Deepfake Music and Listener Sentiments: A Large-Scale Analysis of YouTube Comments.

Francisco Tigre Moura
Marketing and Communication
IU International University of Applied Sciences
Bonn, Germany
Email: francisco.tigre-moura@iu.org

Visieu Lac
Business Management
IU International University of Applied Sciences
Bad Honnef, Germany
Email: visieu.lac@iu.org

***Abstract*—This study investigates emotional responses to deepfake music by analyzing 31,363 YouTube comments using 'twitter-roberta-base-sentiment-latest' model. The research addresses an important gap in the literature by focusing on user sentiments towards Artificial Intelligence (AI)-generated deepfake songs that mimic human voices, contrasting them with non-deepfake AI music (instrumental or relaxation genres). Findings reveal a surprising predominance of positive and neutral sentiments towards deepfake music, particularly in genres like Rock/Pop and Cartoon, though negative reactions are more pronounced when renowned human voices are mimicked. The study also identifies genre-specific patterns and a longitudinal decline in novelty-driven engagement, especially within Rap/Hip-Hop. Compared to non-deepfake AI music, which consistently triggers positive sentiments, deepfake music evokes more mixed responses, suggesting that voice mimicry remains a critical aspect. The findings contribute to the understanding of how AI creativity influences listener emotions and perceptions, while raising timely questions regarding authenticity and acceptance of deepfake art.**

Keywords—Artificial Intelligence; Sentiment analysis; Deepfake; Deepfake Music.

I. INTRODUCTION

The creation of novel songs using Artificial Intelligence (AI) "*deepfake*" tools (tools that allowed users to, for example, create unauthorized tracks mimicking voices of renown artists or celebrities) has gained great popularity and is now impacting the music sector. For example, in 2023 a deepfake song cloning the voices of Drake and The Weekend (titled "Heart on My Sleeve") represented the first deepfake track to become a viral hit (over 20 million streams in 2023 on Apple Music [1]). Since then, deepfake songs have become increasingly popular, with an extraordinary number of tracks being released online, resembling the voices of artists such as Frank Sinatra, John Lennon, Freddy Mercury, Taylor Swift, Kanye West, Kurt Cobain and many others. In view of the various technological, ethical and legal challenges of "*deepfaking*" voices [2], including how to legally protect artists, the music industry has reacted. Universal Music and Google, for example, have started negotiations for the development of tools that allow users to produce deepfakes music legitimately and financially reward the copyright holders [3].

Given that AI, and its unauthorized use in deepfakes, drastically alters the fundamental nature of creative processes in music, and the extraordinary impact and threats it poses to various sectors [4], it is crucial to further understand the human response towards deepfake music. However, although the current literature has addressed many aspects regarding human acceptance towards AI generated music [5]-[7][27], to our knowledge, it has yet to investigate human response to deepfake music, thus representing a significant gap. In view of the ever-growing amount of deepfake art, including music, and tools available for their creation, and the relevance of addressing this phenomenon, this study aims to investigate two main research questions:

- **RQ1:** *What sentiments do YouTube users hold towards deepfake music?*
- **RQ2:** *How have the sentiments of YouTube users towards deepfake music changed over time?*

To address both questions, in the next section the paper debates human response towards AI generated or co-created music, followed by section on deepfake and music. Afterwards, Section 4 describes the methodology of the study, while Section 5 reveals the findings. Finally, Section 6 concludes the paper by presenting the conclusions and limitations of the study.

II. HUMAN RESPONSE TOWARDS AI GENERATED OR CO-CREATED MUSIC

The understanding of human response towards innovation and technology generated or co-created outputs is, overall, complex and multifaceted. Several factors such as demographic variables, gender, and educational levels all play a role in acceptance [8][9]. Concerning human response towards AI generated or co-created music, recent literature found also that music professionals often hold different attitudes when compared to listeners in general [5]. Also, despite the rather overall negative perceptions towards AI generated music [10] and perception biases [11], a controlled experiment revealed that effects are weakened if respondents hold a positive perception towards the song they listen to [5]. Moreover, the listener's level of involvement with music and the context of their involvement should also be considered

[7]. Hong et al. also further emphasized that the perception of AI as an independent creative agent affects how its music is received. For instance, those who view AI as a musician tend to appreciate its music more than those who do not [12]. This emphasizes the importance of the creative process in influencing acceptance and response.

Another example of composer bias (preference for human composer in contrast to AI as composer) was revealed through a series of experiments developed by [10]. The studies revealed that due to the "AI-sounding" of electronic music, participants were more accepting of AI as composer and the music it generates, when compared to classical music, as it resembles a more "human-sound". Furthermore, a recent systematic review of 30 studies investigated human emotional responses elicited by AI-composed music (e.g., arousal, enjoyment, interest and liking), with a focus on understanding the emotional authenticity and expressivity of such compositions [13]. The review suggested that, while AI can help explore emotional authenticity in music, there remains significant skepticism and preference for human-composed music among listeners and music professionals. The review also emphasized that factors, such as music genre, cultural context, and age group confound the understanding of emotional responses, and suggests the development of advanced analytical methods, for example using machine learning and deep learning (adopted in this paper), to enhance the comprehension and effectiveness of AI in music composition [13]. Finally, all studies mentioned here have focused on human responses towards novel or authentic compositions (non-deepfakes) created independently by or co-created with AI.

III. DEEPFAKES AND MUSIC

Although there is not a consensus regarding the definition of deepfakes, [14] defends that "as its name implies, the term "deepfake" is derived from the combination of "deep" (referring to Deep learning (DL)) and "fake". It is normally used to refer to manipulation of existing media (image, video and/or audio) or generation of new (synthetic) media using DL-based approaches" (p.2). Importantly, as [14] defend, deepfakes enhance the naturalness of artificial agents (regardless of format), improving their ability to generate empathy and emotional connection with humans that are exposed to or interact with them.

The recent developments in machine and deep learning technologies have enabled an extraordinary increase in deepfakes use for authentic purposes, for instance, entertainment [16], but also for various malicious purposes, such as misinformation. In view of the advances in deepfakes, the extreme challenge of identifying it, and its high level of persuasiveness, many authors have raised serious concerns regarding the social impact it already causes and may cause in the future [17][18].

The current literature involving deepfake music is largely focused on voice identification, rather than human response. In this regard, recent studies have suggested that, different to speech voices, identifying deepfake singing voices is particularly more challenging, due to the role of melody, rhythm, and the broader range of timbre in singing [19][20]. Moreover, detecting deepfake singing voices presents a unique challenge due to the interference of background music, which can mask the artifacts used to identify synthesized voices [20]. For example, unlike speech deepfakes, where the vocal track is often isolated, singing voices are typically surrounded by musical arrangements that include instrumental accompaniments and digital effects. This layering of sounds makes it difficult to discern the subtle cues of synthesis, as the instruments and other musical elements can mask or mimic these artifacts. Additionally, the artistic nature of music production, with its wide range of timbres and dynamic variations, adds another layer of complexity to the detection process. Traditional speech countermeasure systems, when applied to these mixed audio tracks, often fail to accurately distinguish between authentic and fake singing, leading to significantly higher error rates.

Next, we discuss the methodology of our study, which aimed to investigate sentiments of listeners towards deepfake music.

IV. METHODOLOGY

The methodology section will first describe the sample of YouTube channels used in the study, followed by the process for sentiment analysis.

A. Sample of YouTube Channels and Type of Songs

The first step consisted of collecting a large sample of YouTube channels containing deepfake music. Importantly, to ensure that sentiments analyzed were specific to deepfake music, it was necessary to contrast them with AI generated music that did not mimic human-like voices. The inclusion criteria of YouTube channels consisted of: (a) videos explicitly labelled as deepfake music or AI generated music; (b) contained a minimum of 10 videos, and (c) most videos contained large number of comments. The initial selection consisted of 62 channels. After further screening, a final sample of 44 channels was used for the analysis, as displayed in Table I.

TABLE I. OVERVIEW OF CHANNELS, VIDEOS AND COMMENTS OF DEEPFAKE AND NON-DEEPFAKE AI MUSIC.

Type	Voice Mimicking	Genre	No. of Channels	No. of Videos	No. of Comments
Deepfake Music	Renown Human Voices	Rap/Hip-hop	6	190	7,558
		Rock/Pop	23	1,151	13,544
	Fictional Characters	Cartoon	5	249	5,937
Non-Deepfake AI Music	Does Not Apply	Instrumental	2	247	5,396
		Relaxation	8	697	724
		Total	**44**	**2,534**	**33,159**

Next, channels were classified into two broad categories:

(1) "Deepfake Music": Included songs which used AI to mimic voices. This category was split into two sub-categories, according to the type of voice mimicking:

(1a) "Renown Human Voices": Included songs that mimicked voices of famous and recognizable human artists or celebrities (e.g., Kurt Cobain, Kanye West, Taylor Swift, Adele, Frank Sinatra, Donal Trump, Barack Obama and Joe Biden) performing popular songs from other artists, or novel compositions. Examples included Kurt Cobain (former lead singer of the grunge band Nirvana) singing “Wonderwall” (Originally composed and recorded by Oasis), and Freddie Mercury (former lead singer of British rock band “Queen”) singing “Hey Jude” (originally composed and recorded by “The Beatles”). This sub-category was composed of two music genres: Rap/Hip-Hop, Rock/Pop.

(1b) “Fictional Characters”: Included deepfake songs which did not mimic voices of renown human artists, but instead, of fictitious characters. The main genre derived from this category is “Cartoon”. For example, the cartoon character “Bluey” singing “Bumble Bee”.

(2) “Non-Deepfake AI music”: Included instrumental songs, or atmospheric sounds, composed by AI, and that did not include voices. Two genres comprise this category: Instrumental and Relaxation (not songs, but AI generated soundscapes for relaxation and focus, for example).

By using the YouTube Application Programming Interface (API), comments were extracted from the videos published by the channels.

B. *Sentiment Analysis*

Prior to the sentiment analysis, comments underwent pre-processing consisting of: (1) elimination of punctuation marks and stop-words, and (2) normalization through lemmatization. This preprocessing filtered out records lacking meaningful textual content, resulting in a higher-quality dataset with reduced noise. The cleaned dataset served as input for sentiment analysis, which classified comments into three categories: positive, negative, or neutral based on the emotions expressed. This step is essential as it reveals the overall emotional tone of the comments, offering valuable insights into user perceptions and opinions [21].

To achieve this, we utilized the Twitter-roBERTa-base model developed by CardiffNLP [22], which has been validated in previous sentiment studies [23][26]. While many sentiment analysis models have been created using various datasets, such as movie reviews [24], we chose a Twitter-based model due to the nature of user comments on YouTube, which are typically brief, written in a social media style, and often include emojis.

This model classifies comments into 'positive', 'neutral', and 'negative' categories. The sentiment analysis was performed to understand the general sentiment of the comments and to identify any prevalent trends or patterns in the audience. Also, one of the key advantages of the CardiffNLP model is its ability to provide a detailed numerical breakdown of sentiments, showing the distribution of positive, neutral, and negative sentiments at the corpus level as well as the individual comment level. The CardiffNLP Twitter-roBERTa-base model (specifically, cardiffnlp/twitter-roberta-base-sentiment-latest) was trained on approximately 124 million tweets from January 2018 to December 2021 and fine-tuned for sentiment analysis, effectively incorporating emojis [25]. Furthermore, to validate the effectiveness of the Cardiff model, we benchmarked it against several sentiment analysis models, including OpenAI’s GPT-3.5-turbo. Our independent tests using the dataset [29] showed that the Cardiff model achieved the best performance, with an accuracy score of 0.72, outperforming OpenAI’s GPT-3.5-turbo, which had an accuracy score of 0.66. Thus, the higher accuracy score of the Cardiff model indicated its greater reliability in accurately classifying the sentiments of the comments in our dataset and therefore was adopted.

V. RESULTS

Prior to the analysis, comments were pre-processed as described in section IV.B of this paper. This text processing reduced the initial dataset by roughly 6.8% on average. Table II, shown below, depicts the reduction in the number of comments per genre, after the text processing, which resulted in a final sample of 31,363 comments used for the study.

TABLE II. OVERVIEW OF CHANNELS, VIDEOS AND COMMENTS OF DEEPFAKE AND NON-DEEPFAKE AI MUSIC.

Type	**Voice Mimicking**	**Genre**	**No. Comments Before Text Processing**	**No. Comments After Text Processing**
Deepfake Music	Renown Human Voices	Rap/Hip-hop	7,558	7,178
		Rock/Pop	13,544	13,003
	Fictional Characters	Cartoon	5,937	5,604
Non-Deepfake AI Music	Does Not Apply	Instrumental	5,396	4,931
		Relaxation	724	647
		Total	**33,159**	**31,363**

Next, the results of the sentiment analysis are provided through two perspectives: overall sentiment and a longitudinal analysis.

A. *Overall Sentiment Analysis*

The CardiffNLP model results are illustrated in Figure 1, which presents the sentiment distribution across different genres.

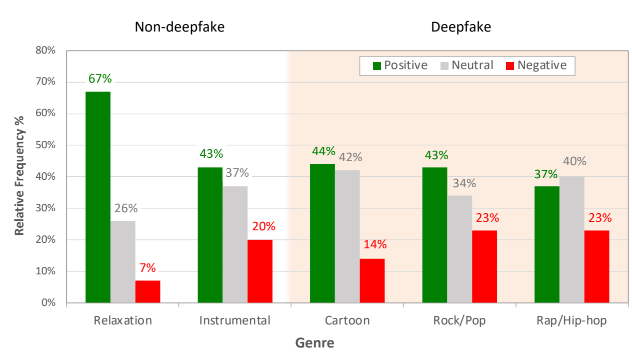

Figure 1. Overall display of sentiments by genre.

 ISBN: 978-1-68558-330-9

Figure 1 illustrates that, overall, sentiments towards non-deepfake AI music (*Relaxation* and *Instrumental*) are predominantly positive, with 67% and 43% of positive sentiments respectively. *Relaxation* also indicated the lowest negative sentiment among all genres (7%), reinforcing the potential for AI use in such music applications. Regarding deepfake music, *Cartoon* and *Rock/Pop* revealed most positive sentiments (44% and 43% respectively). *Rap/Hip-Hop* represented the main exception, where neutral sentiment exceeds positive sentiment (40%). The predominance of positive sentiments represents a surprising finding, as it conflicts with findings from Shank et al. (2023), who noted that listeners tend to be biased against music they believe was created by an AI, especially if music does not meet their expectations of what an AI could produce. Concerning negative sentiments, the deepfake music genres *Rock/Pop* and *Rap/Hip-hop* exhibited the highest proportion of negative sentiment (23% each) of all genres. This represents an important finding, as these genres normally involve renown human voices, suggesting the sensitivity of listeners towards the use of deepfake technology towards people they recognize. This issue is discussed later in this paper for future research agenda.

B. *Longitudinal Analysis of Sentiments*

Furthermore, we analyzed the data as a function of time, thus allowing the visualization of listeners' sentiment trend. Results are reported per week of each year within the dataset, as seen in Figure 2, which shows a distribution on a 100% scale.

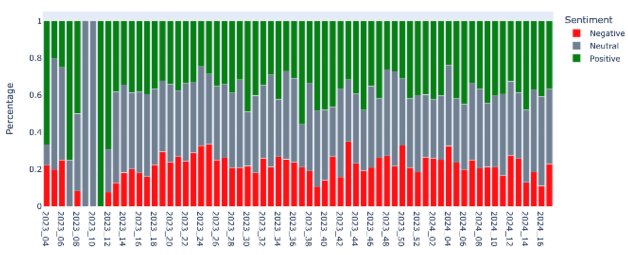

Figure 2. Distribution of sentiments for Rap/Hip-Hop on a 100% frequency scale

Results regarding *Rap/Hip-Hop* indicate two important findings. First, although not displayed in the graph, we identified in the data a high incidence of comments at the start of 2023, followed by steady decline. This may suggest an indication that, for this genre, the novelty effect quickly diminished, an aspect that requires further investigation in future studies. Secondly, and importantly, that the distribution of sentiments towards deepfake *Rap/Hip-Hop* music remains stable over time, with a greater predominance of neutral and positive sentiments (Figure 2).

Regarding the genre of *Rock/Pop*, results indicate that (differently to *Rap/Hip-Hop*), the volume of deepfakes comments has been increasing considerably over time, which may indicate a greater implementation of deepfake technology within this genre. Nevertheless, the general sentiment shown on Figure 3 reveals a growing negative sentiment trend, a finding that requires further future investigation.

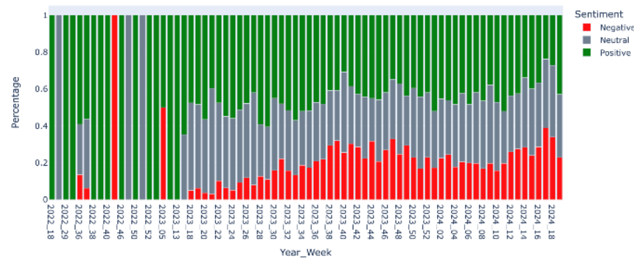

Figure 3. Distribution of sentiments for Rock/Pop within a 100% frequency scale.

Furthermore, the temporal analysis for the genre of *cartoon* indicated a stable trend regarding the overall sentiment of comments, displaying a predominance for positive and neutral sentiments, and very low negative sentiment (Figure 4), in comparison to other genres.

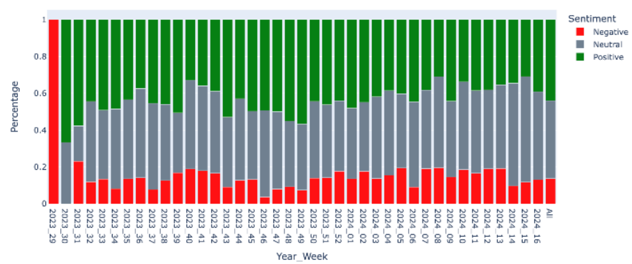

Figure 4. Distribution of sentiments for Cartoon within a 100% frequency scale.

Further results reveal that the range of comments for AI music (non-deepfakes) is considerably broader (when compared to deepfake music categories), potentially because the technology was earlier accessible for independent creators. Overall, the results for the *Instrumental* genre suggest a stable trend regarding volume of comments and sentiments, which indicate largely positive and neutral responses (Figure 5).

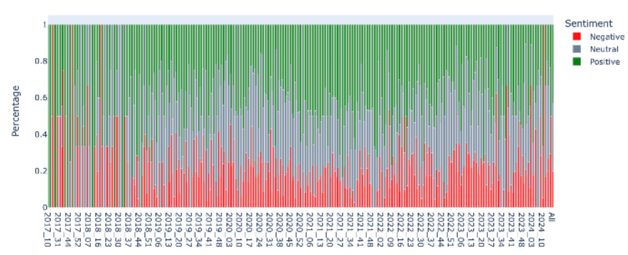

Figure 5. Distribution of sentiments for Instrumental within a 100% frequency scale.

Finally, the analysis for the genre of *Relaxation* also indicated largely displays positive sentiments towards this type of AI generated music (Figure 6).

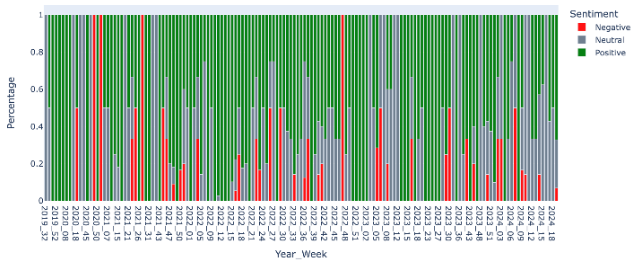

Figure 6. Distribution of sentiments for Relaxation within a 100% frequency scale.

The next sections will address the limitations of the study and provide a critical reflection on the results.

VI. LIMITATIONS

Results from the study should be interpreted in light of a few limitations. First, the analysis is based on a sample of YouTube comments, which may not be representative of the broader population and potential biases related to the cultural and geographic distribution of comments must be considered. Also, channels used for the analysis displayed varying levels of quality, based on the artistic and technological skill of the developer. For example, knowing what key and pitch the artist normally sings can be a factor when creating deepfake music. Such variation of deepfake quality certainly impacted the sentiment of comments. Thus, this issue should be addressed in future studies. Secondly, some music genres might be more accepting than others towards the use of AI, as rap-hip-hop, for example, often uses Autotune [24], which resembles robotic voices often found in AI in deepfake songs. Furthermore, the sentiment analysis models (Twitter-roBERTa-base model developed by CardiffNLP) used may contain limitations in accurately capturing the nuances of human emotions, particularly in mixed or ambiguous comments. Finally, focusing only on channels that self-label as "deepfake" may exclude secret or non-discrete deepfakes, and potentially over represent novelty seeking audiences.

VII. CONCLUSION AND FUTURE WORK

The analysis revealed that while the general sentiment towards *deepfake music* is rather positive, a comparison with *non-deepfake AI music* provides more revealing insights. Overall, listeners tend to be more negative towards *deepfake music*, in contrast to non-deepfake. One interpretation, which requires further investigation, is that it may be due to concerns regarding AI's ability to replicate human emotional depth and authenticity. However, based on observations of comments during the analysis, positive sentiment prevails when the genre or focus of the song involves entertainment factors (e.g., humor or satire) and emotional connection. Despite these concerns, the broader acceptance of AI-generated music, particularly *deepfake music*, is evident, with neutral and positive comments making up a significant portion of user feedback. Specifically, 37% of comments were neutral, indicating curiosity or ambivalence as the audience adapts to AI as a creative agent. Moreover, positive sentiments accounted for nearly 40% of the comments in deepfake genres like Rap, Hip-hop, Rock, and Pop, suggesting a promising level of acceptance. Overall, the combined neutral and positive sentiments (77%) outweigh the negative sentiments (23%), indicating a potential shift in public perception towards AI-composed or co-created music.

Regarding sentiments towards *non-deepfake AI music*, results revealed a largely positive sentiment. This contradicts previous studies which have shown a rather negative attitude towards AI music [5]. This finding may suggest that this issue is genre dependent. Listeners may be generally positive about non-deepfake AI music as it does not mimic human voice, thus being skeptical about deepfake AI-generated music using renown human voices. Reference [24] further elaborates on this skepticism, highlighting that emotional engagement with AI-composed music is a complex issue influenced by factors such as music genre, cultural perspective, and age group. Their research highlights the ongoing doubt and preference for human-made music, despite AI's potential to explore emotional authenticity. Thus, our findings align with these insights, emphasizing that the genre of music significantly influences listeners' attitudes towards AI-generated compositions.

Moreover, the longitudinal analysis also revealed relevant insights. First, regarding deepfake music, the genres of *Rap-Hip-Hop* and *Cartoon* indicated stable sentiment trends, with generally low negative sentiments, and predominance of positive and neutral ones. Both genres have also revealed a steady decrease in the total volume of comments, potentially suggesting that the novelty effect may be vanishing. On the other hand, the analysis indicated a different trend for *Rock/Pop*. In this genre, the increasing volume of comments and growing negative sentiment reinforces the need for further investigation towards deepfake music contrasting further genres. Lastly, the non-deepfake genres (*Instrumental* and *Relaxation*) indicated very stable trends, of low negative sentiments, and number of comments. This strengthens the notion that the mimicking of human voices through AI is the main factor to trigger sentimental responses towards AI music.

Finally, the limitations and the conclusions of the study indicate future directions for this investigation. First, future studies should extend the genre analysis, contrasting further genres (e.g., electronic, blues, jazz, country) to gain a more holistic understanding of the acceptance of deepfake music. Second, regarding results, no inferential statistics were reported due to formatting restrictions. Thus, not allowing the reporting of whether the differences found across sentiments are statistically significant, which represents a limitation of the paper. This will be addressed in the future. Third, future research should explore and further compare models using larger and more diverse datasets, including comments from different platforms and cultural contexts.

REFERENCES

[1] M. Mendez II, “The Drake AI song is just the tip of the iceberg,” *Time Magazine*, 2023. [Online]. Available: https://time.com/6273529/drake-the-weeknd-ai-song/. [Accessed: May 20, 2025].

[2] H. H. S. Josan, “AI and Deepfake Voice Cloning: Innovation, Copyright and Artists’ Rights,” *Artificial Intelligence*, 2024.

[3] A. N. Nicolaou and M. Murgia, “Google and Universal Music negotiate deal over AI ‘deepfakes’,” *Financial Times*, 2023. [Online]. Available: https://www.ft.com/content/6f022306-2f83-4da7-8066-51386e8fe63b. [Accessed: May 20, 2025].

[4] M. Sareen, “Threats and challenges by DeepFake technology,” in *DeepFakes*, Boca Raton, FL: CRC Press, 2022, pp. 99–113.

[5] F. Tigre Moura and C. Maw, “Artificial intelligence became Beethoven: how do listeners and music professionals perceive artificially composed music?,” *Journal of Consumer Marketing*, vol. 38, no. 2, pp. 137–146, 2021, doi: 10.1108/JCM-02-2020-3671.

[6] H. Chu, et al., “An empirical study on how people perceive AI-generated music,” in *Proc. 31st ACM Int. Conf. on Information & Knowledge Management (CIKM)*, 2022, pp. 304–314.

[7] F. Tigre Moura, C. Castrucci, and C. Hindley, “Artificial Intelligence Creates Art? An Experimental Investigation of Value and Creativity Perceptions,” *Journal of Creative Behavior*, vol. 57, pp. 534–549, 2023, doi: 10.1002/jocb.600.

[8] G. Tellis, E. Yin, and S. Bell, “Global consumer innovativeness: cross-country differences and demographic commonalities,” *Journal of International Marketing*, vol. 17, no. 2, pp. 1–22, 2009, doi: 10.1509/jimk.17.2.1.

[9] I. Tussyadiah and G. Miller, “Perceived impacts of artificial intelligence and responses to positive behaviour change intervention,” in *Information and Communication Technologies in Tourism 2019*, Cham: Springer Verlag, 2018, pp. 359–370. [Online]. Available: www.tussyadiah.com/ENTER2019_TussyadiahMiller.pdf.

[10] D. B. Shank, C. Stefanik, C. Stuhlsatz, K. Kacirek, and A. M. Belfi, “AI composer bias: Listeners like music less when they think it was composed by an AI,” *Journal of Experimental Psychology: Applied*, vol. 29, no. 3, p. 676, 2023.

[11] R. Zenieris, *Perception and Bias towards AI-Music*, Bachelor’s thesis, University of Twente, 2023..

[12] J. W. Hong, Q. Peng, and D. Williams, “Are you ready for artificial Mozart and Skrillex? An experiment testing expectancy violation theory and AI music,” *New Media & Society*, vol. 23, no. 7, pp. 1920–1935, 2021.

[13] P. Fernando, T. V. Mahanama, and M. Wickramasinghe, “Assessment of human emotional responses to AI–composed music: A systematic literature review,” in *Proc. 2024 Int. Res. Conf. on Smart Computing and Systems Engineering (SCSE)*, vol. 7, pp. 1–6, Apr. 2024

[14] E. Altuncu, V. Franqueira, and S. Li, “Deepfake: Definitions, performance metrics and standards, datasets and benchmarks, and a meta-review,” *arXiv preprint arXiv:2206.07788*, 2022.

[15] W. D. Weisman and J. F. Pena, “Face the uncanny: The effects of doppelganger talking head avatars on affect-based trust toward artificial intelligence technology are mediated by uncanny valley perceptions,” *Cyberpsychology, Behavior, and Social Networking*, vol. 24, pp. 182–187, 2021, doi: 10.1089/cyber.2020.0175.

[16] B. Usukhbayar and S. Homer, “Deepfake videos: The future of entertainment,” *ResearchGate*, Berlin, Germany, 2020.

[17] J. T. Hancock and J. N. Bailenson, “The social impact of deepfakes,” *Cyberpsychology, Behavior, and Social Networking*, vol. 24, no. 3, pp. 149–152, 2021.

[18] M. Sareen, “Threats and challenges by DeepFake technology,” in *DeepFakes*, Boca Raton, FL: CRC Press, 2022, pp. 99–113.

[19] D. Afchar, G. M. Brocal, and R. Hennequin, “Detecting music deepfakes is easy but actually hard,” *arXiv preprint arXiv:2405.04181*, 2024.

[20] Y. Zang, Y. Zhang, M. Heydari, and Z. Duan, “Singfake: Singing voice deepfake detection,” in *Proc. ICASSP 2024 IEEE Int. Conf. on Acoustics, Speech and Signal Processing (ICASSP)*, Apr. 2024, pp. 12156–12160

[21] S. Yang and H. Zhang, “Text mining of Twitter data using a latent Dirichlet allocation topic model and sentiment analysis,” *International Journal of Computer and Information Engineering*, vol. 12, no. 7, pp. 525–529, 2018.

[22] CardiffNLP, “Cardiff Natural Language Processing,” 2024. [Online]. Available: https://cardiffnlp.github.io/

[23] M. A. Jahin, M. S. H. Shovon, and M. F. Mridha, “TRABSA: Interpretable sentiment analysis of tweets using attention-based BiLSTM and Twitter-RoBERTa,” *arXiv preprint arXiv:2404.00297,* 2024.

[24] H. Guo, “Comparison of neural network and traditional classifiers for Twitter sentiment analysis,” *Highlights in Science Engineering and Technology*, vol. 38, pp. 1062–1070, 2023, doi: 10.54097/hset.v38i.5996

[25] J. Camacho-Collados et al., “TweetNLP: Cutting-edge natural language processing for social media,” in *Proc. Conf. on Empirical Methods in Natural Language Processing (EMNLP)*, 2022.

[26] M. A. Jahin, M. S. H. Shovon, and M. F. Mridha, “TRABSA: Interpretable sentiment analysis of tweets using attention-based BiLSTM and Twitter-RoBERTa,” *arXiv preprint arXiv:2404.00297*, 2024.

[27] V. Lac and F. T. Moura, “Sentiment analysis of AI generated music using latent Dirichlet allocation (LDA),” in *Proc. AIMC 2024*, Sep. 9–11, 2024.

[28] C. Provenzano, “Making voices: The gendering of pitch correction and the auto-tune effect in contemporary pop music,” *Journal of Popular Music Studies*, vol. 31, no. 2, pp. 63–84, 2019.

[29] M. Yasser, “Twitter Tweets Sentiment Dataset,” *Kaggle*, 2021. [Online]. Available: https://www.kaggle.com/datasets/yasserh/twitter-tweets-sentiment-dataset. [Accessed: May 22, 2025].