Transformations issue 36 (2022)
www.transformationsjournal.org

ISSN 1444-3775

Beyond Machine Vision: How to Build a Non-Trivial Perception Machine

Joanna Zylinska

AUTHOR BIO

Joanna Zylinska is Professor of Media Philosophy + Critical Digital Practice at King's College London. The author of a number of books – including *AI Art: Machine Visions and Warped Dreams* (Open Humanities Press, 2020), *The End of Man: A Feminist Counterapocalypse* (University of Minnesota Press, 2018) and *Nonhuman Photography* (MIT Press, 2017) – she is also involved in experimental and collaborative publishing projects, such as *Photomediations* (Open Humanities Press, 2016). Her own art practice involves working with different kinds of image-based media.

ABSTRACT

Approaching the problem of artificial creativity through the lens of machine vision, this article examines the impact of computer science's model of vision on our socio-political values and institutions. It also proposes a creative experiment in "conceptual engineering," which entails an attempt to build a non-trivial perception machine. This idea references two science papers on the relationship between humans and machines: Heinz von Foerster's "Perception of the Future and the Future of Perception," in which the concept of a "non-trivial machine" was first introduced, and Gerald M. Edelman and George N. Reeke Jr.'s "Is It Possible to Construct a Perception Machine?". Critically engaging with those papers, the author ends by constructing a conceptual scaffolding for the theory and praxis of machine perception, while addressing the wider problem of epistemic and racial (in)justice in the industry focused on getting machines to "see."

KEYWORDS

machine vision, computer vision, perception, imperialism, bias

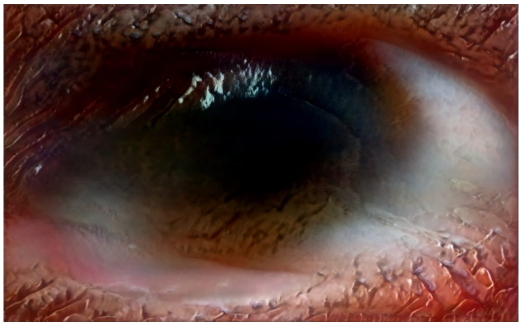


Figs. 1 and 2. Stills from Joanna Zylinska, *Neuromatic*, 2020.

How to see better than humans

This article approaches the problem of artificial creativity through the lens of machine vision (figs. 1 and 2). Foregrounding the artifice involved in the production of the concept and practice of vision in machines *and* humans, [1] I will interrogate the impact of computer science's model of vision on our socio-political values and institutions. As well as offering critical analysis, as part of my argument I will embark on a creative experiment in "conceptual engineering," an approach that links the pragmatism of the machine-building discipline with its Latin etymology in *ingeniare*, meaning "to devise," "create" or "contrive" – and thus also "to play." The title of my article draws on two science papers, separated by a couple of decades, by authors who made a significant contribution to the debate on the relationship between humans and machines: Heinz von Foerster's 1971 paper "Perception of the Future and the Future of Perception," in which the concept of a "non-trivial machine" was introduced, and Gerald M. Edelman and George N. Reeke Jr.'s 1990 article, "Is It Possible to Construct a Perception Machine?". I will engage with those papers in an attempt to construct a conceptual scaffolding for the theory and praxis of machine perception, while addressing the wider problem of epistemic and racial (in)justice in the industry efforts focused on getting machines to "see."

[1] I have previously discussed this idea of human vision being constitutively artificial in *Nonhuman Photography* (2017). Chapter 1 of the book analyses the human '*as part of a complex assemblage of perception in which various organic and machinic agents come together – and apart – for functional, political, or aesthetic reasons*' (14, emphasis in the original).

First, some clarification regarding the key terms and concepts under discussion is in order. The term "machine vision" refers to the systems engineering discipline which works on the automatic extraction of information from digital images to enable machines to perform tasks requiring human sight. Such tasks may include quality control, identification, positioning and measurement, and can be found in applications such as detection systems in self-driving cars, security or space exploration. Machine vision systems rely on cameras with sensors, processing hardware and software algorithms. The software side of those systems is developed by a cognate discipline called "computer vision," a subfield of the broad area of artificial intelligence (AI), which deals with theorising how the extraction of information from digital images actually occurs – although the two terms, "machine vision" and "computer vision," are sometimes used interchangeably. The goal of machine vision is to imitate the way humans see the world, but this imitation attempt occurs only *after human visual processes have been redefined in computational terms* (see Impett). [2] It is also to learn how to see *better*, that is faster and more efficiently, *than* humans.

[2] There are some affinities between my approach in this article and Impett's statement that the algorithms of computer vision have an ideology, and, more importantly, a philosophy – even if both tend to remain latent in the dominant scholarship of computer vision.

The methodological premises for teaching computers how to see were laid in the 1959 paper by two neurophysiologists, David Hubel and Torsten Wiesel. Titled "Receptive Fields of Single Neurons in the Cat's Striate Cortex," it is referenced as foundational to the development of computer vision. The paper was based on a series of experiments conducted by the two researchers on "lightly anaesthetized" (Hubel and Wiesel 589) felines whose eyes had been "immobilized by continuous intravenous injection of succinylcholine" (575) to eliminate unpredictable eye movement. The cats were being shown pictures of dots and different light shapes with a view to assessing their brain activity and thus identifying cortical correlates of vision. While failing to detect any significant changes to the cats' neuronal activity as a result of being exposed to light projections of various shapes, the experimenters had a breakthrough when some of the cats' cortical neurons started firing furiously in reaction to the felines being exposed to slits of light. They eventually realised that what the cats' retinas were reacting to was not any specific shape, such as a dot or a line, but rather a change in light intensity at the edge of a slide frame. Hubel and Wiesel traced this activity to what they ended up calling "simple cells," arranged into columns, in (what is now known) as the primary visual cortex. This experiment, conducted as it was on immobile nonhuman subjects, under laboratory conditions, led to the inauguration of one of the foundational assumptions (or, indeed, myths) of computer vision: the belief that the process of vision is multi-layered and hierarchical, that it is possible to extricate the essence of vision in various animals (including those of a human variety), that the mechanism of edge perception is what lies at the core of vision, that the process is physiological and content-independent, and that machines can be taught to see "like humans" by mimicking this process of pattern perception at the level of pixels.

These formalist parameters were consolidated by the 1982 textbook of computer vision, simply titled *Vision*, by MIT scientist David Marr. For Marr, vision was primarily a phenomenon of information processing (see Marr 3-5). The framework assumes that the mechanism of primary visual processes such as edge detection or binocular vision is computational and works its way from those primary processes upwards in its examination of the neural circuitry supposedly enabling vision, all the way to the brain. It is this understanding of "computer vision," and of its engineering counterpart, "machine vision," that I critically interrogate here. With this, I aim to probe how, or indeed whether, machines can actually see at all. I am thus also interested in what it means for humans to classify *as seeing* machines' ability to differentiate between objects in the world on the basis of light reflected off them and transmitted as data to those machines' processors. This critical investigation will allow me to explore the possibility of building a perception machine, on terms that engage with, but also go beyond, those delineated by biology and computer science.

From machine vision to machine perception

The shift from vision to perception enacted as part of this experiment is not entirely mine: in recent years many AI researchers have gone beyond the explicit ocularcentrism of machine vision to extend the study of machines' data capture operations to other senses, such as hearing, touch and olfaction.

Expanding their data sources from still images to sounds, music and video, Google is now using the term "machine perception" in lieu of "machine vision." Yet its perceptive operations are still premised on object recognition, which involves algorithms trained on processing large, partially-labelled datasets using parallel computing clusters. Google's notion of "machine perception" encompasses a wider set of sensory data, yet it retains the sense of predefined objects to be sensed, with their corresponding categories. [3] The notion of "machine perception" I am proposing here, adopted in a wider sense than the one offered by Google, aims to address one of the key blind spots of computer vision today: its inability to fully account for how our brains actually work and how the translation process from retinal stimulation through to the neural circuits of the brain occurs, while producing a sensation and an awareness of this sensation (recognised in the form of an image that we "see") in humans. The concept can also raise questions for the postulation of a discrete physiological unity called "the brain" as the core organ of perception – and thus for modelling machine vision on human cortical processes. Cognitive scientist J.Y. Lettvin et al.'s 1959 paper, "What the Frog's Eye Tells the Frog's Brain," demonstrated that perceptive activity that was assumed to take place in the brain as a consequence of the retina being stimulated by light in fact had already begun in the eye. The exact location of perceptive processes and the exact working of their operative mechanisms remain difficult to pin down not just in frogs but also in humans. The concept of machine perception thus complicates the model of vision as simply representational – and offers a different way of understanding what it might mean for machines to see. It also postulates that perception occurs in the world as much as it does in the eye or the brain. This repositioning calls for a more embedded, embodied and dynamic understanding of how computers (and, indeed, humans) see the world – and of how they act in it.

[3] See https://research.google/pubs/?area=machine-perception

Perception of the future

The "conceptual engineering" approach of my article departs somewhat from the recent uses of this term in analytical philosophy, where it serves as a more programmatic and less playful thought device. David Chalmers, for example, defines conceptual engineering as "the process of designing, implementing, and evaluating concepts" (2). Fixing and hence stabilising concepts is at least as important to him as the act of inventing them. For me, in turn, conceptual engineering has something of cyborg bricolage about it: it is less permanent, less intent on fortifying edifices – and much more mischievous. Bringing together the two aspects of Chalmers' proposition, "de novo engineering" and "re-engineering," it retraces the patterns and scaffoldings of the established ideas and texts to arrange some new, albeit temporary, constructions – while looking for safety exits.

As indicated earlier, there are good reasons for turning to the two papers by von Foerster, and Edelman and Reeke Jr., as these authors played an important role in redefining physiological phenomena such as perception and vision as programmable. As part of their research, they set out to redraw traditional disciplinary boundaries in an attempt to find a new way of understanding humans and machines, or even humans *as* machines. Von Foerster was an

Austrian physicist and engineer with an active interest in biology. Having spent most of his working life in the US, he was one of the key participants in the Macy conferences, a series of meetings that took place between 1940-1961 which led to the emergence of cybernetics as a transdisciplinary field studying biological and mechanical systems. Von Foerster was a prolific writer, often veering beyond the confines of his discipline or even science as such to offer a wider commentary on society and the world. "Perception of the Future and the Future of Perception," first presented at the opening of the Annual Conference on World Affairs at the University of Colorado, Boulder, in 1971, is an example of this approach.

The paper opens with a turn to systems theory as a way of thinking about socio-cultural change, while also indicating the limitations of the static model of the system, whereby change is seen as an aberration to be corrected rather than an incentive to perform alteration at the systemic level. Taking account of systemic operations, von Foerster goes on to seek an opening within the system towards an execution of human freedom and agency – and it is perception that he turns to for this purpose. Specifically, he advises us that "if we wish to be subjects, rather than objects, what we see now, that is, our perception, must be foresight rather than hindsight" (von Foerster 31). In other words, he implores us not to fall back on the established patterns of seeing things but rather to follow unexpected routes and pathways, taking a lesson from children. So far, so pastoral: he is of course not the first thinker to evoke the idea of original untarnished vision as a supposedly better entry point into the truth of the world: Enlightenment philosophers advocating a return to the "state of nature," Romantic poets and surrealist artists were there before him. Yet von Foerster has some interesting things to say about the current atrophy of perception, for which he blames the commodification of information and "an educational system that confuses the process of creating new processes with the dispensing of goods called 'knowledge'" (33). To offer a remedy for this state of events, von Foerster introduces the concept of a "non-trivial machine," where the "machine" "refers to well-defined functional properties of an abstract entity rather than to an assembly of cogwheels, buttons and levers, although such assemblies may represent embodiments of these abstract functional entities" (40).

While a "trivial machine" is premised on a one-to-one relationship between its "input" and its "output," thus delivering predictable and consistently identical results, at least in theory, in a non-trivial one "its previous steps determine its present reactions" (40). While both systems are deterministic, the second one is unpredictable because its output changes with what it has picked up in the previous cycles of its operation. Von Foerster does appreciate the reliable predictability of various trivial systems, such as toasters or cars, but raises concerns about the application of trivialisation to other systems – and this is where his concept of the machine and his argument become particularly interesting. He takes the example of a student, seen as a potentially non-trivial machine, becoming completely "trivialized" when he [*sic*] enters a higher-level machine, i.e. a university, with its predictable teaching and uncreative testing. "A perfect score in a test is indicative of perfect trivialization: the student is completely predictable and thus can be admitted into society. He will cause neither any surprises nor any trouble" (41). While the human (i.e. the student) is understood here as a machine, still in the process of being constructed in their relatively young age, the higher-level operations of the recursive loops enable the possibility of creativity within the system, made up of other machines (educators, forms of knowledge and ways of their transmission, pedagogies, buildings and infrastructures) – but only if the human agents execute their potential: in other terms, when they embrace the incomputable as part of systemic rationality.

It may seem obvious that educators, or, indeed, most humans would want to do this, and that they would reach for the systemically determined yet ultimately undefined degree of freedom available to them. Yet, as confirmed both by von Foerster (who sees such a reluctance as a form of illness) and philosopher Vilém Flusser (for whom we always operate under the condition of systemic predictability, be it on the level of the machine, society or the inevitably entropic universe) (see *Towards a Philosophy of Photography* 26-30, 82-92; *Into the Universe of Technical Images* 19), such acts of reaching out towards freedom, i.e. attempting to change the system's course of action, are quite rare. Von Foerster's paper ultimately ends with a rather humanist call to recognise and cherish the "troublemakers," whom he says we will recognise "by an act of creation: 'Let there be vision: and there was light'" (43). Interestingly, it is through perception that we (as educators and assessors of students) are said to be able to introduce novelty into the input-output process and thus create a systemic opening. This could be a first step in attempting to build a non-trivial perception machine.

But how *do* we make students, or indeed anyone, see better, without knowing in advance what we want them to see, what kind of frameworks we are looking through and aiming for? Isn't there a danger that this call on the part of von Foerster will ultimately sound banal, or in fact trivial: at best vaguely humanist, at worst a recipe for all sorts of educational libertarians to break with the expertise and care developed within the educational system? Frequently described as a "polymath," von Foerster is one of a long line of (predominantly male) scientists who opine on wider societal issues, making quick excursions to scientific models and protocols in order to develop explanations of the world while demonstrating strange blind spots when it comes to the cultural aspects of the systems they write about, including the terminological (and ideological) determination of their concepts. (Of course, all conceptual frameworks are ideologically determined, but it is the unawareness or obfuscation of this phenomenon, coupled with the ascription of "ideology" to one's intellectual opponents, that I am highlighting as a problem here.)

This model of what we might term "truncated rationality" is very much at work in the very hotbed of cybernetic and post-cybernetic thinking, Silicon Valley, but it also thrives in scholarly disciplines such as analytical philosophy, linguistics, biology, cognitive psychology and neuroscience – a cluster of disciplines which provide intellectual foundations for the field of AI, and which frequently rely on copious amounts of technological and military funding. By way of responding to the omnipresent working of truncated rationality in our increasingly automated and datafied society, Luciana Parisi has come up with what she has termed the "Alien Hypothesis," a proposition which embraces "an abductive, constructionist, experimental envisioning of the working of logic … as part of a speculative image of the alien subject of AI" (32). What is interesting about her idea is that it comes from within the logical parameters of cybernetic thought. Her hypothesis is designed as an alternative to both the Cybernetic Hypothesis, positing that the only possible escape from the constraints of the computational system is a flight into zones of opacity and invisibility involving the total negation of the logic of the system, and the Accelerationist Hypothesis, which advocates dismantling the system's power from within by exhausting its operations. Instead, Parisi proposes to engage with the systemic operations on their own logical terms with a view to discovering "an alien space of reasoning" (30) *within* the system. Drawing on the consequences of mathematician Gregory Chaitin's identification of incomputables within computational systems, she argues for the possibility of finding this space *beneath* the scripted and looped servo-mechanics of the system that make it seem like there is no other way. Parisi's solution lies in reconceptualising the medium of thought beyond its automation and instrumentalisation within the current machinic systems. Mine here, in turn, involves taking a step back alongside our cognitive-sensory spectrum to offer perception rather than thought as a primary mode of engaging with the world through which such an attempt at undoing the system's operations can begin to take place. [4]

[4] In making this strategic distinction, I am mindful of Humberto Maturana and Francisco Varela's intimation developed in *Autopoiesis and Cognition* that perception is *already* a form of cognition – although cognition for them stands for more than pure thought. Indeed, they conceptualise both perception and cognition as distinctly biological phenomena.

Vision beyond the brain

So is it possible to construct a perception machine? This question, as mentioned earlier, was originally posed by biologist and Nobel laureate Gerald M. Edelman and his collaborator George N. Reeke Jr. in their 1990 paper. Edelman and Reeke's work was part of their wider project on the neglect of findings from evolutionary biology in AI research. In an article written two years prior, they chastised AI researchers for remaining too wedded to "epistemological assumptions drawn on the one hand from the arguments of Alan Turing and Alonzo Church about the universal problem-solving capabilities of computers (suggesting that the brain may be understood as a computer) and on the other hand from the reductionism of molecular biology (suggesting that the brain may be understood as a collection of units that exchange chemical signals)" (Reeke and Edelman 143). They also had some rather critical things to say about "neural network computing" [5] as an approach that promised to develop machine vision by modelling it on the operations of the human brain. For them "neural network computing" is a misnomer, because the approach that underpins it is premised on positing a badly conceived analogy between neural networks in the brain and computer networks, and not on the way biology actually understands and works with neural structures. The strictly computational approach, they argue, cannot really tell us much about perception because of its foundational error – namely, the belief that

[5] The concept of "neural networks" was introduced by two researchers at the University of Chicago, Warren McCullough and Walter Pitts, in 1944. Neural networks kept falling in and out of favour in AI research through the last century but came to their own with the recent developments in machine learning.

> objects and events, categories and logic are given and that the nature of the task for the brain is to process information about the world with algorithms to arrive at conclusions leading to behavior…. This "category problem" leads directly to the inability of AI systems to cope with the complexity and unpredictability of the real world. (145)

Recursive neural networks used in machine aka deep learning that consist of layers of nodes can be said to have partly addressed the problem of being unable to deal with the increasing complexity and uncertainty of the world. Indeed, they have had significant successes in identifying patterns and trends in imprecise data, as evident in applications such as face recognition, medical data analysis, weather prediction or natural language translation. Yet neural networks as currently conceived in AI research still do not ultimately challenge the assumption "that information exists in the world," while the organism "is a *receiver* rather than a *creator* of criteria leading to information" (153). I am particularly interested in Edelman and Reeke's critique of the idea of objects and events existing in the world out there, to be seen, grasped and manipulated by us. This critique corresponds to the philosophical position of "conscious realism" espoused by cognitive psychologist Donald D. Hoffman (*The Case Against Reality: How Evolution Hid the Truth from Our Eyes*). This is not to say that there is nothing in the world, only that the pre-conceived and discretised objects that allegedly present themselves for our vision to capture them are in fact outcomes of a creative, dynamic yet ultimately undetermined process. We see what we need to see, argues Hoffman. For him, the nature of this need is biological, or, more precisely, evolutionary, in the sense that the purpose of what we see is our survival – and this is where his intellectual trajectory coincides with Edelman and Reeke's. Indeed, as a riposte to physicists' computational schematism, Edelman and Reeke suggest that any viable theory of categorisation and intelligence to be used in AI research needs to embrace the Darwinian model of selection but adjusted for the working of the neurons of a single organism operating during its lifetime. This conclusion serves as the grounding for their 1990 paper about the possibility of constructing a perception machine.

And, indeed, in the paper they propose such a construct. Named, without any irony, Darwin III, this machine's architecture is premised on their proposition that categorisation (rather than more straightforward object recognition as applied in computer vision systems) is a critical component of perceptual systems. Perception does involve "the adaptive discrimination of objects or events through one or more sensory modalities, separating them from the background and from other objects or events" (Edelman and Reeke 36), but the difference here is that these objects are not predetermined. Instead, for Edelman and Reeke perception is not about grasping representations, as the categories used by the organism or machine do not exist in the world for this organism or machine only to come and recognise them: rather, those categories are actively constructed (i.e. they are meaningful to the perceiver "in a given situation" 37). Perception thus becomes redefined as "an active sensorimotor process, requiring exploration and depending on past experience" (37). In this model categorisation needs to be seen as biological, not mathematical, because it cannot be solely expressed in symbolic form. The biological aspect of these organisms implies that movement through the world is needed for recognition, and hence perception and intelligence, to be enacted. This model is premised on the interesting idea, formulated in the two authors' earlier paper, of the brain as "a selective system operating in somatic time" (Reeke and Edelman 157), one which resonates with the phenomenological approach to AI design, which is increasingly growing in popularity.

Yet Edelman and Reeke's concept of the perception machine itself remains hamstrung in this acultural model, with machinic operations positioned as primarily driven by natural selection. Cultural transmission of information is dismissed by them as not being relevant to the *evolutionary* development of working perceptual systems. Even though Darwin III is supposedly more firmly placed in the world, its orientation, vision and goals are crippled by the temporality of its movement, with "somatic time" requiring an erasure of the said soma's singularity at the expense of the (supposedly) timeless operations of natural selection. There is therefore a danger of Edelman and Reeke's perception machine ending up being rather trivial, because the information it will pick up to develop its categories, even if collected in a more dynamic and embedded way, will by design have been stripped off any cultural specificities – which here only amount to noise. However, designating which aspects of the surroundings can be classified as nature and which belong to the realm of culture is not straightforward.

In his poignantly titled article, "Beyond Biology and Culture: The Meaning of Evolution in a Relational World," anthropologist Tim Ingold argues that the difference between nature and culture is just a matter of temporality: history (which manifests as "culture" in different epochs) simply operates on shorter time scales than evolution (which proceeds according to laws of "nature"). Ingold goes on to show how Darwin's *The Origin of Species* challenged the eighteenth-century view of "man" as a being equipped with special characteristics such as reason and morality, characteristics which were meant to separate him from the other species. Darwin postulated that the difference between the human and other animals was in fact one of degree, not of kind, and that evolution, which he had originally termed "descent with modification," was the manner through which this change towards rational man, with his intellectual and moral faculties, happened. Yet the problem with that view was that it inaugurated another differentiation which today we would (rightly) describe as racist: one between the savage and the civilised man. The only way out of this dilemma was to attribute "the movement of history to a process of culture that differs in kind, not degree, from the process of biological evolution," explains Ingold (213). It was also to put forward the idea of two parallel kinds of inheritance in human populations: one biological (involving the transmission of genetic information encoded in the DNA and referring to the core identity of the human – i.e. walking), the other – cultural (taking place through social learning – e.g. being able to play an instrument). Yet there is nothing "purely" natural about walking – and neither is learning to play an instrument a fully cultural experience, separated from the transmission of embodiment. Instead, "those specific ways of acting, perceiving and knowing that we have been accustomed to call cultural are enfolded, in the course of ontogenetic development, into the constitution of the human organism," which makes them equally "facts of biology" (Ingold 216). At the heart of the problem lies not so much the conflation of the biological with the genetic, Ingold points out. The decoupling of the processes of historical and biological change should therefore rather be treated as a de-temporalisation, a process premised on a fictitious instauration of separate temporalities for certain kinds of changes over others. This fallacious model is still very much with us. It subtends the present consensus with regard to science as a method – and the disciplinary division between the sciences and the humanities.

Machine vision and epistemic (in)justice

Ingold's argument poses a challenge to any attempt to construct perceptual systems which will be capable of passing on their evolutionary traits to their offspring but which will not be able to pass on any cultural influences (or, indeed, which will need to be free from such influences). Any traits that *will* get passed on will always carry both "natural" and "cultural" inscriptions, in a manner that will not allow for their easy decoupling, but the construction of culture as a separate domain of uninheritable features will perpetuate the distinction, while allowing computer scientists to forgo embodied and embedded modes of perception and cognition. This disembodied model of computer vision results in the preservation of one of the biggest science (and computer science) myths: the belief that data bias understood as cultural bias, once eliminated, will result in the data that is both pure and fair. [6] We are regularly presented with consequences of such essentialisation of biology and "the brain," at the expense of "cultural traits," in cognitive and computer science. Two examples from late 2020 include the video conferencing platform Zoom's background algorithm, which removed the head of a Black academic any time he tried to use a virtual background, and the Twitter cropping algorithm, which privileged the showing of white faces in cropped images in its timeline.

[6] As Claudio Celis Bueno and María Jesús Schultz Abarca point put, drawing on Bernard Stiegler's philosophy: "The 'bracketing off' of inherited prejudice to perceive reality 'in itself' is an illusion that conceals the fact that technology permanently modifies our internal senses of perception and memory. Naked human vision too is always already machine vision. Human vision, like machinic vision, depends on the surfaces of inscription that function as an external faculty of imagination" (Celis Bueno and Abarca 1185).

While the computer vision machine reveals itself not to be particularly perceptive, the consequences of its racialised blind spots are anything but trivial. Indeed, the algorithms that run within it are the same ones that make decisions about people's social, financial or legal status, including punitive action at border control, denial of credit, prediction of educational failure or assignation of criminality. While early Google image-recognition algorithm was highlighted in 2015 to auto-tag pictures of Black people as "gorillas," there is an ongoing problem with face recognition of Black females, with the high false match rate explained by industry experts as a combination of the difficulty of lighting a Black face and the makeup worn. Denying explicit bias, Thorsten Thies, director of algorithm development of German company Cognitec which supplies facial recognition systems to governments, explained in a troublingly disarming manner that it is "harder to take a good picture of a person with dark skin than it is for a white person" (Simonite n. pag.). One factor is that the image databases that serve as training sets for the algorithms are not properly representative, being skewed, in terms of volume and quality, towards photographs of white males. But there is a deeper logic at work here, with the whole systemic infrastructure involved in the production of cameras, lighting systems, image-processing software, and the visual and cultural training of photographers and image technicians that produces a particular set of internalised norms that can then be presented as posing an "objective" difficulty in taking a photo of a person with a dark skin. This mode of thinking, embedded in all sorts of technologies that precede the digital, is what Safia Noble has critiqued in her book *Algorithms of Oppression*. Suggesting that "artificial intelligence will become a major human rights issue in the twentyfirst century" (Noble n. pag.), Noble issues a clarion call for us all to understand the architecture and logic of algorithmic decision-making tools in masking and deepening social inequality. As part of her analysis, she puts to rest the belief that unfair systemic decisions are just occasional aberrations which can be easily eliminated for the functionality and efficiency of the supposedly neutral system to be restored.

CDO at Twitter Dantley Davis responded to the 2020 cropping algorithm's debacle discussed earlier with the chest thumping yet predictable: "It's 100% our fault. … Now the next step is fixing it" (@dantley on Twitter, September 19, 2020). Yet Noble leaves us no illusion that "algorithmic oppression is not just a glitch in the system but, rather, is fundamental to the operating system of the web" (n. pag). Mitra Azar, Geoff Cox and Leonardo Impett similarly suggest that

> In a structurally unequal society, it is exceedingly difficult to make a "fair" algorithm; and it is effectively impossible to make an algorithm which is both fair and effective.… In a society which is unfair, a classification-machine will always be unfair (in at least one sense). ("Introduction: Ways of Machine Seeing") [7]

It is therefore not enough to de-bias the data. Instead, we need to ask bigger questions about the forms of injustice embedded in the systems that host it. We also have to ask what it means when the elimination of the glitch, while desirable from a technical point of view, ends up making the punitive surveillance running on this data even more efficient. The correction of the data bias does not correct the violently penetrative and extractivist logic of the computer vision system: it actually strengthens it. We could thus say that the need to identify the unjust operations of the non-trivial *vision* machine to start thinking of building a non-trivial *perception* machine.

With this, as part of the "conceptual engineering" project outlined at the outset, we are now shifting towards a conceptual expansion of the notion of the machine. In line with its cybernetic legacy from von Foerster and colleagues, the term "machine" departs from its strict engineering connotations to embrace any kind of system, be it mechanical or biological, of a varying degree of complexity. In this framework, systems are arranged in a nested manner, from the microscopic to the cosmic, while undergoing internal transformations or even cross-systemic mutations. In the 1960s and 1970s the cybernetically-inflected notion of the machine became a potent concept for philosophers and cultural theorists attempting to articulate different levels of socio-political complexity while taking into account society's and individuals' biological and technical constitution, from Gilles Deleuze and Félix Guattari's war machines and desiring machines (*Anti-Oedipus: Capitalism and Schizophrenia*) through to Michel Foucault's *dispositif* ("The Confession of the Flesh") and Vilém Flusser's apparatus (*Towards a Philosophy of Photography*; *Into the Universe of Technical Images*). Building on this legacy we should recognise that "the perception machine" we are attempting to build will espouse multiple, albeit interwoven, levels of meaning. The concept highlights the interlocking of scientific and cultural discourses in the production of images – and the production of subjectivity and objectivity as functions of all kinds of images.

[7] The special issue of the journal *AI and Society* titled *Ways of Machine Seeing*, edited by Azar, Cox and Impett (2021), explores the problems of (and with) machine vision through a multidisciplinary framework. Departing from John Berger's proposition, made in his canonical *Ways of Seeing*, that "Every image embodies a way of seeing," (Berger 10), it explores how "machines, and, in particular, computational technologies, change the way we see the world."

In my attempt to build a non-trivial perception machine I am therefore not really constructing a device, at least not in any straightforward way, but rather building a conceptual framework for both framing what we term "the world" and seeing this world better. Such a perception machine will have to do more than just identify symbols, avoid bias or even counter with the environment. In its architecture the idea of the brain as a discrete perceptive organ will need to give way to a dynamic interaction between the organism, with its constantly changing embodiment, and the environment – which is not a constant for it but which "exists only in relation to the organisms that inhabit it, and embodies a history of interactions with them" (Ingold 218). Indeed, there is no perception machine outside of its cultural and historical embeddedness. The very gesture of embarking on the task of trying to build a non-trivial perception machine is also meant to serve as a multi-scalar attempt to both rethink our human perception and vision, and challenge the parameters of the emergent vision machine in our globally networked world. (This approach sums up the rationale behind my "conceptual engineering" project.)

These parameters of the vision machine have been poignantly analysed by Paul Virilio in his eponymous book, published in French in 1988, and in English in 1994. Virilio traces back its structural origins to the transformation of vision in the twentieth century, from the technology of warfare in the Great War which turned vision into visualisation, through to the industrialisation of vision which shifted the majority of activities unfolding within the field of visibility from humans to machines. This transformation of vision as primarily machine-based was accompanied by the "automation of perception" as a result of developments in AI, he claims. However, Virilio's argument postulates a clear shift from more organic (and seemingly pre-technological) vision through to its machinic counterpart: "the relative fusion/confusion of the factual (or operational, if you prefer) and the virtual; the ascendancy of the 'reality effect' over a reality principle" (Virilio 36). This shift, initiated with the development of the prosthesis of sight such as the telescope in 1608 but accelerated three centuries later, with the proliferation of "seeing machines," is interpreted by him primarily in terms of confusion and loss. While a political critique of the automation of perception is very much needed – and Virilio does in fact offer it by looking at visual technologies of war and the automation of vision in propaganda and marketing, and the mechanisation of justice in video-enabled courtrooms – his ontological critique of the change to vision ends up conserving the view of vision as both human and humanist, while situating human subjectivity on a metaphysical plane. His vision machine is thus "an autonomous technological system" (Armitage 203), which disturbs the original organic unity and purity of human sight. Borrowing some of his ideas, in this article I embrace an understanding of machines as encapsulating both organic and non-organic components, changing over time while constitutively shaping the human sensorium in different ways. My attempt to build a non-trivial perception machine presented here is aligned with Bernard Stiegler's analysis of automation in pharmacological terms, [8] where the process, although harmful to our individual and social life – as evidenced in overall "ill-being," "nihilism" and "[h]umanity's doubt about its future" (*Automatic Society* 9, 7) – can be worked through to release the machine's curative properties.

[8] Pharmacology, a concept derived from Jacques Derrida ("The Pharmakon"), became operative in many of Stiegler's later works, which were devoted to the analysis of the conditions of computational capitalism – and of the socio-economic misery this particular political formation had brought about. Importantly, Stiegler's work was not just diagnostic: he was committed to finding openings within, and passages through, the desperation and nihilism enacted by the automatic decision-making and the disindividuation of the human under those political and economic conditions. As he put it in *Automatic Society, Vol. I*, "From such a perspective the question of innovation must be taken very seriously – and not just treated as an ideological discourse based on the storytelling of marketing. Innovation clearly has a real economic function: it evidently constitutes a production of negentropy. But what has now become obvious is that this negentropy produced in the short term generates far higher entropy in the long term. The whole question of organology and its pharmacology in the neganthropic field resides in the fact that the *pharmakon* can be toxic and curative only to the extent (and in the excess) that it is both entropic and negentropic" (100).

The vision machine we currently live in acts as both a surveillance machine and a data capture device. It is also, as pointed out by Virilio, a war machine, serving as it does as a digitised battleground featuring multiple operations of capture *and* carnage, with real-life consequences for human and nonhuman lives. A non-trivial perception machine I am interested in building would need to be able to scan through its obscure logic. To do this, such a machine would need to do more than be just neutral or un-biased (although it should be that too): it must also be decisively anti-racist. We could go even further and suggest that, to counter both the racist legacy of the war machine and the capitalist extractivism which fuels it, it must also, following Ariella Aïsha Azoulay, be counter-imperialist. In *Potential History: Unlearning Imperialism* Azoulay challenges the regime of rights and privileges that has shaped mechanical image-making since its inception, whereby "the world is made to be exhibited" and "it is only for a select audience" (Azoulay 4). She goes so far as to claim that "Photography developed with imperialism; the camera made visible and acceptable world destruction and legitimated the world's reconstruction on imperial terms" (5).

It may seem at first glance that Azoulay's argument is only aimed at human-centric photography, which is meant to be displayed and seen by (select groups of) humans – and not at large datasets feeding and shaping machine vision today. Yet the same form of rationality arguably underpins the production, storage and categorisation of all mechanically produced images because their constitutive logic, history and modes of framing still hark back to the imperial mindset that legitimated classification as supposedly neutral, while putting it to work with Empire's goals in mind. The location and format of imperial rule have shifted today: in their eponymous book Michael Hardt and Antonio Negri argue that Empire now has no specific boundaries or territorial centre of power. It has become "a *decentered* and *deterritorializing* apparatus of rule" (Hardt and Negri xii), with global capital flows enacting a form of biopolitics by being involved in the production of social life itself as an overlapping nexus of economic, political and cultural forces. The neo-imperial war machine is thus first and foremost a hegemonic surveillance machine: it conquers by implicit consent and by the scale and invisibility of its penetration. All this is not to say that no nodes of power's concentration can be identified within this new imperial apparatus. However, the increasing shift of domination and decision-making from governments to corporations, from Washington to Silicon Valley, from the US to China, and from humans to algorithms creates an uncertain and fuzzy geopolitics in which political and technological black boxes obscure the location of power as well as its actual operations. It also shifts the onus of responsibility onto individual citizens – whose primary yet already dual identity in this system is that of both internet users and data points. The neo-imperial surveillance machine is thus an updated version of Virilio's vision machine. A non-trivial perception machine could be seen as its conceptual and technical counterpart, one that offers an opening into a new vision of both ourselves and the world.

The Recognition Machine

The Recognition Machine by artists Antje Van Wichelen and SICV (Michael Murtaugh and Nicolas Malevé) can be seen as one possible enactment of such a non-trivial perception machine that is also actively counter-imperialist (figs. 3 and 4). I had an opportunity to interact with the version of the work presented at the Photoszene Cologne festival in May 2019, but the project also has an online counterpart. Looking like a photo booth, *The Recognition Machine* invites gallery visitors to enter and take a digital photo of themselves. The act of taking a photo activates an algorithm that attempts "to establish links between the pixels just recorded and those of images from a database of 19th century anthropometric photographs," which have been transformed by analogue techniques. "The resulting print output links contemporary regimes of surveillance to those of a colonial past." ("The Recognition Machine") [9]

[9] See https://recognitionmachine.vandal.ist/

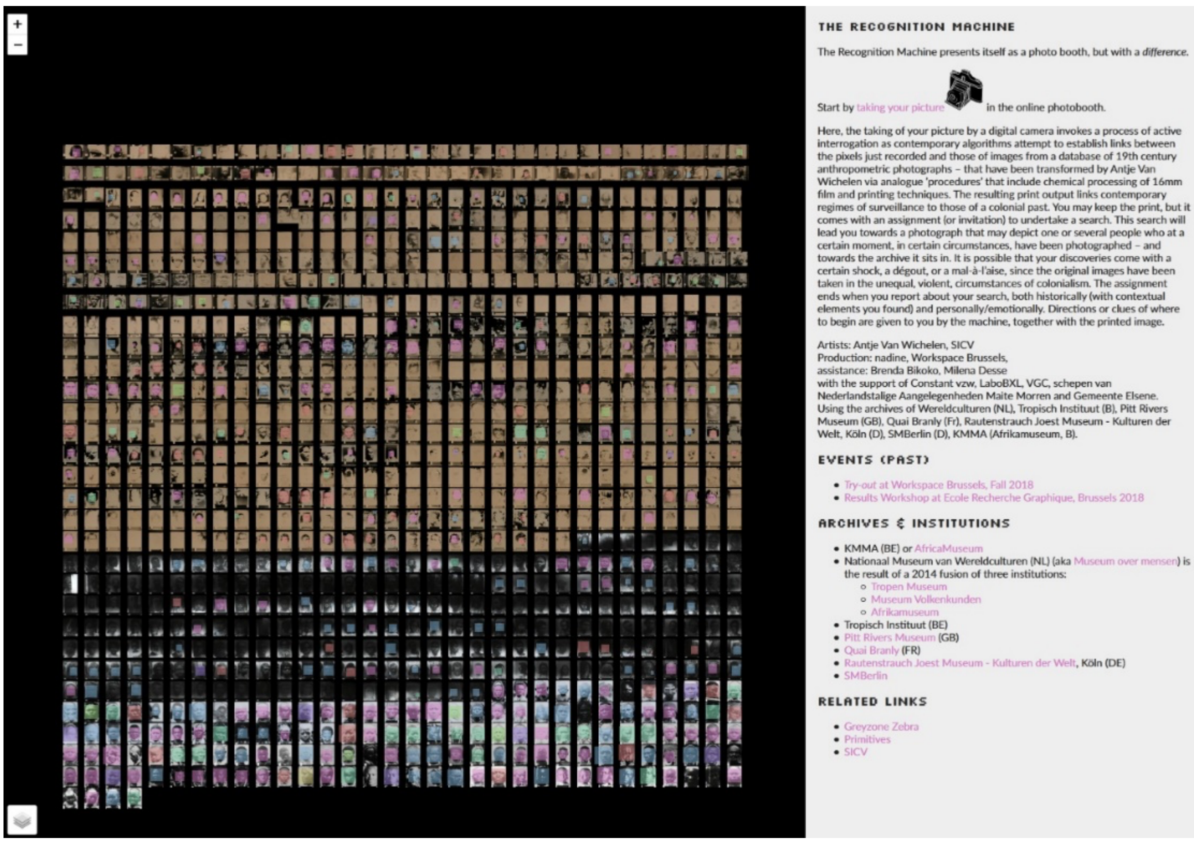

Figs. 3 (this page) and 4 (next page). Antje Van Wichelen and SICV, *The Recognition Machine*, 2018.

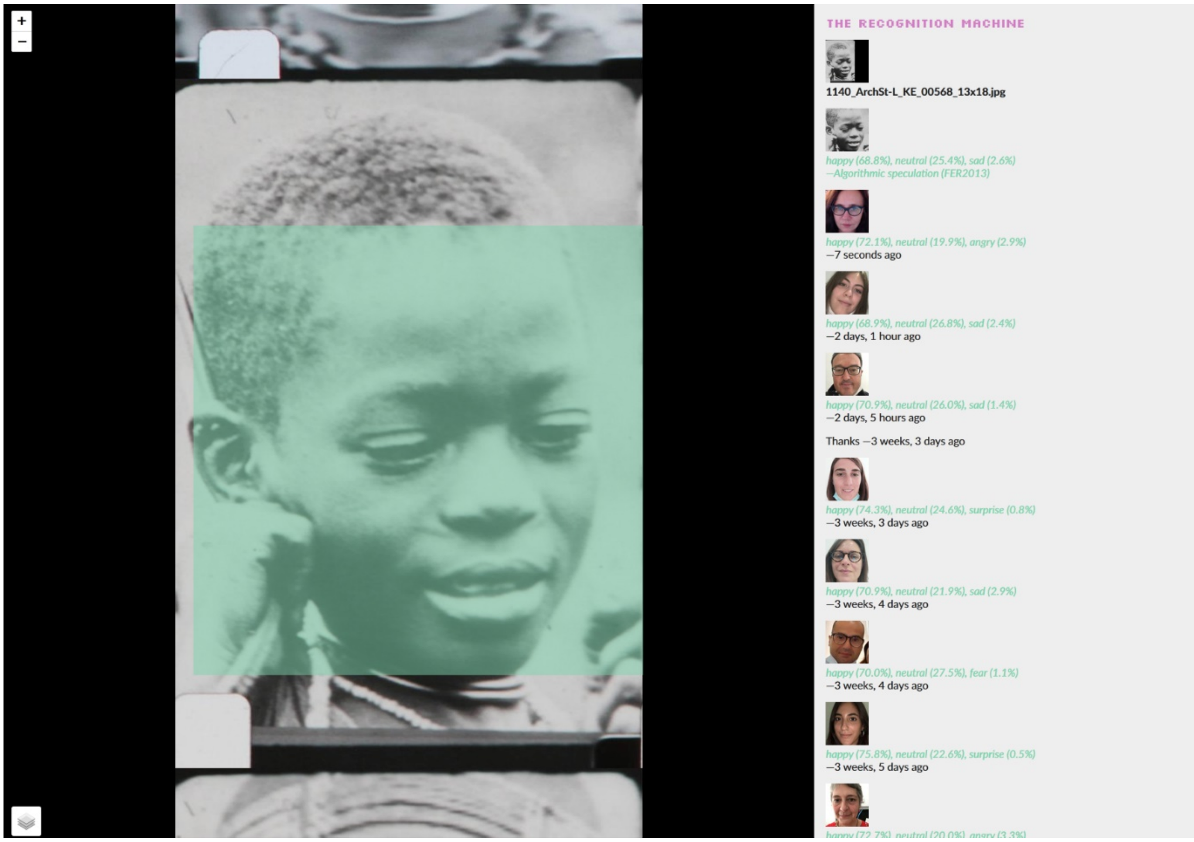

The link between the images pivots around the emotions identified by the algorithm in the viewer's face and linked with the emotions read in the archival photos. The reading was obtained by training the algorithm on the FER-2013 dataset, in which each image had been assigned one of seven emotions: anger, disgust, fear, happiness, sadness, surprise, or neutral. Any possible misrecognition of emotions that occurs as part of the process serves as an alert to the system of consequences that predictive technology is imbricated with: while labelling here is just an innocent game for art audiences, the misrecognition of image links, their wrong categorisation and ascription, has serious consequences for the lived lives of many. The visitor may keep the print obtained, but they are also asked to explore further the posited analogy and thus go deeper both into the archive and the colonial history of portraiture. *The Recognition Machine* also shows us that all images exist as part of the imperial-colonial network of visuality, a network that renders some bodies as visible and proper while deeming others as illegible and/or illegal. What is interesting about this project is that the artists dispense with the idea of a singular image as a stand-alone artefact to be admired, classified and otherwise exploited, showing that *all* images are part of multiple networks of knowledge and data exchange. *The Recognition Machine* thus offers a model of the perception machine as an invitation to study the production of visuality, the image networks and their infrastructures, their underlying data and databases, the algorithms that shape both their production and their networking.

We could therefore conclude that a *non-trivial* perception machine would need to encourage an ethico-political engagement with images, their histories, databases and infrastructures. It should also entail strategies for entering the database on the part of the human, with a view to de-industrialising visuality and vision. It is therefore not just a matter of seeing what is inside the archival machine and how it thinks but also of creating conditions for thinking about human and machine vision otherwise. The human may not be able to see all the available images contained in multiple databases and data clouds, trace all the possible connections between them or take cognisance of all the categories and labels on offer. But what is possible – and indeed imperative – for the human viewer to do is take stock of the logic of opacity and scale that shapes the AI-driven perception machine, to ask questions about its operations and to demand a better – fairer, more historical, and more explicitly anti-racist and counter-imperialist – engagement with the image and data flow. Unlearning imperialism, as pointed out by Azoulay, needs to involve attending to "the conceptual origins of imperial violence, the violence that presumes people and worlds as raw material, as always already imperial resources" (8). The fuzzy borders of today's Empire, coupled with the individual benefits of becoming-data for platform capitalism's servo-technology, make it easier for this form of violence to be seen as consensual – or not as violence at all. Taking first steps towards building a non-trivial perception machine – which is to serve as a blueprint for an alternative version of society and its mode of framing the world – can help us see this form of extractive biopolitical violence for what it is, and then start devising operations for countering it.

Works Cited

Armitage, John. "Accelerated Aesthetics: Paul Virilio's *The Vision Machine*." *Angelaki* 2.3 (1997): 199-209.

Azar, Mitra, Geoff Cox and Leonardo Impett. "Introduction: Ways of Machine Seeing." *AI & Society* (2021) doi:10.1007/s00146-020-01124-6.

Azoulay, Ariella Aïsha. *Potential History: Unlearning Imperialism*. Verso, 2019.

Berger, John. *Ways of Seeing*. British Broadcasting Corporation and Penguin Books, 1972.

Celis Bueno, Claudio and María Jesús Schultz Abarca. "Memo Akten's *Learning to See*: From Machine Vision to the Machinic Unconscious." *AI & Society* (2020) doi:10.1007/s00146-020-01071-2.

Chalmers, David. "What is Conceptual Engineering and What Should it Be?" *Inquiry: An Interdisciplinary Journal of Philosophy* (published online 16 Sep 2020), doi:10.1080/0020174X.2020.1817141, page number refers to the preprint from the author's website.

Deleuze, Gilles and Félix Guattari. *Anti-Oedipus: Capitalism and Schizophrenia*. University of Minnesota Press, 1983.

Derrida, Jacques. "The Pharmakon," *Dissemination*. University of Chicago Press, 1981, pp. 95-119.

Edelman, Gerald M. and George N Reeke Jr.. "Is It Possible to Construct a Perception Machine?" *Proceedings of the American Philosophical Society*, 134.1 (1990): 36-73.

Flusser, Vilém. *Into the Universe of Technical Images.* University of Minnesota Press, 2011.

---. *Towards a Philosophy of Photography*. Reaktion Books, 2000.

Foerster, Heinz von. "Perception of the Future and the Future of Perception." *Instructional Science* 1.1 (1972): 31–43.

Foucault, Michel, "The Confession of the Flesh." Michel Foucault, *Power/Knowledge: Selected Interviews and Other Writings.* Ed. Colin Gordon. Pantheon Books, 1980, pp. 194-228.

Hardt, Michael and Antonio Negri. *Empire*. Harvard University Press, 2000.

Hoffman, Donald D. *The Case Against Reality: How Evolution Hid the Truth from Our Eyes*. MIT Press, 2019.

Hubel, David H. and Torsten N. Wiesel. "Receptive Fields of Single Neurons in the Cat's Striate Cortex." *Journal of Physiology* 148 (1959): 574-591.

Impett, Leonardo. "The Image-Theories behind Computer Vision." Paper presented at *#DHNord2020 – The Measurement of Images*, MESHS, updated on 10 November 2020. https://www.meshs.fr/page/the_image-theory_behind_computer_vision

Ingold, Tim. "Beyond Biology and Culture: The Meaning of Evolution in a Relational World." *Social Anthropology* 12. 2 (2004): 209–221.

Lettvin, Jerome Y. et al. "What the Frog's Eye Tells the Frog's Brain." *Proceedings of the Institute of Electrical and Electronic Engineers* 47 (1959): 1940-1951.

Maturana, Humberto R. and Francisco J. Varela. *Autopoiesis and Cognition: The Realization of the Living.* Reidel, 1972.

Marr, David. *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. W H Freeman, 1982.

Noble, Safiya. *Algorithms of Oppression: How Search Engines Reinforce Racism*. New York University Press, 2018. Kindle edition.

Parisi, Luciana. "The Alien Subject of AI." *Subjectivity* 12 (2019): 27–48.

Reeke Jr., George N. and Gerald M. Edelman. "Real Brains and Artificial Intelligence," *Daedalus* 117.1 (1988): 143-173.

Simonite, Tom. "The Best Algorithms Struggle to Recognize Black Faces Equally." *Wired*. July 22, 2019. https://www.wired.com/story/best-algorithms-struggle-recognize-black-faces-equally/

Stiegler Bernard. *Automatic Society*, *Vol I. The Future of Work*. Polity Press, 2017.

Virilio, Paul. *The Vision Machine*. Indiana University Press, 1994.

Zylinska, Joanna. *Nonhuman Photography*. MIT Press, 2017.