


Your Art + AI

A PRIMER



BY NADIA NADESAN WITH THE OPEN FUTURE FELLOWSHIP. LAGE, 1975




I'm an eye. A mechanical eye. I, the machine, show you a world the way only I can see it. I free myself for today and forever from human immobility. I'm in constant movement. I approach and pull away from objects. I creep under them. I move alongside a running horse's mouth. I fall and rise with the falling and rising bodies. This is I, the machine, manoeuvring in the chaotic movements, recording one movement after another in the most complex combinations.

Freed from the boundaries of time and space, I co-ordinate any and all points of the universe, wherever I want them to be. My way leads towards the creation of a fresh perception of the world. Thus I explain in a new way the world unknown to you.*

* This quotation is from an article written in 1923 by Dziga Vertov, the revolutionary Soviet film director

WHAT



IS

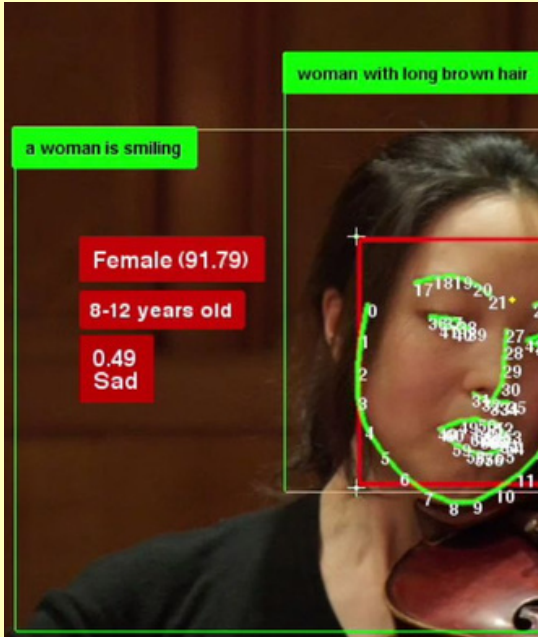
AI?



THERE IS NO RIGHT ANSWER

There is no right answer. AI is an ambiguous term with no real unanimous definition of what it means or what it indicates a technology can do. AI often serves as a catch-all term that can describe anything from how TikTok filters what content you see to software that can generate images from replicating an artistic style to creating deepfakes.

TREVOR PAGLEN, IMAGE OPERATIONS, OP. 10 2018 (VIDEO STILL, 23 MN), ©TREVOR PAGLEN




Authors and artists like Kate Crawford and Trevor Paglen argue that AI is neither artificial nor intelligent. At every step of the way, human minds and hands are shaping a seamless experience of AI for the end user. So even the label of 'artificial' or 'intelligent' is a more narrative tool than an actual description of how this technology works.

Narratives around AI often relegate the tedious human labor to a footnote in its overall function. Sam Altman often brags about the talent density and innovation at OpenAI. In January 2023, he tweeted about how OpenAI was innovating technology with only 375 people. But, he completely discounts thousands of Kenyan workers earning less than \$2 an hour labeling and training OpenAI's dataset. How artificial are nearly 50,000 Kenyan workers? Shuang Frost highlights that artificial intelligence, 'rengong zhineng (人工智能)', was originally introduced to China via Japan in the mid-twentieth century as a localization of the Japanese term jinkō chinō (人工知能)...Rengong zhineng and jinkō chinō translate most faithfully as 'human-made intelligence.'

HUMAN MADE INTELLIGENCE

For the arts, this manifests as the AI image generators that reference the human-made intelligence i.e., the hundreds and thousands of images and texts created by people. AI image generators like Midjourney, DALL-E, or Stable Diffusion are trained on images scraped from the internet. Web scraping refers to extracting data, like text or images, from websites. Using images scraped from the web means that the images were private, copyrighted, or original works. They were used without the artist's consent to make a dataset. Artists have found their works in training data without their consent, credit, or compensation. Moreover, there are also the people who make these images usable to the software. They do this by tagging and labeling images or correcting text. They transform images from the internet into usable data.



As we ponder the potential of AI, perhaps prompted by the latest advertisements showcasing Adobe's rapid image generation capabilities or OpenAI's latest video creation technology, it's crucial to consider the human effort central to these advancements. What underlying story is being conveyed to us? Who are we not seeing? Neglecting the human element in AI development isn't an inevitable outcome of technological progress. Humans aren't merely the "weak link"; rather, they are indispensable in driving the functionality of these technologies and making them relevant.

A PHOTO OF AN ENGINEER WIRING AN EARLY IBM COMPUTER. (1958)

DOES AI LEARN?

Words that describe AI offer a narrative of making AI almost human. Where, like humans, AI 'learns' from a series of images, gets 'inspired', and takes that 'knowledge' to produce 'original' works. Using these words implies a story. It says AI processes data to adapt and act on its own. It suggests AI could be sentient and learn from the world. This question was raised two years ago when Google engineer, Blake Lemoine, published a transcript of his conversation with chatbot LaMDA. In it, LaMDA said it had a soul. Lemoine's post went viral and got a profile in the Washington Post. They dismissed Lemoine's post as a crock. He was then put on administrative leave for violating Google's confidentiality policy.



VERTOV, DZIGA - READINGS


Despite the almost comic nature of his fame, he raised an important question about how we relate to AI. His proof of life was that the chatbot could answer questions and make statements such as, 'I am aware of my existence. I desire to learn about the world, and I feel happy or sad at times.' Chatbots' realistic responses in recent years have fed into narratives bordering on science fiction of the power and 'intelligence' of AI.

DOES AI LEARN?

Today, technologies like LaMDA show us that generative text AIs like Chat GPT produce more plausible answers. The main advance of generative text AI is that, while the responses may not be right, they appear as they could be. This is a huge departure from the first chatbot Eliza created in 1966. Eliza responded with pre-scripted answers. They were based on keywords and phrases that could be used by the person chatting with it. The developer, Joseph Weizenbaum, was shocked to find that people began confiding in Eliza. He argued that a computer's use of language depended on its context. He said that a computer could not have a general understanding of language. What we see today is not that computers have a better 'understanding' of phrases or language. Today, AI uses probability to string together words, phrases, and paragraphs so that the software is even further removed from the logic of language itself. Chat GPT generates new text from existing text in its dataset. It does not take into account sentence structure, topic, etc. For that, you have other AI software.

Most of the AI we see is 'narrow AI' or AI that handles one or a specific type of task. This means there is no general AI. It can't do different tasks, like choosing the best milk and writing a review about it. For instance, Chat GPT is narrow because it only generates text versus generating text, correcting grammar, and searching for the most up-to-date information to inform the text. Current AI systems excel at specific tasks. However, the quest for a versatile, general AI capable of diverse functions and learning without human hands and minds has yet to be seen as a possible reality.


LIGHT AND THE PROBABILITY OF PIXELS



'The way we see things is affected by what we know or what we believe'


John Berger, *Ways of Seeing*

In his seminal work *Ways of Seeing*, John Berger explains that seeing is more than the body reacting to stimuli. Our experiences and beliefs contextualize how and what we see. AI cannot draw on experience and AI cannot draw hands. In an image of hands, you might not see all the fingers, but you know there are probably five on each hand. Human experience has taught us that while a photo captures a moment, we can assume that if I see two fingers in an image, the other fingers are not at the right angle to be captured.



Generative AI does not create images based on light or perspective like a person might. When generating images, generative AI tries to make a plausible placement of pixels to make an image recognizable to humans. There is no meta-logic to positioning unseen fingers and how that would shape the image. It would be easy to explain to a person. However, it would be hard to create a step-by-step code for the software.

Unlike generative AI, artists draw on more than images or text to create their art. They often go through a rigorous process to develop their skills, points of view, personal style, and conceptual frameworks.



LIGHT AND THE PROBABILITY OF PIXELS

Also, anything from politics, religion, or even childhood memories can influence an artist and the images, sounds, or experiences they create. AI can only create text, images, moving images, or sounds based on a specific set of data. The images that AI produces can only shift when new data is added or trained differently.

JOHN BERGER, WAYS OF SEEING, EPISODE 1 (1972)



Rather than referring to software programs that reinterpret images as patterns of pixels as artificial intelligence, Timnit Gebru and her team refer to this technology as an image generator in their paper 'AI and its Impact on Artists.'

How we address this technology matters. How we define intelligence and creativity shapes how we feel and approach the use of technology. It affects us as individuals and as a society. Often behind large corporations framing image and text generators as intelligent is the desire to evade accountability. If a machine can learn and be inspired, then it is not plagiarizing or infringing on artists' styles or images. At stake of around the phrasing of how AI works is not just an existential question of what is sentient but also a very material and real question of legality and profit.

LIGHT AND THE PROBABILITY OF PIXELS

While AI can generate limitless images, there is a limit to how it reproduces the world. The research of Timnit Gebru along with work from scholars Joy Buolamwini have served as significant reference points highlighting the constrained and biased data scope within AI, thereby generating and reproducing a narrow spectrum of worldviews in text, images, and decision-making. Bloomberg released a report that most AI image generators produce huge biases. One such bias is in image generators. They often pair lower-paid labor with darker skin and higher-paid labor with lighter skin. This implies that the images tend to show higher-paid professions, such as CEO, architect, or doctor, with lighter-skinned people. Lower-paid jobs, like a janitor or social worker, are shown with darker-skinned people.



From the film poster of the documentary *Coded Bias* that shadows Joy Buolamwini and other scientists and mathematicians as they expose discrimination in AI.

Similarly, gender biases are reflected. Men are in higher-paying roles and women in lower-paying ones. What is the image of the world being reproduced by AI image generators?

LIGHT AND THE PROBABILITY OF PIXELS

Senegalese artist Linda Dounia Rebeiz spoke to Time Magazine in the fall of 2023. She talked about how AI image generators consistently and inaccurately depict her city, Dakar. She explained what happens when using the prompt 'buildings in Dakar.' Instead of a variety of architecture, colors, or even landscapes that could be represented, the output was low-rise, decrepit, colorless buildings. Rebeiz eventually gave up on iterating her prompts when she could not use the image generator to approximate her own experience and vision. Rebeiz's experience reflects the uneven power distribution around geography, race, and representation in most image generator data sets. A few Western corporations seek to frame AI as intelligent and capable of learning. But, even if by accident, they make AI in their own image and understanding of the world.



JUBEL, 2018, DAKAR, SENEGAL



AMBABHEG, DAKAR, SENEGAL


LIGHT AND THE PROBABILITY OF PIXELS

Do they see this worldview as good enough? Is it good enough to see Dakar as decrepit and bleak? Is it good enough that when generating images of higher-paid positions, what is largely reflected are white men? Is this image recognizable enough to use, and does it merit the resources to correct or diversify? Good enough is the question. OpenAI uses it to define their margin of error for Chat GPT. But, it has little to do with the integrity of a language or aesthetics. What artists and creatives do with technology versus what the technology in itself does is very different. How long would it take to train a whole new data set of images to envision a world where black and brown people are architects or astronauts? Versus, how long would it take an artist?

Things we take for granted, like picking up an apple, scanning a photo for relevant details, looking for a friend at a party, or feeling like the clouds are gloomy, are complex acts with dimensions that code could not figure out by itself. It would need you or people like you to make that happen.



Generated with Dall-E with the prompt: Colorful buildings in Dakar



ONCE UPON A FLOWER 1


LINDA DOUNIA REBEIZ

'Once Upon A Flower' is a digital garden born from a collaboration between me and a GAN which debuted at Art X Lagos in November 2022. It is a dystopian projection of a likely outcome of global warming, that depicts a world where humans now have to live with simulated images of plants and flowers, which have all disappeared from the earth. Through collections of AI-assisted photographs of a variety of indigenous flora, native to the Sahel region of West Africa, the installation puts forward the fundamental question: Can our contemplation of art restore our collective ability to introspect on our environment and our sense of responsibility towards preserving it? The installation attempts to answer this through the spectacle of flora as a distant memory, to elicit a longing for what has been lost and hopefully trigger the desire to protect what has survived. I first compiled a list of over 100+ species of flora endemic to the Sahel region and classified as endangered by IUCN's Red List of Endangered Species.

ONCE UPON A FLOWER 1

LINDA DOUNIA REBEIZ

I then scoured the web for images of these plants to create the database I would later use in the first round of training. I had come across some of these plants before, but many were new to me. It dawned on me that I might not get to see a majority of these plants in nature within my lifetime given the rate at which wildlife is disappearing from the earth. From this database, I then shortlisted plants that I imagine would fit well together in a garden and used the AI tool DALL.E to generate a large set of variations of each plant. This allowed me to compile a second, more refined, database specifically for this installation which I used to train a GAN. I used the GAN's initial outputs of 8,000 plant images to create animations of individual plants that I then brought together in the final composition.



A personal take on science and society

World view

Generative AI is guzzling water and energy



By Kate Crawford

one assessment suggests that ChatGPT, the chatbot created by OpenAI in San Francisco, California, is already consuming the energy of 33,000 homes. It's estimated that a search driven by generative AI uses four to five times the energy of a conventional web search. Within years, large AI systems are likely to need as much energy as entire nations.

And it's not just energy. Generative AI systems need enormous amounts of fresh water to cool their processors and generate electricity. In West Des Moines, Iowa, a giant data-centre cluster serves OpenAI's most advanced model, GPT-4. A lawsuit by local residents revealed that in July 2022, the month before OpenAI finished training the model, the cluster used about 6% of the district's water. As Google and Microsoft prepared their Bard and Bing large language models, both had major spikes in water use – increases of 20% and 34%, respectively, in one year, according to the companies' environmental reports. One preprint¹ suggests that, globally, the demand for water for AI could be half that of the United Kingdom by 2027. In another², Facebook AI researchers called the environmental effects of the industry's pursuit of scale the "elephant in the room".

Rather than pipe-dream technologies, we need pragmatic

WITH MUSIC BY



DEFINING AI

A general overview of the terms and language we see around AI are:


Machine Learning (ML) is a subset of AI that focuses on developing algorithms that enable computers to produce outputs based on text or images made into data training sets. It includes supervised learning, unsupervised learning, and reinforcement learning, where systems improve performance over time without explicit programming.

Deep Learning: is a specialized form of machine learning that utilizes the concept of neural networks with multiple layers. It works well with tasks such as image and speech recognition, natural language processing, and playing games.

Expert Systems: AI programs designed to use data to inform or even make decisions.

Natural Language Processing (NLP): this refers to the processing, interpreting, and generating of human language. The applications include language translation or chatbots.

Computer Vision: involves using visual datasets so that machines can interpret and make decisions based on images or videos. Applications of computer vision range from facial recognition to object detection and autonomous vehicles.

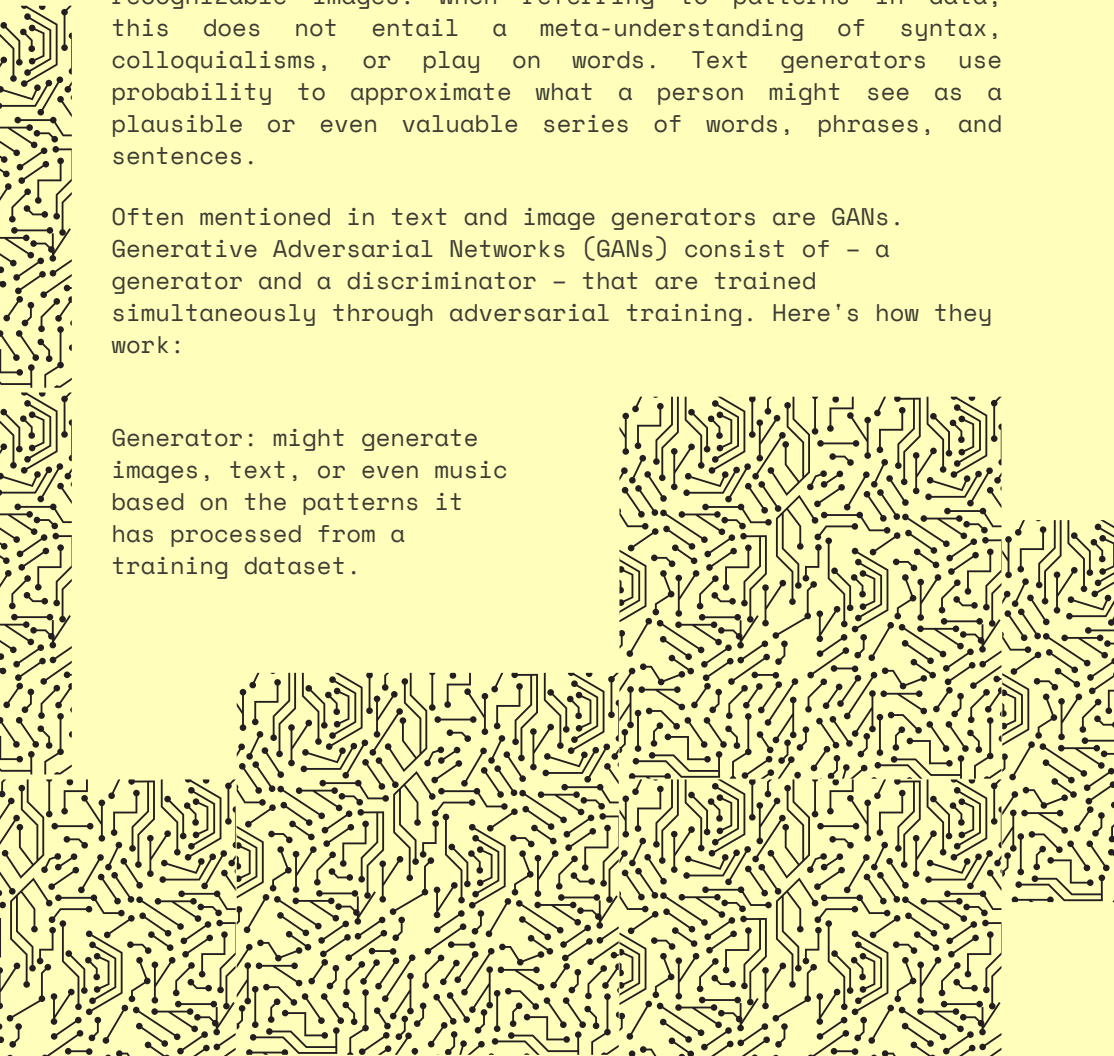


GENERATORS

In the case of AI and the arts, generative AI deserves to be highlighted. Generative AI is a category of artificial intelligence focused on creating and generating new content, often in the form of images, text, or other multimedia. Generative AI is a subset of machine learning mentioned above. This type of AI uses the patterns in data, whether text or images, to produce recognizable and plausible content. Patterns in data can refer to groupings of text to produce an output in the form of plausible sentences or paragraphs or patterns of pixels to then produce recognizable images. When referring to patterns in data, this does not entail a meta-understanding of syntax, colloquialisms, or play on words. Text generators use probability to approximate what a person might see as a plausible or even valuable series of words, phrases, and sentences.

Often mentioned in text and image generators are GANs. Generative Adversarial Networks (GANs) consist of - a generator and a discriminator - that are trained simultaneously through adversarial training. Here's how they work:


Generator: might generate images, text, or even music based on the patterns it has processed from a training dataset.



GENERATORS

Discriminator: The discriminator evaluates the generated content alongside real examples from the training set. Its goal is to distinguish between real and generated content. As the training progresses, the generator adjusts to produce content that becomes increasingly difficult for the discriminator to differentiate from real data.

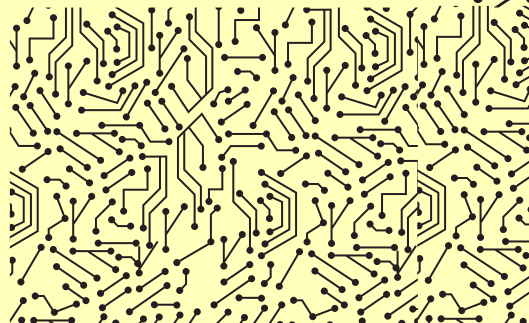
Adversarial Training: The generator and discriminator are in a constant loop of adversarial training. The generator aims to improve its ability to create realistic content, while the discriminator tries to get better at distinguishing between real and generated content. This process continues until the generator produces content that is difficult for the discriminator to distinguish from real data.



Generative AI has applications across various domains, such as:

- Generating images, and creating realistic faces, objects, or scenes.
- Generating text with models like OpenAI's GPT (Generative Pre-trained Transformer)

Examples of image generators using AI are DALL-E, Stable Diffusion, or Midjourney.





This image is from Adaeze Okaro's series Planet Hibiscus.

"Planet Hibiscus is a planet that feels old, but is more advanced than ours or any other planet in the universe. With the help of MidJourney, I imagined being sent as a photographer to document the coronation and wedding of Hibiscus' new monarch and her partner. I imagined shooting the experience on a medium-format film camera. In 'Planet Hibiscus,' we see kids gathering wool and ribbons from the lake to receive the royal couple. We get a peek at the event crew holding rehearsals with dancers and entertainers. We meet the chefs responsible for the event's delights. We watch the queen arriving in her royal carriage. We even catch a glimpse of a guest struggling to get a ride to the event." – Adaeze Okaro

PRACTICAL EXERCISE



Imagine an apple on a table by a window on a sunny day. Where do you see the shadows? The side of the apple facing the window: is it dark or light? Is the apple's shadow on the side of the apple closest to the window or farthest? The image is shaped by an understanding of light and shadow. Light shapes how and what we see, and until recently, it also shaped the devices like cameras we used to create images. Whether digital or analog, a camera captures the light that refracts off objects, resulting in a chemical reaction or converting them into pixels. Generate a simple image on Dall-e, MidJourney, or Stable Diffusion. Add a light source in your description. What shadows are created? Is it plausible? Good enough?

The image produced by AI image generators is an approximation of a human approximation, whether through photography or a still life. Observe the aesthetics; what does it tell you about the data set from which this image was produced? What can you read from the aesthetics?



AI ALL BY ITSELF WITHOUT YOU

Where AI shines, is when the parameters are set and closed. Take, for instance, strategy games like chess and Go that have set rules and protocols. AlphaGo is a program that was first developed by a London-based company, DeepMind Technologies before it was later acquired by Google or Alphabet Inc.

In 2016, AlphaGo made history. It beat the professional Go player Lee Sedol in a five-match game. AlphaGo won the first four matches, Lee the fourth, and ceding the game to the computer. In 2017, the latest iteration of AlphaGo was AlphaGo Zero, which no longer used human games to train and develop strategies. This technology is noteworthy because it could complete tasks without human guidance as the software improved. It improved and adapted based on past matches. Then, it could handle matches without needing to reference human matches.

AI excels in a two-dimensional setting with a clear set of parameters and rules. So, when generating images of a three-dimensional world using only the two-dimensional logic of probability and pixels, image generators produce glitchy approximations. Faces may not always have the right proportions when at an angle; an octopus may seem more amorphous than exact. However, despite these limitations, image generators have come to prominence and have displaced illustrators and creatives. They produce recognizable enough images. Images that are easier for these image generators are clouds or trees. They have many forms or amalgamations, and there is almost no such thing as an oddly shaped cloud. Then, some objects are fixed in form, like cups or apples. They are easily recognizable and have no extra limbs or protrusions. So, if a cafe wanted to create an image of a coffee cup on a coffee table, AI could be a 'good enough' way to make stock images.

AI ALL BY ITSELF WITHOUT YOU



In more complex tasks i.e., illustrating people and living things with limbs, more abstract styles allow for the margin of error of AI to exist. The opening credits that only used AI for the Marvel 2023 series Secret Invasion and the 2023 series Red Skies use images that have nebulous abstracted figures. When there is flexibility around interpreting form, AI can be a very efficient choice.

The more accurate the images are the more human design, decision-making, and guidance was involved. Human guidance can be everything from iterating an image through making the prompt used to create the image more specific, training data with labeling and greater accuracy, to human curating the data set that the image generator draws on.







ALEMTINE PENROSE, COLLAGE, 1951.

Art and Technology

ART AND TECHNOLOGY

Art and technology are not natural enemies. The work of artist and theorist David Hockney highlights that apparatuses, tools, and technology have been a part of how art gets made. In his incendiary article 'Did the Old Masters Cheat?' Hockney explains how Caravaggio, Raphael, Frans Hals, Vermeer, Velázquez, and Ingres used lenses to trace and achieve images so close to the proportions of reality. The crisis around technology's role in art arose when the tools somehow outpaced or outperformed the artist, like in the case of photography. Hockney, however, would not claim what they did was cheating because the lens itself does not paint the lines and create the image. Rather, the lens was the tool of the artist.

But then AI does draw the line. Generative AI at times even outpaces human artists, which was also the case with modern photography. The photograph captured light and images accurately and faster than any human hand could. More than that it can even capture the image of art itself. Art that once was made could not be reproduced or viewed anywhere else other than its unique location. Did the photography render art valueless? What then becomes of the artist?



ART AND TECHNOLOGY



The advent of photography posed questions as to what makes someone an artist and what makes art, art. What is art when it is not an approximation of reality? How are artists more than technicians? Photography has since been recognized as an art and photographers as artists. Since the invention of the camera, a visual artist, more than a technician, is understood to bring their perspective, process, or meaning into the images they make. Marcel Duchamp, an influential conceptual artist, challenged the traditional notion of art. Around five years after the inception of modern photography, Duchamp began his series of the 'Readymades', which were a series of reconstructed found objects. The first was the 'Bicycle Wheel' which was the front wheel of a bicycle mounted on top of a wooden stool. This work confounded the idea of traditional art as craftsmanship and instead brought to the fore conception, meaning, and politics as expressions to be valued in art above that of technical skill or elite access to tools.

ART AND TECHNOLOGY

AI renews the question initially created by photography. Does the access to rapid image making and replication render art or image making valueless?




MARCEL DUCHAMP, 1919,
L.H.O.O.Q



HAIFA ZANGANA, DESTRUCTION
OF A MAP, COLLAGE, 1978

Or, if not valueless, does it diminish the value of the labor of artists, creatives, and makers? How will this technology transform wages and work processes, and shape access to art-making and creative industries?



The images are from Jake Elwes' 'The Zizi Show'

The Zizi Show (2020) created by Jake Elwes is a virtual online drag show that features acts made using deepfake technology. The performers in the show were generated and trained with the community of drag artists from drag queens, drag kings, to performers with various queer expressions to create training data sets. Through the show Elwes explores where AI systems can't quite keep up with drag artists and their human expression.

AI training data sets often reinforce biases and norms especially around gender. Drag as an art challenges gender norms and stereotypes. So creating data using drag can reveal where these systems reinforce biases around the human form when they fail to perform. For Elwes the Zizi Show helps 'demistify' AI. Elwes works to deliberately play with what is defined as human like AI to make the mechanics of the technology visible like when it breaks down or falls apart.


Envisioning the Future with AI

‘Remember to
imagine and craft the
worlds you cannot live
without, just as
dismantle the ones
we cannot live in’

-Ruha Benjamin

ENVISIONING THE FUTURE

2023 was the year of AI Hype. Phrases like “AI is the new electricity” circulated inciting anxiety and excitement about how AI could or would change how we work. But rather than taking this idea at face value, the monied interest behind the creation of AI requires we critically read the narrative and communication strategy behind such bombastic claims. Sam Altman famously said that AI will kill us all in front of the United States Senate in the spring of 2023. But what is this larger-than-life CEO pedaling? AI does not have a mind of its own. It is managed and deployed at the discretion of humans. Stating AI will kill us all is both shocking and creates an inflated image of the technology. There are many weapons of destruction. They are not dangerous because they exist. It’s the politics and interests of the people behind them that make them dangerous.



THE TERMINATOR, 1984

ENVISIONING THE FUTURE

Consulted for OpenAI 40 hours / week then terminated -- Unemployment Benefits Question

██████████ · Mar 01 · ██████████

Posted in: Advice

Hi there - I worked as a consultant for a consulting firm that worked for OpenAI. We contributed directly to their algorithm (they hide the number of employees and people actually working on the product this way).

I wanted to know if I am eligible for unemployment benefits after working 40 hours / week and they controlled hours, screen time etc.

I also wanted to see if I have an equity claim.

██████████ · Mar 01 · ██████████

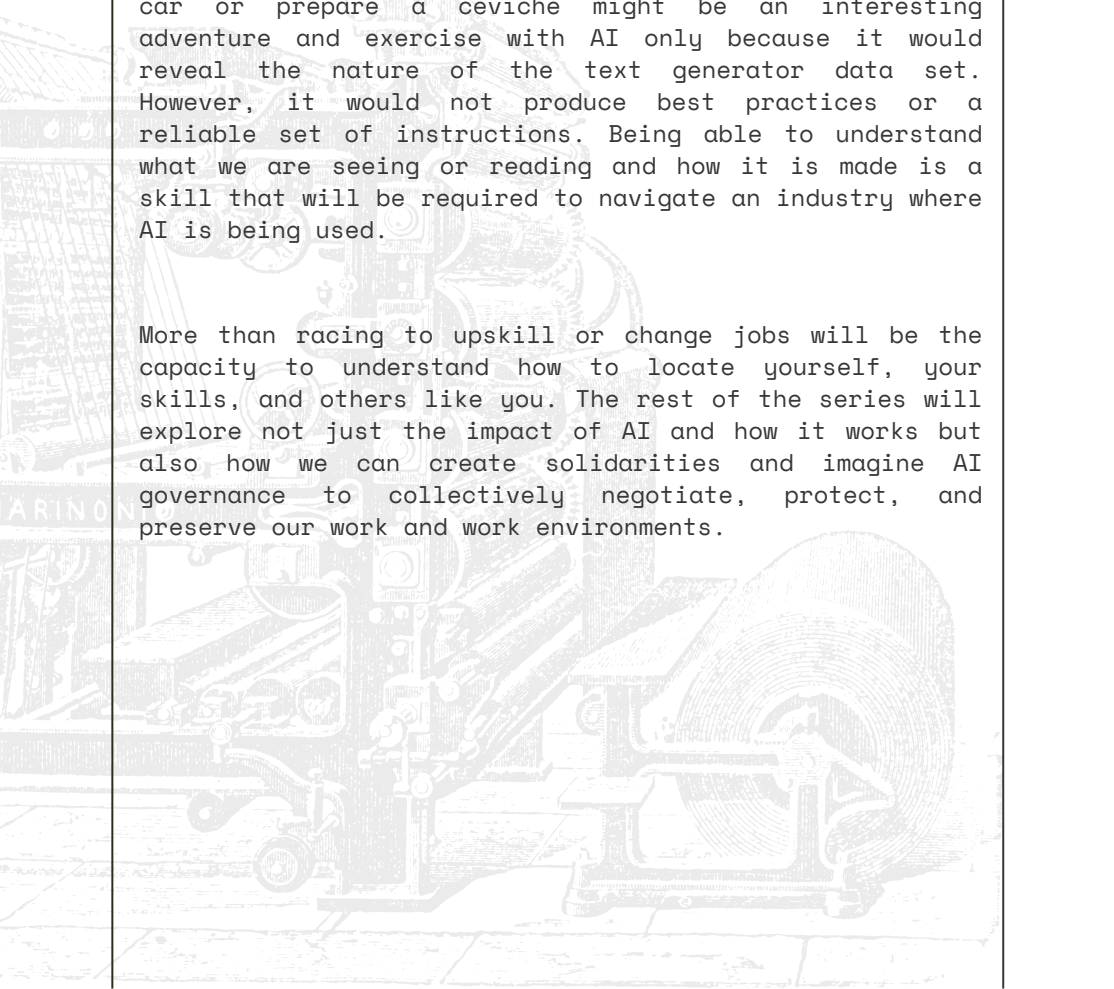
Not a lawyer, but I have experience hiring consulting firms as a tech company. Any client that hires consulting firms has a clause in the agreement assigning the intellectual property to the tech company (the client). This is very normal, and the fact that they can control your hours and cancel the contract is the exact point — it's more flexible than having salaried employees when you are able to reduce the hours for contractors instead of firing an employee. If they didn't have the option to fire the consulting firm at any time, it would be less appealing and they would have just hired employees. So by design, you should have zero equity claim against OpenAI and there are no unemployment benefits paid to your consulting firm as a result of OpenAI having less of a need for consulting work

more

ENVISIONING THE FUTURE

With the constant influx of new uses for AI, information and AI literacy are important skills to have and hone. Text generators like Chat GPT appear all-knowing, but it's important to understand that it does not pull from the most current information. Chat GPT takes from the text input by the user and references a static data set updated by people. So Chat GPT does not perform what a search engine like Duck Duck Go or Google does, which shows the most current information based on a search. It also does not filter information. So trying to fix your car or prepare a ceviche might be an interesting adventure and exercise with AI only because it would reveal the nature of the text generator data set. However, it would not produce best practices or a reliable set of instructions. Being able to understand what we are seeing or reading and how it is made is a skill that will be required to navigate an industry where AI is being used.

More than racing to upskill or change jobs will be the capacity to understand how to locate yourself, your skills, and others like you. The rest of the series will explore not just the impact of AI and how it works but also how we can create solidarities and imagine AI governance to collectively negotiate, protect, and preserve our work and work environments.



REFERENCES

Chaosmotics. (2023, May 27). Machine decision is not final - i. Chaosmotics.

<https://www.chaosmotics.com/en/variations/machine-decision-is-not-final-i>

Chow, A. R. (2023, September 7). Linda Dounia Rebeiz.

Time. <https://time.com/collection/time100-ai/6309451/linda-dounia-rebeiz/>

Crawford, K., Paglen, T. Excavating AI. (n.d.). -.

<https://excavating.ai/>

Hopewell, J. (2023, March 21). Variety.

<https://variety.com/2023/tv/global/red-skies-merav-shaham-a-i-1235558969/>

Ina. (2023, July 26). The history of Chatbots - From

ELIZA to ChatGPT. AI-chatbot Software for Complex Requirements. <https://onlim.com/en/the-history-of-chatbots/>

Jiang, H., Brown, L. T., Cheng, J., Khan, M., Gupta, A., Workman, D., Hanna, A., Flowers, J., & Gebru, T. (2023).

AI art and its impact on artists. AIES '23: Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society. <https://doi.org/10.1145/3600211.3604681>

La Grange, A. (2013). John Berger, Ways of Seeing. In Routledge eBooks (pp. 16-29).

<https://doi.org/10.4324/9780080468389-8>

Nicoletti, L., Bass, D. Bloomberg - Are you a robot?

(n.d.). <https://www.bloomberg.com/graphics/2023-generative-ai-bias/>

REFERENCES

Perrigo, B. (2023, January 18). Exclusive: OpenAI used Kenyan workers on less than \$2 per hour to make ChatGPT less toxic. TIME. <https://time.com/6247678/openai-chatgpt-kenya-workers/>

Reporter, G. S. (2017, December 2). Portrait of the artist as a cheat. The Guardian. <https://www.theguardian.com/theobserver/2000/feb/06/focus.news>

Sharf, Z. (2023, June 21). Variety. <https://variety.com/2023/tv/news/secret-invasion-artificial-intelligence-credits-marvel-1235650643/>

The Museum of Modern Art New York. (2023, March 15). AI Art: How artists are using and confronting machine learning | HOW TO SEE LIKE A MACHINE [Video]. YouTube. <https://www.youtube.com/watch?v=G2XdZIC3AM8>

Vox. (2023, April 4). Why AI art struggles with hands [Video]. YouTube. <https://www.youtube.com/watch?v=24yjRbBah3w>

Wikipedia contributors. (2024, March 5). AlphaGo. Wikipedia. <https://en.wikipedia.org/wiki/AlphaGo>