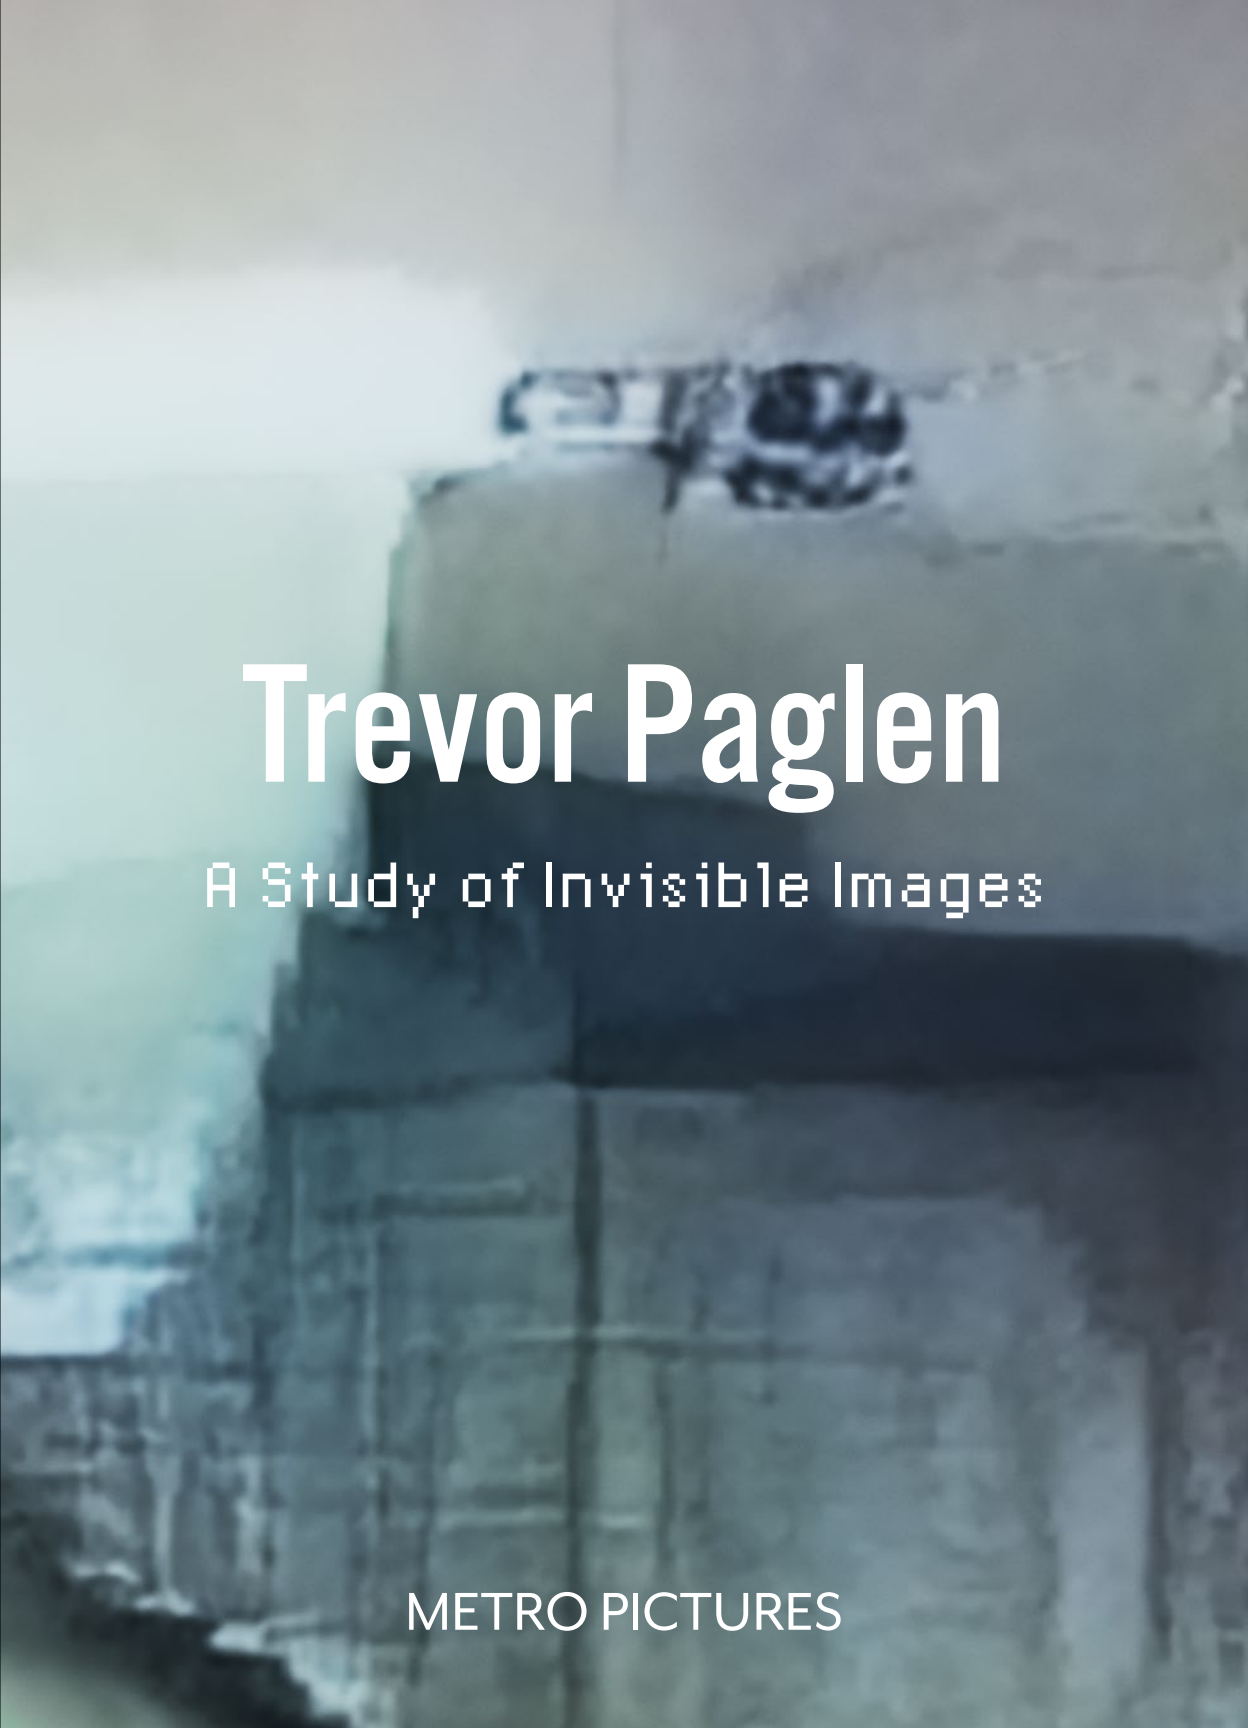

Trevor Paglen

A Study of Invisible Images

METRO PICTURES

Trevor Paglen

METRO PICTURES

A Study of Invisible Images

Produced on the occasion of
"Trevor Paglen: A Study of Invisible Images" at Metro Pictures, New York
September 8 - October 21, 2017

Metro Pictures
519 West 24th Street
New York, NY 10011
www.metropictures.com

Trevor Paglen's *A Study of Invisible Images* is the first exhibition of works to emerge from his ongoing research into computer vision, artificial intelligence (AI) and the changing status of images. This body of work has formed over years of collaboration with software developers and computer scientists and as an artist-in-residence at Stanford University. The resulting prints and moving images reveal a proliferating and otherwise imperceptible category of "invisible images" characteristic of computer vision.

Paglen's exhibition focuses on three distinct kinds of invisible images: training libraries, machine-readable landscapes and images made by computers for themselves. For *Machine-Readable Hito*, for example, Paglen took hundreds of images of artist Hito Steyerl and subjected them to various facial recognition algorithms. This portrait of Steyerl presents the images alongside metadata indicating the age, gender, emotional state and other signifiers that the algorithms have interpreted from the images. In another portrait in the show, Paglen trained facial recognition software to read the face of philosopher Frantz Fanon. A ghostly image of Fanon shows the facial signature–the unique qualities of a face as determined by biometric recognition software–used by computer vision to identify an individual.

To make the prints in *Adversarially Evolved Hallucinations*, Paglen trained an AI to recognize images associated with taxonomies such as omens and portents, monsters, and dreams. A second AI worked in tandem with the first to generate the eerie, beautiful images that speak to the exuberant promises and dark undercurrents characterizing our increasingly automated world.

The video installation *Behold These Glorious Times!* brings together hundreds of thousands of training images routinely used for standardized computer vision experiments and pairs them with visual representations of an AI learning to recognize the objects, faces, expressions and actions. A loose narrative begins to emerge about the collapsing distinctions between humans, machines and nature. Electronic musician Holly Herndon composed a soundtrack using libraries of voices created to teach AI networks how to recognize speech and other acoustic phenomena.

A Study of Invisible Images builds on *Sight Machine*, a work staged in San Francisco earlier this year in collaboration with the Kronos Quartet. As the quartet played a concert, cameras attached to computer vision algorithms used in self-driving cars, guided missiles, spy satellites and other autonomous vision systems projected what they "saw" onto a screen behind the musicians.

Paglen has upcoming exhibitions at the Smithsonian American Art Museum and Museo Tamayo in Mexico City. He has had one-person shows at Secession, Vienna; the Berkeley Art Museum; Kunsthall Oslo; and the Frankfurter Kunstverein. His work has been exhibited at the Metropolitan Museum of Art; Tate Modern; San Francisco Museum of Modern Art; BALTIC Centre for Contemporary Art; Museo Reina Sofía, Madrid; MIT List Visual Arts Center; Haus der Kunst, Munich; and the Walker Art Center. He participated in the 2009 Istanbul Biennial, 2012 Liverpool Biennial, 2013 ICP Triennial and the 2016 Gwangju Biennale. He has received numerous awards, including the 2014 Electronic Frontier Foundation Pioneer Award and the 2016 Deutsche Börse Photography Foundation Prize. He has written *The Last Pictures*, a critical compendium of his Creative Time project to launch an ultra-archival disc micro-etched with one hundred photographs into orbit; *Blank Spots on the Map: The Dark Geography of the Pentagon's Secret World*; and *I Could Tell You but Then You Would Have to Be Destroyed by Me: Emblems from the Pentagon's Black World*. Paglen's article "Invisible Images: Your Pictures are Looking at You" about computer vision and artificial intelligence was published by *The New Inquiry* in December 2016.

Four Clouds

The four skyscapes shown here are overlaid with strokes and lines that show what four different computer vision algorithms are seeing in the images. The algorithms are looking for unique keypoints—areas of interest—and are attempting to simplify the underlying photograph into a series of sections. The algorithms shown here are used in simplified object-recognition contexts and in technologies such as guided missiles, autonomous surveillance systems, and 3-D modeling.

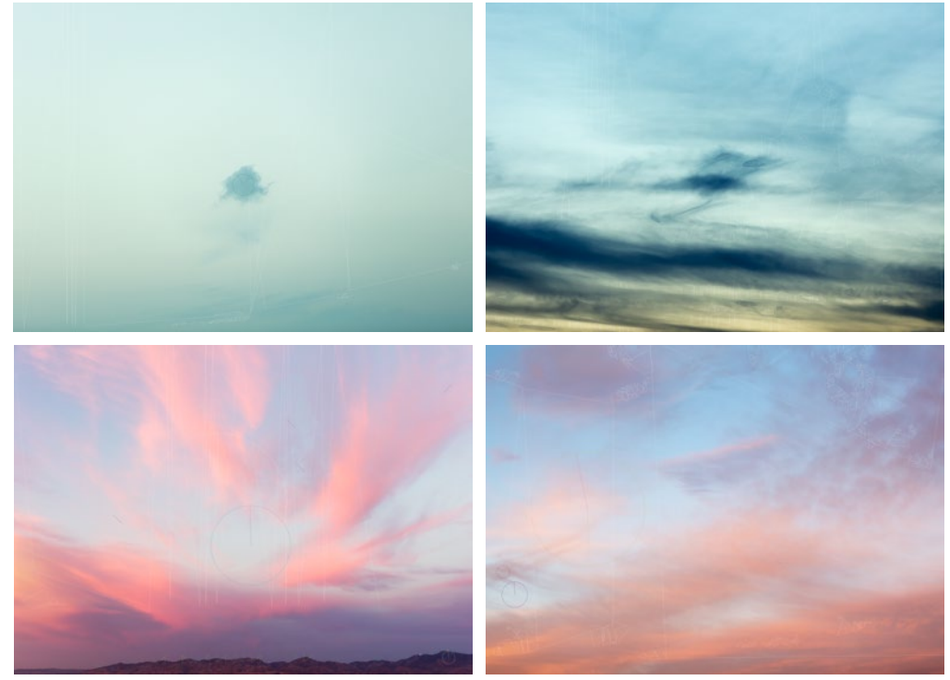

Four Clouds
Scale Invariant Feature Transform; Maximally Stable Extremal Regions;
Skimage Region Adjacency Graph; Watershed, 2017
4 pigment prints

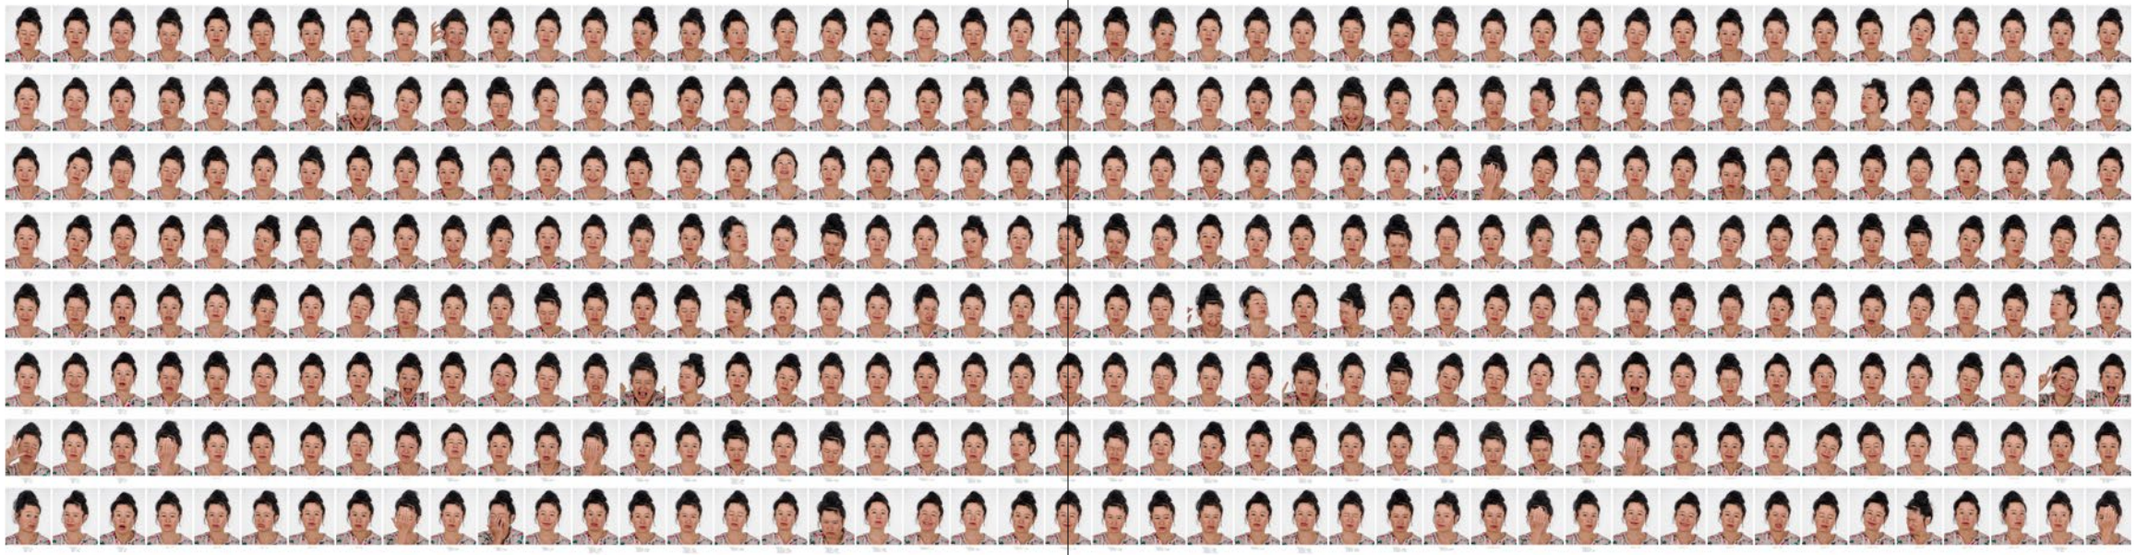

Machine-Readable Hito

This piece is made out of hundreds of portraits of artist Hito Steyerl that have been analyzed by various facial-analysis algorithms. Below each picture is the output of algorithms attempting to detect her age, gender, and emotional state. Other algorithms attempt to determine whether, for example, she is wearing glasses, is smiling, or has a beard.

Machine-Readable Hito, 2017
adhesive wall material

Eigenface

A standard technique in facial recognition software is to use an algorithm to create a "faceprint" of a given person and to use that faceprint to try and match a person's face with photos. To grossly oversimplify, if you want to teach an algorithm how to distinguish a particular person (say, philosopher Frantz Fanon) from a collection of other people, you need to have a big collection of photographs of people's faces, with everyone's faces labeled. You then take all the faces of Fanon, align them so their eyes and mouths are in the same place, and average them together. Then you take the average of all the other faces in the collection and average them together. If you subtract the average image of all the other people from the average of Fanon's face, you end up with a faceprint for Fanon showing what distinguishes him from everyone else in the dataset. You can then use this faceprint to identify any future images of Fanon's face that you might come across. This portrait translates that faceprint (which in its native form is a mathematical abstraction) into an image that human eyes can recognize as a face.

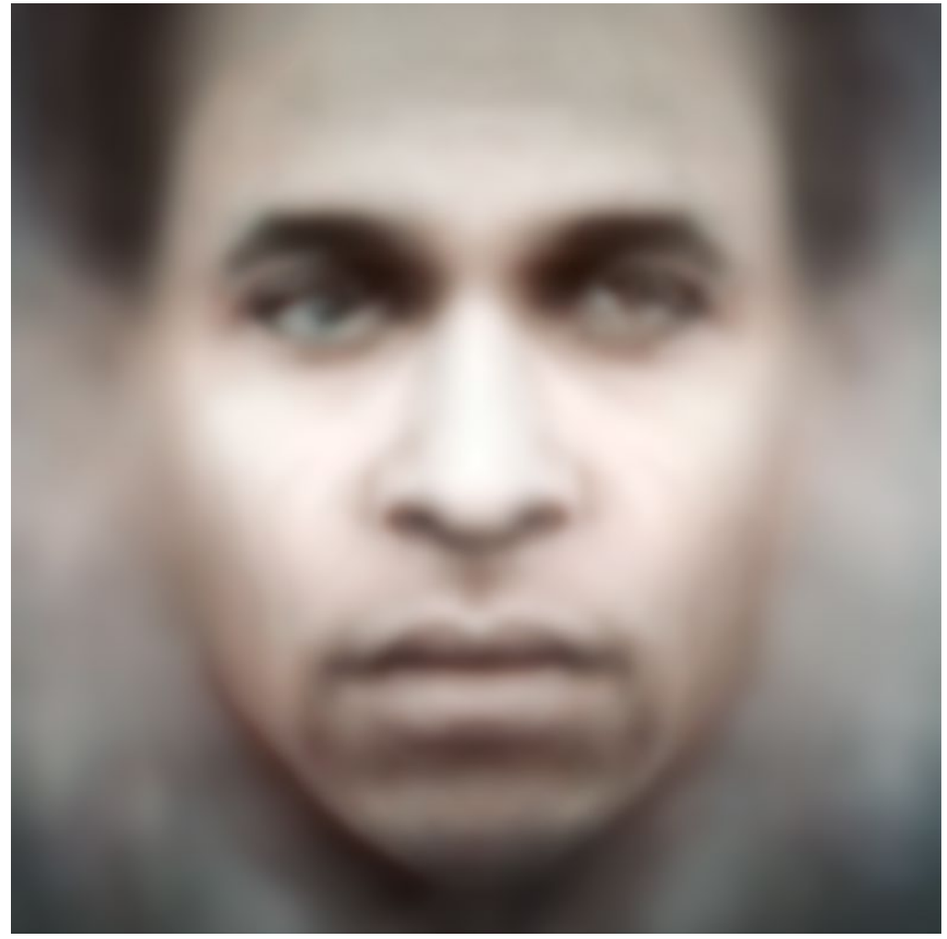

"Fanon" (Even the Dead Are Not Safe)
Eigenface, 2017
dye sublimation metal print

Megalith

One of the earliest tasks that neural networks and artificial intelligence could do reliably well was to recognize written numbers. These sorts of number-recognition systems are ubiquitous, as anyone who's ever had an ATM automatically read the handwritten numbers of a deposited check knows.

Megalith is made out of nearly 70,000 handwritten digits that represent one of the original collections of images that number-recognition systems were built upon.

One can think of this piece as a kind of Rosetta stone, an interface between two languages: that of written human numbers, and the language of artificial intelligence. But in this piece, the translation is an invisible, machinic interpretation of these numbers, and is inaccessible to our human senses.

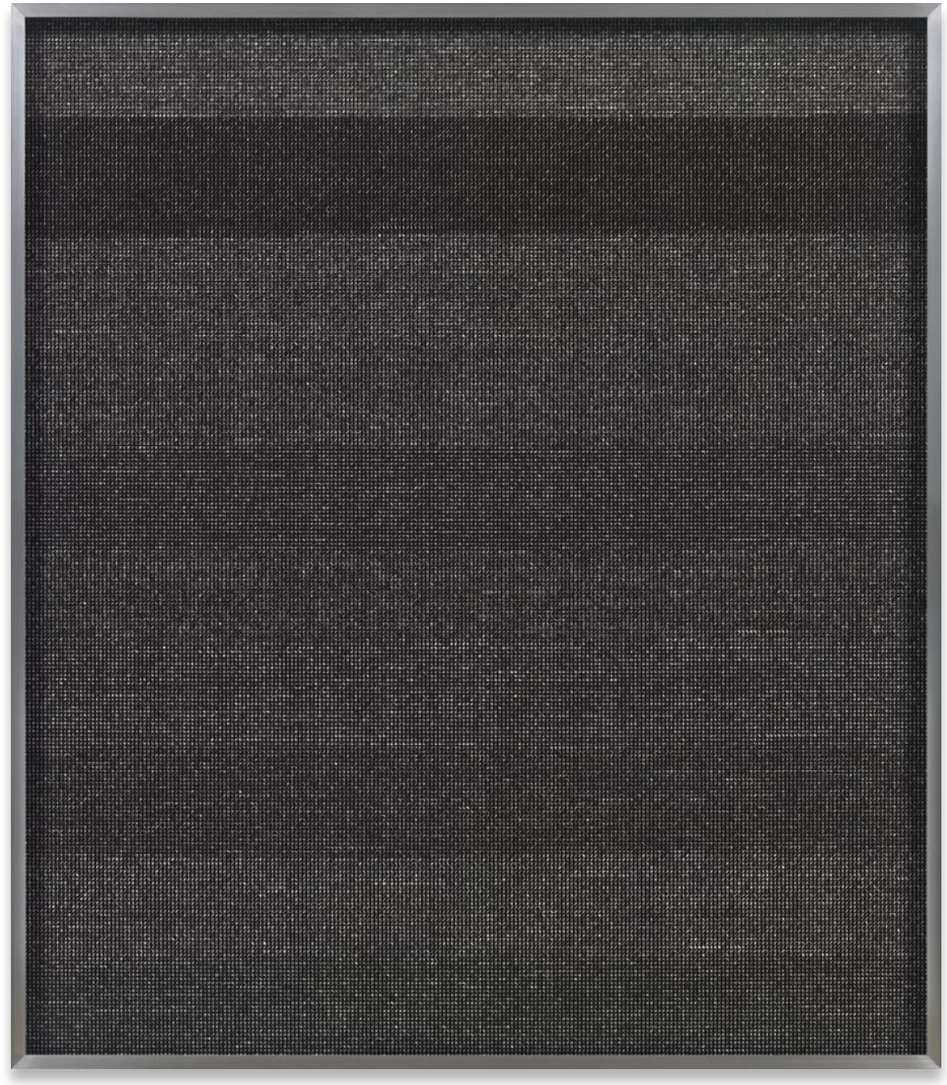

Megalith, 2017
pigment print

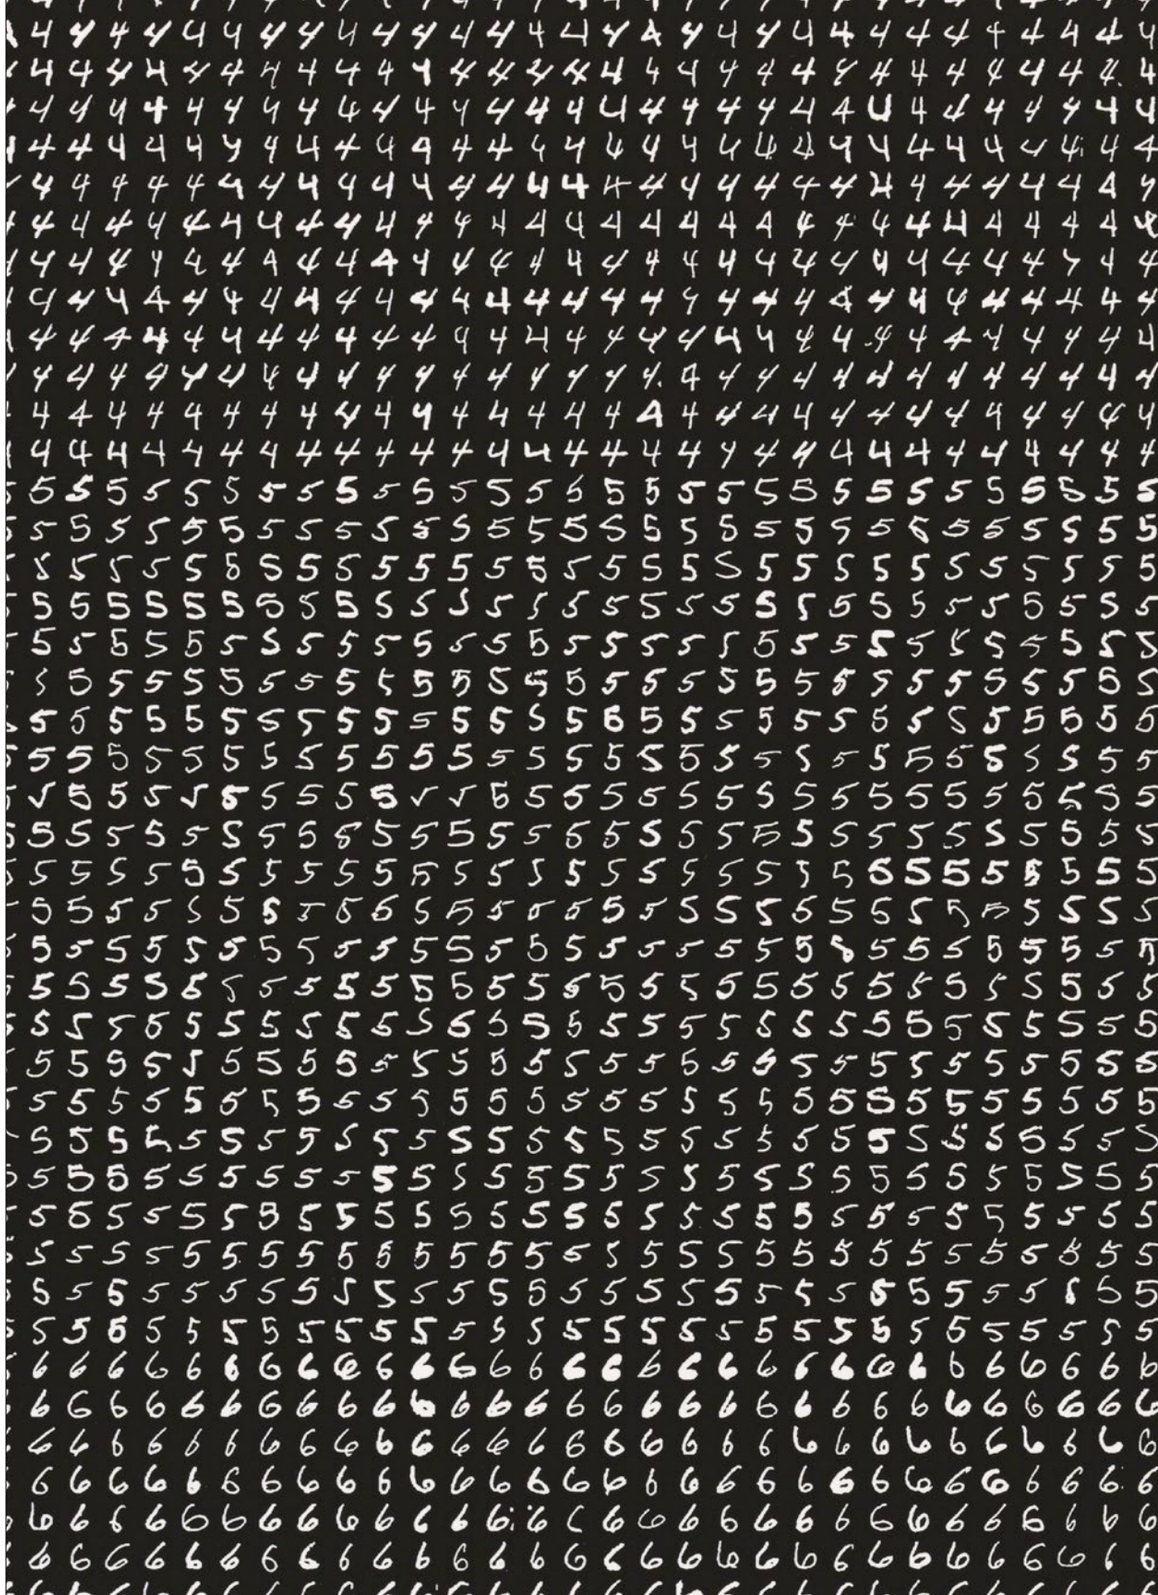

Megalith, 2017
(detail)

It Began as a Military Experiment

Contemporary research into facial recognition technology began in earnest in the mid-1990s at the behest of DARPA, the Defense Advanced Research Projects Agency. The military wanted facial recognition to exist, so DARPA began funding researchers in computer science and computer vision to work on the problem. The military realized that to do facial recognition, researchers would need to have access to thousands of images of people's faces. So in the early 1990s the military funded the creation of something called the FERET database, which included tens of thousands of pictures of many thousands of people, most of whom worked at a military base in Maryland. This work was made by combing through the FERET database over many months to "curate" a selection of portraits, to retouch and color correct them, and to run them through an algorithm that identifies the keypoints in their faces.

It Began as a Military Experiment, 2017
set of ten pigment prints

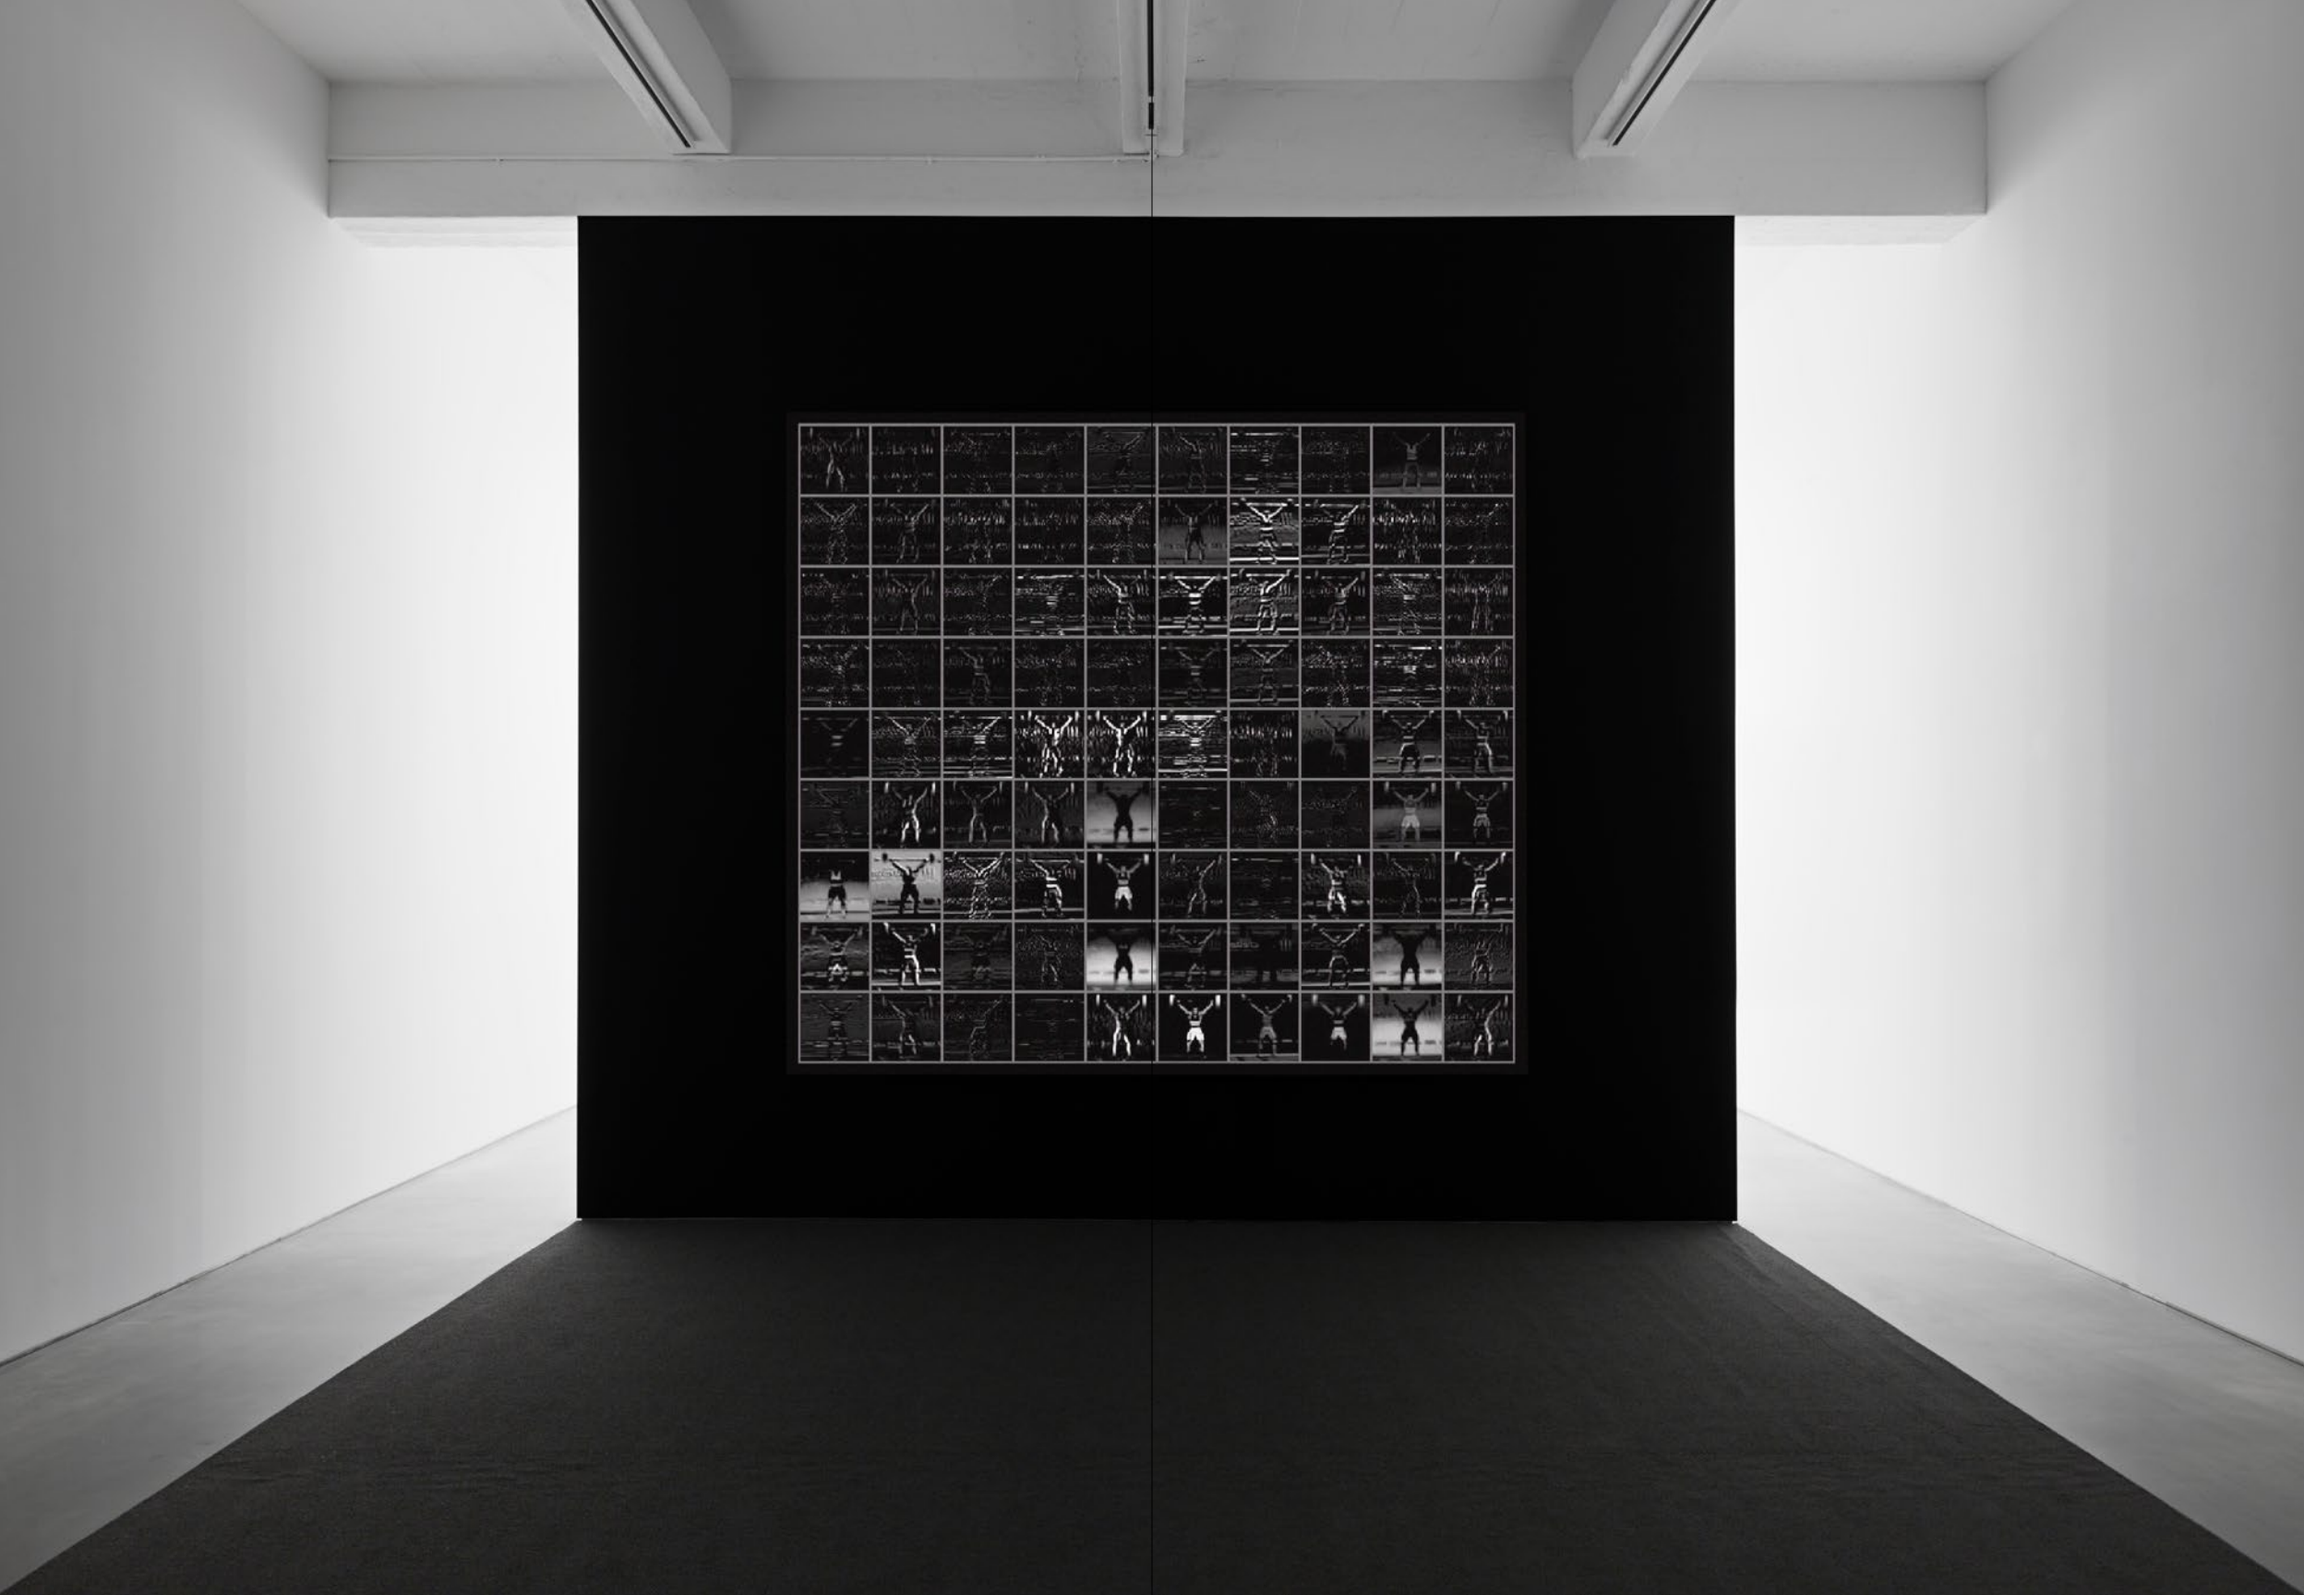

Behold these Glorious Times!

This video installation is composed of images from two sources. The photographic images in the video are parts of training libraries used to teach artificial intelligence networks how to recognize objects, faces, gestures, relationships, emotions, and much more. They are images designed to teach machines how to see. The other type of images in this video installation show what a Deep Neural Network (an artificial intelligence architecture) is actually seeing when it ingests these images. In order to learn how to recognize images, the AI breaks images into hundreds of component parts and tries to put them back together. In the black-and-white grids and images in this video, we are seeing the various ways that the AI is pulling the images apart in order to try to make sense of them. Overall, we are seeing images being used to teach AI, and seeing what the AI is seeing when it looks at them and tries to make sense of them.

The music for this piece, composed by Holly Herndon, was crafted in part by using neural networks designed to synthesize voices and instruments, and from training libraries used in speech recognition and other auditory machine learning applications.

Behold these Glorious Times!, 2017
single channel color video projection, stereo mix

Hallucinations

One of the most common applications of artificial intelligence is to do automatic object-recognition and image-captioning. When you upload an image to Facebook or other social media, powerful artificial intelligence algorithms can recognize the identities of people in images, the objects, the products and even the places depicted in those images.

AIs are taught how to recognize objects by giving them training sets. A training set will consist of thousands or even millions of images organized into pre-sorted classes that correspond to each of the kinds of objects that the AI will eventually be able to distinguish. For example, if you want to train an AI to recognize all the objects in a kitchen, you might give it a thousand pictures of a fork, a spoon, a knife, a countertop, a frying pan, a pot, etc.... Once that AI is trained, you can give it a picture of a fork that it has never seen before and it should be able to recognize it as a fork.

For this body of work, Paglen created massive training sets based on literature, philosophy, folk-wisdom, history, and other "irrational" things, then taught the AIs to recognize things from those corpuses. Some examples include: "Interpretations of Dreams," an AI that has been trained to see – and only see – symbols from Freudian psychoanalysis; "Omens and Portents," an AI that can only see things like comets, eclipses, and other signs of bad things to come; and "American Predators," an AI that sees various predatory animals, plants, and humans indigenous to the United States as well as military hardware like predator drones and stealth bombers.

Once an AI has been trained to see all the objects in a particular corpus, Paglen tries to get it to hallucinate an image of something it's been trained to see. This is done by creating a second AI network, whose job it is to draw shapes. The two AIs then play a little game. The drawing AI (also called a Generator) tries to draw pictures that will fool the AI that's been trained to see or to discriminate between particular objects (this is the AI Paglen trained; we can call it the Discriminator.) The two AIs go back and forth thousands or millions of times, until the Generator has learned how to make images that can reliably fool the Discriminator. The images that come out of this process are called Hallucinations. Together, the AIs have evolved an image that is entirely synthetic and has no referent in reality, but that the pair of AIs believe are examples of things they've been trained to see.

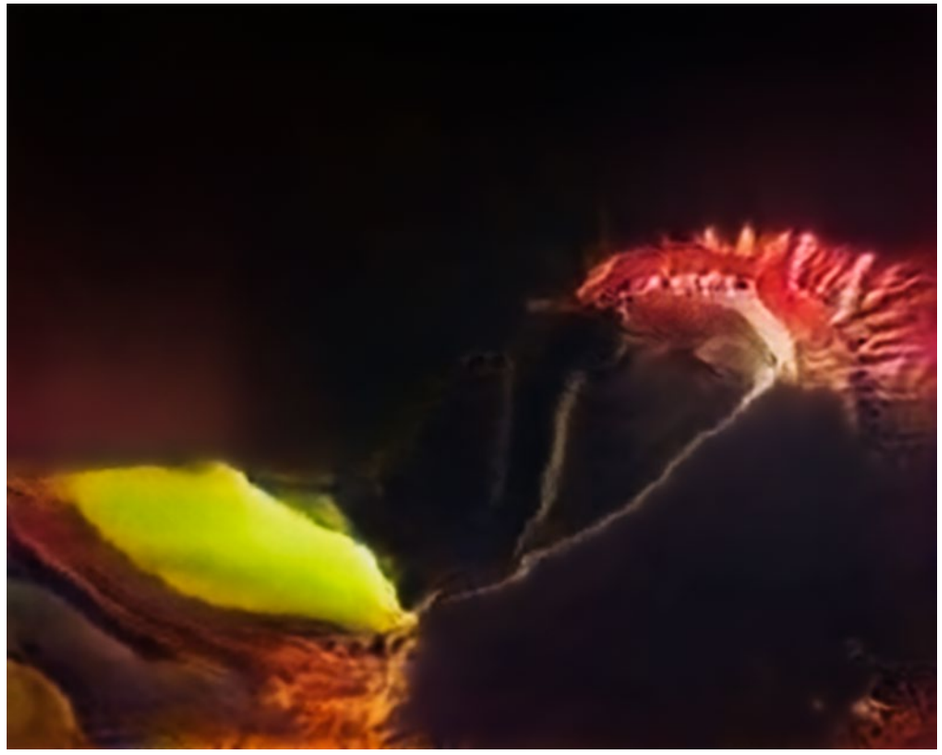

Venus Flytrap (Corpus: American Predators)
Adversarially Evolved Hallucination, 2017
dye sublimation metal print

A Man (Corpus: The Humans)
Adversarially Evolved Hallucination, 2017
dye sublimation metal print

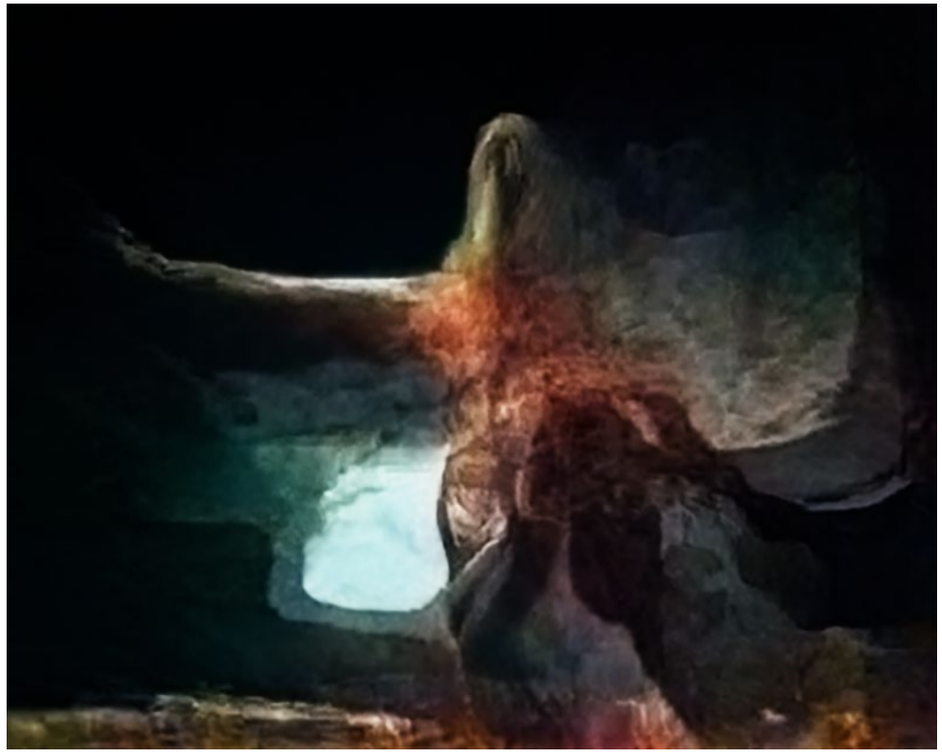

Octopus (Corpus: From the Depths)
Adversarially Evolved Hallucination, 2017
dye sublimation metal print

Comet (Corpus: Omens and Portents)
Adversarially Evolved Hallucination, 2017
dye sublimation metal print

Highway of Death (Corpus: The Aftermath of the First Smart War)
Adversarially Evolved Hallucination, 2017
dye sublimation metal print

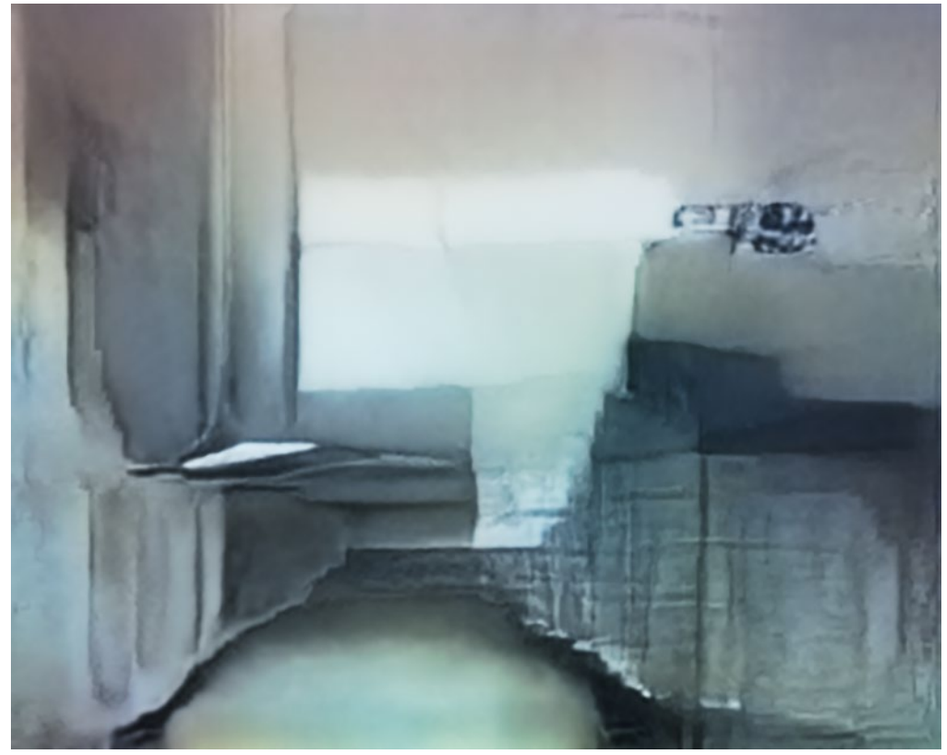

A Prison Without Guards (Corpus: Eye-Machines)
Adversarially Evolved Hallucination, 2017
dye sublimation metal print

Vampire (Corpus: Monsters of Capitalism)
Adversarially Evolved Hallucination, 2017
dye sublimation metal print

Porn (Corpus: The Humans)
Adversarially Evolved Hallucination, 2017
dye sublimation metal print

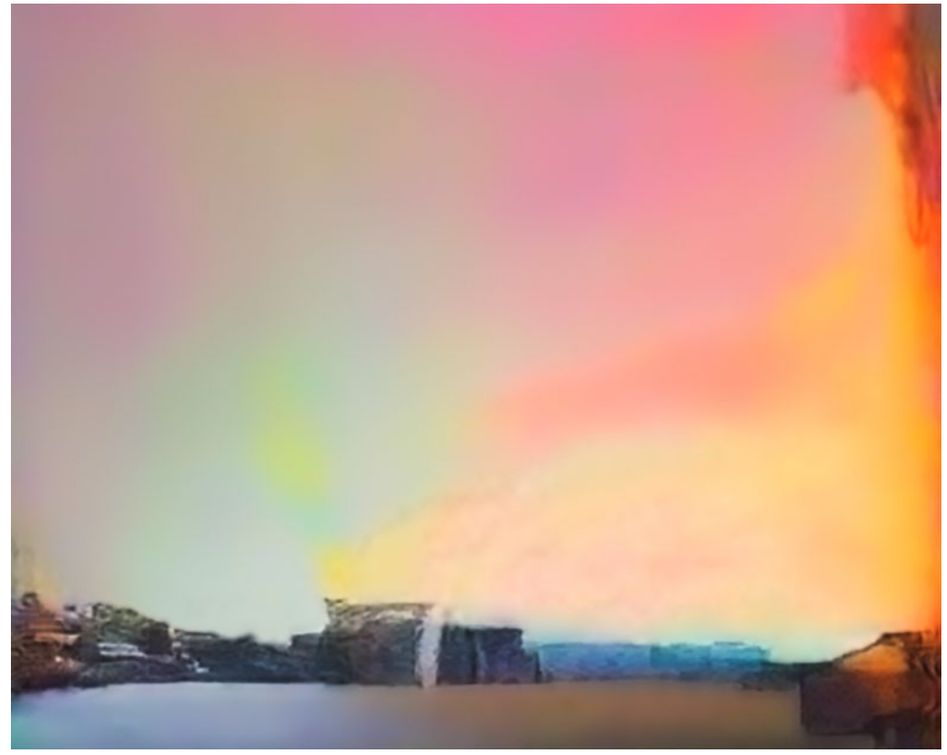

Rainbow (Corpus: Omens and Portents)
Adversarially Evolved Hallucination, 2017
dye sublimation metal print

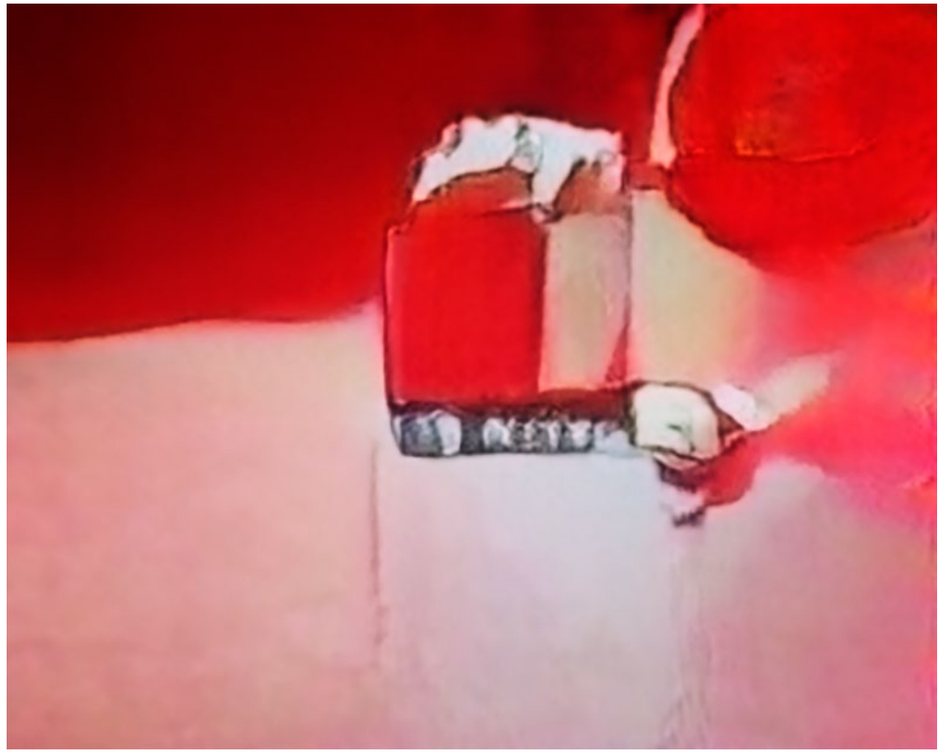

False Teeth (Corpus: Interpretations of Dreams)
Adversarially Evolved Hallucination, 2017
dye sublimation metal print

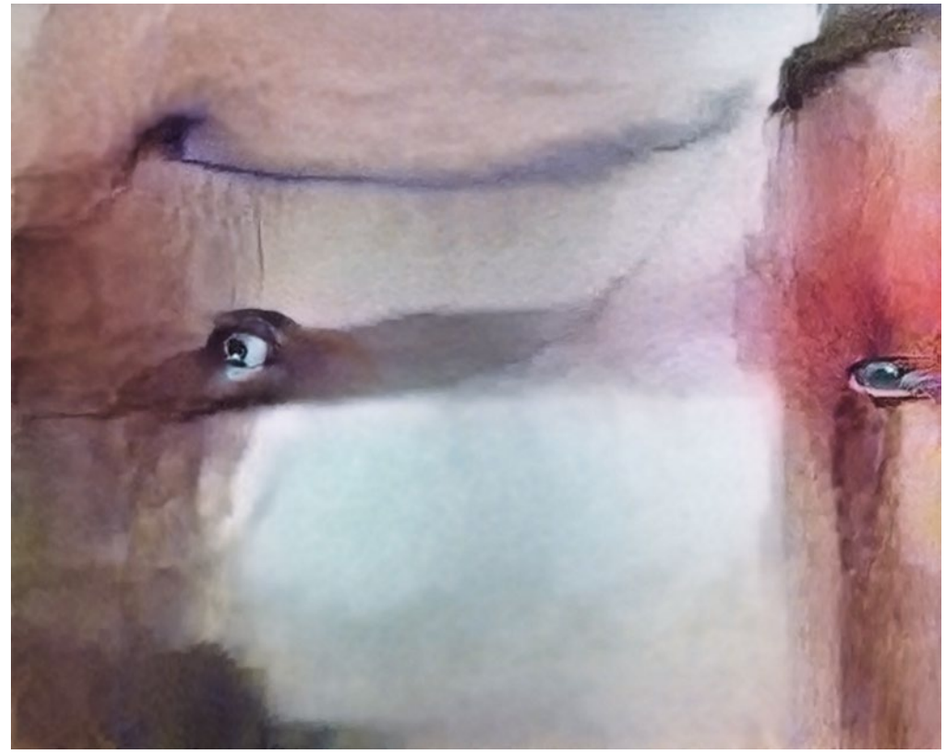

Human Eyes (Corpus: The Humans)
Adversarially Evolved Hallucination, 2017
dye sublimation metal print

List of Works

1. *Four Clouds*
Scale Invariant Feature Transform; Maximally Stable Extremal Regions; Skimage Region Adjacency Graph; Watershed, 2017
quadriptych; each pigment print
each: 33 x 46 inches; 83.8 x 116.8 cm

2. *Machine-Readable Hito*, 2017
adhesive wall material
193 x 55 1/8 inches; 490.2 x 140 cm

3. *"Fanon" (Even the Dead Are Not Safe)*
Eigenface, 2017
dye sublimation metal print
48 x 48 inches; 121.9 x 121.9 cm

4. *Megalith*, 2017
pigment print
72 x 63 inches; 182.9 x 160 cm

5. *It Began as a Military Experiment*, 2017
set of ten pigment prints, glossy finish
each: 13 5/8 x 10 1/2 inches; 34.6 x 26.7 cm

6. *Behold these Glorious Times!*, 2017
single channel color video projection, stereo mix
duration: 10'

7. *Venus Flytrap (Corpus: American Predators)*
Adversarially Evolved Hallucination, 2017
dye sublimation metal print
21 1/2 x 26 7/8 inches; 54.6 x 68.3 cm

8. *A Man (Corpus: The Humans)*
Adversarially Evolved Hallucination, 2017
dye sublimation metal print
48 x 60 inches 121.9 x 152.4 cm

9. *Octopus (Corpus: From the Depths)*
Adversarially Evolved Hallucination, 2017
dye sublimation metal print
21 1/2 x 26 7/8 inches; 54.6 x 68.3 cm

10. *Comet (Corpus: Omens and Portents)*
Adversarially Evolved Hallucination, 2017
dye sublimation metal print
21 1/2 x 26 7/8 inches; 54.6 x 68.3 cm

11. *Highway of Death (Corpus: The Aftermath of the First Smart War)*
Adversarially Evolved Hallucination, 2017
dye sublimation metal print
32 x 40 inches; 81.3 x 101.6 cm

12. *A Prison Without Guards (Corpus: Eye-Machines)*
Adversarially Evolved Hallucination, 2017
dye sublimation metal print
32 x 40 inches; 81.3 x 101.6 cm

13. *Vampire (Corpus: Monsters of Capitalism)*
Adversarially Evolved Hallucination, 2017
dye sublimation metal print
60 x 48 inches; 152.4 x 121.9 cm

14. *Porn (Corpus: The Humans)*
Adversarially Evolved Hallucination, 2017
dye sublimation metal print
48 x 60 inches; 121.9 x 152.4 cm

15. *Rainbow (Corpus: Omens and Portents)*
Adversarially Evolved Hallucination, 2017
dye sublimation metal print
21 1/2 x 26 7/8 inches; 54.6 x 68.3 cm

16. *False Teeth (Corpus: Interpretations of Dreams)*
Adversarially Evolved Hallucination, 2017
dye sublimation metal print
21 1/2 x 26 7/8 inches; 54.6 x 68.3 cm

17. *Human Eyes (Corpus: The Humans)*
Adversarially Evolved Hallucination, 2017
dye sublimation metal print
48 x 60 inches; 121.9 x 152.4 cm