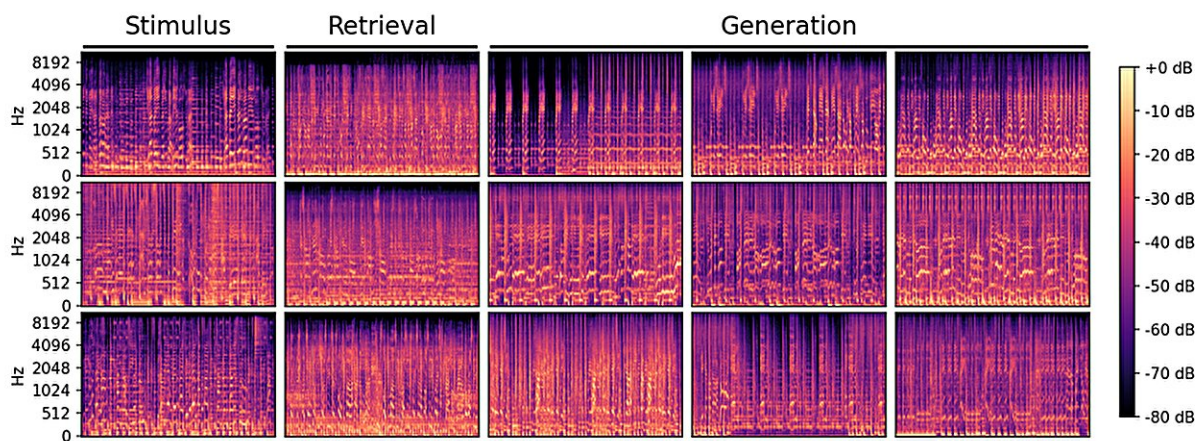


# Brain2Music taps thoughts to reproduce music

July 27 2023, by Peter Grad



Spectrograms of different music clips: The left-most column contains the stimulus which subjects were exposed to. To the right is the music retrieved from FMA and three clips sampled from MusicLM. Both generation and retrieval are done via MuLan embeddings. It is visually perceptible that spectrograms in the same row resemble similarities. Audio examples (randomly sampled, one per genre) can be found at [google-research.github.io/seanet/brain2music#ret-vs-gen](https://research.github.io/seanet/brain2music#ret-vs-gen). Credit: *arXiv* (2023). DOI: 10.48550/arxiv.2307.11078

Legendary Stones guitarist Keith Richards once said, "Music is a language that doesn't speak in particular words. It speaks in emotions, and if it's in the bones, it's in the bones."

Keith knows [music](#), but researchers at Google and Osaka University in Japan know brain activity and they have reported progress in reconstructing music, not from bones but from human brain waves observed in the laboratory. The team's paper, "Brain2Music: Reconstructing Music from Human Brain Activity," was published on the preprint server *arXiv* on July 20.

Music samples covering 10 genres including rock, classical, metal, hip-hop, pop and jazz were played for five subjects while researchers observed their [brain](#) activity. Functional MRI (fMRI) readings were recorded while they listened. (fMRI readings, unlike MRI readings, record metabolic activity over time.)

The readings were then used to train a [deep neural network](#) that identified activities linked to various characteristics of music such as genre, mood and instrumentation.

An intermediate step brought MusicLM into the study. This model, designed by Google, generates music based on text descriptions. It, like the fMRIs, measures factors such as instrumentation, rhythm and emotions. An example of text input is: "Meditative song, calming and soothing, with flutes and guitars. The music is slow, with a focus on creating a sense of peace and tranquility."

The researchers linked the MusicLM database with the fMRI readings, enabling their AI model to reconstruct music that the subjects heard. Instead of text instructions, [brain activity](#) provided context for musical output.

"Our evaluation indicates that the reconstructed music semantically resembles the original music stimulus," said Google's Timo Denk, one of several authors of the recent paper.

They aptly named their AI model Brain2Music.

"The generated music resembles the musical stimuli that human subjects experienced, with respect to semantic properties like genre, instrumentation and mood," he said. They further identified [brain regions](#) reflecting information originating from text descriptions of music.

Examples provided by the team demonstrate remarkably similar sounding excerpts of music interpreted by Brain2Music based on brainwaves of subjects.

One of the sampled songs was one of the earliest Top 10 hits of the year 2000, "Oops!... I Did It Again," by Britney Spears. A number of the song's musical elements such as the sound of the instruments and the beat closely matched, although lyrics were unintelligible. Brain2Go focuses on instrumentation and style, not lyrics, the researchers explained.

"This study is the first to provide a quantitative interpretation from a biological perspective," Denk said. But he acknowledged that despite advances in text-to-music models, "their internal processes are still poorly understood."

AI is not yet ready to tap into our brains and write out perfectly orchestrated tunes—but that day may not be too far away.

Future work on music generation models will bring improvements to "the temporal alignment between reconstruction and stimulus," Denk said. He speculated that ever-more faithful reproductions of musical compositions "from pure imagination" lie ahead.

Perhaps future songwriters will need only imagine the chorus of a song

while a printer wirelessly connected to the auditory cortex prints out the score. Beethoven, who began going deaf in his 20s, ruined pianos by pounding the keys so he could better hear the music. He'd get a kick—and spare more than a few pianos—out of Brain2Music.

And Paul McCartney, who authored "Yesterday," voted in a BBC poll in 1999 the best song of the 20th century, famously explained the idea for the song came to him in a dream, yet it took a year and a half for him to fill in all the missing parts. If a future McCartney comes up with a potential global hit in a midnight dream, a Brain2Go-type model would likely ensure a complete, speedy, accurate rendering awaiting the author at the breakfast table.

**More information:** Timo I. Denk et al, Brain2Music: Reconstructing Music from Human Brain Activity, *arXiv* (2023). [DOI: 10.48550/arxiv.2307.11078](https://doi.org/10.48550/arxiv.2307.11078)

Brain2Music: [google-research.github.io/seanet/brain2music/](https://google-research.github.io/seanet/brain2music/)

© 2023 Science X Network

Citation: Brain2Music taps thoughts to reproduce music (2023, July 27) retrieved 27 April 2024 from <https://techxplore.com/news/2023-07-brain2music-thoughts-music.html>

<p>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.</p>
--