

Homework 7: Illustrations of AI

Due April 8th at 10pm

1 Stable Diffusion

In the past year, text-to-image generation models have improved dramatically, and gained a lot of attention, both positive and negative. For instance, [this New York Times article](#) discusses the controversy surrounding an AI-generated piece of art that won an art competition at a state fair.

This week, we will be running **our own art competition**. However, all of the entries will be AI-generated. Your task is to design a series of illustrations for the short story “When Robot and Crow Saved East St. Louis” by Annalee Newitz.

First, you will need to get used to using your art generation tool: Stable Diffusion.

Stable Diffusion is a state-of-the-art image generation model: given a natural language description, the model tries to generate matching images. These models cannot be run on a laptop, but luckily, Hugging Face has a free web version: [Hugging Face’s Stable Diffusion demo](#).

You can also use [DreamStudio](#), which is a faster model, but it requires you to create a free account to use. **You can use either model for your assignment, but please note which you are using.**

1.1 CONTENT WARNING

The dataset that Stable Diffusion was trained on contains offensive, disturbing, and pornographic images. In my experiments with Stable Diffusion, I have not generated anything concerning; however, the model is capable of (re)producing upsetting images. Please take care!

You can read more about these issues in [Birhane, Prabhu, & Kahembwe \(2021\)](#). Note that the paper also comes with a content warning, as it discusses many disturbing examples.

1.2 Generating images with Stable Diffusion

To generate images, you will need to enter a **prompt** into the textbox. There are many tricks to writing prompts that generate different kinds of output. Take a look at the [DALL-E 2 prompt book](#)¹ for some ideas about how to prompt text-to-image models.

Stable Diffusion allows both *positive* prompts, which describe what the image should be like, and *negative* prompts, which describe what it should *not* be like. You can read more about how the negative prompt option works here: [Negative Prompts](#). Positive prompts are more common, but you’re also welcome to use negative prompts.

Generate some images with different prompts. Make some observations about what works.

¹DALL-E 2 is another text-to-image generation model; however, it has a monthly query quota system.

2 *When Robot and Crow Saved East St. Louis*

Your goal is to produce a series of illustrations for the short story “**When Robot and Crow Saved East St. Louis**” by Annalee Newitz. **Start by reading the short story, which is part of Slate’s Future Tense series.**

2.1 Reading response questions

The robot in “When Robot and Crow Saved East St. Louis” is of course more advanced than any current technology. But science fiction has always played a role in inspiring future technology. **Reflect on the way that technology is presented in the story, and respond to the following questions:**

Question 1

Robot is able to work together with a number of different agents to prevent a serious outbreak of disease. Analyze Robot’s actions in the narrative, and identify the skills that Robot uses to overcome the various problems it faces.

Question 2

Over the course of the story, Robot builds social ties to a number of characters. Do you think current AI technology is on a trajectory to succeed in these kinds of settings? What kinds of lessons do you draw from this story about social skills and AI?

2.2 Shane’s Response

“When Robot and Crow Saved East St. Louis” was published as part of Slate’s Future Tense series, which brings together sci-fi writers and technologists. **Read the response piece** written by Janelle Shane, and then **answer the following questions.**

Question 3

Shane’s response was written three years ago. Much has changed in AI since then. Identify some places in Shane’s response that seem out-of-date to you. Do recent AI advances lead you to disagree with Shane’s broader views on the story, or do you think they reinforce her original points?

Question 4

In Week 1, we read Jordan (2019), which argued for an emphasis on Intelligent Infrastructure (II) over human-imitative AI. Discuss the technology presented in *Robot and Crow* in relation to the distinction between II and human-imitative AI. Do you think the story illustrates the importance of one over the other?

3 Illustrations

You will submit a set of 10 illustrations for the short story. Your images should work together to tell the story, just like illustrations in a picture book. Your images should cover different events and characters in the story.

Here are the categories for our competition:

1. Best illustration of a character
2. Best illustration of an event in the story
3. Best representation of technology
4. Best representation of the story's atmosphere
5. Best photo-realistic illustration
6. Best illustration in a particular aesthetic
7. Best foreground illustration
8. Best background illustration
9. Most creative prompting technique
10. Most surprising prompt

You must submit 10 images, but you need not nominate them for a particular category.

Each image must be submitted along with the prompt that generated it. You should also say which model you used.

4 Ethics of Image Generation Models

Text-to-image generation models have sparked many creative applications and broadened the accessibility of digital illustration. However, they have also raised a number of ethical issues.

Read [this Ars Technica piece](#) on ethical issues surrounding Stable Diffusion. Discuss your opinions about the ethical implications of Stable Diffusion and other text-to-image generation models. Talk about each of the following issues:

- Deepfake generation
- Artist imitation / copyright issues
- Reproduction of cultural biases