



Mission AI

notebook

Mission AI - Introduction

In the academic year 2023-2024, we were asked to research generative AI for the Research Group Material Practices at Willem de Kooning Academy, Rotterdam University of Applied Sciences, the Netherlands.

For this project, we used a research-by-making method. We kept a notebook of the information we gathered during our research. The notebook reads like a journal, providing an overview of how we have developed possible approaches to integrating AI in our education program.

Now we would like to share our notebook with you, our dear students, and colleagues.

Han & Brigit

Han Hoogerbrugge is an artist, illustrator, and animator whose work spans diverse media. He is regarded as one of the pioneers of internet art and began his online artistic career with a series of interactive web animations. His work has been exhibited at: Stedelijk Museum Amsterdam, Boymans van Beuningen Rotterdam and Centraal Museum Utrecht in the Netherlands, Museo Tamayo (Mexico City), MOCA (Tapei), the Design Museum London, Sala Parpello (Valencia) and at the 54th Venice Biennale at the Danish Pavilion. Hoogerbrugge made illustrations for The New York Times, various Dutch newspapers, and Playboy magazine. Hoogerbrugge is a teacher at the Willem de Kooning Academy where he focuses on concept and idea development with a special interest in mixing traditional and modern technologies.

Brigit Lichtenegger is a media artist and software engineer, interested in the relationships and interactions between science, technology, and art. She is currently employed by the Willem de Kooning Academy, where she teaches and conducts research at its Interaction Station, a meeting place that offers programs and facilities for people who like to experiment with new tools, technologies, and hardware, taking a critical eye to developing technologies and seeing how they can be appropriated for creative uses. Since 2015 Brigit has developed a variety of educational programs for the arts around the topic of AI, ranging from lectures about the history and current developments in the field, to creating chatbots and generating images.

Mission AI – Introduction

How much of a professional illustrator's (or any other field's) practice can be outsourced to AI while maintaining a personal signature and how can AI contribute on a creative level?

A few years ago, most articles about AI talked about the future, and how it would change everything. Today we can see the changes happening. AI is being introduced to the public on many levels. People don't need illustration skills to produce (technical) high-quality illustrations using AI. But does that make them illustrators? What distinguishes them from professionally trained illustrators?

AI raises many ethical, philosophical, technical, and legal questions. These critical issues should be included in art education because they directly affect the future of our students. As AI becomes an integral tool in the daily practices of artists and designers, students should be encouraged to contemplate how they can use this tool in their future work.

For *Mission AI* however, our aim was to explore the creative potential of these technologies beyond workflow optimization, automation, and addressing copyright concerns. Our focus was on discovering what unique possibilities AI offers beyond those found in other media forms. What kind of art projects will they enable, and how can these innovations be implemented in an educational program?

To navigate our research, we adopted the point of view of illustration students. Using a research-by-making methodology our objective was to create a "children's book" using AI technologies.

This approach has led us to believe that there is a significant potential for AI technologies in the field of interactive storytelling. By combining techniques such as generative custom machine learning models for style, images, 3D models, text, voice, and music, with object recognition, natural language processing, and sensor input, new real-time adaptive and personalized interfaces are becoming possible. These advances offer new opportunities for artists and their audiences to engage and interact in fresh, innovative ways.

Therefore our research was not about the individual technologies, but rather about their connections and effects. The “children's book” served as an example, but the fundamentals can be used for a variety of art and design projects that will be relevant not only to illustration, but to all visual design disciplines.

Mission AI – Contents

Quick Google Search	5
Brainstorm 1	6
Brainstorm 2	7
Technical	8
Children’s Book Idea	9
Demo	10
Technical Questions	12
Style Consistency Experiment: Embeddings	12
Style Consistency Experiment: Hypernetworks	14
Style Consistency Experiment: LoRA’s	15
Training Han’s Style LoRA	17
Training Results	18
Technical Conclusions	19
Additional Control	20
Presentation Animation Students	21
Elective	22
One Day Workshop	26
Conclusion	30
Resources	31
Works cited	32
Afterword	33
Colophon	35

Mission AI – Quick Google Search



Are any professional illustrators already working with AI?

We choose to do research through making. We aim to make our research valuable to all practices, and we take the practice of illustration as our starting point.

Using generative AI, people no longer need illustration skills to produce (technical) high-quality illustrations. But does that make them illustrators? What distinguishes them from professionally trained illustrators?

The first thing that pops up when doing a quick Google search for AI illustrators is Fiverr.

Fiverr is an online marketplace that connects freelancers with people or companies looking to get a job done.

Fiverr has an AI section with subcategories like ChatGPT, Midjourney Artists, DALL-E Artists, and Stable Diffusion Artists. There are 10,619 services available in the Midjourney section. It is unclear how professional the artists are. Most of the pictures shown have an ultra-clean, polished look. Prices are around \$20 for one image.


Fiverr shows a lot of technically impressive illustrations, but it is unclear if these illustrations are being used for commissions. It is unclear if the featured artists maintain a professional practice with AI-generated illustrations.

An A.I.-Generated Picture Won an Art Prize. Artists Aren't Happy, article from the New York Times: www.nytimes.com/2022/09/02/technology/ai-artificial-intelligence-artists.html

This is one of many articles about Mr. Allen winning a painting contest in 2022 with an AI-generated image. Mr. Allen claims he was curious how the new generation of AI image generators would compare to human artists. (Roose)

June 2022. Allegedly the first AI magazine cover for Cosmopolitan, a collaboration between members of the Artificial Intelligence company OpenAI, and digital artist Karen X. Cheng. (Liu)

Cristina Beller works in traditional illustration and techniques and used the AI image generator Midjourney for this work. Her 'classic' Instagram account has 3500 followers, her Midjourney Instagram account has 7000 followers.



Beller, Christina [[@asunderbazaar](https://www.instagram.com/asunderbazaar)]. [Illustration], *instagram*, 30 Aug. 2023, https://www.instagram.com/asunderbazaar/p/CwjekATAqHC/?img_index=1



Beller, Christina [[@asundera](https://www.instagram.com/asundera)]. [Illustration], *instagram*, 24 Dec. 2023, https://www.instagram.com/asunderbazaar/p/CwjekATAqHC/?img_index=1



A quick Google search for 'illustrators working with AI' did not return any professional illustrators with a practice based solely on AI for commissioned work. Some illustrators, such as Cristina Beller, see potential in AI, but they are still figuring out how to use it.

Mission AI – Brainstorm 1



How will we use generative AI, and for what kind of project?

We decided to do research by making. We started talking about a project in related to illustration practice. Illustration practice is very diverse, from advertising, editorial, tattoos, animation, and murals to package design and social projects. Between illustration students, making a children's book is always a popular subject, so we decided to take that as a starting point.


First idea: we don't have children ourselves, let's create a story about us getting children using AI.

We asked ChatGPT to describe what a possible child could look like, based on the description of the father, and to create a prompt we could use for AI image generators.

"Create a punk illustration of Pim, the 5-year-old boy showing his full body. His gray eyes should have a mischievous glint, and his blonde hair should be tousled. He's wearing a ripped punk t-shirt, green shorts, and mismatched socks and black boots. Capture his energetic and playful nature in an expressive ugly style."

We used several 'free' online AI image generators that returned a wide variety of interpretations of the same prompt.

It is an interesting process to see how differently the AI image generators process the prompt. It might be useful to get a quick start in designing a character, but it mostly felt unsatisfying, because why would an illustrator let AI take over the fun part of illustration?



Deepai.org (Sept. 2023)




Openart.ai (Sept. 2023)



dreamstudio.ai (Sept. 2023)



Picsart.com (Sept. 2023)



clipdrop.co (Sept. 2023)



stablediffusionweb.com (Sept. 2023)

Mission AI – Brainstorm 2




Can we make our children’s book interactive using of Ai?

Once we decided that we didn’t want to use AI to completely take over the visual aspects of the children's book, we thought about adding an interactive element. Since AI can process data very quickly, using the reader’s input to interact with the story might add something interesting.



Are there examples of classic children’s books that use interactivity?

To find out, we did a Google search and found the TIME magazine list of [the 100 best children’s books of all time](#).



Among the 100 best children’s books, we found “Press here” by Hervé Tullet. The book contains images of colored dots and readers are invited to press the dots and see what happens on the next page.

We ordered the book and much to our surprise we found that, although its interactivity was technically “fake”, it still felt like you were interacting with the book.



This led us to the idea of thinking about interaction without changing the story.

Mission AI – Technical



Which fields of AI are interesting for adding interactivity to our children's book project?

To answer this question, we have taken an inventory of the ways in which AI can be applied.

Image Segmentation

Image Segmentation is the field of image processing that deals with separating the image into multiple subgroups or regions (such as a set of pixels, also known as image segments) that represent distinctive objects or their subparts.

Learning to see: Gloomy Sunday by Memo Akten (2017)

Learning to See: Gloomy Sunday is a video and an interactive installation in which the live camera image of a table covered with objects is analyzed by a series of neural networks trained on different data sets (including: oceans, fire, clouds and flowers).

Image to Image

Image-to-image translation allows images to be converted from one form to another while retaining essential features. This allows operations such as style transfer and colorization.

Neuralangelo

This AI model, built by NVIDIA Research in 2023, transforms 2D video into lifelike 3D replicas, capturing the fine details of complex materials such as roof shingles, panes of glass, and smooth marble. (Li)



Image Segmentation




Image to Image

Mission AI – Children’s Book Idea

We came up with a setup for an interactive story involving characters in a linear story. As the creators of the book, we are responsible for the story, the design, and the style of the story. Readers of the book can create their own characters, in the style of the illustrator, using generative AI.

Help the illustrator!



This is Han.
He is an illustrator.
He is working on a story but cannot decide what
the main character should look like!
Can you help him?

Let's make a face.
You need two eyes, two ears, one nose, and one mouth.

(technical instructions)

The AI will be used to generate characters based on simple instructions. As the story progresses, the reader will be asked to adjust or create new characters.

Let's go for a walk



You have given the character a face, great!
Our story will begin with a walk in a forest
so we need to put the character's face on a body.

Let's make a body.
You need two legs, two arms, and a torso.

(technical instructions)


Mission AI – Demo



We created a demo to show how the interaction could work.

In our demo, the audience would be asked to find objects and make a configuration to create a face. This configuration is then used to generate a child's face in a style fitting the book.

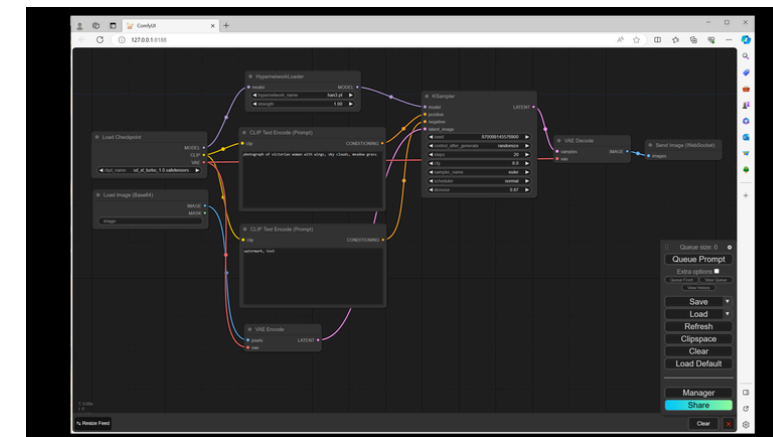
We used a live camera feed, but the same mechanism could easily be used with drawings, and it could be an installation or an app for example.




Stills from the demo. On the left the camera input, and on the right the generated image.



The demo was made with [ComfyUI](#) and [TouchDesigner](#) and used [TDComfyUI](#) to have an interface to ComfyUI in TouchDesigner. The image from the camera is the input for [image2image](#) in combination with a prompt to generate the output image.



ComfyUI is a node-based modular GUI and back-end for [Stable Diffusion](#)




TouchDesigner is a node-based visual programming language for real-time interactive multimedia content.

Stable Diffusion is a generative AI model that produces images from text and image prompts. It can also be used to create video and animation.

With image2image, the position and color of the supplied image is used in combination with the text prompt to generate a new image.

Mission AI – Demo



The [demo video](#) can be seen on YouTube.

Mission AI – Technical Questions



Which software or platform to use?

There are many platforms and software applications that allow you to generate images. They differ in the quality of the images they generate, user-friendliness, how close the results are to the prompts, whether it's a cloud service running on a remote company server or an application you can run on your own device, and in other features.



We decided to work with Stable Diffusion from Stability AI. Stable Diffusion comes with an open RAIL License (Responsible AI License), which is very similar to an Open Source license, but with some restrictions on use. This provides transparency in how the software works, and flexibility because we and others can modify and extend it.

It can be installed on our own computers and servers, giving us more control and a safer environment for our students, because their images will only be on our school computers and not be uploaded to company servers. There are also a lot of additional tools and extensions that can be used in conjunction with Stable Diffusion.

We mostly used a locally installed web interface developed by AUTOMATIC1111. These advantages outweigh the disadvantages of a somewhat clunky user interface and image quality that is perhaps a bit behind commercial competitors. Stability, flexibility and safety first.



Can we achieve style consistency?

We don't want our character to look different on every page. Ideally, we will be able to teach our AI about style, characters, and other concepts that are important to the story, while still retaining all the common knowledge that has already been entered into the system.

Mission Ai – Style Consistency Experiment: Embeddings



In a first attempt, we tried to train an “embedding”. With an embedding, you can add keywords to a base model, and it makes it possible to add objects or styles with just a small number of images, without modifying the base model. Can we train a “man in a suit” character in the style of Han Hoogerbrugge, without having to teach it the concept of a beach?


For training, we used the Stable Diffusion v1.5 model as the base model. It has been trained on over two billion images from the open LAION dataset. These images come mostly from user-generated content platforms like Pinterest and WordPress, but also shopping sites and stock image sites. (Bajo)

It is possible to use other models that are more specialized, for example in anime, or photorealistic models, but the StableDiffusion model is quite good for general purposes.



Results of the prompt “A man in a suit at the beach”, using only Stable Diffusion's base model.

Mission AI – Style Consistency Experiment: Embeddings



We used 18 frames from an animation made by Han to train the “man in a suit” character. A caption was generated for each frame.



bald, facial hair, formal, greyscale, long sleeves, male focus, monochrome, mustache, old, old man, pants, shirt, simple background, standing, suit, white background



images generated during the training process



images generated with the embedding, using the prompt: “Man in a suit at the beach”

Training up to 10,000 runs of the dataset through the machine learning algorithms – i.e., in AI lingo: completing 10,000 “epochs” – took about 3.5 hours, on a PC with a AMD Ryzen 9 3900X 12 Core CPU and Nvidia GeForce RTX 2080 T1 GPU running Stable Diffusion on Linux. The results were interesting. Though they were not consistent or did not look much like Han’s illustrations, the system had clearly learned the pose, and the suit. Also, most of the pictures looked more like illustrations than photos.

Mission AI – Style Consistency Experiment: Hypernetworks



Next, we tried to train a hypernetwork. A hypernetwork is a small neural network that can be used to modify the style of the base model.

For training, we again used the Stable Diffusion v1.5 model as the base model.



We used only 11 images for training. The images are basically the same as those used for the embedding, but cropped and scaled to 512 x 512 for more variety.



blurry, blurry background, blurry foreground, chain, depth of field, greyscale, long hair, monochrome, motion blur, shirt, simple background, white background

Again, captions were generated for each image. Because of the cropping and scaling, there was a lot of blurring.



Images generated during the training process




Image generated using the hypernetwork with prompt: "Man in a suit at the beach".

Mission AI – Style Consistency Experiment: LoRA's



After experimenting with embeddings and hypernetworks, we tried to train a LoRA. LoRA (Low-Rank Adaptation) is a training technique for fine-tuning Large Language Models (LLMs) and Stable Diffusion models without the need for full model training. They are fast to train, don't require large amounts of data, and are small in size. Multiple LoRAs can be used in the same prompt, and they can be combined with different compatible base AI models for Stable Diffusion (which are known as "checkpoints").



More and better quality images used in our first attempt to train a LoRA

We used [Kohya's GUI](#) to generate the captions and train our LoRA. Again, we used the Stable Diffusion v1.5 model as our base model ("checkpoint"). The training process took only twenty minutes.



images generated during the training process



image generated after training, with the prompt: "Man in a suit at the beach"

We were then able to use our self-made model in the Stable Diffusion GUI. The results were interesting, but also puzzling. To get a better handle on the results, we ran a series of tests using the same prompt with different settings. We found that we obtained the best results with a handful of what in Stable Diffusion are called "sampling methods", and that it was important to find a good balance between Stable Diffusion's so-called CFG scale and the weight of our self-made LoRA, as explained on the next page.

Mission AI – Style Consistency Experiment: LoRA’s



These images were generated using the prompt: “A man at the beach”. From top to bottom, the four most promising sampling methods, and from left to right, different values of the CFG Scale. In the top image the weight of the LoRA is 1.0, and in the bottom image it is 0.7.

The weight of the LoRA determines how close the result should be to the LoRA, and the CFG scale determines how true the generated image stays to the prompt.

Mission AI – Training Han's Style LoRA



After many experiments, our understanding of how to train Machine Learning Models improved significantly.

Brigit selected 15 images from Han's collection of 405 illustrations he had originally created for the Dutch newspaper *Volkskrant*. The aim was to train a LoRA with Han's style. The images were chosen to be a combination of black and white and color, to use large color blobs, and to have no text.

The Kohya software program was used to generate the captions and train the LoRA. Captions were generated with the program Blip, and preceded with "h.hoogerbrugge". This word can be used in a prompt to trigger his style. For the class of images, "cartoon" was chosen. This is a generic description for the type of image to be created via a prompt.

The actual training was set to 26 repetitions and 10 epochs (learning cycles) and took about 30 minutes.



The generated caption for this image was:
"h.hoogerbrugge, a cartoon tree with a clown face on its face and hands in the air, with a green background"

After every epoch sample images were created for each of the following prompts:

h.hoogerbrugge, a cartoon of a man in a forest

h.hoogerbrugge, a cartoon of a tree with the Eiffel tower as it's stem

h.hoogerbrugge, a cartoon of keanu reeves at the beach

h.hoogerbrugge, a cartoon of a scene from pulp fiction



Mission AI – Training Results



Mission AI – Technical Conclusions

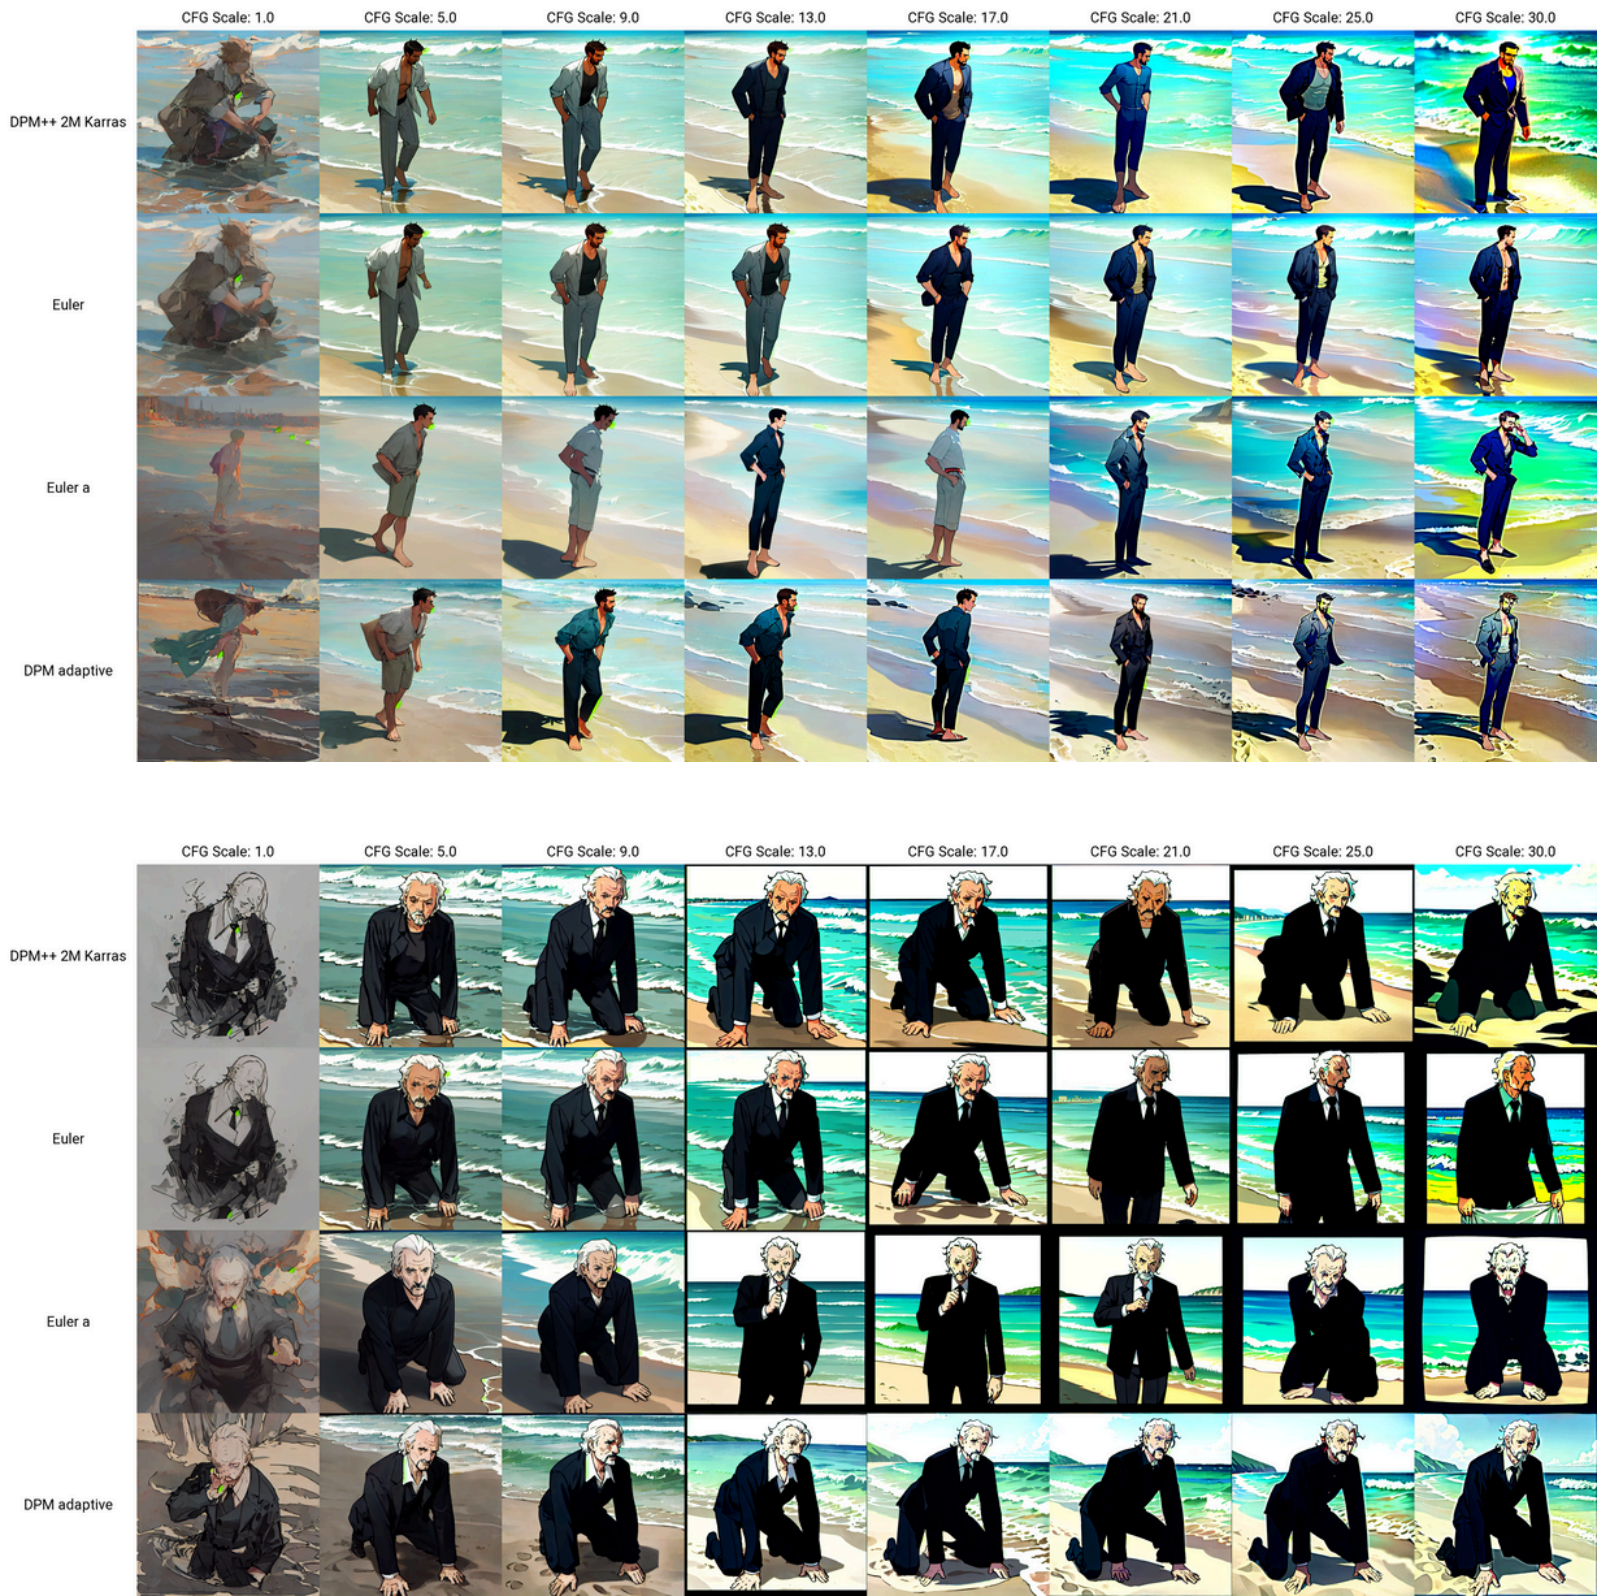


There are many ways to train models. Embeddings, hypernetworks, LoRA's, and dreambooth for example. These can also be trained on different large base models. The results also depend very much on the data used for training. This includes the images, as well as the captions. There are many many settings. Once a model is created, the results also depend on the prompt and again many more settings.

After trying different methods and learning more about all the options, LoRA's seemed to be the way to go. They are fast to train, don't require huge amounts of data, and are small in size. This makes training machine learning models achievable for our students, who typically don't have a lot of data. It is something that can be done in the time frame of a workshop. The possibility to combine multiple LoRAs in the same prompt, and to train and combine them with different base models, is very useful to meet the individual needs of our students.



These images are generated using our first LoRA in combination with the epiCRealism checkpoint model. Prompt: "A portrait of a man".



These images are generated using the Counterfeit checkpoint model (a popular Stable Diffusion base model for generating anime images). Both images have the same prompt: "A man on the beach", but the bottom image uses our first LoRA with a weight of 0.7.

Mission AI – Additional Control



How to get more control over the generated images?

In addition to having control over style, characters, and objects, we want control over composition, colors, perhaps some details, and more. Not everything can be achieved by writing a prompt. Techniques such as image-to-image, inpainting, segmentation maps, and more allow artists to have significant control over the image that is generated.




The images on the right were generated using the Scribble extension for ControlNet. ControlNet is a neural network structure that can augment large diffusion models like Stable Diffusion with conditional inputs such as edge maps, segmentation maps, key points, etc.

The first image is given as the input with the prompt: “A drunken man shouting”. The second image used the epiCRealism checkpoint model, the third image GhostMix, and the last image Stable Diffusion’s base model v1.5.



We ended up spending much more time experimenting with training machine learning models and tools like Scribble and image-to-image than we thought we would. This was partly because it took more time to really get a handle on things, but mostly because our colleagues and students expressed a lot of interest in these techniques when we presented our research as it progressed. We felt encouraged that the combination of LoRAs and ControlNet tools would enhance the sense of creative authorship in generating images.



Mission AI – Presentation Animation Students

We presented our AI research to second-year animation students.
During the presentation, we asked the students to draw three cartoons.

At the beginning, we asked the students to draw themselves with a speech bubble telling what they think of AI.



Halfway through, we asked the students to draw us with a speech bubble telling something they remembered we had said.



At the end, we asked the students to draw themselves with a speech bubble telling what they think of AI after our presentation.



Students were critical of AI. Even if they are interested in AI, they find it hard to think of ways to use it in their practice. After the presentation, some of the students were more positive about the creative possibilities.

The cartoons also told us that we needed to make our presentation more focused, since most of the students only remembered trivial information.

Mission AI – Elective

We organized an elective course that focused on exploring ways to use AI to create original visual work. Rather than using AI as a tool to replace human creativity, we wanted to explore ways to use AI as a collaborator in the creative process.



Giuseppe Arcimboldo 1526-1593



Students were critical of AI, but chose the elective because they felt it was unavoidable, so they might as well learn more about it. After the elective course, they felt more positive that it could be useful and fun for them.

At the beginning of the elective course, we gave a presentation about collections of visual elements that, when put together, form a new visual. For the course, we asked the students to use this as a starting point for experimenting with AI. The students were taught how to train a LoRA and how to use the Stable Diffusion graphical user interface in combination with ControlNet. We also gave a lecture on the history of AI, and other creative techniques coming from its field, to provide more context.



Nicolas de Larmessin
1640-1725



Charles Allan Gilbert
1873 – 1929)



Matthäus Merian the Elder 1625

We asked about ownership. Did they feel that the generated work was theirs? Some students definitely did, while some did not. Looking at all the different student presentations we could clearly see authorship.


Mission AI – Elective

Artúr Voznickij, first year Bachelor Graphic Design student
Grotesque Self-Portrait

Collection – Selfies pulling faces



Experiments – Trained LoRA and ControlNet
Scribble



Final result




Mission AI – Elective

Justė Dovainytė, 1st year Bachelor Animation student
Portrait trained on her own style

Collection – Her own artwork



Experiments – Style LoRA




Final result




Mission AI – Elective

Felix Döringer, second year Bachelor Transformation Design student
A sexual creature mixing the digital realm with the physical realm from the past

Collection – photos from a 90's porn magazine



Experiments

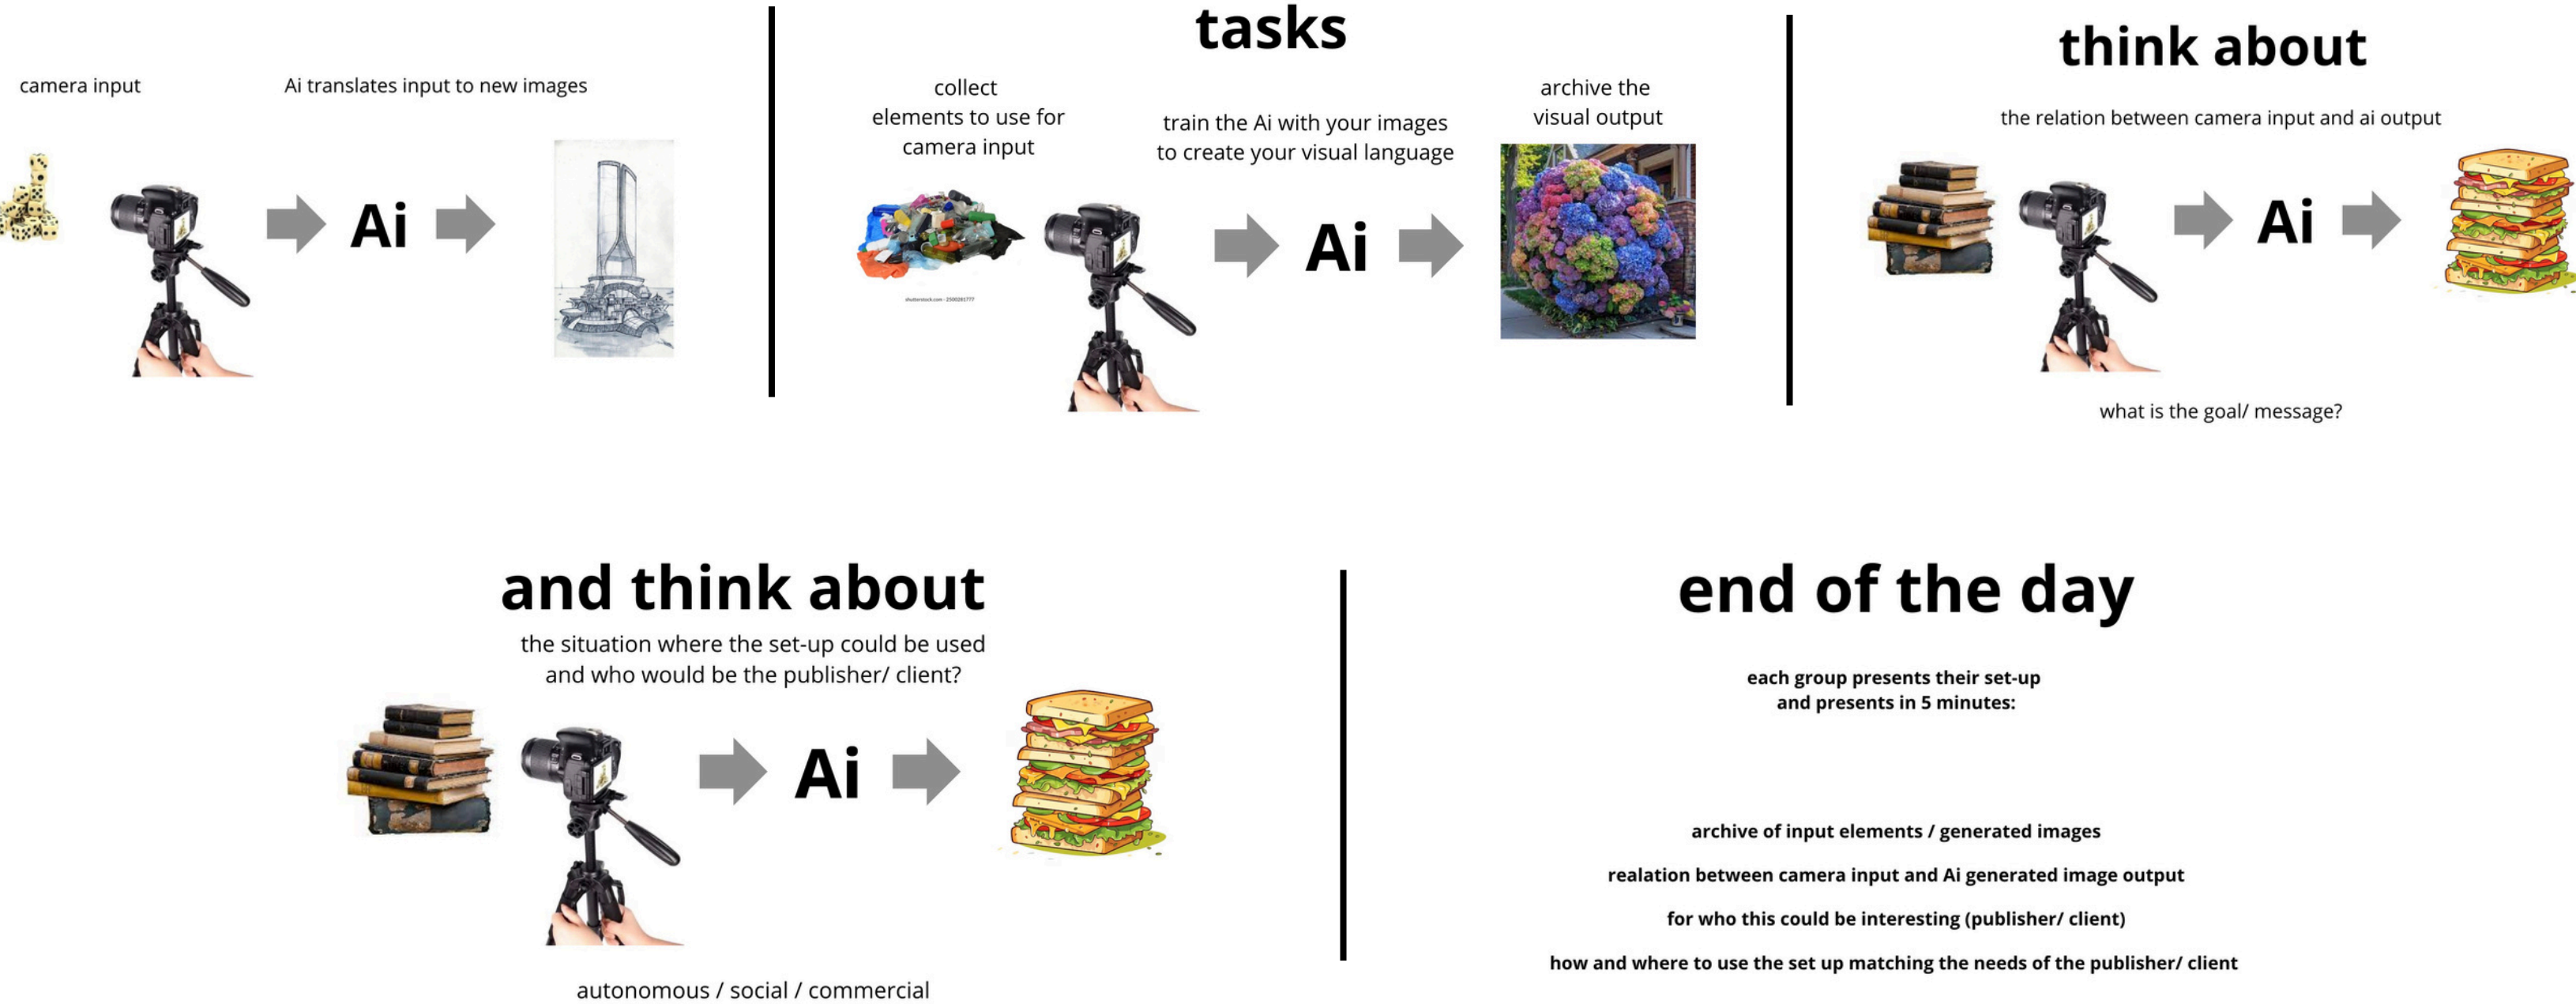


Final result



Mission AI – One Day Workshop

We organized a one-day workshop with third-year Bachelor Illustration students.
We instructed the students to work with a camera setup similar to the one we used
for our interactive children's book.



Mission AI – One Day Workshop

Students worked in small groups. We pre-installed all the software they needed. Students trained the AI with 15 images they had collected ranging from drawings, photographs and images collected from the internet.



The students enjoyed working with AI. They were fascinated to see how the camera input generated images based on their trained model, which turned out to be their main focus. There was little discussion of the relationship between camera input and generated images and the situations in which this setup could be used.

To make the workshop more efficient, it may be better to spread it out over a longer period of time, giving the students more time to experiment and think about possible applications in their practice.





Mission AI – One Day Workshop

This group of students trained the LoRA on style. The input for the machine learning and, ultimately, these images were they themselves posing in front of the camera, a mug, the camera pointed outwards, and emojis.



Mission AI – One Day Workshop

The LoRA from this group was trained on knots and entangled human bodies.



Mission AI – Conclusion

The approach we took for our *Mission AI* was very hands-on, research by making. We focused on finding ways to use generative AI on a creative level and avoid simple content generation. We came up with some interesting ideas for using AI while respecting the authorship of the designer. We also translated our findings into an educational model on that we successfully put into practice with students.

We set out to explore how current technological advancements offer new avenues for artists and their audiences to engage and interact in fresh, innovative ways. Therefore, we did not want to focus on individual technologies, but on connections.

As the technology of generative AI is changing rapidly, we didn't want to focus on specific tools, but rather on concepts. In practice, however, we ended up exploring specific technologies.

Looking back, we could summarize that our energy was concentrated on machine learning and 2D content creation, as our students and colleagues expressed their interest in these areas.

Mission AI has been a learning process, not only on a technical level, but also on a conceptual level. Generative AI challenges our views of creativity and ownership. Students, and teachers as well, seem to be mostly aware of the negative elements of AI and therefore unwilling to explore the creative potential of the technology.

Our colleagues, the teachers at Willem de Kooning Academy in Rotterdam, usually work with the knowledge they have gained in their individual specialization as designers. They know about the challenges in their specific field because they have experienced them first hand. These experiences are an integral part of our education program. Through AI, most teachers also become students. We all are in a process of learning what AI can bring to our practice.

Mission AI – Resources

Reading

[How does Stable Diffusion work?](#)

[What are Diffusion Models?](#)

[Stable Diffusion Samplers: A Comprehensive Guide](#)

[Exploring 12 Million of the 2.3 Billion Images Used to Train Stable](#)

[Diffusion's Image Generator](#)

[Comparing Stable Diffusion Models](#)

[What is a hypernetwork in Stable Diffusion?](#)

[ControlNet: A Complete Guide](#)

Models

[Stable Diffusion v 1.5](#)

[epiCRealism](#)

[Counterfeit](#)

[Ghostmix](#)

Software

[ComfyUI](#)

[Kohya's GUI](#)

[Stable Diffusion](#)

[Stable Diffusion WebUI](#)

[TDComfyUI](#)

[TouchDesigner](#)

Tutorials

[Stable Diffusion prompt: a definitive guide](#)

[How to use negative prompts?](#)

[How to use image prompt in Stable Diffusion](#)

[How to Train an Embedding in Stable Diffusion](#)

[Hypernetwork Style Training, a tiny guide](#)

[The Ultimate Stable Diffusion LoRA Guide \(Downloading, Usage, Training\)](#)

[Create INSTANT Renders from your Sketches with AI \(Control Net Tutorial\)](#)

https://interactionstation.wdka.hro.nl/wiki/LoRA_training

These links were last accessed and checked on 19th December 2024. If some of them should no longer work, we recommend accessing them through the Wayback Machine of archive.org .

Mission AI – Works Cited

Baio, Andy. "Exploring 12 Million of the 2.3 Billion Images Used to Train Stable Diffusion's Image Generator." Waxy.Org, 30 Aug. 2022, [waxy.org/2022/08/exploring-12-million-of-the-images-used-to-train-](https://waxy.org/2022/08/exploring-12-million-of-the-images-used-to-train-stable-diffusion/)

Li, Zhaoshuo, et al. "Neuralangelo: High-Fidelity Neural Surface Reconstruction." Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. arXiv.Org, 12 June 2023, arxiv.org/abs/2306.03092.

Liu, Gloria. "Dall-E 2 Makes Its First-Ever Magazine Cover for Cosmopolitan," www.cosmopolitan.com/lifestyle/a40314356/dall-e-2-artificial-intelligence-cover/. Accessed 19 Dec. 2024.

Roose, Kevin. "An a.i.-Generated Picture Won an Art Prize. Artists Aren't Happy." The New York Times, The New York Times, 2 Sept. 2022, www.nytimes.com/2022/09/02/technology/ai-artificial-intelligence-artists.html.

Mission AI – Afterword

This publication was released just after the peak of the hype around text-, image- and sound-generating AI systems. It began with the first releases of the image generator Dall-E in 2021 and the text bot ChatGPT in 2022 (both by the company OpenAI, which has Microsoft as a major shareholder). For art and design schools, the obvious question was – and still is – the extent to which these technologies will disrupt existing visual design professions, from illustration, graphic design and photography, animation and moving images, to product, fashion and spatial design.

In this debate and hype, AI is typically discussed as a single buzzword, as if it were one thing, with little or no differentiation between the sometimes different, sometimes overlapping technologies and approaches that make up ‘artificial intelligence,’ such as: symbolic AI, machine learning, generative AI and large language models (LLMs).

Most debates also make little or no distinction between (a) out-of-the-box, shrink-wrapped, corporate-developed, -trained and -owned AI generators (such as the currently popular ChatGPT text bot and the Midjourney image generator) that run as “cloud” services in corporate data centres, and (b) Open Source systems that run as personal computer software on people’s own computers, allowing them to train their own AI models. This is very much like painting by numbers or cooking an industrially produced instant meal – accepting the aesthetic and content choices of the producing company – versus painting with self-chosen colours on a self-chosen medium, or cooking with self-chosen ingredients.

This publication deals only with the latter, demonstrating the agency that artists and designers can reclaim by not accepting off-the-shelf AI.

As a practical experiment, it also shows the limits of that agency. If closed generative systems can be likened to kaleidoscopes, then Open Source systems break open their box by allowing artists and designers to bring their own visuals into the system. However, it remains an open question to what extent the algorithmic remixing of these elements may still suffer from kaleidoscopic limitations. This publication allows readers – especially designers and artists – to learn about all these practicalities, the gory details of its use, even if they do not have the time and capacity to delve as deeply into the technology.

One issue that couldn’t be covered in this publication, but remains urgent, is the resource consumption of generative AI. For Mission AI, the training of a custom model (“LoRA”) on a high-end PC (“AMD Ryzen 9 3900X 12-core CPU and Nvidia GeForce RTX 2080 T1 GPU”) takes half an hour and about 200 watt hours.

If we were to scale up Han and Brigit’s experiment and introduce intensive work with custom AI models for all of the approximately 2500 art and design students at our school, our energy consumption would increase substantially. We would also have to upgrade our entire computer infrastructure, with all the associated costs and environmental impact. Our goal of becoming an environmentally sustainable art school that teaches ecologically responsible art and design practices would be

Mission AI – Afterword

undermined. Already today, generative AI applications on server machine farms are a significant factor in increasing global electricity consumption.*

On the practical side, the lack – or difficulty – of control over AI-generated outcomes is an important takeaway from Han and Brigit's experiment. My anecdotal observation is that AI image generators are immensely popular in university research departments where academics aren't trained image makers, while professional image makers often find it faster and easier to create something manually in conventional design software than to keep tweaking AI prompts to get the desired result. Conversely, art and design students who struggle with writing make extensive use of AI text generators, while academics and other professional writers tend to agree that they produce mediocre writing and are only helpful in more limited, very well-defined usage scenarios.

All of these observations are snapshots of the year 2024/2025. This publication will soon become a historical document, similar perhaps to a graphic designer's guide to using desktop publishing software from the late 1980s, when DTP was a new and disruptive technology. It is my hope that this publication will also be of value as future media history.

Florian Cramer

- "The increase in electricity consumption by data centres, cryptocurrencies and AI between 2022 and 2026 could be equivalent to the electricity consumption of Sweden or Germany," according to Ligozat, Anne-Laure and Alex De Vries. "Generative AI: energy consumption skyrockets." *Polytechnique Insights*, 13 November 2024, accessed 16 January 2025, <https://www.polytechnique-insights.com/en/columns/energy/generative-ai-energy-consumption-soars/>.

Colophon

Published by Willem de Kooning Academy,
Rotterdam University of Applied Sciences, The Netherlands
www.wdka.nl

Authors: Han Hoogerbrugge and Brigit Lichtenegger
Additional text editing: Florian Cramer, Aldje van Meer

All images created by Han Hoogerbrugge and Brigit Lichtenegger,
partly using the Stable Diffusion AI generator, except for images by
Christina Beller, B. Pallac, Giuseppe Arcimboldo, Nicolas de Larmessin,
Charles Allan Gilbert, Matthäus Merian the Elder, Anna Sandri, Artúr
Voznickij, Justė Dovainytė, and Felix Döringer

...

Graphic design by Han Hoogerbrugge and Brigit Lichtenegger

An Open Access electronic publication
© Willem de Kooning Academy,
Rotterdam University of Applied Sciences, 2025
Freely shareable and reusable according to the terms of the
Creative Commons Attribution-NonCommercial-ShareAlike 4.0
International License

