

AI and Musical Interpretation

Wolfgang Marx
Univerzitetni kolidž v Dublinu
University College Dublin

Introduction

It was only a question of time until generative artificial intelligence (AI) would move on from creating texts and pictures to the creation of music. While general AI content generators are expanding more and more into the music area, there are now also more and more music-specific generative AI systems such as Udio,¹ Soundraw² and DeepAI Music Generator.³ Many of them are accessible for free, although using the full range of capabilities usually requires signing up for a paid subscription. What is the potential impact of AI-generated music on questions of musical interpretation? My argument will be developed in three steps: firstly, a number of examples of AI-generated songs are presented in order to outline what AI can do at the moment. Secondly, the way in which large language models (LLMs) – the underlying systems of many AI generators – work is discussed, showing that LLMs have a tendency to standardise their results, while also removing any quirk or idiosyncrasy that human creators would (unconsciously) include. AI outputs are “depersonalised”, as it were. Thirdly, the concept of “model collapse” is introduced. It leads to further homogenisation as AI systems will inevitably include more and more data in their processes that are already AI-created, setting in motion a spiral of self-referentiality that

1 Udio Beta, <https://www.udio.com/>.

2 Soundraw, <https://soundraw.io/>.

3 Deep AI Music Generator, <https://deepai.org/music>.

is likely to lead to a collapse of the systems in the long run. It is here that the potential relevance of AI for musical interpretation will become clear.

AI Recreations Of Human Voices

Today generative AI is already used in many areas to create music. Yet this involves not just creating new pieces, or Vocaloids (voice-synthesising software) such as Hatsune Miku.⁴ The way AI “learns” from existing data in order to generate its output (discussed in the next section) lends itself to recreating or imitating the sound of existing voices and copying real-life artists. This can be done with the purpose of demonstrating the, by now, astonishing abilities of the AI systems. Yet another (and far more likely) utilisation occurs in order to make money by duping unsuspecting listeners who may think they are listening to a new track by a singer they love when they have actually encountered a fake AI recreation of that singer’s voice.

Such an AI-generated recreation of a singer’s voice first made major headlines in April 2023 when “Ghostwriter977” released the song “Heart on My Sleeve” on a range of streaming platforms, including TikTok, YouTube, Spotify and others. “Heart on My Sleeve” featured the voices of Canadian musicians Drake and The Weeknd, yet had been created without their knowledge or consent with the help of AI. Ghostwriter openly acknowledged the AI contribution to the song’s creation. It emerged later that the song was not entirely AI-generated: The still unknown “Ghostwriter” (who wears a bedsheet and sunglasses in his videos) stressed in an anonymous interview that “[i]t is *definitely my songwriting, my production and my voice*.”⁶ AI’s contribution thus consists of changing Ghostwriter’s voice so that it sounds like those of Drake and The Weeknd – what is called a “deepfake” of their voices.⁷ Within weeks the track was removed from all platforms at the behest of Universal Music Group, the company represent-

4 Hatsune Miku Official Website, <https://piapro.net/intl/en.html>.

5 Ghostwriter, “Ghostwriter – Heart on my Sleeve Drake ft. The Weeknd (Remastered AI song),” Buko, uploaded on 20 April 2023, YouTube video, 2:11, <https://www.youtube.com/watch?v=rQssjhX31Zo>.

6 Kristin Robinson, “Ghostwriter, the Mastermind Behind the Viral Drake AI Song, Speaks For the First Time,” *Billboard*, November 11, 2023, <https://www.billboard.com/music/pop/ghostwriter-heart-on-my-sleeve-drake-ai-grammy-exclusive-interview-1235434099/>.

7 Aiden Kenway explains in this YouTube video how the song was likely created with the help of AI: Aiden Kenway, “How the VIRAL Drake Song ‘Heart on my Sleeve’ was Made,” uploaded on 18 May 2023, YouTube video, 9:55, <https://www.youtube.com/watch?v=pkbQV71ZvK4>.

ing both Drake and The Weeknd. Yet the AI coup had already created a huge media echo.⁸ Joe Coscarelli called the episode in the *New York Times* “a harbinger of the headaches that can occur when a new technology crosses over into the mainstream consciousness of creators and consumers before the necessary rules are in place.”⁹ By now many corporations have engaged in the creation of AI versions of current singers’ voices in order to present songs that those vocalists have never recorded. Google actually released “Dream Track”, a system creating songs which imitate the voices of nine well-known singers including Demi Lovato, Jon Legend, Charlie Puth and others. All nine singers have agreed to their voices being used as part of this “experimental AI tool”. Apart from the voices Dream Track also provides “a stylistically relevant backing track.”¹⁰

All the examples mentioned thus far relate to singers currently active. They are by and large better known and have a larger fan base, so that the option to make money is greater. It is also easier to explain the emergence of a new song recorded by a singer who is still active. Yet AI has also been used to recreate the voices of much earlier famous vocalists. Hence we can now hear Frank Sinatra singing songs that he actually never recorded. One of them is “Smells Like Teen Spirit”, the most successful track from Nirvana’s 1991 album *Nevermind*.¹¹ Frank Sinatra was still alive when the Nirvana song was released (he died in 1998), yet never recorded it himself. In the AI version Nirvana’s alternative rock style is replaced by a big band backing, while converting Kurt Cobain’s grunge style and voice colour to Sinatra’s colour and crooning style. Getting both voice colour and singing style right is quite a challenge, and different recreations of voices can be more or less successful in that respect. AI Music World, the YouTube channel featuring Sinatra’s “Smells Like Teen Spirit”, offers many other AI-generated

8 See, for example: Eli Glastner, “AI-generated Drake/Weeknd Deep Fake Song Rattles Music Industry,” CBC News: The National, uploaded on 21 April 2023, YouTube video, 2:20, <https://www.youtube.com/watch?v=COok3YO1L5w>.

9 Joe Coscarelli, “An A.I. Hit of Fake ‘Drake’ and ‘The Weeknd’ Rattles the Music World,” *The New York Times*, April 24, 2023, <https://www.nytimes.com/2023/04/19/arts/music/ai-drake-the-weeknd-fake.html>.

10 Chloe Veltman, “Google’s Latest AI Music Tool Creates Track Using Famous Singers’ Voice Clones,” *NPR*, November 17, 2023, <https://www.npr.org/2023/11/17/1213551049/googles-latest-ai-music-tool-creates-tracks-using-famous-singers-voice-clones>.

11 “Frank Sinatra – Smells like Teen Spirit (AI Cover),” AI Music World, uploaded on 11 August 2023, YouTube video, 2:40, <https://www.youtube.com/watch?v=Numoq-l-lldc>. The original Nirvana song can also be found on YouTube: Nirvana, “Nirvana – Smells Like Teen Spirit (Official Music Video),” uploaded 16 June 2009, YouTube video, 4:38, <https://www.youtube.com/watch?v=hTWKbfoikeg>.

Sinatra “adaptations”. These include: “Say Something” by A Great Big World (2013);¹² “Still D.R.E.” by Dr. Dre (1999; here the unsuspecting listener can hear “Sinatra” rapping);¹³ “Bad Habits” by Ed Sheeran (2021; in this adaptation “Sinatra” curiously stays off-pitch much of the time);¹⁴ and others. The last two examples are perhaps less convincing than the Nirvana clone as Sinatra here sings in a style that he would probably never have adopted in real life, even if his voice may be imitated in a passable way. Not all AI-imitated voices are fully convincing though; the YouTube channel “Elvis Presley AI” presents a range of AI Elvis clones which often do not manage to reproduce the fullness of the singer’s rich baritone voice, as exemplified by “Baby Got Back”.¹⁵

How Large Language Models Work

The breakthrough in the development of a generative AI that can be used by anyone to create texts, pictures and music was the development of “Large Language Models” (LLMs) which can cope with enormous amounts of data. However, it has to be stressed that LLMs are not databases but instead algorithms that create texts, pictures, music and other media on the basis of examples in a database it has access to, coupled with a set of rules it has been programmed to apply to them. The larger the amount of data the LLM has access to the more convincing its output will be. We will look here at the creation of texts, because this is where the principle was developed, and it is a bit easier to understand. Later we will move on to the creation of music though generative AI.

When creating a text, an LLM checks a vast amount of data to find out which word is usually used following to the previous one in a chain of words (or next to a previous phrase, or sentence, or paragraph – and all that in the topical context of what has been requested of it), adding that word, and repeating the process all over again. The broader the contextual information is the more convincing the result will appear. Thus the task is to

12 “Frank Sinatra – Say Something (AI Cover),” AI Music World, uploaded on 13 September 2023, YouTube video, 2:40, <https://www.youtube.com/watch?v=6Rgv-BAoent4>.

13 “Frank Sinatra – Still D.R.E. (AI Cover),” AI Music World, uploaded on 2 September 2023, YouTube video, 3:38, <https://www.youtube.com/watch?v=VrsNCIZ8F38>.

14 “Frank Sinatra – Bad Habits (AI Cover),” AI Music World, uploaded on 20 August 2023, YouTube video, 4:11, <https://www.youtube.com/watch?v=jbbyGF1t5Eg>.

15 “Baby Got Back – Elvis Presley,” Elvis Presley AI, uploaded on 5 June 2023, YouTube video, 0:52, https://www.youtube.com/watch?v=_uy2NQWal-4.

train a neural network (the LLM) to predict the next word in a given sequence of words, no matter if that sequence is long or short, in German or in English or in any other language, whether it's a tweet or a mathematical formula, a poem or a snippet of code. All of those are sequences that we will find in the training data.¹⁶

To achieve this, words are “translated” into word vectors, i.e., numbers that allow for those statistical operations to be undertaken more smoothly.¹⁷ This means the result of an LLM operation is what is statistically most likely to be convincing on the basis of the model's extensive database.

The LLM also has certain pre-programmed rules (or “parameters”), such as when to use a noun, a verb or a participle (or perhaps certain chords, or voice-leading rules), or which words are more likely to appear in a certain thematic context (beyond the purely grammatical and syntactical context). There is also “instruction tuning”, which teaches the system what human operators are likely to expect as a response to a given input. Finally, there is learning through human feedback – the algorithm constantly interacts with human operators and changes its rules on the basis of the responses it gets from them.¹⁸ This means that those rules are not static – they can be changed by the system itself over time on the basis of what it learns from an expanded data set. This is the reason why after a while even the algorithm's developers do not know in detail any more what the set of rules is that the algorithm is applying.

The questions or requests put to an LLM are called “prompts”, and their “design” is crucial to ensure that the response fully covers what the prompter expects it to engage with. Good prompt design is a central aspect of the successful use of generative AI, and it is often advisable to vary prompts several times to see how the responses change (at a basic level this may be comparable to the use of search engines, where a similar principle applies).

It is important to note that the algorithm's goal is to create a text that appears as “human” as possible, that is, a convincing answer that passes the

16 Andreas Stöffelbauer, “How Large Language Models Work. From Zero to ChatGPT,” *Data Science at Microsoft*, October 24, 2023, <https://medium.com/data-science-at-microsoft/how-large-language-models-work-91c362f5b78f>. This article also discusses the AI's analysis of pictures and music.

17 Timothy B. Lee, “A Jargon-Free Explanation of How Large Language Models Work,” *Ars Technica*, July 31, 2023, <https://arstechnica.com/science/2023/07/a-jargon-free-explanation-of-how-ai-large-language-models-work/>.

18 See: Stöffelbauer, “How Large Language Models Work.”

Turing test – and it has by now become very good at that. However, it is incapable of determining whether what it says is actually correct or true. This criterion is just not built into the system, and there is no way that the statistical method it is based on could achieve this result. Most of what it says is likely to be correct, yet any text on a topic I know something about virtually always includes one or several so-called “hallucinations”, that is statements which are incorrect (I do not like the term “hallucinations” very much as it seems to indicate the presence of a conscious “something” in the machine, which is not the case). Generative AI systems can thus be described as the ultimate bullshitters – they sound incredibly convincing, yet create their output without any concern for what is correct or true.¹⁹ Another potential issue with generative AI is that it is difficult – if not impossible – for it to be really creative or original. This is because what it creates is based on what already exists. Postmodern combinations of unusual and distant stylistic traits can be regarded as original and creative, of course, yet generative AI does not go for this kind of “outlier” – instead the texts, pictures and music generated by AI often appear to be cursory and stereotypical. If they are given the same prompt repeatedly the results are never exactly identical, yet often very similar with the differences being of a superficial nature that affect (in the case of texts) wording more than content. Texts as well as pictures often come in a very anodyne style lacking the quirks and idiosyncrasies that most human authors and creators have, and that give a text a personal note. The nature of the algorithm ensures that such a personalisation cannot occur – instead it is at risk of evolving towards the so-called “model collapse” which will be discussed in the next section.

Model Collapse And The Standardisation Of AI Results

How can AI relate to issues of musical interpretation? Since AI is still relatively new, the answer has to some degree be speculative, particularly with regard to music. While the number of AI-generated texts is already vast, the amount of AI-generated music is much lower – partly because there is more demand for AI-generated texts, yet also because LLMs started generating pictures and music later than texts. Yet since, as we have seen, all LLMs

19 I use the term “bullshitting” here in the sense in which the philosopher Harry G. Frankfurt defined it in his essay *On Bullshit* (Princeton, NJ: Princeton University Press, 2005). He distinguishes bullshitting from lying by outlining that bullshitting doesn’t require any knowledge of the truth and isn’t concerned with it at all. This makes bullshitters more dangerous than liars; populist politicians are prime examples of bullshitters in today’s world.

operate following the same principles the resulting problems are likely to be very similar, too.

The main issue causing problems is called “*model collapse*”.²⁰ It was the subject of an article in *Nature*²¹ in July 2024 which triggered a number of other articles in popular science journals and on tech websites discussing the phenomenon. However, some experts already issued warnings about it a year ago, when applications such as ChatGPT were still relatively new. Carl Franzen, for example, asked in June 2023: “*What happens as AI-generated content proliferates around the internet, and AI models begin to train on it, instead of on primarily human-generated content?*”²²

As discussed in the previous section, LLMs operate by scouring a very large database – ideally the entire world wide web – in order to determine what the most convincing response to a prompt might be. Model collapse occurs if a larger and larger section of the data consulted consists of data previously created by AI, as pointed out by Shumailov et al.: “*indiscriminately learning from data produced by other models causes ‘model collapse’—a degenerative process whereby, over time, models forget the true underlying data distribution.*”²³ This means

that, over time, models start losing information about the true distribution, which first starts with tails disappearing, and learned behaviours converge over the generations to a point estimate with very small variance.²⁴

20 Devin Coldewey, “‘Model Collapse’: Scientists Warn Against Letting AI Eat Its Own Tail,” *Tech Crunch*, July 24, 2024, <https://techcrunch.com/2024/07/24/model-collapse-scientists-warn-against-letting-ai-eat-its-own-tail/>.

21 Ilia Shumailov, et al., “AI Models Collapse When Trained on Recursively Generated Data,” *Nature* 631 (24 July 2024): 755–9, <https://www.nature.com/articles/s41586-024-07566-y>. Shumailov et al. discuss the impact of model collapse not just on LLMs yet also on other AI systems such as Gaussian Mixture Models (GMMs) and Variational Autoencoders (VAEs), yet these don’t need to be discussed further in our context.

22 Carl Franzen, “The AI Feedback Loop: Researchers Warn of ‘Model Collapse’ as AI Trains on AI-Generated Content,” *Venture Beat*, June 12, 2023, <https://venturebeat.com/ai/the-ai-feedback-loop-researchers-warn-of-model-collapse-as-ai-trains-on-ai-generated-content/>. Franzen’s text was based on an earlier article by the Shumailov team that was already submitted on 27 May 2023: Ilia Shumailov, et al., “The Curse of Recursion: Training on Generated Data Makes Models Forget,” *arXiv* 2305, no. 17493 (27 May 2023), <https://doi.org/10.48550/arXiv.2305.17493>.

23 Shumailov, et al., “AI Models Collapse.”

24 Ibid.

What does this mean in practice? As Maggie Harrison Dupré points out,

when you feed synthetic content back to a generative AI model, strange things start to happen. Think of it like data inbreeding, leading to increasingly mangled, bland, and all-around bad outputs.²⁵

Devin Coldewey presents an interesting practical example in relation to pictures of different species of dogs (which are reproduced in his article).

models gravitate toward the most common output. It won't give you a controversial snickerdoodle recipe but the most popular, ordinary one. And if you ask an image generator to make a picture of a dog, it won't give you a rare breed it only saw two pictures of in its training data; you'll probably get a golden retriever or a Lab.

Now, combine these two things with the fact that the web is being over-run by AI-generated content and that new AI models are likely to be ingesting and training on that content. That means they're going to see a *lot* of goldens!²⁶

Hence the output of AI operations based on more and more content already generated by AI is likely to be more and more standardised – or, as Shumailov et al. put it as quoted above, with a “*very small variance*”. Coldewey juxtaposes pictures created by an AI image generator fed only with “real” data with another one influenced significantly by AI-generated pictures, which does indeed result in a very obvious lack of variety. The degenerative effect comes to the fore even more if the process is repeated. AI now trains itself on data already produced by AI, on the basis of a previous AI picture and so on. The results become less and less distinctive, until they don't look like a dog – any dog – at all anymore. As Tor Constantino elaborates,

[a]fter the first two prompts, the answers steadily miss the mark, followed by a significant quality downgrade by the fifth attempt and a complete devolution to nonsensical pabulum by the ninth consecutive query.²⁷

25 Maggie Harrison Dupré, “When AI is Trained on AI-generated Data, Strange Things Start to Happen,” *Futurism*, February 8, 2023, <https://futurism.com/ai-trained-ai-generated-data-interview>.

26 Coldewey, “Model collapse’.”

27 Tor Constantino, “Is AI Quietly Sabotaging Itself – And the Internet?” *Forbes*, August 26, 2024, <https://www.forbes.com/sites/torconstantino/2024/08/26/is-ai-quietly-killing-itself-and-the-internet/>.

And, as Shumailov et al. stress, once the AI starts training itself on AI-generated content model collapse becomes ultimately inevitable.

Carl Franzen presents another instructive animal-based example of model collapse:

a machine learning model is trained on a dataset with pictures of 100 cats — 10 of them with blue fur, and 90 with yellow. The model learns that yellow cats are more prevalent, but also represents blue cats as more yellowish than they really are, returning some green-cat results when asked to produce new data. Over time, the original trait of blue fur erodes through successive training cycles, turning from blue to greenish, and ultimately yellow. This progressive distortion and eventual loss of minority data characteristics is model collapse.²⁸

The basic problem in the examples of both dogs and cats is that rare yet existing cases are more and more neglected and eventually weeded out completely in the process of generating new output. Hence Franzen insists that

it's important to ensure fair representation of minority groups in datasets, in terms of both quantity and accurate portrayal of distinctive features. The task is challenging due to models' difficulty learning from rare events.²⁹

So far cases of LLMs operating entirely on the basis of input already generated by AI only exist in artificially constructed situations like the ones described in the articles referenced above. Yet the percentages of fully AI-generated or at least partly AI-influenced data on the web is continuously increasing.

I could not find examples of model collapse exemplified through music, yet given that generative AI creates texts, pictures and music based on the same principles, it is to be expected that the concept applies to AI-generated music in the same way. In the real world, within the data sets that are training the AI, fully AI-generated content may not be prevalent yet, but the percentage of partly AI-generated content it is certainly higher than many of us may guess. Constantino quotes a study by Amazon Web Services according to which 57 % of all content on the internet that is available in three or more languages has very likely been translated by AI, also referring

28 Franzen, "The AI Feedback Loop."

29 Ibid.

to (yet not referencing) a prediction that by 2025 more than 90 % of all web content may be AI-generated to some extent.³⁰ In September 2024,

[t]he creator of an open source project that scraped the internet to determine the ever-changing popularity of different words in human language usage says that they are sunseting the project because generative AI spam has poisoned the internet to a level where the project no longer has any utility.³¹

How can model collapse affect AI-generated music, and in particular issues of musical interpretation? Taking the example of AI-generated Frank Sinatra songs, if an AI system trains itself for the task of creating a “new” Sinatra by checking any Sinatra song that has been uploaded to the web this is fine as long as those are all “real” Sinatra recordings. If, however, some of them already are “AI Sinatras”, this means model collapse is starting to set in. If the number of AI Sinatras keeps increasing while the number of “real” Sinatra songs remains the same (after all, he is not producing new recordings anymore), we enter the spiral of AI training itself more and more on its own previous creations, with the likely result that a homogenised, “standard” Sinatra – the golden retriever of Sinatras, so to speak – is going to emerge, and the specific idiosyncrasies or traits that only occur in a few songs will be more and more suppressed. A “one and only” AI way of how Sinatra sings will be the result, ultimately followed by Constantino’s *“complete devolution to nonsensical pablum.”*

Some AI-created songs – usually those meant to demonstrate the abilities of the AI systems – are publicly designated as such, although I doubt whether AI systems are able to distinguish those from the recordings by the real artists. However, most AI-generated music is not identified as such anyway. This particularly applies to functional music used in areas like advertising, or background music of any kind such as contemporary music. Here the risk of music entering the model collapse spiral is real, and if researchers such as Shumailov are correct we will all witness the ever-increasing blandness and eventual degeneration of these musics in the not-too-distant future (yes, those musics are already bland today, yet it can still get worse). Even in the case of AI imitating well-known artists like Sinatra it is likely

30 Constantino, “Is AI Quietly Sabotaging Itself.”

31 Jason Koebler, “Project Analyzing Human Language Usage Shuts Down Because ‘Generative AI Has Polluted the Data,’” *404 Media*, September 19, 2024, <https://www.404media.co/project-analyzing-human-language-usage-shuts-down-because-generative-ai-has-polluted-the-data/>.

that model collapse will negatively affect songs the later they are being produced. In those cases model collapse will probably occur at a much slower pace since songs of this kind are created in much smaller numbers than advertising or background musics, yet it will eventually become noticeable too. And if eventually a homogeneous AI standard of Sinatra's singing emerges there is even a risk that people who did not hear his recordings before the emergence of AI (necessarily an ever-increasing cohort over time) may take it to be the real thing. All of this can, of course, be applied not just to voices, yet also to the sounds and styles of instrumentalists or ensembles (such as specific big bands, for example) in all genres.

Imitating the voice of a specific singer runs into copyright problems, and Tennessee has recently become the first US state to restructure its copyright law and ban the unauthorised use of AI to imitate the work and style of copyright-protected artists.³² However, there will always be states and territories which have not signed up to international copyright agreements. Servers located there can be used to distribute unauthorised imitations that are illegal elsewhere, and attempts to block them will probably be ultimately unsuccessful.

Conclusion

Generative AI has the potential to limit the range of musical interpretation. At least for the time being "deepfake" imitations of performers' voices are often not entirely convincing. While this aspect is likely to get better over time, two other problems cannot be overcome that easily. Firstly, the AI systems "compose" music by "averaging out" the data they study while operating, meaning that a certain homogenisation of the results is virtually inevitable. This applies both when the style of specific pieces or performers is imitated, yet also when music in certain genres or styles is requested. Secondly, it is very likely that – as the number of AI-created pieces of music increases – an onset of model collapse, that is a spiral of self-referentiality, will be set in motion that will initially reinforce the homogenisation effect before potentially leading to a collapse of the systems in the long run. Generative AI has parameters that give it a certain leeway when selecting results, yet comparing different texts generated in response to the same

32 Rebecca Rosman, "Tennessee Becomes the First State to Protect Musicians and Other Artists Against AI," *NPR*, March 22, 2024, <https://www.npr.org/2024/03/22/1240114159/tennessee-protect-musicians-artists-ai>. The new law is known as the ELVIS act, an acronym for Ensuring Likeness Voice and Image Security Act.

prompt indicate that this mainly affects the surface level, while at a deeper level it is most likely that very similar results will be the outcome. Human performers do not have to follow those standardised leads offered by AI, yet as the percentage of AI-generated music increases so does the risk that the listening habits of many will be more and more affected by it. While this may not be inevitable it poses enough of a risk to give us food for thought and spend more time assessing the danger posed to interpretative variety by generative AI.

Bibliography

Literature

Frankfurt, Harry G. *On Bullshit*. Princeton, NJ: Princeton University Press, 2005.

Periodicals

Coldewey, Devin. "‘Model Collapse’: Scientists Warn Against Letting AI Eat Its Own Tail." *Tech Crunch*, July 24, 2024. <https://techcrunch.com/2024/07/24/model-collapse-scientists-warn-against-letting-ai-eat-its-own-tail/>.

Constantino, Tor. "Is AI Quietly Sabotaging Itself – And the Internet?" *Forbes*, August 26, 2024. <https://www.forbes.com/sites/torconstantino/2024/08/26/is-ai-quietly-killing-itself-and-the-internet/>.

Coscarelli, Joe. "An A.I. Hit of Fake ‘Drake’ and ‘The Weeknd’ Rattles the Music World." *The New York Times*, April 24, 2023. <https://www.nytimes.com/2023/04/19/arts/music/ai-drake-the-weeknd-fake.html>.

Franzen, Carl. "The AI Feedback Loop: Researchers Warn of ‘Model Collapse’ as AI Trains on AI-Generated Content." *Venture Beat*, June 12, 2023. <https://venturebeat.com/ai/the-ai-feedback-loop-researchers-warn-of-model-collapse-as-ai-trains-on-ai-generated-content/>.

Glastner, Eli. "AI-generated Drake/Weeknd Deep Fake Song Rattles Music Industry." CBC News: The National. Uploaded on 21 April 2023. YouTube video, 2:20. <https://www.youtube.com/watch?v=COok3YO1L5w>.

Harrison Dupré, Maggie. "When AI is Trained on AI-generated Data, Strange Things Start to Happen." *Futurism*, February 8, 2023. <https://futurism.com/ai-trained-ai-generated-data-interview>.

Kenway, Aiden. "How the VIRAL Drake Song ‘Heart on my Sleeve’ was Made." Uploaded on 18 May 2023. YouTube video, 9:55. <https://www.youtube.com/watch?v=pkbQV71ZvK4>.

Koebler, Jason. "Project Analyzing Human Language Usage Shuts Down Because ‘Generative AI Has Polluted the Data.’" *404 Media*, September 19, 2024.

- <https://www.404media.co/project-analyzing-human-language-usage-shuts-down-because-generative-ai-has-polluted-the-data/>.
- Lee, Timothy B. “A Jargon-Free Explanation of How Large Language Models Work.” *Ars Technica*, July 31, 2023. <https://arstechnica.com/science/2023/07/a-jargon-free-explanation-of-how-ai-large-language-models-work/>.
- Robinson, Kristin. “Ghostwriter, the Mastermind Behind the Viral Drake AI Song, Speaks For the First Time.” *Billboard*, November 11, 2023. <https://www.billboard.com/music/pop/ghostwriter-heart-on-my-sleeve-drake-ai-grammy-exclusive-interview-1235434099/>.
- Rosman, Rebecca. “Tennessee Becomes the First State to Protect Musicians and Other Artists Against AI.” *NPR*, March 22, 2024. <https://www.npr.org/2024/03/22/1240114159/tennessee-protect-musicians-artists-ai>.
- Shumailov, Ilia, et al. “The Curse of Recursion: Training on Generated Data Makes Models Forget.” *arXiv* 2305, no. 17493 (27 May 2023). <https://doi.org/10.48550/arXiv.2305.17493>.
- Shumailov, Ilia, et al. “AI Models Collapse When Trained on Recursively Generated Data.” *Nature* 631 (24 July 2024): 755–9. <https://www.nature.com/articles/s41586-024-07566-y>.
- Stöffelbauer, Andreas. “How Large Language Models Work. From Zero to ChatGPT.” *Data Science at Microsoft*, October 24, 2023. <https://medium.com/data-science-at-microsoft/how-large-language-models-work-91c362f5b78f>.
- Veltman, Chloe. “Google’s Latest AI Music Tool Creates Track Using Famous Singers’ Voice Clones.” *NPR*, November 17, 2023. <https://www.npr.org/2023/11/17/1213551049/googles-latest-ai-music-tool-creates-tracks-using-famous-singers-voice-clones>.

Other Sources

- “Baby Got Back – Elvis Presley.” Elvis Presley AI. Uploaded on 5 June 2023. YouTube video, 0:52. https://www.youtube.com/watch?v=_uy2NQWal-4.
- Deep AI Music Generator*. <https://deepai.org/music>.
- “Frank Sinatra – Smells like Teen Spirit (AI Cover).” AI Music World. Uploaded on 11 August 2023. YouTube video, 2:40. <https://www.youtube.com/watch?v=Numoq-l-ldc>.
- “Frank Sinatra – Bad Habits (AI Cover).” AI Music World. Uploaded on 20 August 2023. YouTube video, 4:11. <https://www.youtube.com/watch?v=jbbyGF1t5Eg>.
- “Frank Sinatra – Still D.R.E. (AI Cover).” AI Music World. Uploaded on 2 September 2023. YouTube video, 3:38. <https://www.youtube.com/watch?v=VrsNCIZ8F38>.

“Frank Sinatra – Say Something (AI Cover).” AI Music World. Uploaded on 13 September 2023. YouTube video, 2:40. <https://www.youtube.com/watch?v=6RgvBAoent4>.

Ghostwriter. “Ghostwriter – Heart on my Sleeve Drake ft. The Weeknd (Remastered AI song).” Buko. Uploaded on 20 April 2023. YouTube video, 2:11. <https://www.youtube.com/watch?v=rQssjhX31Zo>.

Hatsune Miku Official Website. <https://piapro.net/intl/en.html>.

Nirvana. “Nirvana – Smells Like Teen Spirit (Official Music Video).” Uploaded on 16 June 2009. YouTube video, 4:38. <https://www.youtube.com/watch?v=hTWKbfoikeg>.

Soundraw. <https://soundraw.io/>.

Udio Beta. <https://www.udio.com/>.