


AN-ICON



On the altered  
states of machine  
vision. Trevor Paglen, Hito  
Steyerl, Grégory Chatonsky

by Antonio Somaini Machine learning

Digital images

Paglen

Steyerl

Chatonsky

AN-CON Studies in  
Environmental  
images

Issue No1 Year 2022

→ Altered states

Edited by Giancarlo Grossi  
and Andrea Pinotti

# On the altered states of machine vision.

Trevor Paglen, Hito Steyerl,  
Grégory Chatonsky



by Antonio Somaini

## Abstract

The landscape of contemporary visual culture and contemporary artistic practices is currently undergoing profound transformations caused by the application of technologies of machine learning to the vast domain of networked digital images. The impact of such technologies is so profound that it leads us to raise the very question of *what we mean by “vision” and “image” in the age of artificial intelligence*. This paper will focus on the work of three artists – Trevor Paglen, Hito Steyerl, Grégory Chatonsky – who have recently employed technologies of machine learning in non-standard ways. Rather than using them to train systems of machine vision with their different operations (face and emotion recognition, object and movement detection, etc.) and their different fields of application (surveillance, policing, process control, driverless vehicle guidance, etc.), they have used them in order to produce entirely new images, never seen before, that they present as *altered states* of the machine itself.

[Machine learning](#)

[Digital images](#)

[Paglen](#)

[Steyerl](#)

[Chatonsky](#)

To quote this essay: A. Somaini, “On the altered states of machine vision. Trevor Paglen, Hito Steyerl, Grégory Chatonsky”, *AN-ICON. Studies in Environmental Images*, no. 1 (2022): 91-111

The landscape of contemporary visual culture and contemporary artistic practices is currently undergoing profound transformations caused by the application of technologies of machine learning – one of the areas of so-called “artificial intelligence” – to the vast domain of networked digital images. Three strictly interrelated phenomena, in particular, are producing a real tectonic shift in the contemporary iconosphere, introducing new ways of “seeing” and new “images” – we’ll return later to the meaning of these quotation marks – that extend and reorganize the field of the visible, while redrawing the borders between what can and what cannot be seen.

These three strictly interrelated phenomena are:

- the new technologies of *machine vision* fuelled by processes of machine learning such as the Generative Adversarial Networks (GAN);
- the ever-growing presence on the internet of trillions of networked digital images that are *machine-readable*, in the sense that they can be seen and analyzed by such technologies of machine vision;
- the entirely *new images* that the processes of machine learning may generate.

Considered from the perspective of the *longue durée* of the history of images and visual media, the appearance of these three phenomena raises a whole series of aesthetic, epistemological, ontological, historical and political questions. Their impact onto contemporary visual culture is so profound that it leads us to raise the very question of *what we mean by “vision” and “image” in the age of artificial intelligence*.

What is “seeing” when the human psycho-physiological process of vision is reduced, in the case of machine vision technologies, to entirely automated operations of pattern recognition and labelling, and when the various applications of such operations (face and emotion recognition, object and motion detection) may be deployed across an extremely vast visual field (all the still and moving images accessible online) that no human eye could ever attain? In speaking of “*machine vision*”, are we using an anthropomorphic

term that we should discard in favor of a different set of technical terms, specifically related to the field of computer science and data analysis, that bear no connection with the physiological and psychological dynamics of human vision? Artists-researchers such as Francis Hunger and scholars such as Andreas Broeckmann (with his notion of “optical calculus” as “an unthinking, mindless mechanism, a calculation based on optically derived input data, abstracted into calculable values, which can become part of computational procedures and operations”),<sup>1</sup> Adrian MacKenzie and Anna Munster (with their ideas of a “platform seeing” operating onto “image ensembles” through an “invisual perception”),<sup>2</sup> Fabian Offert and Peter Bell (with their idea according to which the “perceptual topology” of machines is irreconcilable with human perception)<sup>3</sup> have argued for the necessity of moving beyond anthropocentric frameworks and terms, highlighting the fact that machine vision poses a real challenge for the humanities.

Can we still use the term “image” for a digital file, encoded in some image format,<sup>4</sup> that is machine-readable even when it is not visible by human eyes, or that becomes visible on a screen as a pattern of pixels only for a small fraction of time, spending the rest of its indefinite lifespan circulating across invisible digital networks? Can concepts such as that of “iconic difference” [*ikonische Differenz*],<sup>5</sup> which highlights the fundamental perceptual difference between an image and its surroundings (its “*hors champ*”) be still applied to machine-readable images?

And what is the status of the entirely new images produced by processes of machine learning? These

---

1 A. Broeckmann, “Optical Calculus”, paper presented November 2020 at the conference *Images Beyond Control*, FAMU, Prague, video, 5:01, <https://www.youtube.com/watch?v=FnAgBblnMfA>.

2 A. MacKenzie, A. Munster, “Platform Seeing: Image Ensembles and Their Invisualities”, *Theory, Culture & Society* 36, no. 5 (2019): 3-22.

3 F. Offert, P. Bell, “Perceptual bias and technical metapictures: critical machine vision as a humanities challenge”, *AI & Society* (2020), <https://link.springer.com/article/10.1007/s00146-020-01058-z>.

4 On the theory of formats, see M. Jancovic, A. Schneider, A. Volmar, eds., *Format Matters: Standards, Practices, and Politics in Media Cultures* (Lüneburg: Meson Press, 2019).

5 On the notion of “iconic difference”, see G. Boehm, “Ikonische Differenz”, *Rheinsprung 11. Zeitschrift für Bildkritik* 1 (2011): 170-176.

are images that are not produced through some traditional form of lens-based analog or digital optical recording, nor through traditional computer-generated imagery (CGI) systems, but rather through processes belonging to the wide realm of artificial intelligence that either *transform* existing images in ways that were impossible until a few years ago, or *create* entirely new images, never seen before. What do such images represent, what kind of agency do they have, how do they mediate our visual relation to the past, the present, and the future? And why have such new images generated by processes of machine learning been so often considered, both in popular culture and in the work of contemporary artists, to be the product of some kind of *altered state* – a “dream”, a “hallucination”, a “vision”, an “artificial imagination” – of the machine itself?

Before we analyse the way in which this last question is raised, in different ways, by popular computer programs such as Google’s DeepDream (whose initial name echoed Christopher Nolan’s 2010 film “*Inception*”), and by the recent work of contemporary artists and theorists such as Trevor Paglen, Hito Steyerl and Grégory Chatonsky, let us begin with a quick overview of the current state of machine vision technologies, with their operations and applications, and of the new images produced by processes of machine learning that are increasingly appearing throughout contemporary visual culture.

### **The impact of machine learning technologies onto contemporary visual culture**

First tested in the late 1950s, with image recognition machines such as the Perceptron (developed at the Cornell Aeronautical Laboratory by Frank Rosenblatt in 1957), and then developed during the 1960s and 1970s as a way of imitating the human visual system in order to endow robots with intelligent behavior, machine vision technologies have entered a new phase, in recent years, with the development of machine learning processes,

and with the possibility of using immense image databases accessible online as both training sets and fields of application.<sup>6</sup> If in the 1960s and 1970s the goal was mainly to extract three-dimensional structures from images through the localization of edges, the labelling of lines, the detection of shapes and the modelling of volumes through feature extraction techniques such as the Hough transform (invented by Richard Duda and Peter Hart in 1972, on the basis of a 1962 patent by Paul Hough), the recent development of machine learning techniques and the use of vast image training sets organized according to precise taxonomies – such as ImageNet, in which 14 millions of images are organized according to 21,000 categories derived from the WordNet hierarchy (a large lexical database of English nouns, verbs, adjectives and adverbs)<sup>7</sup> – have allowed a rapid increase in the precision of all the operations of machine vision.

Among such operations, we find pixel counting; segmenting, sorting, and thresholding; feature, edge, and depth detection; pattern recognition and discrimination; object detection, tracking, and measurement; motion capture; color analysis; optical character recognition (this last operation allowing for the reading of words and texts within images, extending the act of machine “seeing” to a form of machine “reading”).

For a few years now, such operations have been applied to the immense field of machine-readable images. A field whose dimensions may be imagined only if we understand that *any networked digital image* – whether produced through some kind of optical recording, or entirely computer-generated, or a mix of the two, as it is often the case – may be analysed by machine vision technologies based on processes of machine learning such as the Generative Adversarial Networks (GAN).<sup>8</sup> Starting from vast training sets containing images similar to the

---

6 For a general overview of computer vision and computer imagery, with its historical developments, see S. Arcagni, *L'occhio della macchina* (Turin: Einaudi, 2018).

7 “ImageNet”, accessed November 3, 2021, <http://www.image-net.org/>.

8 I. Goodfellow *et al.*, “Generative adversarial nets”, *Advances in neural information processing systems* (2014): 2672-2680.

ones the system has to learn to identify, and feeding such training sets into an ensemble of two adversarial neural networks that act as a Generator and a Discriminator that are in competition against one another, the GAN-based machine vision systems have gradually become more and more precise in performing their tasks. All the main smart phone producers have equipped their devices with cameras and image processing technologies that turn every photo we take into a machine-readable image, and internet giants such as Google and Facebook, as well as a host of other companies, have developed machine vision and face recognition systems capable of analysing the immense quantity of fixed and moving images that exist on the internet and that continue to be uploaded every day, raising all sorts of ethical and political issues and highlighting the need for a broader legal framework that for the moment is largely missing.<sup>9</sup>

Considered together, such machine vision systems are turning the contemporary digital iconosphere and the entire array of contemporary screens, with their various dimensions and degrees of definition,<sup>10</sup> into a vast field for data mining and data aggregation. A field in which faces, bodies, gestures, expressions, emotions, objects, movements, places, atmospheres and moods – but also voices and sounds, through technologies of *machine hearing* – may be identified, labelled, stored, organized, retrieved, and processed as data that can be quickly accessed and activated for a wide variety of purposes: from surveillance to policing, from marketing to advertising, from the monitoring of industrial processes to military operations, from the functioning of driverless vehicles to that of drones and robots, from the study of the inside of the human body through the analysis of medical imaging all the way up to

---

9 As I complete the final revisions of this essay, on 2 November 2021, Facebook just announced its decision to stop using facial-recognition software that could automatically recognize people in photos and videos posted on the social network: a massive shift for a company that is currently trying to reposition itself, also through the new company name Meta, adopted in October 2021.

10 On the aesthetic, epistemological, historical and political implications of the high and low definition of images, see F. Casetti, A. Somaini, eds., *La haute et la basse définition des images. Photographie, cinéma, art contemporain, culture visuelle* (Milan-Udine: Mimesis, 2021).

the study of the surface of the Earth and of climate change through the analysis of satellite images. Even disciplines that might seem to be distant from the most common current applications of machine vision technologies, such as art history and film studies, are beginning to test the possibilities introduced by such an automated gaze, capable of “seeing” and analyzing, according to different criteria, vast quantities of visual reproductions of artworks or vast corpuses of films and videos in an extremely short time.<sup>11</sup>

In order to fully understand the impact of machine learning onto contemporary visual culture, we need to add, to the vast field of machine vision technologies that we just described, the *new images* produced by processes of machine learning – often, the same GAN that are used to train and apply machine vision systems – that either *transform* pre-existing images in ways that were impossible until a few years ago, or *create* entirely new images, never seen before.

In the first case, we are referring to processes of machine learning capable of *transforming* existing images that can have very different applications: producing 3D models of objects from 2D images; changing photographs of human faces in order to show how an individual’s appearance might change with age (as with the app Face-App) or by merging a face with another face (Faceswap);<sup>12</sup> animating in a highly realistic way the old photograph of a deceased person (Deep Nostalgia, developed by My-Heritage);<sup>13</sup> creating street maps from satellite imagery;<sup>14</sup> taking any given video, and “upscaling” it, by increasing its frame rate and its definition. An emblematic example

---

11 See for example the various experiments being developed at the Google Arts & Culture Lab: “Google Arts & Culture”, accessed November 3, 2021, <https://artsandculture.google.com/>, or the way in which the EYE Film Museum in Amsterdam is testing new ways of accessing its collections through a program fuelled by machine vision systems: “Jan Bot”, accessed December 2, 2021, <https://www.jan.bot/>.

12 The website of Faceswap, which announces itself as “the leading free and OpenSource multiplatform Deepfakes software”: “Faceswap”, accessed November 3, 2021, <https://faceswap.dev/>.

13 “Deep Nostalgia”, accessed November 3, 2021, <https://www.myheritage.fr/deep-nostalgia>.

14 R. Matheson, “Using artificial intelligence to enrich digital maps. Model tags road features based on satellite images, to improve GPS navigation in places with limited map data”, *MIT News*, January 23, 2020, <https://news.mit.edu/2020/artificial-intelligence-digital-maps-0123>.

of this last application, which may alter significantly our experience of visual documents of the past, would be the videos realized by Denis Shiryayev<sup>15</sup> in which, through a process of machine learning, a Lumière film such as *L'Arrivée d'un train en gare de La Ciotat* (1896) is transformed from the original 16 frames per second to 60 frames per second, from the original 1.33:1 format to a contemporary 16:9 format, and from the original, grainy 35mm analog film to a 4K digital resolution.<sup>16</sup> In other examples of images transformed by machine learning, the transformations are much more radical, as it happens with the so-called “deepfakes”: videos that use neural networks in order to manipulate the images and the sounds of pre-existing videos – in some cases a single, pre-existing image – producing new videos that have a high potential to deceive. Among the many examples that can now be found across the internet in different domains such as pornography, politics and social media, pornographic videos in which faces of celebrities are swapped onto the bodies of porn actors, a TikTok account with a whole series of odd videos by a “Deep Tom Cruise”,<sup>17</sup> or speeches by public figures such as Barack Obama<sup>18</sup> and Queen Elizabeth<sup>19</sup> whose content has been completely altered in such a way that the movements of their mouths perfectly match the new, invented words they are uttering. And among the applications of deepfakes in the musical realm, the “new” songs by long deceased singers, whose style and voice are reproduced in a highly realistic way by applications of machine learning such as Jukebox, developed by OpenAI<sup>20</sup>: a “resurrecting” function

---

15 “Denis Shiryayev”, Youtube channel, accessed November 3, 2021, <https://www.youtube.com/c/DenisShiryayev/videos>. For an online platform offering video enhancement powered by AI, see: “Neural Love”, accessed December 2, 2021, <https://neural.love/>.

16 The video of the upscaled version of *L'Arrivée d'un train en gare de La Ciotat* (1896) can be found all across YouTube, in various black-and-white and colored version. For a version in color, see: Deoldify videos, “[DeOldified] Arrival on a Train at La Ciotat (The Lumière Brothers, 1896)”, Youtube video, February 4, 2020, <https://www.youtube.com/watch?v=EqbOhqXHL7E>.

17 “Deep Tom Cruise”, TikTok account, accessed November 3, 2021, <https://www.tiktok.com/@deeptomcruise>.

18 BuzzFeed Video, “You Won’t Believe What Obama Says in This Video!”, Youtube video, April 17, 2018, <https://www.youtube.com/watch?v=cQ54GDm1eL0>.

19 Channel 4, “Deepfake Queen: 2020 Alternative Christmas Message”, Youtube video, December 25, 2020, <https://www.youtube.com/watch?v=lvY-Abd2FfM>.

20 “Jukebox”, accessed November 3, 2021, <https://openai.com/blog/jukebox/>.

that in the animated photographs of deceased persons of Deep Nostalgia.

In the second case, the use of machine learning processes leads to the *creation* of entirely new images or sections of images: modelling patterns of motion (for example, crowd motion) in video, thereby leading computer generated imagery, in some cases fuelled by artificial intelligence, to produce new kinds of “contingent motion”;<sup>21</sup> producing highly photorealistic images of objects and environments for different kinds of advertising; inventing perfectly realistic faces of people that actually do not exist through open source softwares such as StyleGan and then make them accessible through projects such as Philip Wang’s *This Person Does Not Exist*.<sup>22</sup>

To these widespread applications of machine learning we may add the hybrid, unprecedented imagery produced by the software DeepDream, created in 2015 by the Google engineer and artist Alexander Mordvintsev: a program that uses Convolutional Neural Networks in order to enhance patterns in pre-existing images, creating a form of algorithmic *pareidolia*, the impression of seeing a figure where there is none, which is here generated by a process which repeatedly detects in a given image patterns and shapes that the machine vision system has been trained to see.<sup>23</sup> The result of such a recursive process, in which every new image is submitted again to the same kind of pattern and shape recognition, are images that recall an entire psychedelic iconography that spans through cinema, photography, the visual arts and even so-called *art brut*: images that are here presented as a sort of dream – a hallucinogenic, psychedelic dream – of the machine itself.

---

21 J. Schonig, “Contingent Motion: Rethinking the ‘Wind in the Trees’ in Early Cinema and CGI”, *Discourse: Journal for Theoretical Studies in Media and Culture* 40, no. 1 (2018): 30-61.

22 “This Person does not exist”, accessed November 3, 2021, <https://thispersondoesnotexist.com/>.

23 The program is now open source: see “Deep Dream Generator”, accessed November 3, 2021, <https://deepdreamgenerator.com/>. See also: “Alexander Mordvintsev”, accessed November 3, 2021, <https://znah.net/>.



Fig. 1




Fig. 2



Fig. 3  
The original image (Fig. 1) has been modified by applying ten (Fig. 2) and then fifty (Fig. 3) iterations of the software DeepDream, the network having been trained to perceive dogs.

## Exploring the “altered states” of machine vision through Generative Adversarial Networks

The idea that lies at the basis of the DeepDream software – the idea that machine learning technologies, and in particular the Convolutional Neural Networks (CNN) and Generative Adversarial Networks (GAN), may be used in order to explore and reveal *the altered states of machine vision* – can also be found in the recent works of artists (often active also as theorists) such as Trevor Paglen, Hito Steyerl, and Grégory Chatonsky.<sup>24</sup> By insisting on the “creative”, image-producing potential of Generative Adversarial neural networks – rather than on their standard application for the training of machine vision systems – the three of them explore a vast field of images that they consider to be “hallucinations”, “visions of the future”, or the product of an “artificial imagination”, characterized by a new form of realism, a “disrealism”.

Trevor Paglen’s entire work as an artist and theorist, since 2016, has been dedicated to the attempt of understanding and visualizing the principles that lie at the basis of machine vision technology. Through texts (sometimes written in collaboration with the researcher Kate Crawford),<sup>25</sup> exhibitions, performances, and works made of still and moving images, Paglen has tried to highlight not only the social and political biases that are inherent in the way machine vision systems are structured, but also the way in which such systems diverge profoundly from human vision.<sup>26</sup>

In an article published in December 2016 in *The New Inquiry* with the title “Invisible Images (Your Pictures Are Looking at You)”,<sup>27</sup> Paglen discusses the new challenges that arise in a context in which “sight” itself has become machine-operated and separated by human eyes,

---

24 Among the first artists who started working with GANs, one should remember Helena Sarin and Mike Tyka. See “Helena Sarin”, AI Artist, accessed November 3, 2021 <https://aiartists.org/helena-sarin>, and “Mike Tyka”, accessed November 3, 2021, <https://miketyka.com/>.

25 K. Crawford, *Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence* (New Haven: Yale University Press, 2021).

26 See: “Trevor Paglen”, accessed November 3, 2021, <https://paglen.studio/>.

27 T. Paglen, “Invisible Images (Your Pictures Are Looking at You)”, *The New Inquiry*, December 8, 2016: <https://thenewinquiry.com/invisible-images-your-pictures-are-looking-at-you/>.

participating in a vast field of image operations. Arguing that what we are currently experiencing is part of a *vast transition from human-seeable to machine-readable images* – a new condition in which “the overwhelming majority of images are now made by machines for other machines, with humans rarely in the loop” – Paglen writes that

if we want to understand the invisible world of machine-machine visual culture, we need to unlearn to see like humans. We need to learn how to see a parallel universe composed of activations, keypoints, eigenfaces, feature transforms, classifiers, training sets, and the like.<sup>28</sup>

*We need to unlearn to see like humans.* But how can we not see like humans, how can we step out of our human point of view? Accomplishing this apparently impossible task – a task which echoes the recurrent philosophical problem of how to step out of one’s own socio-historical position, of one’s own cognitive and emotional framework – has been the main goal of Trevor Paglen’s artistic practice during the last few years, as we can see in a body of works that was initially produced in 2017 through various collaborations with computer vision and artificial intelligence researchers as an artist-in-residence at Stanford University, and was first exhibited at the Metro Pictures Gallery in New York in an exhibition entitled *A Study of Invisible Images* (September 8 – October 21, 2017)<sup>29</sup>, before being presented at various other galleries and museums such as the Osservatorio of the Fondazione Prada in Milan, or the Carnegie Museum of Art in Philadelphia, in an exhibition entitled *Opposing Geometries* (2020).

The works in the exhibition at Metro Pictures present a possible response to the challenge of how to penetrate within systems of machine vision that tend to expel the human gaze from their processes. Among them, we find the attempt to master the machine learning techniques

---

<sup>28</sup> Ibid.

<sup>29</sup> A series of images of the works presented in the exhibition can be found at the following address: “Trevor Paglen. A study of Invisible Images”, Metro Pictures, accessed January 19, 2020, <http://origin.www.metropictures.com/exhibitions/trevor-paglen4>.

that are commonly used for machine vision applications, in order to hack them and lead them to produce entirely new images, never seen before, that may be considered as a form of *hallucination* of machine vision.

This is what happens in a series of still images entitled *Adversarially Evolved Hallucinations*, which Paglen developed through a non-standard application, in three steps, of Generative Adversarial Networks.<sup>30</sup>

The first step consisted in establishing new, original training sets. Instead of using the usual corpuses of images that are used to train machine vision systems in recognizing faces, objects, places and even emotions – corpuses that are often derived by pre-existing and easily available image databases such as the already mentioned ImageNet – Paglen established new training sets composed by images derived from literature, psychoanalysis, political economy, military history, and poetry. Among the various taxonomies he used in order to compose his training sets we find “monsters that have been historically interpreted as allegories of capitalism”, such as vampires, zombies, etc.; “omens and portents”, such as comets, eclipses, etc.; “figures and places that appear in Sigmund Freud’s *The Interpretation of Dreams*”, a corpus which includes various symbols from Freudian psychoanalysis; “Eye-Machines”, a series of images clearly inspired by Harun Farocki’s videoinstallations *Eye-Machine I, II, III* (2001-2003) and containing images of surveillance cameras or of spaces under surveillance; “American Predators”, a corpus containing various predatory animals, plants, and humans indigenous to the United States, mixed with military hardware like predator drones and stealth bombers.

The second step consisted in feeding these unusual training sets into the two neural networks of the GAN system: the Discriminator and the Generator. These two networks begin interacting with one another in an adversarial, competitive way, in such a way that the Discriminator,

---

30 Information on the *Adversarially Evolved Hallucinations* can be found in Trevor Paglen’s website: T. Paglen, “Hallucinations”, accessed November 3, 2021, <https://paglen.studio/2020/04/09/hallucinations/>.

after having received the initial training set, has to evaluate the images that it receives from the Generator, deciding whether they resemble or not to those of the training set. As the process unfolds through reiterated exchanges between the two neural networks, the Discriminator becomes more and more precise and effective in evaluating the images that are submitted to it.

The third step consists in the artist intervening in the process and choosing to extract, at a given moment, one of the images produced by the Generator: an image that emerges from the sequence of the adversarial exchanges, and that is the result of one of the countless attempts by the Generator to test the precision of the Discriminator, trying to fool it. In the case of the series of the *Adversarially Evolved Hallucinations*, all the images selected by Paglen seem to bear some kind of resemblance to the ones contained in the original training sets – even though we cannot really assess the degree of this resemblance, because the training sets are not accessible to us – while displaying at the same time different forms of deviations and aberrations that recall a sort of psychedelic imaginary.

Among Paglen's *Adversarially Evolved Hallucinations*, we find images with titles such as *Vampire (Corpus: Monsters of Capitalism)*, *Comet (Corpus: Omens and Portents)*, *The Great Hall (Corpus: The Interpretation of Dreams)*, *Venus Flytrap (Corpus: American Predators)*, *A Prison Without Gards (Corpus: Eye-Machines)*. In the case of *Vampire (Corpus: Monsters of Capitalism)*,



Fig. 4. Trevor Paglen, *Vampire* (*Corpus: Monsters of Capitalism*), dye sublimation metal print, 2017

the Discriminator was trained on thousands of images of zombies, vampires, Frankensteins, and other ghosts that have been at some point – be it in essays, literary texts, or films – associated with capitalism. Paglen then set the Generator and Discriminator running until they had synthesized a series of images that corresponded at least in part to a specific class within the given corpus. From all of the even slightly acceptable options that the GAN generated, Paglen then selected the one we see in the series exhibited at Metro Pictures as the “finished work”.

There may be multiple reasons behind Trevor Paglen’s decision to call these images “hallucinations”. To begin with, such images recall a type of imagery that we might consider to be “psychedelic” or “surrealist”: in some of them, we definitely see echoes of Max Ernst, or Salvador Dalì. A second reason may lie in the attempt to emphasize the fact that the result of this non-standard application of the processes of machine learning – a process which unfolds within a closed machine-to-machine space, the invisible space of the back-and-forth between the Generator and the Discriminator from which human eyes are excluded – produces images that, just like human hallucinations, have no footing in exterior reality, or may merge in unpredictable

way with shapes and forms stemming from the perception of the outer world. Finally, the term “adversarially evolved hallucinations” may underline the fact that these images are the result of a machine learning process gone astray: a process which has been hacked and led to drift away from its original, standard applications.

The reference to the term “hallucinations”, though, should not be misleading. What Trevor Paglen tries to show us with his *Adversarially Evolved Hallucinations* has really nothing to do with a disruption of the orderly functioning of human consciousness. What they highlight, rather, is the radical *otherness* of machine vision, if compared with human vision. A radical *otherness* based on operations that have nothing to do with human, embodied vision, which we may just try to grasp through the impossible attempt of “unlearning to see like humans”.

We find a different application of images produced by machine learning in *This is the Future*, an installation by Hito Steyerl which was presented at the Venice Biennale in 2019, and which was conceived as an expansion of the exhibition *Power Plants* at the Serpentine Gallery in London the previous year. In the Venice installation, Hito Steyerl arranged onto different platforms a series of nine videos in which one could see images resembling to some kind of “vegetal” time-lapse imagery: flowers quickly blossoming and spreading out, plants and bamboo shoots growing in height and width.




Fig. 5a



Fig. 5b



Fig. 5c  
Figg. 5a, 5b and 5c. Hito Steyerl, *This is the Future: A 100% Accurate Prediction*, stills from the single channel HD video, color, sound, 16', 2019

What interested Steyerl in the use of neural networks in this installation was the *predictive* nature of machine learning, and the status of “visions of the future” of its imagery: the fact that neural predictive algorithms operate through statistical models and predictions based on immense, “big data” databases, and are therefore related to the vast spectrum of predictive systems (be they financial, political, meteorological, environmental, etc.) that are present in contemporary “control societies”, while at the same time being part of the *longue durée* of the history of prediction systems elaborated by human cultures.

The main video presented in Hito Steyerl’s installation, entitled *This is the Future: A 100% Accurate Prediction*, consists of images produced through a collaboration with the programmer Damien Henry, author of a series of videos entitled *A Train Window*<sup>31</sup> in which a machine learning algorithm has been trained to predict the next frame of a video by analyzing samples from the previous image, in such a way that, as in a perfect feedback loop, each output image

---

31 The videos, programmed with Tensorflow, are available at: D. Henry, “A train window”, *Magenta*, October 3rd, 2018, [https://magenta.tensorflow.org/nfp\\_p2p](https://magenta.tensorflow.org/nfp_p2p). I thank both Hito Steyerl and Damien Henry for useful information on the coding used in *This is the Future*.

becomes the input for the next step in the calculation. In this way, *after intentionally choosing or producing only the first image*, all the following ones are generated by the algorithm, without any human intervention. In Hito Steyerl's work, this idea of a video entirely made up of images "predicted" by neural networks and located, as the video says, "0.04 seconds in the future", is presented as a new kind of *documentary* imagery: an imagery that can at the same time *predict* and *document* the future, as paradoxical as this may seem. The video begins with white text on a black background that reads: "These are documentary images of the future. Not about what it will bring, but about what it is made of". The next five sections of the video – *Heja's Garden*, *The Future: A History*, *Bambusa Futuris*, *Power Plants* and *Heja's Prediction* – lead us to a psychedelic landscape of images that morph sample images stemming from categories such as "sea", "fish", "flower", "rose" or "orchid": each of them is produced by a neural network that, as the electronic voice accompanying the video tells us, "can see one fraction of a second into the future".

*Second Earth* (2019) by Grégory Chatonsky takes another route into the iconosphere produced by GAN-driven machine learning. What interests him is the idea of an "artificial imagination", capable of visualizing, through the means of artificial intelligence, "the hallucination of a senseless machine, a monument dedicated to the memory of the extinct human species".<sup>32</sup> Himself in charge of the coding which lies at the base of the various elements and the various media mobilized in his work, Chatonsky works in particular on what he calls the "*chaînage*", the "sequencing" of different artificial intelligence systems that, taken all together, produce a whole cascade of new forms: among them, neural networks capable of generating new texts (read by synthetic voices) starting from some given text databases, or capable of generating images from given texts, and texts from given images,

---

32 See the presentation of *Second Earth*: G. Chatonsky, "Second Earth / Terre Seconde", accessed November 29, 2021, <http://chatonsky.net/earth/>.

with a new kind of AI-powered *ekphrasis*. The metamorphical universe that we see in the videos of *Second Earth*,



Fig. 6a



Fig. 6b



Fig. 6c




Fig. 6d  
Figg. 6a, 6b, 6c and 6d. Grégory Chatonsky, *Second Earth*, stills from one of the videos in the installation, 2019

a work that Chatonsky presents as an “evolving installation”, evokes the idea of a form-generating power that used to be rooted in nature and which is now taken over by machines which are incorporating and re-elaborating the trillions of

images that humans have uploaded on the internet as a sort of hypertrophic memory. The text accompanying the work reads: “Accumulate data. Outsourcing our memories. Feed software with this data so that it produces similar data. Produce realism without reality, become possible. Disappear. Coming back in our absence, like somebody else”.<sup>33</sup>

As products of a “realism without reality”, what Chatonsky calls a “disrealism”,<sup>34</sup> the images produced through Generative Adversarial Networks in *Second Earth* do have a hallucinatory, oneiric, “surrealistic” quality, that bears a strange kind of “family resemblance” to the ones that we find in Paglen’s *Adversarially Evolved Hallucinations*, Steyerl’s *This is the Future*, and in the work of other artists who have recently explored the GAN-generated imagery, such as Pierre Huyghe in his installation *Umwelt* (2018). In Chatonsky’s *Second Earth*, these images do refer to some kind of “outer reality”, but the status of this outer reality is highly unclear.

On the one hand, they bear traces of the images contained in the training sets that have been employed in order to activate the GANs: in the cases of the stills from one of the videos in the installation that are here reproduced, such training sets referred probably to categories such as “birds”, “faces”, “eyes”, etc., and the images contained in the training sets – be they actual photographs, or still from videos – do in most cases refer to some profilmic reality.

On the other, extracted as they are from the “latent space” of a process of machine learning in motion from the pole of absolute noise to the pole of a perfect resemblance to the images of the training set, the images of *Second Earth* refer to another kind of reality, one that does not exist yet. We are here in the domain of “anticipation”, rather than “prediction”,<sup>35</sup> in the perspective of an explora-

---

33 G. Chatonsky, *Second Earth*, <http://chatonsky.net/earth/>.

34 Grégory Chatonsky has begun to use and theorize this term in recent lectures held at the Jeu de Paume and at Campus Condorcet in Paris, in the framework of the lecture cycle *L'esthétique à l'heure du pixel* (September 2021 – May 2022) and the seminar *L'image à l'épreuve des machines. Reconfigurations du visible* (25-26 October 2021). See for example “*Le disréalisme (le pixel perdu de l'espace latent)*”, Jeu de Paume, accessed november 29, 2021, <https://jeudepaume.org/evenement/seminaire-esthetique-pixel-1/>.

35 Unpublished conversation with Grégory Chatonsky, 2019, whom I thank for the useful information on the different software used in *Second Earth*.

tion of a non-human “artificial imagination”, rather than in the denunciation of the pervasive presence of systems of control and surveillance, as it was the case in the work of Hito Steyerl. At the basis of Chatonsky’s “Second Earth”, we find the observation that

the machine was becoming capable of automatically producing a phenomenal quantity of realistic images from the accumulation of data on the Web. This realism is similar to the world we know, but it is not an identical reproduction. Species metamorphose into each other, stones mutate into plants and the ocean shores into unseen organisms. The result: this “second” Earth, a reinvention of our world, produced by a machine that wonders about the nature of its production.

Over fifty years ago, in his seminal *Understanding Media. The Extensions of Man* (1964), Marshall McLuhan formulated the idea that art could become, in some decisive historical moments, a form of “advance knowledge of how to cope with the psychic and social consequences of the next technology”, and added that new art forms might become in these moments “social navigation charts”, helping us find some orientation across a sensorium entirely transformed by new media and new technologies. Today, while we witness the first signs of what promises to be a massive impact of artificial intelligence onto all areas of our psychic, social, and cultural life, the works of artists such as Trevor Paglen, Hito Steyerl and Grégory Chatonsky do appear like “navigation charts”: their exploration of the *altered states* of machine vision through the appropriation and the *détournement* of technologies such as the Generative Adversarial Networks help us better understand the aesthetic, epistemological and political implications of the transformations that such technologies are producing within contemporary visual culture. Taken together, they highlight the fact that what is at stake is the very status of what we mean by “image” and by “vision” in the age of artificial intelligence.

**AN-ICONOLOGY**  
**History, Theory, and Practices**  
**of Environmental Images**



UNIVERSITÀ DEGLI STUDI DI MILANO  
DIPARTIMENTO DI FILOSOFIA  
"PIERO MARTINETTI"



AN-ICON has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme. Grant agreement No. 834033 AN-ICON. The project is hosted by the Department of Philosophy "Piero Martinetti" – Department of Excellence at the State University of Milan.