 200 to 400. An outlier where the number of words is more than 1000 can also be observed but it only occurred once.




Fig. 2. Frequency distribution of the number of words in country song lyrics.

#### 4.4 Markov Chains

A Markov chains model is a statistical tool that identifies the pattern dependencies in different kinds of systems, especially pattern recognition system [15]. As characters or words are normally characterized by dependencies between patterns, the Markov chain theory is suitable for implementation in the domain of natural language processing.

Markov chains is selected in our study as it is one of the basic methods for text generation. The core idea of Markov chains is a simple assumption that the next word is dependent on the previous word.

First, the song lyrics is tokenized into each token. Then, a dictionary is initialized to hold all the words and next words. After that, all the words will pair up with the next

word and they will be stored in the previously created dictionary. Finally, a function can be created to generate consecutive words upon receiving an input text by referring to the dictionary iteratively. For Markov chain, the output will be measured based on the readability and density score.

#### 4.5 Long Short-Term Memory (LSTM)

Long short-term memory (LSTM) is a type of recurrent neural network (RNN) that is able to learn the order dependencies that exists in sequence prediction problems [16]. These networks were introduced by Sepp Hochreiter et al. in 1997 [17]. A memory unit known as a “cell state” is introduced in the LSTM to address the existing failure of RNN in learning the presence of past observations that is greater than 5–10 discrete time steps between relevant inputs and their target signals [18]. The cell state acts as a carrier to transfer information or context over longer discrete steps, hence allowing adjustment of the network gradient descent in the information flow.

The layers of our trained LSTM model that output the best results with the limitation in hardware specification are as described in Table 3.

**Table 3.** Summary of LSTM model for pop genre.

Layer	Output	Param#
Embedding	(None, 756, 40)	117160
LSTM (bidirectional)	(None, 200)	112800
Dropout	(None, 200)	0
Dense	(None, 2929)	588729

The model trained with 30 epochs and achieved a range of accuracy from 0.6446 to 0.773 based on different kinds of genres. Then, the model is implemented to predict the class from the generated token list and the output will become the newly generated song lyrics.

#### 4.6 Gated Recurrent Units (GRU)

In 2014, Kyunghyun Cho et al. introduced gated recurrent units (GRU), which is an improvement of the standard RNN [19]. This is a relatively new method compared to RNN and LSTM as it is an improved version of them. GRU able to perform well in sequence learning tasks and handling the vanishing gradients problem seen in traditional RNN [20].

Compared to LSTM, GRU implements gates to control the flow of information and abandons the usage of cell states. GRU consists of only one hidden state and has a simpler architecture, thus it will shorten the training time of the model [21].

The layers of our trained GRU model that output the best results with the limitation in hardware specification are as detailed in Table 4.



**Table 4.** Summary of GRU model for pop genre.

Layer	Output	Param#
Embedding	(None, 719, 40)	102520
GRU (bidirectional)	(None, 200)	85200
Dropout	(None, 200)	0
Dense	(None, 2563)	515163

All models trained with 30 epochs with one exception of the EDM genre, which experienced early stopping in 25 epochs. They achieved a range of accuracy from 0.6993 to 0.8198 based on different types of genres. Then, the model is implemented to predict the class from the generated token list and the output will become the newly generated song lyrics.

## 5 Evaluation Criteria

Three evaluations were performed in this paper, namely model performance, average readability, and rhyme density score. In addition, we have also included sample of generated song lyrics from Markov chain, LSTM and GRU.

### 5.1 Model Performance

Based on Table 5, the GRU models for every genre slightly outperformed the LSTM models in terms of accuracy after 30 epochs. Due to hardware limitations (which will be further elaborated in the discussion section), the epoch is set to 30 as the maximum value. Therefore, it is believed that the LSTM models require more epochs to achieve higher accuracies based on the theory stated above.

**Table 5.** Comparison of model performance.

Genre	LSTM accuracy	GRU accuracy
Pop	0.7420	0.8068
Rock	0.7426	0.8198
Country	0.6945	0.7738
EDM	0.7773	0.8142
Hip-Hop	0.6517	0.6993
R&B	0.6698	0.7271

## 5.2 Average Readability

Readability is the ease with which a reader is able to understand a written text and is measured by the complexity of the text's vocabulary and syntax [22]. In this experiment, the average readability of the generated lyrics is obtained by using the textstat library in Python. The higher the average readability, the better the generated song lyrics.

Based on Table 6, the average readability of generated lyrics for Markov chains are the highest in every genre. As a result of the stored dictionary that is implemented in Markov chains, the fixed structural and grammatical rules in the Markov chains approach enable it to obtain high scores in average readability. In the meantime, LSTM model outputs managed to score better than GRU model outputs in 4 different genres such as pop, rock, country, and R&B. However, GRU model outputs score better for the EDM and Hip-Hop genres that have huge number of tokens. As the GRU model trains faster along the epochs, the model is determined capable to handle the higher complexity and huge dimension dataset.

**Table 6.** Average readability of generated lyrics.

Genre	Average readability (Markov chains)	Average readability (LSTM)	Average readability (GRU)
Pop	21.42	20.92	20.32
Rock	22.88	22.10	20.62
Country	21.84	21.62	20.04
EDM	22.60	18.9	23.12
Hip-Hop	22.52	19.78	20.60
R&B	22.30	20.52	19.62

## 5.3 Rhyme Density Score

Rhyme density score referred to the total number of rhymed syllables that divided by total number syllables in the corpus or song lyrics in our case [23]. It is part of evaluation criterion to determine whether which approach able to generate the best lyrics as the output. For this measurement, the higher the rhyme density represent the better the generated song lyrics.

Referring to Table 7, the GRU model has the highest score for Rhyme density score in overall. In the meantime, the Markov chains score the lowest due to the randomness retrieval from the stored dictionary and form the lyrics. Besides that, the pop genre songs more likely to score higher compared to the other genre. It could be due to the chorus and word repetition in the pop genre songs.

**Table 7.** Rhyme density scores of generated lyrics.

Genre	Rhyme density score (Markov chains)	Rhyme density score (LSTM)	Rhyme density score (GRU)
Pop	0.6250	0.7114	0.7482
Rock	0.5652	0.5833	0.6715
Country	0.4385	0.5431	0.6419
EDM	0.5349	0.7634	0.5082
Hip-Hop	0.5556	0.6207	0.6275
R&B	0.6667	0.4715	0.5912

## 5.4 Sample Output of Generated Song Lyrics

### Markov Chains Sample Output (Pop Genre)

*Breathing just rub it never wanna keep you first baby, let's get your bad 'cause I got it, got me be alone in my records on everything seems like you leave me, girl? not the things that I never does why? you are you so I see one is that.*

### LSTM Sample Output (Pop Genre)

*Happy for me out of myself I am I think I'm gonna get so I've been thinking I know what you know that I was born to run I don't belong to everybody but you're not to me I don't deserve someone loyal to me I don't want to be a*

### GRU Sample Output (Pop Genre)

*Sad so don't say oh woah oh but yeah I hate you I don't wanna be my spot I've been work out baby it's just like this might be so bitter ooh ooh ooh ooh ooh just sayin' this what you know that you're hiding something I know it's true it's*

## 6 Discussion

### 6.1 Models

Throughout the processes, all the methods are compared to each other based on their differences, time required, and the output of the generated song lyrics.

First, LSTM model retains even more information further down the sequence when it compared to GRU model. Meanwhile, Markov chains approaches implemented a simple method to generate dictionary on top of the corpus to generate the song lyrics randomly based on the stored dictionary.

Besides, Markov chains took the shortest time to implement among all the approaches as it doesn't involve complex model training process. Then, GRU model is faster than

LSTM model due the number of gates in the neural network architectures. LSTM has three gates, but GRU only has two gates in the network.

Despite the Markov chains are fast, however the average readability of the generated song lyrics outputs is highest among all the methods but due to its randomness in generating the lyric. Thus, it is not suitable to select as the right approach for lyrics generation. In the meantime, when comparing the outputs of the GRU and LSTM models, LSTM scored better in the average readability index in overall. However, the GRU have the overall highest Rhyme Density score.

Overall, GRU is the most favorable approach for small datasets that was applied in our paper as it has fast computational speed and better output.

## 6.2 Limitations

In our experiments, LSTM is the model that required high computation power and long hours to train. In the first few trials in training the LSTM, the time taken to complete for a model took around 8 h. Due to that issue, different kinds of approaches being implemented to improve the overall training time or speed One of the approaches is instead of using CPU in the tensorflow library, the CUDA and GPU driver are installed to enable the tensorflow-gpu. The GPU that being applied in this experiment are NVIDIA GeForce GTX 1650. There is an obvious improvement in the training time which reduced to 3 to 4 h for training the LSTM model. It has been very challenging for us to train to the models for LSTM and GRU models for every genre in total 12 models as the training model are time intensive.

Other than that, the huge dataset also is one of the limitations for our experiments. Apparently, our hardware insufficient RAM to train huge dataset that exceeded around 2GB. Thus, the dataset required to limit down to 2GB so that it can fit into the model and carry out training process. For example, due to large dimension for EDM genre in our dataset, therefore it reduced to 80 song tracks in order to train.

## 7 Conclusion and Future Work

In this paper, three different algorithms, specifically Markov chains, long short-term memory (LSTM), and gated recurrent units (GRU) have been implemented to generate song lyrics. Our experimental results show that the GRU has the best output based on the song lyrics. Based on our trials in training all the stated model, a larger dataset is required to produce a better outcome. However, our hardware resources are limited, and the GPU memory is unable to support a bigger dataset. Therefore, our future work includes collecting more data, using upgraded hardware to train the models, and observing the outcome.

## References

1. Montagu, J.: How music and instruments began: A brief overview of the origin and entire development of music, from its earliest stages. *Front. Sociol.* **2**, 8 (2017)

2. Levinowitz, L.M.: The importance of music in early childhood. *General Music Today* **12**(1), 4–7 (1998)
3. North, A.C., Hargreaves, D.J., O’Neill, S.A.: The importance of music to adolescents. *Br. J. Psychol.* **70**(2), 255–272 (2000)
4. Cohen, A., Bailey, B., Nilsson, T.: The importance of music to seniors. *Psychomusicol. J. Res. Music Cogn.* **18**(1–2), 89–102 (2002)
5. Barradas, G.T., Sakka, L.S.: When words matter: a cross-cultural perspective on lyrics and their relationship to musical emotions. *Psychol. Music* **50**(2), 650–669 (2022)
6. Astor, P.: The poetry of rock: song lyrics are not poems but the words still matter; another look at Richard Goldstein’s collection of rock lyrics. *Pop. Music* **29**(1), 143–148 (2010)
7. Russell, E.: 25 pop songs with social messages. In: *PopCrush* (2015). <https://popcrush.com/pop-songs-social-messages>. Accessed 11 Aug 2022
8. Pudaruth, S., Amourdon, S., Anseline, J.: Automated generation of song lyrics using CFGs. In: 2014 Seventh International Conference on Contemporary Computing (IC3), pp. 613–616. IEEE, Noida, India (2014)
9. Gill, H., Lee, D.T., Marwell, N.: Deep learning in musical lyric generation: an LSTM- based approach. *Yale Undergrad. Res. J.* **1**(1), 1–7 (2020)
10. Barbieri, G., Pachet, F., Roy, P., Esposti, M. D.: Markov constraints for generating lyrics with style. In: *Proceedings of the 20th European Conference on Artificial Intelligence (ECAI 2012)*, pp. 115–120. IOS Press, Amsterdam, Netherlands (2012)
11. Fernandez, A.C.T., Tarnate, K.J.M., Devaraj, M.: Deep rapping: character level neural models for automated rap lyrics composition. *Int. J. Innov. Technol. Explor. Eng. (IJITEE)* **8**(2S), 306–311 (2018)
12. Lowder, J.: The 10 most popular music genres. In: *Audio Captain* (2021). <https://audiocaptain.com/most-popular-music-genres>. Accessed 11 Aug 2022
13. Clark, B.: The top 10 genres in the music industry. In: *Musician Wave* (2021). <https://www.musicianwave.com/top-music-genres>. Accessed 11 Aug 2022
14. Settles, B.: Computational creativity tools for songwriters. In: *Proceedings of the NAACL HLT 2010 Second Workshop on Computational Approaches to Linguistic Creativity*, pp. 49–57. Association for Computational Linguistics, Los Angeles, California, USA (2010)
15. Al-Anzi, F.S., AbuZeina, D.: A survey of Markov chain models in linguistics applications. In: *Computer Science and Information Technology (CS & IT) Conference Proceedings*, vol. 6, no. 13, pp. 53–62. AIRCC Publishing Corporation, Chennai, Tamil Nadu, India (2016)
16. Brownlee, J.: A gentle introduction to long short-term memory networks by the experts. In: *Machine Learning Mastery* (2017). <https://machinelearningmastery.com/gentleintroduction-long-short-term-memory-networks-experts>. Accessed 11 Aug 2022
17. Hochreiter, S., Schmidhuber, J.: Long short-term memory. *Neural Comput.* **9**(8), 1735–1780 (1997)
18. Rathore, A.S.: LSTM — Introduction in simple words. In: *Medium* (2020). <https://medium.com/nerd-for-tech/lstm-introduction-in-simple-words-fe544a45f1e7>. Accessed 11 Aug 2022
19. Cho, K., et al.: Learning phrase representations using RNN encoder–decoder for statistical machine translation. In: *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 1724–1734. Association for Computational Linguistics, Doha, Qatar (2014)
20. Shen, G., Tan, Q., Zhang, H., Zeng, P., Xu, J.: Deep learning with gated recurrent unit networks for financial sequence predictions. *Proc. Comput. Sci.* **131**, 895–903 (2018)
21. Saxena, S.: Introduction to gated recurrent unit (GRU). In: *Analytics Vidhya* (2021). <https://www.analyticsvidhya.com/blog/2021/03/introduction-to-gated-recurrent-unit-gru>. Accessed 11 Aug 2022

22. Banga, S.S.: Readability index in Python (NLP). In: GeeksforGeeks (2021). <https://www.geeksforgeeks.org/readability-index-pythonnlp>. Accessed 11 Aug 2022
23. Hirjee, H., Brown, D.G.: Using automated rhyme detection to characterize rhyming style in rap music. *Empir. Musicol. Rev.* **5**(4), 121–145 (2010)