

Teach Computer How to Improvise Musically

Preparation

- Read through the activity plan before the session. Here is the original link to the news article on [Project REACH](#)
 - Here are a few things to look for when reading
 - What is the limitation of the current AI models?
 - The end goal of the project
 - How does it help AI improve through improvised interactions
- Send a follow-up email to students to make up for the technical aspects of the articles mentioned in the session
 - In case if students want to get into the technical details of the research
 - Two machine learning models in the first article
 - <https://medium.com/cindicator/music-generation-with-neural-networks-ga-n-of-the-week-b66d01e28200>
 - <https://www.ibm.com/cloud/learn/recurrent-neural-networks>
 - REACH project proposal can be found [here](#)
 - Songs played
 - <https://soundcloud.com/deepjazz-ai/sets>
 - https://www.youtube.com/watch?v=Emidxpkyk6o&ab_channel=Aiva
- Organize [discussion strategies](#) to use to facilitate conversation

Goals and Outcomes: PACE-internal

- Develop student's interests in UCSD-affiliated Artificial Intelligence research
- Introduce possibilities of CS via cross-subject research (music and computer science)

Goals and Outcomes: Student-facing

- I will learn about what creative AI is about and the implication of its success for AI research
- I will get to know the cutting-edge research project at UCSD that aims to understand human interaction and creative behavior through the lens of AI.

Script

This activity is a bit longer, feel free to pick, add, and choose what you think your cohort would be interested in discussing:

1. Play the three songs [through comparison activity](#)
2. Quickly go over the [summary of the Scientific American article](#) as an introduction to today's topic
3. Transition to talking about the newest [project on creative AI in music at UCSD](#)
4. Go over the proposal and summary of the study
5. [Reflection on the materials and discussion](#)

Song comparison

- Let's listen to short clips of two jazz songs, one of which are composed by AI and one made by an actual musician. See if you can tell the difference
 - <https://soundcloud.com/deepjazz-ai/sets>
 - deepjazz On Metheny 128 Epochs
 - Original "And Then I Knew" by Pat Metheny
 - Play each song for 15 - 20 seconds
- Ask students to tell them apart and briefly describe why
 - Sample answer: the AI-generated song lacks the feel of ambiance; feels bland...
- Play one more song. This time the song is not jazz
 - https://www.youtube.com/watch?v=Emidxpkyk6o&ab_channel=Aiva
 - I am AI - AI Composed Music by AIVA
- Ask students whether this is composed by AI or not
 - Answer: no
 - follow-up question: is this sample less obvious than the two we just heard?
 - Maybe music genre is a big factor

Summary of Scientific American Article

Introduction to [the Scientific American article](#) on creative AI

- The current AI research on music generation and art performance
- **How deep the learning has to be?**
 - Two neural networks used to train music AI
 - Recurrent neural networks (RNNs) are popular in today's AI composition because they learn from previous input by looping and thus backpropagate on the fly.
 - Long-term short memory recurrent neural networks (LTSM) come into play because they're able to engage more memory and work over the

course of an entire song, tackling the overall structure, verses, bridges, refrains and so on.

- **The caveat of AI: dealing with elements besides the notes**
 - Regardless of musical genre, the ad hoc, real-time communication that takes place between musicians during the collective improvisation of live jams simply isn't there yet between machines and humans. The acoustics of a room or performance venue that affect sound, the energy of the audience and, of course, the visual cues shared between musicians cannot be accounted for by any current technology.
 - It would require sophisticated audio recognition that allows machines to hear and interpret the other instruments, advanced computer vision to pick up on varied and subtle visual cues, and some way to signal and communicate with the human musicians—all synced up with a real-time improvisatory algorithm.
 - Computational power is achieved but no one has put together an AI equivalent to a live jazz musician
- **What else is needed in terms of training the AI**
 - To recreate an improvisation of a certain musician
 - You'd probably need to create your own data set from scratch: Get new jazz musicians to play each instrument in every possible song from that musician and then train algorithms on those recordings.
- **Why do we need a creative AI that can jam? (what are other applications)**
 - For students learning music
 - it would be useful to someday have an on-demand musical partner to jam with in a seamlessly realistic way any time of the day or night, not just for pure enjoyment, but also for learning
 - After all, how much better could music students hone their talents if they had AI-based teachers that could provide instruction and feedback on their playing anytime?

Introduction to Project REACH

- [A new project](#) to answer the question of whether computers can have the same capacity as musicians to improvise live music.
- Led by Professor Shlomo Dubnov in the departments of Music and Computer Science and Engineering at UCSD and a researcher Gerard Assayag from Paris.
- **Premise**
 - Co-creativity between humans and machines will lead to the emergence of distributed and performative information structures involving the cooperation of artificial and human agents with a significant impact on human practices, knowledge, and cultures. The generative learning of symbolic representations from physical and human signals, and the understanding of the artistic and social strategies of improvised interactions will help us to better understand the dynamics of cooperation (or conflict) inherent in cyber-human networks.

- **Objectives and directions**
 - Understanding, modeling, implementing music generativity and improvised interaction as a general template for symbiotic interaction between humans and digital systems (cyber-human systems)
 - Creating the scientific and technological conditions for mixed reality musical systems, based on the interrelation of creative agents and active control in physical systems.
 - Achieving distributed co-creativity through complex temporal adaptation of creative agents in live cyber-human systems, articulated to field experiment in musical social sciences.
- **Significance**
 - Music is a fertile ground for developing models and tools of creativity that can be generalized to other activities, due to the richness of its structures and constraints that foster the emergence of cooperative behavior and highly integrated courses of action. REACH will study "shared musicality" at the intersection of the physical, human and digital spheres, as an archetype of distributed intelligence, and will produce models and tools to better understand and encourage human creativity in a context where it is increasingly intertwined with computation.
 - Revealing the mechanisms of co-creativity in music will also be a remarkable projector on creativity in general, as music is one of the most highly organized, interactive and complex human activity — at the same time an abstract, sensitive and physical one, yet profoundly shaped by communication and improvisation, and as such a powerful metaphor of human creative interactions.

Discussion and reflection

Facilitate a [discussion around the following question\(s\)](#), pick ones you see fit, feel free to add others:

- The first article makes the example of Jazz and the project, although it does not completely rely on jazz music in the proposal, uses Jazz music as a hook to demonstrate the significance of having a capable AI to resemble the social aspect of creativity.
 - Go back to the beginning of the session when we are listening to the three songs and left off with a question that why the first two songs are easier to tell which is generated by AI [Because we think about improvisation in music I believe jazz is definitely one of the first few answers we'd have in our mind.] Jazz requires more human effort (emotions, interactions, careful listening to other instruments) compared to other songs ==> more human touch to the music essentially. So it is more difficult to mimic. But at the same time, it more comprehensively reflects human creativity. That's why getting AI to improvise Jazz or anything that doesn't possess a lot of patterns will be beneficial to push AI to another level.

- Question to ask: from the two articles we mentioned today, what do you think will be the challenges to this type of research? If you were the researcher on the REACH project, what would you think would be the hardest thing to tackle, and what is missing from the current research goals?
 - One quote from the first article that might help answer “Deep learning in jazz has similarly downplayed the crucial rhythmic, timbral, and textural aspects of music,” says University of California, San Diego, music professor and saxophonist David Borgo

Post-session

- Check in with the student(s) who you noticed might need more support during session (e.g. Looked overwhelmed/lost etc)
- Fill out the [LINK]
 - What activities did you do
 - How did they go
 - Could ask what their favorite part of the session was
 - Any concerns