My inspiration for my artwork is Machine Hallucinations – Space by Refik Anadol. I'm studying a lot of machine learning and Anadol's work using machine learning hallucinations , which was an example in my machine learning class, seemed very interesting to me. Using an LLM like ChatGPT or Gemini to create AI art is one thing, being able to recreate that from scratch is a different challenge. I planned to recreate the project using visual data from NASA and have an algorithm create its own visualization of space. Ideally, the algorithm will generate a small set of images after training on a reasonably large dataset. Anadol's project is on a larger scale and his collection features images that shift and swirl around. He had access to larger and better data through his partnership with NASA and much better GPUs to train his model, so I decided something simpler and on a smaller scale what he did for my project. Growing up, I spent a lot of time figuring out how computers worked, teaching myself how to code around middle school. During high school, I decided I wanted to work with machine learning, since it was growing quickly around that time. This interest stayed with me into college, and I'm currently majoring in Data Science. I've taken a few machine learning classes, but have only done projects involving predicting future data. While this project is for this class, I also want to see what I can do with machine learning. I thought this would be an interesting intersection of two fields usually considered to be unrelated.

I originally planned to create and train a StyleGAN2-ADA[1] model from scratch on a dataset sourced from NASA using Python. Since I couldn't find a quick way to download a large enough dataset from NASA images, I used an SDSS dataset of around 100,000 images (linked) from Kaggle and resized all images to 256x256 pixels. I tried training it locally on my PC using my GPU but ran into many technical issues. I couldn't import specific required libraries, leaving many parts of my model incompatible with each other. I also encountered version issues with PyTorch, the deep learning library I was using, and other further conflicts stemming from constantly switching versions of Python to meet the requirements of the libraries. Though I was able to generate images after a few hours, these images were nowhere close to what I needed for a machine hallucination. They were mostly just a mix of psychedelic neon colors (see project notes). This was also putting a large strain on my GPU and computer in general, causing it to be extremely slow outside of the training model and constantly freeze. I kept trying the StyleGAN2-ADA model but couldn't get it working so I started looking at a different method that could generate the images under the idea of the project.

The next model I tried using was a Stable Diffusion model[2] on Google Colab[3], which provided me access to Google's GPUs and TPUs, which were much stronger than the one on my computer. One problem I ran into here was my dataset. Since Colab is cloud-based, I needed to upload the dataset, which was way too large and my GPU access would run out before I could even start building the model. Because of this, I decided on using a much smaller dataset of 31 images since Stable Diffusion already had access to hundreds of millions of images that included

[1] StyleGAN2-ADA is a generative model type used for image generation
[2] Stable Diffusion is another image generation model but focuses more on using text prompts to generate images
[3] Cloud-based code editor for Python by Google

space images. This allowed me to give it a small dataset to point it in the right direction along with a prompt when generating images. I first tested with text-to-image hallucinations, which is just your regular AI image generation through an LLM using only a prompt. These 8 images were colorful, swirling, and kaleidoscopic, reminding me of the wormhole scene in 2001: A Space Odyssey. I then tried image-to-image hallucination using my uploaded dataset, allowing the model to blend real images with its own abstract creations. It created 3 images for each of the 30 images I uploaded. Since most of the images in the dataset were of a single cosmic entity, these generated images were not as colorful and chaotic as the ones from the text-to-image generation, but were instead closer to the original images. These text-to-image and image-to-image creations are the ones I'll be submitting.

Overall, I learned a lot about how machine learning works in the context of image generation. I also gained a better understanding of generative AI and how it actually works because of this. My machine learning class only spent a lecture on how media was generated, and it wasn't very in-depth for any form of media. This, in combination with another project I worked on this semester involving sound design, was a great learning experience and opportunity to explore machine learning on my own.