Beauty and the Prompt

Joanna E. Kabala, MFA, MSc
Eindhoven, The Netherlands
vangdansk.com/kabala.html
joanna.kabala@vangdansk.com

Abstract

This paper presents observations made during image generation using text-to-image and image-to-image prompts in several commercially available Generative AI Art generators. Images from various sessions have been compiled into a short animation that is part of an ongoing art practice, conceptually touching on the issue of consciousness. Without delving into philosophy, consciousness is viewed as the essential feature differentiating humans from machines. Thanks to consciousness, the domain of art belongs to humans. AI, even if it behaves intelligently, is perceived merely as a tool that artists can efficiently use to create content. The observations have been made by a visual artist who, years ago, was involved in an experimental generative art design-research project and created some simple graphical shapes for a system run by genetic algorithms.

The main impulse to investigate Generative AI Art generators and report on them, arose from the avalanche of media headlines proclaiming the end of art, due to the advent of the AI era. Many content creators feel uncertain if the threat is real. Will Generative AI eliminate artists?

1. Introduction

The flow of media messages has been less intense in the last few months. Still, many people are concerned with the frightening vision of the complete replacement of humans by artificial intelligence in some professions. Is artificial intelligence also a threat to artists and designers? Is imagery generated by AI intelligent? Is it art?

We are now living in a *post-ChatGPT* era, a time marked by the accelerated spread of awareness about generative tools, which until recently were known mainly to specialists. Millions of people now use AI editors (*) daily. Some claim they no longer need artists or writers. Many artists feel anxiety when they see images that once required teams of designers and months of rendering time, now generated within seconds, three or four variants at a time.

It seems it took the entire lifetime of Salvador Dalí [1], or some longer effort even for an artist as fast as Picasso [2], to create what renders now in front of our eyes in no time.

In addition, it is loudly advertised as free, although it is not because creating a work of art requires many more than just a few random sample renders. But is there anything to fear, or should artists feel

blessed by the opportunities brought by the broad access to generative tools?

2. Text-to-Image

The initial experience with Generative AI Art generators has been intentionally simplified. Prompted traditional themes like landscapes and flower fields have been captivating and even slightly breathtaking, evoking afterimages of beautiful dreams. While beautiful, they remained impersonal, offering too few cues to suggest any human-like intelligence behind their creation.

It is also possible that the textual prompt, which needs to be given to the artificial intelligence to start rendering, posed a challenge to a visual artist confronted for the first time with the task of creating images solely with words. Imagery creation via natural language is unnatural to visual artists who create forms and compositions organically, *while* observing them on a specified plane and only name the images with full sentences *after* they come into existence — *if* they name them at all, as many visual artworks remain untitled.

Writing a sentence to express a concept should not be a difficult task, as concepts usually precede visualizations. And if the urge to create is an impulse and not a concrete thought, the natural language feature of AI editors may expand artists' ability to articulate ideas, requiring them to put visions into words, rather than loosely thinking about them. Thus, it is already a useful tool that enhances art. Other features may seem overwhelming, especially the unreal beauty and astonishing speed at which AI generates images. It's easy to think that humans can't match this level of performance. The visual quality is exceptionally high and often immediately recognizable as AI renderings. These images primarily differ in their surreal use of color. It seems that AI employs a range that humans can't even perceive. While humans may not notice the difference between shades that lie next to each other on a scale of millions of colors, artificial intelligence can differentiate and efficiently utilize them.

Nevertheless, this is not the end of human art. This undoubtfully helpful tool does not essentially change the creative process, it speeds it up and enhances it. But what follows is a day as usual, as every author knows well. It is sometimes satisfactory but could be exhausting and may not yield the desired results. Long rendering sessions are needed to obtain a meaningful image, especially when aiming for a sequence of images. Sometimes corrections by hand are necessary as these intelligent generators still lack some accuracy. The most striking are anatomical errors that AI makes. Especially, if prompts did not include any request to generate defective figures. Of course, such random software behavior is not surprising for generative designers.

What may surprise in the whole context is that people may believe the view spread by the media, which has already broadly announced AI as a sentient being that even possesses free will. We could hear such declarations about on-screen software. Also, we could see demonstrations of robots, such as Sophia [3], whose successful attempt to obtain citizenship made a huge impression in 2017. However, it is disputable if this was

a result of *free will*.
And if AI is so intelligent and free in its creations, why does it repeatedly make certain errors? Characters rendered by this artificial competitor to the human author are often anatomically incorrect, it seems as if they were created by a student, not a very talented one. Some body parts are let go as if they couldn't be drawn accurately. It is not only human anatomy rendered by AI that is often neglected. Other forms, objects, or architectures are frequently incomplete, illogical, or, in other words, not well-designed. Should they be accepted as, for example, futuristic or cubistic even if such a style was not included in the prompt?

Another attempt was made to communicate with the Generative AI Art generator through natural language. This time, the prompted image was supposed to be a portrait of a woman with red hair sitting at a wooden table (*ill. 1*), in Andy Warhol's style [4]. When writing the prompt, the goal was to create a graphical, almost stencil character with a hint of Marilyn Monroe. After some trials, such a character appeared. Or at least a character close to the original imagination, though it is certainly not Warhol's colorful graphical Marilyn, and she does not sit at the table. Rewriting the prompt did not help. All important details, such as a wooden table, were completely ignored. Additionally, the composition is tight, resembling a talking head sitting in a chair on TV, with more chairs in the frame than needed. Nevertheless, it was accepted.
If AI interprets the human vision written in a sentence this way, a closer look can be taken, just out of curiosity, even if it's not exactly what was envisioned.

Illustration 1a
Mary is OK - AI imagery composed into a short animation. Without AI, this image would not exist; without the artist's mind, it would not be Mary.

Illustration 1b
Mary is OK - set of variants. Some images were edited by hand to obtain a clear stencil style.

There is more to this artwork than the struggle for adequate visualization [5]. An important aspect of this work was a specific scent, intended to evoke a thought that artificial intelligence has no sense of smell. AI cannot smell or taste and will never experience such sensations because it is not a biological process. Moreover, biology alone is not what makes us human. Without consciousness, we would not create and

lead discussions about creativity. As for consciousness, people cannot implement into machines such phenomena that they do not understand themselves. To make AI human-like, their creators (software developers) would first need to understand who humans are. While such an understanding is emerging nowadays [6], the way to make it common is long. When it becomes common, it will be obvious that thinking about and making efforts to imitate humans makes no sense.

Another issue that goes beyond the topic of this work is how artificial intelligence may be used against humans. This paper does not neglect the real threats; it just stays in the domains of art and design.

3. Image-to-Image

Months of rapid development in Generative AI Art generators have passed, yet, the shape of characters' hands and feet still needs corrections — or perhaps it should be accepted as a *signum temporis*.

The common opinion is that AI creates. But is this truly the case? AI searches the data space, makes associations (some may be considered intelligent), and then renders images, fully blown, hyper-realistic, or any other known style if required. So, in a way, it does create, although the result is a kind of *artificial* art. The results of AI creation processes are chaotic, by the very nature of generative algorithms. The first images in the series are random. Some are very beautiful, but their randomness often reveals a lack of understanding on the part of the tool that created them. During longer rendering sessions, an increase in chaos can be observed. With more specific prompts, the generators quickly get lost and fill the images with unnecessary details that are difficult to remove. The generator needs a reset to resign from a specific style or abandon a particular object. Upon reset, all the previous work is gone. Repeating the prompt will not recall the former stage, it will be a whole new journey.

Chaos arises not only due to the nature of generative algorithms. Deliberating on human-like creation, it is important to note that artificial intelligence creates chaos because it "looks" but does not see. It multiplies the variants mathematically and recognizes the images only as far as it was taught to read data but does not see their meaning. Extracting semantic or figurative meaning from images or giving such meaning to images is not a mathematical process. Will this ever change? Although it cannot be ruled out, it remains a fantasy that Consciousness would one day inspire AI. For now, semantic and metaphoric understanding of natural language remains in the human domain.

Meaningful images can be generated more efficiently with image-to-image generators than with text-to-image generators, although in most cases, these two methods work together. Often, the generative process must start with a written prompt, after which the option to use a reference image may become available. Some generators offer the option to use two reference images, one for composition and another for style. This enhances creative possibilities and provides some level of control over the renderings.

Illustration 2a
Love Victory = √1%
On the left: the initial reference image
On the right: AI rendering, edited

The main theme of this work is a woman showing the sign of victory, the sign of freedom. She is the symbol of peace, and her name is Love Victory. The Jackdaw-Phoenix flying next to her represents intelligence and social skills. Let the woman and the city jackdaw become a symbol of the interaction of an intelligent society, capable of calming itself internally and radiating peace externally. (**)

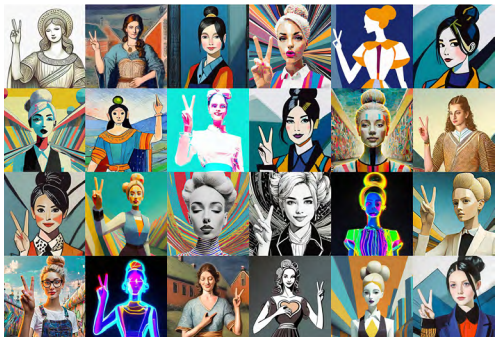

Illustration 2b. Love Victory = √1%
24 snapshots from the animation.

Illustration 2a shows the initial hand drawing, made with a digital pen, and one of more than 500 renderings that followed. Throughout the process, some AI renderings were also used as reference images. The selected set of final images has been composed in a short animation according to the author's storyline.

4. Conclusions

Generative AI Art generators mark a new epoch in human creativity. Nevertheless, those currently available, especially text-to-image generators, are merely filters comparable to tools like the Photoshop Filter tool. Creating artworks according to concepts conceived by the human mind often requires broader toolsets. Current generators do not allow for control over scenes, even in basic aspects such as the number and placement of forms and characters in the composition. It is all chaos, all random, as befits generative tools. Yet, to maintain the meaning of the author's storyline, a certain level of control is useful. This can to a certain degree, be achieved with image-to-image generators - but this is not the end of human art. Contrarily, this is the dawn of new types of tools. And as for tools, these are intelligent indeed. They allow individual artists to produce artworks that previously required whole teams and studios.

Broad access to a quickly evolving variety of Generative AI Art tools is good news in the domain of visual arts. Although artists should not expect that artificial intelligence understands poetry, they may efficiently use AI in creating poetical and symbolic imagery.

As for exceeding humans, Generative AI Art generators still have a long way to go to match what they advertise. Despite unreal beauty and astonishing rendering speed, there exists a problem with the datasets that AI uses. Even the largest ones are finite and limited. Often, they also infringe on copyrights. Additionally, they seem to be governed by semantic coincidence. This observation comes from the point of view of an artist, who is not satisfied to see *anything* resembling a piece of art. The artist wants to see something pleasing to the human mind and satisfactory to consciousness.

We say that AI creates, but it merely generates and renders. It pulls data randomly and the answers are not very intelligent when it lacks information. In this respect, it resembles a human being a bit indeed. However, such productions can hardly be called art because they contain no conscious or subconscious intention. The effort needed to control or at least direct the generative process according to the artist's vision is comparable to the effort required to visualize a concept in any other visual content creation software. The difference lies in the quantity and quality of the images.

Generative AI Art generators are brilliant tools for image creation. However, do not fear that generative artificial intelligence will replace humans in art, because neither humans nor their arts are mathematical processes, as AI is. Additionally, biology is not the only barrier between humans and artificial intelligence. Human art is not a product of the brain alone. The brain controls the arm, which paints, but notice that some disabled people have developed the capability to paint with their feet or mouth. This was possible because the brain controls the body. However, the impulse to create comes from the mind, not the brain. Artists create thanks to biological impulses and spiritual processes in the brain-mind area. The human mind is the domain of consciousness that is unavailable and not transferable to machines.

Artists create under the influence of (un) consciousness, in which phenomena still largely unrecognized, such as dreams, play a role. Art is uncountable and unpredictable, while AI renderings are relatively predictable, as they always depend on datasets.

To the incomprehensible aspects of human existence, which are not transferable to machines and are essential in the creation process, we can also add semantics. This is another threshold that cannot be crossed by AI, as it is the human minds of the authors that give meaning to art. From turning on devices to providing prompts, to selecting or rejecting AI products, AI would create nothing without humans.

Art exists only thanks to the artists' will. It exists only when it is perceived by consciousness. Artificial Intelligence does not perceive; it reads data. It is said that AI creates, but it only renders data supplied to it. It merely imitates. Artists may appreciate AI renderings and incorporate them into the creative process. Yet, conscious human effort is needed to obtain specific images in a particular style. AI art editors are just tools, and they remain so. They are not and will not be independent entities that could replace humans in the field of art. Not today, not tomorrow, not even in our lifetime, and possibly never, will we be able to create *artificial consciousness*.

Notes

(*) ChatGPT was launched at the end of 2022. Since then, many AI chatbots have been created. There are also many Generative AI Art generators, such as Leonardo, DALL-E, Midjourney, Stable Diffusion, Firefly, Runway, LeiaPix, and many others, with new ones appearing continuously.

(**) The Generative AI animation titled Love Victory = √1% conceptually addresses the notion of consciousness, which has started to be acknowledged as a “force” even by mainstream physics. In 2012, Peter Higgs discovered the God particle, and recently, Sir Roger Penrose, referring to the issue of consciousness, admitted that for the universe to exist, a process that cannot be described mathematically is essential. Speaking colloquially, miracles happen in the bosonic quantum field. Additionally, their elusive nature sometimes brings tangible results.

For example, scientists have researched the so-called Extended Maharishi Effect for 43 years, since 1974. They have collected enough data to conclude that it only takes a small group of people to initiate a change in social consciousness. They counted it carefully: the square root of one percent is enough. When a group of people the size of the square root of 1% of a given community experiences pure awareness through extended transcendental meditation, that is, in simple words, calms down internally, there appears a greater cohesion in the collective consciousness, contributing to conflict resolution, such that the number of recorded incidents drops. The research was carried out in various places on Earth. A decrease in incidents and even a temporary easing of armed conflicts occurred near meditating groups, in areas equal to districts of larger cities [7].

The square root of one percent isn’t much. For example, in a community of 2,500 people, just 5 people are enough to calm down this community, at least temporarily. A sustained action would require a sustained group effort, which is not feasible. But this thought and the scientifically confirmed fact seemed worthy of creating a symbolic drawing and editing the results of generative rendering sessions into an animation.

Love Victory = √1% is available on YouTube, https://youtu.be/MX1XyvyQxW The contextual article is available on LinkedIn, https://bit.ly/3YlHbe2.

References

[1] Robert Descharnes, Gilles Néret. Dalí. The Paintings, Benedikt Taschen Verlag 1997.

[2] Henri-Georges Clouzot. The Mystery of Picasso, 1956. A fragment of the movie with time-lapsed, “generative” images is available on YouTube http://bit.ly/4g2JcDo

[3] Sophia, a human-like robot https://www.hansonrobotics.com/sophia/

[4] Andy Warhol, Marilyn Monroe, 1967 https://www.moma.org/collection/works/portfolios/61240

[5] Joanna E. Kabala. An image spoke to me: Mary is OK. October 2023, available on LinkedIn https://bit.ly/4eVS1hJ

[6] Faggin, F. Irreducible, Essentia Books, an imprint of John Hunt Publishing, 2024.

[7] David W. Orme-Johnson and Lee Fergusson. Global Impact of the Maharishi Effect from 1974 to 2017: Theory and Research, Journal of Maharishi Vedic Research Institute, 2018