Generative Algorithms for Art and Architecture

A collaborative interdisciplinary course structure nurturing transdisciplinary GenAI-supported design

By Sam Keene and Benjamin Aranda, The Cooper Union

Apollonian Gasket

By Joshua Hsu, Sean Lee, Alexa Jakob, and Alsu Flare

This student project reimagines the common fractal pattern Apollonian gasket in three dimensions.

Editor's introduction:

The rise of generative AI tools presents new possibilities for supporting highly complex work that could help solve humanity's most vexing problems: That is one of the hopes projected onto rapidly evolving AI tools, seen in AI promotional materials for for-profit product development – such as Microsoft AI YouTube video promotions of its generative AI projects – as well as in non-profit ventures – such as United Nations' (UN) International Technology Union (ITU) *AI For Good* global summits which aim to encourage the use of AI tools for supporting the UN's 17 Sustainable Development Goals (SDGs).

One of the necessary steps to realizing such laudable hope for AI is developing imaginative applications of generative AI capabilities for solving complex challenges, or as *AI For Good* goals state: "Identify practical applications of AI, accelerate progress…, scale solutions for global impact" (AI For Good Global Summit).

Supporting transdisciplinary work – which brings together expertise from multiple disciplines to create solutions greater than the sum of their parts

– can be part of the process towards identifying, accelerating, and scaling solutions to complex problems through AI tools.

This contribution by Keene and Aranda to the special issue on *Artificial Intelligence and Possible Futures for the Arts* illustrates concrete steps for creating conditions conducive to teaching transdisciplinary design skills. Keene and Aranda brought together engineering and art/architecture students to transcend disciplinary boundaries and create the diverse projects highlighted here. Their collaborative project-based courses – honed over several years – helped multidisciplinary student teams develop a shared language and integrate diverse perspectives, methods, and insights. The students identified tools and accelerated design and production of transdisciplinary projects through creative and productive collaborations.

Student projects in these courses pushed the limits of AI tools, exposing their boundaries and biases. These projects demonstrate that successful transdisciplinary work can reveal aspects of AI technologies that might otherwise remain hidden. Through their students' work, Keene and Aranda illustrate their belief in the importance of fostering a robust feedback loop between technologists and creatives for productive AI tool development.

Nurturing productive interdisciplinary collaborations between technologists and creatives presents significant challenges, especially in the age of generative AI. Educators must teach disciplinary expertise across multiple areas while creating a supportive environment for interdisciplinary collaboration. As AI technologies advance rapidly, teaching foundational knowledge in computational disciplines becomes even more critical, adding to the challenge of imparting centuries-old expertise in the arts.

Learning design in the era of AI has high stakes for the current generation of students. AI could either augment or replace their creative work, depending on how these technologies are developed and used. An anonymous student from one of Keene's and Aranda's interdisciplinary courses reflects on concerns and possibilities:

> ***"In order to avoid blind reliance to the AI design tools that are currently being developed by a few companies, design practitioners should become interdisciplinary generalists, learning higher level concepts of how the tools that they use work, what they are made of, and eventually understand the boundary of their creative agencies to critically use the tools as augmentations to their design intent."*** - Anonymous student

To help students resist "blind reliance" on AI tools, Keene and Aranda share lessons learned from their engineering/art/architecture education collaboration aimed at dismantling disciplinary expectations. They demonstrate how students from disparate fields can develop a shared language and a mutually supportive collaborative space—conditions essential for nurturing the next generation of creatives as they navigate transdisciplinary design in the age of rapidly developing AI tools.

- Yvonne Houy, Ph.D., Editor, *Tradition-Innovations in Art, Design, and Media Higher Education*

—

Emotions

Student project by Sue Zhou, Nicole Joseph, Anne He

The user sits inside a mirrored chamber, where an AI detects their emotion. These cues are fed to a generative algorithm which creates images illuminating the chamber.

"Our course model intentionally aims to dismantle disciplinary expectations and instead encourages a more collective, group authorship.

While architecture students learn to apply machine learning (ML) and AI to the design process, engineering students learn to critically question the impacts of technologies used in design processes."

- Sam Keene and Benjamin Aranda

"At first, working along disciplines meant understanding a whole new set of workflows and ambitions for a project. While at the start, most of the designers were often more interested in the conceptual and aesthetic implications of the work while a lot of engineers were tasked with making the models work.

The most interesting collaborations took place when architects and designers interfered with and began to understand the workflow of the engineers, and the engineers added extremely valuable insights on the conceptual implications of the projects to designers who tend to be seduced by the imagery that is produced through generative learning.

It took some time to understand how to create a true collaboration between the two disciplines without delegating them with the usual tasks they would perform in their respective fields."

- Anonymous student comment on challenges in collaborative work

Cellular Automata

Student project by Daniel Matallana, Leslie-Fairuz Abad-Neagu, Gavri Kepets

This video uses 3-dimensional cellular automata and generative images to envision the growth and decay of buildings.

Cellular Automata video / YouTube link

"Being in an interdisciplinary class was a bit of cacophony in the beginning, with each student trying to stubbornly argue each other's discipline's design process with no mature understanding of what their own discipline is supposed to be.

However, as internal and external identity crisis continued, I started to have a better understanding of what architecture design education could offer - with coming up with design thesis question that eventually had to be solved through technical challenges, which team members with engineering background had much better process for.

It was continuous back and forth conversations about why we are doing something and how we could realize it that at the end offered me to see what the world outside of school could be."

- Anonymous student comment

"We aim to emulate a creatives-technologists feedback loop in our classes, by putting creatives and technologists in direct conversation.

We support this feedback loop by ensuring our students have the tools to communicate with each other.

We build a shared language by discussing creative works – and exploring the technology that enabled them.

We then guide our students through the process of developing collaborative work."

- Sam Keene and Benjamin Aranda

Through the Looking Glass AI Narrator

Student project by Azhin Nam, Leon Fang, Ravindra Bisram

This is an attempt to develop an AI narrator, annotating one's day.

The students quickly realized that image captioning methods only recognize objects (nouns), and not verbs. As they tried to use video captions as inputs to a text generation method, things did not go as expected. The video begins with a mundane errand of buying a toothbrush, then pushes the limits of comprehension by the AI narrator - by including scenes the AI narrator had difficulty comprehending - and the narration goes completely off track.

Through the Looking Glass AI Narrator video / YouTube Link

"Similar to how the recent AI image filling function of photoshop has showed [sic] the creative industry a glimpse of how AI tools could optimize the workflow to reduce tedious manual labor (traditionally designated to entry level jobs), it is fully expected that AI tools will eventually be integrated into day to day workflow of a design practitioner, buying more time for a practitioner to invest in higher level design concepts." - Anonymous student advice for the next generation

"We found that students often defy expectations defined by their discipline. For instance, engineers developed highly artistic creative concepts, while artists and architects suggested solutions to technical challenges."

- Sam Keene and Benjamin Aranda

—

Textures

Student project by Juan Cardona

This piece used generative algorithms in a feedback look to create new textures, and then realized these textures as objects in alien worlds.

"With the advent of chatGPT and Dall-E there has been a lot of talk about how AI in the design field is a threat to the discipline but art has gone through this phase many times in its history as new technology such as the printing press and photography were developed.

If we can think of generative models as yet another tool in our production of work, and not just a quick fix to automate image/text or sound generation very exciting things are possible...

Another question which has emerged in the field is one of datasets - at the moment most of the open source models use stolen work to generate their final products. Most datasets as they exist now are also a reflection of human bias - often the white man's bias...

So as we continue to develop design practices using these methods it is important to think of how we can be mindful of the datasets we are tapping into- can we create our own? How do we decolonize and reduce the existing biases that exist on the internet?"

- Anonymous student advice for the next generation

Replacement

Student project by Aiden Shankman, Garret Ingman, and Samuel Finley

This piece classifies the users age, gender, race, and emotion. It then replaces the user's face with a generated version of themselves as the AI sees them. It can be deeply unsettling to see yourself replaced, particularly when the biases of AI are revealed in the process.

"The goal of our project was to show non-AI savvy individuals how biased commonly used AI models can be. Along with this, we wanted to demonstrate just how easy and fast these AI models can be used and abused by someone who has just a small amount of coding experience.

The first step we wanted to nail down was the feeling of being analyzed and judged wrongly by an AI model...

The whole process was done in less than 400 lines of code and all the software cost less than 10 dollars."

- Anonymous student comment

"Our work shows that creative applications of AI are not only important new tools for creativity but that creative interdisciplinary applications of AI are essential for exploring and understanding the broader societal impact of these new methods.

Creative applications of AI reveal much about the technology that would be unobserved otherwise."

- Sam Keene and Benjamin Aranda

"Often the individuals would be misrepresented and reveal bias in the data used for the software. Many of the female minorities that used the project would be misgendered or have drastically different ages, while Caucasian and Asian males would often have fairly accurate outcomes.

The project would represent the creepy truths and stereotypes about Ai [sic] and facial recognition made over the years.

Seeing what our computer supposedly sees of us from this project should create some sense of unease especially considering that this might be how we are represented by online companies or the government.

The exhibition design is a depiction of this broken reality."

- Anonymous student comment

"We have found that generative AI tools have been developed under very constrained conditions and assumptions.

Pushing tools away from their pre-determined constraints reveals much about their biases, weaknesses, vulnerabilities, and threats."

- Sam Keene and Benjamin Aranda

—

Intelligent Archaeometric Reconstruction

Student project by Kevin Chow, Maren Speyer, Taesha Aurora, Theo Jaquenoud

This project shows how multiple techniques can be combined for a larger work. A language model was trained using only myths and stories from indigenous, non-Western sources. Landscapes were then generated from these myths.

Long before the world was created, there was an island floating in the sky upon which the Sky People lived. They lived quietly and happily. No one ever died or was born or experienced sadness. Their sky was full of stars, every planet a rainbow.

One moment they came of age, the next, they had to survive far beneath the sky.

They slept anywhere they wanted. It was as if the sun never really shone and everything in town looked like ice and snow.

The only light that shone came from the houses and the trees. They were a lot more peaceful than we are today.

All the families were well known among the Sky People and they seemed to get along well. One particular couple was very close to their beloved brother and friend, Miao.

Together, they went to the mountains, which had remained unexplored. At first they didn't know where they had come from anymore, so they decided to settle within the stone. Miao grew sick in the coming days and his brother decided to travel back to the village to get medicine, but when he came back, Miao had begun to turn to stone and become one with the mountain. His brother tried to carry him home, but the stone was too heavy. His brother declared "no, I can not carry you home, but the mountain is yours now and your name will always be there".

> ***"A key insight we have observed when working with AI methods is that weaknesses, biases, and flaws are revealed at the boundaries...***
>
> ***In the Intelligent Archaeometric Reconstruction, much care had to be taken to fine-tune a model with non-Western texts."***
>
> ***- Sam Keene and Benjamin Aranda***

Tu lived in the clouds, in a round hut made of rainbow arches.

*This place was a place filled with the natural beauty and natural culture of *place*. It was one where men and women of every age could dream as well as to be a part of.*

They had all found the most beautiful scenery, the most amazing plants the most fabulous trees and the most amazing flowers, and now, they were finally living at this place called Kannav.

Even if they had come here alone, had they not come here in droves, they would probably have lived at this very same place.

—

They spent several months in the hut. The water slowly flowed, and it ran at good and slow speed into a little stream. It never felt like they were standing straight out of a hole, that they could go and fall between trees.

Everything was fine and quiet, however, not all was still. After breakfast, we would go and watch the stream move. It was so beautiful and wild. The river rose with its own white rays. The air was bright and green, though we never saw the sun coming, only the clouds. It was like standing on the edge of a mountain, you never knew what was on top.

For a long time I sat still, my heart pounding but always holding the light. We came back here one day to visit a friend. He was a farmer. He asked, "Why do you want to see this river?" Before I could answer, "I cannot," my friend said to me, "it is impossible for the earth to support a body of water from above." The next morning the river started to go, and so I followed it. In this way I gained my perspective and got into the habit of being the river-dweller. No one gave me a more careful observation.

> ***"This world building project involved creating fictional landscapes and fictional myths relating to them as a commentary on how geography and culture are inextricably linked and how the myths we form also form us.***
>
> ***Part of the ambition of this project was to challenge the typical vertical perception of machine learning work, where the format of production follows the guidelines of harvesting a dataset, tweaking a model and producing images/text or using an open source application.***
>
> ***With this work we were interested in utilizing machine learning processes as a step or tool in a larger process - considering the capabilities of machine learning not just for singular productions but more so for immersive generation."***
>
> ***- Anonymous student comment about this project***

LightNest Inhabitable Screen

Student project by Rhyna Chen, Michael Sluchevsky, River Friedman, Allister Liu, Victor Zhang

The video shows a 3-dimensional "inhabitable" screen, with generated images projected upon it.

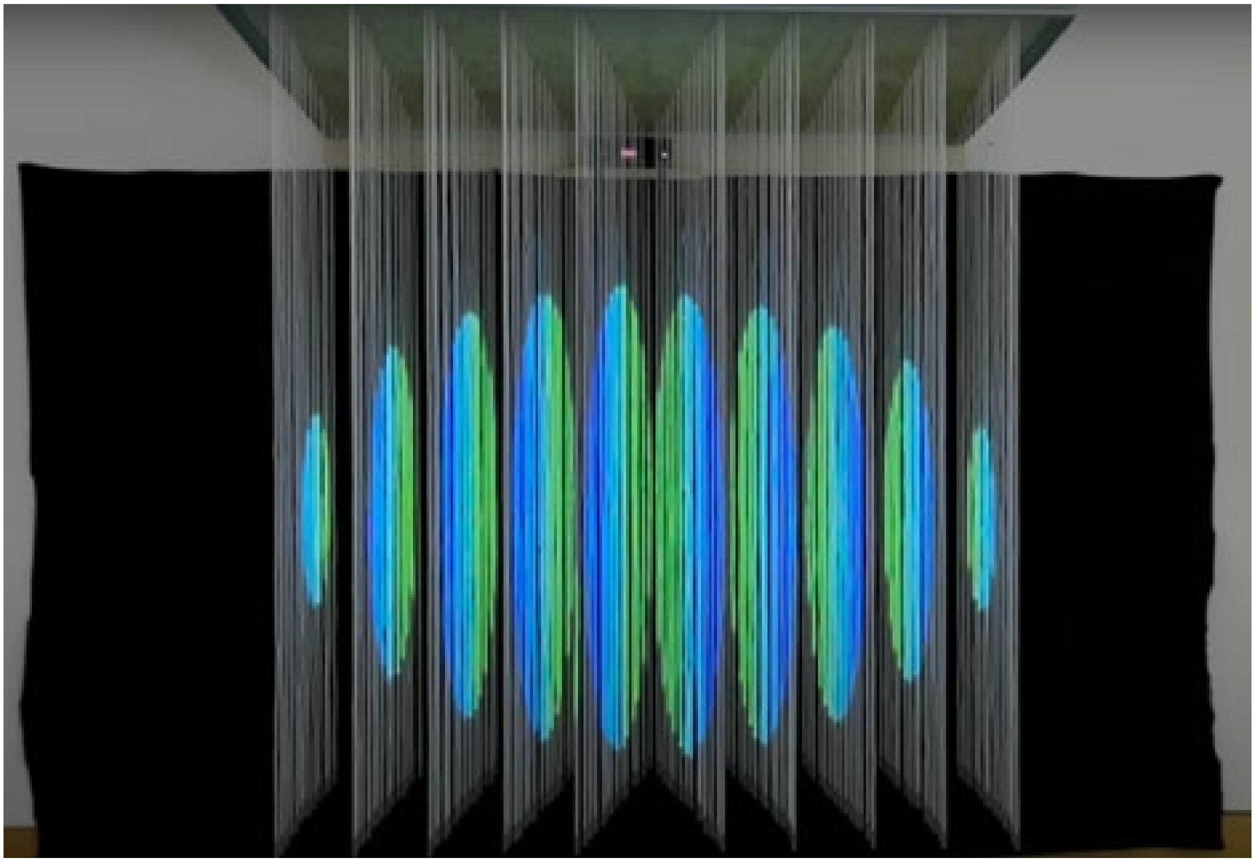

LightNest Inhabitable Screen video / YouTube link

> *"The main challenges were comprehending what the other majors were capable of and what was considered impossible to do in a certain time frame.*
>
> *There would be artistic/design aspects that an engineer thought could be done quickly but actually would take a while. On the other side, an artist or architect would have technical ideas they thought would be possible but would be practically undoable to do within the budget or timeframe we had." - Anonymous student comment about collaboration*

"... we believe next-generation creative practitioners need to engage with technologists to understand how new tools work.

Creatives and developers need a shared language, with tools enabling a feedback loop for improving tools."

- Sam Keene and Benjamin Aranda

—

Generative Algorithms for Art and Architecture

Notes on a collaborative interdisciplinary course structure for nurturing transdisciplinary AI-supported design

By Sam Keene and Benjamin Aranda, The Cooper Union

Introduction

Generative art is a rapidly growing field. The potential is enormous, and risks abound. While creatives can now interact with generative AI tools such as Chat-GPT (OpenAI, 2022), Dall-E (OpenAI, n.d.) and Stable Diffusion (Mostaque, 2023) and a growing host of others without much technical training, we believe next generation creative practitioners need to engage with technologists to understand how new tools work.

Creatives and developers need a shared language, with tools enabling a feedback loop for improving tools. This feedback loop has been demonstrated across many fields such as photography, animation, and music (Hertzmann, 2018). Examples of this abound, starting perhaps most obviously with the field of photography, which was enabled by scientific advances. To be a practicing photographer in the early days of the field, it became necessary to tinker with the scientific processes involved. This pattern continues in filmmaking and animation, where creatives and technologists must collaborate on the design of lenses, cameras, audio, and special effects. 2-D and 3-D modeling has enabled studios like Pixar to completely change the concept of an animated film. Pixar has a well-known collaborative model between creatives and technologists (Catmull, 2019). In music, Jimi Hendrix was a voracious consumer of new guitar effects and worked directly with an electrical engineer to create the fuzz pedals which define Jimi's guitar tone (Scott, 2022).

We aim to emulate a creatives-technologists feedback loop in our classes, by putting creatives and technologists in direct conversation. We support this feedback loop by ensuring our students have the tools to communicate with each other. We build a shared language by discussing creative works – and exploring the technology that enabled them. We then guide our students through the process of developing collaborative work. We do not restrict our discussion to only generative AI, but we consider generative AI to be the most current version of the field of generative art.

Creative applications of generative artificial intelligence can help expose weakness and bias in the underlying algorithms of AI tools. An important recent example of this is the work of Joy Buolamwini and Timnit Gebru (Joy Buolamwini, 2018). This project started out not asan investigation into machine learning bias, but a creative application of machine learning.

While trying to create an art piece to project masks onto faces, the authors noticed that the algorithms worked much better on white faces than black faces. The resulting investigation showed that facial recognition algorithms have significant performance differences based on the race of the face being recognized. This work, which started as a class project in the MIT Media Lab's Science Fabrication Course, was a critical moment in demonstrating the significant bias present in facial recognition algorithms, a perfect example of how a creative application of AI can reveal weaknesses in technology.

In this article we describe a successful course model that trains creative and technical students to collaborate and discuss technical, design, and social issues in interdisciplinary teams, preparing them to be future leaders in this field. Our institution, like many others, offers computer science courses in generative machine learning for engineers, while a different part of campus might offer creative applications of machine learning for artists and designers. What makes our course unusual is that we engage students from many disciplines in a collaborative way.

Our course model intentionally aims to dismantle disciplinary expectations and encourages a more collective, group authorship. While architecture students learn to apply machine learning (ML) and AI to the design process, engineering students learn to critically question the impacts of technologies used in design processes.

Course Structure

In this section, we present a general model for the development of an interdisciplinary course, including steps necessary prior to the semester, and an outline for a 15-week delivery.

We believe that a collaborative model is essential for an interdisciplinary course, and this requires that two (or possibly more) instructors are involved in development and delivery. We found that the inclusion of multiple instructors ensures that the contributions from all disciplines will be given equal weight. Additionally, the presence of an instructor from a student's discipline is an important signal that they will be welcome, and their perspective valued. Finally, the instructors themselves model the kinds of collaboration between disciplines that the students are expected to participate in during the course. In our delivery of this course, there were two instructors, one from the School of Architecture, and one from the School of Engineering.

In our interdisciplinary courses the students ranged from 3^{rd} year to 5^{th} year, with the 5^{th} year students being architecture thesis students or Master's degree level engineering students. We placed enrollment caps to ensure a rough balance of between 12-15 art/architecture students and 12-15 engineering students. Separate course codes for each group ensure that students receive credit applicable to their degree.

The engineering students are expected to have strong coding skills and some familiarity with machine learning. The art and architecture students are expected to have a corresponding level of sophistication in their creative fields that comes with having completed at least two years of foundational instruction.

We recommend the instructors' work begin 6-9 months before the start of the semester, to select readings, design projects and rubrics, solicit guest speakers and any other related tasks.

Broadly speaking, the course has three distinct phases, which vary in length, and may have some overlap.

Phase 1: Developing a shared language

The goal of the first phase is developing a shared language across the various disciplines present in the course. This may consist of lectures from the faculty members on the relevant aspects of their discipline, selected readings and class discussions, or group-based activities.

We found that at this stage, common terms in one discipline mean very different things in another. For example, a "model" to an engineer might consist of mathematical equations, but to an artist or architect, a "model" could be a physical or digital sketch. More generally, however, across all fields, a "model" can be understood as an abstraction of the real world that we use as a tool. It is crucial to ensure that students can communicate effectively with each other throughout the course, and therefore time must be devoted to creating a shared language.

Phase 2: Short projects to develop teamwork and collaboration skills

The second phase of the semester consists of short, introductory projects to continue to build upon the shared language, and to develop teamwork/collaboration skills. Each project typically lasts 2-3 weeks. A group discussion is held to help understand the difficulties of interdisciplinary collaboration.

Example prompts for introductory exercises:

> *As a group, select a recent generative artwork, examine the technical aspects required and critique the work.*
>
> *Use machine learning methods to transform a piece of media (images, sounds, videos, objects, etc.) from one form to another. This can be the same medium (i.e., transform one image into a different image or try and make a text into a different text) or going from one medium to another (ie. transform a text into an audio work). The methods, models and data are left to the students.*
>
> *Create a generative work that results in a physical medium.*

A grading rubric is provided in Appendix A at the end of this document.

At the conclusion of the second phase is a full-class critique of the projects. The experience of a formal critique, while common for architecture and art students, is often quite new for the engineering students. The process of critique, where choices are explained and defended – and indeed even the word “critique” for this process – are further examples of the development of shared vocabulary between disciplines.

Phase 3: A large-scale project

The final phase of the semester consists of a long-term, large-scale project. By this point, the students should have been introduced to the basic concepts of the course through lectures and assigned readings. They have also completed short, introductory projects and have received feedback on those projects.

The instructors sort the students into groups based on common interests, while taking care to ensure that all teams have representation from multiple disciplines. The groups, under the guidance of the instructors, will define and implement a project to be carried out for the remainder of the semester. There may be additional lectures, guest speakers and critiques as the semester progresses. However, the instructors gradually step back until they are serving primarily in an advisory role, allowing the students to direct their own projects.

An example final project prompt:

> *Create an installation from a generative concept, either one completely new or elaborated from previous exercises. An installation is understood as a physical manifestation of a work in a space. It can be a physical object or series of objects, composed of projections or monitors, or a site-specific intervention. While the medium may vary, the installation is conceived at the scale of a room or part of a room and frames the work in an intentional way.*

For final deliverables, and assessment purposes, we have different expectations for the artists/architects than for the engineers. The art and architecture students are typically expected to document their work as a contribution to a professional portfolio. The engineers are expected to submit code (typically in the form of a repository such as Github), as well as documentation detailing the technical aspects of the work.

However, there are no set roles for any student in any of the projects. Our expectation is that students will bring their respective skill sets and learn new techniques as their project requires. While it is common that the engineers tend to do more coding, and the artists/architects do more design work, by the end of the semester, we have seen many students branch out beyond their usual roles. We have seen several cases of artists/architects who started to code, and engineering students learn architectural design tools such as Rhino and Grasshopper and partake in fabrication. We encourage the students to teach and learn from each other. No specific tools, methods or approaches are dictated by the instructors. Over the course of the semester, the final deliverables and methods used to achieve them are discussed and agreed upon between the student groups and the faculty.At the beginning of this article, we presented several examples of student responses to the final prompt. These works were the result of approximately half a semester. Each student group proposed and conceptualized the projects. For brevity, we showed only the results of the projects and not intermediate outputs. As the final prompt required the work to stand alone, we presented these with only minimal descriptions.

The faculty meets weekly with each group to review and critique the emerging projects. The technological challenges and potential solutions are discussed, with all students present. In addition to feedback meetings with the faculty, there are also pinup sessions and invited guests for additional critiques.

Sample outline for a 15-week semester

- Week 1
    - ◇ First class session: Class topics are introduced by faculty, and an icebreaker in-class exercise introduces students to each other.
    - ◇ First assignment: Random groups are formed for a warm-up assignment to develop the shared vocabulary necessary for interdisciplinary work.
- Weeks 2-3
    - ◇ Lectures continue to develop a shared vocabulary.
    - ◇ Faculty provide critiques of the first assignment.
- Weeks 4-6
    - ◇ Additional exercises are assigned to create a shared vocabulary and expand technical and design expertise.
    - ◇ Current literature is assigned and discussed in class.
- Midterm
    - ◇ Interdisciplinary groups are finalized.
    - ◇ Project proposals are presented and critiqued by the class, faculty, and invited guests.
- Weeks 8-15
    - ◇ Faculty conduct focused meetings with each group, sometimes with invited guest evaluators.
- Final
    - ◇ An exhibition showcases student projects.
    - ◇ Students give short presentations introducing their projects.
    - ◇ Faculty and invited guests provide critiques.

Student Perspectives

We solicited perspectives from students via an anonymous survey, asking them to share their narratives, discuss the challenges they faced, and offer some guidance for what the current and future generations should learn and do in response to AI tools. Some quotes from the survey were included in the section on student projects.

Faculty Observations and Insights

In this section, we discuss some of the challenges we faced as instructors, as well as some insights.

The primary challenges were adjusting to the different expectations of students from varied backgrounds. For example, engineering students are not used to the concept of having their work critiqued, where it is judged subjectively by a group of people. It was also challenging to get engineering students to change their mindset around AI and machine learning (ML). In particular, the students needed to learn to think more critically about applications of AI/ML, rather than simply designing a system for optimal performance. For many engineering students, this course was the first time they were asked to engage critically with technology. Conversely, art and architecture students may be used to more free form work, and might not like the notion of "deliverables" even if they are self-defined. Care must be taken to set expectations among all students and reinforce them over the course of the semester.

An additional challenge we noticed was that in some cases it can be difficult to tell which students are contributing, particularly in larger groups. Occasionally, individual students would lose interest, perhaps because they were unsure how to contribute. We felt that in groups sized larger than 4 or 5 students this problem was more likely to occur. We recommend that if there is significant student interest in one topic area to split the larger interest group into multiple project groups.

The scaffolding of early assignments served the students well in terms of building a shared language, and in developing their confidence to approach a larger project as the semester progressed.

Once larger projects started, shifting to a group meeting-based approach also served the students well, as the instructors were able to continuously monitor progress, and make suggestions for improvement on a regular basis. The regular meetings also worked well to identify any group dynamic problems, so that they could be rectified quickly.

We observed every week how the group members communicated with each other. We would witness art students talking about statistics, engineering students discussing graphic design, etc. During group meetings, when we observe the students discussing with each other, it becomes clear if the collaboration is successful. When this occurred throughout the semester, the projects were coherent and impactful. When group meetings were disjointed, such as individual students discussing their own parts, rather than a cohesive presentation, it became clear they are not communicating well. When this persisted, the quality of the projects frequently suffered.

One significant challenge we faced was institutional support. While our institution was willing to financially support multiple faculty members teaching one course, there were an array of bureaucratic challenges. For example, we needed to use numerous course codes to ensure students received proper credit that counted towards their respective degrees. This requires a lot of complicated work to understand multiple curricular goals, and possibly the approval of several curriculum committees.

Recruiting students was also a challenge. In some years, we had an imbalance with more engineering students than art/architecture, or vice-versa. This can somewhat affect the balance and equity of the course and make some students feel unwelcome. We found that it is important for faculty members to personally recruit students for the courses. This can take several forms, such as reaching out personally to students, giving presentations at student meetings, and making other faculty/advisers aware of the existence of new courses.

On pushing GenAI tools

The process of "provoking the tool" described by Fisher (2024) in response to Vermillion's (2023) AI-Human hybrid creative process resonates with our work:

> "Vermillion discusses how, with each engagement, each provocation with the GenAI, his language shifts as he seeks to realize a creative vision. Only by provoking the tool in different ways can the 'most interesting results […] subvert our preconceptions, forcing us to look at our own ideas in very different ways, and extending our imaginations.'" (Fisher, 2024)

We have found that generative AI tools have been developed under very constrained conditions and assumptions. Pushing tools away from their pre-determined constraints reveals much about their biases, weaknesses, vulnerabilities, and threats. It is from these challenges we pose to the AI systems that we often see the most creative work.

On GenAI prompt engineering

We recognize the value of learning to use a tool intentionally, and the method proposed by Fisher (2024) is a valid way to achieve that intentionality:

> "Choosing, accepting, and rejecting GenAI outputs from different prompts as an AI provocateur can help students understand the underlying computational models generating the content. When the same patterns frequently occur—whether in content, structure, or form—the way the generative systems of AI work becomes more transparent. Understanding these patterns allows the creative student to use them intentionally in their processes." (Fisher, 2024)

However, we do not typically approach generative AI through prompt engineering: We discourage our students from focusing their work on engaging with prompts. In the context of our course, interaction only via the prompt API as created by AI developers would not lead the kind of transdisciplinary collaboration we seek to foster.

On autonomy, authorship, and critical agency in relation to AI

We discuss different kinds of generative work through the notion of autonomy: To what degree is there an autonomous system at work and how does it affect a claim to authorship and agency by the author?

As instructors, we think all of the relational models between humans and GenAI referenced in Joshua Fisher's (2024) article - actor, material, and provocateur (Steinfeld, 2021) - are valid. While we have not used this explicit framework, we do present examples of work to our students that would be compatible with this framing.

Earlier versions of generative artists such as Sol Lewitt, who developed rule-based instructions for carrying out installations (Sol LeWitt, 2022), might fit into the Actor model. There are many current artists whose work might fit more into the material model. For example, Mario Klingemann has established himself as an expert at using generative AI models to expertly craft images in a style that is unique (Klingemann, M. , 2020). Finally, as an example of the provocateur model, one might consider an artist such as Tom White, whose work Perception Engine used "neural networks to create abstract representations from collections of real-world objects"(White, 2018).

We strive not to foster a particular approach, but to ensure our students are thinking critically about the choices they make in the work they produce. Our job as instructors is to provide students with the tools and knowledge to create their own works. The work they produce is extraordinarily varied, and we have had work that could be categorized into each of these categories: actor, material, and provocateur (Steinfeld, 2021).

Conclusion

In these interdisciplinary courses we found that students often defied expectations defined by their discipline. For instance, engineers developed highly artistic creative concepts, while artists and architects suggested solutions to technical challenges. We felt that all students learned about techniques outside their discipline, and how to translate their mode of thinking and problem-solving to others with different backgrounds. The evidence is in the projects.

The guest faculty in the critique process were very impressed with the student projects in general. We felt many of the projects were the best examples of combining different disciplines that we've seen in our experience as teachers. These interdisciplinary classes produced creative work that is unusual for being both interdisciplinary and critical of the technology used. When we've observed projects in purely disciplinary courses where creatives have used AI/ML without engineers, there often seems to be a tenuous grasp of technical concepts, and the work is unable to make deeper connections about the implications of data, algorithms and bias. Conversely, when engineering students have attempted artistic work without creatives, the result often feels facile, shallow, or superficial, due to the inexperience of making artistic choices and being critical about them. The faculty all felt that the quality of work produced in these consciously interdisciplinary courses, be it computer code, data analysis, infographics, visualizations or installations, was significantly better than work done in purely disciplinary courses.

As we met with all students weekly, their curiosity about other disciplines was obvious to us. We believe this led to significantly stronger student engagement in the material. Finally, since the semester culminated in a critique or exhibition, there was an obvious sense of accomplishment and satisfaction among the students. As they were able to share their work with a larger community, they gained a realization that their work could have a broader impact than they previously imagined.

A key insight we have observed when working with AI methods is that weaknesses, biases and flaws are revealed at the boundaries. We see this in the student projects. For example, the Replacement application seemed to work much better for white men, and often would misgender or otherwise inaccurately characterize non-white faces. In the *Intelligent Archaeometric Reconstruction*, much care had to be taken to fine-tune a model with non-Western texts. The *AI Narrator* video shows the limitations of using object-identifying methods to make sense of the world around us.

Our work shows that creative applications of AI are not only important new tools for creativity but that creative interdisciplinary applications of AI are essential for exploring and understanding the broader societal impact of these new methods.

Creative applications of AI reveal much about the technology that would be unobserved otherwise.

References

Catmull, E. (2019, February 25). How pixar fosters collective creativity. Harvard Business Review. https://hbr.org/2008/09/how-pixar-fosters-collective-creativity

Fisher, J. A. (2023). Teaching Creatives to be A.I. Provocateurs: Establishing a Digital Humanist Approach for Generative A.I. in the Classroom. Tradition Innovations in Arts, Design, and Media Higher Education, 1(1), Article 5. DOI: https://doi.org/10.9741/2996-4873.1002

Hertzmann, A. (2018). Can Computers Create Art? Arts. https://www.mdpi.com/2076-0752/7/2/18

Joy Buolamwini, T. G. (2018). Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification. Proceedings of the 1st Conference on Fairness, Accountability and Transparency. Proceedings of Machine Learning Research. https://proceedings.mlr.press/v81/buolamwini18a.html

Mostaque, E. (2023, November 8). Stable diffusion public release. Stability AI. https://stability.ai/news/stable-diffusion-public-release

Klingemann, M. (2020, August 29). Quasimondo. https://quasimondo.com/

LeWitt, S. (2022, February 3). Sol LeWitt: A wall drawing retrospective: Mass moca. MASS MoCA | Massachusetts Museum of Contemporary Art. https://massmoca.org/sol-lewitt/

OpenAI. (2022, November 30). Introducing ChatGPT. OpenAI. https://openai.com/blog/chatgpt

OpenAI. (n.d.). Dall·E 2. DALL·E 2. https://openai.com/dall-e-2

Scott, J. (2022, April 5). The technology of Jimi Hendrix. The JHS Show. https://thejhsshow.com/articles/the-technology-of-jimi-

Steinfeld, K. (2021). Significant others: Machine learning as actor, material, and provocateur in art and design. The Routledge Companion to Artificial Intelligence in Architecture (pp. 3–12). Routledge.

Vermillion, J. (2023). Giving Up Control: Hybrid AI-Augmented Workflows for Image-Making. Tradition Innovations in Arts, Design, and Media Higher Education. 1(1), Article 3. https://doi.org/10.9741/2996-4873.1009

White, T. (2018b, September 3). Perception engines. Medium. https://medium.com/artists-and-machine-intelligence/perception-engines-8a46bc598d57

Appendix 1

Grading Rubric

	Integrated 3	Developing 2	Emerging 1
Concept	Deliberate and well-developed concept. Materials/media selected thoughtfully and intentionally, creating formal coherence in the work.	Concept present but not clearly developed or deliberate. Materials/ media selected seem arbitrary or sloppy.	No clear concept present. We have no idea why you made these choices and you can't explain this to us.
Methodology and Documentation	Technical processes used are coherently documented (in written statements and code comments). Elegant solutions employed. Someone with little programming background could be given a really basic README for this and basically understand how it works or even run the code themselves.	This mostly makes sense, but feels a little slapped-together and could be simplified. Would require an in-person walkthrough to really make sense.	Sloppy code, process incoherent, we have no idea how the work you presented emerged from this pile of spaghetti you've sent us.
Collaboration	Collaborators demonstrate a shared ownership of the resulting work based on documented statements and their ability to articulate their process. All collaborators can speak fluidly about the technical approach and the conceptual underpinnings.	It doesn't seem like there was a lot of meaningful dialogue between collaborators; work feels disjointed or lopsided more heavily toward one person's technical or creative interests. In talking about the work it feels very "now I'm talking about the math", "now I'm talking about the art."	Only one of you seems to be saying anything about the work, or lots of talking over each other. It's not clear that you worked together at all.