

Something from anything: guiding random numbers to music

by Christian M. Denis

Introduction: Recently we've seen a rise in the number of AI-generated contents. It is now common to stumble upon an AI-generated image or have someone casually say that they've used ChatGPT for a certain task. One area of AI generation that is not yet quite prominent is that of music generation. Some tools do exist, but they have not been injected and adopted by the public masses. This is only a matter of time.

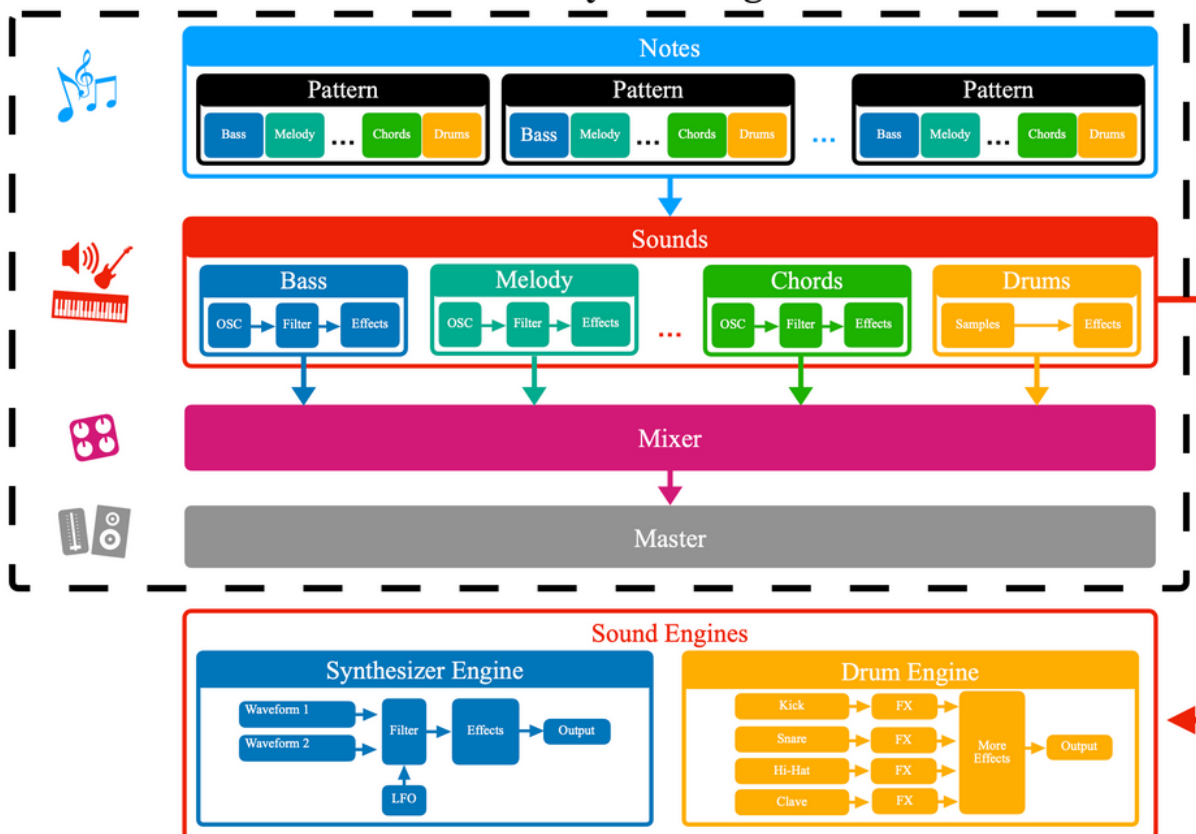
Machines creating music is not a new idea, and in fact, has been around for a long time. There is a distinction to be made in the way a machine synthesizes music. One type of "brute-force" method can use machine learning (popularly referred to as AI) and a large data set of existing music to "learn" how to make music. In this scheme, the music that was used to train the machine will influence the output of the program. It is to be expected that the output music will sound similar to the one used for training. A different method for machine-generated music is through simpler hard-coded algorithms. For

instance, we can imagine an algorithm that was designed to randomly choose notes from a specific scale. This is hard-coded in the program because the programmer has the knowledge that using notes in the same scale will sound generally harmonious (and of course, that is a subjective/cultural perspective). Without going too much into detail, we can imagine how someone writing a program to make music might use different rules related to chord progressions, melodies, and rhythms for composing.

I've mentioned two different methods for machine-assisted music-making. The one this project focuses on is the latter: an algorithmic hard-coded method. In other words, the project does not use AI, but because of the familiarity people have with AI-generated content, I find it useful to involve it in an introduction. The end result/goal of the project is to produce an original musical audio file every time the program is run. From the user's perspective, this is akin to producing AI-generated content, but "under the hood" this is not done through AI.

How it works: The way this program works is two-fold. In the code, there are two interconnected modules that talk to each other. The first generates the notes. By notes, I mean the raw instructions of what frequencies and their timing in space. You can think of this as the sheet-music writing stage. The program starts by writing the sheet music, but doesn't create the sound. The second module reads off the note data (the sheet music) and produces the sound information. You can think of the first module as the composer who writes the notes in sheet music, and the second module as the orchestra that plays

Anatomy of a song



these notes. The way this is actually implemented is slightly more complicated and interconnected than this modular approach, but in terms of the flow of information, this is an accurate analogy.

The notes module has many subcomponents for all the different instruments. To make sure all the instruments play in tune and on time, there is initially a chord pattern generated. This is done through different algorithms. For instance, one of the algorithms is to simply select a random chord out of a bank of chords, another is to go through the circle of fifths, and so on. Each of these different methods has a probability of occurring. Once the chord pattern is determined, then different instruments create sequences of notes differently. For example, the bass line might just be playing constant notes following the root of the chord pattern, or possibly play octave arpeggios around the root note. The melodic instrument might be playing notes within the arpeggios of the chord patterns or random notes in the scales associated with the chord pattern. Drum patterns must also be created, and different drum hits (bass drum, snare drum, etc.) are added to the sheet music with different probabilities. Hence, for each instrument in the arrangement, the notes are generated. Once this is done, it is time to oscillate these to life.

The second module acts as an interpreter of the generated notes. Not only must it "play" the notes, but it must also play them with an original sound. Therefore, this module must be able to synthesize new sounds every time that the script is run. To do so, the script allows each instrument to randomly select its sounds. For melodic instruments, such as basses, melodies, and chords, an architecture similar to classic (east coast) subtractive synthesis is used. Constraints on the kind of wave shapes, filter parameters, and envelopes are constraints within certain ranges. All of these ranges depend on the instrument. For instance, basses have a lower pitch than other instruments. Furthermore, they often have shorter attacks and releases. The constraints on the randomly generated parameters for each synthesizer are different for every instrument. Things are a bit different for non-melodic instruments, such as the drum. The drum set is made up of different sounds (whose number and type are also randomized within constraints). For example, there are bass drums, snare drums, and hi-hats. Each of these sounds is generated independently of the others, and they all have their own "parameter space" they can exist in. For example, bass drums have been hard-coded to have frequencies low enough so that they sound like typical bass drums.

Once all the audio tracks are synthesized, they are put together and then, what could be called a 3rd module, adds some extra effects to the resulting signal. That is it. The result is a WAV or MP3 file containing a somewhat melodic piece.

Discussion: Over the course of the conception and coding of this project, I have had good discussions with friends and family about what this means and what can such technology be used for. One of the points that arose often was the one

about copyright. Who owns the copyright to this music? Personally, I don't have a strong stance on copyright issues and I don't enjoy focusing on the legalities surrounding art. However, I would compare this program to me throwing a piano off the second floor's window and recording the sound it makes as it lands on the ground. Am I the composer of such a piece? I have no idea. However, I should probably take credit for having brought the instrument to the second floor and throwing a piano off a window, which were most likely not easy tasks.

Having worked on this kind of project and witnessing the evolution of the state of the music industry, I make the following prediction: In the near future, a great proportion of musical content (and possibly of other artistic media) will be generated through AI. I have seen how an amateur like myself can go about creating an infinite number of (granted, crappy) musical pieces. In my opinion, it is extremely likely that, in the coming years, companies like Spotify start producing some machine-made *muzak*. In my case, I was alone, working part-time on the project, did not resort to AI, or particularly sophisticated algorithms, and obtained interesting results that I know very well I can improve by spending a little time more time on the project. Now imagine a multi-billion corporation with thousands of full-time trained engineers. That corporation has a financial incentive to save money by not having to pay artists for their music because they can generate similar content. On top of that, let's face it, the amount of people listening to some random Spotify playlists without even knowing whom they're listening to is astronomical. I don't really know what to think morally of such a forecast because I know that, at the end of the day, people will have the choice to listen or not to this content and I know what mine is.

About me: My name is Christian, and I am a physicist and musician. I am a very nerdy person and when I think of a project I think I can tackle, like the present one, I just have to do it. I try whenever possible to mix up physics and math with music whenever possible and not be gimmicky. I make music under the alias Chris Mauden. Furthermore, I've scored some short film projects as well as released multiple studio albums, which you can find on all streaming platforms. This project started in December 2022 and was pursued within Building 21 during the Winter 2023 semester after graduating from Honors Physics at McGill. I am soon starting a master's degree at the Université de Montréal in May 2023 in Biophysics.

Project GitHub Repo:



Other music:

