

## Critical Response #1: What do you (really) want from AI music generation?

Music 356 / CS 470 [“Music and AI”](#)

Instructor: [Ge Wang](#) | Stanford University | Winter 2023

Do a critical reading and response for MusicLM (or a similar system), the recently announced prompt-based text-to-musical audio generation model, and more broadly probe the question, "*what do you (really) want from AI music generation?*"

FYI while the model for MusicLM is not (yet) released—some examples, a paper, and a dataset are available: <https://google-research.github.io/seanet/musiclm/examples/>

The MusicLM paper: <https://arxiv.org/pdf/2301.11325.pdf>

**Things to think about.** Gain a sense of the merits and limitations of MusicLM (or a specific similar system) for what it aims to do (the examples are a good place to start for this). At the same time, as a critical exercise, think as broadly as possible about its implications for culture, artistic practice, its promise and its perils, and who might be impacted and in what ways. Question the premise of the stated goals of the system—possibly starting with a critical (not cynical) pondering of the question, "is this even a good idea?" Favor lens-building over opinion sharing; favor well-framed questions over verdicts. One fun and potentially telling exercise is to imagine a version of the world where the system in question is in prevalent use. In what ways does that version of the world align with the world you would want to see? In what ways does it not?

Pick a particular aspect of your reflections on this question and write a ~500-word response (yes, you can optionally use chatGPT to help with the task, but the viewpoint must be *yours*).

Happy reflecting. For your deliverable, please turn in a webpage URL that contains your reflections.