

Artificial Intelligence and Institutional Critique 2.0: Unexpected Ways of Seeing with Computer Vision¹

Gabriel Pereira² and Bruno Moreschi³

Author's Accepted Manuscript of an article to be published on the "Ways of Machine Seeing" Special Issue of the AI & Society Journal (edited by Mitra Azar, Geoff Cox, and Leonardo Impett). Final version available on: <https://doi.org/10.1007/s00146-020-01059-y>

Introduction: Duchamp's Fountain is a urinal

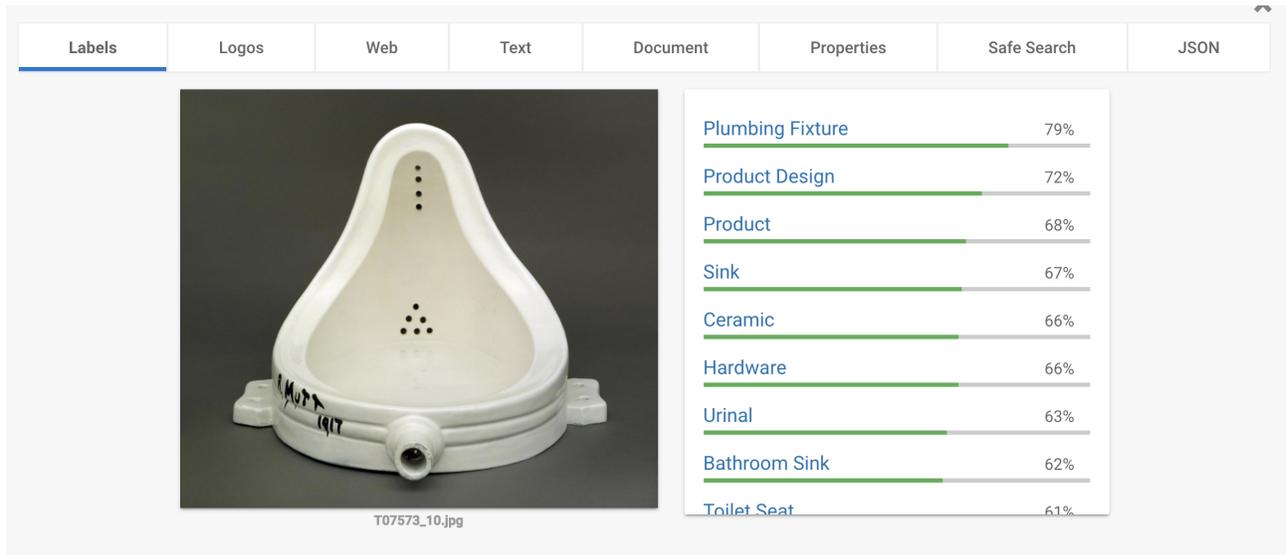


Fig. 1: Print from the e-mail that Moreschi received from Pereira, with the reading of Duchamp's Fountain by the AI Google Cloud Vision. None of the results considered the image a photograph of an artwork.

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.

¹ Portions of this article appeared in a different version in the Van Abbemuseum's "Deviant Practice Research Programme 2018-19" electronic publication (CC BY-NC-ND 4.0). We'd like to thank the Special Issue editors, reviewers, and others that have contributed and supported this research project. Special thanks to Giselle Beiguelman, the staff of the Van Abbemuseum, and the Center for Arts, Design and Social Research.

² PhD Fellow in Information Studies and Digital Design at Aarhus University (DK): <http://gabrielpereira.net/>

³ Postdoctoral fellow at the Faculty of Architecture and Urbanism at the University of São Paulo (FAUUSP): <https://brunomoreschi.com/>

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 may be familiar with this controversial story: after having lunch 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, the juror George Bellows cried out: "It's indecent! It's indecent! We can't show this. This thing is nothing more than what it is."⁴

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. His act of transforming 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. (...) 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 the famous photographer Alfred Stieglitz, to be

⁴ The story of the *Fountain* has been under dispute in recent years. Research by historian Irene Gammel indicates there is evidence that the piece was actually created by dada artist Baroness Elsa, although Duchamp was that ultimately proposed it to the jury.

photographed using the same method as a sculpture. 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 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 are not embedded with what these artworks mean, their context, history and subtext. But this is not necessarily a problem in itself, as 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. 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, was 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. 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.⁵

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 open-source 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) through this [link](#).

⁵ 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.

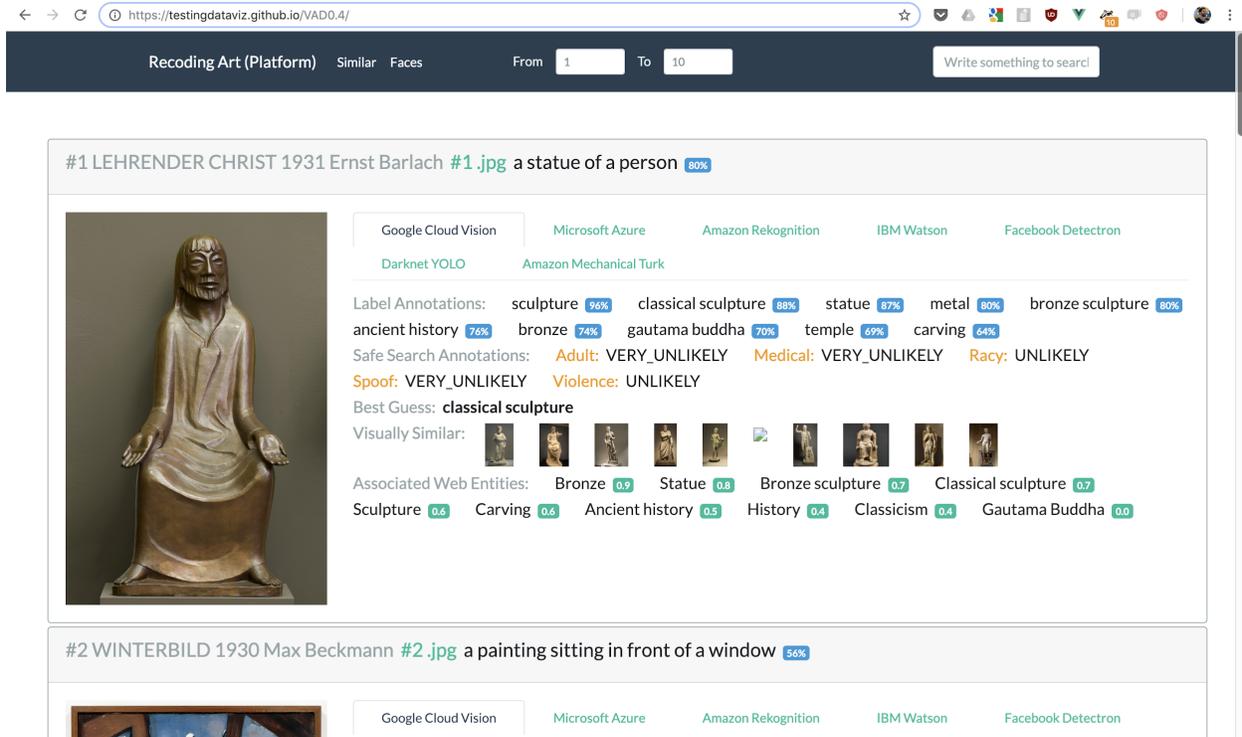


Fig. 3: Recoding Art, an open-source website with the Van Abbemuseum’s art works reading by AIs.

This centralized AI results interface was created prior to the official start of Moreschi's residency at Van Abbemuseum. 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.

“Honor thy error as a hidden intention.” This card, from the set of cards *Oblique Strategies* created by Peter Schmidt and Brian Eno to aid in the artistic process, epitomizes our methodological approach. This 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. Minimal attention was put on results that were “true” or “correct.” On the contrary: we decided to value the unexpected outcomes.

Although the struggle for algorithmic auditing, accountability, and ethics is very important, given the amount of problems AI is already causing and how these errors affect people (especially underserved minorities and marginalized communities), focusing exclusively on “solving bias” may serve 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 as a way of critically and imaginatively speculating on the machinations of the

systems of both AI and Art.⁶ It is about critiquing the operating logic of algorithms and showing that they are neither a “given,” nor “certain,” and thus complicating the “mythical, objective omnipotence” (ibid) that they so often evoke.

During the analysis process, Moreschi classified works by results with similar characteristics. The same work could be included in more than one of the groups, which are discussed in the next section. 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*, a short film which integrates the outcomes of this research.

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: A possible 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, 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-

⁶ It is worth noting that the subject matter of this research (artworks from a collection) is particularly suitable for this approach, since, unlike predictive policing and other egregious algorithmic systems, these errors do not directly cause harm.

called aura from the art object and transform a very important art collection into an assortment of easily recognizable objects.

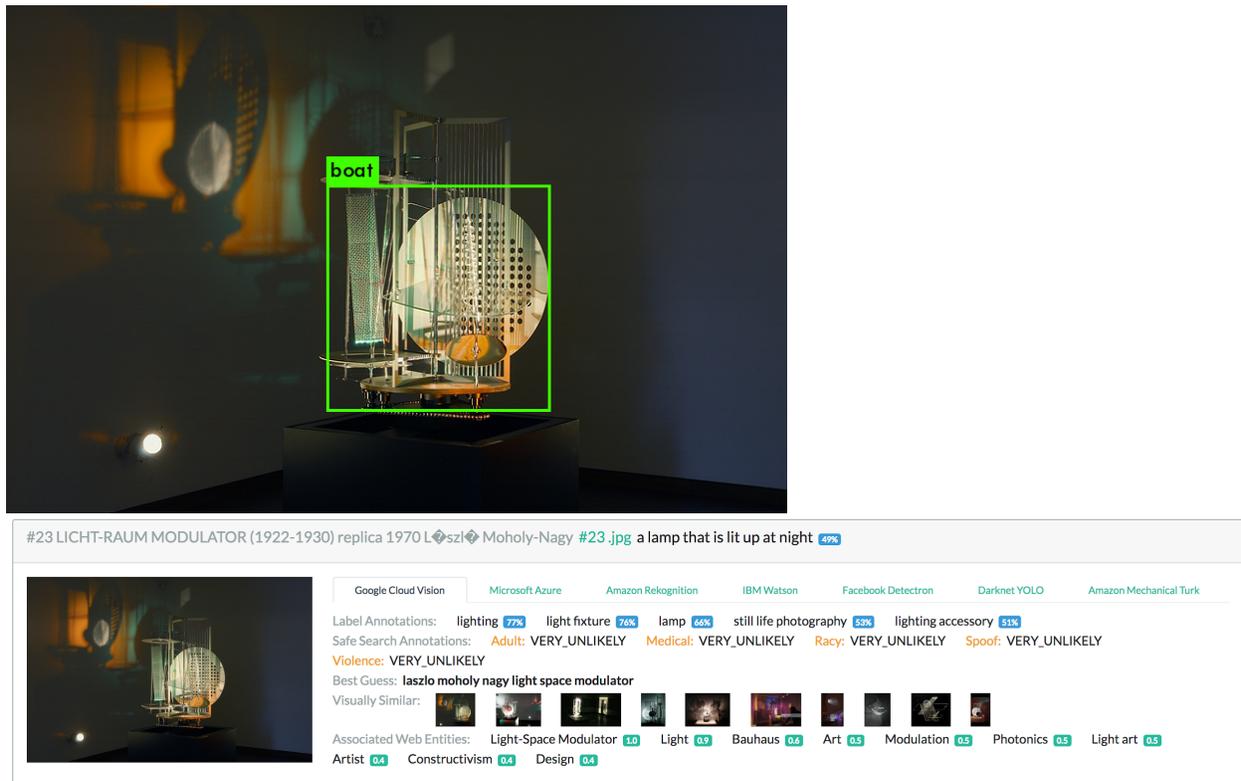


Fig. 4: *Licht-raum Modulator* (1922-1930, replica 1970), by László Moholy-Nagy, described by Microsoft's AI as "a lamp that is lit up at night," a similar result to Google's AI ("lighting," "lamp," "light fixtures"). Darknet YOLO (open-source AI) goes further and sees the work as a possible "boat", which helps us to construct an interesting speculative hypothesis: that, from the interpretive logic of AIs, the circular reflection on the wall can be a full moon in the high sea.

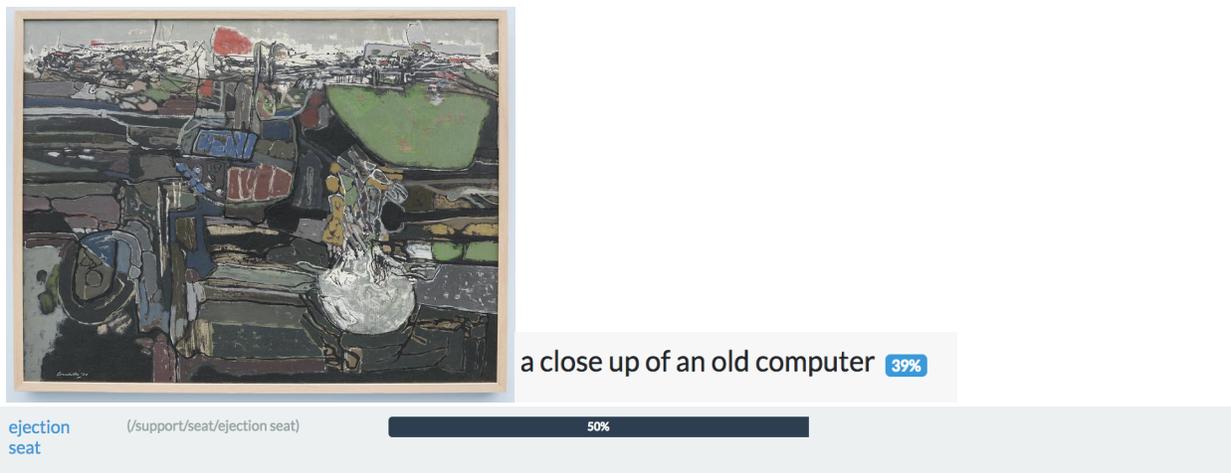


Fig 5: *Aux Abords De La Grande Cité* (1960), by Corneille, read as "ejection seat" and "a close up of an old computer." A considerable part of the works analyzed by Microsoft Azure's AI are understood as approximations of something. Since the AI's gaze does not operate from the human logic of physical 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 99% art 96% orange 92% painting 86% 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

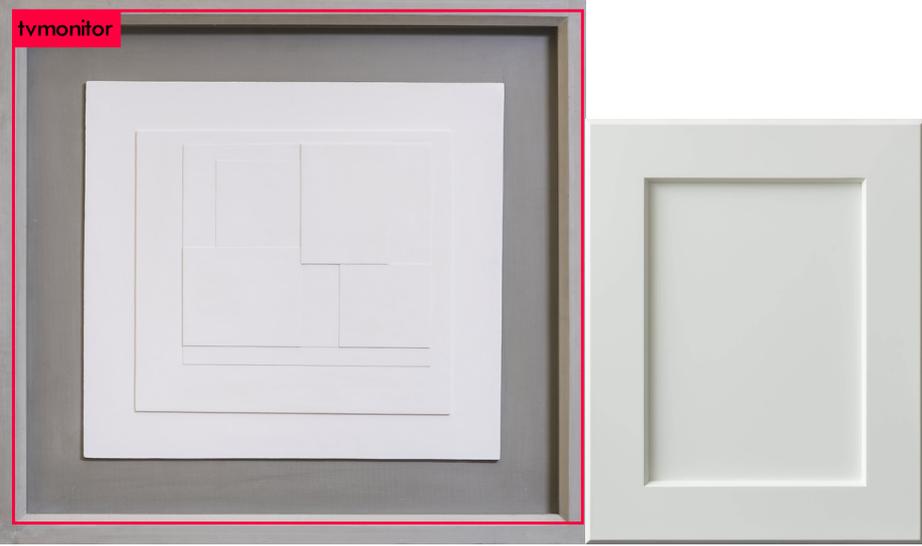
Visually Similar: 

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 a United States 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-artistic state.



tvmonitor

passepoutout (frame) (/supporting structure/framework/passepoutout (frame)) 60%

frame (/supporting structure/framework/frame) 60%

Label Annotations: picture frame 77% window 64% paint 54% door 52% rectangle 52% molding 51%

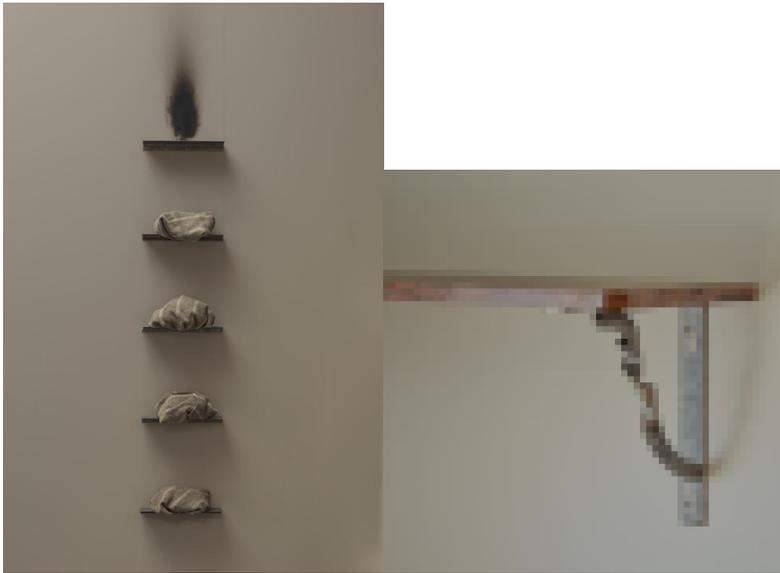


Fig. 7: Some images received interpretations that prioritize the physical structures that protect or support the works ("picture frame," "framework," "supporting structure") rather than their legitimized artistic contents. This occurred with *White Relief* (1936), by Ben Nicholson, often understood by 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 90%) 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

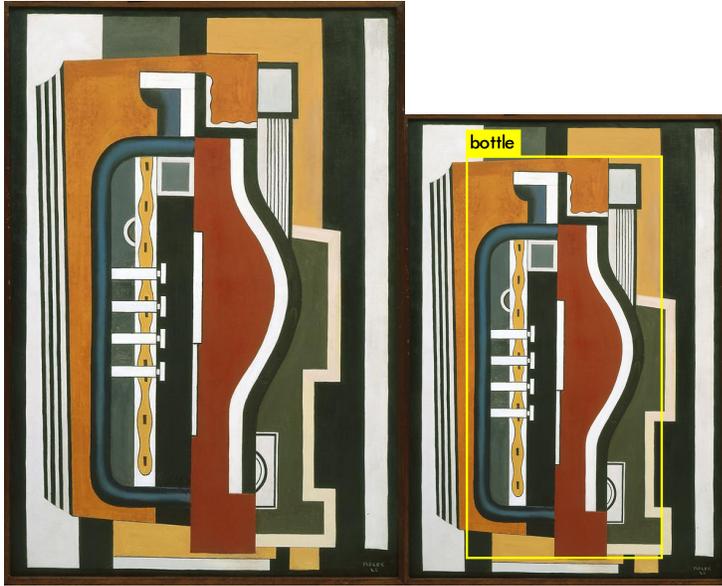
Accent Color: [redacted]

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 77% 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%

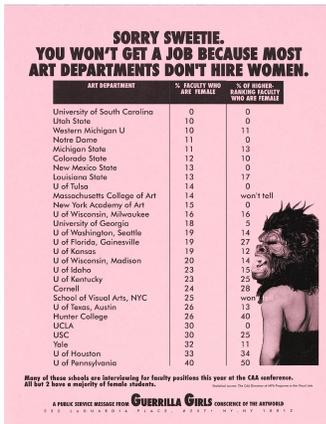
tetraskelion 0

66%



Label Annotations: font 78% product design 42% brand 42%
 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 0.4 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 en

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.



#46 GLEICHZEITIGKEITSSTÜCK (NR. 23, 1. WERKSATZ) 1967 Franz Erhard Walther #46.jpg a white shirt 60%

military uniform

()

55%

Label Annotations: white 96% product 82% product 73% beige 62% product design 59% handbag 51%



Label Annotations: green 95% khaki 84% product 76%



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 AIs, which makes some sense, since many of them were worn by artists and/or the public. Here, we have an interesting moment where the AIs 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. 1932-1937), 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 (and its contemplation) 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%

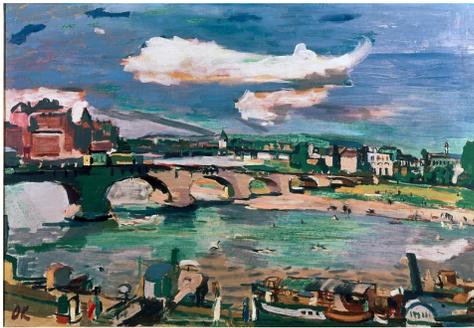
Fig. 13: *Javaanse Danser* (ca.1921-1922), by Isaac Israëls, described as “a group of people posing for a photo.” *Slapende Boer* (1936), by Hendrik Chabot, as a skater doing tricks.

4. New titles

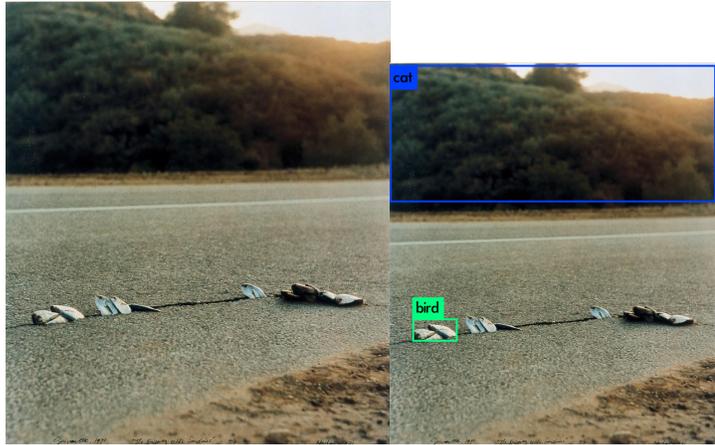
Microsoft Azure Computer Vision is an AI service that describes images in short sentences. During our experiment at Van Abbemuseum we performed an exercise in detachment with regard to the artist and their intentions: we began to use these descriptions as new titles for works in the collection. Procedures 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: სანჯიბ ka Alphabet 89% Text 89% Art 74%



water next to the ocean 71%

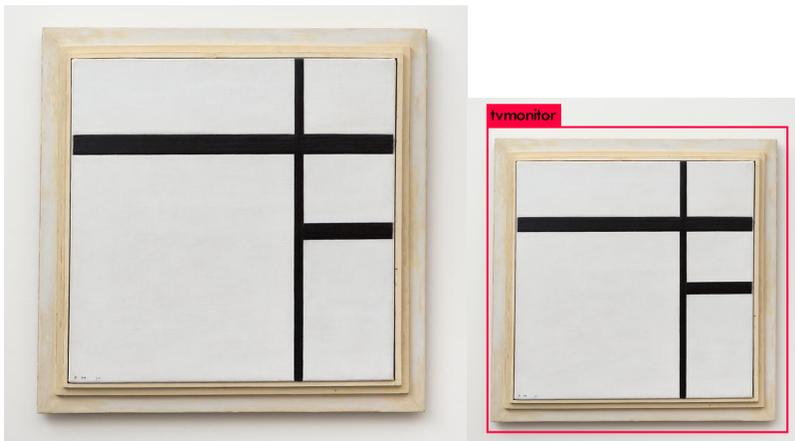


a bird flying over a body of water 78%

Fig. 14: According to Google Cloud Vision, the painting *Nature Morte* (1920), by Juan Gris, contains the Georgian word "საჩუბი", 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.



Label Annotations: sculpture 96% classical sculpture 88% statue 87% metal 80% bronze sculpture 80% ancient history 76% bronze 74% gautama buddha 70% temple 69% carving 64%

Labels: Art 95% Sculpture 95% Buddha 88% Worship 88%

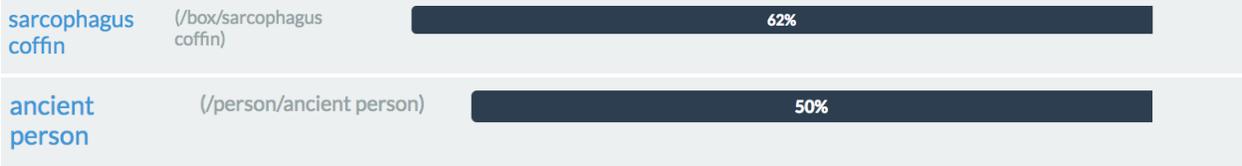


Fig. 16: *Lehrender Christ* (1931), by Ernst Barlach, read as “Buddha,” “sarcophagus coffin,” and associated to images of Ancient Greek sculptures.



© 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.

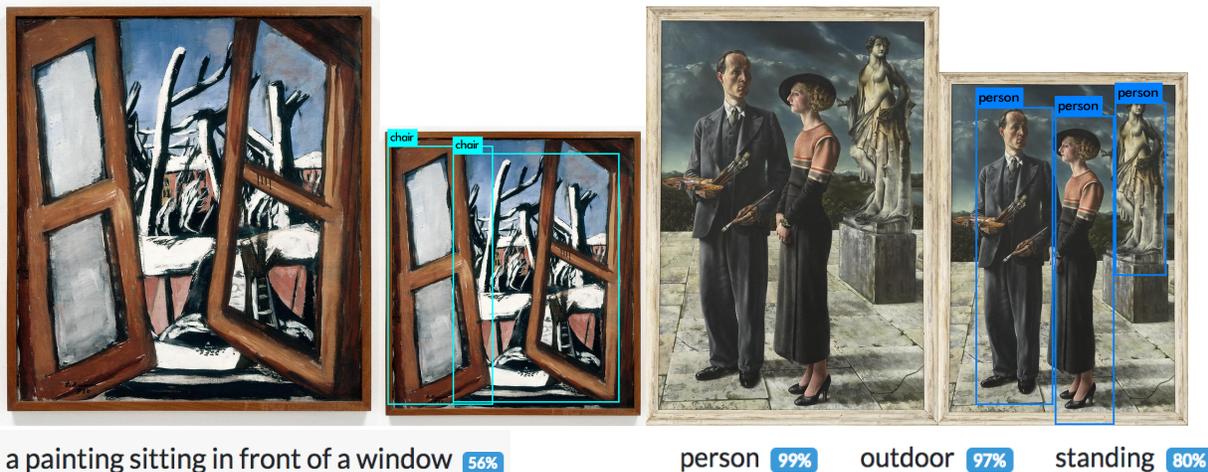


Fig. 18: In the painting *Winterbild* (1930), by Max Beckmann, the readings of different AIs complement each other. For Microsoft Azure Computer Vision, the work is “a painting sitting in front of a window.” Seated where, exactly? Probably in one of the two chairs read by Darknet YOLO’s AI. Similarly, 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 databases 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* (2010). 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 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. The term is related to a series of procedures that seek to reveal the structures that make the art system function. Historically, the institutional critique practice operates from the critical repertoire of conceptual art and conceptualisms of the 1960s and 1970s, especially in the expanded concept of art (Freire 2006).

According to Andrea Fraser (2005), this mode of analytical approximation of art and its elements follows the premise of considering the social context as intrinsic to art – to her, 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 elements of this construction: when we took photographs of works away from the context of a museum and into that of Computer Vision algorithms, art seemed to lose its support, and the results obtained were almost never related to the art system. The highly specialized and elitist codes that permeate the artistic field – well protected and validated by powerful actors such as art institutions, curators, gallery owners, specialized critics and even artists – were visibly ignored by Computer Vision, which instead offered different paths for understanding the artworks.

In other words, when looked at without prejudices and with an open mind, seeing art through the glitchy results of Computer Vision allowed us to distance ourselves from specialized meanings and create relevant materials for the critical study of artistic works and the system they are inserted in. Bringing art and AI together in a critical way serves not only to reveal the latent power structures of the artistic field, but also to democratize it, opening its meaning and significations to the people who engage with art (as espoused by institutional critique). This way of seeing and its potential should be understood and appreciated in its partiality, though, as it intentionally doesn't engage with the wide context of art history and its specialized discourses.

The unexpected ways Computer Vision sees art, both by levelling and expanding the potential meanings of artworks, is also particularly innovative and relevant in a moment when visual culture has changed form to hybrids of human-machine cognition and "machine-to-machine seeing" (Paglen, 2016), with a plethora of limitations and problems which we address in the following section.

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 inefficacies, 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.⁷

The commercial AI systems we used, as any other algorithm, 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). The AIs did 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. Moreover, they do not expose how they actually work under the rig, being presented as inscrutable black boxes: their processes are not directly interpretable to the user. Although some of them, i.e. Facebook Detectron and Darknet YOLO, are open-source, they still operate in a highly specialized way that is not inviting to a deeper understanding of the system. And so, as we looked at the results and tried to make sense of them, 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”?

These questions become interesting as they expose the issues that underlie contemporary machine learning training datasets. 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 by making use of large-scale data. In the case of image classification, which is our focus here, “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 classes within the training datasets, along with the images that compose these classes, are responsible then for defining what the AI can “see.” What’s quite interesting is how two very disparate things (or “classes”) may become approximated with each other, whether they look like each other or not for our human eyes. 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, which becomes visible after engaging with the AI’s interpretation.

⁷ For more scholarship critically exploring the limitations of computer vision’s ways of seeing see, e.g.: Mintz et al (2019), Buolamwini & Gebru (2018), Crawford & Paglen (2019), and other articles in this Special Issue.

The prototypical construction of a dataset to enable Computer Vision occurred through Fei-Fei Li's *ImageNet* initiative.⁸ This project was responsible for gathering a huge number of 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 took 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, the ways of seeing of these computer vision AIs are directly tied to the origin of their datasets.

As can be noted in our results, art is not a "material" that commercial AIs are extensively trained on. 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.

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 AIs to identify, categorize, and see the world as commodities.⁹ Other frequent results, such as business

⁸ 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 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.

⁹ 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.

cards, advertisements and billboards 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.¹⁰ 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.

Als also often identify "raciness" in abstract sculptures, and where there are images of women, undressed or not.¹¹ 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). When we look at *ImageNet* as a prototypical training database, the nude female body appears in a sexualized way, an object of the male gaze (see Fig. 21), not too differently than throughout much of the History of Art (see Pollock & Parker, 2013). This understanding of the female body, among other problematic categories/classes, trickles down to the systems that are used for image recognition in social media and other platforms, through the training data they use (see Crawford & Paglen 2019). This is just one way that AI systems flatten images of artworks by ignoring their context, or what their images mean in a broader sense.

¹⁰ A possible consequence of this is: why are museums using these same AIs, in so many projects with Google Arts & Culture, for example?

¹¹ 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.

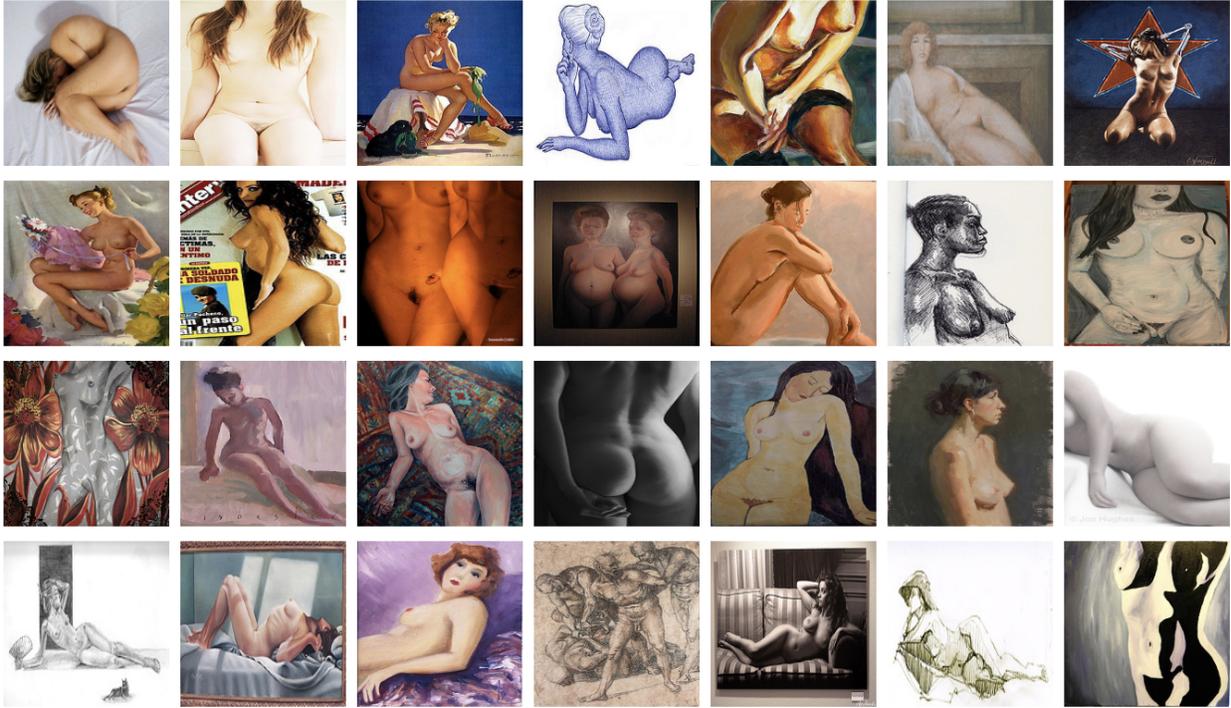


Fig 21: Screenshot of some of the images that are under the synset "Nude, nude painting" ("A painting of a naked human figure") in ImageNet. There are a total of 1229 images, most of which represent a sexualized, naked, thin, white woman.

Taking this even further, we argue that the obsession of these AIs with nakedness through categories such as “racy” (from Google Cloud Vision) and “pornographic” conceals their lack of contextual comprehension through a numerical result, thus normalizing problematic categories. Star and 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).

AI’s categories have a lot of power, as they become embedded in our everyday world, as previously discussed by many other scholars (see Noble 2018; Eubanks 2018; O’Neill 2016; D’Ignazio & Klein 2020; Crawford 2018). This certainly happens as AI systems become carelessly incorporated into public services, such as policing (Brayne 2017), but also more subtly through image recognition in our 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 visual culture. But even when the machine is not sure, it generates results; leaving no possibility of not knowing, contrary to art. As discussed by

Amoore (2019), by transforming doubt into weighted probabilities, a number between 0 and 1, “the single output of the machine learning algorithm is rendered as a decision placed beyond doubt; a risk score or target that is to be actioned.” 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.

Ais 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 “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 human: 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 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).



Fig 23: A scene of *Recording Art*, a short film resulted from this research, wherein turkers analyze works of art. The film premiered at IDFA Competition for Short Documentary 2019, Amsterdam.

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." (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 and critique their limited perspectives.

Art is historically formed through internal deconstructions, some of which become paradigmatic like Duchamp's *Fountain*. These processes of self-critique and analysis are essential for the field and its relation to the world, as notably proposed by institutional critique. The intermingling of Art and AI described in this article continues this self-critical artistic practice in a time of 'ways of machine seeing.' Likewise, this research also aims at incorporating the creative practices from the artistic world into the AI field, provoking it to be more (self-)critical and experimental.

We present our research in this article, alongside different artistic outputs, as a way of experimenting with research methodologies and interdisciplinary positions. This research was also presented at the museum whose archive we analyzed, creating a crack for reflexivity within an elitist, codified space. This can be seen as a contribution to art education in museums: we propose that people, as exhibition visitors, be invited to experiment with the distanced look of AI as a way of critically thinking about the art system. We hope such practices help to create new relationships, openings, and connections for those who are non-specialists to explore art critically. Instead of museums using commercially available AI from Big Tech, in an uncritical way, why not make more radical and creative uses of technology?

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 continually train their 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

- Amoore L (2019) Doubt and the Algorithm: On the Partial Accounts of Machine Learning. *Theory, Culture & Society* 36(6):147-169. <https://doi.org/10.1177/0263276419851846>.
- Berger J (2008) *Ways of Seeing*. Penguin, London.
- Bowker G C, Star S L (2000) *Sorting things out: Classification and its consequences*. MIT Press, Cambridge.
- Brayne S (2017) Big Data Surveillance: The Case of Policing. *American Sociological Review* 82(5):977-1008. doi:10.1177/0003122417725865
- Buolamwini J, Gebru T (2018) Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification. 1st Conference on Fairness, Accountability and Transparency. PMLR 81:77-91
- Cheney-Lippold J (2011) A New Algorithmic Identity. *Theory, Culture & Society* 28(6):164-181. doi:10.1177/0263276411424420
- 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. <https://youtu.be/a2IT7gWBfaE>. Accessed 19 Feb 2019.
- Crawford K, Paglen T (2019) Excavating AI: The Politics of Training Sets for Machine Learning. <https://www.excavating.ai/>. Accessed 1 Jun 2020.
- Deng J, Dong W, Socher R., Li L-J, Li K, & Fei-Fei L (2009) Imagenet: A large-scale hierarchical image database. 2009 IEEE Conference on Computer Vision and Pattern Recognition. doi: 10.1109/CVPR.2009.5206848.
- 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.
- D'Ignazio C, Klein L F (2020) *Data Feminism*. MIT Press, Cambridge.
- 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
- Eubanks V (2018) *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor*. St Martin's Press, NY.
- Finn E (2017) *What Algorithms Want*. MIT Press, Cambridge.
- Fraser, A (2005) Was ist institutionskritik? *Texte Zur Kunst* 59.
- Freire C (2006) *Arte conceitual*. Jorge Zahar Editora, Rio de Janeiro.

- 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*. MIT Press, Cambridge.
- Gillespie T (2016) Algorithm. In B. Peters (Ed.), *Digital Keywords: A Vocabulary of Information Society and Culture* (pp. 18-30). Princeton University Press, NJ.
- Google. AI & Machine Learning Products: Cloud Vision. <https://cloud.google.com/vision/>
Accessed 18 Feb 2019.
- 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. <http://imagenet.stanford.edu/synset?wnid=n04225987>. Accessed 7 Mar, 2019.
- 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.
- 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*. MIT Press, Cambridge.
- Mintz A, Silva T, Gobbo B, Pilipets E, Azhar H, Takamitsu H, Omena J, Oliveira T (2019) Interrogating Vision APIs. SMART Data Sprint: Beyond Visible Engagement. <https://smart.inovamedialab.org/smart-2019/project-reports/interrogating-vision-apis/>.
Accessed 1 Jun 2020.
- Noble S U (2018) *Algorithms of Oppression: How Search Engines Reinforce Racism*. NYU Press, NY.
- Olah C, Satyanarayan A, Johnson I, Carter S, Schubert L, Ye K, & Mordvintsev A (2018) The Building Blocks of Interpretability. *Distill*. doi:10.23915/distill.00010
- O'Neil C (2016) *Weapons of Math Destruction*. Crown Books, Largo.
- Paglen T (2016) *Invisible Images (Your Pictures Are Looking at You)*. The New Inquiry. <https://thenewinquiry.com/invisible-images-your-pictures-are-looking-at-you/>. Accessed 1 Jun 2020.
- Parker R, Pollock G (2013) *Old Mistresses: Women, Art, and Ideology*. I.B. Tauris, London.
- Powles J, Nissenbaum H (2018) The Seductive Diversion of 'Solving' Bias in Artificial Intelligence. <https://medium.com/s/story/the-seductive-diversion-of-solving-bias-in-artificial-intelligence-890df5e5ef53>. Accessed 1 Jun 2020.
- Simanowski R (2016) *Data Love: The Seduction and Betrayal of Digital Technologies*. Columbia University Press, NY.

Tomkins C (1998) *Duchamp: A Biography*. Holt Paperbacks, NY.
Warburg A (2010) *Atlas Mnemosyne*. Akal, Madrid.