AI program writes music and lyrics

May 1 2020, by Peter Grad

Credit: CC0 Public Domain

The Rolling Stones summed up popular music in a simple phrase and title in their Eighties hit "It's Only Rock 'n Roll."

But, as Silicon Valley is finding, teaching machines to create the music and lyrics is is anything but simple.

OpenAI this week announced the creation of an open source system called Jukebox that creates unique melodies and harmonies along with lyrics and vocalizations in the styles of popular artists from a large field of musical genres.

The results are far from perfect, but they are nevertheless impressive. In some instances, JukeBox created credible renditions of songs that resemble the hard-driving guitar-centric style of the Rolling Stones, the golden vocalizations of Ella Fitzgerald, the intonations of Elton John and a near spot-on raspy Rod Stewart.

But there is a notably bizarre concoction of an Elvis Presley-style tune called "Mitosis," where AI Elvis croons, "From dust we came with humble start; From dirt to lipid to cell to heart. With [mitosis] with [meiosis] with time; At last we woke up with a mind." And "Merry Christmas Baby" sung in the style of the great Chuck Berry may well frighten little children.

There remains much work, to put it mildly, to be done. As Devin Coldewey of TechCrunch described it, some of Jukebox's efforts sound like "good, but drunk, karaoke heard through a haze of drugs."

Still, the project is a significant and commendable step forward in the field of AI music generation.

To train computers to generate songs, OpenAI created a massive dataset composed of 1.2 million songs.

"We show that our models can produce songs from highly diverse genres of music like rock, hip-hop, and jazz," OpenAI stated in a report posted on its blog. "They can capture melody, rhythm, long-range composition and timbres for a wide variety of instruments, as well as the styles and voices of singers to be produced with the music."

Computer programs for decades have tackled music generation. One of the more impressive efforts is PG Music's Band in A Box, which lets users input chords and then creates unique melodies and multi-part harmonies in any of dozens of musical styles. In 2018, Google unveiled Musical Transformer, which generate songs with recognizable, catchy repetition. And OpenAI last year released MuseNet, which created original melodies, but did not include lyrics.

Jukebox appears to be the first program that can create songs with lyrics and vocals.

The highly technical specs behind the program are available online. A key challenge is processing massive amounts of musical data. A typical 4-minute song at CD quality (44 kHz, 16-bit) has over 10 million timesteps, according to OpenAI. An auto-encoder was used to compress data, remove irrelevant bits of information and upsample the results back into the program. With that compression, it still takes approximately nine hours to render one minute of audio.

Advances in the program are sure to come. OpenAI acknowledges current limitations of the program: "While Jukebox represents a step forward in musical quality, coherence, length of audio sample, and ability to condition on artist, genre, and lyrics, there is a significant gap between these generations and human-created music."

As an example, Jukebox does not yet compose larger musical themes such as choruses and repeating phrases.

Observers expect the next iteration of the program to be available in April 2021.

OpenAI is a research laboratory based in San Francisco. It says its mission is "to ensure that artificial general intelligence benefits all of

humanity."

More information: openai.com/blog/jukebox/

cdn.openai.com/papers/jukebox.pdf