

Meta's AI MusicGen makes music from text prompts

June 13 2023, by Peter Grad



Credit: Unsplash/CC0 Public Domain

The 19th century philosopher Thomas Carlyle once declared, "Music is well said to be the speech of angels."

One wonders what Tom would have thought about a later era serving up protopunk, death metal and gothic rock.

Moreover, what would he have had to say about current approaches to [music](#) composed not by angels but by CPUs, RAM and large language models?

Meta announced last week that it has developed an AI music processor that generates music based on natural language descriptions.

Meta's [MusicGen](#), which follows on the heels of Google's January release of MusicLM that generates music based on text prompts or humming, was trained on 20,000 hours of music. Its [open source code](#) is available on Github and the model can be tested online at Hugging Face.

Users enter a brief description of music they want to hear such as "A dynamic blend of hip-hop and orchestral elements, with sweeping strings and brass, evoking the vibrant energy of the city." Or the description could be simpler: "90s rock [song](#) with a guitar riff."

Optionally, a song may be uploaded to help guide the creation of desired content.

MusicGen generates a 12-second clip within a couple of minutes.

We present MusicGen: A simple and controllable music generation model. MusicGen can be prompted by both text and melody.

We release code (MIT) and models (CC-BY NC) for open research, reproducibility, and for the music community:

<https://t.co/OkYjL4xDN7> pic.twitter.com/h114LGzYgf

— Felix Kreuk (@FelixKreuk) [June 9, 2023](#)

According to their evaluations of the program, Meta found MusicGen did better on both objective and subjective measures than other comparable programs such as MusicLM, Diffusion and Noise2Music.

"MusicGen yields high quality samples which are better melodically aligned with a given harmonic structure, while adhering to a textual description," Meta reported in a paper published June 8 on the *arXiv* preprint server.

The tool is seen as a potentially invaluable aid to composers and performers.

Meta tested three versions of their model. They varied in the amount of music detail provided: 300 million, 1.5 billion and 3.3 billion parameters.

In evaluations, Meta found humans preferred the results from the middle range (1.5 billion parameter) model. This perhaps reflects the French electronic and music producer Rone's theory for successful music production: "Less is more."

The model with the greatest number of parameters, however, generated output that most accurately reflected text and audio input.

Inevitably, as with so much of AI projects in other fields, concerns will arise. Foremost are [legal issues](#) surrounding the use of copyrighted material.

Meta says all songs used in training are cleared by legal agreements with copyright holders.

But users who add a song or artist's name to their description may open the door to potential copyright infringement. Google's MusicLM bars

users from including artist's names, but Meta doesn't.

And since Meta allows users to also upload a song to be used to help shape the final output, the line between original and copied content is blurred.

Australian singer and composer Nick Cave, addressing the topic of AI music earlier this year, left no doubt where he stands on AI music.

"Songs arise out of suffering ... they are predicated upon the complex, internal human struggle of creation," he said. "As far as I know, algorithms don't feel. Data doesn't suffer. ChatGPT has no inner being, it has been nowhere, it has endured nothing, it has not had the audacity to reach beyond its limitations, and hence it doesn't have the capacity for a shared transcendent experience, as it has no limitations from which to transcend."

Similarly, The Guardian's Ben Beaumont-Thomas asserted recently, "AI will always be a tribute act. It may be a very good tribute act, the type that, were it a human, would get year-round bookings on cruise ships and in Las Vegas casinos."

But Spotify's CEO Daniel Ek has a more upbeat outlook on AI-generated music.

"This could be potentially huge for creativity... That should lead to more music [which] we think is great culturally," he said.

More information: Jade Copet et al, Simple and Controllable Music Generation, *arXiv* (2023). [DOI: 10.48550/arxiv.2306.05284](https://doi.org/10.48550/arxiv.2306.05284)

- MusicGen: Simple and Controllable Music Generation: ai.honu.io/papers/musicgen/

- Code and models: github.com/facebookresearch/audiocraft
- Demo: huggingface.co/spaces/facebook/MusicGen
- Felix Kreuk's Twitter thread:
twitter.com/FelixKreuk/status/1667086356927901696

© 2023 Science X Network

Citation: Meta's AI MusicGen makes music from text prompts (2023, June 13) retrieved 29 April 2024 from <https://techxplore.com/news/2023-06-meta-ai-musicgen-music-text.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.