NEWS@UNITED_AI

APRIL 2023

THIS EDITION

ART with AI

"HEART" - AI MERCH

GPT-4
DID YOU KNOW...?

We're guessing you've heard of GPT-4, right? It's the latest and greatest AI tool that can give you some seriously impressive answers. Even if you're not a tech whiz, it's worth giving it a shot. To get started, either use it directly with Bing or create an account on OpenAI. Our favourite way to use it? Instead of just chatting, try uploading a picture and talking about it with GPT. It's a fun and unique way to test out this groundbreaking technology!

Learn more here.

Can artificial intelligence produce art?

This question led to United AI's first completed student project: AI Merch. In this article, we take a look at how the AI Merch student team built a model that was able to generate impressive pieces of artwork.

The impetus behind this project was to leverage AI-produced art in club merchandising. For example, these could be club-branded items such as mugs, fashion apparel, and posters with interesting designs. The AI Merch team predicted that AI-generated designs would be a huge selling point for buyers. As an added bonus, licensing fees and commission costs could be avoided, as no human artists were involved in creating the art. The club can then use the saved money to finance further projects.

As the project was carried out alongside the team's regular studies at the DIT, every member could only devote a limited amount of his or her time to the project. Thus the decision was made early on to restrict the generated images to a certain theme: AI itself. But what exactly does that mean, artistically? Artificial intelligence is commonly associated with robots and the human brain. Most of the images that appear when one searches for AI have either one or both of these characteristics. Thus it is important that people can identify certain patterns in the generated images that could be related back to AI.

"DEMENTOR" - AI MERCH

The images were created with the help of state-of-the-art Generative Adversarial Networks (GANs). This model architecture requires 2 neural networks trained in parallel. The "Generator" NN creates the images, while the "Discriminator" NN selects the images. The finer details are outside the scope of this article, but feel free to refer to this excellent **TensorFlow tutorial** for more information.

Any good machine learning model requires good input data. In this particular case, the data needed were images. These were collected programmatically from Google Images using a web scraper written in Python. The web scraper utilises the Google Custom Search API, which allows search keywords to be defined. Based on these keywords, the API returns a number of relevant images which the web scraper then downloads.

After collecting the images, they needed to be preprocessed. The downloaded images had a few quirks: first – all of them had different formats and resolutions. Second – there were a ton of duplicate images. These were removed using an image comparison algorithm. The images were then resized to a predefined resolution suitable for later training and converted to JPEGs. The last step was to filter out irrelevant images, as some of the images we scraped did not fit a certain criteria. Remember, the resulting images needed to evoke the essence of AI; thus this should also apply for the input images.

As manually filtering the videos was a monotonous task, the team developed a simple Tinder-like web-app that helped make the filtering more engaging. It works as such: a picture is shown, and the user needs to decide if the picture is suitable for inclusion in the training set. By swiping right on the picture, it is marked for inclusion. Swiping left excludes the picture.

The next step was training the neural networks. There was a lot of hyperparameter tuning involved, and thus the team ran the model around 5 times in total, each model having around 200,000 epochs. The whole process took around 2 weeks.

Training the models wasn't always smooth sailing. The Python kernel crashed frequently during training, forcing frequent restarts to the process. Another limitation was the GPU. When the resolution of the output images is increased, the training time increases. With the GPU used, it would take an unreasonable amount of time. Thus the solution here was to reduce the output resolution. This meant that the resulting images were blurry. This was alleviated by running the images through an image upscaling algorithm. The final result was what the team actually expected at the outset of the project.

A selection of the generated pictures are featured in this newsletter issue. Thank you to the AI Merch team, who devoted their free time and effort to work on the project. We are also grateful to the Deep Learning Lab at the DIT for letting us use their rooms and borrow the necessary computing resources needed for the AI Merch project.

"ASTRONAUT" - AI MERCH

PROJECT STACK

HARDWARE

CPU
8x 3.6 GHz

GPU
NVIDIA RTX A5000

RAM
128 GB

SOFTWARE

LANGUAGE
PYTHON

LIBRARIES
TENSORFLOW
GOOGLE-IMAGE-SEARCH