

HUMAN-MACHINE COLLABORATION VISUALIZES MUSIC

PhD student Andrew Paley blends his two passions—artificial intelligence and music—to add a unique dimension to his compositions and performances.

Andrew Paley is a computer scientist. As a PhD student at Northwestern Engineering, he works to democratize access to information through the power of artificial intelligence (AI) and human-centered design.

Paley is also a musician. As a solo artist and as front man for Vermont-born The Static Age, Paley has released multiple albums and performed at venues worldwide from Japan to Brazil to Europe and North America.

Paley has brought these two distinct interests together in an art-generator platform, Pixie, an AI system that leverages human-computer collaboration to produce unique visualizations for his music.





Bringing Pixie to life

While Paley has devoted the past two years to working on his PhD in the computer lab, he has also remained deeply involved in writing and recording a new solo album, *Scattered Light* (Highwires 2020), in the music studio. These dual pursuits have fostered an interest in how human-machine collaboration can support artistic endeavors.

Inspired by the idea of creating a visual accompaniment to his new music, Paley pitched the concept of Pixie—an AI platform that builds unique, blended images synced to music at the direction of the user—as an exploratory project in Professor Han Liu’s Statistical Machine Learning course. After receiving positive feedback, he went to work building it.

“It felt like an opportunity to take advantage of everything I was working on,” Paley says. “It was separate from my core research, but there was a thread connecting my work in these two different parts of my life.”

In designing Pixie, Paley leveraged BigGAN, a conditional AI model developed by Google DeepMind that creates realistic looking images. He says his work has been influenced by Artbreeder, a collaborative flow model that generates synthetic images users can combine to create something totally new.





Pixie's approach echoes Artbreeder's concept of collaborative exploration, allowing Paley to tag model-generated images he likes for different portions of a song. Unlike other platforms, however, Pixie generates a script based on those tagged keyframes, an analysis of the song, and additional human-set preferences. It then creates animations synced to the music.

The results are hypnotic. One of Paley's videos, for his song "Give Up," is generated entirely with Pixie's architecture. Another, "Stay Safe," experimentally melds Pixie functionality—song analysis, scene collections, and script generation—with aspects of Artbreeder. Both present lush, dreamlike visuals that transition seamlessly from one form to another, all perfectly timed to the music.

"It's a process of collaboration and discovery. I tell Pixie what imagery I like and give it a song file, and then see what it dreams up," Paley says. "You can get very granular with the decisions, or you can let it be more random. Each run of the engine produces a wholly new visualization."

What's next?

With touring sidelined by the COVID-19 pandemic, Paley continues to create more Pixie-produced music videos and expand the platform's functionality, exploring new models and forms of generative animation. He hopes to incorporate Pixie's visuals into his shows once he can resume live performances.

Paley says working at the intersection of his two interests feels natural. "It's been fun to work on something at the border between the two halves of my life," he notes. "It's a reminder that the underpinnings of these two halves are not all that different—it's all a creative process."

ALEX GERAGE

"IT'S A PROCESS OF COLLABORATION AND DISCOVERY. I TELL PIXIE WHAT IMAGERY I LIKE AND GIVE IT A SONG FILE, AND THEN SEE WHAT IT DREAMS UP."

ANDREW PALEY
PhD student, Northwestern Engineering

