Artificial Intelligence

Practice

An easy way to use AI locally, on your computer is comfyUI. The only trouble is that you need a GPU installed. You **can** use cpu, but it will output images really slow. On Mac, an m1 or m2/3 will work as well. The system is cross platform. Another option is AUTOMATIC1111 (Windows/Linux), AUTOMATIC on Apple Silicon.

Image Generators (semi payed, works with credit systems)

1. **Craiyon (formerly DALL-E Mini)**:
    - Free AI image generation based on user prompts. It's a smaller version of OpenAI's DALL-E.
    - Craiyon
2. **Stable Diffusion (via platforms like DreamStudio or HuggingFace)**:
    - Stable Diffusion is a popular open-source AI model that can generate high-quality images. Platforms like DreamStudio (by Stability AI) and HuggingFace allow you to use this model for free, with some limitations.
    - DreamStudio
    - Hugging Face's Stable Diffusion
3. **Artbreeder**:
    - Artbreeder is a collaborative platform where you can blend images and create unique AI-generated art.
    - Artbreeder

NightCafe, https://creator.nightcafe.studio

Video Generators

1. **Runway ML**:
    - Runway offers free trials for video generation with AI models like Stable Diffusion for video (video-to-video generation) and text-to-video tools.
    - Runway ML
2. **Pictory**:
    - A free online AI video generator that can turn text content into videos, using a simple drag-and-drop interface.
    - Pictory
3. **Synthesia**:
    - While primarily a paid platform, Synthesia offers a limited free tier for generating AI videos with digital avatars.
    - Synthesia
4. **Deep Dream Generator**:

- Focused more on style transfer and creating surreal art from images, it’s great for video if you are looking for a trippy and artistic style.
- Deep Dream Generator
- Sora
- Music: Udio

Most of these tools have free versions with some limitations, like watermarked outputs or a limited number of credits per month. If you're looking for completely unrestricted use, you might need to explore open-source options like Stable Diffusion that you can run locally on a capable PC.

Makers:

1. Hito Steyerl (Germany)

- A filmmaker and digital artist based in Berlin, Steyerl critically examines digital culture, AI, and their sociopolitical ramifications.
- Her work often critiques the role of technology in surveillance and capitalism, exploring how AI shapes our understanding of reality.

2. Anna Ridler (UK)

- A London-based artist known for her work with AI, specifically focusing on how datasets shape AI outputs.
- Projects like *Mosaic Virus* use hand-curated data sets to explore how biases are embedded within AI systems.

3. Memo Akten (UK/Turkey)

- Based in London, Akten is a computational artist exploring the philosophical implications of AI and machine learning.
- His work, such as *Learning to See*, critiques how AI perceives and interprets the world, comparing it to human perception.

4. Mario Klingemann (Germany)

- A pioneer in the AI art field, based in Germany, who critically engages with the ideas of creativity and authorship in the context of machine learning.
- He often uses generative adversarial networks (GANs) to create art that questions the role of the artist and the value of AI-generated imagery.

5. Joana Moll (Spain)

- A Barcelona-based artist who critically explores the environmental and societal impacts of digital technologies, including AI.
- Her work often deals with the hidden costs of AI and data processing, like the energy consumption of data centers.

6. Sofia Crespo (Germany/Spain)

- Originally from Argentina but now based between Berlin and Spain, Crespo uses AI to explore the relationship between biological systems and artificial life.
- Her projects critically address how AI-generated imagery can reshape our understanding of nature and biodiversity.

7. Tega Brain (Germany)

- Based in Berlin, Tega Brain is known for her critical work that intersects art, ecology, and technology.
- Her projects question the role of AI in environmental systems, often highlighting the limitations and biases of algorithmic solutions.

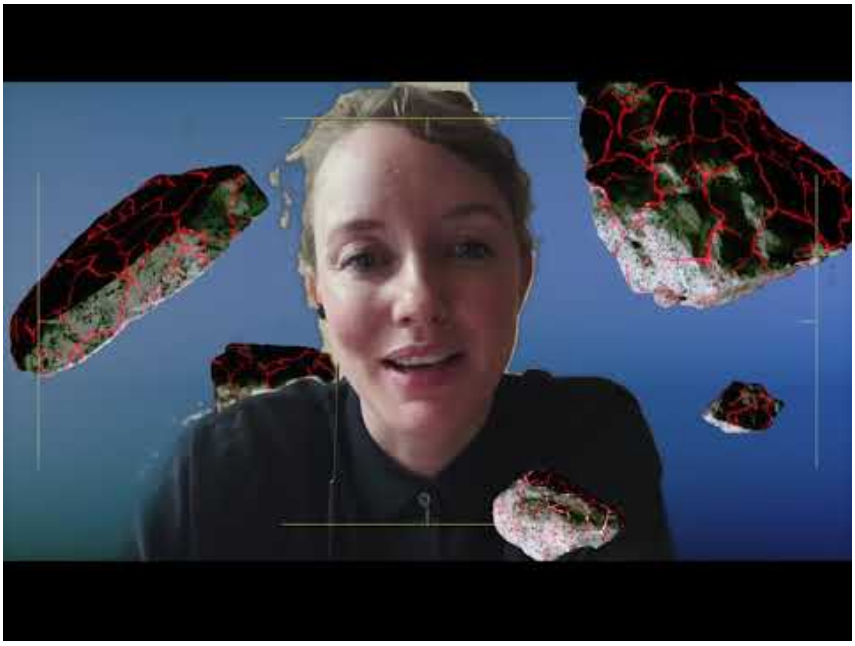

8. Zach Blas (UK)

- Though originally from the United States, Blas is currently based in London, where he engages critically with issues of AI, surveillance, and identity.
- His works, such as *Face Cages*, critique biometric technologies and the ways AI is used for control and categorisation.

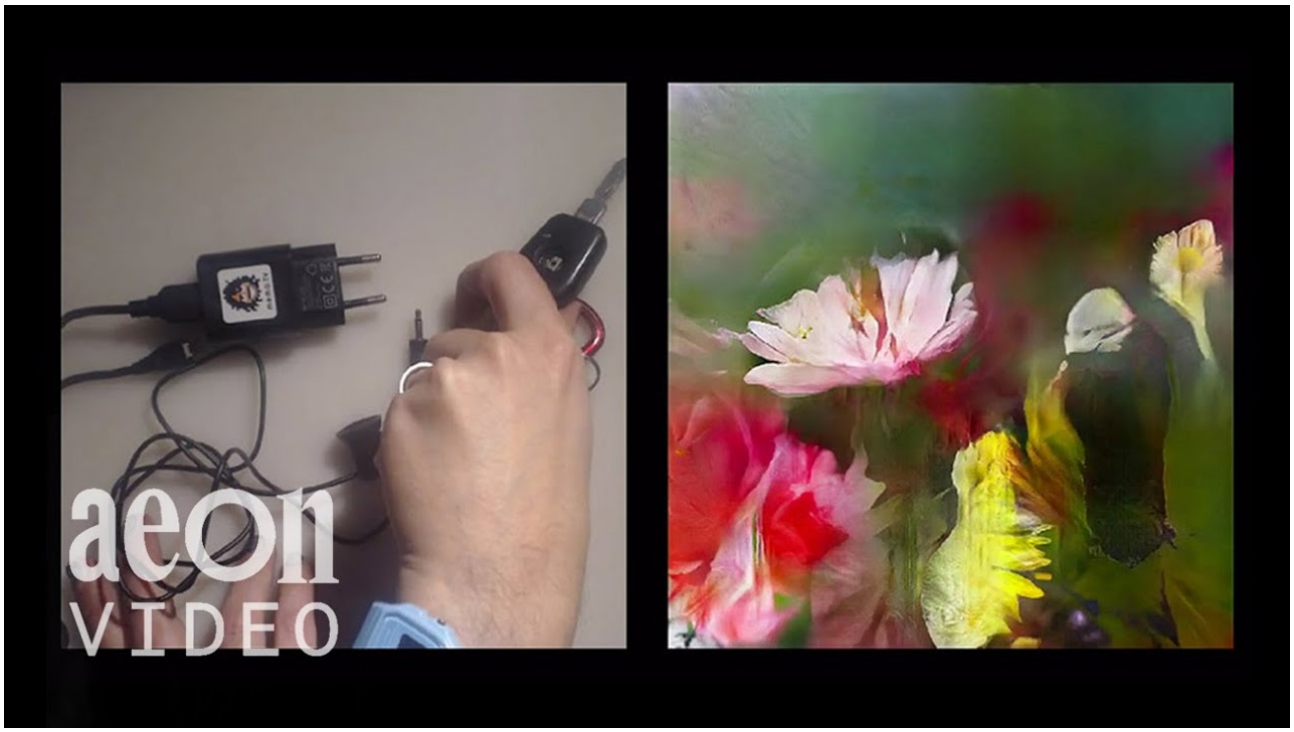

9. Memo Akten

- Focuses on the philosophical and social implications of AI, exploring how these technologies shape human perception.
- His project *Learning to See* examines how AI learns from data and how this process mirrors human vision and understanding.

10. Trevor Paglen

- Focuses on the hidden aspects of surveillance and data collection, often incorporating AI to reveal the hidden structures of technology.
- Works like *ImageNet Roulette* critically examine the biases in AI training datasets.

11. Sougwen Chung

- An artist and researcher who collaborates with AI systems to create art, particularly in the realm of drawing and performance.
- Her work often questions the boundaries between human and machine creativity, emphasising the collaborative aspect.

AI HISTORY

1. **1950s-1960s - The Birth of AI:**
    - The field of artificial intelligence (AI) was officially born in the 1950s, with the Dartmouth Workshop in 1956 often considered the birthplace.
    - Early AI researchers like Alan Turing and John McCarthy laid the theoretical groundwork for AI. Turing believed it would be possible to design circuitry which could adapt itself to new inputs leading to a kind of conditioned reflex, similar to that of animals, leading to learning. Turing suggests that intelligence might be multiple and relational, that it might exist between, rather than within beings of all diverse kinds.
    - Early AI programs focused on tasks like playing chess and proving mathematical theorems.
2. **1970s-1980s - AI Winter:**
    - Progress in AI research was slower than anticipated, leading to what is known as the "AI winter."
    - Funding for AI projects decreased as expectations were not met, and interest waned.
3. **1980s-1990s - Expert Systems and Narrow AI:**
    - AI research shifted towards "expert systems" that could solve specific problems by emulating human expertise.

- Systems like MYCIN (diagnosing bacterial infections) and Dendral (chemical analysis) gained attention.
- Machine learning techniques like neural networks continued to develop.

4. **1990s-2000s - Rise of Machine Learning:**
    - The concept of machine learning, particularly supervised learning, gained prominence.
    - Researchers developed algorithms for tasks like speech recognition, computer vision, and natural language processing.
    - The Internet's growth provided vast amounts of data for training AI systems.
5. **2010s - Deep Learning and AI Resurgence:**
    - Deep learning, using neural networks with many layers, led to breakthroughs in image and speech recognition.
    - Companies like Google, Facebook, and Microsoft invested heavily in AI research.
    - AI applications in self-driving cars, healthcare, and recommendation systems became increasingly common.
6. **Present and Future - AI Everywhere:**
    - AI is now deeply integrated into everyday life, with voice assistants, recommendation algorithms, and autonomous systems.
    - Research into general AI (AGI), which can perform any intellectual task humans can, continues, with companies like OpenAI and DeepMind at the forefront.
    - Ethical and societal concerns surrounding AI, including bias and job displacement, are central topics of discussion.

This history provides a broad overview of AI development, but it's important to note that AI is a dynamic and evolving field, and advancements continue to be made at a rapid pace, making it a central force in technological innovation and societal change.

A brief overview of different types of AI systems:

1. **Supervised Learning**:
    - **Description**: The AI learns from labeled training data, where each input has a corresponding correct output. It aims to learn the mapping between input and output.
    - **Applications**: Image classification, speech recognition, spam detection.
    - **Example**: A model learns to classify emails as "spam" or "not spam" by studying a dataset where each email is already labeled.
2. **Unsupervised Learning**:
    - **Description**: The AI is provided with data that doesn't have labeled responses and finds patterns or structures in the data on its own.
    - **Applications**: Clustering (like grouping similar customers), anomaly detection, dimensionality reduction.
    - **Example**: Identifying customer segments based on purchasing behaviour without prior knowledge of segment labels.
3. **Reinforcement Learning (RL)**:
    - **Description**: The AI learns by interacting with an environment and receiving feedback through rewards or penalties. It aims to maximise cumulative rewards.
    - **Applications**: Game AI (like playing chess), robotics, autonomous driving.

    - **Example**: Training a robot to walk by rewarding forward movement and penalising falls.
4. **Semi-Supervised Learning**:
    - **Description**: A mix between supervised and unsupervised learning. It uses a small amount of labeled data and a large amount of unlabelled data for training.
    - **Applications**: When labeling data is expensive or time-consuming, like in medical image analysis.
    - **Example**: Using a few labeled medical scans to help classify a large set of unlabelled scans.
5. **Self-Supervised Learning**:
    - **Description**: A form of supervised learning where the system generates its own labels from the input data. The goal is often to predict part of the data from other parts.
    - **Applications**: Natural language processing (NLP), computer vision.
    - **Example**: Predicting the next word in a sentence, as seen in models like GPT.
6. **Transfer Learning**:
    - **Description**: This involves taking a pre-trained model on a related task and fine-tuning it for a new task. It allows leveraging knowledge from a larger problem to a smaller one.
    - **Applications**: NLP, computer vision, speech recognition.
    - **Example**: Using a model trained on a large image dataset to classify medical images with limited labeled data.
7. **Generative Adversarial Networks (GANs)**:
    - **Description**: Consists of two models: a generator that creates data and a discriminator that distinguishes between real and generated data. They train together to improve.
    - **Applications**: Image generation, style transfer, deepfake creation.
    - **Example**: Creating realistic-looking images from random noise.
8. **Deep Learning**:
    - **Description**: A subset of machine learning that uses neural networks with many layers (deep neural networks) to model complex patterns in large datasets.
    - **Applications**: Computer vision, NLP, speech recognition.
    - **Example**: Recognising objects in images using convolutional neural networks (CNNs).

Statements about AI

EGO

Only algorithms could create such fluid and unstable imagery like AI. These AI images are moulded from the internet hive mind. They are impossible to track down. It is impossible to establish which particular node in the network they originate from. This is not necessarily a bad thing, this instability establishes a reduction of ego behind content. Much like John Cage used the I Ching to erase individual tastes from his work, AI images are tasteless and unstable. Their instability is more a result of their mixed origins and uncertain outcomes than by pure artist intend, the human ego. Cage's philosophical

stance on embracing the inherent beauty and unpredictability of the world made him question his own tastes and therefore alter his creative process to render it more inclusive of unforeseen elements. Cage aimed to create music that was not predetermined by the composer's will but was instead shaped by random processes and the guidance of the I Ching.

Question: Is the human ego evaporating behind AI image content desirable?

THIS MIND IS POOR BUT IT HAS A MEMORY

Early image generators like StyleGANS could be qualified as ‘poor’ images, in the way that Hito Steyerl describes poor images in her essay ‘In defence of the poor image’ from 2009. The redistribution of early (poor) internet images with their crappy dimensions is however done post-creation. In the case of AI training sets, the ‘redistribution’ is done before the actual image is generated. The prompt merely delves into the existing wired up well of image training data, to conjure up some kind of relevance to it. One could say that the reproduction of the AI image is done well before the output of the image. These early StyleGANS are somewhat foreclosed to artists, because they have been used in abundance by now. But who knows, somebody might find a new use for them in the future. Artists might find new opportunities to challenge these patterns and imagine new possibilities. This could especially be done in retrospect, when the AI medium has become less hot to handle.

In Walter Benjamin’s ‘The work of art in the age of mechanical reproduction’ from 1935 a machine (in this case a camera) is used to create images, but they are still images of, reflections of reality. AI images on the other hand are reflections of the ‘internet reality’ or of the internet spectacle. The machine in more recent times has evolved into a network of machines. Machines now reside in their own ‘hive minds’. Even if this network is not intelligent, it does possess a memory and a capability to process and mix imagery from all kinds of backgrounds. In the form of memes it can even reflect upon its own cringe state of being. Memes are poor images that reflect upon their own cringe state of poorness.

Question: Is AI ‘taking pictures’ of a non-historic past? (Historic images become fluid because they are thrown in together, in a mix.) What are the consequences of a visual history that is unstable?

FAKING AN AURA

Benjamin talks about the loss of the aura with mechanically, photographically produced images and films. With AI images there is not only a loss of the physical aura of an artwork (for instance, an actual painting), but also **a loss of human aura**. The human presence is reduced to almost the mere state of providing fodder for the algorithms to feed from. Again, this is not per se a bad thing, because human centric approaches are shifting in all domains, in different cultures. Human centricity is something that has not aged well, resulting in all kinds of global, ecological and economical problems. The AI image is almost autonomous, but without realising this.

Benjamin describes the aura as follows: "Even the most perfect reproduction of a work of art is lacking in one element: its presence in time and space, its unique existence at the place where it happens to be. This unique existence of the work of art determined the history to which it was subject throughout the time of its existence." (Benjamin p4) AI images therefore exist outside of history. See also Debord and history of social life below.

Question: Do you think there is a specific aura to AI content emerging?

RICH & SPECTACULAR

Recently AI images are trying to shed their ‘poor’ state (in contrast with ‘cold’ state by McLuhan?). Image generators like midjourney or Dall-E are doing their best to produce images that are high res and are capable of capturing some form of reality in a spectacular way. They are securing their place in the society of the spectacle. Guy Debord describes the spectacle as follows: "All that once was directly lived has become mere representation.” Debord argues that the history of social life can be understood as "the decline of *being* into *having*, and *having* into merely *appearing*." AI images are becoming seemingly lived. They are the opposite of the individual. They represent a mix of realities and find similar happiness there. They excel in everything, you can basically throw anything at them, or train them with whatever data, and they’ll have a go at it. AI has gone from poor to rich images that have become prime examples of the Debord spectacle.

Byung Chul Han

Question: Is current AI just a spectacle? A form of content created to make us need it?

IS AI NEW? IS IT PER DEFINITION RETRO? DO WE FINALLY HAVE A NEWLY LOST FUTURE?

Because of the used training data, most AI generated images seem familiar or even retro. Mark Fisher's concept of "hauntology" is a term he popularised in his writings, particularly in his book "Ghosts of My Life: Writings on Depression, Hauntology and Lost Futures.” Mark Fisher's concept of *hauntology* invites us to critically examine the cultural and psychological consequences of living in a world where the promises of a brighter future have been shattered, leaving us haunted by the past and trapped in a sense of melancholic stasis. AI image generators interact with the past, they produce uncanny or nostalgic imagery. AI image generators often use large datasets of existing images to create new ones. In this process, they repurpose and remix elements from the past. This can be seen as a form of *hauntology* because it involves the reemergence of visual motifs, styles, and cultural references from the past. The generated images may carry echoes of previous eras, contributing to a sense of nostalgia or uncanniness. By relying on established patterns and historical data, AI image generators can contribute to a feeling that we are stuck in a loop of recycling past aesthetics and ideas rather than forging new creative frontiers.

Question: Are we stuck in a loop of recycling past aesthetics and ideas rather than forging new creative frontiers with AI?

AI ORGANISMS

Could we make AI less human centric? Train AI’s to incorporate more non-human intelligences. AI as it stands now, is primarily based on the human brain. What about other forms of biology like (animal) bodies, algae, octopuses and plant life? Lynnn Margulis, Donna Haraway… Is it so bad to have another intelligence in the mix, besides our own? We haven’t really performed that well in reference to for instance wars and ecology. The current craze that AI will become more intelligent than us, is yet another acknowledgement of humans not being the centre of things, comparable to the Copernican revolution or recent research in biology and quantum mechanics. Could we make AI less certain of itself?

Question: Should we make AI less human centric?

SPECTERS OF COPYRIGHT

Fluidity of AI images require an equally fluid approach to copyright. Is there still such a thing as copyright when the copy has been copied innumerable times over? Images have been taken up in a gigantic hive. They have been mixed, shredded, (inter)changed meaning. The right of the copy is thus muddied and has lost its purpose in the process. In art there has been a long tradition in copying images, appropriating them, giving images new life, etc. Since Andy Warhol, the copy has taken on a life of its own. The copy has lost its negativity. The limiting idea of the positive and the negative image has been overcome. This freedom resulted in images to be freely used by everyone.

Specters of copyright: In some cases, AI image generators challenge traditional notions of creativity and authorship. They operate in a way that echoes *hauntological* idea of the past, haunting the present. AI algorithms draw from a vast repository of historical images and styles, making them creators that are, in a sense, haunted by the artistic and visual history they have learned from.
Mark Fisher's hauntology provides a lens through which we can analyze how AI image generators engage with the past, memory, and cultural production. These technologies can evoke feelings of nostalgia and uncanniness, blur the lines between creativity and imitation, and contribute to a broader discourse on the haunting effects of the past on our contemporary visual culture. At the same time, they also offer opportunities for artists and thinkers to challenge these patterns and imagine new possibilities.

Question: Is there still such a thing as copyright? Should it still exist?