

Can Chat-GPT Rap: Analyzing the ability of Chat-GPT to replicate and Analyze the Musical Style of a Song Hug.

Eli France, Jon Chun, and Katherine Elkins
IPHS 300 AI for Humanity (Spring 2023) Prof Elkins and Chun, Kenyon College

Abstract

Initially, it was believed that while automation would quickly replace manual labor, it would be more difficult for artificial intelligence to replicate human creativity and produce art. However, in the last few years, this assumption has been proven quite wrong. Art, music, and literature are all burgeoning genres in their own right. While there are questions regarding the creative merit of their work, the ascendancy of these programs seems to indicate that AI can produce art that replicates that created by humans. Nevertheless, artists have been typically resistant to embrace these innovations in artificial intelligence out of fear that they may be replaced by these programs.

While AI art is rapidly gaining popularity, one medium which will not yet be infiltrated by AI is songwriting. As of the time of this work, there are no AI programs designed to engage in songwriting. However, in my research, I saw that through the proper usage of one-shot learning and prompt engineering, LLMs possess the ability to replicate the songwriting style of rapper Young Thug accurately. These results suggest that AI may not replace artists but rather make their jobs easier.

Introduction

While musical creation was considered one of the pinnacle of human creativity, AI has made significant advances in the field. While the first known AI music experiment was conducted by Alan Turing in 1951, it would not be until when Google Magenta and Weston Behl could generate musical compositions. While these initial compositions were elementary and often one-note, many have taken AI music generation, such as AIVA, now allow users to generate entire songs. A recent study found that the AI music generation could create compositions indistinguishable from those of humans. The paper further predicted that by 2030, AI music would dominate the industry. While it remains to be seen if their prediction is validated, an AI-generated song, "Lil Nas X," created by anonymous user @hug9777, which utilized the voice of popular Canadian musician, The Weeknd, and Drake, was streamed millions of times in just a week before being removed from major platforms due to copyright issues on the request of Universal Music Group. However, while this song is no longer accessible, it sparked larger debates about the creation and legality of AI music, especially as it pertains to the creation of the style of a specific performer.

Opponents of AI-generated music claim it potentially devalues the music industry. AI music, especially when used to replicate pre-existing songs, can potentially devalue human artists. Considering the capability of artificial intelligence (AI) to generate music at a significantly faster rate and lower cost compared to humans, it is reasonable to suggest that the music industry, using AI-generated songs, could result in a decrease in the number of jobs available in the industry. The program also has capabilities regarding innovation and creativity. Finally, there are questions about whether AI can truly create music that encapsulates the nuances and subtleties of human emotion and expression, as expressed in a 2022 research paper that claimed AI struggles most to make music humans consider to be emotionally resonant. Thus, they suggest that humans must be behind or provide feedback.

Defenders of AI music see it as a replacement for traditional artistry but rather as a new tool artists can use to enhance their work. Proponents of this technology argue that AI music can be created more efficiently and cost-effectively than traditional methods of music creation. In this way, AI music can provide an opportunity to express musical creativity for people who may, for various reasons, such as lack of formal training or limited monetary assets, have been limited from doing so. Moreover, AI music can also allow for new possibilities, such as collaborations that may not have otherwise been possible or the creation and release of new posthumous music. Finally, as Carlos Hernandez-Olivan argued in his 2022 paper, "A Survey on Artificial Intelligence for Music Generation: Artists, Sounds, and Perspectives," individuals creating a finite resource, and thus, AI music generation may allow for the creation of innovative new sounds and compositions. Finally, AI lyrics can assist in ending respect culture. Respect culture, a term defined by popular music publication Pitchfork, refers to when artists preview unfinished versions of their songs never release them as they struggle to finish the song due to writer's block. Thus, Chat-GPT and other LLMs can be used to help artists complete songs.

Currently, AI music creation largely revolves around the generation of original compositions. Services such as AIVA, Amper, and OpenAI MuseNet allow users to compose royalty-free AI-generated instrumental tracks. These programs allow the user to create technically proficient compositions in any style, and some also allow the user to provide "inspiration" in the form of a pre-existing song. However, songs often contain not only music but also lyrics. While significant AI innovations have occurred regarding the former, far less attention has been paid to the latter. As of the time of this writing, there are no songwriting AI programs, and the existing music generation programs do not include lyrics in their compositions. Many of the most popular pieces of AI music, such as AI-generated songs by "AI-Flow" track, utilized AI vocal and music composition but human-written lyrics. Nevertheless, Large Language Models (LLMs) provide an exciting opportunity for the generation of lyrics. LLMs, including OpenAI's Chat-GPT, generate responses to prompts or questions by utilizing natural language processing and deep learning algorithms. These models undergo a training process utilizing a vast corpus of textual data, enabling them to learn common patterns and associations within language. The training data is tokenized, which means that it is divided into units of meaning, such as individual words or phrases, which AI can comprehend. When a user inputs a prompt, the LLMs employ its knowledge of contextual language to predict the most probable word or phrase to follow. In this way, it generates textual responses through a series of predictions. This process allows the AI to simulate natural communication and produce outputs often indistinguishable from human authors.

Consequently, LLMs have shown the ability to create lyrics that conform to standard songwriting practices, including consistent rhyme schemes and verse-chorus format. Moreover, given the vast textual data on which these programs were trained, they can often create lyrics in the style of certain artists. This feature is the one that interests me most. The act of songwriting has always held a deep fascination for me. Although the composition of music is often viewed as a technical process, it is also a highly creative one. The act of writing a single aspect of a song, if it is often the lyrics of my favorite artists that truly captivate me. While genre conventions may restrict instrumental pieces, the lyrical aspect of a song provides abundant room for creativity and, in this way, allows artists to stand out. My son, Sawyer, can express their individuality through a unique vocabulary, thematic focus, flow, and rhyme scheme and cultivate an inimitable style that sets them apart from their peers. In this way, the lyrics of a song serve as a medium for conveying an artist's personality and artistic vision. Therefore, to write and discern the extent to which LLMs can successfully replicate the compelling nature of a specific, especially those rappers, as they are far more likely than other artists to write their lyrics rather than relying on a team of songwriters. Unlike traditional writing, rap music is characterized by its rhythmic vocal delivery, which provides a distinct and diverse approach to songwriting. The use of rhythmic patterns in rap allows for greater creativity and experimentation. Rappers often employ complex rhyme schemes, repeating words or phrases in ways that go beyond traditional definitions. Consequently, as songs do not have to follow a much-to-a-lack-of-a-rhyme, more recognizable patterns can be discerned in their songwriting.

I believe that the results of the project have valuable implications concerning the future of AI music. The extent to which LLMs can accurately generate lyrics similar to those of existing rappers will help revealing regarding the ability of AI music to replicate human-made music. Comparing the generated lyrics will also insight into the legality of AI-generated lyrics. Specifically, it can be seen if the songs created (paraphrase lyrics or produce content that, while abiding by the standards of a certain artist, has recognizable material be classified as fully unique. I am aware of certain limitations to this project. First, LLMs must abide by certain content policies that prevent them from generating content that is considered to be harmful or illegal. Second, the AI-generated lyrics may not always be as creative or as distinct as activity common in rap music. Due to content policies that limit potentially offensive or controversial language, AI-generated lyrics may prioritize conformity to mainstream standards of acceptability and therefore lack the distinctive creativity and originality that characterizes the work of the human artists that are used in this comparison. Moreover, since that LLMs work in a word-by-word or phrase-by-phrase basis, they lack long-term foresight and may thus struggle to create intricate rhymes or cohesive narratives. Nevertheless, I believe many of these limitations can be avoided through strategic prompt engineering, and thus it is worthwhile to investigate the songwriting capabilities of LLMs. Therefore, my project will answer the question: To what extent can LLMs replicate the songwriting style of rapper Young Thug?

QR Codes For Resource



Methodology

The first step in my research would be to select a rapper whose style was to be replicated. After researching various artists, I ultimately decided that Young Thug, real name Jeffrey Williams (phonetic), referred to as "Thug" or "Wesley", would be the focus of my study. Williams was selected due to his unique and highly recognizable songwriting style, which Pitchfork has described his writing as "viciously distinctive". Moreover, Williams' has been categorized as the most influential rapper of the 21st century by the BBC, who noted that many young rappers emulate his style. Thug's unique use of slang, wordplay, and flow. My criteria for selecting the artist whose sentences would replicate were that I had to include hallmarks of Williams' style, a clear narrative and had to have been released after the knowledge cutoff date for Chat-GPT. Accordingly, from his most recent 2021 album "Punk," the song "Hate The Game" was chosen due to its hallmark use of slang, clear narrative, and triumphant message, which aligns with the themes which I determined to be most present in Young Thug's music. It is important to note that the song was released after the knowledge cutoff date for Chat-GPT, meaning that the AI would not have any knowledge of the lyrics.

Eight different prompts were used to simulate different possible ways in which AI could be implemented into songwriting practice. The first prompt utilized zero-shot learning, which allowed the AI to create lyrics with as little human involvement as possible. This prompt simply asked Chat-GPT to create an entire "Hate the Game" in the style of Young Thug. The lyrics created from this prompt are classified as "Standard," and from here on out, this will be referred to as the "standard" prompt. Subsequently, I attempted to use one-shot learning, providing an example of real Young Thug lyrics which the LLM was to replicate for vocabulary and phrasing. Next, I explored the capabilities of different prompt engineering, providing details of the mood, themes, and topics with which to reference in addition to the standard prompt. Thus, I engaged in human-AI collaboration, providing the LLM with lyrics to the chorus of "Hate The Game" and asking it to write two additional verses.

The attributes of these prompts were not mutually exclusive. Consequently, I combined them to form a different combination prompt. The first was a prompt with both an explanation and an example, the second was a prompt with an explanation and chorus, the third was a prompt with a chorus and example, and the fourth included all three.

I judged the generated songs not on their ability to replicate one aspect of Young Thug's songwriting but rather on encapsulating the various features which make his style unique. The process of manually assessing the similarities of the outputs to the real song began with me analyzing the vocals lyrics to "Hate The Game." I identified several of the most apparent sentiments, themes, and topics/disco in Chorus and I read assessed and recorded the frequency at which these themes or sentiments were apparent in the lines of the songs. This process was repeated for each of the eight songs. To calculate the average disparity percentage (the metric I created to assess similarity) for each category, I took the absolute value of the difference between the frequencies with which each theme or sentiment appeared in the real and generated songs. For each prompt attribute type, I added the absolute value scores together and divided them by the number of categories such that I could determine the average disparity concerning the frequency with which topics and themes appear between the real and generated songs. For example, if a prompt attribute type, Percentage of AFD, I then repeated this process with a smaller set of categories so I could focus specifically on sentiment and topical disparity percentage. These were then referred to as SAFP and TAFP, respectively. Finally, I identified key attributes of Thug's music, such as multi-syllabic rhyme scheme, adlibs, and his distinct phrases, and subtracted 2% from the scores of each of the generated lyrics for how many times they appeared while adding 2% for every lyric which I deemed unusual. Finally, I recorded and charted each of the three for each similarity type.

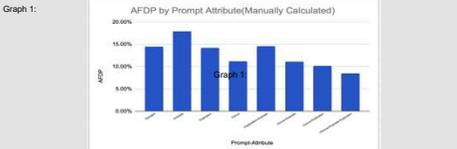
To get a second opinion on my results, I asked Chat-GPT to generate lyrics similar to Young Thug's. I also had Chat-GPT assess the generated lyrics. To achieve that goal, I asked Chat-GPT for the most common sentiments, themes, and topics of Young Thug songs/disco in Chorus 1. Subsequently, I asked GPT to rate the lyrics of "Hate The Game" by Young Thug on the basis of how present each sentiment, theme, and topic was represented in the lyrics generated in the prompt and then again took the average difference between their scores and the score of the original song. I then analyzed and compared the results to see which prompts produced the most similar text and how my evaluations differed from those of the AI.

Chart 1:	Chart 2:
Manually Chosen Themes or Topics	Manually Chosen Themes or Topics
Triumphal	Triumphal
Sad/Neutral	Sad/Neutral
Bragadocious Statements	Bragadocious Statements
Relationship Toxicity	Relationship Toxicity
Violence	Violence
References to Brand Names	References to Brand Names
Weapons	Weapons
Transportation	Transportation
Cops	Cops
Drugs	Drugs
Clothing, Jewelry and Accessories	Clothing, Jewelry and Accessories
References to Women or Sex	References to Women or Sex
Lines about Thug	Lines about Thug

Results and Analysis Part 1

The AI showed a remarkable propensity to create words quite similar to those of Young Thug. Each piece GPT created incorporated a standard rhyme scheme through-out and was also elements of songwriting and adlibs. From my own personal analysis, a vast majority of lyrics created to align with those of Thug, as on the standard prompt (but more than three lines which I marked as being unusual). Most of these unusual lines used dated vernacular, suggesting that Chat-GPT may be limited in its ability to understand the nuances of language usage fully. Nevertheless, the AFDP (the metric defined in the Methodology section) revealed below 20% among all eight different prompts utilized. As illustrated in Graph 1, the standard prompt, which aimed to provide the LLM with minimal information, achieved an AFDP below 15. Surprisingly, the LLM seemed to learn from the learning prompts that included more information, resulting in a higher AFDP of 17.9. The higher amount of information provided in the learning prompts, as I noticed when analyzing the output that I repeated a few of the lyrics in the prompt nearly verbatim. Explaining to the LLM what to include proved much more successful at replicating the lyrics of Young Thug than simply providing the lyrics of Young Thug. As seen in Graph 1, the most detailed explanations attached achieved an AFDP of 14.3, showing improvement over the standard prompts. However, the lowest AFDP of any of these one-attribute prompts, indicating the highest similarity, belonged to the prompt, which included the song's chorus and simply asked the LLM to write verses. As seen in Graph 1, the AFDP of these prompts was 12.2%. The fact that the chorus performs the best suggests that the usage of human-AI collaborations might be an effective means of songwriting, allowing artists to complete unfinished lines.

This results from multi-attribute prompts further support the idea that LLMs work most effectively when asked to complete songs rather than create them from scratch. Prompts that included both an explanation and examples produced an average disparity percentage of 14.6, a score that was significantly higher than prompts with only the chorus. Moreover, as shown in Graph 1, why prompts with both explanations and chorus, as well as examples and chorus, scored lower than those with just the chorus or just examples, and the prompts with only the chorus outperformed the prompt with both a sample and explanation by a notable margin 1.1. Notably, the prompt that included both explanations and examples, along with the chorus, achieved the lowest score of all prompts, with 8.5. This score was the average difference in the frequency with which the themes or sentiments were used in the real and generated songs. This result is particularly interesting as it suggests that AI can write a song with a nearly identical pattern to an existing artist, with only a slight variation in topics and mood.



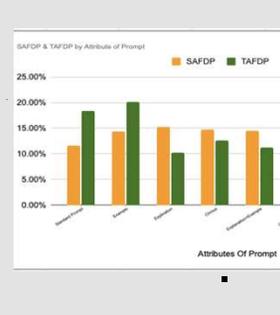
Results and Analysis Part 2

Graph 2 illustrates that while the standard prompt could generate outputs that captured the rapper's sentiments, it struggled to replicate the topics without specific details. Thus, it achieved a SAFP score of 11.6 but a TAFP score of 18.4. The same was true for the prompt utilizing one-shot learning through a Young Thug verse example. This suggests that while one-shot learning can generate output that "feels" like a Young Thug song (with a SAFP of 14.4%), it faces difficulties in addressing the song's themes based on limited examples, as seen by the TAFP of 20.2%. However, the inclusion of an explanation appeared to help the LLM more accurately replicate the frequency of topics discussed in Young Thug's lyrics. Graph 2 demonstrates that this prompt achieved one of the lowest TAFP (10.2%), likely due to the suggestion of certain themes. While it faced challenges in replicating the mood of the chorus, with a SAFP of 15.2%, these findings highlight the value of advanced prompt engineering in songwriting. However, providing the chorus resulted in relatively low scores for both SAFP and TAFP, at 14.8 and 12.0%, respectively. While the chorus had achieved the lowest AFDP, it neither achieved the highest SAFP nor TAFP. This indicates that while prompts like this that specifically engage in human-AI collaboration can replicate lyrics to the highest degree of accuracy of any single-attribute prompt, they are as strong at replicating sentiments as one-shot learning prompts or topics as prompts that use advanced engineering.

The multi-attribute prompts seemed to balance SAFP and TAFP with more ease than the single-attribute prompts. The strength of multi-attribute prompts largely lies in their ability to complement each attribute's strengths. As shown in Graph 2, the prompt with both explanation and example achieved one of the lowest SAFP (14.5%) and a relatively low TAFP (14.5%). Consequently, this prompt combined the advantages of prompt engineering and one-shot learning. The same pattern emerged in prompts with a chorus and explanation, which successfully replicated topic frequency (8.1%) but faced challenges with replicating sentiment (16.8%). These findings suggest that advanced prompt engineering is valuable for replicating thematic usage styles in songwriting. Consequently, this indicates that without the use of one-shot learning, Chat-GPT has difficulty discerning emotions. This suggests that LLMs may better be able to understand the emotive qualities of songs through examples than specific directions.

Amazingly, the scores for the prompt with the chorus and examples indicate that one-shot learning is most successful in capturing the mood of specific lyrics. However, the prompt, which included an example, an explanation, and a chorus, had both the lowest average topic and sentiment disparity percentages of 8 and 8.2, respectively. This suggests that AI songwriting can most accurately replicate the lyrical style of artists through a combination of one-shot learning, prompt engineering, and human-AI collaboration.

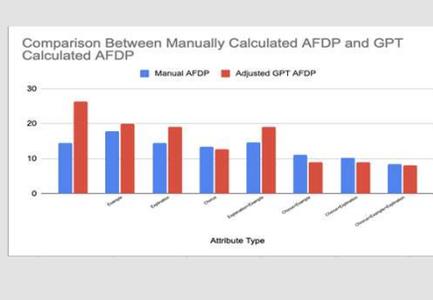
Graph 2:



As can be seen from Graph 3, for the most part, Chat-GPT assessed the different prompts similarly to me. One notable variation was our assessment of the standard prompt, which Chat-GPT found to be far more dissimilar from the actual song than I believe this results from this prompt creating more unique lyrics due to lack of information, which Chat-GPT could have more difficulty analyzing for themes or sentiments. Chat-GPT's score was the same rankings to the lyrics produced by prompts with explanations and examples. While I had determined the version with examples to be the most notably similar to the original song, Chat-GPT apparently considered them to be equally similar. The version believed to be more similar was generated by Chat-GPT for being as violent or sexual as the original song. However, in my own analysis, I found many lines that could be classified as violent or sexual, even if they did not explicitly reference violence or sex. The further points to the fact that Chat-GPT has difficulty analyzing subtle lyrics. Nevertheless, like me, Chat-GPT gave the lowest disparity score of any single-attribute prompt to the prompt with the chorus, indicating the highest similarity.

With the exception of the example and explanation version, which I also considered to be the worst of the advanced prompts, Chat-GPT gave high multi-attribute prompts similar scores within 1 of each other. As I have noted, it seems that prompts with explanations allow for better replication of topics. Prompts with examples produce lyrics that better encapsulate the emotions of the song. Given that Chat-GPT gave both the explanation and chorus, as well as the chorus and example prompts, the same score, it seems that Chat-GPT values thematic similarity equally to topical similarity. The fact that the explanation and example-only prompts also got the same score provides further evidence for this theory. Personally, as I identified more topics than mood, topical disparity affected my rating system the most. This difference might explain the disparity between my calculations and those of Chat-GPT. However, both Chat-GPT and I graded the prompt, combining all attributes the highest. Consequently, this means that Chat-GPT assesses the attributes of lyrics similarly to humans. Nevertheless, just as Chat-GPT did not fully replicate the styles of actual artists, there were also irregularities in its ability to imitate mood. Not only did Chat-GPT miss subtleties, but it would occasionally make up lyrics in its explanations for its scores. Consequently, this casts doubt on the reliability of Chat-GPT to assess lyrics, suggesting that it needs the guidance of a human.

Graph 3:



Conclusion

These results suggest that LLMs can understand and replicate the lyrical styles of rappers such as Young Thug. Even the standard prompts created a song that expressed sentiments and emotional topics similar to those in "Hate the Game" by Young Thug. Moreover, the experiment shows how different types of prompts help replicate different attributes of an artist's lyrical style. For example, advanced prompt engineering results in a high topical similarity, while one-shot learning helps the LLM to replicate the sentiments expressed by an artist. Consequently, it can be seen how combining these different attributes into one prompt produces the best results. Additionally, the fact that the chorus-only prompt outperformed all of the other single-attribute prompts suggests that the best results will emerge from human-AI collaboration. On a broader scale, this suggests that LLMs will be another tool in their kit. Rather than replace songwriters, I foresee this helping artists and writer's block to complete song ideas and consequently address the issues of output culture.

From comparing my assessment of the songs with those of Chat-GPT, it appears that LLMs still do not understand lyrics like a human. While I cannot say that my scores are objectively more accurate, I will note that the divergence between my scores and the scores of the AI seems to have several implications. First, my scores seemed to vary more, suggesting that Chat-GPT is less sensitive to small changes in the lyrics or presents a less numerically-based analysis. Moreover, it seems that Chat-GPT has trouble detecting more subtle themes. This was because it often tended to rate the songs very high or low on their inclusion of a certain category's themes. Therefore, it struggled to identify song themes or discuss the topics discussed in some less prominent lines. Nevertheless, the fact that the LLM agreed with me regarding which song was the most similar to "Hate The Game" shows that, on some level, it can understand lyrics similarly to a human. Thus, while the ability of Chat-GPT to understand lyrical themes may well be behind its ability to generate music, it is still capable of discerning those lyrics to some extent.

Limitations, Future Research and Considerations

The findings of this study are subject to limitations due to time constraints and available technology during the project's completion as well as the sample size. Had I had more access to conduct this experiment, I would have been able to experiment more with the creation of different types of prompts. This could have allowed for further insights into the best practices for this experiment. Moreover, experimenting with more types of prompts would have allowed me to further explore the ability of AI-generated artists. Lastly, I believe that it will have an easier time creating lyrics to complete this experiment, I could have created multiple sets of lyrics to better understand how each prompt's different attributes generated different lyrics. Furthermore, while running this experiment, I only had access to Chat-GPT 3.5 rather than the most updated version, Chat-GPT 4. Therefore, my results and their implications on the ability of LLMs to recreate lyrical styles may be slightly out-of-date. Moreover, I only had Chat-GPT recreate the lyrical style of one specific artist. Therefore, these results may not be reflected when analyzing the ability of LLMs to recreate other artists' or lyrics in genres besides Rap. A final limitation is that my manual assessment of themes and emotions is purely my own interpretation of the songs. While I tried to research an existing explanation for the lyrics of the song "Hate the Game," where these were unavailable, I relied upon my subjective analysis of what I believe was being discussed. This limitation also affected my evaluation of the AI-generated song.

Ultimately, while my results show promising signs regarding the ability of AI to produce songs lyrics in the style of Young Thug, I believe that further research should be conducted concerning the capability of Chat-GPT to generate lyrics. Specifically, I believe further research should be done to explore the ability of Chat-GPT to replicate the style of various artists. Lastly, I believe that it will have an easier time creating lyrics to complete this experiment, I could have created multiple sets of lyrics to better understand how each prompt's different attributes generated different lyrics. Furthermore, while running this experiment, I only had access to Chat-GPT 3.5 rather than the most updated version, Chat-GPT 4. Therefore, my results and their implications on the ability of LLMs to recreate lyrical styles may be slightly out-of-date. Moreover, I only had Chat-GPT recreate the lyrical style of one specific artist. Therefore, these results may not be reflected when analyzing the ability of LLMs to recreate other artists' or lyrics in genres besides Rap. A final limitation is that my manual assessment of themes and emotions is purely my own interpretation of the songs. While I tried to research an existing explanation for the lyrics of the song "Hate the Game," where these were unavailable, I relied upon my subjective analysis of what I believe was being discussed. This limitation also affected my evaluation of the AI-generated song.

While I used an example of one-shot learning, I believe that future studies should attempt to incorporate many-shot learning. In doing so, these studies will provide a better understanding of how Chat-GPT learns about lyrics from examples. Moreover, I believe future studies should delve deeper into prompt engineering. Due to time and technology constraints, I only provided one prompt, which included in-depth song creation details. By experimenting with different advanced prompts, I believe more could be learned about the effective use of prompt engineering to create lyrics. Specifically, it will be interesting to see how different levels of specificity affect the generation of distinct lyrical patterns. By designing prompts that call for different degrees of attention to detail, it can be examined whether highly specific prompts lead to a more accurate replication of an artist's lyrical patterns. Additionally, research on this field can also identify if more open-ended prompts allow for greater creativity and innovation. This exploration has immense value because it would help reveal the optimal balance between providing specific guidance and allowing room for the LLM to explore its own interpretations.

Finally, I believe that future studies should address innovation in technology. Most importantly, I believe that there should be research on the ability of Chat-GPT to produce lyrics. Not only will the results of this study better research the ability of current generation AI to write lyrics, but also, by comparing them to the results of my study, much can be gleaned about the differences between GPT 3.5 and GPT 4. Moreover, as AI technology advances, I believe research should be conducted on the ability of AI to create a novel and advanced lyrical style. These results will have massive implications regarding the ability of LLMs to produce works of genuine creativity and innovation.

Acknowledgements

I would like to extend my heartfelt gratitude to Professor Elkins and Professor Chun for their unwavering support and guidance throughout my academic journey. Their dedication to empowering students and fostering a passion for learning has equipped me with invaluable skills that enable me to thrive in today's rapidly evolving world. I am sincerely grateful for their mentorship and the impact they have had on my personal and professional development.

References

Cole, Ben. "The Problem with AI Music: Song and Culture Created by the Digital Age." *Produce Music*. Cambridge, UK, August 17, 2022. <https://www.produce-music.com/ai-music/>

Crabtree, Matthew. "Generative AI: Music and Songwriting." *AI Music: The Future of Sound*. Cambridge, UK, August 17, 2022. <https://www.cambridge.org/core/books/ai-music-the-future-of-sound/ai-music-the-future-of-sound>

Diaz, Adrián, and Kai R. Agha. "AI-Based Affective Music Generation: A Review of Methods and Challenges." *AI*. January 2023. <https://arxiv.org/abs/2301.08860>

Groves, Robert A. "Humanity, Utopia, AI by Tyler, The Creator." *AI Music: The Future of Sound*. Cambridge, UK, August 17, 2022. <https://www.cambridge.org/core/books/ai-music-the-future-of-sound/ai-music-the-future-of-sound>

Hernandez-Olivan, Carlos, Juan Hernandez-Ortiz, and Josef M. Balboa. "A Survey on Artificial Intelligence for Music Generation: Artists, Sounds, and Perspectives." *AI*. November 1, 2022. <https://arxiv.org/abs/2210.13944>

Jee, The. "The AI of Artificial Intelligence." *Music*. Limer, January 20, 2023. <https://www.themusicologist.com/the-ai-of-ai-music-2/>

Thompson, Paul A. "Physicist Carl's Rap: Beyond Language." *Wired*. October 15, 2011. <https://www.wired.com/2011/10/physicist-carl-rap/>

Lorenson, Paul A. "Physicist Carl: While Limer Reels." *Pitchfork*. January 5, 2021. <https://pitchfork.com/news/while-limer-reels/>

Wang, Jiefu. "Why Young Thug Is the 21st Century's Most Influential Rapper." *BBC Culture*. January 23, 2023.

<https://www.bbc.com/culture/articles/culture-2023-01-23-why-young-thug-is-the-21st-century-s-most-influential-rapper>