

Music 356 / CS 470

Music and AI

Stanford University | Winter 2023

T/Th 10:30am-12:20pm

Syllabus (v1.03)

Instructor

Ge Wang

Associate Professor, Center for Computer Research in Acoustics (CCRMA)

Department of Music (also Computer Science, by Courtesy)

Senior Fellow & Associate Director, Stanford Institute for Human-centered AI

ge@ccrma.stanford.edu

<https://ccrma.stanford.edu/~ge/>

Teaching Assistant

Andrew Zhu Aday

Ph.D. Candidate

CCRMA, Department of Music

Department of Computer Science

azaday@ccrma.stanford.edu

Course Website

<https://artful.design/music356/>

Course Overview

How do we make music with artificial intelligence? What does it mean to do so (and is it even a good idea)? How might we design systems that balance machine automation and human interaction? More broadly, how do we want to live with our technologies? Are there—and ought there be—limits to using AI for art? (And what is Art, anyway?) In this “critical making” course, students will learn practical tools and techniques for AI-mediated music creation, engineer software systems incorporating AI, HCI and Music, and critically reflect on the aesthetic and ethical dimensions of technology.

Coursework will span the practical (“how?”), the philosophical (“why?”), and the social (“for whom?”). Students will use these techniques and lenses to think as broadly as possible about the implications of AI in our world, and to design interactive AI tools that keep human judgment in the loop. Through these exercises, we will explore AI might augment, not replace, human creativity.

Prerequisites

1-2+ years of programming; completion of Music 256A/CS 476a or Music 220B is recommended but not necessary; (the most important requisite: willingness to show up with some heart and compassion)

Teaching Philosophy

We firmly believe that anyone can learn anything to which they put their earnest effort and thought. In this course, we also believe the answers to questions are secondary and sometimes even irrelevant. What truly matters here are the *thoughtfulness* of the questions we frame and the *effort* we put into the craft of designing things. Above all, our aim is for each student to acquire for themselves both “things to create with” and “things to think with” as tools that will stay and grow with them for years to come.

Course-specific Learning Goals

As a student, you will be able to (in ranked order of priority):

- 1) understand the various roles AI plays or could potentially play in music creation—and to be able to frame critical questions about Music and AI, considering technical, aesthetic, ethical, and cultural contexts.
- 2) conceptualize and prototype software systems, integrating AI and music.
- 3) demonstrate working knowledge of the tools used in this course, including audio and music synthesis in the ChucK programming language, Python/C++/PyTorch, Wekinator for interactive machine learning and music, and the newly developing ChAI (ChucK for AI).

Textbook

Course readings will consist of a curated list of articles probing into the technical, philosophical, social, and cultural dimensions of AI and music/art. Students will also be tasks to bring readings of interest to share.

Workload Expectations

Students are expected to devote at least five hours of preparation out of class for each 80 minutes of class time, for a minimum of ten hours of preparation per week.

Grading Expectations

Your work will be evaluated on completion, thoughtfulness, effort—in short, the personal integrity with which you carry out your work in the following aspects:

- Weekly reading responses + programming etudes (20%)
- Two significant programming assignments (30%)
- Final project: “design your own interactive AI-based system” (25%)
- Milestone Presentation and general participation (25%)

Class meetings are synchronous; attendance is required (within reason)

Attendance at lectures and milestones are synchronous and required. If a student falls ill or have a personal situation that would affect their ability to attend, the student should contact the instructors before missing lecture or section. Under certain conditions (including illness), a student may be provided an opportunity to make up the work missed in section.

Course Policies

1. Participation in milestone presentations and discussion is a significant part of your grade (25%). Please join class meetings on time, showing evidence of having done the work (readings, assignments, milestones) by sharing your questions, analyses, critiques, and informed comments.
2. As a rule, no late work will be accepted except for family or health accommodations. Please contact the instructors as early as possible in these cases.
3. Please think of this course in lectures and sections as a safe place to be yourself and while respecting others. As the Platinum Rule states: always treat others as they would have you treat them.

The Honor Code

Violating the Honor Code is a serious offense, even when the violation is unintentional. The Honor Code is available at the URL below. Students are responsible for understanding the University rules regarding academic integrity. In brief, conduct prohibited by the Honor Code includes all forms of academic dishonesty, among them copying from another's exam, unpermitted collaboration and representing as one's own work the work of another. If students have any questions about these matters, they should contact their fellow.

<https://studentaffairs.stanford.edu/communitystandards/honorcode>

FERPA: Student Record Privacy Policy

<http://studentaffairs.stanford.edu/registrar/students/ferpa>

Students with Documented Disabilities

Students who may need an academic accommodation based on the impact of a disability must initiate the request with the Office of Accessible Education (OAE). Professional staff will evaluate the request with required documentation, recommend reasonable accommodations, and prepare an Accommodation Letter for faculty dated in the current quarter in which the request is being made. Students should contact the OAE as soon as possible since timely notice is needed to coordinate accommodations. The OAE is located at 563 Salvatierra Walk (phone: 650-723-1066; web: <http://oae.stanford.edu>).

Acknowledgements

This course and its design are supported by Stanford HAI (research and teaching), Computer Science and Music Departments, Symbolic Systems, and CCRMA.

Music and AI: Three “Pillar” Topic Areas

A Practical Survey of AI Techniques for Music

- GOF AI techniques (constraint-based systems; algorithmic composition)
- Machine learning techniques for musical information retrieval (k-NN, SVM, HMM, etc.)
- Survey of deep artificial neural networks and beyond for music and art (RNN, GANs, VAEs, diffusion models, transformers and LLMs, etc.)
- Interactive machine learning: combining AI and HCI

System Design

Architecting software systems for AI and music

Discriminative systems; building tools for analysis

Generative systems; designing creativity support tools

Ethical-aesthetic considerations (What is good? for whom? accessibility and inclusive design)

Interactive machine learning

Philosophical issues of Music and AI

- “How might we best balance machine automation and meaningful human interaction?”
- “Ought there be limits? What should they be?” / “Is this even a good idea?”
- “How do we reconcile the meaning of AI-generated aesthetic output?” / “Is it Art if generated by a machine?” / “What is Art, anyway?”
- “By whom and for whom? Who does this serve and who is left out?”

Assignments and Evaluation

Reading Responses + Programming Etudes (weekly)

- **Written responses** should be 400-500 words in length, and clearly demonstrate understanding of the material. They should favor reflection over synopsis. Feel free to pose well-framed questions in your responses.
- **Programming Etudes** are small exploratory coding exercises designed to experiment with a particular idea or technique, often to produce a brief musical or otherwise creative statement. You are expected to comment thoughtfully and briefly (150-200 words) on your outcome in these etudes.
- **Significant Programming Assignments** (2): these directed programming projects involve observation, reflection, and software design, as well as studio-like milestone peer critique. The general topics of the (two) assignments include
 - 1) designing, training, and evaluating a system for musical genre/sentiment classification
 - 2) a generative music system, using various techniques presented in the course so far.
- **Final Project**: the medium-scale project asks students to conceptualize and prototype an artfully designed interactive AI music software system, integrating the technical, philosophical, and system-building dimensions of the course.