Ways of Seeing with Computer Vision: Artificial Intelligence and Institutional Critique

Gabriel Pereira¹ and Bruno Moreschi²

This preprint has been significantly changed for publication on AI & Society (Special Issue on Ways of Machine Seeing). [Please read the Accepted Manuscript here](#) or the final authenticated version in the AI & Society journal: <https://doi.org/10.1007/s00146-020-01059-y>.

Summary

During 2018, as part of a research funded by Deviant Practice Grant, artist Bruno Moreschi and digital media researcher Gabriel Pereira worked with the Van Abbemuseum collection (Eindhoven, NL), reading their artworks through commercial image-recognition (Computer Vision) Artificial Intelligences from leading tech companies.

First, Pereira created a platform that analyzes the images of the collection using six different commercial AIs, generating a myriad of results. Then, Moreschi went through the results of 654 works, and meticulously analyzed a large portion of them. He iteratively created analytical categories for the results. After that, in a process of understanding and organizing the collected data, Pereira and Moreschi looked for glitches, errors and other unexpected readings by the AI (from their human perspective, evidently). The aim of this was to look for new ways of understanding artworks, by both levelling and expanding their meanings. The images of artworks were often read without their inherent contexts – much different than a traditional human reading. Lastly, the research speculated on how the AIs are constructed, and what these “mistakes” reveal about them and their training data. This led to a brief analysis on the crucial invisible labor of Amazon Mechanical Turkers.

The main takeaways were: somewhat as expected, AI is constructed through a capitalist and product-focused reading of the world (values that are embedded in this sociotechnical system); that this process of using AI is an innovative way for doing institutional critique, as AI offers an untrained eye that reveals the inner workings of the art system through its glitches. This research, unlike some already existing critical analyses of AI, does not focus on accountability and bias. It aims to regard these glitches as potentially revealing of the art system, and even poetic at times. They reveal the inherent fallibility of the commercial use of AI and machine learning to catalogue the world: it cannot comprehend other ways of knowing about the world, outside the logic of the algorithm. But, at the same time, due to their “glitchy” capacity to nivelate and reimagine, these faulty readings can also serve as a new way of reading art; a new way for thinking critically about the art image in a moment when visual culture

¹ PhD Fellow in Information Studies and Digital Design at Aarhus University (DK): <http://gabrielpereira.net/>

² Visual Artist, PhD in Visual Arts at Unicamp (State University of Campinas/BR): <https://brunomoreschi.com/>

has changed form to hybrids of human-machine cognition and “machine-to-machine seeing” (Paglen, 2016).

Finally, the results of this research are presented in different ways. Along with this article, the researchers have produced the video *Recoding Art: Van Abbemuseum Collection*, relating non-traditional footage of the museum with some of the conclusions here discussed; two proposals to intervene in the exhibition space; and a draft for a possible publication that presents the museum's collection through their readings by AI. We find it important to speak through these different outputs, as they show how this project uses both art and social research methodologies, occupying an interdisciplinary position.

Introduction: Duchamp's Fountain is a urinal

Fig. 1: Print from the e-mail that Moreschi received from Pereira, with the reading of Duchamp's Fountain by the AI Google Cloud Vision. None of the results considered the image a photograph of an artwork.

MORESCHI: On 23 October 2017, Gabriel sent me an email. In it, there were two images – a painting of Christ and Duchamp's Fountain. The email went on with a series of graphics, percentages and keywords analyzing these two images. At no point were they interpreted as art. Duchamp's Fountain was described as a plumbing fixture, product design and as... a urinal. Behind this reading was Google's state-of-the-art AI: Google Cloud Vision. As someone not specialized in Artificial Intelligence, I was fascinated by these results.

The image of Duchamp's *Fountain* is especially relevant as a starting point. With this artwork, as Calvin Tomkins (2014) details in his biography of Duchamp, the French artist aimed to test how democratic the New York Society of Independent Artists was in the selection process for their salon. Those who work with art are often familiar with this story: after having

lunch (and probably lots of wine) with Walter Arensberg and Joseph Stella, Duchamp invited them to accompany him to J. L. Mott Iron Works, a store in New York that specializes in sanitary equipment. There, he bought a porcelain Bedfordshire urinal. Later, in his studio, Duchamp flipped the urinal upside down and signed the bottom left side with the name R. Mutt and the year (1917). He then submitted the piece to the exhibition, without a return address. According to Tomkins, when the package was opened by the salon's jury (Duchamp was one of the jurors, but he did not speak), the juror George Bellows cried out: "It's indecent! It's indecent!" One of the members tried to calm him down, insisting that their main rule was to accept all submissions, without restrictions. But Bellows reacted by saying "We can't show this. This thing is nothing more than what it is."

Today, Duchamp's *Fountain* is considered one of the most influential artworks of the twentieth century, an iconic image that is widely known by the popular imaginary. It is essential to understanding what contemporary Western art became, and much of the art produced in the past century. Duchamp's transformation of the urinal into a fountain by placing it in the artistic space is a crucial ontological shift proposed by contemporary and conceptual art. But, one hundred years later, when Google's Cloud Vision looked at the image of this artwork, its "eyes" didn't see anything different than Bellows': "ceramic, product, urinal."

Google Cloud Vision is the result of much human and machine labor. It is based on Google's recent expansion into the new and promising field of Computer Vision: using algorithms, Machine Learning, and a lot of data to train smart machines to see and understand the world around us. Fei-Fei Li, a Professor at Stanford and also former Chief Scientist at Google, is one of the most prominent voices of the Computer Vision research field. In her widely watched Ted Talk *How We're Teaching Computers to Understand Pictures*, she goes on about what this means:

Just like to hear is not the same as to listen, to take pictures is not the same as to see, and by seeing, we really mean understanding. In fact, it took Mother Nature 540 million years of hard work to do this task, and much of that effort went into developing the visual processing apparatus of our brains, not the eyes themselves. So, vision begins with the eyes, but it truly takes place in the brain. (Li, 2015)

This difference between taking pictures, seeing, and understanding is intriguing, especially in a moment when the utopian computer vision discourse (as seen above) claims that it is possible "to teach the machines to see just like we do: naming objects, identifying people, inferring 3D geometry of things, understanding relations, emotions, actions and intentions" (Li, 2015). But, as we first experimented at that moment, commercially available Computer Vision algorithms from leading tech companies are not trained to "understand" artworks – they do not understand the context, the subtext, the emotion. When they do interpret it correctly, from the point of view of the human observer, they read them in their superficiality: the "thing as nothing more than what it is."

Going back to Duchamp's *Fountain*, the jury of the salon finally decided not to exhibit the piece. When the artist got it back, he took it to famous photographer Alfred Stieglitz, to be photographed using the same method as a sculpture. According to artist Beatrice Wood (Tomkins, 2014), for at least one hour, Duchamp helped Stieglitz with the illumination and decided that the best position would be one that made the urinal look like some sort of bride's veil. The *Fountain* itself probably had the same destination of many other of Duchamp's ready-mades (the trash bin) – replicas of the work now abound in many prominent museums. But the image of the artwork by Stieglitz remains, and is now the cover of books, magazines and widely available on Google Image Search. This story, although somewhat specific, reveals the way in which art and its history are constructed and experienced through images (rather than through the works' materiality). Now, using Artificial Intelligence, like the one by Google, it is also clear that these images alone are not enough. The images themselves are not embedded with what these artworks mean, their context, history and subtext. There is a rich possibility in all of this: because commercial AIs are not at all familiar with works of art, we have interpretations of their images that are almost always devoid of a more "subjective" sense of art and its context – which were so vital to Duchamp and most other artists.

PEREIRA: I've just watched the first episode of "Ways of Seeing," the influential BBC TV show by art critic John Berger (Bruno made me watch it, to be honest). I'm particularly impressed by the scene where he asks kids to describe a painting of Christ by Caravaggio. The kids (very adorably) speak incessantly, conjecturing what the subjects are doing (maybe stealing the food or about to kiss?), and who they are (male or female? Jesus?). Berger points out how these kids 'demystify' the artistic work by looking at it "very directly," from their own experiences, ignoring the context in which the images live in.

Fig 2: Frame of episode 1 from *Ways of Seeing*, TV show written and presented by British art critic John Berger in 1972 (BBC Two) and adapted from a book of the same title.

The AIs that power Computer Vision can only understand the world based on their own “experiences” when making assessments about the world. AI requires data to be trained on, from which it generates a model; for example, pointy ears means cat, floppy ears means dog. This model is most often not visible or interpretable by humans, and frequently involves patterns that are not visible to the human observer (Olah et al, 2018). Could AI then be used to bring a fresh, denaturalized set of “eyes” to art, expanding and levelling what artworks are (or could be)? But also, if art is such a new thing to the “eyes” of computers, could it also reveal what its underlying experiences are, much like the kids use their experience to talk about Caravaggio’s painting? It was with these questions that we began this research. How revealing are these “eyes” that have never entered a museum, and how can we make use of them to see for ourselves? How could we use these nonhuman perspectives to “illuminate our understanding of the world,” thus unsettling “the relations between what we see and what we know in new ways”? (Cox, 2017, p.14)

We were invited to work with the Van Abbemuseum (NL) collection and proceeded to reading their images using commercial image-recognition (Computer Vision) Artificial Intelligences from leading tech companies. The museum’s collection, consisting of conceptual and contemporary art, as well as some older and lesser-known works, would be particularly suitable for our proposed inquiry.

In the first part of this article, we describe our methodology and how we approached the creation of the platform for reading images through AI, how we analyzed the results, and what categories came out of it – if you are in a hurry, you can skip the methodology and go straight to the results on page 8. In the second part, we look at the question: what happens to art (i.e. images of artworks) when it is read through commercial AIs? This means putting art through a system that is not specifically made for it, where it is not protected by the elements of the artistic apparatus. Drawing from Institutional Critique, we investigate how this may work to expand and level the meaning of artworks. In the third part, we follow with: how does commercial AI react to art, a content that is different from what it is used to? And, what does this procedure reveal about the underlying values and epistemologies of popular AI tools? Here we borrow from literature on critical AI/algorithm studies and other academic research on technology/data, where algorithms are considered in their sociotechnical complexity. We conclude by stating the contributions we offer, as well as our considerations for future research and practice.

Methodology: “Honor thy error as a hidden intention”

The Van Abbemuseum collection is generally made up of contemporary and conceptual art. We received all the images of the collection (about 2500 high-resolution photographs of

artworks), but decided to focus on two of the permanent exhibitions, *The Making Of Modern Art* and *The Way Beyond Art*. This resulted in working with a total of 654 images.³

To create a new way of interpreting this set of images, Pereira created a script to send the images of the artworks to six of the most commonly used commercial AI services from Google, Microsoft, Amazon, IBM, Facebook, and the widely used YOLO library. The results obtained for each artwork are shown through a custom web interface, which is accessible and open-source (enabling other readings and analyses by other people) through this [link](https://testingdataviz.github.io/VAD0.4/).

Fig. 3: Recoding Art, an open-source website with the Van Abbemuseum's art works reading by AIs.

This centralized AI results interface was created prior to the official start of Moreschi's residency at Van Abbemuseum – when he was still in São Paulo (BR), where he lives. This meant that our first contact with the works in the collection was digitally-mediated. Following this logic of physical detachment, even though he stayed only a couple of blocks away from the museum and its collection, the first two weeks of Moreschi's stay in Eindhoven did not focus on the museum itself (and its physical works). Instead, he dedicated all his time to the analysis of the approximately 55,000 results obtained from the analyses of 654 works (available on the Recoding Art interface), and to the construction of a method capable of organizing the results through identified patterns.

³ They were chosen because the first exhibition deals directly with the changes in the status of art in modernity, especially its reproducibility, and the second is dedicated almost exclusively to contemporary art, much of which dissociates what is seen from its signification. In other words, our initial research material was a set of images that, to be understood as art, may require a level of expertise that many people and Artificial Intelligences do not possess.

To understand the construction of this methodology, there is nothing better than to analyze the soundtrack used by Moreschi in this process of screening results – as sung by Bjork, “My headphones / They saved my life.” Much of the selection of AI data happened while listening to Brian Eno, Peter Schmidt and Edgar Varèse, three artists open to non-traditional sound results that many would call “noise.”⁴

Following this logic of appreciating what is generally overlooked, longtime friends Peter Schmidt and Brian Eno created, in 1975, a set of cards called *Oblique Strategies*, designed to aid in the artistic process. These cards had three limited editions before Schmidt died in the 1980s. Today, the deck is a luxury item among collectors. The entire project, including the contents of the cards, is documented on a [website](#) created by musician and educator Gregory Alan Taylor with the permission of Brian Eno.

One of the letters from *Oblique Strategies* epitomizes our methodological approach. It states: “Honor thy error as a hidden intention.” The advice was valuable in a selection process involving interpretations that at first seemed like blatant misunderstandings by dumb machines. We decided to steer away from a feeling of superiority, in relation to technological systems.

In general, the cards in the deck (as well as Varèse's strategy of composing music using sounds that usually go unnoticed) were fundamental to a kind of methodology that would make Moreschi look carefully at the interpretations by the AIs, minimizing a search for results that were “true” or “correct.” On the contrary: we decided to value the unexpected outcomes.

This stance goes against a significant part of the social studies of AI, which largely focuses on the mistakes and biases of AI (and how to avoid them).⁵ This quest for algorithmic accountability and ethics is indeed important, given the amount of problems AI is already causing and how these errors affect people (especially underserved minorities and marginalized communities). Focusing on “solving bias” may serve, however, as a diversion from critically interrogating these systems and understanding them in their complexity (Powles and Nissenbaum, 2018). Here, we turn to commercial Computer Vision systems’ failures not as something that needs to be improved for increasing the system’s overall efficacy (i.e. fixing the algorithm),⁶ but as a way of speculating on the machinations of the systems of both AI and the art system. It is about showing that these systems are neither a “given,” nor “certain.”

⁴ In the case of Varèse, there was a happy coincidence. In Eindhoven, Moreschi selected the results in a house down the street from where Varèse lived for nine months, when he was financed by the Philips company to compose *Poème Électronique*. The eight-minute song was to become the soundtrack of the pavilion created by Le Corbusier at the 1958 Brussels World's Fair. Much of the musical piece was created from ordinary sounds collected around the city, including its first sound, a church bell ringing. To this day, it is possible to hear the bell when visiting the exhibition halls of the Van Abbemuseum.

⁵ For works related to this, see for example O’Neil (2016) and Eubanks (2018).

⁶ Of course, the subject matter of this research (artworks from a collection) is also particularly suitable for this approach, since, unlike predictive policing and other egregious algorithmic systems, these errors do not directly harm others.

MORESCHI: *In the 14 days I spent analyzing all the results of our Art Recoding interface, I tried to put away my rational side, trained to deal with art after at least a decade of working as an artist. Many of those results had a freshness that I had not felt in a long time. It was two weeks of an intense and fascinating process of denaturalization: it felt as if I was coming into contact with contemporary art for the first time in my life.*

During the screening process, Moreschi classified works by results with similar characteristics. The same work could be included in more than one group: images of works read as products, particularly objects for household use; those that have been associated with pop culture, and may even contain celebrities; those interpreted as moments of human self-promotion (people posing and/or taking selfies, for example); works with official historical periods expanded through the AI readings; works with results which Moreschi considered poetic, including many “Dadaist” results; interpretations that reinvent the idea of human representation, like many sculptures and paintings (some abstract) that are read as people; artworks associated with “similar images” that radically shift their official discourse; images in which the AIs identified letters, words or phrases (whether or not they were there).

It was only after coding and categorizing the collection's images that Moreschi visited the museum's two exhibits in person. During the third week of work, this approach to the artworks happened in the most traditional way: walking through the exhibition space like any other visitor. However, during the fourth and final week of research, the mediation with the works was again denatured during the filming of *Recoding Art: Van Abbemuseum Collection*, which integrates the outcomes of this research. A considerable part of the scenes of this film were done with partially or completely shut eyes, deconstructing the persuasive power of the exhibitions' displays.

PEREIRA: *The most interesting thing about this research, which may not be clear to readers at a first, is that Moreschi's eyes, though very trained to the art system itself, did not fully understand the inner workings of AIs. We had extensive conversations, where I explained to him how these systems were “trained,” how they are thought, and what their intended uses are. In his curiosity, he asked me: “What are the human layers of AI?” That's when I introduced him to Turkers.*

During the last week of work, Pereira and Moreschi worked together in person. Quite familiar with the new collection of artworks that emerged from the AI analyses, we decided to interact with Amazon Mechanical Turkers, as a way to better understand the human layers of AI and avoid the oversimplified idea that AI is completely automatic. These workers are responsible for doing tasks that are still impossible for computers, such as classifying images inside of predefined categories, thus creating the training data for AIs. We surveyed a random sample of Turkers, asking them for descriptions of some of the artworks from the collection, and if they considered these so-called artworks to be art.

Denaturalizing art through AI: Institutional Critique 2.0

The glitches and mistakes of AI help us to denaturalize the art system and its functioning. The art system is highly codified and embedded with power/value systems. When we start considering the AI results that do not necessarily follow the structures of specialized meanings of art, we are unmasking much of what specialized discourses attempts to disguise. In this sense, many of the results obtained from the AIs invited us to think about important issues regarding the art system, which are not always apparent. In addition to that, many of these results can help mediate these works to non-specialized audiences, initiating a more accessible relationship with these objects. Among the results are:

1. Art as everyday objects

Interpretations such as these show that artworks are, beyond their discourse, made of materials that are also found outside of the museum context, in everyday life. As is the case when Duchamp's *Fountain* is read as an actual urinal, even with his signature, these readings invite us to see works of art in a way that is disconnected from the idea of authorship. To analyze these results is to think about the process of symbolic transformation of artworks, one of the processes that underpin contemporary art. These results, much like the results in the next pages, help to remove the so-called aura from the art object and transform a very important art collection into an assortment of easily recognizable objects.

Fig. 4: *Licht-raum Modulator* (1922-1930, replica 1970), by László Moholy-Nagy, described by Microsoft's AI as "a lamp that is lit up at night," a similar result to Google's AI ("lighting," "lamp," "light fixtures"). Darknet YOLO

(open-source AI) goes further and sees the work as a possible "boat", which helps us to construct an interesting hypothesis: that, from the interpretive logic of AIs, the circular reflection on the wall can be a full moon in the high sea.

a close up of an old computer 39%

ejection
seat

(/support/seat/ejection seat)

50%

Fig 5: *Aux Abords De La Grande Cité* (1960), by Corneille, read as "ejection seat" and "a close up of an old computer." A considerable part of the works analyzed by Microsoft Azure's AI are understood as approximations of something. Since the AI's gaze does not operate from the human logic of physical distance between the observer and the observed, the concepts of approach and depth radically change here – anything not recognizable at first may indeed be the detail of an everyday object.

#11 INTERROGATION (WHAT KIND OF BIRD ARE YOU?) 1956-1958 Max Ernst #11.jpg a close up of a decorated tree 26%

Google Cloud Vision Microsoft Azure Amazon Rekognition IBM Watson Facebook Detectron Darknet YOLO Amazon Mechanical Turk

Label Annotations: red 94% art 94% orange 92% painting 84% acrylic paint 78% modern art 77% visual arts 75% illustration 75% paint 71% artwork 62%

Safe Search Annotations: Adult: VERY_UNLIKELY Medical: LIKELY Racy: VERY_UNLIKELY Spoof: VERY_UNLIKELY Violence: UNLIKELY

Best Guess: max ernst what kind of bird are you

Visually Similar:

Associated Web Entities: Painting 1.2 Interrogation (what kind of bird are you?) 1.0 Art 0.7 Artist 0.6 Visual arts 0.5 Painter 0.5 Sculptor 0.5 Modern art 0.5 Drawing 0.5 Image 0.5

jack-o'-
lantern
(halloween)

(/lamp/jack-o'-
lantern
(halloween))

50%

Safe Search Annotations: Adult: VERY_UNLIKELY Medical: LIKELY Racy: VERY_UNLIKELY Spoof: VERY_UNLIKELY Violence: UNLIKELY

Fig. 6: *Interrogation (What Kind of Bird Are You?)* (1956-1958), by Max Ernst, read by IBM's AI as a "jack-o'-lantern," the pumpkin that is traditionally carved on Halloween in the United States – an example that shows how the interpretations by AIs are constructed from American ethnocentric logic. This example also shows that the fact that works by well-known artists in the collection were read as art does not prevent them from being interpreted as things that are unrelated to the artistic context. It is curious to see the multifaceted ability of AIs to offer, within the same set of results, both the legitimated layer of the image as well as its pre-art state.

Fig. 7: Some images received interpretations that prioritize the physical structures that protect or support the works ("picture frame," "framework," "supporting structure") rather than their legitimized artistic contents. This occurred with *White Relief* (1936), by Ben Nicholson, often understood by the AIs as merely a frame, and *Untitled* (1980), by Jannis Kounellis, which, for Google, has to do with the image of a shelf, which in fact is something necessary for exhibiting the work. Results such as these de-structure the hierarchy between layers of the art object that are considered to be artistic and non-artistic, and invite us to view envelopes and bases as part of the artistic structure that is often

indispensable in the legitimation of what is art. Framed works are also often read as television monitors, which leads us to the second group of results.

2. IKEA shopping cart

The vast majority of the works (almost 80%) were read, in at least one of their results, as consumer products that are easily found in department stores. Such results are valuable in critical art studies for reinforcing the fact that works of art are essentially commodities – even if much more expensive than curtains – and placing our current understanding of what art is within the context of capitalism and a consumer society.

Gargoyle 60% Ornament 60%

Tags (from caption): building statue sitting cat bird old table graffiti street covered man standing phone perched

#13 CONCETTO SPAZIALE: ATTESE 1960 Lucio Fontana #13.jpg

Google Cloud Vision Microsoft Azure Amazon Rekognition IBM Watson Facebook Detectron Darknet YOLO Amazon Mechanical Turk

Is Adult Content? false 3%

Is Racy Content? false 4%

Is Black and White? true

Dominant Colors: White

Accent Color:

Categories: abstract_ 0% abstract_shape 16% others_ 11% outdoor_ 0%

Tags (from caption): indoor white refrigerator sitting kitchen photo black door man table standing cat side counter room wearing young water woman stove

Tags (with confidence): wall 97% white 67%

Fig. 8: *Femme en vert* (1909), by Pablo Picasso, as a “gargoyle,” “ornament” and “phone.” And *Concetto Spaziale: Attese* (1960), by Lucio Fontana, as a “refrigerator,” “stove” and “kitchen.”

Labels: Alphabet 90% Text 90% Logo 53% Trademark 53% Emblem 52% Accessories 51% Art 51% Modern Art 51%

tetraskelion 0

66%

Label Annotations: font 70% product design 62% brand 62%
Safe Search Annotations: Adult: VERY_UNLIKELY Medical: VERY_UNLIKELY Racy: UNLIKELY Spoof: VERY_UNLIKELY
Violence: VERY_UNLIKELY
Best Guess: Studio South
Visually Similar:
Associated Web Entities: Brand 0.7 Apartment 0.6 Studio South 0.6 monk mackenzie. 0.5 Corporate identity 0.5
Studio apartment 0.5 Graphic design 0.5 Business 0.4 Printing 0.4
Text: LAT. 31° 25' N LONG. 8° 41' E em

Labels: Business Card 68% Paper 68% Text 68% Logo 66% Trademark 66% Word 59%

License 54% advertising 68% joint 64%

Fig. 9: Cubist works and those with textual content tend to be related not only to marketable objects ("product design," "bottle"), but also to specific companies or more general ideas of the business world. This is the case with *L'accordéon* (1926), by Fernand Léger, associated with Tetraskelion Softwares, a company in Jaipur (IN) that offers technological solutions for travel agencies. The same is true for *LAT. 31°25'N, LONG. 8°41'E* (1965), by On Kawara, ("brand," "business," "corporate identity") and the poster *Sorry, Sweetie, Way To Go, Dude!* (1994), by Guerrilla Girls, ("license," "advertising," "joint"). This demonstrates that the capitalist logic in AI readings is broader than just interpreting images as products – it also includes notions and practices not necessarily material to the consumer society.

fabric

0

50%

Fig. 10: *Gleichzeitigkeitsstück* (Nr. 23, 1. Werksatz) (1967), by Franz Erhard Walther as “a white shirt,” “military uniform,” “handbag” and a lot of t-shirt images as visually similar. *Balance* (Nr. 26, 1. Werksatz) (1967), by Franz Erhard Walther as “a bag of luggage,” “clothes” and a lot of trouser images. *Politisch* (Nr. 36, 1. Werksatz) (1967), by Franz Erhard Walther, as “fabric.” Performance fabrics are almost always read as fashion clothes or accessories by the Als, which makes some sense, since many of them were worn by artists and/or the public. Here, we have an interesting moment where the Als actually agree with contemporary art, since most artists, curators and art critics do not consider these fabrics actual works of art either, but documentary remnants of a previous artistic experience.

Art 89%

Modern Art 89%

Cushion 52%

Home Decor 52%

Cushion 56%

Fig. 11: As with *La Roche-guyon* (1909), by Georges Braque, and *Vaas met Bloemen* (1929), by Jan Sluijters, colorful paintings tend to be read as cushions, which reminds us how the visual content of works of art can expand

beyond the museum and fit into more popular, household products. Results such as these also relate to museum shops and their practices of transforming images of artworks into souvenirs.

Dance 66% Dance Pose 66%

Leisure Activities 66% Clothing 51% Evening Dress 51% Gown 51% Robe 51%

#10 OOGST ca. 1932-1937 Victor Dolphijn #10.jpg a group of people posing for the camera 83%

Fig. 12: As with *Oogst* (ca. 1932-193), by Victor Dolphijn, images containing people are interpreted based on the objects that appear within them. In virtually every case with human representations, there were results related to their clothing and other personal objects – including moments in which only those objects were identified and not the humans holding them. Results such as these are a reminder of how part of building an individual's identity in capitalist society is formed with the help of the objects they possess, and the properties of such objects. The same painting was also described as "a group of people posing for the camera" and as a possible "dance pose," which brings us to the idea of displaying these products, and to the following category.

3. Self-promotion

In figurative paintings, AI tends to read people as posing for the camera, which poetically shows how art is a space for human exhibitionism – including selfies and people practicing sports. These results invite us to think of art as an essentially social and egoic practice by human beings, a process of constant self-affirmation.

a group of people posing for a photo 54%

a man doing a trick on a skateboard 26%

Fig. 13: *Javaanse Danser* (ca.1921-1922), by Isaac Israëls, described as “a group of people posing for a photo.” *Slapende Boer* (1936), by Hendrik Chabot, as a skater doing tricks.

4. New titles

Microsoft Azure Computer Vision is an AI service that describes images in short sentences. During our experiment at Van Abbemuseum we performed an exercise in detachment with regard to the artist and their intentions: we began to use these descriptions as new titles for works in the collection. Procedures like this help to demystify the authorship and origin of art objects, creating less fetishized paths of comprehension. Because they are almost always funny, phrases such as these can be valuable material for art classes for non-specialists and young students. Also regarding textual results: Google's AI sometimes identifies texts where there actually are none, thus creating curious descriptions.

Text: 6a83j5 ka Alphabet 89% Text 89% Art 74%

water next to the ocean 71%

a bird flying over a body of water 78%

Fig. 14: According to Google Cloud Vision, the painting *Nature Morte* (1920), by Juan Gris, contains the Georgian word "ნაჩვენს", which translated into English by Google Translate becomes "display." *Augustusbrücke Dresden* (1923), by Oskar Kokoschka, was summed up by Microsoft's AI as "water next to the ocean," adding more poetry to the scene. The *Discovery of the Sardines* (1971), by Ger van Elk, is described as "a bird flying over a body of water," completely reversing the image's idea of aridity.

5. Passages: windows, doors and (why not?) some tables

Poetically, this shows that the space contained by the frame of an artwork creates a space that follows different rules than the space outside of the frame, and that goes on beyond the wall where the artwork is placed—a microsystem that has values and significations of its own. Almost every time there was an interpretation of a "window", there was also a "TV monitor." Although it is a typical case within the first and second groups of this list (Art as objects and IKEA shopping cart), identifying monitors also suggests a depth expansion of the exhibited work. The works read as tables certainly had these results because a framed painting may visually look like a table when seen from above. This recurring result can be seen as an invitation to view paintings from other perspectives, not only face to face or at eye level.

door cabinet shelf blue sink room table

Tags (with confidence): wall 98% window 91% indoor 89% building 88%

#24 COMPOSITION EN BLANC ET NOIR II 1930 Piet Mondriaan #24 .jpg a close up of a window 86%

window frame	(/supporting structure/framework/window frame)	78%
window sash	(/supporting structure/framework/window sash)	63%

#9 COMPOSITIE XXII 1922 Theo van Doesburg #9 .jpg a close up of a door 63%

Fig. 15: *Composition En Blanc Et Noir II* (1930), by Piet Mondriaan, was read as “a close up of a window,” “window frame,” “window sash” and “table.” *Compositie XXII* (1922), by Theo van Doesburg, was read as “a close up of a door.”

6. New temporalities

When AIs do not understand the historical context of an artwork, it allows us to look at art as another kind of object – stripping it away from authorship and historicity. Readings such as these can help in the construction of new narratives of Art History, helping to build new associations between societies from different regions and/or periods. Moreover, some of the new temporalities offered by the AIs do not transform the image into something of a different time but suggest more recent or later moments of what is represented there.

Fig. 16: *Lehrender Christ* (1931), by Ernst Barlach, read as “Buddha,” “sarcophagus coffin,” and associated to images of Ancient Greek sculptures.

© Can Stock Photo

© Can Stock Photo

Fig 17: *Dorp in de lente* (1936), by Constant Permeke, was associated by Google's AI to images of firewood, which suggests a later moment for the trees represented in this bucolic painting. The same AI related the frame of the video *Martha Rosler Reads Vogue* (1982), by Martha Rosler, with the image of a younger woman – it could be a younger Rosler, but in fact is another artist, the Spaniard Cristina Garrido.

7. Personification processes

Often images of artworks were read as flesh and blood people, or as performing human tasks. Images read as people show how the AI's understanding system does not differentiate between the concepts of representation and presence. There were also many cases of sculptures (not necessarily human bodies) that were read as people, which emphasizes the physical strength of large works. It was also interesting to note the human attributions related to some works, such as a sit-down painting—a typical process of prosopopoeia.

Fig. 18: In the painting *Winterbild* (1930), by Max Beckmann, the readings of different AIs complement each other. For Microsoft Azure Computer Vision, the work is “a painting sitting in front of a window.” Seated where, exactly? Probably in one of the two chairs read by Darknet YOLO’s AI. Similarly, the sculpture in the background of the painting *Schilder Met Zijn Vrouw* (1934), by Carel Willink, was read as a person, standing alongside the couple.

8. Visual similarities, new and more democratic possibilities

The fact that AIs associate museum artworks with other images of similar visual forms in their databases⁷ results in a maximized mode of experiments that have long characterized the study of artistic images. Many of the associative processes of these “intelligences” have to do with practices developed by historians such as Aby Warburg and his *Mnemosyne Atlas*.⁸ Due to this, considering these results may be important for expanding this field.

Fig. 19: Some associations of images created approximations of consecrated works with artistic manifestations that are not considered “museum art.” This is the case of *Zomer* (1932), by Constant Permeke, compared by Google's AI to an amateur painting by an unknown artist. *Landschap* (1910), by Jan Sluijters, was interpreted as a possible painting by a child, corroborating with the idea that modern art was interested in abstract and unconscious experiences, as opposed to the academic realism of the late nineteenth and early twentieth centuries.

9. Incomprehensible results, that are very poetic (and that we really like)

Many of the results of the AIs were not fully categorizable into homogeneous groups of results, as is always the case with some works in any museum collection. This shows that art and AI have in common a high load of unpredictability. These results also suggest a possible use of the

⁷ A considerable part of the visually similar images the Google Cloud Vision points to are in low resolution, which brings us to a document that is typical of our times. In *In Defense of the Poor Image*, artist and researcher Hito Steyerl (2009) writes about the importance of analyzing these images: “Poor images are the contemporary Wretched of the Screen, the debris of audiovisual production, the trash that washes up on the digital economies’ shores. They testify to the violent dislocation, transferrals, and displacement of images – their acceleration and circulation within the vicious cycles of audiovisual capitalism. Poor images are dragged around the globe as commodities or their effigies, as gifts or as bounty. Poor images show the rare, the obvious, and the unbelievable – that is, if we can still manage to decipher it. (...) The circulation of poor images feeds into both capitalist media assembly lines and alternative audiovisual economies. In addition to a lot of confusion and stupefaction, it also possibly creates disruptive movements of thought and affect.”

⁸ UCLA professor Christopher D. Johnson writes about this Warburg project: “Begun in 1924 and left unfinished at the time of his death in 1929, the *Mnemosyne Atlas* is Aby Warburg’s attempt to map the “afterlife of antiquity,” or how images of great symbolic, intellectual, and emotional power emerge in Western antiquity and then reappear and are reanimated in the art and cosmology of later times and places, from Alexandrian Greece to Weimar Germany (...) In its ‘last version,’ the *Mnemosyne Atlas* consisted of sixty-three panels (Tafeln). Using wooden boards, measuring approximately 150 x 200 cm and covered with black cloth, Warburg arranged and rearranged, in a lengthy combinatory process of addition and subtraction, black and white photographs of art-historical and cosmographical images.” [More here](#).

AI readings in the expansion of the poetic layers of art, contrary to the productivist and efficiency-focused logic of those who argue that AIs must necessarily provide precise results.

Fig. 20: In the painting, *Moeder en Kind* (1922), by Gus de Smet, an elephant (marked in blue) is read in the room by Facebook's AI. This was also one of the beautiful cases in which a work was read as a "mirror," referring to the idea that the understanding of an artwork is a reflection, a consequence of the way of thinking of those who look at it. Microsoft's AI went beyond the idea of object and added in the conceptual work *B 12000030 = 25 = 16X17 = NOIR BLANC BLEU* (1975), by André Cadere, the information "air" – the true context of art and all other things of this world. But, of course, since not everything is poetry in the AIs, Google has associated this conceptual work to the image of a... lamp.

As made evident by the above examples, our experience in using AIs to interpret images of artworks can be seen as new mode of a practice known as Institutional Critique.⁹ Very widespread since the 1970s, the term is related to a series of procedures that seek to reveal the structures that make the art system function. As the artist Andrea Fraser points out in *What is Institutional Criticism?*, this mode of analytical approximation of art and its elements follows some premises. The first of these is to always consider social context as something of great importance for understanding of what is seen as art. This practice:

(...) takes care of sites above all as *social sites*, structured sets of relationships that are fundamentally social relations. To say that they are social relations does not mean to oppose them to intersubjective or even intra-subjective relations, but to say that a *site* is a social field of these relations. (Fraser, 2014)¹⁰

⁹ The use of AIs to reveal processes of the art system also connects to Conceptual Art practices, especially in the premises written by artist Sol Lewitt (Paragraphs on Conceptual Art, 1967) that defends non-logical and non-specialized artistic understandings. Some of his sentences include: "Irrational judgements lead to new experience"; "The artist may not necessarily understand his own art. His perception is neither better nor worse than that of others" and "There are many elements involved in a work of art. The most important are the most obvious."

¹⁰ Text translated freely from the Portuguese. We used the text as published in the Brazilian magazine *Concinnitas* (2014). This is an expanded version of a German text by Fraser for *Texte zur Kunst magazine* (September 2005).

According to Fraser, art is never the object of art, but rather a network that is interconnected with this object of socially constructed elements. Our AI experiments were successful in revealing this construction. When we take photographs of works away from the context of a museum, art seems to lose its support and the results obtained are almost never related to the art system.

An anecdote by British writer Rudyard Kipling illustrates this detached stance, which is common to both Institutional Critique and to the process of reading images by AIs. Kipling (2013), says that Adam, the first man in the world, sat next to a tree and scratched something in the clay with a small twig. It was "the first rough drawing the world has ever seen" and "a jubilation for the vigorous heart of this man." That's until the Devil appears, walks up and whispers behind the foliage: "It's beautiful, Adam, but ... is this Art?" Because they are not trained to read images of contemporary art, AIs used as practices of Institutional Critique do as the devil did with Adam. Thinking about their results without prejudices and with an open mind is a possible way to distance ourselves from the specialized codes of art and create relevant materials for the critical study of art and its system.

Denaturalizing AI through art: looking critically at algorithms

Another possible course of action is to use all the glitches we have just seen to denaturalize AI's gaze. The results: a list of analyses of every single image of the art collection, when looked at carefully, can work like a reverse engineering of these systems, exposing some of how Computer Vision "sees" the world. Beyond pointing inefficiencies, we can interrogate AI "not only as modes of adjudicating in the world, but also and in their very essence, modes of knowing about the world" (Elish & boyd, 2018, p. 74). They help question AI's positioning as a magic "view from nowhere" (Haraway, 1988), and the power of their epistemologies and ontologies of understanding the world. This critical reading of AI understands it as one of the many ways of understanding the world, that privileges certain values and renders other things invisible. So, then, how do these AI algorithms work?

But what is an algorithm anyway? In an analysis of algorithms as a keyword, Tarleton Gillespie points to how hard it is to precisely define what algorithms are, as the word is used with completely different meanings by computer scientists, artists, social researchers and the broader public: "the term algorithm will turn out to be important in part because it is vacant, a cypher, a ghostly placeholder upon which computational systems now stand" (2016, p.19). There are elements of algorithms though, that seem constant now: "they are designed to work without human intervention, they are deliberately obfuscated, and they work with information on a scale that is hard to comprehend" (Gillespie, 2014, p. 192).

Were we to take Gillespie's observations as a starting point, it is easy to see how the AIs we used give back results without direct human intervention. They spit out their results in milliseconds. In fact, uploading an image to their server takes much longer than for AIs (running

on servers in data centers in, e.g., Oregon/USA) to categorize and analyze artworks. AIs do not need to stare at an image for seconds, minutes, or even days to assess what it means. Instead, they offer multiple results (often conflicting), alongside “confidence ratios”: a percentage of how much the prediction can be trusted, or the “risk of getting it wrong.”

The AI systems we investigated were inscrutable black boxes, and not very inviting to people digging deep into them to understand how they actually work under the rig. Some of them, Facebook Detectron and Darknet YOLO, are open-source, but they still operate in a highly specialized way that is not inviting to a deeper experimentation with the system itself. Moreover, these systems, based on Machine Learning, are not necessarily interpretable to humans; they operate based on the data they have been trained on.

And so, as we looked at the results, a few questions kept coming up (some of which were considered in the previous section):

MORESCHI and PEREIRA: 1. Why so many windows? 2. Why so many tables? 3. Why so many cushions? 4. Why so many close ups? 5. Why so many elephants? 6. Why so many cats? 7. Why so many things related to skate? 8. Why so many computers? 9. Why so many umbrellas? 10. Why so many “Sky plc - company tv cables”?

As previously mentioned, all of these AIs make use of machine learning techniques, whereby algorithms are trained in data on a large scale, which makes it impossible for human comprehension to completely grasp in the same way that machines do. Indeed, the basis of machine learning, and what makes it different from traditional AI, is the idea that algorithms can, with enough data input, build themselves. In the case of image classification, which is what we are discussing, “a dataset is used to train a typical machine learning device, a neural net, and the neural net classifies subsequent images probabilistically” (Mackenzie, 2017, p.4). A “reading” by the AI must be understood as a prediction, which necessarily “depends on classification, and classification itself presumes the existences of classes, and attributes that define membership of classes” (Mackenzie, 2015, p.433).

The questions we ask become very interesting because they deal with how prediction associates two very disparate things with each other, that may or may not look like each other, meaning two completely different “classes.” For example, a 1936 painting of a man sleeping by Hendrik Chabot (see Fig. 13) is read by Microsoft Azure as “a man doing a trick on a skateboard,” which is not what the image depicts. There is, however, some similarity because of the body’s position.

The prototypical construction of a dataset to enable Computer Vision occurred through Fei-Fei Li’s *ImageNet* initiative.¹¹ This project was responsible for gathering a huge number of

¹¹ To be clear, not all of the commercially available AIs we used are based on *ImageNet*, but the project was responsible for triggering a spark. By providing plenty of data about objects and their properties, and creating multiple

images (3.2 million images in total), which were originally organized into categories: 12 subtrees with 5247 synsets (Deng et al, 2009). To define what these categories would be, the project made use of a previously existing structure called *WordNet*. Created in 1985 by Stanford psychology professor George Armitage Miller, with funding from the military and DARPA, *WordNet* was devised to work like a dictionary, but one in which words existed in relation to others (not in alphabetical order). This index of words in a machine-readable logic would become the categories used for all images: dogs, pudding, tracks, excavation. Behind *ImageNet* is a desire to have “more data” for training Machine Learning algorithms, thus allowing them to recognize more objects in images. In the dataset, for example, there are 1289 images of skateboards: “A board with wheels that is ridden in a standing or crouching position and propelled by foot.” (ImageNet, n.d.)

ImageNet takes images from the photo-sharing website Flickr, where users upload their personal photos and often choose to keep them copyright-free, allowing others to use them. It is interesting to think about the origin of dataset’s images and how they make direct use of user-generated data, crowd-produced by all of us. As Cheney-Lippold warns, the “algorithm ultimately exercises control over us by harnessing these forces through the creation of relationships between real-world surveillance data [Flickr, in this case] and machines capable of making statistically relevant inferences about what that data can mean” (2011, 178). In the end, we produce the prediction that produces us.

There is also another argument that can be made considering the origin of these images: art is not a “material” that these AIs are extensively trained on.¹² Images of art and artistic works are only a tiny fraction of what users upload to Flickr, which mainly consists of imagery of contemporary life and social practices, mostly from the United States, where Flickr is most popular (estimated over 25% of its content). We have seen previously how AIs read an immense number of artworks as products, especially department store or home decor products. These are highly accessible, but diverse in value (which is slightly ironic when talking about artworks): tables, shelves, curtains, refrigerator, furniture, cushions, clothing, mobile phones, computer/laptop, TV monitors, etc. These products are part of the wide catalogue of stores such as Ikea, but also present in the modern home imagery. Besides products, the high frequency of sports, cats, dogs, selfies, mirrors, are other indications of the origin of these images in contemporary day-to-day life. The fact that artworks are read as that (although they seldom represent these things) attests to the high frequency of these things in the original training databases.

competitions around it, the field became legitimated and useful for the industry. If an AI does not use it, it is certainly made in connection to it.

¹² Interestingly, this is quite unlike many previous historical attempts of organizing the world: “The *artes memoriae* have been visual techniques of memorization from the rhetorics of antiquity to the Renaissance. Museums – collections, images of picture galleries, catalogs – have since then always dealt with programming material image banks. The struggle for visual knowledge in (literally) the age of enlightenment in the eighteenth century led to visual encyclopedias and their visualizations (like the *planches*, i.e., the visual supplements to the large French *Encyclopédie* edited by Diderot and d’Alembert). Photography, then, has been the medium for switching from perception to technology, creating the first technical image archives, and movies themselves (Hollywood and the rules of image sequences) have been archives” (Ernst, 2012, ebook).

These readings point not only to the origin of images, but to the way images become useful for the clients of commercial AIs. Google, Amazon, Facebook and others build their AI to identify, categorize, and see the world as commodities.¹³ Other frequent results, such as business cards, advertisements and billboard are not products per se, but are also related to contemporary capitalist life. All of this reasserts the way these AI systems embed values that are ideologically capitalist and focused on value-production for the companies that use it.¹⁴ A client featured on Google Cloud Vision's home page, Urban Outfitters, an American multinational lifestyle retail corporation (i.e. clothes shop), uses the AI system to "automate the product attribution process by recognizing nuanced product characteristics like patterns and neckline styles" (Google, n.d.). No wonder the performance fabrics seen before (Fig.10) are almost always read as fashion clothes or accessories, metaphorically transformed into Urban Outfitter's hipster turtleneck long sleeve tees.

AIs also often identify "raciness" in abstract sculptures, and where there are images of women, undressed or not.¹⁵ Although these images come from a context that is considered completely different from a pornographic image, they nonetheless fall under the same category and classification. The consequence of this is felt when these same machine learning systems are used for content moderation, often without subsequent human analysis, thus frequently determining artworks as pornographic (see Gillespie, 2018, introduction).

When we look at *ImageNet* as a prototypical training database, the nude female body, as previously pointed out by John Berger, appears in a sexualized way, an object of the male gaze (see Fig. 21). This understanding of the female body (and naked bodies in general) trickles down to the systems that are used for image recognition in social media and other platforms, through the training data they use. These AI systems further flatten images of artworks by ignoring their context, or what they actually mean.

¹³ And to support the military, but this arguably happens through other systems based on the commercially available ones; or through military grants, which also underlie the whole system.

¹⁴ A possible consequence of this is: why are museums using these same AIs, in so many projects with Google Arts & Culture, for example?

¹⁵ Images with nude women, as in the painting *Liggend Naakt* (1931), by Jan Sluijters, or even dressed, as in *Moeder en Kind* (1922), by Gust de Smet, and *Boerderij* (1919), by Heinrich Campendonk. The same has also happened with images of more abstract sculptures, perhaps because of possibly phallic shapes, such as in *My neck, my back curve silently* (1930), by Karin Arink.

Fig 21: Screenshot of some of the images that are under the synset "Nude, nude painting" ("A painting of a naked human figure") in ImageNet. There are a total of 1229 images, most of which represent a sexualized, naked, thin, white woman.

Taking this even further, there is a more radical argument to be made regarding the obsession of these AIs with nakedness through the “racy” category itself. Susan Leigh Star and Geoff Bowker, in their classic book *Sorting Things Out* (2000), discuss classification as a way of seeing the world and naming it. They are interested in “how basic categories and standards are formed, and how they are formed as ordinary,” (p. 35) in order to examine the power these systems embody. They state:

The advantaged are those whose place in a set of classification systems is a powerful one and for whom powerful sets of classifications of knowledge appear natural. For these people the infrastructures that together support and construct their identities operate particularly smoothly (though never fully so). For others, the fitting process of being able to use the infrastructures takes a terrible toll. To “act naturally,” they have to reclassify and be reclassified socially (p. 225).

As further discussed by Kate Crawford (2018), AI’s categories have a lot of power, as they become embedded in our everyday world. This certainly happens as AI systems become carelessly incorporated into public services, such as policing (Brayne, 2017), but also more subtly through social media, image recognition, and other elements of what we consider day-to-day life. It is imperative to recognize that categories such as racy, pornographic etc. do not engage with any context or meaning behind images, and thus can ever only be of partial use to really engaging with images. But even when the machine is not sure, it generates results; leaving no possibility of not knowing, contrary to art. This epistemology confounds correlation

with causation: formal similarities are the way to understanding, and if something looks like something else, it must be that. Jesus Christ being read as a man (or even a woman) shows exactly how the surface-level reading does not engage with a more embodied, contextual reading of reality. This is a crucial way of showing that AI represents only one of the many ways of knowing and working through the world.

Als and the exploitation of labor in technological capitalism

MORESCHI: On the last night of work, I asked Gabriel an essential question: What, in this process, was not machine, but human? That's when he told me about Turkers. People who are paid meager cents of a dollar to categorize images in systems such as ImageNet. We wanted to talk to them to understand how they made their choices. We requested a survey of the collection's images, sent to them without any additional information. We asked them to describe the images and whether or not they considered those images to be art.

In our experiment with these “lower levels” of the AI stack – the Turkers – we were able to get a glimpse of how datasets such as *ImageNet* are built. Although AI systems are branded as an external, non-human, objective, “view from nowhere” (Haraway, 1988), as its mathematical and statistical analysis claims its way of understanding the world as the only possible way (rationalization), they in fact have their origins in these workers, spread out throughout the world (75% in the US, 16% in India, and many other countries in a lesser degree; see Difallah et al, 2018), who are being paid meager salaries. It is imperative to make clear that AI's backbone is, in fact, constituted by human thought, labor and clicks on a screen: these systems are built on “workers' invisibility” (Irani & Silberman, 2013). Platforms such as Amazon Mechanical Turk “commercialize the thesis that humans are important cogs in computational machines” (Finn, 2017, ebook) and treat them as such.

More than invisible, Gray and Suri (2019) define the work of Turkers as “ghost work” in their book-long review of on-demand digital work performed by this vast, invisible human labor force. This term highlights the central irony of how on-demand work is prevalent today, hidden in the shadows of so-called artificial intelligence. It is this “ghost” aspect that makes it so that we often cannot see what is behind the scenes, hiding AI's materiality (see Crawford and Joler, 2018). As Gray and Suri explain, “MTurk workers are the AI revolution's unsung heroes. (...) Humans trained an AI only to have the AI ultimately take over the task entirely” (2019, p. 8).

What do you think	Description	Is this Art?	Have you been paid fairly?	Completion time (seconds)	Reward
I think it is a classy nude painting.	<ul style="list-style-type: none"> nude woman laying on bed with back facing viewer woman in the nude on a bed with back facing viewer painting of a nude woman on a bed with back facing viewer 	yes	yes	115	\$0.40
It is very well composed and has a great palate of color	<ul style="list-style-type: none"> woman lies prone in bed woman resting, soundly a woman at rest 	yes	yes	132	\$0.40
I think that the woman is ugly.	<ul style="list-style-type: none"> There is a woman naked, facing the bed, laying down. Laying down on white sheets is a naked woman on a bed. The woman without clothing is holding her arm over a pillow, while naked. 	yes	no	159	\$0.40
Very sexual Sleeping	<ul style="list-style-type: none"> Sleeping on the bed. Very Nude showing. Without dress sleeping on the bed. 	yes	yes	161	\$0.40
It's bland.	<ul style="list-style-type: none"> Woman sleeping on couch Naked woman sleeping on couch Nude form from behind 	yes	no	124	\$0.40

Fig 22: Painting *Liggend Naakt* (1931), by Jan Sluijters, as described by 5 different Turkers. One of them thinks it is “very sexual,” while another says “the woman is ugly.” All of them take just over two minutes to complete their task. They all agree: the image is Art.

When asked to describe images of artworks from the collection, Turkers presented quick, direct, and not particularly analytical readings of the visual material. As they attempt to complete the HIT (Human-Intelligence Task) in the shortest time possible, which, in this case, ended up being just over a couple of minutes, they engage with the image with a particular distance, not looking for or relating it to any context that is not offered, or even from their own perspective. The descriptions attempt to maintain a detachment and objectivity. Their answer to what they think of the artwork is also quick and direct, exposing the very brief relation with the image: it is either bland, sexual, ugly, classy, or something else.

The “artificial artificial intelligence”, as Turkers are described by Amazon (Finn, 2017, ebook) give many of the same responses as we have seen with the AIs that were previously analyzed: they seem to treat things based on what they look like, a “first-impression,” made in the haste of their race for generating enough income. What is essential in these responses, though, is that Turkers “at the heart of this system not only take on the challenge of endless micro-tasks managing ambiguity—they also take on the affective work of acting as a human element inside of a computational application” (Finn, 2017, ebook). By this, it is meant that HITs, as Human Intelligence Tasks, expect from Turkers “a response that is both mechanically reliable and reliably human... constantly negotiating between computational and cultural regimes of meaning.” In the end, these descriptions and opinions about artworks may sound like they are coming from the AIs, as they are made manifest through a direct, objective, distant format that fits into the expected computational tone, but they are generated by a real human seeing, their human mind, culture, idiosyncrasies, and context.

Artificial artificial intelligence is only valued because it is intrinsically different from artificial intelligence. Unlike AI, Turkers are humans: they have feelings, opinions, agency, families, and bills to pay. Although directly participating in a precarious, neoliberal, pro-employer marketplace, where they have very little say or protection, and are identified as an alphanumeric code, their responses must be understood as such: they come from a different position of power than that of AI, a position which must be understood critically. When Irani & Silberman (2013) asked Turkers what their major concerns were with the platform, workers pointed to Amazon and how it does not really care about them, exposing a clear power imbalance where the needs of workers are not prioritized. As put by one of their Turker respondents, “*I don’t care about the penny I didn’t earn for knowing the difference between an apple and a giraffe, but I’m angry that MT will take requester’s money but not manage, oversee, or mediate the problems and injustices on their site.*” (p. 615) The difference between an apple and a giraffe, between an artwork or something else, when analyzed by Turkers, and later by the AIs that are built from their labor, must all be understood as part of this larger system of worker invisibility, low pay, lack of rights, instability, and Amazon’s political economy of monopoly.

PEREIRA: Thinking about non-specialists in AI, we found it important to show the human labor behind these machines, as a way of raising awareness of what AI actually is. In the short video we made, we tried to use an accessible language (unlike much of the video art produced today) to show this infrastructure, explain the role of Turkers, and use their own voices to read descriptions they would give to images of artworks, thus foregrounding their contribution to AI systems. At the same time, we denaturalize artificial intelligences as both AI and human Turker readings are shown to have similarities (as well as differences). This relates to academic calls for increasing literacy on datafication and its processes.

This process not only demonstrated the importance of bringing these discussions of critical infrastructure to society in general, but also the importance of creating interdisciplinary research teams that can operate from different backgrounds and perspectives. In this moment

where image recognition algorithms become embedded across society through social media platforms, as well as through the use of technology by museums and archives, we must be attentive as to how these algorithms operate, what they obfuscate, and which kinds of invisible labors they rely on. And, more radically, we must be concerned on the increasing pervasiveness of the logic of the algorithm, the "if-then" causal understanding of categories that confuse seeing with understanding, correlation and causation, nudity and pornography. As described by Simanowski:

If reason is reduced to formal logic, then any discussion becomes unnecessary because if-then relations do not leave any room for the "but" or the "nevertheless," nor for ambivalence, irony, or skepticism; algorithms are indifferent to context, delegating decisions to predetermined principles. (...) If one clears away the gaps and strategic ambivalences from communication, thinking becomes calculation, and all that remains are determinative decisions – or simply mistakes in the combinatorial calculations (2016, p. 55).

Our experiment with denaturalizing AI through art points exactly to this intrinsic instability of the "artificial" in "artificial intelligence": the lack of seeing more than what things are. The artistic space, with its openness to different logics and "ways of seeing," serves as fertile ground for exposing AI's binary, capitalist, and value-laden gaze. At the same time, however, it also provokes us to reconsider, to be more creative and explorative. What are we really seeing in an artwork? And how could we ever be so sure?

Conclusion

The process of seeing paintings, or seeing anything else, is less spontaneous and natural than we tend to believe. A large part of seeing depends upon habit and convention. (Berger, 2008)

We have explained in this article how commercially available AIs can work, through their glitches, to level and reimagine artworks, giving us a fresh set of eyes for understanding art. At the same time, these very glitches can serve as a peek down the stack of ever-present algorithmic image recognition systems, helping us speculate on their inner workings.

We present our research in this article, alongside different artistic outputs (video, exhibition, publication, and intervention), as a way of experimenting with research methodologies and interdisciplinary positions. This research is also presented at the museum whose archive we analyzed, creating a space for reflectivity within an elitist, codified space. It is a crack inside of the museum for self-reflectivity, through a way that is typical of institutional critique.

This research can also contribute to art education in museums: we propose, through this text and the other results of this research, that people, as exhibition visitors, experiment with the distanced look of AI as a way of critically thinking about the art system. We hope it helps to create new relationships, openings, and connections for people who are non-specialists to

explore it critically.¹⁶ Instead of museums using commercially available AI from Big Tech, in an uncritical way, why not make more radical and creative uses of technology? It is undeniable we need more critical and creative work, and less superficial gimmicks, shallow interactions and flashy lights.

As AI continues to expand, change and “improve,” we understand these results in their limitations: they are a snapshot of how they worked when we tested them. As more and more data is produced in our everyday lives and interactions, and as companies hoard and process increasingly greater swaths of data to continually train these models, it also continually changes how art is read. In our experience throughout our research, we have seen both minor and major changes, which we think point to the simultaneously productive and critical instability of AI and Art.

References

- Berger, J. (2008). *Ways of Seeing*. London: Penguin.
- Bowker, G. C., & Star, S. L. (2000). *Sorting things out: Classification and its consequences*. Cambridge: MIT Press.
- Brayne, S. (2017). Big Data Surveillance: The Case of Policing. *American Sociological Review*, 82(5), 977-1008. doi:10.1177/0003122417725865
- Cheney-Lippold, J. (2011). A New Algorithmic Identity. *Theory, Culture & Society*, 28(6), 164-181. doi:10.1177/0263276411424420
- Cheney-Lippold, J. (2017). *We Are Data: Algorithms and The Making of Our Digital Selves*. New York: NYU Press.
- Cox, G. (2017). Ways of Machine Seeing: an introduction. *A Peer-Reviewed Journal About*, 6(1). Retrieved from <http://www.aprja.net/ways-of-machine-seeing-an-introduction/>
- Crawford, K. (2018, March 6). *AI Now: Social and Political Questions for Artificial Intelligence*. Distinguished Lecture presented at the Tech Policy Lab/University of Washington, Seattle, WA. Retrieved 19 Feb 2019, from <https://youtu.be/a2IT7gWBfaE>
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., & Fei-Fei, L. (2009). Imagenet: A large-scale hierarchical image database.

¹⁶ Before the Van Abbemuseum experience, Moreschi had already carried out experiments that prioritized understandings not specialized in works of art. *Painters* (2014-2017) is a site-specific and collaborative installation, where he visited construction sites around the exhibition space and invited wall painters to collaborate on the creation of an artwork on the museum wall. In 2018, for the first time, we include the opinion of non-specialized people and the interpretations of AIs in the same artistic investigation. In *Another 33rd Biennial*, a project commissioned by the 33rd São Paulo Biennial, we created alternative ways of understanding this exhibition. One of the project's questions was "What do the nonspecialists have to say?" and included audio guides with Biennial staff (forklift operator, cleaning staff and museum attendants), as well as galleries with interpretations of the Biennial art system read by a set of AIs.

- Difallah, D., Filatova, E., & Ipeirotis, P. (2018). *Demographics and Dynamics of Mechanical Turk Workers*. Proceedings from Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining - WSDM '18, New York, New York, USA.
- Elish, M. C., & boyd, d. (2018). Situating methods in the magic of Big Data and AI. *Communication Monographs*, 85(1), 57-80. doi:10.1080/03637751.2017.1375130
- Ernst, W. (2012). *Digital Memory and the Archive (Electronic Mediations)*. Univ of Minnesota Press.
- Eubanks, V. (2018). *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor*. NY: St. Martin's Press.
- Finn, E. (2017). *What Algorithms Want*. Cambridge: MIT Press.
- Fraser, A. (2014). O que é Crítica Institucional? *Concinnitas*, 24(2). Retrieved from <https://www.e-publicacoes.uerj.br/index.php/concinnitas/article/viewFile/18731/13645>
- Gillespie, T. (2014). The relevance of algorithms. In T. Gillespie, P. J. Boczkowski, & K. A. Foot (Eds.), *Media technologies: Essays on communication, materiality, and society*. Cambridge, MA: MIT Press.
- Gillespie, T. (2016). Algorithm. In B. Peters (Ed.), *Digital Keywords: A Vocabulary of Information Society and Culture* (pp. 18-30). Princeton University Press.
- Google. AI & Machine Learning Products: Cloud Vision. Retrieved 18 Feb 2019, from <https://cloud.google.com/vision/>
- Haraway, D. (1988). Situated Knowledges: The Science Question in Feminism and the Privilege of Partial Perspective. *Feminist Studies*, 14(3), 575. doi:10.2307/3178066
- ImageNet. (n.d.). Skateboard. Retrieved 7 Mar, 2019 from <http://imagenet.stanford.edu/synset?wnid=n04225987>
- Irani, L. C., & Silberman, M. S. (2013). *Turkopticon*. Proceedings from Proceedings of the SIGCHI Conference on Human Factors in Computing Systems - CHI '13, New York, NY, USA.
- Johnson, C. D. *About the Mnemosyne Atlas*. Retrieved 3 Jan 2019, from <https://warburg.library.cornell.edu/about>
- Kipling, R. (2013). The complete works of Rudyard Kipling. Birmingham: Inktree.
- Lewitt, S. (1967). *Paragraphs on Conceptual Art*. New York: Artforum
- Mackenzie, A. (2015). The production of prediction: What does machine learning want. *European Journal of Cultural Studies*, 18(4-5), 429-445. doi:10.1177/1367549415577384

- Mackenzie, A. (2017). *Machine Learners: Archaeology of a Data Practice*. Cambridge: The MIT Press.
- O'Neil, C. (2016). *Weapons of Math Destruction*. Crown Books.
- Olah, C., Satyanarayan, A., Johnson, I., Carter, S., Schubert, L., Ye, K., & Mordvintsev, A. (2018). The Building Blocks of Interpretability. *Distill*, 3. doi:10.23915/distill.00010
- Paglen, T. (2016). Invisible Images (Your Pictures Are Looking at You). *The New Inquiry*. Retrieved from <https://thenewinquiry.com/invisible-images-your-pictures-are-looking-at-you/>
- Powles, J., & Nissenbaum, H. (2018). The Seductive Diversion of 'Solving' Bias in Artificial Intelligence. Retrieved from <https://medium.com/s/story/the-seductive-diversion-of-solving-bias-in-artificial-intelligence-890df5e5ef53>
- Simanowski, R. (2016). *Data Love: The Seduction and Betrayal of Digital Technologies*. Columbia University Press.
- Steyerl, H. (2009). In Defense of the Poor Image. *e-flux*, 10. Retrieved from <https://www.e-flux.com/journal/10/61362/in-defense-of-the-poor-image/>
- Tomkins, C. (1998). *Duchamp: A Biography*. New York: Holt Paperbacks.
- Warburg, A. (2010). *Atlas Mnemosyne*. Madrid: Akal.