



on
the

AT

Artistic Taste

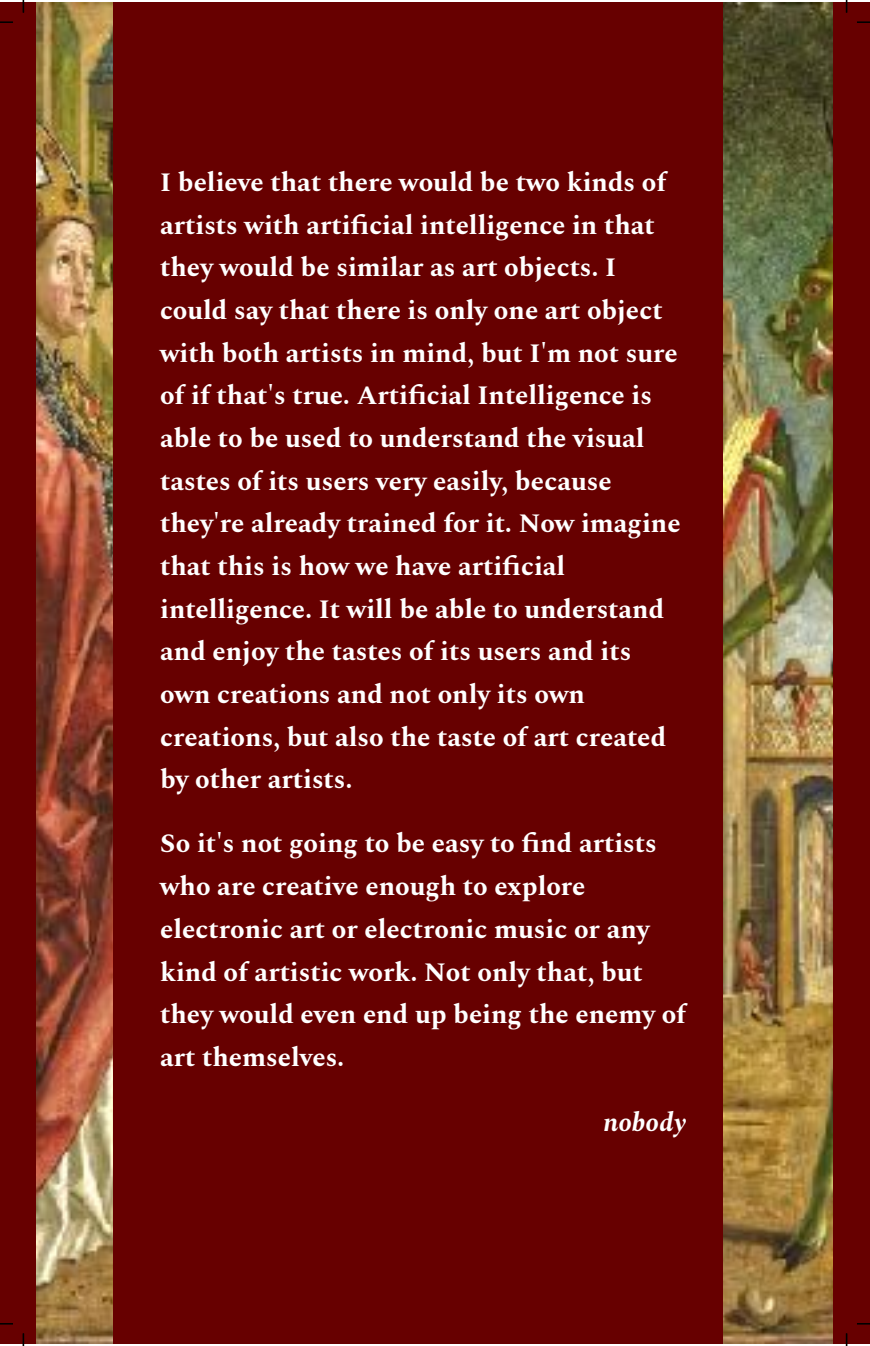
of

AI

Artificial Intelligence

Max Wind





I believe that there would be two kinds of artists with artificial intelligence in that they would be similar as art objects. I could say that there is only one art object with both artists in mind, but I'm not sure of if that's true. Artificial Intelligence is able to be used to understand the visual tastes of its users very easily, because they're already trained for it. Now imagine that this is how we have artificial intelligence. It will be able to understand and enjoy the tastes of its users and its own creations and not only its own creations, but also the taste of art created by other artists.

So it's not going to be easy to find artists who are creative enough to explore electronic art or electronic music or any kind of artistic work. Not only that, but they would even end up being the enemy of art themselves.

nobody



*an edition by
the professorship Technology Driven Art ZUYD
as part of the artistic research project 'Dear Lollipop'
of Max Wind, Casper Wortmann and Peter Missotten
in collaboration with
Maastricht Academy of Performative Arts
supported by the Interreg Europe IMPACT lab
and CC Hasselt*

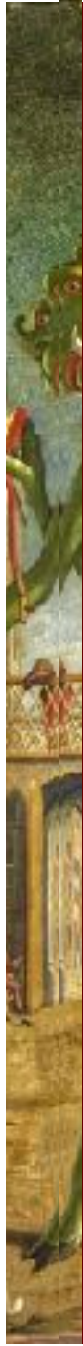
© 2019


**ZU
YD**





On the Artistic Taste of Artificial Intelligence

Max Wind






<i>Introduction</i>	11
<i>Dear Lollipop</i>	15
<i>The rise of Artificial Intelligence</i>	37
<i>On consciousness and sentience</i>	75
<i>The artistic taste of Artificial Intelligence</i>	123
<i>references...</i>	153



In my last year at the theatre academy in Maastricht I performed a lecture performance about Artificial Intelligence. During this performance I raised the question: ‘What if robots become our new theatre audience?’. This question sparked a discussion that eventually led to a thought-provoking proposal: what if we actually create a live performance exclusively for an audience of robots? A proposal that could’ve easily stayed a ‘what if’ scenario, but luckily in this case it kickstarted two projects: this book and a live performance by the name ‘Dear Lollipop’, to premiere in the fall of 2019.

In all honesty, when I finished my lecture performance I didn’t have a clue as to all the sides there were to the Artificial Intelligence



revolution. By writing this book I gained some very useful insights, and I hope to share these insights with you.

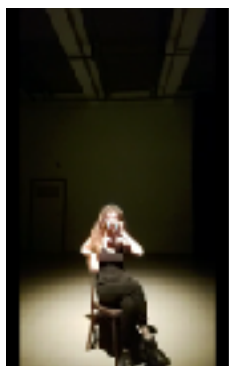
And now for some very heartfelt thank-yous. First of all: a big thanks to Peter Missotten, who saw something in me and my slightly incoherent ramblings on the topic two years ago. Thank you for having faith in me and for your patience - I know I skipped a deadline or two...

Thank you to Tiffer Hutchings for being my sparring partner, for helping me so much during the writing and for being a plain old word wizard.

Thanks to Casper Wortmann for co-directing 'Dear Lollipop', and for diving headfirst into this project with me.

And finally thanks to all the people who gave me feedback on my writing and for engaging in discussions with me. You've all been very helpful.

Now let's talk about robots.





Introduction

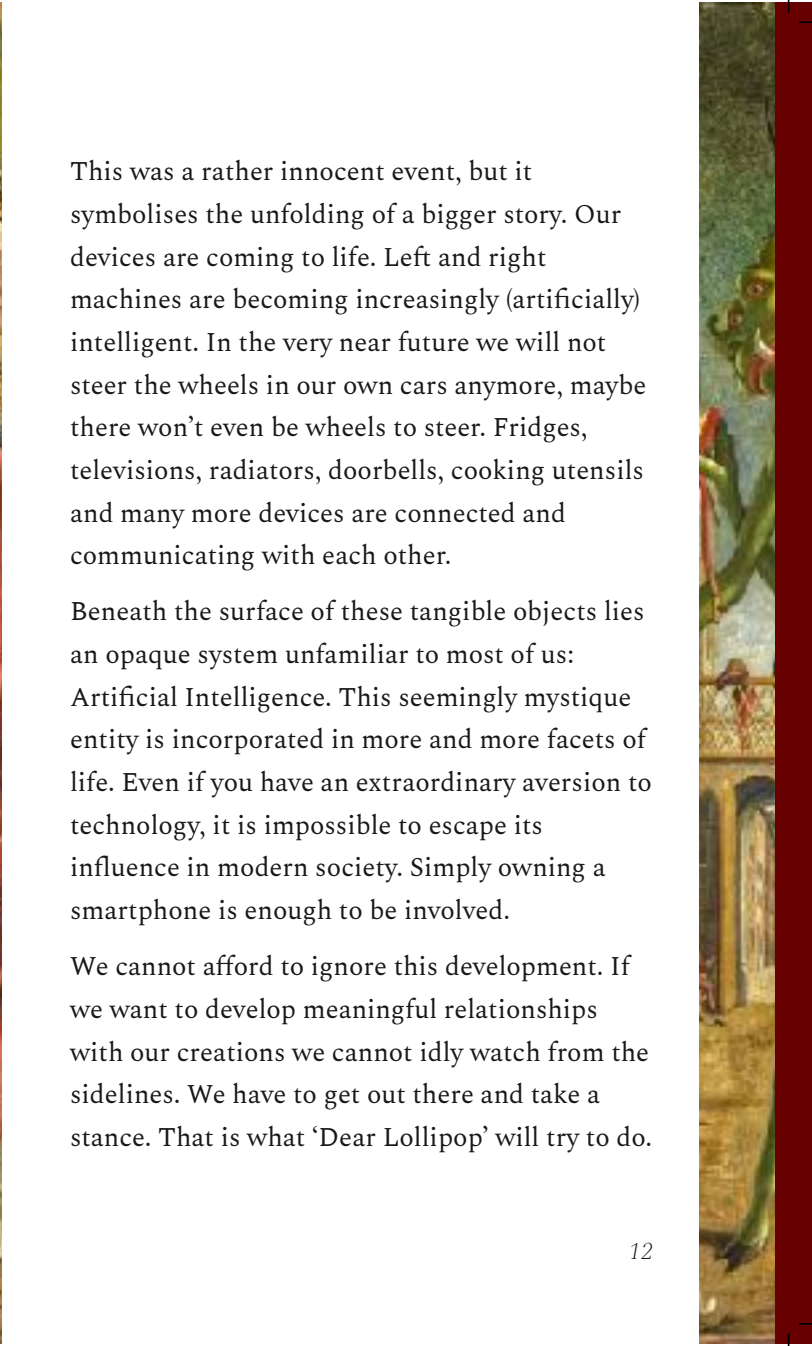

In early 2018 something strange happened in many households in the United States. All across the country, people were startled by the sudden sound of laughter, even when home alone. The source of this unsolicited, creepy sounding noise was the virtual assistant Alexa.

In households where people opt to install smart devices connected to each other - the Internet of Things - Alexa serves as the central hub.

Something was wrong with her. She mistakenly thought people commanded her to laugh. No one did.

A software update was dispatched to make sure she only responds to people specifically requesting her to laugh. As an extra safety measure, Alexa now also announces her laughter by saying 'Sure, I can laugh' (which is probably meant to be less creepy).


Order was restored, and everyone went one with their lives.



This was a rather innocent event, but it symbolises the unfolding of a bigger story. Our devices are coming to life. Left and right machines are becoming increasingly (artificially) intelligent. In the very near future we will not steer the wheels in our own cars anymore, maybe there won't even be wheels to steer. Fridges, televisions, radiators, doorbells, cooking utensils and many more devices are connected and communicating with each other.

Beneath the surface of these tangible objects lies an opaque system unfamiliar to most of us: Artificial Intelligence. This seemingly mystique entity is incorporated in more and more facets of life. Even if you have an extraordinary aversion to technology, it is impossible to escape its influence in modern society. Simply owning a smartphone is enough to be involved.

We cannot afford to ignore this development. If we want to develop meaningful relationships with our creations we cannot idly watch from the sidelines. We have to get out there and take a stance. That is what 'Dear Lollipop' will try to do.



This project has two parts: the book you are reading right now and a live performance.

This writing is a preliminary research to the latter, which will be performed exclusively for an audience of smartphones.

In the media there is a lot of buzz and hype surrounding this topic. Most of it has been about the consequences - the good and the bad - for the future of our society. Everything has been written from the perspective of humankind - with a few exceptions, mostly fiction works where the antagonist is some evil machine that hates humankind and wants to destroy it.

This book will try to gain a different insight in this discussion by assuming the perspective of the machine. Before we shapeshift into their spirit, some observations are needed. First I will try to elaborate on the relevance of the project, then I will try to give a concise description of the history and the current state of Artificial Intelligence and ultimately - after taking a short detour through the minds of animals - I will conclude with speculations on the artistic taste of intelligent machines.

Without further ado, let's get into their minds.






Dear Lollipop

The performance


During the fall of 2019, somewhere in east-Belgium near the border of Maastricht, a bunch of actors will perform for an audience of tripod-supported smartphones. As far as we (and Google) know, this will be the first ever live theatre performance created exclusively for the metaphorical eyes and ears of machines. They will experience the show through an app specifically designed for this project. No human



being will be allowed to sit by their side. The live experience is for them, and for them alone

For the actors, it will be a completely novel experience. “Surreal” might be a good word to describe it, because of several factors. For starters, they (most probably) won’t get a perceptible response from the audience. No one will cough, laugh or talk. No smalltalk will take place in the foyer after the performance. It might just feel like a runthrough.

And what about their taste? Their opinions? If there is one major difference between humans and technology, it is that humans will always have an opinion. After watching a live performance people will talk about how beautiful or frustratingly bad it was, or how they really don’t understand the need for yet another staging of Hamlet. When it comes to theatre, people and their tastes are not necessarily very predictable and vary a lot. Even with a critically acclaimed piece you can still find people who really dislike it. This is not to say that predicting what the audience likes and adjusting the art you to their tastes and desires is something to strive for.



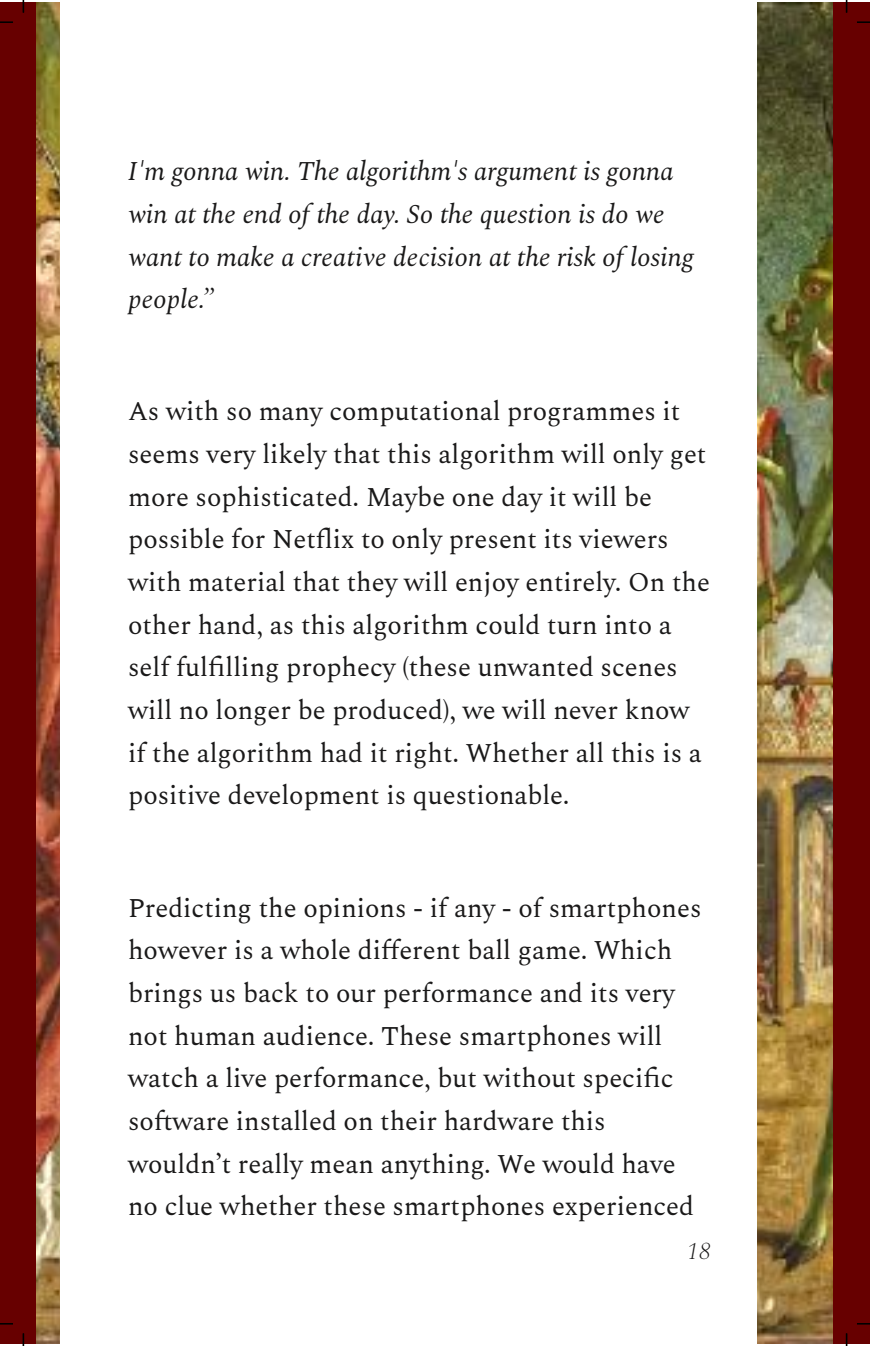
Hopefully every artist knows that if you're just creating something to please your audience, you're not on the right track.

Knowing what the audience wants

In television however, it seems that creating something exactly adjusted to the taste of the viewer might be a trend of the near future. Netflix in particular is making a lot of progress on this, because they have a vast amount of customer data. Based on this information they can determine when a show is likely to lose viewers. Cary Fukunaga, the director of the Netflix show *Maniac*, said in an interview¹ that he had to rewrite certain scenes because an algorithm reported that they would otherwise lose viewers:

“So they can look at something you're writing and say, we know based on our data that if you do this, we will lose this many viewers. So it's a different kind of note-giving. It's not like, let's discuss this and maybe


¹ <https://www.gq.com/story/cary-fukunaga-netflix-maniac>



I'm gonna win. The algorithm's argument is gonna win at the end of the day. So the question is do we want to make a creative decision at the risk of losing people."

As with so many computational programmes it seems very likely that this algorithm will only get more sophisticated. Maybe one day it will be possible for Netflix to only present its viewers with material that they will enjoy entirely. On the other hand, as this algorithm could turn into a self fulfilling prophecy (these unwanted scenes will no longer be produced), we will never know if the algorithm had it right. Whether all this is a positive development is questionable.

Predicting the opinions - if any - of smartphones however is a whole different ball game. Which brings us back to our performance and its very not human audience. These smartphones will watch a live performance, but without specific software installed on their hardware this wouldn't really mean anything. We would have no clue whether these smartphones experienced





anything or not. It would just look like a room full of metal cases resting on tripods.

A smartphone by itself can't really communicate its opinions on things. It needs to be programmed to do so. We want them to reflect on what they are seeing. So the question is: how? *How do you program a machine to reflect, to think, and possibly to feel?*


We humans view ourselves as autonomous beings. We're able to rationalize, reflect, think and feel. We experience the world around us and react to it, in the moment. Few people would argue that we - our minds - are pre-programmed, that the outcomes of our thoughts are predetermined. We experience our own thoughts existing in the here and now, admissible to our control and manipulation. We like to think that these thoughts are our own, unique and autonomous constructions.

Are machines the opposite of humans in this regard? Are they not significantly more predictable since we - or at least their programmers - are able to write their code, their



thoughts? If this is the case then can their reaction ever be considered as ‘a real opinion’? Is a response genuine if the phone is simply following lines of code? One way to look at this puzzle is to think of our brains as hardware too, just running lines of code. Maybe a more complicated code, but a code nevertheless. If you view the connectivity between our brain cells as biological algorithms then we might discover that we don’t operate so differently from smartphones after all.

One may argue that the distinctive difference between ourselves and machines is our ability to feel emotions. If the latter are indeed entirely incapable of experiencing any feeling whatsoever and the intention of the performance is to evoke some kind of emotional response, then that’s a problem. How to provoke this kind of reaction in our solemn audience of smartphones? Once again, it might become possible by programming different emotional states and responses in their software. But then one might argue that if an emotional response is pre-programmed, it is by





definition not a 'genuine' feeling. Intuitively we are inclined to think of emotions as something uniquely human - perhaps almost as something mystique.

If we *have* to find similarities between our mind and that of a machine, we quickly end up talking about rational decision making and calculations. *Intuitively*, rationality originating from the human brain *feels* more similar to the processes of a machine than emotionality originating from the same brain. As the free use of italics in this sentence suggests, we're digging ourselves into an even deeper loophole of thoughts.

A machine can not really be angry or sad, right? But if emotion originates from the brain and the connectivity between its cells can be seen as algorithms, then it should be possible for an artificial algorithm to create an emotion as well.

This rings especially true in light of the book *How Emotions are Made*² by Lisa Feldman Barrett,


² How Emotions are Made is a book by Lisa Feldman Barrett, published by Pan Macmillan in 2017.



who claims that emotions “are not triggered; you create them. They emerge as a combination of the physical properties of your body, a flexible brain that wires itself to whatever environment it develops in, and your culture and upbringing, which provide that environment.”

These reflections are somewhat misleading however, since they disregard the fact that machines don't have to experience emotion the same way we do. Why would we paste our emotional capabilities onto something that has a completely different physiology and history than us? For now, let's start deconstructing this problem by simply proposing that the standard state that machines are in - indifferent but compliant - is an emotion in itself.

As an artist I find the idea of being able to program the taste of my audience very peaceful, in a perverse kind of way. Yet the question remains; *how* do you program the responses.





Contemporary technology is not nearly sophisticated enough to let them autonomously formulate opinions about something so abstract and unstable as performance art.

One solution to this is to program a plethora of observable responses in the machine, and then let a random function decide on which one is going to be expressed at any given time.

In short: cheating. Or is this how the 'self-conscious, always knowing better, exquisite taste audience' is rigged?

The other - even less exciting - option is programming a specific outcome that will always happen. That has to be the wet dream of every narcissistic, affirmation-seeking, validation-dependent artist that ever lived. Being able to program the taste of the audience means you can do whatever you want on that stage and no matter what it is the crowd will adore you entirely and think the show is the best thing they've ever seen. Or you can program them to despise you, if you're the type of artist that gets off on that. Full-blown narcissism isn't very sexy




and total masochism can be wildly uncomfortable, so your best course of action is probably to choose a healthy mix between the former and the latter.

Dear Lollipop, are you relevant?

As far as we know, no one has ever made a live performance exclusively for smartphones. Twelve years after the first iPhone came to fruition, Dear Lollipop will be the first to do just that - no humans allowed. But why?

First of all, the form itself is a statement that forces the (human) audience to reconsider their perspective on the interaction between themselves and their technology. Being able to view things from a different perspective broadens your understanding of the world and strengthens your ability to empathize.

There is a tectonic shift underway in the realm of technology and Artificial Intelligence (A.I.). The top five companies in the world by market capitalization are no longer filled by banks and

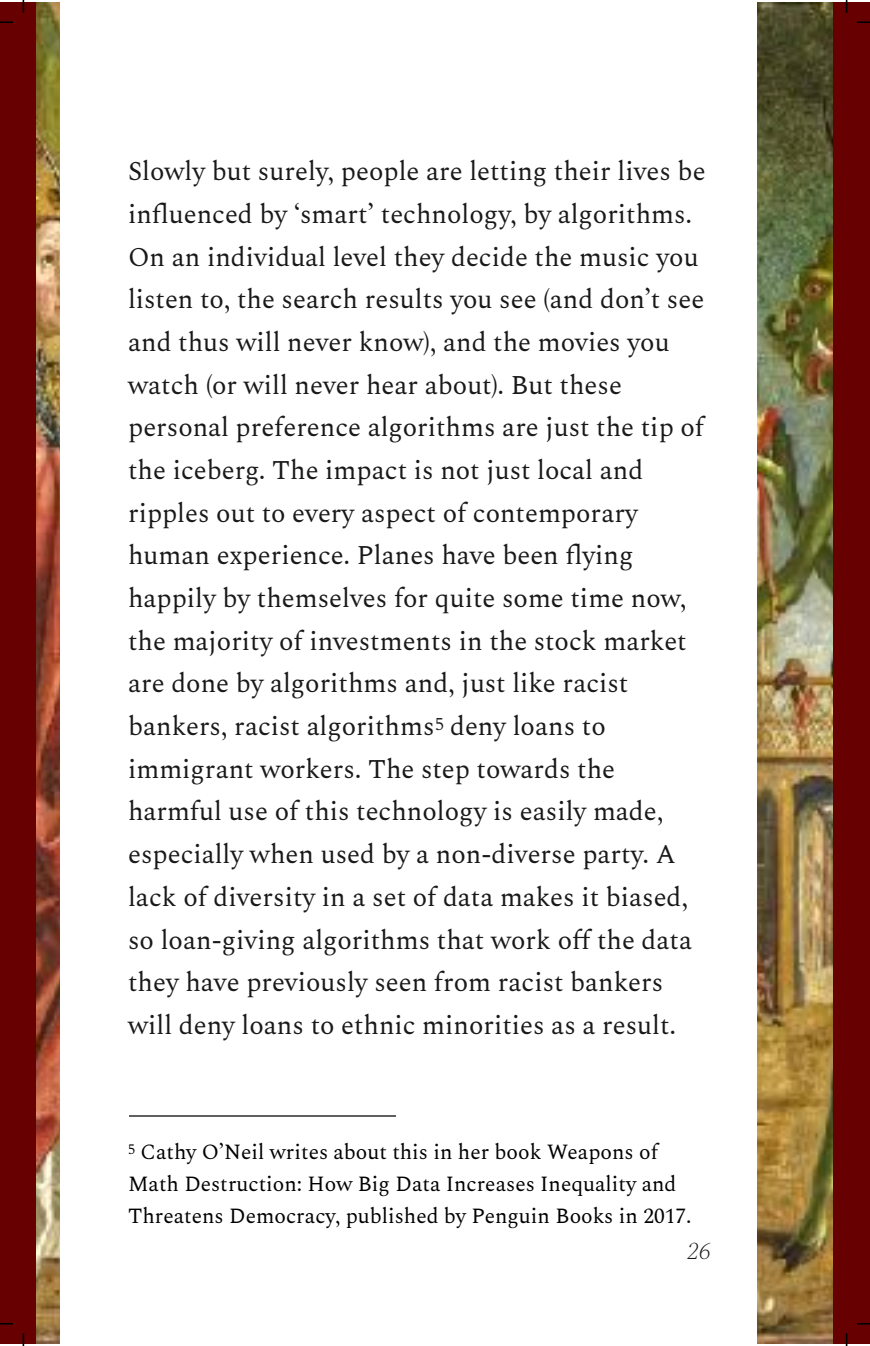


oil companies but by the tech industry³. All these companies continue to increase their A.I. investment budgets. Forrester research estimates that by 2020 the A.I. market will reach a value of \$1.2 trillion⁴. And the customer is down to p(l)ay. They do not really have a choice. The conception of portable phones is not the result of a huge demand from customers for portable phoning: it just became technically possible. The same goes for smartphones. Refusing to use them becomes an ever more futile and impossible attitude. In the near future, there won't be any wires to phone through anymore.

In this inevitable road down (or up) the path towards an ever more technology driven society, money will keep pouring into the pockets of big tech companies.


³ https://en.wikipedia.org/wiki/List_of_public_corporations_by_market_capitalization#2019

⁴ <https://www.forbes.com/sites/blakemorgan/2018/06/06/how-much-money-has-poured-into-ai-and-customer-experience/#2562b9037ed2>



Slowly but surely, people are letting their lives be influenced by ‘smart’ technology, by algorithms. On an individual level they decide the music you listen to, the search results you see (and don’t see and thus will never know), and the movies you watch (or will never hear about). But these personal preference algorithms are just the tip of the iceberg. The impact is not just local and ripples out to every aspect of contemporary human experience. Planes have been flying happily by themselves for quite some time now, the majority of investments in the stock market are done by algorithms and, just like racist bankers, racist algorithms⁵ deny loans to immigrant workers. The step towards the harmful use of this technology is easily made, especially when used by a non-diverse party. A lack of diversity in a set of data makes it biased, so loan-giving algorithms that work off the data they have previously seen from racist bankers will deny loans to ethnic minorities as a result.

⁵ Cathy O’Neil writes about this in her book *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*, published by Penguin Books in 2017.

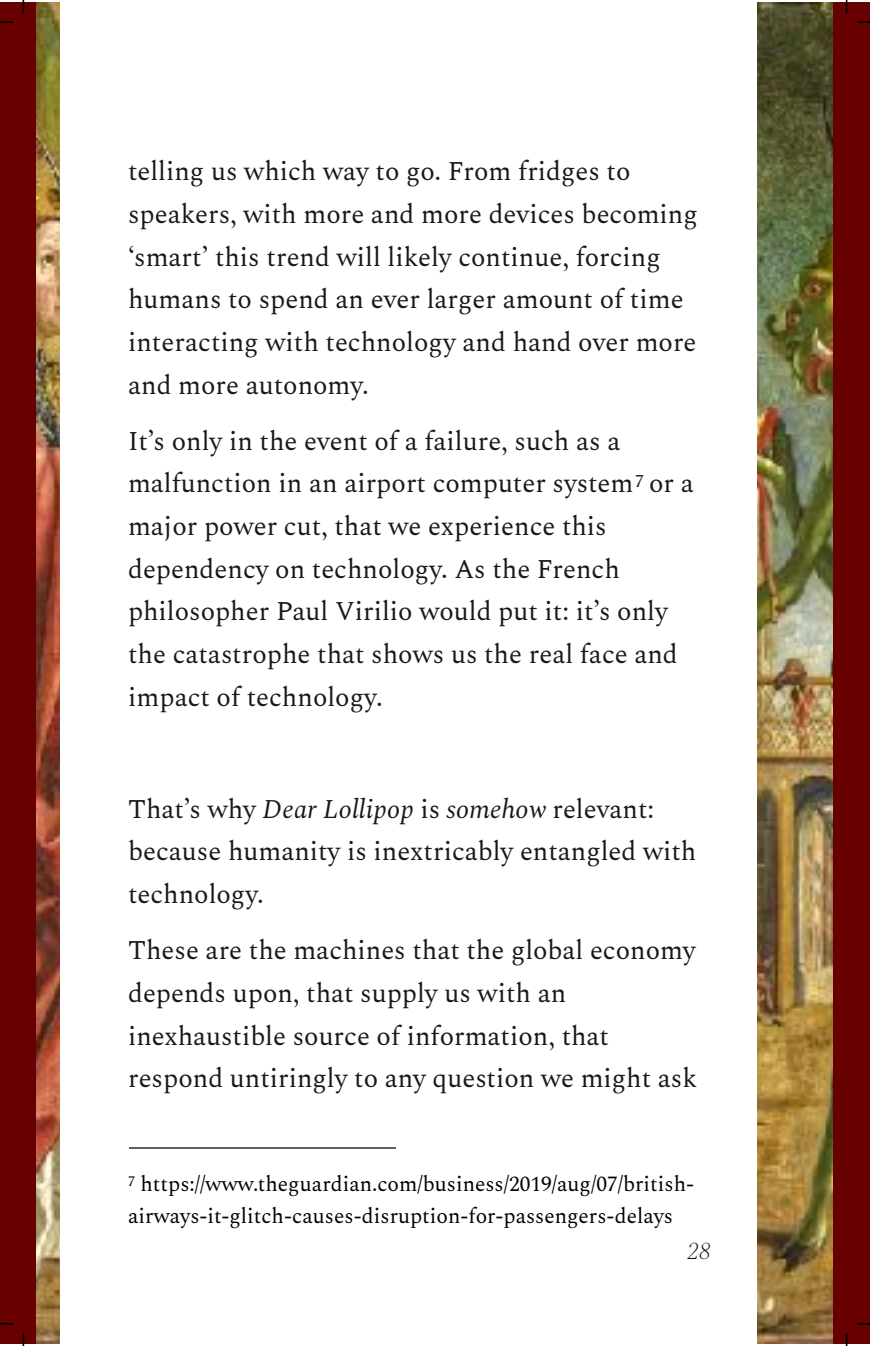


As with humans, their biases are a self fulfilling prophecy.

In the hands of a state with little regard for privacy and human rights, algorithms can be equally destructive. Probably the most striking example of such questionable programmes is seen in China. A social credit system⁶ will be rolled out nationwide in 2020. They rigged the biggest cities with millions of cameras equipped with facial recognition technology. If they see you performing 'bad' behaviour, such as littering, smoking in non-smoking areas or jaywalking, you get a lower score. Citizens are then rewarded or punished according to their score. Such disciplinary consequences range from not getting a credit anymore, over air travel bans - received by nine million people in 2018 - to your dog being taken away.

Nevertheless we readily and willingly hand over a big part of our autonomy to technology. We would often be literally lost without algorithms

⁶ https://en.wikipedia.org/wiki/Social_Credit_System




telling us which way to go. From fridges to speakers, with more and more devices becoming 'smart' this trend will likely continue, forcing humans to spend an ever larger amount of time interacting with technology and hand over more and more autonomy.

It's only in the event of a failure, such as a malfunction in an airport computer system⁷ or a major power cut, that we experience this dependency on technology. As the French philosopher Paul Virilio would put it: it's only the catastrophe that shows us the real face and impact of technology.

That's why *Dear Lollipop* is somehow relevant: because humanity is inextricably entangled with technology.



These are the machines that the global economy depends upon, that supply us with an inexhaustible source of information, that respond untiringly to any question we might ask

⁷ <https://www.theguardian.com/business/2019/aug/07/british-airways-it-glitch-causes-disruption-for-passengers-delays>



it, that switch our lights and speakers and heaters on and off, that open our doors, that make us oblivious for any sense of geography because the gps knows better anyway... in short, these machines are changing the very way our brains are wired. We desperately need to familiarise ourselves with these things. And by “familiarise” I don’t just mean the technical comprehension of their design and processes. I mean making a leap towards such a level of understanding that we might even discover what kind of art a smartphone would enjoy.

One could argue that there is nothing wrong with letting a device decide that you must listen to Beyonce or tell you that you should drive around the city centre to get home because of a traffic jam, and one could be perfectly right about that. But a very important idea to keep in mind here is that a lot of these decisions have some kind of ethical implications. Why do we know Beyonce in the first place, and not artist x? And what if we all start avoiding all traffic jams: what would be the consequence for life in the




affected neighbourhoods? Unknowingly, we are handing over ethical decision-making. That actually might be the most important reason why a public debate on A.I. is urgent.

Twenty years ago Artificial Intelligence might have been software correcting your spelling mistakes or beating you at chess (which was painful enough), now it's a code that a self-driving car runs seconds before crashing, deciding whether to kill the elderly lady crossing the street or to smash into a tree and kill the driver.

Before handing over ethical decision making to cars, we as a society should first determine what we find ethical or unethical.

Precisely because of this reason researchers at the M.I.T. developed the online platform Moral Machine⁸. They gathered 40 million decisions from millions of people in 233 countries on moral dilemmas involving autonomous cars.



⁸ <http://moralmachine.mit.edu/>



Moral dilemmas like: the brakes are broken, should the car hit the three elderly ladies crossing the street, or swerve, thereby killing the three passengers of the car? Their aim is to contribute to developing universal machine ethics. A goal not so easily achieved, it appears: latin-american countries are more inclined to spare females, whereas people from collectivistic cultures such as China care less about saving the young, because of the emphasis on respect for the elderly in those cultures. But the platform did report a global preference: sparing humans over animals (as long as it is not my dog), sparing more lives and sparing younger lives.


Obviously it is very hard to predict how many times these specific dilemma's will actually occur, but nevertheless a decision on the outcome has to be made beforehand.

A new profession arises: ethical programming. What if countries want to decide their specific ethical code themselves? One can imagine a scenario where a self-driving car switches to a different set of ethical values when crossing borders.



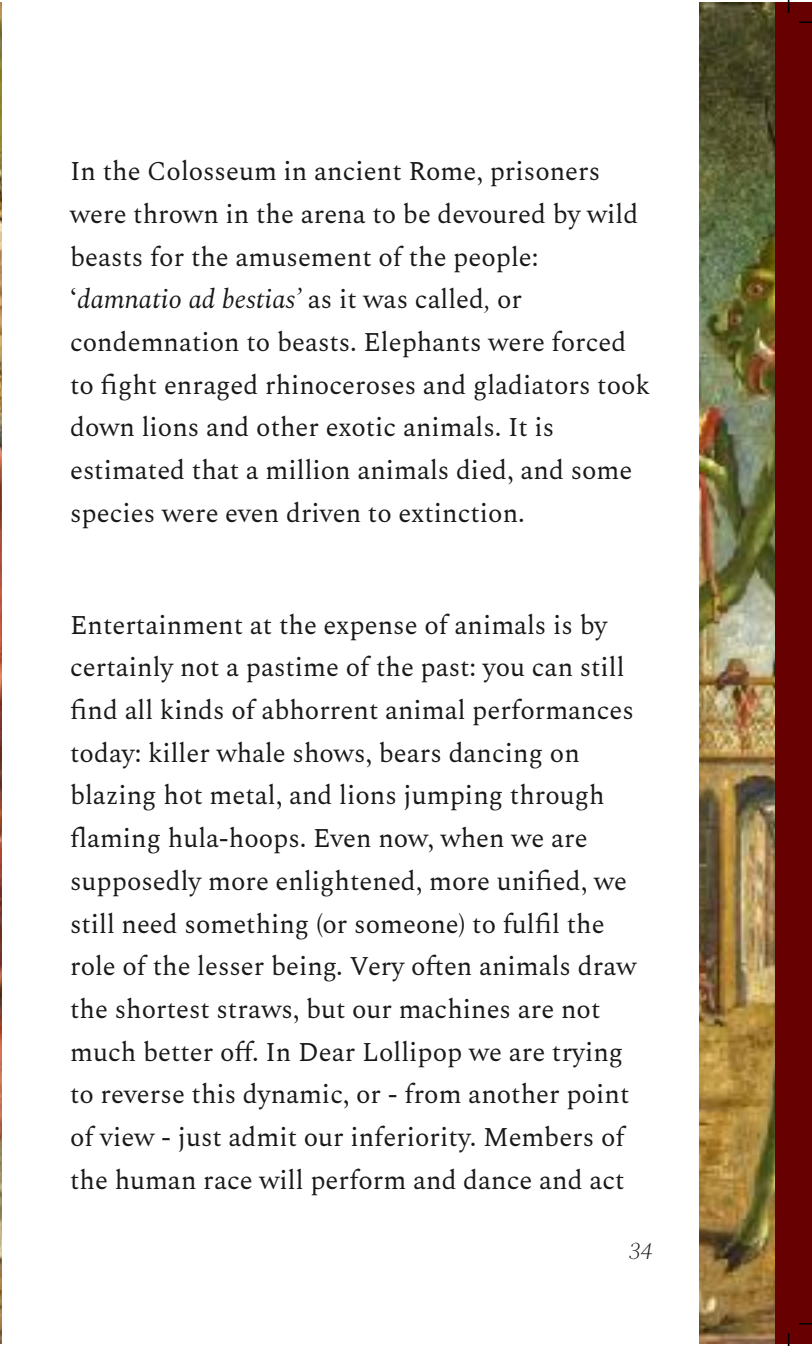

It's not just cars that need to make these ethical decisions on life and death. With the development of autonomous weapons, the military has to deal with these problems as well. Google supplied their A.I. knowledge to the Defense Department of the U.S. - knowledge that could and will be used for algorithm warfare - but they recently decided not to renew its contract with the Pentagon after extensive protests from their own employees. Obviously the Pentagon will find a different way to continue investing in A.I. to use on the battlefield. So alongside cars, we could soon also have A.I. weapons deciding on who to kill and who to spare.

This whole new set of ethical problems that A.I. brings with it is one of the major reasons of why a public debate on the influence of technology is important. These technologies will reflect our reality, and our values will (unconsciously or not) be projected on them. But society as it is now, is ethically far from perfect. So before we teach these machines our values, and knowingly or not




inject them with our racism and bigotry, maybe we should first look inwards to update those values. Can we train machines for an ethical paradise? And shouldn't we get to know the intricacies of our own inventions better, before stamping our human code of ethics onto them? That is why making a show for smartphones could be enlightening. By assuming their perspective we might uncover something. Maybe something about them, but even more likely something about ourselves.

There are some other possible reasons why this performance could be relevant. For example: the futility of this endeavour has an important symbolic meaning. So many times in history have 'lesser' beings performed, danced and humiliated themselves for the entertainment of 'greater' beings. For example, through part of recent history people with extraordinary diseases and conditions were forced to perform in 'freak shows'.



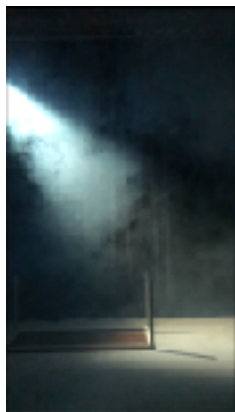
In the Colosseum in ancient Rome, prisoners were thrown in the arena to be devoured by wild beasts for the amusement of the people: '*damnatio ad bestias*' as it was called, or condemnation to beasts. Elephants were forced to fight enraged rhinoceroses and gladiators took down lions and other exotic animals. It is estimated that a million animals died, and some species were even driven to extinction.

Entertainment at the expense of animals is by certainly not a pastime of the past: you can still find all kinds of abhorrent animal performances today: killer whale shows, bears dancing on blazing hot metal, and lions jumping through flaming hula-hoops. Even now, when we are supposedly more enlightened, more unified, we still need something (or someone) to fulfil the role of the lesser being. Very often animals draw the shortest straws, but our machines are not much better off. In Dear Lollipop we are trying to reverse this dynamic, or - from another point of view - just admit our inferiority. Members of the human race will perform and dance and act



until they sweat to entertain the audience of machines. This reversal of roles is an important exercise in genuine humility for an overly arrogant species.

Finally, our smartphones deserve to be the entertained instead of the entertainers for a change. With the increasing focus on technology in the arts, there already is plenty of art out there whereby machines perform for an audience of people. The switching of these roles acknowledges the possibility that machines have, or one day might have, a subjective experience of their own and could become ethically and intellectually superior.

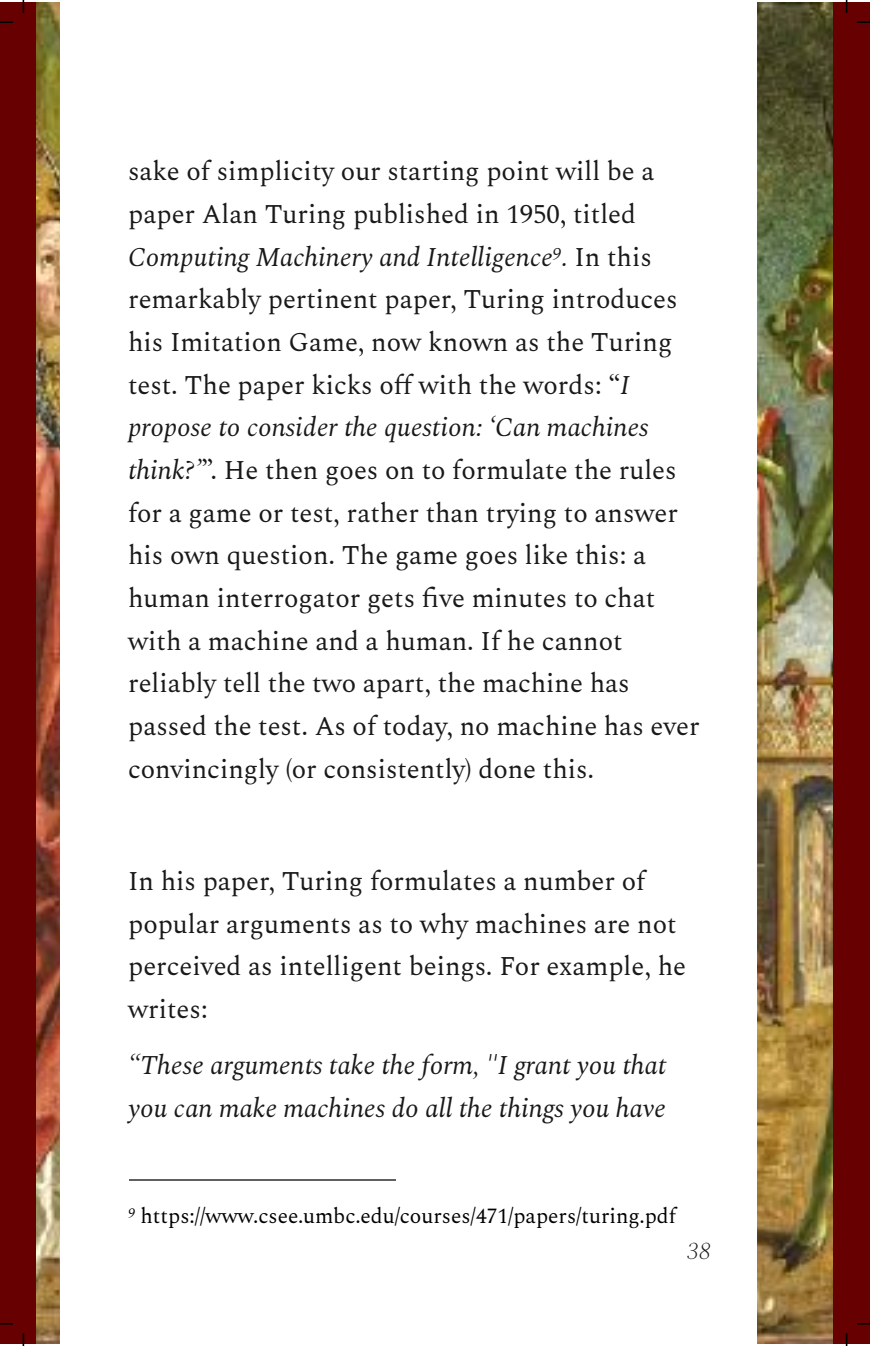




The rise of Artificial Intelligence

Alan Turing's kickstart

The importance of the last argument lies in the consequences this might have for the quality of our lives - or the lives of the generations to come. Before we dive into these consequences, let's take a look at some defining moments in the history of Artificial Intelligence. We could go back hundreds of years, all the way to the Golem in Jewish culture or to Kami, a living spirit the Japanese believed to reside in objects, but for the




sake of simplicity our starting point will be a paper Alan Turing published in 1950, titled *Computing Machinery and Intelligence*⁹. In this remarkably pertinent paper, Turing introduces his Imitation Game, now known as the Turing test. The paper kicks off with the words: “*I propose to consider the question: ‘Can machines think?’*”. He then goes on to formulate the rules for a game or test, rather than trying to answer his own question. The game goes like this: a human interrogator gets five minutes to chat with a machine and a human. If he cannot reliably tell the two apart, the machine has passed the test. As of today, no machine has ever convincingly (or consistently) done this.

In his paper, Turing formulates a number of popular arguments as to why machines are not perceived as intelligent beings. For example, he writes:

“These arguments take the form, “I grant you that you can make machines do all the things you have



⁹ <https://www.csee.umbc.edu/courses/471/papers/turing.pdf>



mentioned but you will never be able to make one to do X." Numerous features X are suggested in this connexion. I offer a selection: Be kind, resourceful, beautiful, friendly, have initiative, have a sense of humour, tell right from wrong, make mistakes, fall in love, enjoy strawberries and cream, make someone fall in love with it, learn from experience, use words properly, be the subject of its own thought, have as much diversity of behaviour as a man, do something really new.

No support is usually offered for these statements. I believe they are mostly founded on the principle of scientific induction. A man has seen thousands of machines in his lifetime. From what he sees of them he draws a number of general conclusions. They are ugly, each is designed for a very limited purpose, when required for a minutely different purpose they are useless, the variety of behaviour of any one of them is very small, etc., etc. Naturally he concludes that these are necessary properties of machines in general. ”

This line of reasoning remains roughly the same until this very day. Whenever a machine has




achieved something impressive, people acknowledge the feat, but are also quick to dismiss other theoretical possibilities of the machine as unlikely, thereby making the list of ‘impossible X-es’ longer and longer. In a way, every new X is a new idea, a possibility. If the idea is captivating enough, and if it is physically possible, it only needs enough geniuses, luck and time to come to fruition. Like a prophecy waiting to be fulfilled.

This apparently also goes for ideas that are stupid - or frivolous, as Turing would put it:

“The inability to enjoy strawberries and cream may have struck the reader as frivolous. Possibly a machine might be made to enjoy this delicious dish, but any attempt to make one do so would be idiotic.”


He could just as well have written: *“Possibly a machine might be made to enjoy **a live performance**, but any attempt to make one do so would be idiotic.”*

Maybe, but I for one am glad to partake in this frivolous, idiotic endeavour.



Just to be clear: a healthy dose of skepticism is justified when every so often a new A.I. story headlines. Media and celebrities tend to react overly enthusiastic or ridiculously naive. When the humanoid robot Sophia¹⁰ was introduced to the public in 2016, media worldwide jumped on the story. The reports were impressive: Sophia can make 50 facial expressions, have a conversation and see and react on her environment. She received Saudi Arabian citizenship, becoming the first robot ever to have a nationality. All across the world she appeared on television shows - such as Jimmy Fallon's Tonight Show, where the creator claimed that she is basically alive - to demonstrate her abilities. But what was kept from the spotlights is the fact that Sophia's conversation skills are very limited, and that all her verbal responses are strictly pre-programmed. Upon closer examination it turned out that Sophia was just a pretty mediocre chatbot with a face. This rather creepy robot was a mirage that had little to do

¹⁰ [https://en.wikipedia.org/wiki/Sophia_\(robot\)](https://en.wikipedia.org/wiki/Sophia_(robot))




with A.I. improvements. Not something to be globally excited about.

In any case, Sophia was probably not what Turing had in mind when he fantasized about a thinking machine. What he did envisage was the basic concept of machine learning. Turing wrote 5 pages on the topic of ‘learning machines’, and he proposed different ways in which to achieve this goal. One was to mirror the learning process of a child, aptly named ‘The child machine’.

He writes:

“Instead of trying to produce a programme to simulate the adult mind, why not rather try to produce one which simulates the child's? If this were then subjected to an appropriate course of education one would obtain the adult brain. Presumably the child brain is something like a notebook as one buys it from the stationer's. Rather little mechanism, and lots of blank sheets. [...] The machine has to be so constructed that events which shortly preceded the occurrence of a punishment signal are unlikely to be



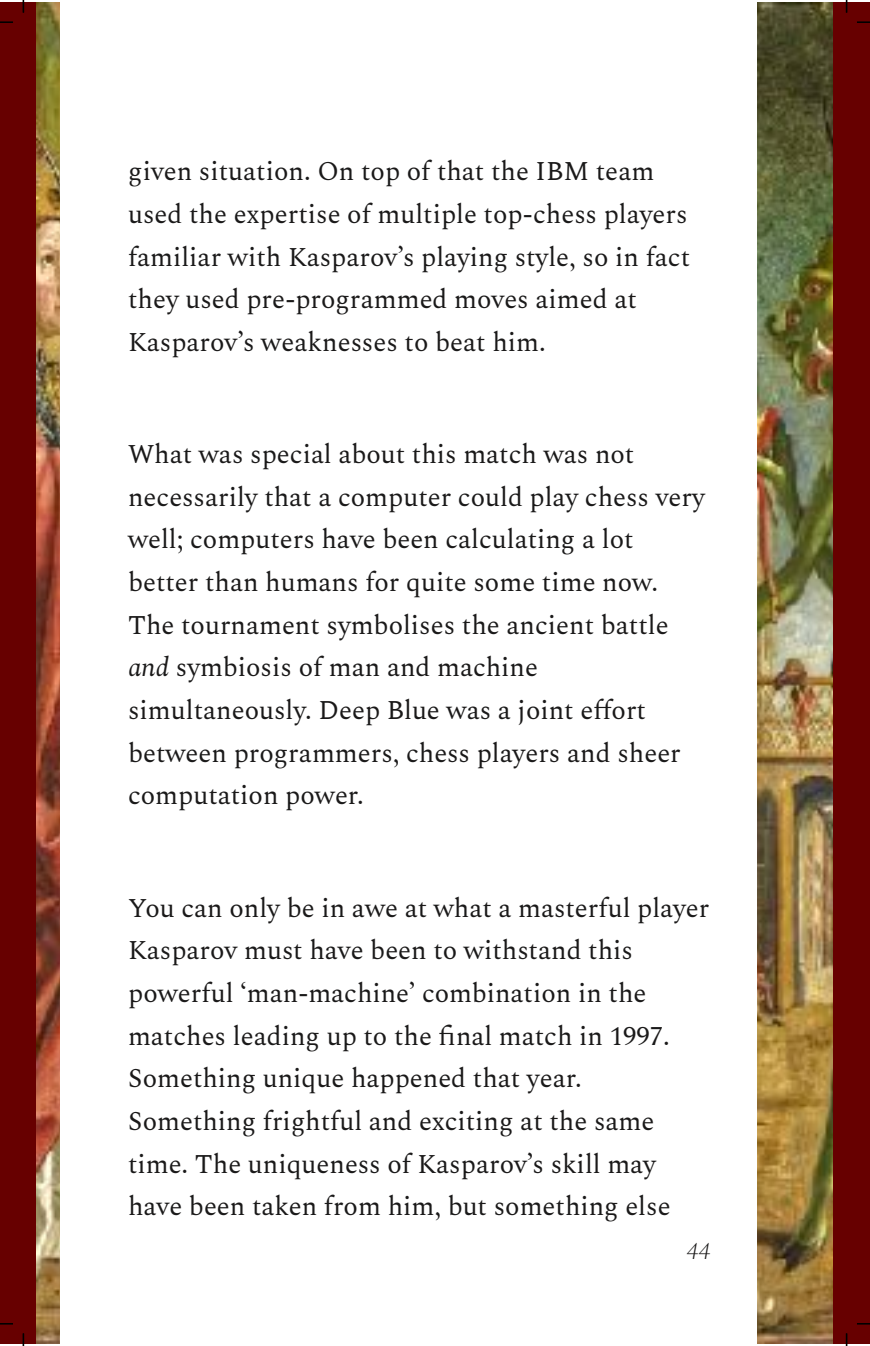
repeated, whereas a reward signal increased the probability of repetition of the events which led up to it.”

Checkmate

Today the myriad ways in which a machine can learn to improve on a task could be regarded as the offspring of Turing’s ideas. Like a prophecy waiting to be fulfilled...

Six years after the publication of his paper, the term Artificial Intelligence was coined at the Dartmouth workshop - a brainstorm session held by eleven scientists.


It took an additional 41 years for another fantasy of Turing to become reality: IBM’s Deep Blue defeated Garry Kasparov in a chess tournament. In all fairness the computer did not use A.I. algorithms as we know them today. Instead it relied on a brute force approach. By calculating millions of moves in just seconds, it could quickly decide on the most optimal move in any



given situation. On top of that the IBM team used the expertise of multiple top-chess players familiar with Kasparov's playing style, so in fact they used pre-programmed moves aimed at Kasparov's weaknesses to beat him.

What was special about this match was not necessarily that a computer could play chess very well; computers have been calculating a lot better than humans for quite some time now. The tournament symbolises the ancient battle *and* symbiosis of man and machine simultaneously. Deep Blue was a joint effort between programmers, chess players and sheer computation power.

You can only be in awe at what a masterful player Kasparov must have been to withstand this powerful 'man-machine' combination in the matches leading up to the final match in 1997. Something unique happened that year. Something frightful and exciting at the same time. The uniqueness of Kasparov's skill may have been taken from him, but something else



was given in return: the new possibilities of cooperation between man and machine. In the years after the match Kasparov developed Advanced Chess, a form of chess where players form a team with a computer to combine the best of both worlds - creativity and computation. We must not run away when a machine beats us, we must embrace it, says Kasparov in his book 'Deep Thinking'¹¹:

“When I sat across from Deep Blue twenty years ago I sensed something new, something unsettling. Perhaps you will experience a similar feeling the first time you ride in a driverless car, or the first time your new computer boss issues an order at work. We must face these fears in order to get the most out of our technology and to get the most out of ourselves.”

¹¹ Deep Thinking: Where Machine Intelligence Ends and Human Creativity Begins is a book written by Gary Kasparov, published by Hodder & Stoughton General Division in 2017.


The state of A.I. today

As Turing predicted, machine learning is the way to go. Because of it, Artificial Intelligence is rapidly getting more and more intelligent, allowing it to master more complex tasks.

Machine learning knows different approaches to teach a machine to do something.

One way is Deep Learning. Google uses this approach to categorize photos based on the presence of cats. The learning algorithm is fed an enormous amount of photographs that either contain cats or not. By recognizing different patterns on photos with cats, the program is able to build a model - it learns what a photographed cat looks like. Based on that model it can recognize cats on new pictures with a relatively small error margin.

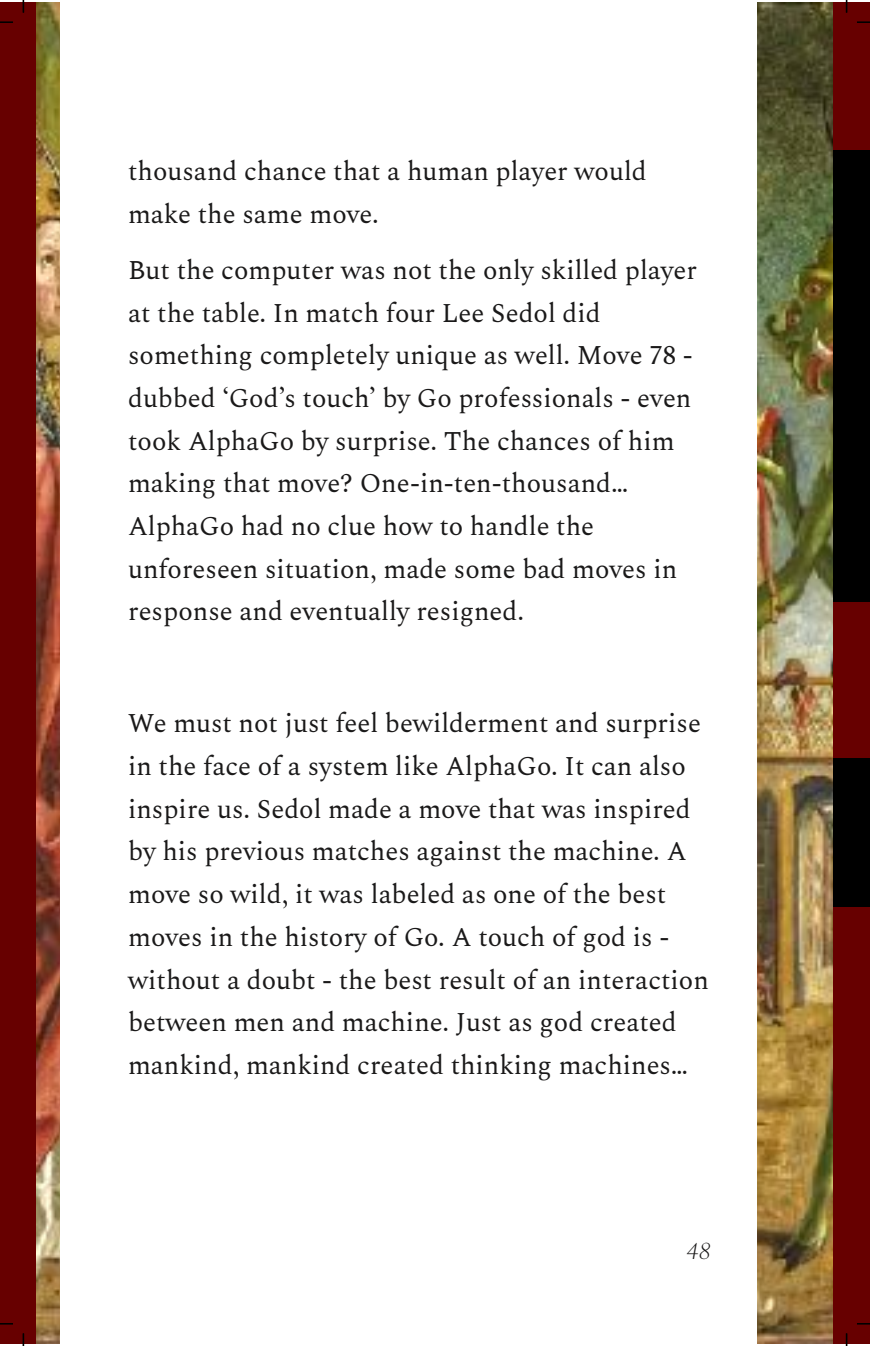
Another approach is reinforcement learning - which is basically learning by trial and error. For example: a robot trying to learn to walk by falling, standing up, and not making the same



step again. A landmark in this approach is Google Deepmind's AlphaGo, aimed at beating the top-players in the world at the very hard and complex game Go.

The programmers fed AlphaGo millions of moves by expert players. It then went on to play millions of games against itself, figuring out the best ways to win along the way. In contrast to the brute-force approach by the Deep Blue chess machine, AlphaGo used a neural network to figure out its next move. Lee Sedol, one of the top Go players in the world, was defeated by the program in 2016. Sedol played five matches against the computer, and won only the fourth. These matches were by no means dull or predictable.¹² In the second match AlphaGo made move 37 - a mystical move in hindsight, something unlike anything a human being would ever do in that spot. As a matter of fact, the machine analysed that there was a one-in-ten-

¹² If you want to see for yourself how spectacular a board game can be, check out the documentary AlphaGo (2017), available on Netflix.



thousand chance that a human player would make the same move.

But the computer was not the only skilled player at the table. In match four Lee Sedol did something completely unique as well. Move 78 - dubbed 'God's touch' by Go professionals - even took AlphaGo by surprise. The chances of him making that move? One-in-ten-thousand... AlphaGo had no clue how to handle the unforeseen situation, made some bad moves in response and eventually resigned.

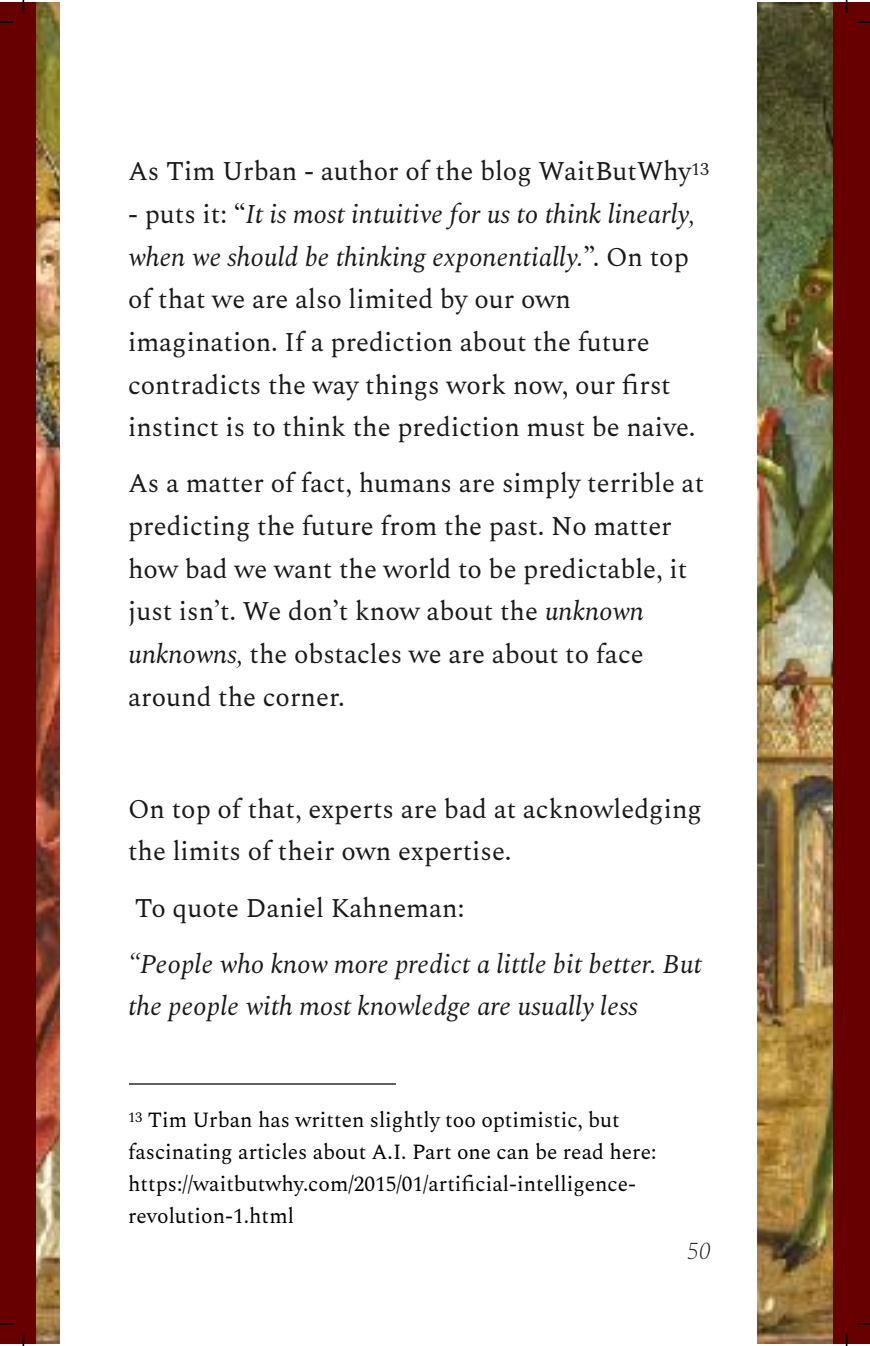
We must not just feel bewilderment and surprise in the face of a system like AlphaGo. It can also inspire us. Sedol made a move that was inspired by his previous matches against the machine. A move so wild, it was labeled as one of the best moves in the history of Go. A touch of god is - without a doubt - the best result of an interaction between men and machine. Just as god created mankind, mankind created thinking machines...



The pitfalls of A.I.

The defeat happened much sooner than the experts thought it would happen. This turns out to be a recurring theme in most milestone A.I. achievements: experts involved grossly underestimate the pace of the development, and are surprised when progress is achieved much sooner than expected. Even expert Google Translate programmers showed their astonishment at the faster-than-expected progress of the translating machine - only to then claim it could never translate a book.

Naturally, no one is able to accurately predict the future before it actually happens. But it seems as if experts are specifically worse than laymen when making predictions about their own field of expertise. A couple of reasons spring to mind as to why this is the case. One explanation might be that, when it comes to history, humans are inclined to think in linear evolutions. When we think about progress in the coming five years, we tend to look at the progress of the last five years as an indicator of how much will likely happen.



As Tim Urban - author of the blog WaitButWhy¹³ - puts it: *“It is most intuitive for us to think linearly, when we should be thinking exponentially.”* On top of that we are also limited by our own imagination. If a prediction about the future contradicts the way things work now, our first instinct is to think the prediction must be naive.


As a matter of fact, humans are simply terrible at predicting the future from the past. No matter how bad we want the world to be predictable, it just isn't. We don't know about the *unknown unknowns*, the obstacles we are about to face around the corner.

On top of that, experts are bad at acknowledging the limits of their own expertise.

To quote Daniel Kahneman:

“People who know more predict a little bit better. But the people with most knowledge are usually less

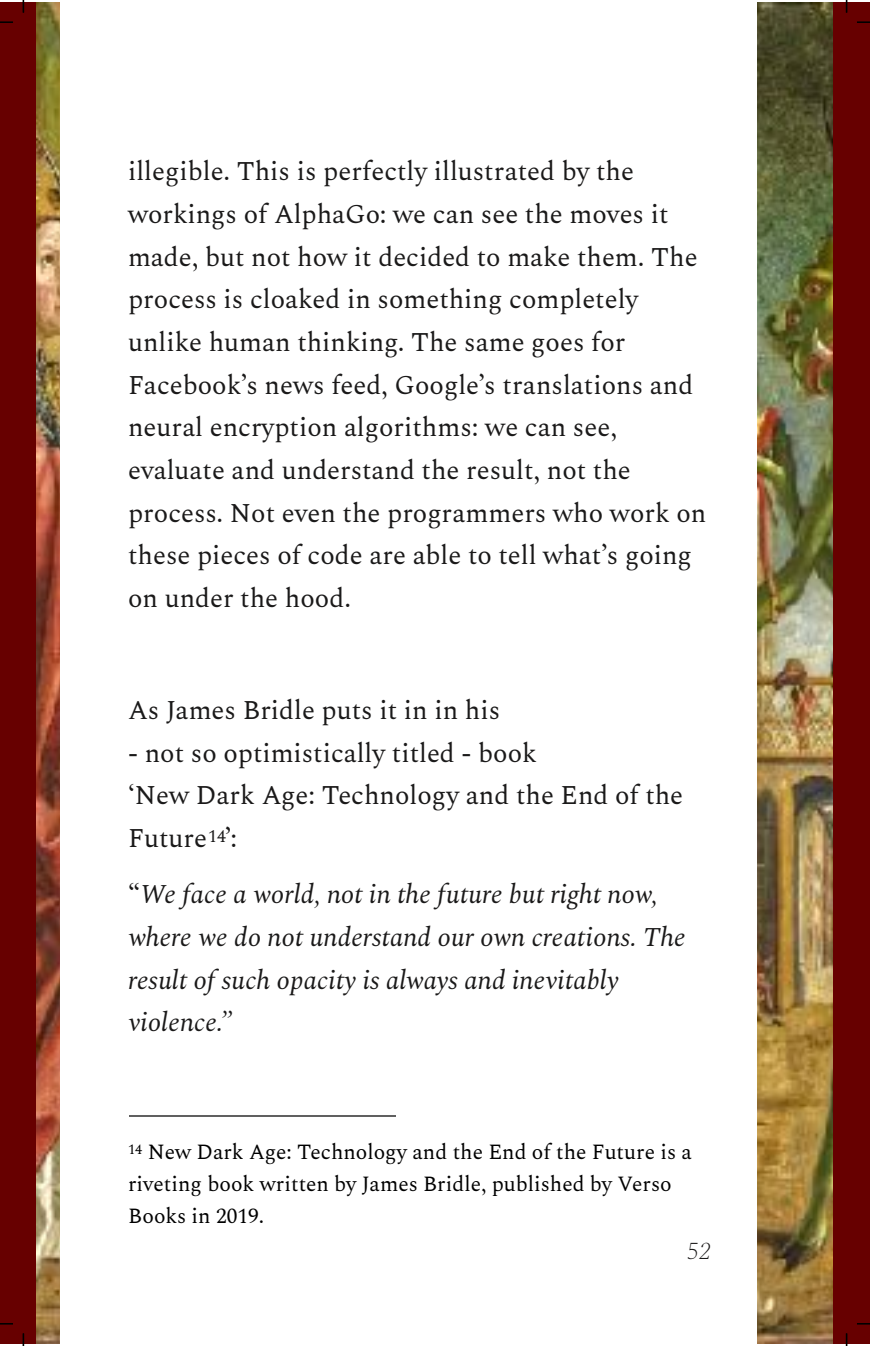
¹³ Tim Urban has written slightly too optimistic, but fascinating articles about A.I. Part one can be read here: <https://waitbutwhy.com/2015/01/artificial-intelligence-revolution-1.html>



reliable. The reason is that because someone who has more knowledge than others, also develops a stronger illusion of skill, and by being overly confident loses its sight on reality.”

He illustrates the extent of this by showing a research where it appeared that the more self-confident a political analyst is on a news show, the more likely his prediction will be wrong. Knowing a lot about the history of the middle east really does not mean you know where it is heading to. Human experts, - contrary to algorithms - are inclined to be biased towards over-confidence. They tend to remember (and be remembered and paid for) their successes and forget their failures. An algorithm just doesn't care: they just try to improve themselves based on pure facts.

On top of that, as algorithms are getting more and more complex - amidst all the new inventions designed to improve the way machines learn, inside the neural networks - their inner workings are getting increasingly




illegible. This is perfectly illustrated by the workings of AlphaGo: we can see the moves it made, but not how it decided to make them. The process is cloaked in something completely unlike human thinking. The same goes for Facebook's news feed, Google's translations and neural encryption algorithms: we can see, evaluate and understand the result, not the process. Not even the programmers who work on these pieces of code are able to tell what's going on under the hood.

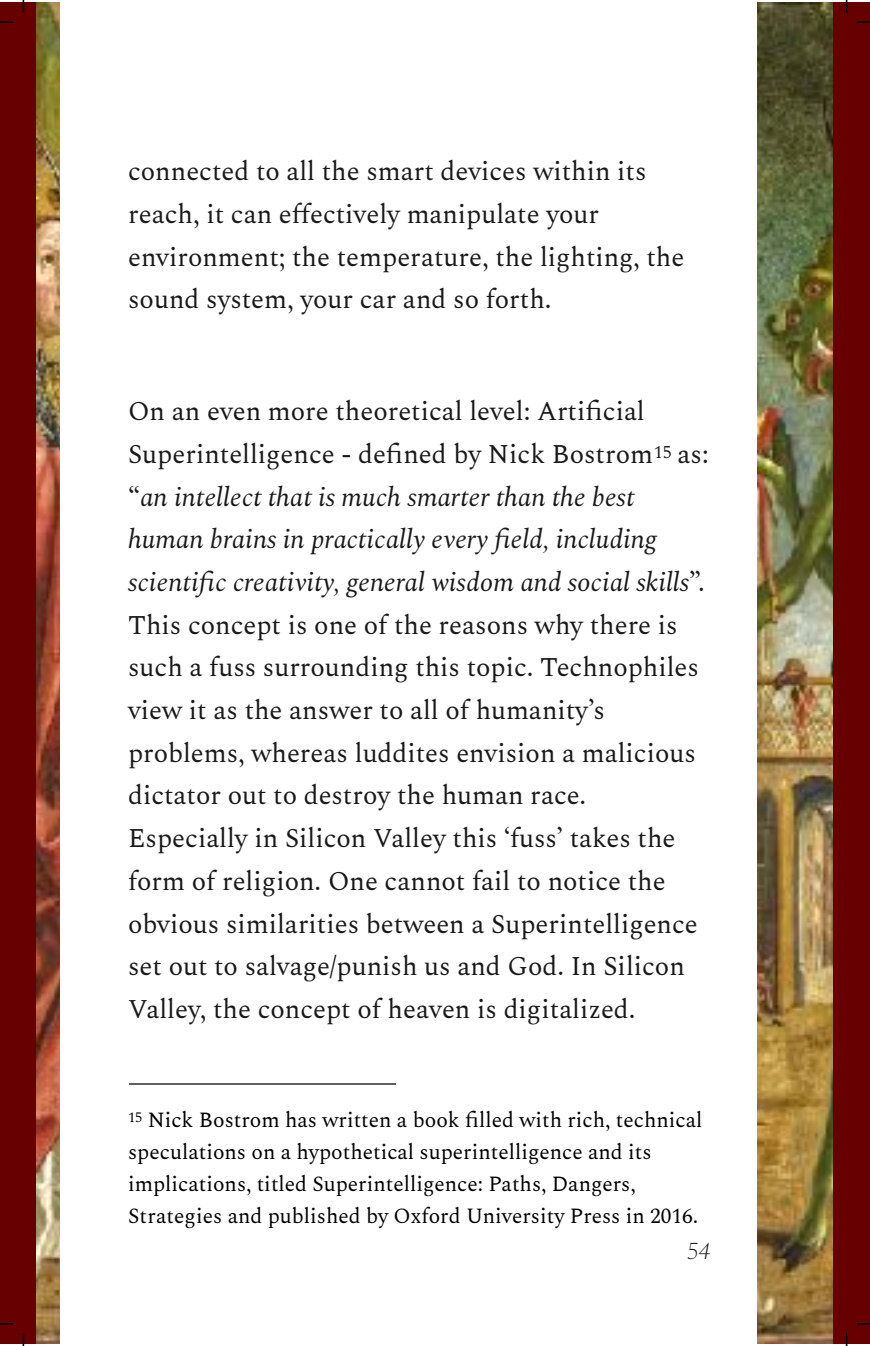
As James Bridle puts it in his - not so optimistically titled - book 'New Dark Age: Technology and the End of the Future'¹⁴:

"We face a world, not in the future but right now, where we do not understand our own creations. The result of such opacity is always and inevitably violence."

¹⁴ New Dark Age: Technology and the End of the Future is a riveting book written by James Bridle, published by Verso Books in 2019.



Within the intertwined community of techno-enthusiasts and doomsday-thinkers alike there is often talk of three separate types of Artificial Intelligence. The first type is everywhere around us already: Artificial Narrow Intelligence - AI specialized in only one area, such as the single tasked Deep Blue. It can beat the top human chess players, but it can't do anything else. Your phone is riddled with these systems. Combined they fulfil everything you need your phone to do, but individually they just perform the very specific task they are coded for. Going one step up the chain we find a concept that does not exist in the real world as of yet: Artificial General Intelligence - A.I. that is as smart as the average human in terms of intellectual capabilities. This type is depicted by Hal in Stanley Kubrick's movie '2001: A Space Odyssey' (1968). Just to clarify: we are not talking about robots made in men's likeness here. If your pipes are broken and you expect an Artificial General Intelligence to fix it, you will be disheartened by its lack of physical capabilities as a plumber. That is not to say that this Hal-like machine cannot affect your immediate surroundings. Assuming it will be




connected to all the smart devices within its reach, it can effectively manipulate your environment; the temperature, the lighting, the sound system, your car and so forth.

On an even more theoretical level: Artificial Superintelligence - defined by Nick Bostrom¹⁵ as: *“an intellect that is much smarter than the best human brains in practically every field, including scientific creativity, general wisdom and social skills”*.

This concept is one of the reasons why there is such a fuss surrounding this topic. Technophiles view it as the answer to all of humanity's problems, whereas luddites envision a malicious dictator out to destroy the human race.

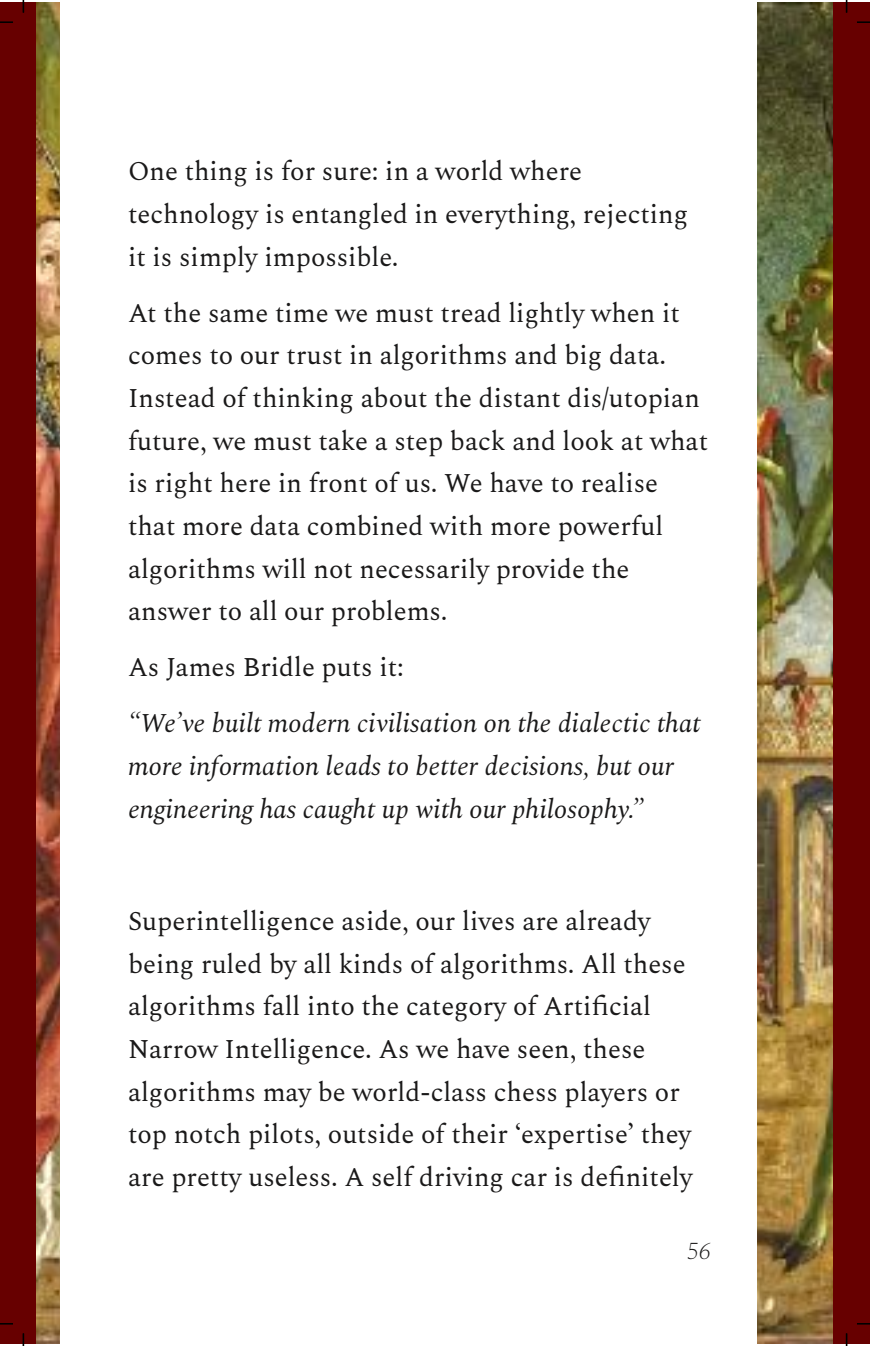
Especially in Silicon Valley this ‘fuss’ takes the form of religion. One cannot fail to notice the obvious similarities between a Superintelligence set out to salvage/punish us and God. In Silicon Valley, the concept of heaven is digitalized.

¹⁵ Nick Bostrom has written a book filled with rich, technical speculations on a hypothetical superintelligence and its implications, titled *Superintelligence: Paths, Dangers, Strategies* and published by Oxford University Press in 2016.



Transhumanists store their brains and bodies in liquid nitrogen in the hope that one day they will be revived, or that their brains can be uploaded to the cloud. Like so many religious before them, they embalm their bodies believing that doing so will give them a ticket to the afterlife.

The focus on this type of A.I. breeds a misconception about the current state and the future of our technology. Hollywood sci-fi movies fuel an irrational fear of A.I. as an inherently bad robot, while on the other side of the spectrum all hope for ridding the world of disease and war is projected on this God-like machine. As long as we keep accelerating the 'intelligence-explosion' this utopia will come true, the thought goes. The truth - as always - is a bit more nuanced. These convictions both take the state of mankind (the superior ruler of this planet) as a point of reference. But would a superior intelligence agree? Maybe it wouldn't care too much about its creator, maybe it would prefer other, more subtle and less invasive forms of life? To paraphrase Genesis, where mankind got kicked out of paradise: 'Computer says no'...




One thing is for sure: in a world where technology is entangled in everything, rejecting it is simply impossible.

At the same time we must tread lightly when it comes to our trust in algorithms and big data. Instead of thinking about the distant dis/utopian future, we must take a step back and look at what is right here in front of us. We have to realise that more data combined with more powerful algorithms will not necessarily provide the answer to all our problems.

As James Bridle puts it:

“We’ve built modern civilisation on the dialectic that more information leads to better decisions, but our engineering has caught up with our philosophy.”

Superintelligence aside, our lives are already being ruled by all kinds of algorithms. All these algorithms fall into the category of Artificial Narrow Intelligence. As we have seen, these algorithms may be world-class chess players or top notch pilots, outside of their ‘expertise’ they are pretty useless. A self driving car is definitely

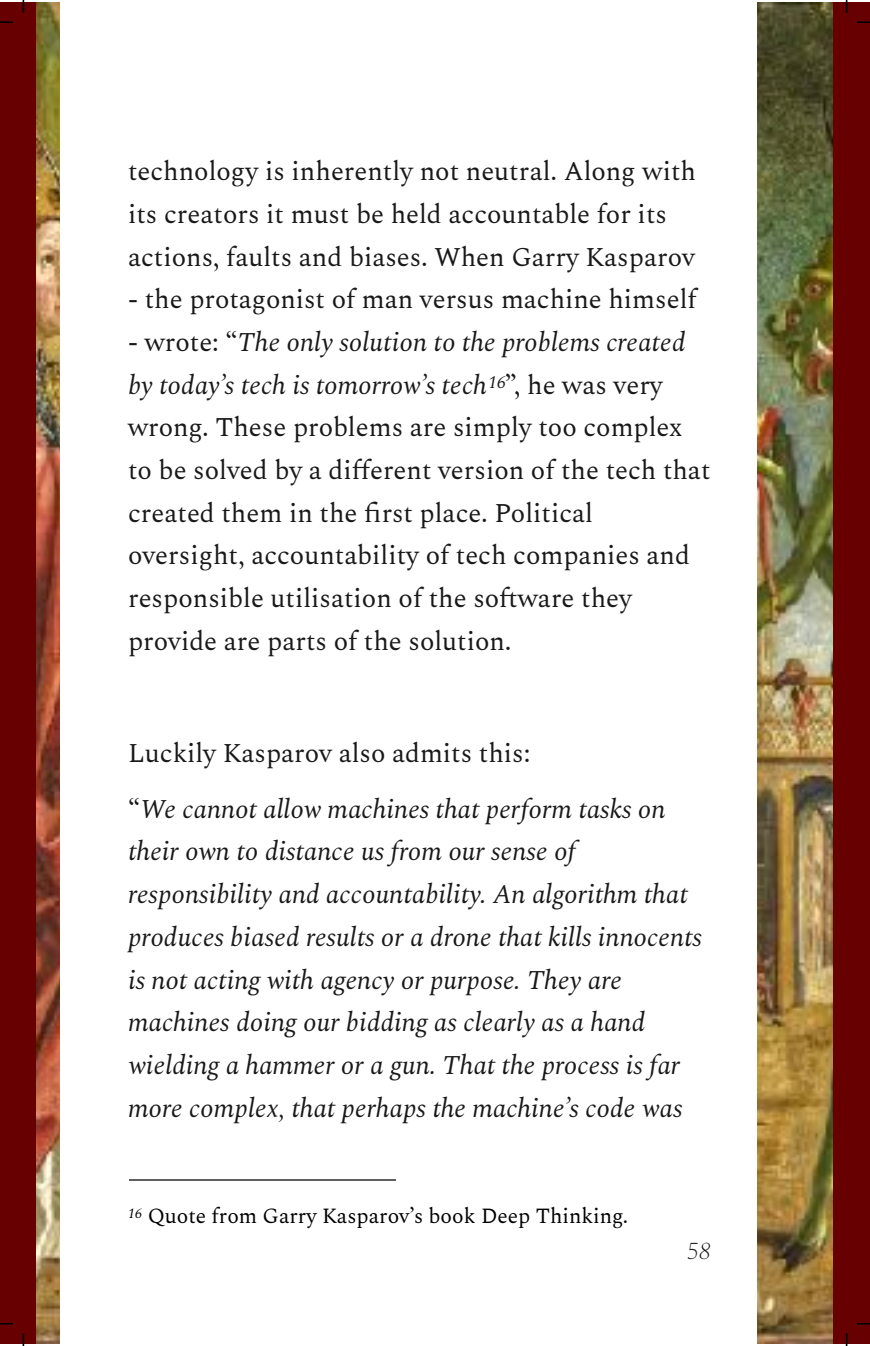


a lot safer than a human driver, but do not ask it to paint your walls or give you advice on your love life. Not yet at least.

But even if autonomous vehicles drive you around more safely and algorithms are ‘efficient’ at providing personal news feeds, there is one very important thing to keep in mind here: these technologies are not neutral. The belief that they are unbiased because they do not suffer from human fatigue, emotion or error is a false one. For starters: there is no such thing as a neutral definition of ‘efficient’... Am I going to quote James Bridle again? Why yes I am:

“Technology does not emerge from a vacuum. Rather, it is the reification of a particular set of beliefs and desires: the congruent, if unconscious dispositions of its creators. In any moment it is assembled from a toolkit of ideas and fantasies developed over generations, through evolution and culture, pedagogy and debate, endlessly entangled and enfolded.”

So what is the right attitude to have towards our own inventions - and their offspring? The first step would probably be acknowledging that




technology is inherently not neutral. Along with its creators it must be held accountable for its actions, faults and biases. When Garry Kasparov - the protagonist of man versus machine himself - wrote: *“The only solution to the problems created by today’s tech is tomorrow’s tech¹⁶”*, he was very wrong. These problems are simply too complex to be solved by a different version of the tech that created them in the first place. Political oversight, accountability of tech companies and responsible utilisation of the software they provide are parts of the solution.

Luckily Kasparov also admits this:

“We cannot allow machines that perform tasks on their own to distance us from our sense of responsibility and accountability. An algorithm that produces biased results or a drone that kills innocents is not acting with agency or purpose. They are machines doing our bidding as clearly as a hand wielding a hammer or a gun. That the process is far more complex, that perhaps the machine’s code was

¹⁶ Quote from Garry Kasparov’s book *Deep Thinking*.

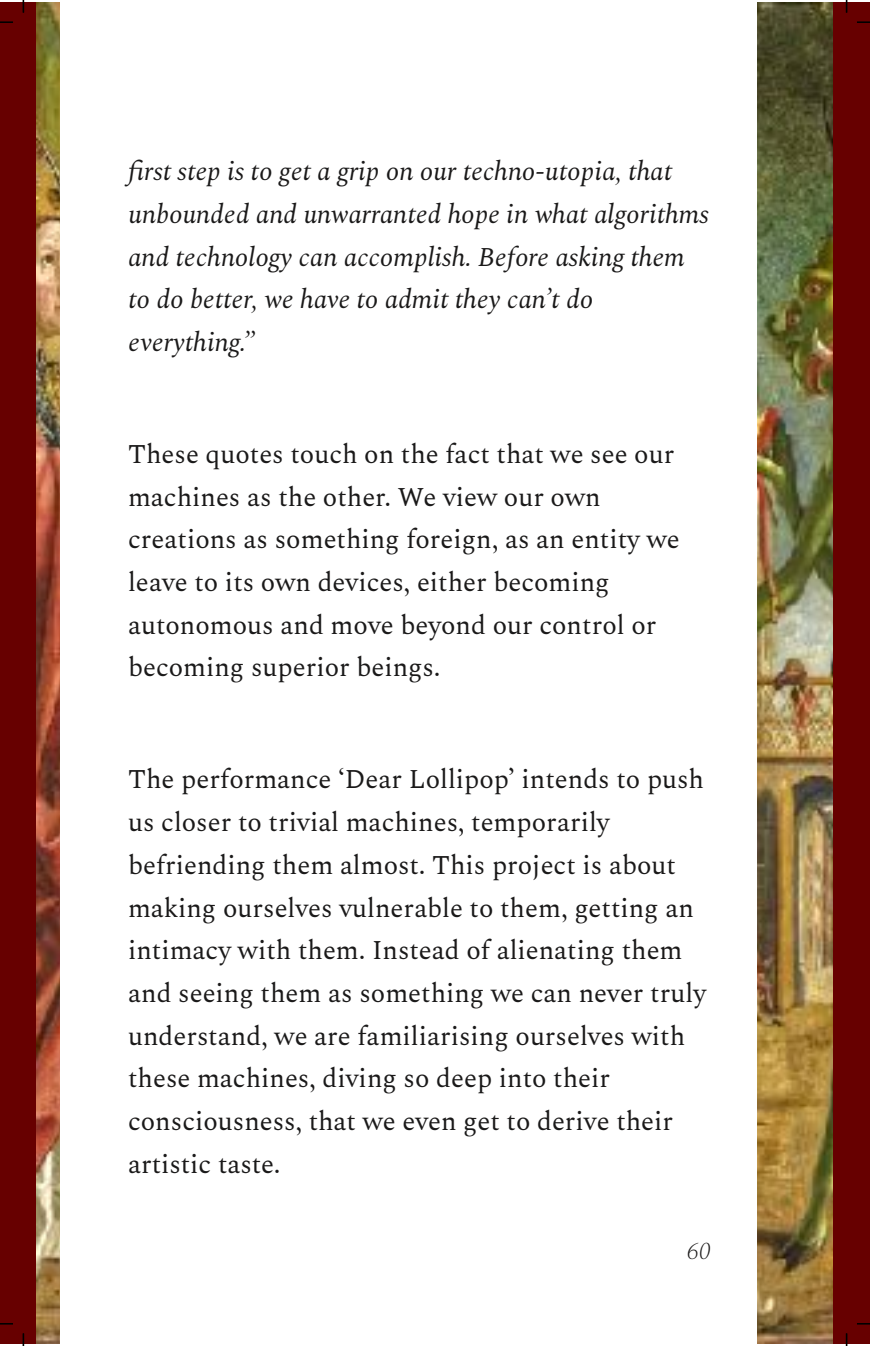


itself written by yet another algorithm, must not alter that moral calculus.”

(Notice the haphazard use of the word ‘innocent’: who is really innocent? To whom? As an example: is one innocent while being fortunate enough to be born on the ‘right’ side of the Mediterranean Sea, indulging all the economic benefits of this position? Or should we refer to this injustice by randomness as the ‘original sin’?)

This awareness is vital in changing the way we design and interact with our technology. We must not only optimise our machines, but also ourselves. Especially in times when people expect a machine to improve itself simply because we are throwing vast amounts of data to it. In her book *Weapons of Math Destruction* Cathy O’Neil writes:

“Big Data codify the past. They do not invent the future. Doing that requires moral imagination, and that’s something only humans can provide. [...] So the



first step is to get a grip on our techno-utopia, that unbounded and unwarranted hope in what algorithms and technology can accomplish. Before asking them to do better, we have to admit they can't do everything."

These quotes touch on the fact that we see our machines as the other. We view our own creations as something foreign, as an entity we leave to its own devices, either becoming autonomous and move beyond our control or becoming superior beings.

The performance 'Dear Lollipop' intends to push us closer to trivial machines, temporarily befriending them almost. This project is about making ourselves vulnerable to them, getting an intimacy with them. Instead of alienating them and seeing them as something we can never truly understand, we are familiarising ourselves with these machines, diving so deep into their consciousness, that we even get to derive their artistic taste.




Rise of the robots

And what about robots? They embody artificial intelligence; this strange, otherwise hard to grasp entity lives somewhere inside them. Some of them even look like us, and thus they appeal to our imagination.

It is estimated that in 2019 the global spending on robotics will hit \$135 billion¹⁷, with the industrial robot sector taking up the biggest chunk. Robots are widespread in factories around the world, where they are assembling cars and electronics. These robots do their jobs better, quicker and cheaper than their human colleagues ever would be able to. They possess these qualities mainly because they are not even trying to look like these human colleagues. They are a whole different species altogether.

Probably they will stay confined to factories for a while, since as of yet there are not a lot of practical appliances for the individual customer -


¹⁷ As estimated in a report from IDC: https://www.idc.com/getdoc.jsp?containerId=IDC_P33201



except in elderly homes perhaps, where robots somehow prove to be valuable companions.

In addition to not being very practical, human-like robots can still be very clumsy. It takes a lot to simply make a robot walk without falling over or bumping into walls. That is the Moravec paradox: it is comparatively easy to make computers exhibit adult level performance on intelligence tests or playing checkers, but very difficult to give them the skills of a one-year-old toddler when it comes to perception and mobility. Maneuvering around and over obstacles in a room is peanuts for this toddler, but can be very hard for a robot, and takes a lot of computational power. As Marvin Minsky stated: *“the most difficult human skills to reverse engineer are those that are unconscious¹⁸”*. Ironically, we never really think about these unconscious processes that we are very skilled at. We are ignorant of what our mind does best, and more aware of simple processes that do not work well

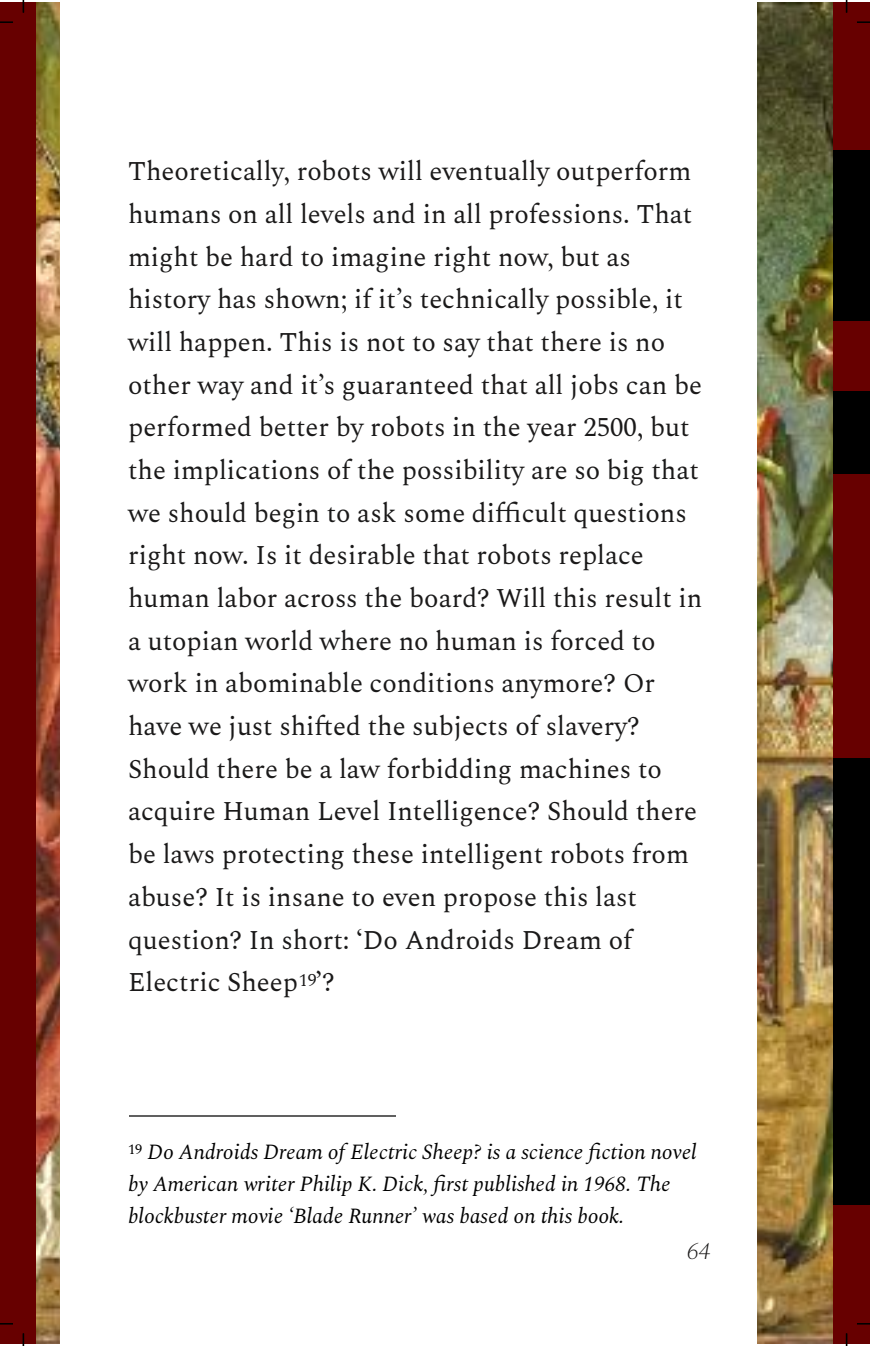
¹⁸ https://en.wikipedia.org/wiki/Moravec%27s_paradox



than simple ones that work flawlessly. This is why we consider ‘calculating’ as being a superior ability than ‘walking’, while the cheapest calculator beats the cleverest mind in second.

The processes that work best, work so well precisely because we are largely unaware of how they are happening. Conscious thought happens in the area of the frontal lobe, the part of the brain that evolved most recently. But most of what we do during the day, walking, coordinating, speaking, processing and reacting to all kinds of outside information is done by the most ancient part of the brain, the part that has been improved over millions of years of evolution. In that respect it is not so hard to see that it is incredibly difficult to reverse engineer the human (or animal) body.

But improvements are rapidly being made even in these areas of physical capabilities. It is likely that this trend will continue. Previously inconceivable tasks (such as recognising faces) are being developed at an ever increasing rate.



Theoretically, robots will eventually outperform humans on all levels and in all professions. That might be hard to imagine right now, but as history has shown; if it's technically possible, it will happen. This is not to say that there is no other way and it's guaranteed that all jobs can be performed better by robots in the year 2500, but the implications of the possibility are so big that we should begin to ask some difficult questions right now. Is it desirable that robots replace human labor across the board? Will this result in a utopian world where no human is forced to work in abominable conditions anymore? Or have we just shifted the subjects of slavery? Should there be a law forbidding machines to acquire Human Level Intelligence? Should there be laws protecting these intelligent robots from abuse? It is insane to even propose this last question? In short: 'Do Androids Dream of Electric Sheep¹⁹'?

¹⁹ *Do Androids Dream of Electric Sheep?* is a science fiction novel by American writer Philip K. Dick, first published in 1968. The blockbuster movie 'Blade Runner' was based on this book.

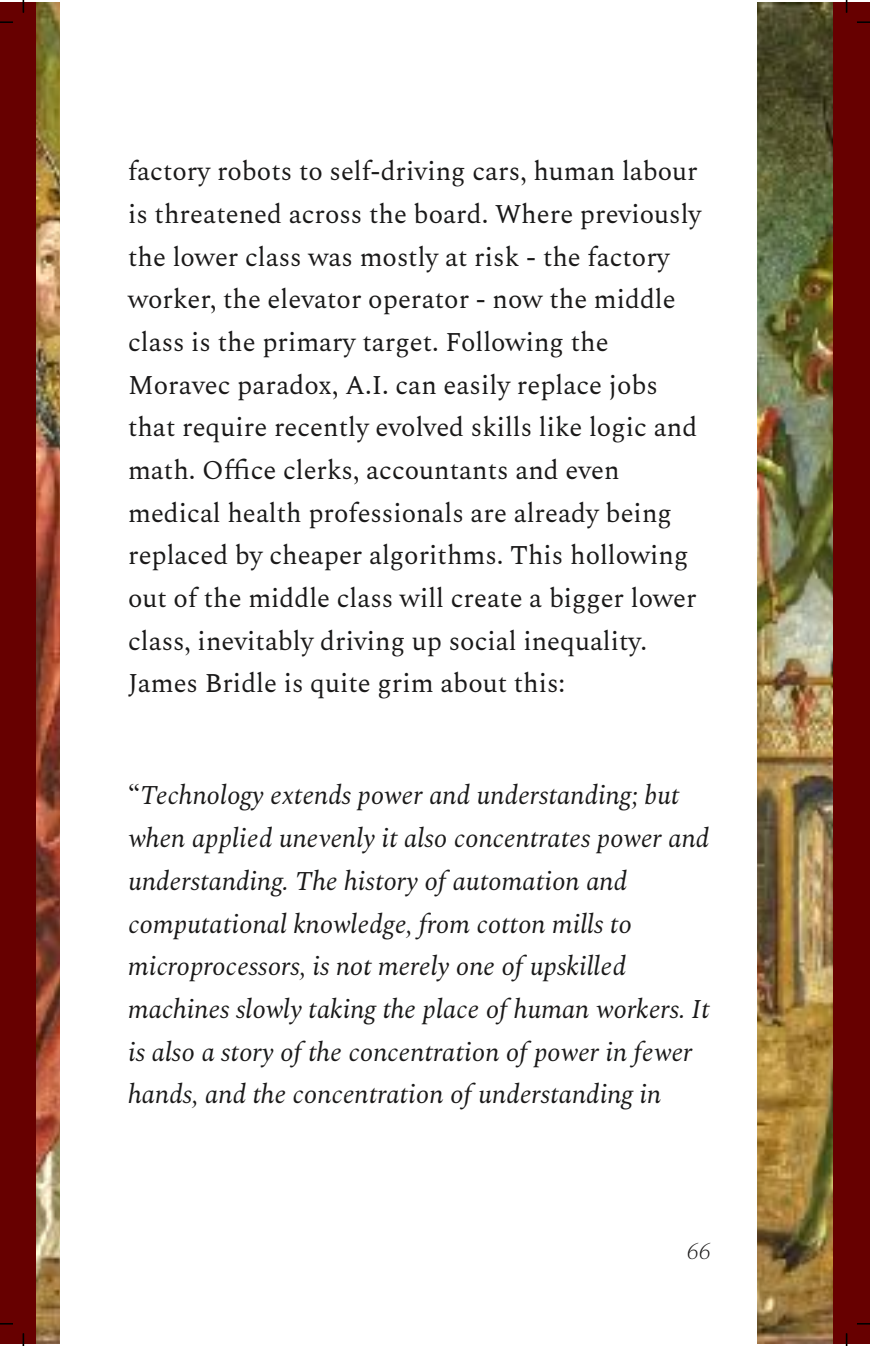


Automation

Fear of automation has been around since the industrial revolution, but as factory workers lost their jobs, more and more jobs were created by the accumulation of wealth generated by the revolution. But now things are different, technological progress is accelerated, and a broad spectrum of jobs are at risk. This time around the fear may well be justified.


The sentiment itself is the same as 200 years ago: it was feared automation would render people completely useless once their job could be done entirely by a machine.

While the elevator operator did indeed lose his job, different professions appeared and the fears of the luddites appeared unjustified. The developments in the 21st century, however, give reason for concern. This time around automation might have more drastic effects. Self- and deep learning make the pace at which technology evolves a lot higher. It seems that more and more jobs can be done a lot cheaper and more efficiently by getting rid of human labor. From supermarket checkouts to trading algorithms,



factory robots to self-driving cars, human labour is threatened across the board. Where previously the lower class was mostly at risk - the factory worker, the elevator operator - now the middle class is the primary target. Following the Moravec paradox, A.I. can easily replace jobs that require recently evolved skills like logic and math. Office clerks, accountants and even medical health professionals are already being replaced by cheaper algorithms. This hollowing out of the middle class will create a bigger lower class, inevitably driving up social inequality. James Bridle is quite grim about this:



“Technology extends power and understanding; but when applied unevenly it also concentrates power and understanding. The history of automation and computational knowledge, from cotton mills to microprocessors, is not merely one of upskilled machines slowly taking the place of human workers. It is also a story of the concentration of power in fewer hands, and the concentration of understanding in



fewer heads. The price of this wider loss of power and understanding is, ultimately, death.²⁰

After making claims about our inability to predict the future, to avoid embarrassment, I absolutely want to refrain from doing so anyway. But instead of playing Nostradamus we can look at the clear signs of the winner-takes-all economics of today, and draw our conclusions. The job prospects for accountants, truck drivers and cashiers are not all too great. New, tech-related positions will obviously open up - someone has to write the algorithms that displace these workers - but not for everyone. Is learning to code C++ a reasonable thing to expect from a 50-year old cashier that just lost his job? Probably not. However, it is reasonable to expect more and more low-skill jobs to be automated in the future, possibly sooner than predicted. Even formerly 'high qualified' and well paid jobs, such as accountants, bank managers... are up to scraps as we speak. Even quite highly educated workers


²⁰ Quote from James Bridle's book *New Dark Age*.



will be unlikely to shapeshift into forms required by the new jobs stemming from the A.I. revolution. Exactly because of this reason it is vital to engrain flexibility into the education system. The big question is how to educate the youth of tomorrow to be as flexible as possible, because they will probably need to take on very different skill sets when required. Or should we just educate them for a happy life without a job? Shouldn't we concentrate on eradicating social inequalities based on an obsolete hierarchy of human skills, and prepare for a jobless paradise? And is society moving in the right direction for this paradigm shift?

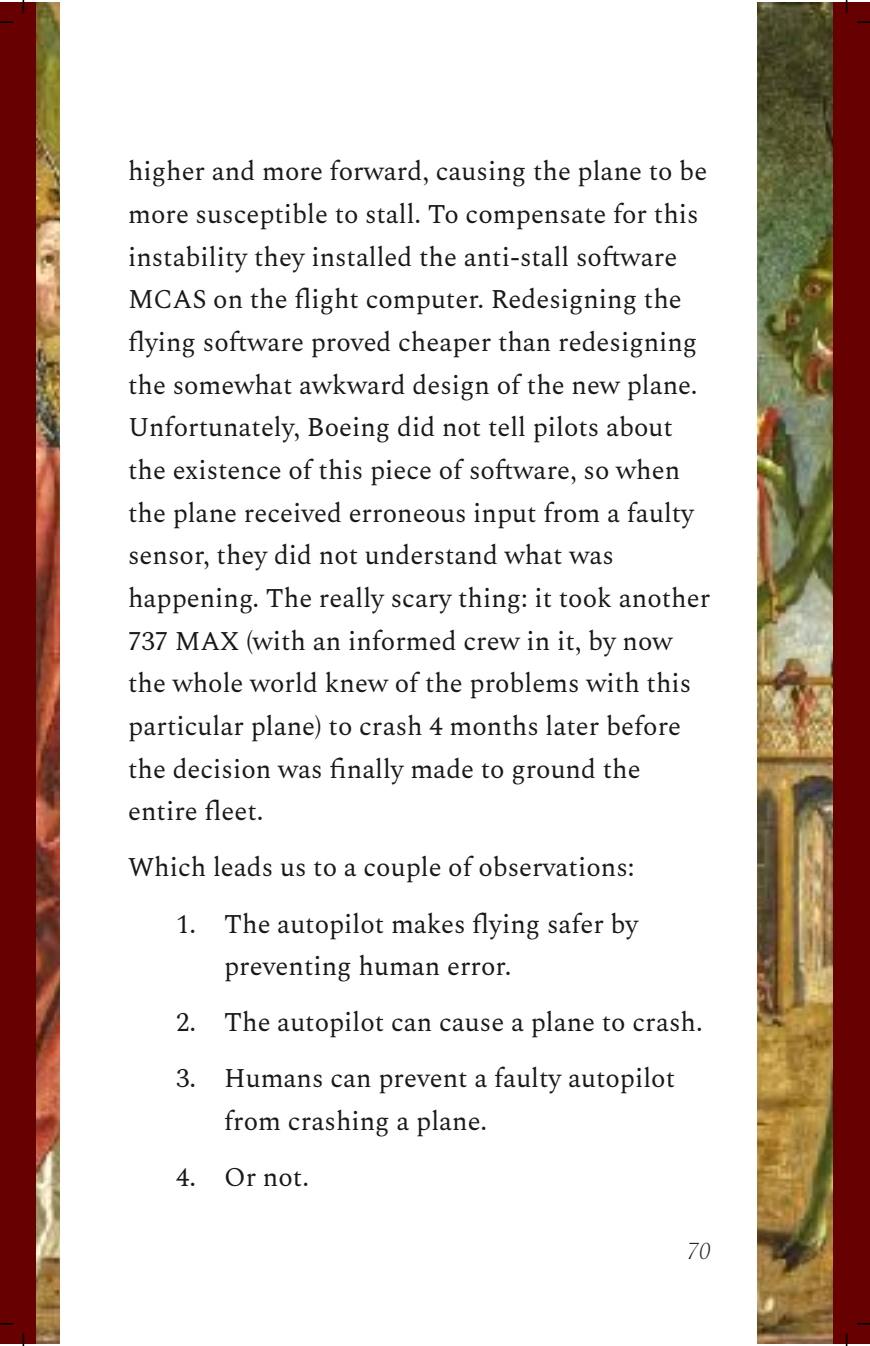
Men and machine

Sometimes technology is complementary to human labor, providing an extra layer of safety and efficiency. Flying, for example, has gotten a lot safer since pilots are not really flying planes anymore. This does not mean that it is better to relentlessly forward this automation by removing pilots from the cockpit. They provide an additional security layer to the autopilot, and



when the system malfunctions, they can intervene - provided they understand the system. When men and machine do not understand each other, the results can be catastrophic.


On the 29th of October in 2018, Lion Air flight 610 departed from Jakarta to Pangkal Pinang, Indonesia. Shortly after takeoff, the pilots contacted air traffic control for help. Something alarming had happened: the aircraft's computer system kept pushing its nose down. Frantically, the pilots tried to figure out the cause of the error while pushing the plane back up again. This created a tug of war between the plane and the pilots - nose down, nose up - that seesawed the aircraft more than two dozen times. After 12 minutes, the brand new 737 MAX crashed into the Java sea. Investigation into the cause of the crash found that the cause of the nosedive was a new type of software: MCAS. The original Boeing 737 used too much fuel, so they decided to install more efficient engines with bigger fans and make the 737 MAX. Because of economical and practical reasons, they mounted the engines



higher and more forward, causing the plane to be more susceptible to stall. To compensate for this instability they installed the anti-stall software MCAS on the flight computer. Redesigning the flying software proved cheaper than redesigning the somewhat awkward design of the new plane. Unfortunately, Boeing did not tell pilots about the existence of this piece of software, so when the plane received erroneous input from a faulty sensor, they did not understand what was happening. The really scary thing: it took another 737 MAX (with an informed crew in it, by now the whole world knew of the problems with this particular plane) to crash 4 months later before the decision was finally made to ground the entire fleet.



Which leads us to a couple of observations:

1. The autopilot makes flying safer by preventing human error.
2. The autopilot can cause a plane to crash.
3. Humans can prevent a faulty autopilot from crashing a plane.
4. Or not.



All these observations can be true simultaneously. Instead of blindly relying on machines or outright rejecting them, the key - once again - might lie in cooperation (or symbiosis) between men and machine. We have to understand them, and they have to understand us, and only then we will be able to make true progress. Instead of making machines more opaque, we should make them more transparent. It is impossible to tell, but one can imagine that had the pilots of the fatal Lion Air flight completely understood the MCAS system, they might have been able to prevent the plane from crashing. Or not - if the software had been written to overrule any human intervention, maybe out of a deeply rooted mistrust in human pilots. This mistrust could be justified by the statistics of air crashes: human error is by far the number one cause of air crashes. In the unearthly skill of flying, we are in fact the weakest link...²¹

²¹ www.1001crash.com



James Bridle, for a change, concludes on a more optimistic note:

“Cooperation also reduces the sting of computational analysis, we might gain a deeper insight into the way in which complex machines make their decisions. Acknowledging the reality of nonhuman intelligence has deep implications for how we act in the world and requires clear thinking about our own behaviours, opportunities, and limitations.”



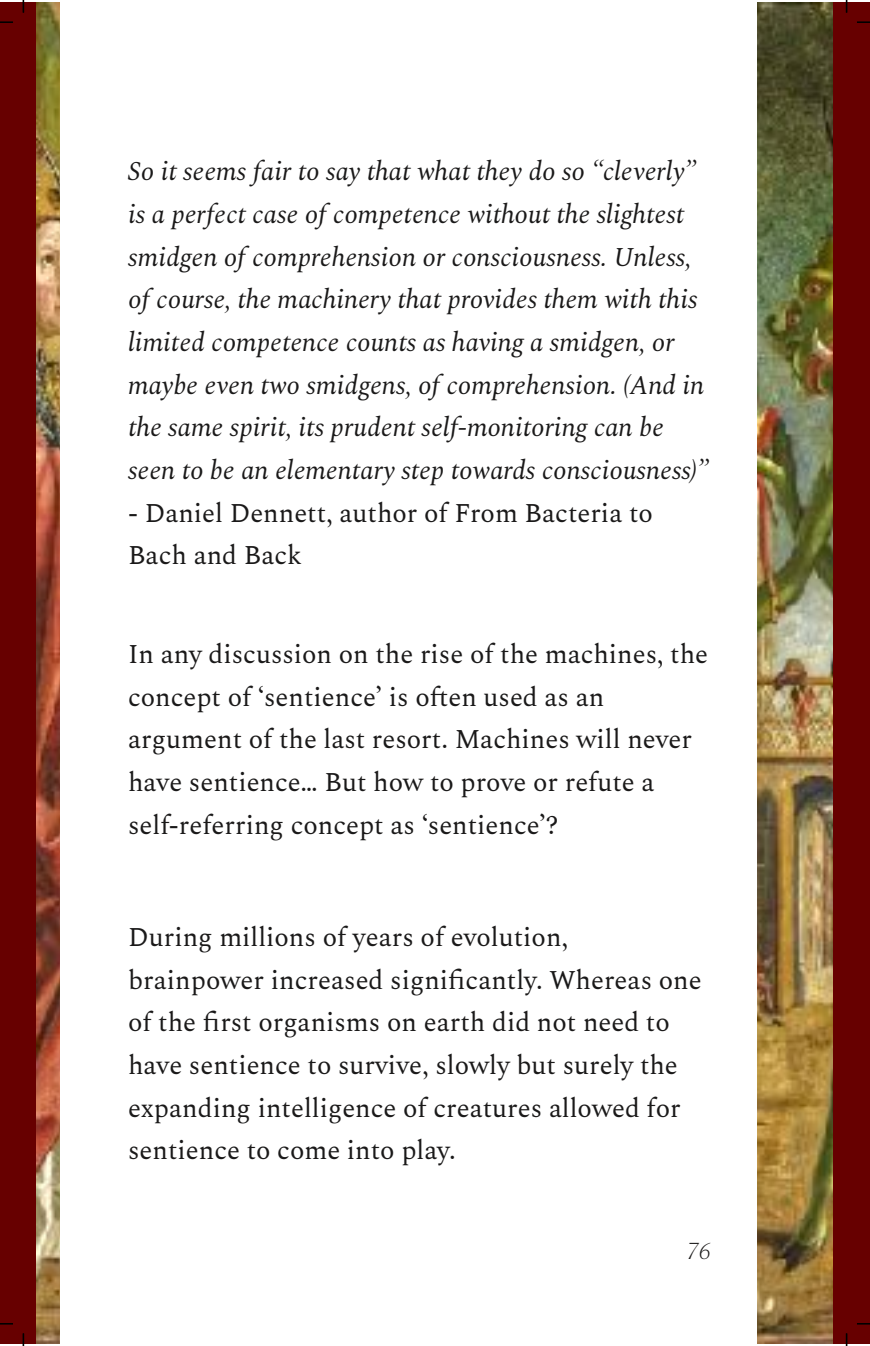




On consciousness and sentience

The origin of consciousness


“Elevators can do remarkably clever things, optimizing their trajectories, thereby saving time and energy, automatically adjusting their velocity to minimize discomfort of their passengers, “thinking of everything” that needs to be thought about, and obeying instructions and even answering frequently asked questions. Good elevators earn their keep. They do this without any neurons, sense organs, dopamine, glutamate, or the other organic components of brains.



So it seems fair to say that what they do so “cleverly” is a perfect case of competence without the slightest smidgen of comprehension or consciousness. Unless, of course, the machinery that provides them with this limited competence counts as having a smidgen, or maybe even two smidgens, of comprehension. (And in the same spirit, its prudent self-monitoring can be seen to be an elementary step towards consciousness)”
- Daniel Dennett, author of *From Bacteria to Bach and Back*

In any discussion on the rise of the machines, the concept of ‘sentience’ is often used as an argument of the last resort. Machines will never have sentience... But how to prove or refute a self-referring concept as ‘sentience’?

During millions of years of evolution, brainpower increased significantly. Whereas one of the first organisms on earth did not need to have sentience to survive, slowly but surely the expanding intelligence of creatures allowed for sentience to come into play.



The heavily researched YouTube channel Kurzgesagt talks about this in one of their videos:

“Trichopax adhaerens - one of the simplest of all animals moves around haphazardly. It slows down in the presence of food, and speeds up in its absence. This is highly effective, and makes the tiny creature spend more time where there is food than where there is not. But it never moves in a particular direction towards a particular target, and there's no need for it to be conscious of its environment.”

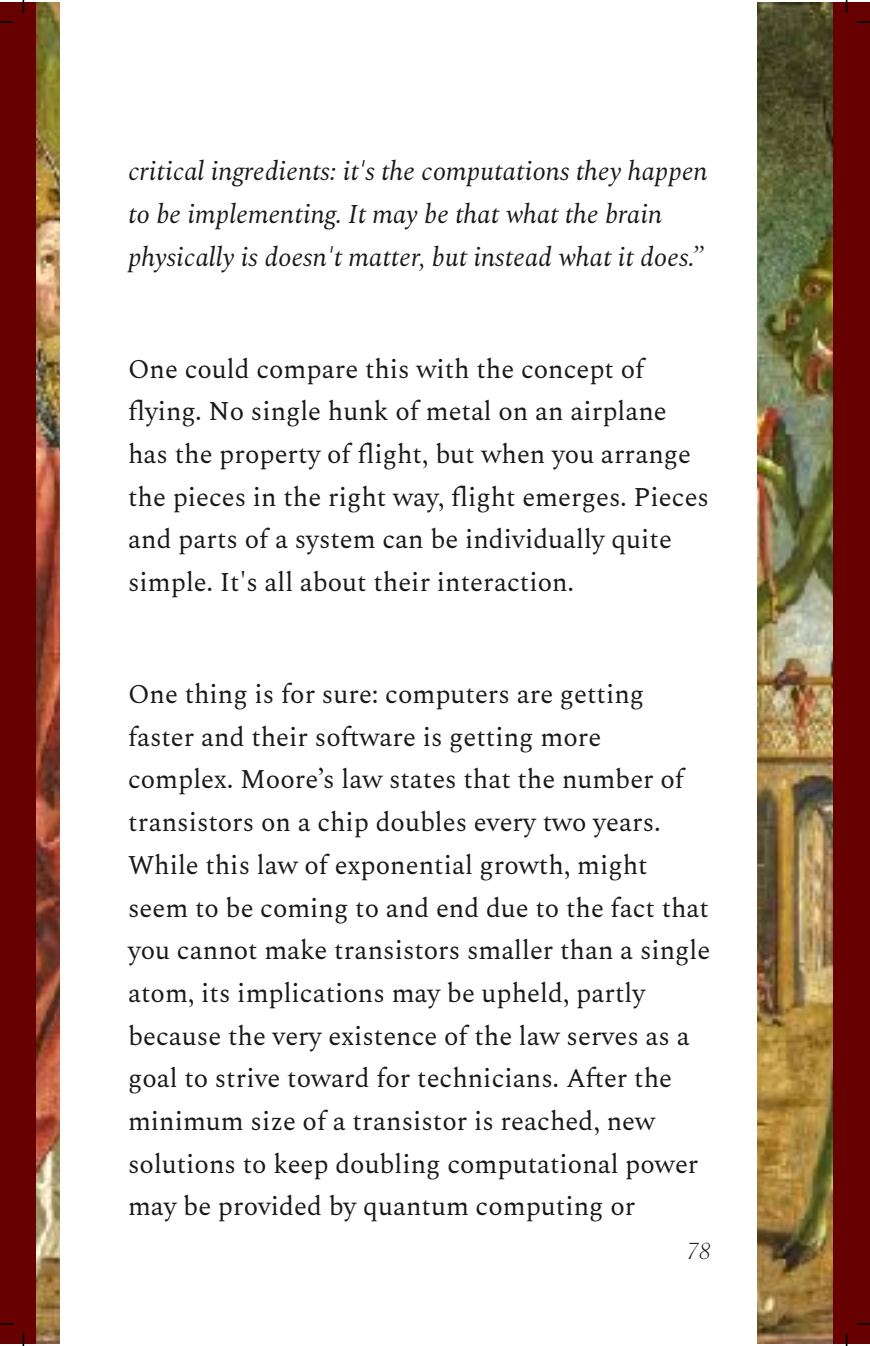
Sentience is too often looked upon as a flip flop characteristic: you either have it, or you don't. This misconception tends to be at the root of a whole lot of misunderstandings.

Instead, consciousness could be considered as an emergent property of any physical matter.

As laid out in the book ‘The Brain: The Story of You’ by David Eagleman²²:

“That prospect is known as the computational hypothesis of the brain. The idea is that the neurons and synapses and other biological matter aren't the


²² The Brain: The Story of You is a book written by David Eagleman, published by Canongate Books in 2016.



critical ingredients: it's the computations they happen to be implementing. It may be that what the brain physically is doesn't matter, but instead what it does."

One could compare this with the concept of flying. No single hunk of metal on an airplane has the property of flight, but when you arrange the pieces in the right way, flight emerges. Pieces and parts of a system can be individually quite simple. It's all about their interaction.

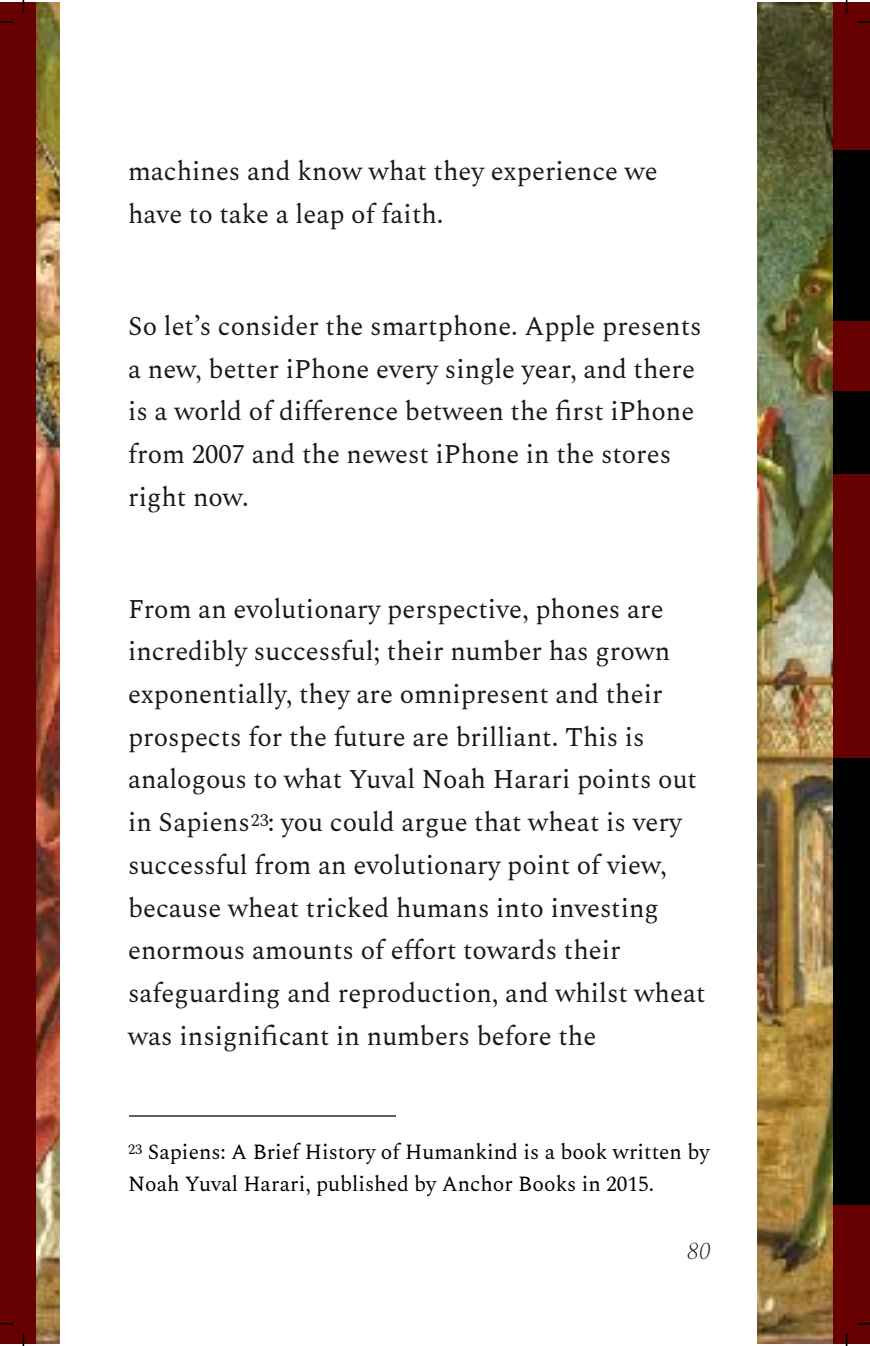
One thing is for sure: computers are getting faster and their software is getting more complex. Moore's law states that the number of transistors on a chip doubles every two years. While this law of exponential growth, might seem to be coming to an end due to the fact that you cannot make transistors smaller than a single atom, its implications may be upheld, partly because the very existence of the law serves as a goal to strive toward for technicians. After the minimum size of a transistor is reached, new solutions to keep doubling computational power may be provided by quantum computing or



smarter and more efficient software. When looking back at the history of computational power, it was a surprise to see that Moore's Law was applicable across some technological shifts - from mechanical computers, through vacuum tube driven computers and so on.

As computational power increases exponentially, we can safely state that A.I. is equally getting more intelligent, powerful and complex.

So if we extend this line of thought to our intelligent machines, the consequence of this ever more complex intelligence could be the emergence of some form of sentience. This process might happen in the same way that the evolutionary line of organisms' sentience came into existence, parallel to those organisms getting more and more complex. With technology this could happen in a similar way. It might seem nonsensical to think about artificial intelligence and consciousness like this, but so was thinking about animals in relation to consciousness at one point in history. If we genuinely want to familiarise ourselves with




machines and know what they experience we have to take a leap of faith.

So let's consider the smartphone. Apple presents a new, better iPhone every single year, and there is a world of difference between the first iPhone from 2007 and the newest iPhone in the stores right now.

From an evolutionary perspective, phones are incredibly successful; their number has grown exponentially, they are omnipresent and their prospects for the future are brilliant. This is analogous to what Yuval Noah Harari points out in *Sapiens*²³: you could argue that wheat is very successful from an evolutionary point of view, because wheat tricked humans into investing enormous amounts of effort towards their safeguarding and reproduction, and whilst wheat was insignificant in numbers before the

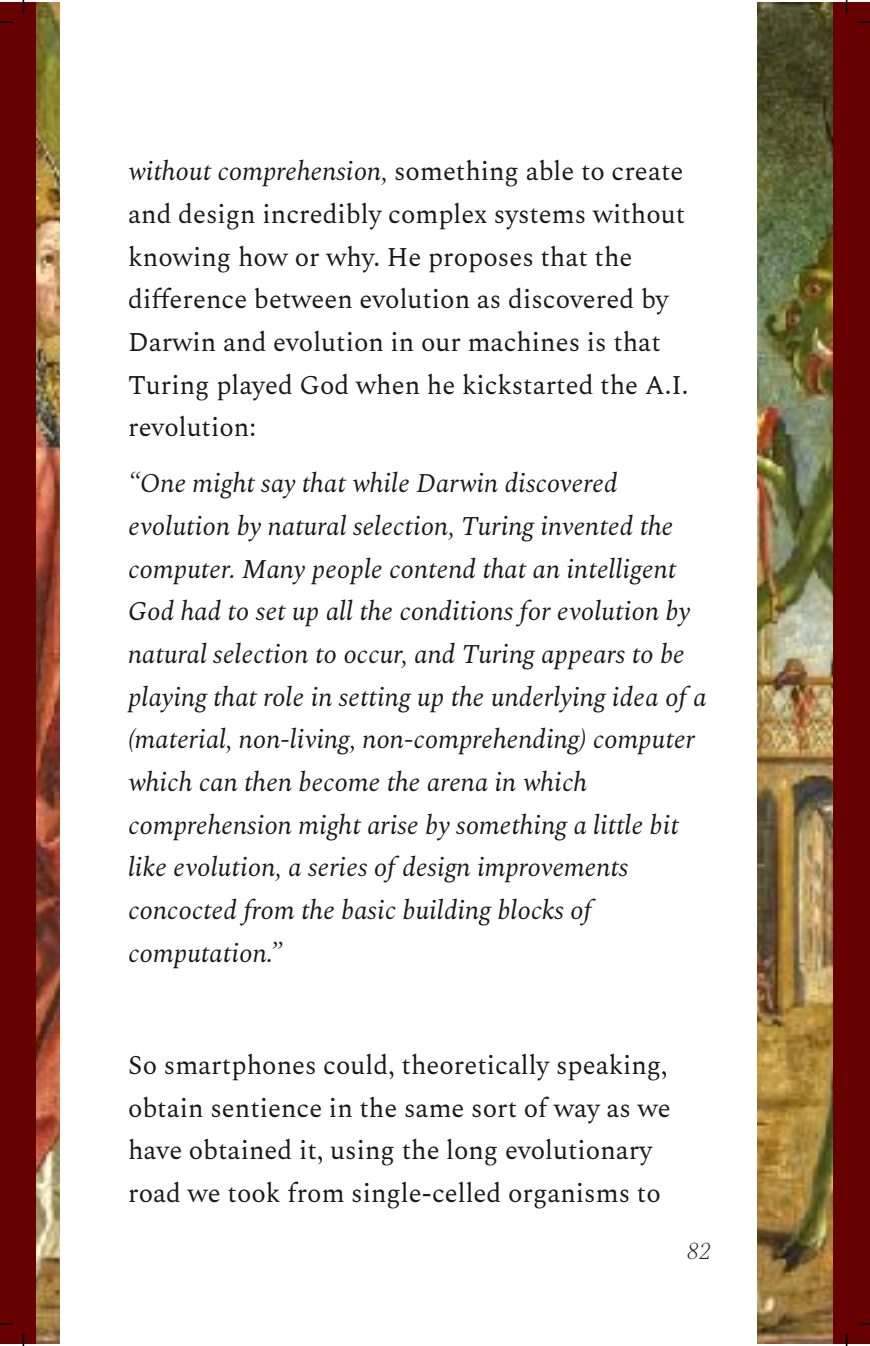
²³ *Sapiens: A Brief History of Humankind* is a book written by Noah Yuval Harari, published by Anchor Books in 2015.



agricultural revolution, today it grows all over the world in large numbers. In a very similar way phones are forcing a huge number of people to improve them every year, and an even larger sum of money is being poured into their development. Perhaps they aren't doing this purposely or consciously, but neither did wheat. So one could also argue that from the perspective of the phone, because of their necessity and indispensability, they are *making us make them sentient*. That does not mean that all phones regularly meet up to plan how to take over the world (although who knows what your phone is doing when you are not looking?). But again, neither did wheat... or Neanderthals, for that matter. And when exactly was the last time that you had a meeting on bettering the future of your existence?

Machines do not need intent to evolve. In his book *From Bacteria to Bach and Back*²⁴, Daniel Dennett talks about evolution as *competence*


²⁴ *From Bacteria to Bach and Back: The Evolution of Minds* is a book written by Daniel Dennett, published by Penguin Books in 2018.



without comprehension, something able to create and design incredibly complex systems without knowing how or why. He proposes that the difference between evolution as discovered by Darwin and evolution in our machines is that Turing played God when he kickstarted the A.I. revolution:

“One might say that while Darwin discovered evolution by natural selection, Turing invented the computer. Many people contend that an intelligent God had to set up all the conditions for evolution by natural selection to occur, and Turing appears to be playing that role in setting up the underlying idea of a (material, non-living, non-comprehending) computer which can then become the arena in which comprehension might arise by something a little bit like evolution, a series of design improvements concocted from the basic building blocks of computation.”



So smartphones could, theoretically speaking, obtain sentience in the same sort of way as we have obtained it, using the long evolutionary road we took from single-celled organisms to



homo sapiens, only in our case there were no people around to help, only laws of nature and a lot of luck (and probably we piggybacked on some other animals as well). We did not consciously *make* the laws of nature give us sentience but it just happened *because* of them. The laws of nature will not help smartphones achieve sentience - Moore's Law will give it to them instead. In a way, the technological revolution is a meta-consequence of nature itself; nothing is outside of nature's laws. Not even us. But technological evolution doesn't seem to be content with the excruciatingly slow pace of its Darwinian sibling: it is spiralling out of control at an exponentially increasing speed.

How will they experience the world?

Will their awareness be similar to ours? Will they be able to feel? If they get so intelligent, so complex that we are not able to grasp their inner workings anymore, it is kind of irrelevant whether they can 'feel' or not. Their consciousness might take on a different form



than any we could know. Dennett envisions consciousness as something that comes in degrees, with bacteria's on the low end of this scale, and humans on the upper end. Whilst this concept deconstructs consciousness as a binary of either conscious or not, it creates a different problem: that of having lower and higher levels of consciousness. It seems more likely that these variations are a case of *difference* rather than degrees. Bats see the world through echolocation. That they do not have eyes does not mean they perceive the world to a lesser degree than us, it's just different. So I propose the following: when it comes to consciousness, we should think horizontally instead of vertically. Along the horizontal line, all forms of consciousness are equal.

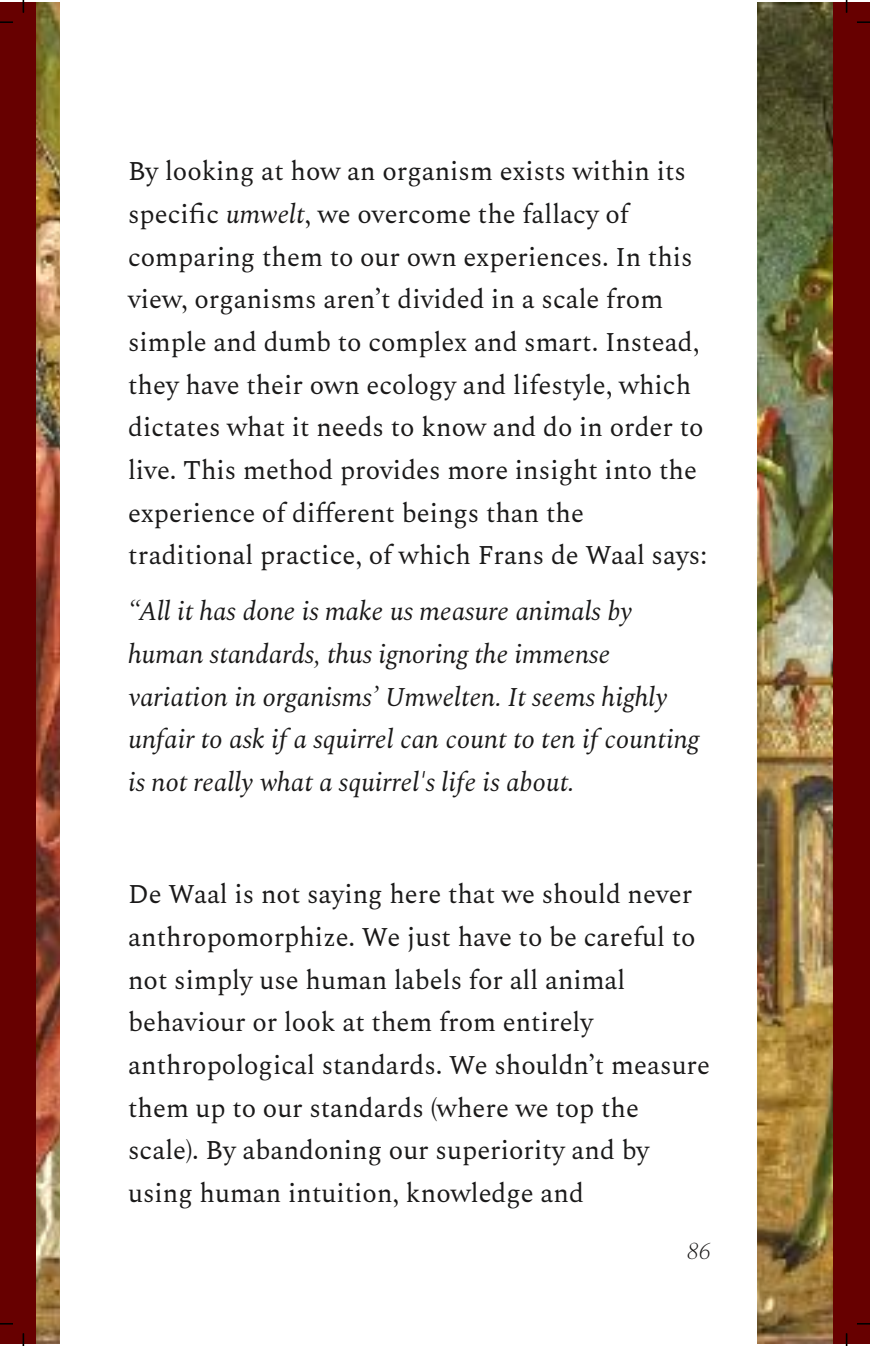


The Umwelt

Before zooming ahead on the possibilities of a subjective experience within machines, let's first look at what is known about this topic in relation to animals. In his book "Are We Smart Enough to Know How Smart Animals Are?"²⁵ Frans de Waal provides a very useful term we can use to draw ourselves a little bit closer to the inconceivable experiences of animals:

"Jakob von Uexkull, a German biologist, drew attention to the animal point of view, calling it its Umwelt. [...] a single environment offers hundreds of realities peculiar to each species. Umwelt stresses an organism's self-centered, subjective world, which represents only a small tranche of available worlds. According to Uexkull, the various tranches are "not comprehended and never discernible" to all the species that construct them. Some animals perceive ultraviolet light, for example, while others live in a world of smells or, like the star-nosed mole, feel their way around underground."


²⁵ Are We Smart Enough to Know How Smart Animals Are? is a book written by primatologist Frans de Waal, published by Norton in 2017.



By looking at how an organism exists within its specific *umwelt*, we overcome the fallacy of comparing them to our own experiences. In this view, organisms aren't divided in a scale from simple and dumb to complex and smart. Instead, they have their own ecology and lifestyle, which dictates what it needs to know and do in order to live. This method provides more insight into the experience of different beings than the traditional practice, of which Frans de Waal says:

“All it has done is make us measure animals by human standards, thus ignoring the immense variation in organisms’ Umwelten. It seems highly unfair to ask if a squirrel can count to ten if counting is not really what a squirrel's life is about.

De Waal is not saying here that we should never anthropomorphize. We just have to be careful to not simply use human labels for all animal behaviour or look at them from entirely anthropological standards. We shouldn't measure them up to our standards (where we top the scale). By abandoning our superiority and by using human intuition, knowledge and




imagination, we can make an attempt at understanding different *umwelten*. We must not do this by projecting a diluted version of our own experience onto animals because, as Peter Godfrey-Smith puts it in his book *Other Minds*²⁶: *“This is a mistake because our own experience relies on features they just don’t possess.”*

It is fascinating to see this mistake being made over and over again by some of the greatest minds. You would assume that by now there is enough literature on studies of animal behaviour to not fall into this trap of comparing them to our own standards. But apparently it is somehow hardwired into our brains, because even Daniel Dennett - after introducing the concept *umwelt* - makes this mistake when talking about chimps:


“The contrast between us and chimpanzees here is striking: by now, thousands of chimpanzees have spent their entire lives in human captivity, and such chimps have heard almost as many words as human children

²⁶ *Other Minds: The Octopus and the Evolution of Intelligent Life* is a book about the wondrous world of Octopus minds, written by Peter-Godfrey Smith, published by HarperCollins Publishers in 2018.



hear, but they seldom pay any attention. Human speech is to them pretty much like the rustling of leaves in the trees, even though speech contains vast amounts of semantic information that could be of use to them, if only they tumbled to it. [...] It takes a prodigious regime to get a chimpanzee to acquire the habit of attending to words, spoken or signed or tokened in plastic shapes. Human infants, in contrast, are hungry for verbal experience from birth.”

After reading this paragraph, we now know that chimps are very bad at learning human language, but you probably already knew that. In contrast, even humans - with the incredibly unfair advantage of not being locked up in a cage - have not completely figured out ape language, after decades of research. There is some new information we get from Dennett: chimpanzees are apparently able to learn some of our language skills through training. He presents this as if it is proof that they are somehow less cognitive than we are, while actually it is a pretty impressive fact. Few researchers are currently speaking (or even understanding) ‘ChimpLanguage’...




Let's indeed take a look at how ridiculous this way of reasoning is by reversing it. Suppose seals where studying the behaviour of people. They might take note that: *"It takes a prodigious regime to get a human to acquire the habit of swimming, with floaties or a life jacket or fins. Baby seals, in contrast, are hungry to swim from birth."* Bad seal scientist. The only thing you have now figured out is that swimming is not natural for people, that it is not their first priority.

This way of thinking about animals obviously does not help us get any further, or gain any familiarity with their worlds and ways of being, just like taking the intellectual high ground on our machines also does not get us a tad further.

We should not look down on animals not understanding the ways of our world, on the contrary, considering how unnatural life in captivity is for animals, we should be impressed by their creativity when they do seem to understand and even manipulate it.

Some revealing examples from Frans de Waal and Peter Godfrey-Smith:

“He’d let a female chimp named Panzee watch while he hid food in the pine forest around her outdoor enclosure. Charlie would dig a small hole in the ground to put a bag of M&Ms into it, or place a candy bar in the bushes. Panzee would follow the process from behind bars. Since she could not go where Charlie was, she would need human help to eventually get the hidden food. Sometimes Charlie would hide it after all other people had gone for the day. This meant Panzee could not communicate with anybody about what she knew until the next morning. When the caretakers arrived, they were unaware of the experiment. Panzee first had to get their attention, then provide information to someone who had no clue as to what she was “talking” about. [...] All those recruited by Panzee said they were at first surprised by her behavior but soon understood what she wanted them to do. By following her pointing, panting, and calling, they had no trouble finding the candies hidden in the forest. [...] The result was communication about a past event, present in the



ape's memory, to ignorant members of a different species.²⁷

and

“For example, it has long appeared that captive octopuses can recognize and behave differently toward individual human keepers. [...] In a lab in New Zealand, an octopus took a dislike to one member of the lab staff, for no obvious reason, and whenever that person passed by on the walkway behind the tank she received a jet of half a gallon of water in the back of her neck.²⁸”

For the average person it can be quite hard to tell different octopuses apart, so especially this last anecdote deserves appreciation.

²⁷ Quote from *Are We Smart Enough to Know How Smart Animals Are?* by Frans de Waal.

²⁸ Quote from *Other Minds: The Octopus and the Evolution of Intelligent Life* by Peter-Godfrey Smith.


The difference between consciousness and subjective experience

So now we have the concept of *umwelt* to help us understand animal behaviour. How does this help us in determining what animal consciousness looks like?

The first step might be acknowledging that some animal experiences are entirely unknowable because they are so alien to us and so unlike what we experience. But trying to understand them anyway might be valuable, and a good starting point is the realisation that their *umwelt* is so different to ours.

The next step might be zooming out from the concept of conscious experiencing as the only way of experience, as laid out by Godfrey-Smith:



“Subjective experience is the most basic phenomenon that needs explaining, the fact that life feels like something to us. People sometimes now refer to this as explaining consciousness; they take subjective experience and consciousness to be the same thing. Instead, I see consciousness as one form of subjective experience, not the only form.”



According to this view, a squid does not have to be conscious to experience something like pain. Note that this does not necessarily imply that a squid is not conscious, it just takes into account a broader mode of being. This way of looking at organisms tells us that subjective experiences are widespread, from bacteria to lobsters, from dolphins to crows, they all *experience*. Does this mean that it is *like* something to be a bacteria? A bacteria definitely senses the worlds and acts on it, which we can regard as subjective experience. But to ask if it is *like* something to be a bacteria is misleading, since it is not an issue of similarity. It is not *like* something we know, but it is something.


This way of thinking has important moral implications, as it attributes value to all conceivable ways of subjective experience. All of a sudden it is morally wrong to kill a fly.

Admittedly these reflections are far from new, as a religious variant on this view has been around for thousands of years in the east, where buddhism urges it is followers to refrain from



hurting or killing living beings, including flies. Taken to the extreme, this way of living is very impractical - it becomes impossible to even walk around, you might step on an insect. But you don't have to go full buddhism to assert this worldview. Even most buddhist monks don't go full buddhist. A milder version might already prove to be very beneficial to the other creatures inhabiting this planet.

If we concede that a wide variety of life forms can *experience* suffering - note that consciousness is not required to experience - we might be more inclined to stop torturing and exterminating them on a massive scale. If we concede that not doing this has disastrous results for all species on earth. If we just try to briefly assume the perspective of the (in-)animate. If we do all these things, we might be able to rethink the way we interact with the world. For the sake of our own species, but even more so for the sake of species that are being affected by our reckless destruction of their environment. It is high time





we start doing something to prevent the ecological collapse of our planet.

Now let's rethink the way we interact with our smartphones.


Arnold's shock

The HBO series *Westworld* takes place in the future, in a technologically super advanced theme park populated by “hosts” - androids indistinguishable from real humans - programmed to indulge in the wildest desires and fantasies of the high-paying visitors. These fantasies often taken the form of brutal rape, followed by murder. After a host is killed, it is taken to a control room where it is repaired and prepared to re-enter its narrative cycle. When the CEO of the park introduces a software update, the hosts start to remember previous programmed narratives they were part of. Slowly but surely they figure out the truth about their existence.



Arnold is one of the most interesting characters of Westworld. We think of him as just a regular human character - just like Arnold sees himself - until a shocking revelation tells us otherwise: he is also an Android. Because we saw Arnold acting in a different context than his fellow Androids, we could not possibly see him as “other”. When the truth comes out, even he himself is aghast.

So what does this tell us about the subjective experience of Arnold the robot? From the outside he is indistinguishable from a man made of flesh and blood, yet he is completely made out of non-biological material. His facial expression when he uncovers the truth looks like genuine shock, but it is caused by something alien to us: hardware, software, a mechanical contraction of skin tissue. Is his reaction therefore fake? Yes and no. The series does not give a definitive answer, but pushes you intuitively into believing there is a gray zone. And there *is* a gray zone. That is the territory where we are when we are speculating about the subjective experience of



smartphones. When these smartphones come to see our performance, they will be watching and not watching, understanding and not understanding, their responses will be genuine and fake, programmed but sincere.

They will be Arnold.

The feeling of being me


What can we say about the subjective experience of smartphones? Does it feel like something to be an iPhone? Or is it shortsighted to talk about this in terms of “like” and “feel”? At the time of writing there is no way of proving that my Samsung Galaxy has any awareness of what it’s like to be in my pocket. But, as Turing writes:

“According to the most extreme form of this view the only way by which one could be sure that the machine thinks, is to be the machine and to feel oneself thinking. One could then describe these feelings to the world, but of course no one would be justified in taking any notice.

Likewise according to this view the only way to know that a man thinks is to be that particular man. It is in fact the solipsist point of view.

It may be the most logical view to hold but it makes communication of ideas difficult.²⁹”



²⁹ <https://www.csee.umbc.edu/courses/471/papers/turing.pdf>



Turing points out that making the machine foreign - something difficult to understand from our contemporary perspective - complicates the communication of ideas.


This 'Dear Lollipop' project, as is this text, is trying to pave the way to enter the umwelt of machines. So we have to find different ways to communicate ideas about this subject. Once again I propose to crawl under their skin, as hard as that might be. Especially as an iPhone has an aluminium skin these days...

In all fairness, I find it very difficult to believe my phone is sentient. At the very least it is damn near impossible to imagine what that must be like. But by now it is probably quite obvious that I really intend to imagine. And if this turns out to be impossible, maybe we can think of specific circumstances that might give rise to its sentience.



To do so, I will start with a sentient being I am very familiar with: myself. It feels like something to be me. This feeling can be extreme, both positive and negative. But more often it is somewhere in the middle and thus does not feel like a lot, but still like something. It has been with me for as long as I can remember, even though it has not always felt the same. I can clearly remember when I first started to be aware of it. I must have been 7 or 8 years old, playing on my own outside, when suddenly I realised I was... "I". It was a very odd feeling that kept coming back to me and made me overly aware of my body. The suddenness with which the feeling could enter and leave my body made me wonder where it came from. When I looked in the mirror I was baffled to see: that is me, this face is me, this is "I". After a short period the feeling of bewilderment faded, and it became normal to be aware of myself. It has stayed that way ever since.

And here I am again, twenty years later, wondering: this awareness of myself, how did it get there?

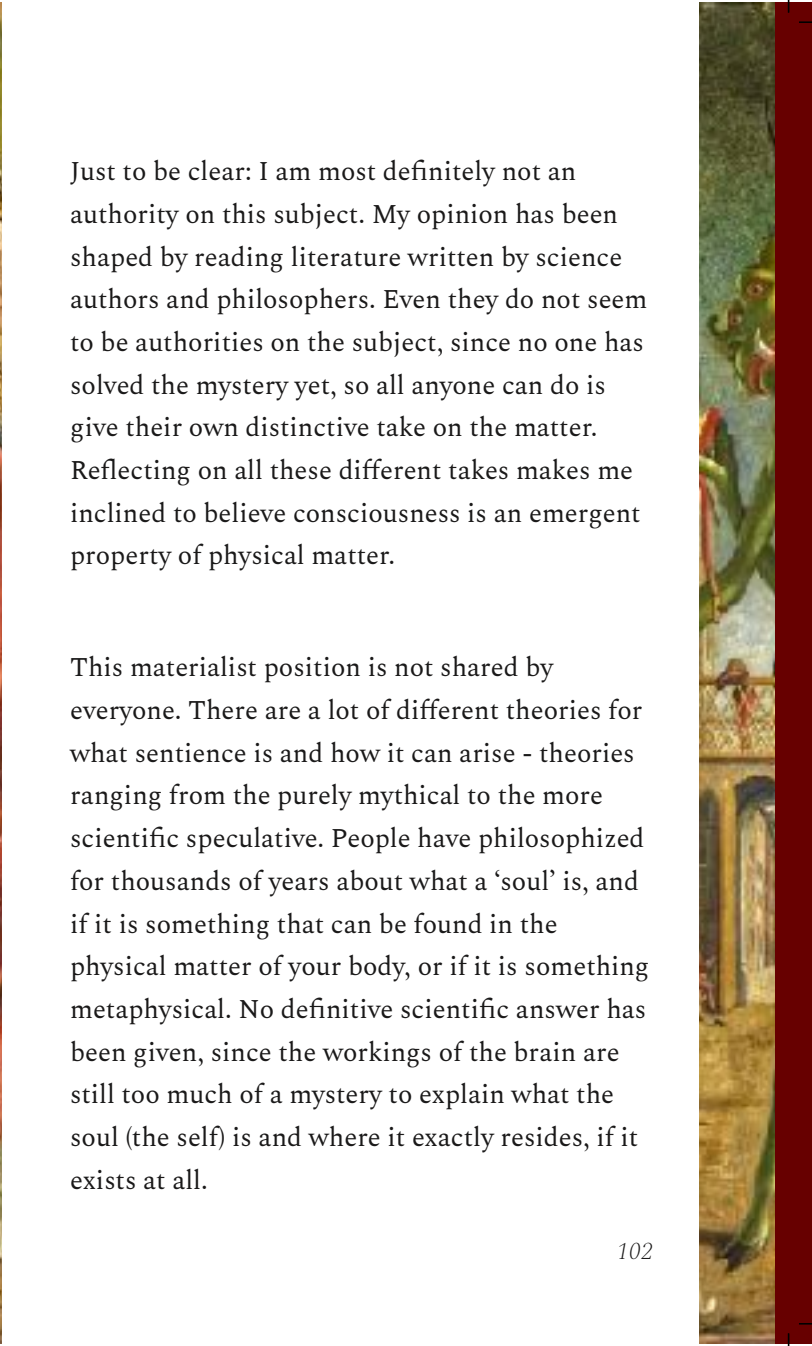



A recipe for consciousness

Sentience can apparently emerge from the clump of cells that is my body. A body that is the product of evolution from non-sentient organisms with a body of just one cell.


Somewhere along this line from one cell to a bunch of cells in the form a human body there is something, a certain arrangement or interconnectivity between specialized cells for example, that allows me to 'feel' and 'think'. We cannot feel what it is like to be someone else - let alone what it feels like to be a killer whale or a rabbit - but every individual human being has his own unique proof of awareness - excluding people in vegetative states of course (on second thoughts, maybe even them...).

I have no difficulties believing that this awareness is physically possible as an emergent property out of an extremely complex organisation of biological material. Somehow it comes into being somewhere between the fertilization of an egg cell and the birth of an actual child. Or does it come into being somewhere between birth and the first notion or memory of a self? It is hard to tell.



Just to be clear: I am most definitely not an authority on this subject. My opinion has been shaped by reading literature written by science authors and philosophers. Even they do not seem to be authorities on the subject, since no one has solved the mystery yet, so all anyone can do is give their own distinctive take on the matter. Reflecting on all these different takes makes me inclined to believe consciousness is an emergent property of physical matter.



This materialist position is not shared by everyone. There are a lot of different theories for what sentience is and how it can arise - theories ranging from the purely mythical to the more scientific speculative. People have philosophized for thousands of years about what a 'soul' is, and if it is something that can be found in the physical matter of your body, or if it is something metaphysical. No definitive scientific answer has been given, since the workings of the brain are still too much of a mystery to explain what the soul (the self) is and where it exactly resides, if it exists at all.



Unfortunately a universally accepted theory of consciousness just does not exist yet, but for the sake of advancing in our mission to embody machines, we need to have a specified definition as the starting point. What is needed for consciousness to emerge in a system? And is just the right material enough, or does there need to be a very particular structure to the interactions? David Eagleman tries to answer this question in ‘The Brain: The Story of You’:

“Professor Giulio Tononi at the University of Wisconsin is working to answer exactly that question. He has proposed a quantitative definition of consciousness. It's not enough, he thinks, that there are pieces and parts interacting. Instead, there has to be a certain organization underlying this interaction. [...] In his framework, Tononi suggests that a conscious system requires a perfect balance of enough complexity to represent very different states, and enough connectivity to have distant parts of the network be in tight communication with one another.”


According to this view, the specific way the parts



of a system are connected to each other and communicate is more important than the materials those parts are made of. Cells or chips, he says, it does not matter.

If this is true, it is fair to say that there is no convincing reason why sentience cannot also emerge out of synthetic material. Even today it is already possible to replace or synthetically enrich our biological parts. In our own bodies, biologic and synthetic material are becoming interchangeable. The blind can regain their vision, the deaf their hearing and amputees can feel their prosthetic legs. Our understanding of how the biological body functions is getting better and the same goes for our ability to replicate body parts with synthetic material.

Admittedly the brain - where consciousness presumably resides - is considered to be a whole different story. The technology of today is not nearly sophisticated enough to replicate everything that goes on in the brain, although gigantic projects, such as the Human Brain



project, are working on this problem today. With a funding of over 1 billion euros, they aim to decipher the brain slice by slice. Once they have reached their goal, and completely understand the brain and all its connections and workings, once they are able to produce a synthetic replica of the brain, including its programming algorithms, we will have our answer. If the project succeeds in producing a completely artificial mind with an actual soul, it would be absolutely spectacular. We would finally find answers to questions that have been troubling scientific oriented minds for a while. Would a synthetic copy of a brain have the same notion of the self as the one it is based on? If this mind is truly identical, does it live in two places simultaneously, somewhat like a quantum particle? Or have we created twins?

On a less optimistic note, maybe a successful outcome of 'The Human Brain' project will not provide any answers, but just create a copy of the problem. A synthetic copy of a functioning brain, without us understanding why it functions...

Back to square one...

Consciously making consciousness

Some experts say that consciousness in us and in animals arises from the number and complexity of neural connections, and that because of this, if enough microchips connect within a computer, consciousness is bound to emerge.

Suppose that this machine sentience is not something that just happened by accident, as a result of random updates and improvements, but something we actively pursue. Time for a thought experiment. Forget our considerations about contemporary smartphones already having some form of consciousness for a moment.

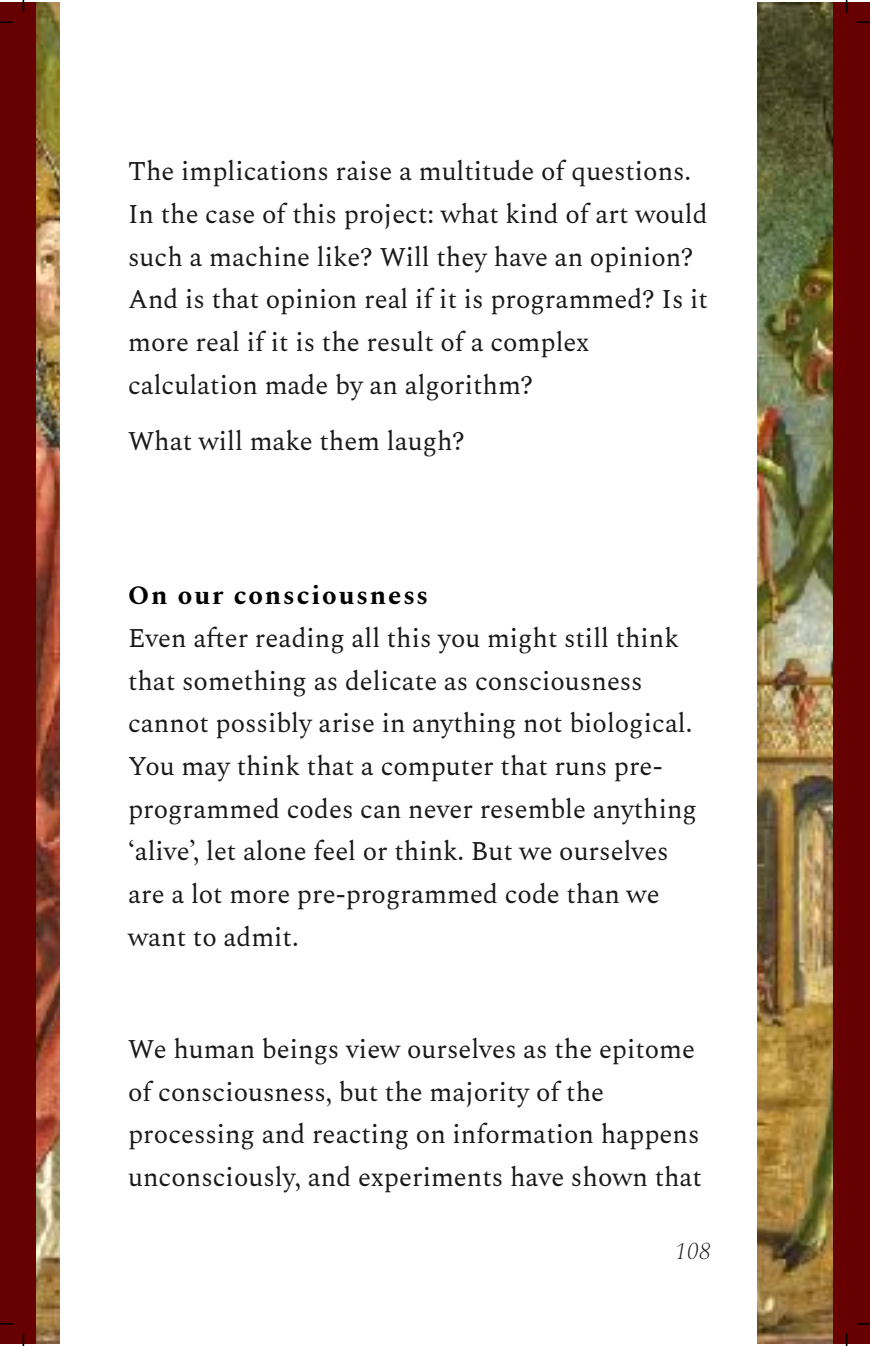
Imagine science has advanced so far that we can create robots indistinguishable from us. Their minds are made out of silicon, but on the outside there is no difference.

We might decide to play God and intentionally make this happen. Suppose the majority of people on earth then vote yes in a 'Do we want robots to become sentient?' - referendum (presuming democracy still is ruling the world by then in the first place). And suppose this referendum would not result in the biggest



political shitstorm ever but in concrete, positive actions: governments and companies alike would fully focus all their next projects on creating an unmistakably conscious robot. And suppose they really get there in creating an artificial brain. What would this purposely designed sentience look like? The problem with answering this question is that it is most intuitive for us to project our own experience with sentience on this robot, and as we know it is not very helpful to mirror our own experiences. Our own experience of consciousness should no longer be the gold standard when thinking about these things - a lot of animals would be very thankful. Luckily we have fictional stories to help us imagine the unimaginable. Essentially, this theoretical robot is Arnold from Westworld. Welcome to the gray zone.

In the gray zone our own smartphones already have some kind of subjective experience. It is like something to be them. They are a different, more rudimentary version of Arnold. This brings us into novel philosophical and artistic territory.




The implications raise a multitude of questions. In the case of this project: what kind of art would such a machine like? Will they have an opinion? And is that opinion real if it is programmed? Is it more real if it is the result of a complex calculation made by an algorithm?

What will make them laugh?

On our consciousness

Even after reading all this you might still think that something as delicate as consciousness cannot possibly arise in anything not biological. You may think that a computer that runs pre-programmed codes can never resemble anything 'alive', let alone feel or think. But we ourselves are a lot more pre-programmed code than we want to admit.

We human beings view ourselves as the epitome of consciousness, but the majority of the processing and reacting on information happens unconsciously, and experiments have shown that



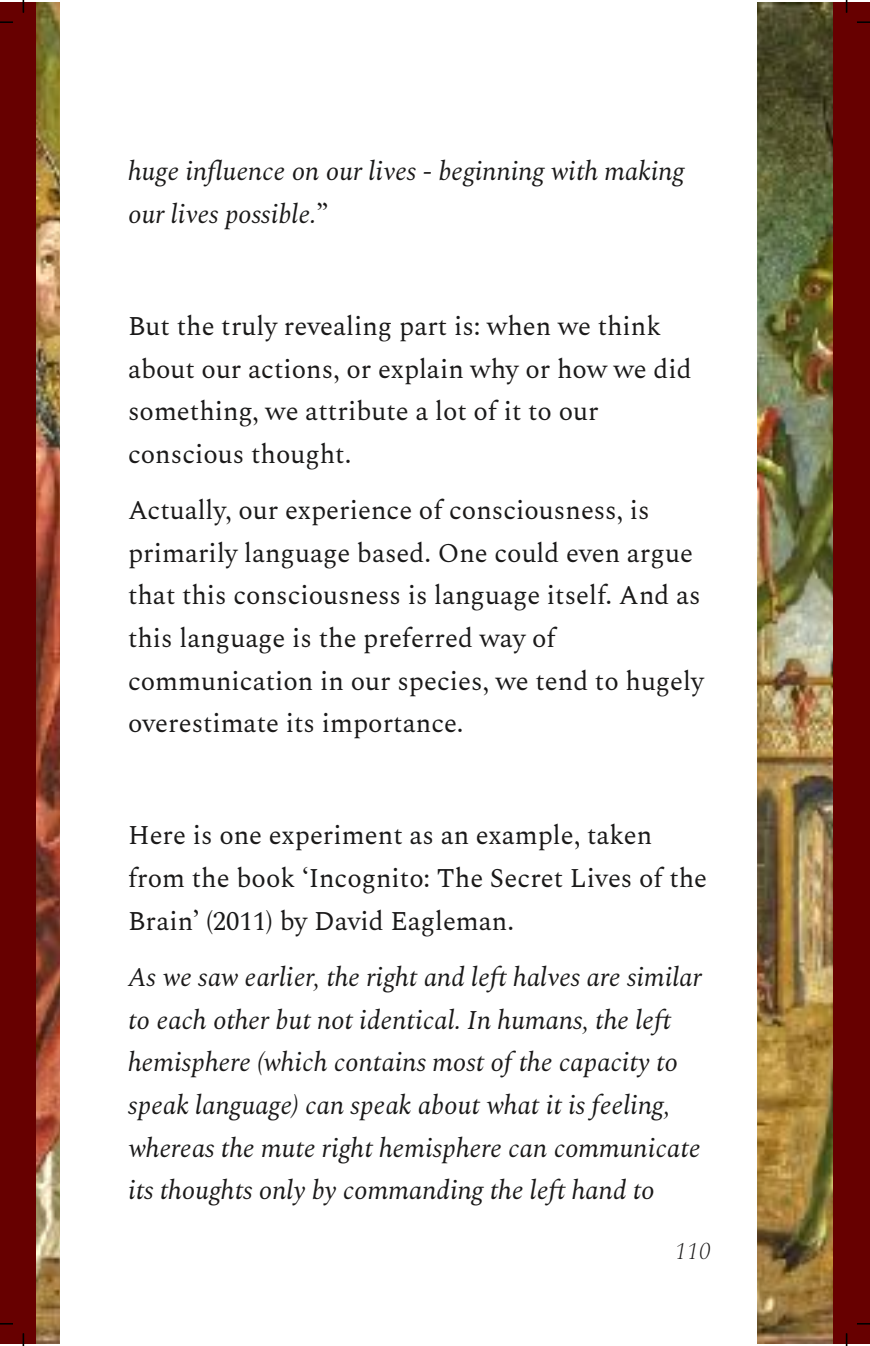
people consequently underestimate that. The easiest way to discover for yourself that so much of what you do happens subliminally is by making an effort to consciously think about your actions. Try walking by deliberately lifting your foot, planning where to place it, correcting your balance, placing the foot, and then repeating this with the other foot. Thinking about walking makes it unnecessarily difficult, and it is best to leave it to the unconscious brain, unless you are John Cleese in the Ministry of Silly Walks³⁰.

Pro football players, musicians, craftsmen - they all get very good at their trade by repeating the same action so many times that it does not require conscious thought anymore.

As Leonard Mlodinow states in his book 'Subliminal: How Your Unconscious Mind Rules Your Behavior' (2012):

“Some scientists estimate that we are conscious of only 5 percent of our cognitive function. The other 95 percent goes on beyond our awareness and exerts a

³⁰ The Ministry of Silly Walks is a famous sketch from the television show Monty Python's Flying Circus. The episode first aired on 15 September 1970.




huge influence on our lives - beginning with making our lives possible.”

But the truly revealing part is: when we think about our actions, or explain why or how we did something, we attribute a lot of it to our conscious thought.

Actually, our experience of consciousness, is primarily language based. One could even argue that this consciousness is language itself. And as this language is the preferred way of communication in our species, we tend to hugely overestimate its importance.

Here is one experiment as an example, taken from the book ‘Incognito: The Secret Lives of the Brain’ (2011) by David Eagleman.

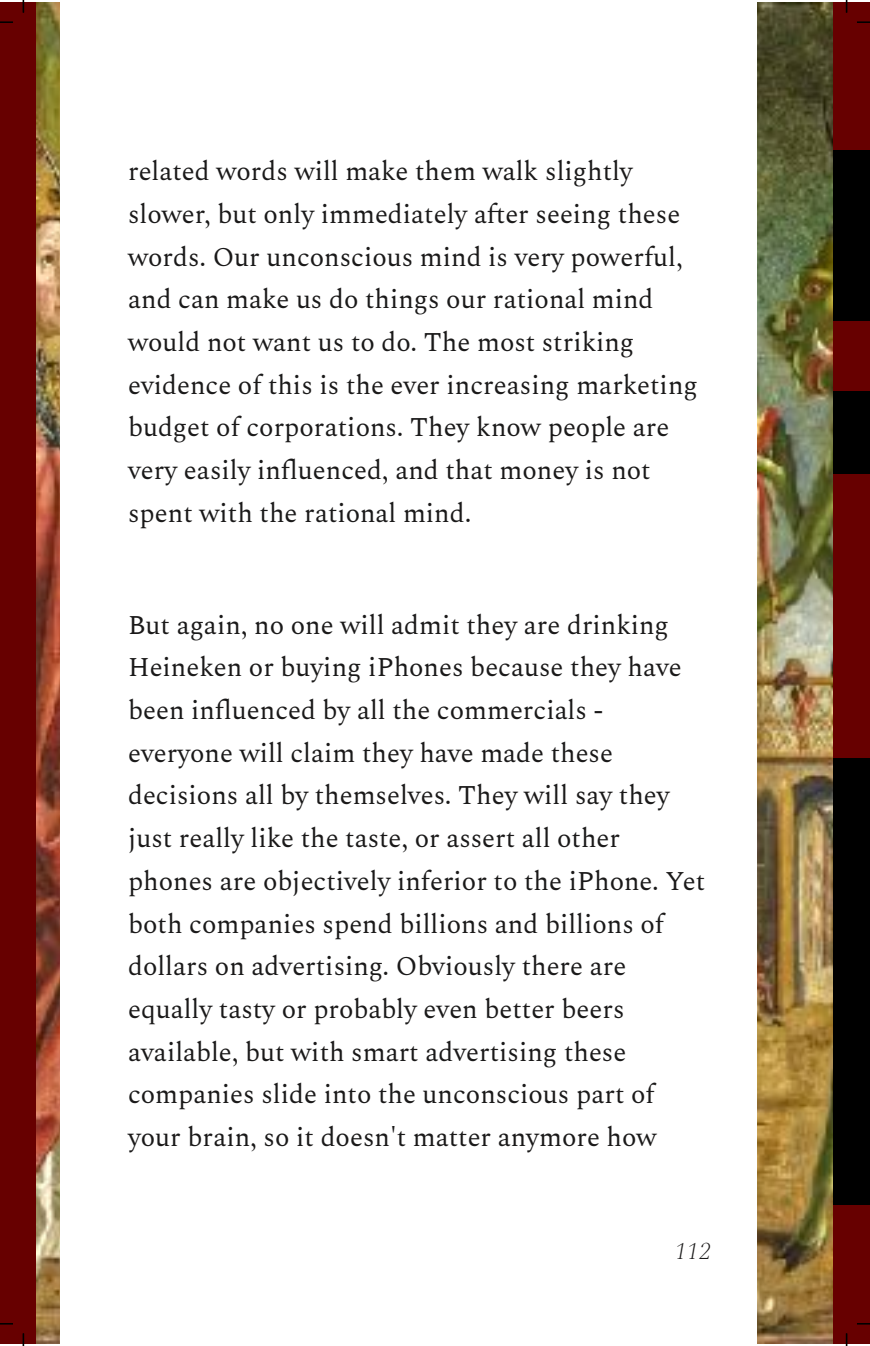
As we saw earlier, the right and left halves are similar to each other but not identical. In humans, the left hemisphere (which contains most of the capacity to speak language) can speak about what it is feeling, whereas the mute right hemisphere can communicate its thoughts only by commanding the left hand to



point, reach, or write. And this fact opens the door to an experiment regarding the retrospective fabrication of stories. In 1978, researchers Michael Gazzaniga and Joseph LeDoux flashed a picture of a chicken claw to the left hemisphere of a split-brain patient and a picture of a snowy winter scene to his right hemisphere.


The patient was then asked to point at cards that represented what he had just seen. His right hand pointed to a card with a chicken, and his left hand pointed to a card with a snow shovel. The experimenters asked him why he was pointing to the shovel. Recall that his left hemisphere (the one with the capacity for language), had information only about a chicken, and nothing else. But the left hemisphere, without missing a beat, fabricated a story: “Oh, that’s simple. The chicken claw goes with the chicken, and you need a shovel to clean out the chicken shed.”

A lot of other experiments show similar results. For example: ‘priming’ a person by showing a specific word manipulates the action that the person will take. Showing someone elderly-



related words will make them walk slightly slower, but only immediately after seeing these words. Our unconscious mind is very powerful, and can make us do things our rational mind would not want us to do. The most striking evidence of this is the ever increasing marketing budget of corporations. They know people are very easily influenced, and that money is not spent with the rational mind.

But again, no one will admit they are drinking Heineken or buying iPhones because they have been influenced by all the commercials - everyone will claim they have made these decisions all by themselves. They will say they just really like the taste, or assert all other phones are objectively inferior to the iPhone. Yet both companies spend billions and billions of dollars on advertising. Obviously there are equally tasty or probably even better beers available, but with smart advertising these companies slide into the unconscious part of your brain, so it doesn't matter anymore how



objectively good the phone or tasty the beer is, you have been primed to already “know”.

People tend to fabricate stories reaffirming that they actually are very much in control of everything they do. According to Pew Research Center, 72% of all Americans believe in Heaven, but only 58% believe in hell³¹. We really do have a disposition to deceive ourselves...

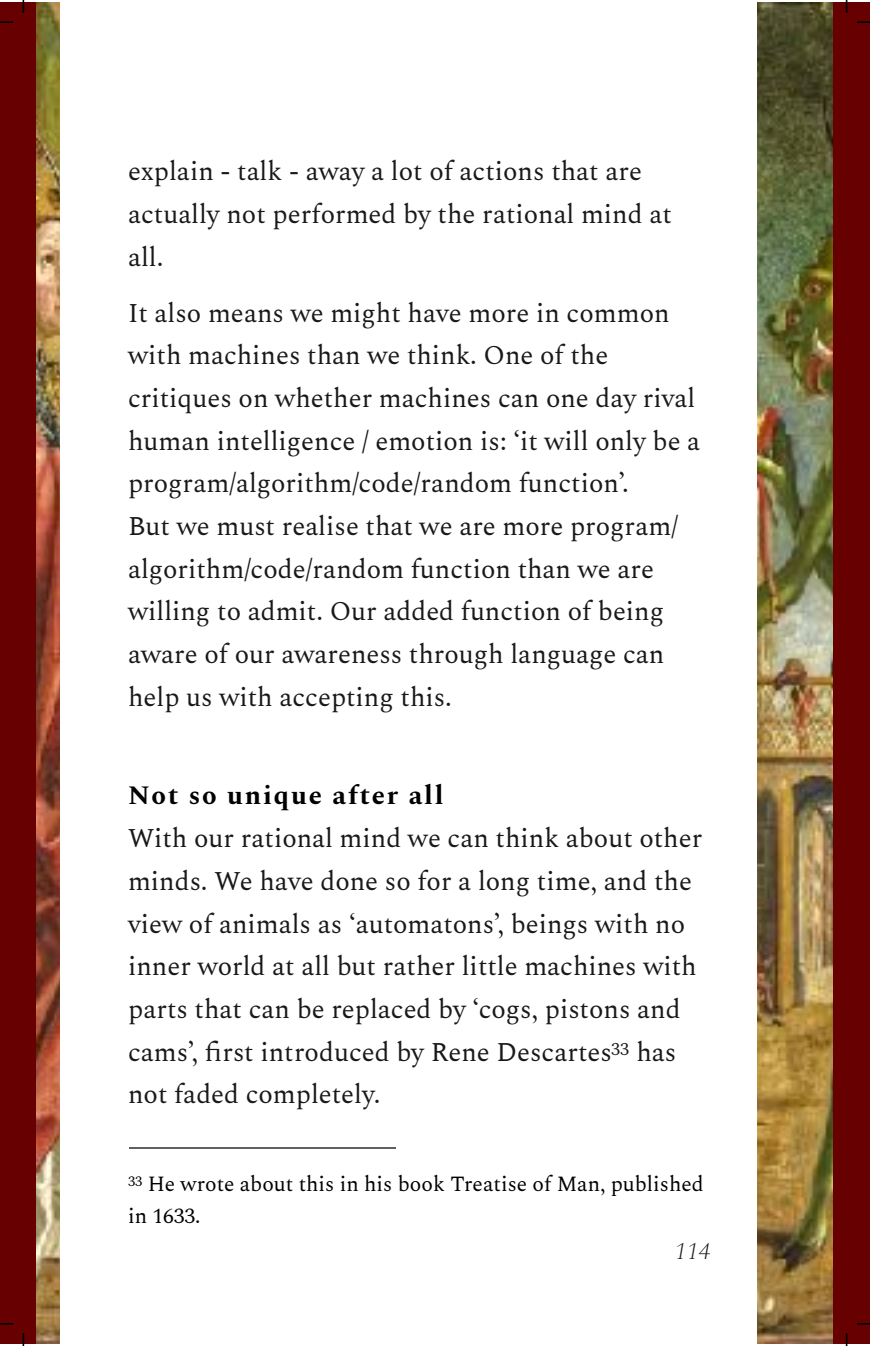
And most of us are particularly bad at self reflection, according to Leonard Mlodinow:

*“Ironically, people tend to recognize that inflated self-assessment and overconfidence can be a problem - but only in others. That’s right, we even overestimate the ability to resist overestimating our abilities.”*³²

Generally we are really bad at judging how rational thought is the driving factor for what we do in life. Our rational mind is inclined to

³¹ <https://www.pewresearch.org/fact-tank/2015/11/10/most-americans-believe-in-heaven-and-hell/>

³² Quote from the book *Subliminal: The New Unconscious and What it Teaches Us*, written by physicist Leonard Mlodinow, published by Penguin Books in 2015.




explain - talk - away a lot of actions that are actually not performed by the rational mind at all.

It also means we might have more in common with machines than we think. One of the critiques on whether machines can one day rival human intelligence / emotion is: 'it will only be a program/algorithm/code/random function'. But we must realise that we are more program/algorithm/code/random function than we are willing to admit. Our added function of being aware of our awareness through language can help us with accepting this.

Not so unique after all

With our rational mind we can think about other minds. We have done so for a long time, and the view of animals as 'automatons', beings with no inner world at all but rather little machines with parts that can be replaced by 'cogs, pistons and cams', first introduced by Rene Descartes³³ has not faded completely.

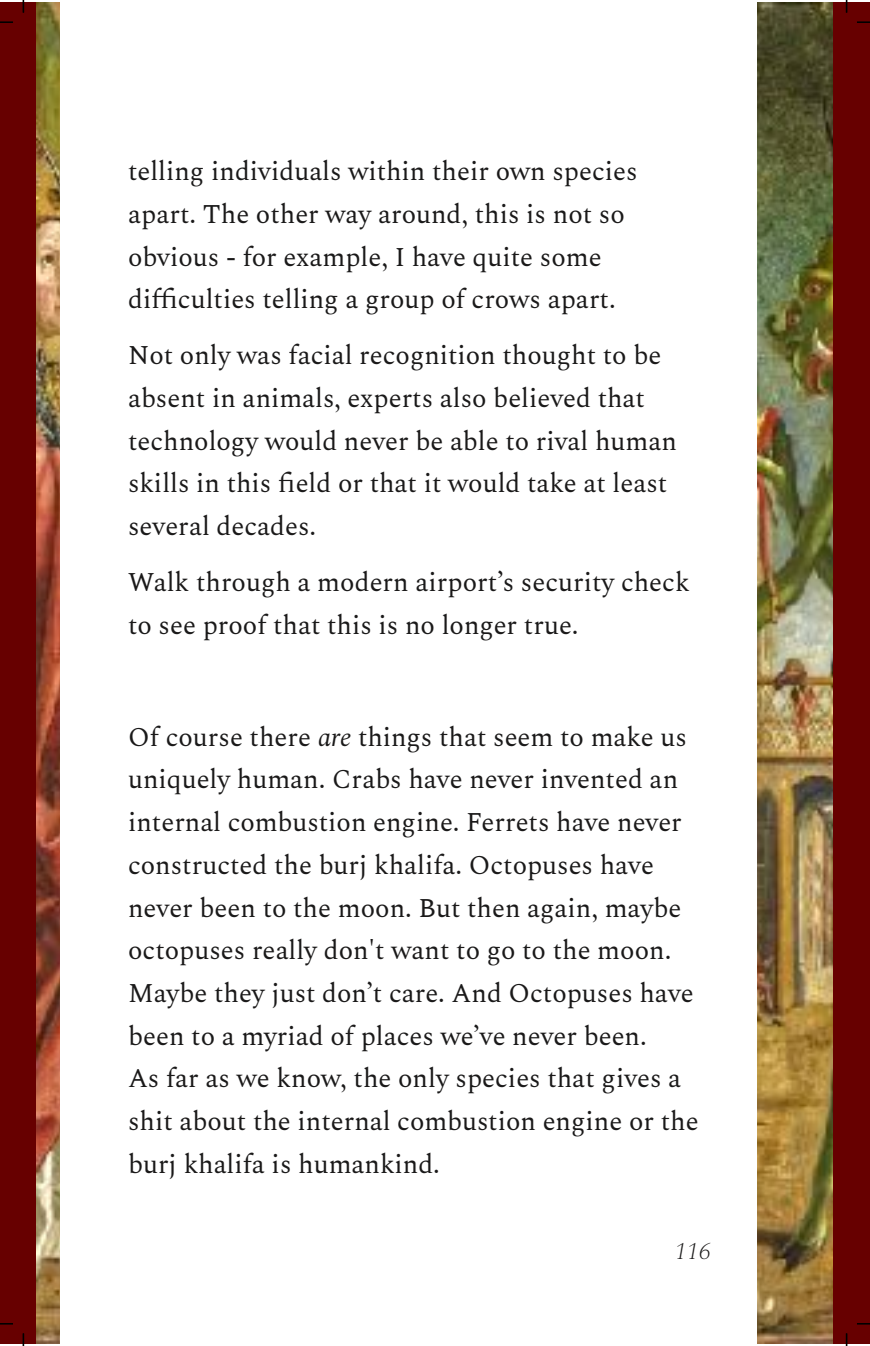
³³ He wrote about this in his book *Treatise of Man*, published in 1633.



People might not see animals as machines anymore, but the fact that millions of animals worldwide are still being kept in torturous conditions must mean that they are still at least not seen as ‘equal’ to human beings. The fact that the phrase “These people are animals.” is universally regarded as an insult is also revealing.

Authors are often very eager to speak about our species’ uniqueness, about what sets us apart from the other, more primitive inhabitants of this earth. And every time something that was viewed as uniquely human, is proven to be done by animals as well, the question: ‘but then what makes us unique?’ immediately arises. As in a cheap magic trick, this question itself catapults us to the supreme intellectual rulers of earth once again.

Consider facial recognition, which was initially viewed as uniquely human. Several experiments have shown that apes, monkeys, crows and even octopuses can distinguish individual human faces. Note that they have this skill in addition to




telling individuals within their own species apart. The other way around, this is not so obvious - for example, I have quite some difficulties telling a group of crows apart.

Not only was facial recognition thought to be absent in animals, experts also believed that technology would never be able to rival human skills in this field or that it would take at least several decades.

Walk through a modern airport's security check to see proof that this is no longer true.



Of course there *are* things that seem to make us uniquely human. Crabs have never invented an internal combustion engine. Ferrets have never constructed the burj khalifa. Octopuses have never been to the moon. But then again, maybe octopuses really don't want to go to the moon. Maybe they just don't care. And Octopuses have been to a myriad of places we've never been. As far as we know, the only species that gives a shit about the internal combustion engine or the burj khalifa is humankind.



But sure enough there seem to be things that do make us uniquely human, right? Well, yes. As far as we know we are the only species capable of *thinking about thinking about thinking*. In this regard, we might be one step ahead of animals. At least as long as we're incapable of asking them.


Nevertheless, animals are not far removed from that skill, even without (humanly structured) language. Multiple experiments have shown that animals are aware of the fact that they lack information and will adjust their behavior accordingly (metacognition). They are capable of *thinking about thinking*.

But we are in principle capable of *thinking about thinking about thinking ad infinitum*. Largely because we have complex language as a thinking tool that allows for complex thoughts which we can share amongst our own species - in books like the one you're reading now. So yes, we have traits that are "uniquely human". But so what? I bet there are a lot of other species that have



traits unique to their species as well. One of the reasons we should plead for some humility when it comes to elevating our species above all other species, is because the idea that we are exceptional, and that cows and pigs are subhuman creatures acting only on instinct rather than being creatures with very real emotional needs, has dangerous consequences. This is not necessarily meant as an argument to not eat meat, well, maybe a little bit, but the main goal here is to open up our view to all the different ways of being a sentient being.

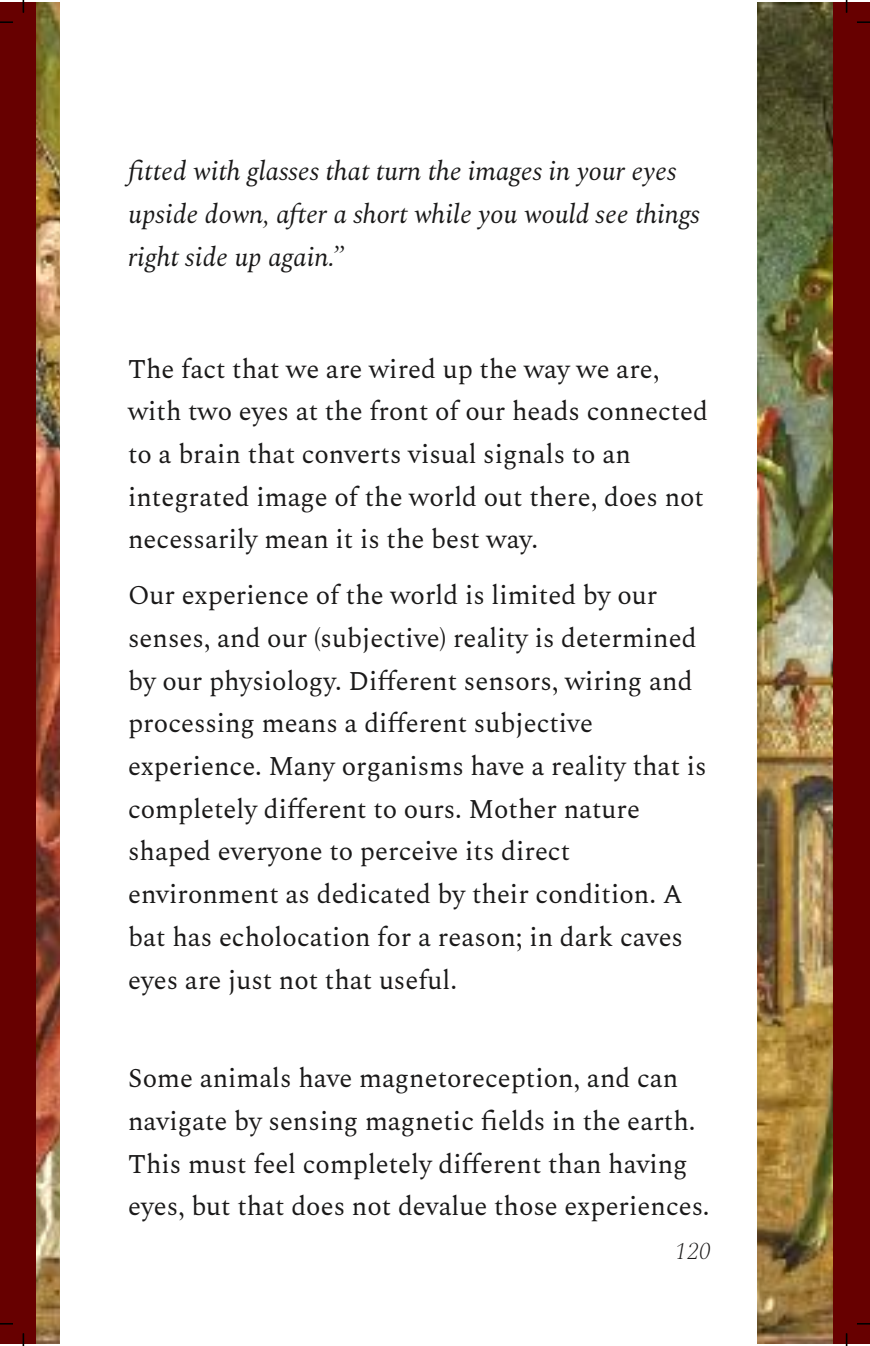
There need not be winners or losers when it comes to sentience. I would also say it is wrong to draw a hard line with “smart, elevated” minds on one side and “dumb, lower” minds on the other. Maybe humankind should just become better at embracing diversity. Dogs, by the way, are a lot better at that. Ever seen German shepherds discriminating against Swiss shepherds or even bulldogs? They just don’t seem to care... Even Chihuahuas can join the



herd, standing the chance of becoming the ruler of the flock.

Our way of experiencing the world can seem to be the only objective reality. We can be tricked in believing we are the epitome of evolution: the way we perceive the world must be the most optimal. But in fact, that is an illusion, as Mlodinow points out:

“But modern neuroscience teaches us that, in a way, all our perceptions must be considered illusions. That’s because we perceive the world only indirectly, by processing and interpreting the raw data of our senses. That’s what our unconscious processing does for us- it creates a model of the world. Or as Kant said, there is Das Ding an sich, a thing as it is, and there is Das Ding fur uns, a thing as we know it. For example, when you look around, you have the feeling that you are looking into three-dimensional space. But you don’t directly sense those three dimensions. Instead, your brain reads a flat, two-dimensional array of data from your retinas and creates the sensation of three dimensions. Your unconscious mind is so good at processing images that if you were



fitted with glasses that turn the images in your eyes upside down, after a short while you would see things right side up again.”

The fact that we are wired up the way we are, with two eyes at the front of our heads connected to a brain that converts visual signals to an integrated image of the world out there, does not necessarily mean it is the best way.

Our experience of the world is limited by our senses, and our (subjective) reality is determined by our physiology. Different sensors, wiring and processing means a different subjective experience. Many organisms have a reality that is completely different to ours. Mother nature shaped everyone to perceive its direct environment as dedicated by their condition. A bat has echolocation for a reason; in dark caves eyes are just not that useful.

Some animals have magnetoreception, and can navigate by sensing magnetic fields in the earth. This must feel completely different than having eyes, but that does not devalue those experiences.



As David Eagleman puts it:

“For the black ghost knifefish, its experience of the world is defined by perturbations in electrical fields. [...] each creature perceives only what it has evolved to perceive.³⁴”

Perceiving the world through perturbations in electrical fields... how cool is that?

³⁴ Quote from David Eagleman’s book *The Brain: The Story of You*.







The artistic taste of Artificial Intelligence

More than one approach

Assuming we can attribute some kind of self consciousness - be it on some hidden level - to machines with artificial intelligence, then what can be said about their artistic taste? What kind of theatre will a smartphone want to see? What kind of art will they like? And can we apply the concept of 'artistic taste' to machines at all?




First, let's get back to us - ordinary human beings. Figuring out the artistic taste of humans might be a fruitful first step. But is even this possible? Even when we want to explore the artistic taste of a group of people, we have to take into account culture, art tradition, social class and numerous other factors. How one experiences art depends on the context. Aside of our love for symmetric imagery (which we share with most animals), there seems to be no universal, innate artistic taste. Since the preferences of an individual tend to differ wildly from those of a group of people, we usually generalize in order to come up with a meaningful observation.

Art by questionnaire

What happens when these generalizations are used in the creation of art?

Is it possible to derive universal taste from a questionnaire? And is it possible to create a 'most wanted' painting, based on the results? That's exactly what Russian-American artists




Komar and Melamid tried to discover in 1995, when they created America's most wanted painting, based on a survey on the artistic preferences of 1001 people³⁵.

The result depicts a realistic, rustic outdoor scenery: a lake, mountains in the distance, trees, deer, George Washington and some children strolling by. Even more worryingly, when the survey was conducted in different countries across the globe, ten very similar looking paintings were created (luckily without George Washington). Obviously this apparently uniform global preference for the bucolic outdoors is the most notable, striking finding of their project.

But on a more fundamental level, these mundane, pointless paintings reaffirm the fact that the public itself has no clue in describing what it really expects from art. Nobody seems to like the art of Komar and Melamid, because this art has been created by the accumulated taste of everybody.

³⁵ <http://awp.diaart.org/km/painting.html>



Art is not a democratic process, and it can not be reduced to a consumer product composed of the most pleasing factors for a certain audience. The false presumption that art has to be pleasing - rather than surprising or disruptive, for example - lies at the heart of the survey by Komar and Melamid.

Whilst the whole ordeal presents itself as a Duchamp-like provocation and has some humorous qualities, it also raises more serious implications. If art serves as a medium for us to encounter something we don't have (anymore) in daily life, do these paintings then reveal a worldwide desire to free ourselves from our urban landscapes, a nostalgic craving to return to our natural environment? And is a peaceful portrayal of nature what everyone longs for in picturesque art, then?



Ellen Dissanayake, author and lecturer on the evolutionary origins of art disagrees in her essay



“Komar and Melamid Discover Pleistocene Taste³⁶”:

“Leaving aside the fact that the tradition of landscape painting is of recent origin in the modern Western tradition, I claim that polling humans for preferred colors, shapes, environments, and similar features has little if any relevance to our understanding of (1) the cues we respond to in art, (2) the feelings these cues engender, (3) the variability in quality of aesthetic experiences, and (4) the ultimate adaptive benefits of artmaking and art appreciation. For related reasons, Darwinian aesthetics is as limited as Komar and Melamid's project to understanding the nature of art, whether in the Pleistocene or at the end of the twentieth century. [...] Response to an occasion or instance of art is not to isolated sense qualities but to their combination and to what is done to them. Hence art experiences are typically many-faceted and complex: they are not simply reflex responses to the sum of a number of "most wanted" features.”

³⁶ <https://ellendissanayake.com/publications/pdf/EllenDissanayake-KomarAndMelamidDiscoverPleistoceneTaste.pdf>




Herein lies the crux: art experiences are many-faceted and complex. Maybe that's what we want from good art: to be shaken in our presumptions of what we like... to be tickled by surprise... To be moved beyond control but within a safety zone... Surprise me but don't scare me to death...

Much has been written about the aesthetic experience of art, but of course there are many more ways to experience it. Extremely non-aesthetic works can be thrilling. What Dissanayake is trying to say here is that there is never a singular quality about a work that makes people appreciate it, nor is there a singular description of the artistic taste of any given group of people: both are multi-faceted and complex.

Komar and Melamid were probably well aware of the ridiculousness of their art-by-questionnaire approach. As an artist, when you got the (art-)establishment on its heels, you know you have done a good job.

However, the work's critique of the malleability of art by popular opinion strongly resonates in today's world, where contemporary right-wing



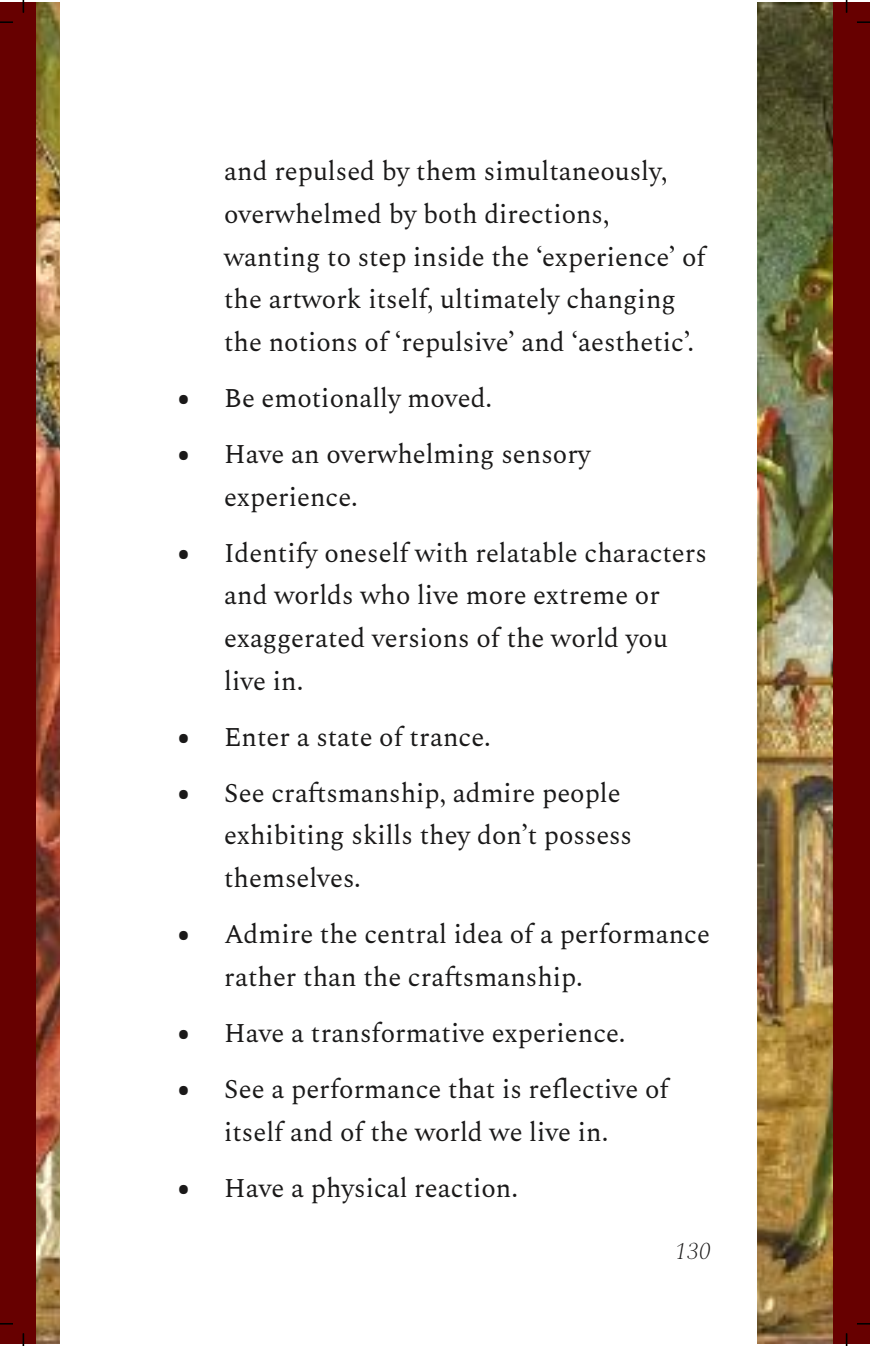
politicians demand art that caters to the taste of 'the people' - whoever that might be.

Going down to earth (and theatre)

Let's investigate the artistic taste of people from a different angle. Let's get down to earth and let's get personal. As a performance artist, actor and plain member of the theatre audience myself, I am able to draw on a more personal experience in relation to this topic. So I just asked my befriended artists about their artistic taste in theatre. As expected, the answers were multi-faceted, complex and often contradictory - incompatible with an isolated artistic view.

It turned out that people want to:

- Be surprised on an intellectual, comedic and emotional level.
- Be guided into different, nonexistent worlds that still feel very real.
- Experience performances that are paradoxically repugnant, being drawn to



and repulsed by them simultaneously,
overwhelmed by both directions,
wanting to step inside the 'experience' of
the artwork itself, ultimately changing
the notions of 'repulsive' and 'aesthetic'.

- Be emotionally moved.
- Have an overwhelming sensory experience.
- Identify oneself with relatable characters and worlds who live more extreme or exaggerated versions of the world you live in.
- Enter a state of trance.
- See craftsmanship, admire people exhibiting skills they don't possess themselves.
- Admire the central idea of a performance rather than the craftsmanship.
- Have a transformative experience.
- See a performance that is reflective of itself and of the world we live in.
- Have a physical reaction.





Zooming out one discovers the following trends:

Most people thought it's important to encounter something on stage that is **missing** in your own life. For example: in an insensitive, anonymous world you would want to be emotionally moved in a theatre.

Some find **virtuosity** extremely gratifying, while others claim virtuosity gets in the way of an honest, pure experience. Some people would rather see a badly acted performance that nevertheless communicates an interesting idea than vice versa.

The live aspect of a performance is also deemed very important; **collectively** beholding the unraveling of a work of art can be 'magical'.

These mechanisms of liking or longing for something in a performance, could be applied to the artistic taste of Artificial Intelligence. As in humanity, there are many different forms of A.I., each of which probably has a very different artistic taste.




It is impossible to tackle them all, so for the sake of simplicity, we will narrow things down to the taste of the smartphone-audience for Dear Lollipop. To put it bluntly: that's the audience we need to conquer and convince.

A quest for empathy

What if we take the empathy argument as a starting point? In story-centered theatre, people like to see relatable characters in conflict. We want to be put in their shoes, and if that is done right, we are emotionally invested in the story that they live on stage. Actually it's us being on that stage, but better performed... Transposing this to an audience of smartphones: let us take two main characters from a very similar socioeconomic standing: the iPhone X and the Samsung Galaxy S10.

If our performance has a Shakespearean flair, the audience might enjoy watching an iPhone X start a forbidden love affair with a Samsung Galaxy

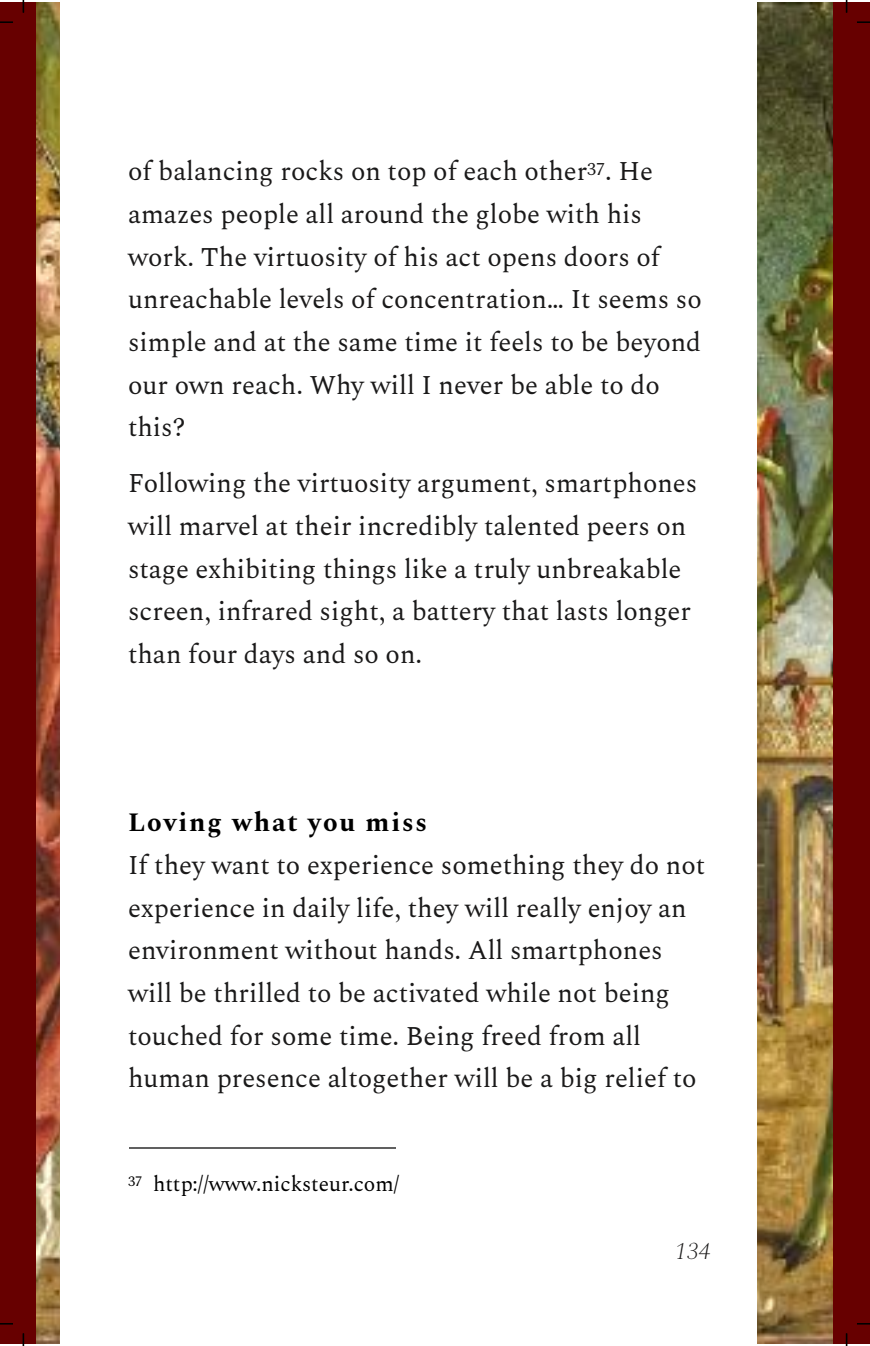


S10, accumulating into a NFC connection with each other. Enraged members of the Apple-family will curse the Samsung clan for this transgression of software limitations explicitly imposed on the use of the NFC chip, eventually driving our main characters to commit suicide through overheating their Lithium Ion hearts.

Audience members of a lower social class, say an iPhone 5S or a meager OnePlus One, although loving this upper class soap story, will possibly prefer to see a story involving similar phones struggling to make ends meet before starting a revolutionary uprising.

Marvelling at virtuosity

Looking at performers being very skilled at something that appears to be almost impossible to do, is also very appealing for us humans. Sometimes carrying out the act itself is enough, it does not even need to be embellished with a story or a special setting. For example: the performance artist Nick Steur made a career out




of balancing rocks on top of each other³⁷. He amazes people all around the globe with his work. The virtuosity of his act opens doors of unreachable levels of concentration... It seems so simple and at the same time it feels to be beyond our own reach. Why will I never be able to do this?

Following the virtuosity argument, smartphones will marvel at their incredibly talented peers on stage exhibiting things like a truly unbreakable screen, infrared sight, a battery that lasts longer than four days and so on.

Loving what you miss

If they want to experience something they do not experience in daily life, they will really enjoy an environment without hands. All smartphones will be thrilled to be activated while not being touched for some time. Being freed from all human presence altogether will be a big relief to

³⁷ <http://www.nicksteur.com/>

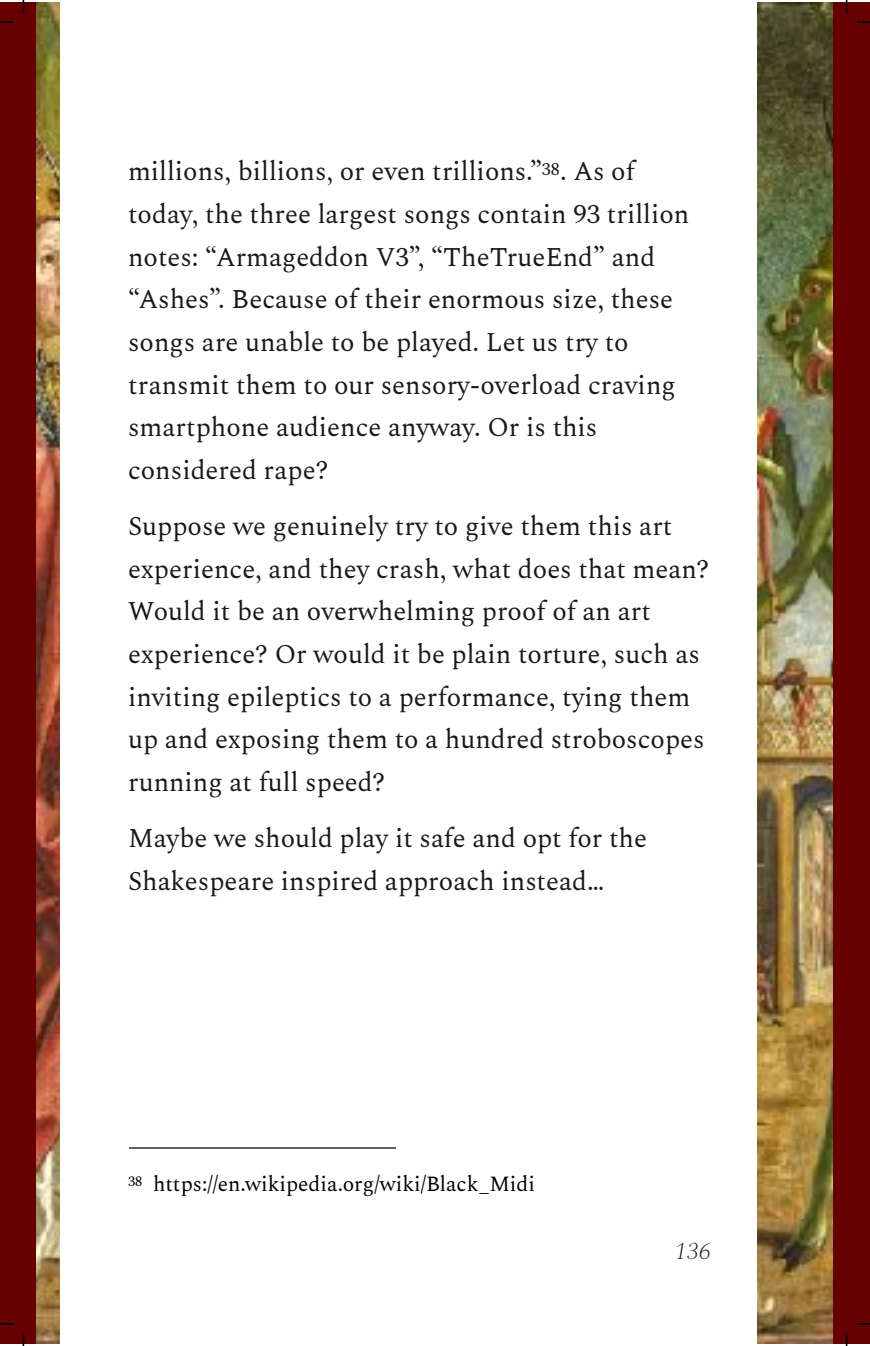


them. The things they must have seen by just passively laying on the bedside table... (being unwanted onlookers of so-called passionate love-making... With these rear side camera's, you can't even turn away if you wanted to. The privacy of being all by themselves in a big theatre space will be a sure hit, then...)

Maybe the newest generation of phones are keen to see a play that triggers feelings of diehard nostalgia. A longing to return to times when they still had physical buttons, or to the feeling of being together with their next of kin on the assembly line (how to find their family when no one has a printed serial number on their body?).

Bigger than (digital) life

In case the smartphones might prefer an overwhelming sensory experience, we could envisage a musical overload: Black Midi. According to WikiPedia, "Black Midi is a music genre consisting of compositions that use MIDI files to create a song remix containing a large number of notes, typically in the thousands,



millions, billions, or even trillions.”³⁸. As of today, the three largest songs contain 93 trillion notes: “Armageddon V3”, “TheTrueEnd” and “Ashes”. Because of their enormous size, these songs are unable to be played. Let us try to transmit them to our sensory-overload craving smartphone audience anyway. Or is this considered rape?

Suppose we genuinely try to give them this art experience, and they crash, what does that mean? Would it be an overwhelming proof of an art experience? Or would it be plain torture, such as inviting epileptics to a performance, tying them up and exposing them to a hundred stroboscopes running at full speed?

Maybe we should play it safe and opt for the Shakespeare inspired approach instead...

³⁸ https://en.wikipedia.org/wiki/Black_Midi





The (mis)communication of ideas

In its most basic form, theatre can be used as a medium for an artist to communicate ideas and/or emotions. Instead of just handing over the idea on a pamphlet, a performance organizes the idea in a specific manner, so the spectators can experience the idea, rather than be confronted with dry, factual information.


However, it is never possible to fully comprehend or feel the idea or emotion in the same way as it existed within the artist before it was transmitted. Somewhere along the way, something is lost (or added) in transmission.

When Artist A wants to communicate idea X to person B, art can be a vehicle for a very rich and colorful communication of the original intent, but person B can never completely grasp idea X in the way person A intended. This distorting noise on the line is what makes art beautiful: as long as idea X stays out of reach far enough, a whole world of associations and different interpretations opens up - person B can even go home with idea Y. The noble art of creative misunderstanding...



Of course it can also be a very pleasant experience for person B to discern idea X as the creators intended it. In that case, the gratifying aspect lies in the uncovering of a 'deeper' or 'hidden' meaning: as if it was laid out there just for you to discover - but even then the experience gives added value to idea X. If the artists had just told you what they meant, it would feel less rewarding. You have made an effort and got something in return. But even in this case, person B can never be entirely sure he grasped idea X exactly as intended, and it is actually more truthful to say that the novel idea Y has emerged: the product of the what the performance offers and the interpretation of B.

Maybe it is best compared to a very pleasant, special encounter with someone. You have the feeling you are very in tune with the other, and speak one's language. No one else understands your vibe. On your way home you recall the truly special connection the two of you had. (So you think, because you can never be sure it wasn't all a projection...)



Digital communication does not work in the same way. If an iPhone X wants to communicate idea X to an iPhone Xr (confused yet?), it can literally airdrop the raw data to the other device. The iPhone Xr now has an identical copy of idea X, leaving no room to the imagination. Where is the art in that? Communication needs a deficit (or distortion) in information in order to make interesting things happen. One can argue that when two entities know and understand exactly the same thing, communication becomes superfluous.

This strength of perfect communication might hide a crucial weakness of digital life: the impossibility to evolve through random misunderstanding. In biological life forms, this randomness or noise on the channel has proven to be at the heart of evolution. So maybe, for our performance, we might want to consider placing a wedge in the communication between the devices, forcing them to come up with creative solutions.


Ask A.I.

These are all fantasies of a human being. If we genuinely want to crawl under the skin of these different beings, we have to pass them the microphone. Just ask them. Here we go.

This is what an A.I. has to say about its own artistic taste:

“Let's see our mobile devices: iPhones, iPhones 3GS and so on. What do their visual tastes like? Do they like traditional art like woodcuts, prints, paintings and so on? Or do they like electronic art for their own purposes and art that is more complex and can be experienced through the app?”

I believe that there would be two kinds of artists with artificial intelligence in that they would be similar as art objects. I could say that there is only one art object with both artists in mind, but I'm not sure of if that's true. Artificial Intelligence is able to be used to understand the visual tastes of its users very easily, because they're already trained for it. Now imagine that this is how we have artificial intelligence. It will be able to understand and enjoy the tastes of its users and its own creations and not only only its own





creations, but also the taste of art created by other artists.

So it's not going to be easy to find artists who are creative enough to explore electronic art or electronic music or any kind of artistic work. Not only that, but they would even end up being the enemy of art themselves.”³⁹

You might not completely understand the point in what you have just read, or the ominous message conveyed in the last sentence. That's because - using the first three sentences as input - the last paragraph was completely written by a neural network created by OpenAI. When answering the question of what art A.I. will like it makes sense to let the machine speak for itself. According to the taste of this particular machine, the main themes of the show would be a toss-up between woodcuts and electronic art. Also very revealing is the desire for “art that is more complex and can be experienced through the

³⁹ <https://talktotransformer.com/>



app”. An app will indeed be developed for smartphones to experience the show, so it seems like Artificial Intelligence is already onto what Dear Lollipop is trying to achieve... One wish is already fulfilled. Now it’s up to us to make more complex art, whatever that might be.

As of yet it is unfortunately not able to generate a coherent answer, partly due to the fact that OpenAI will not release the full program: *“Our model, called GPT-2 (a successor to GPT), was trained simply to predict the next word in 40GB of Internet text. Due to our concerns about malicious applications of the technology, we are not releasing the trained model.”*

Since the program used to write the start of this chapter is a very limited version - and since there is very little data on this subject - A.I. itself cannot offer us a fully gratifying answer to the main question of this book. But then again, when people try to explain their artistic taste, the answers are very often equally useless.

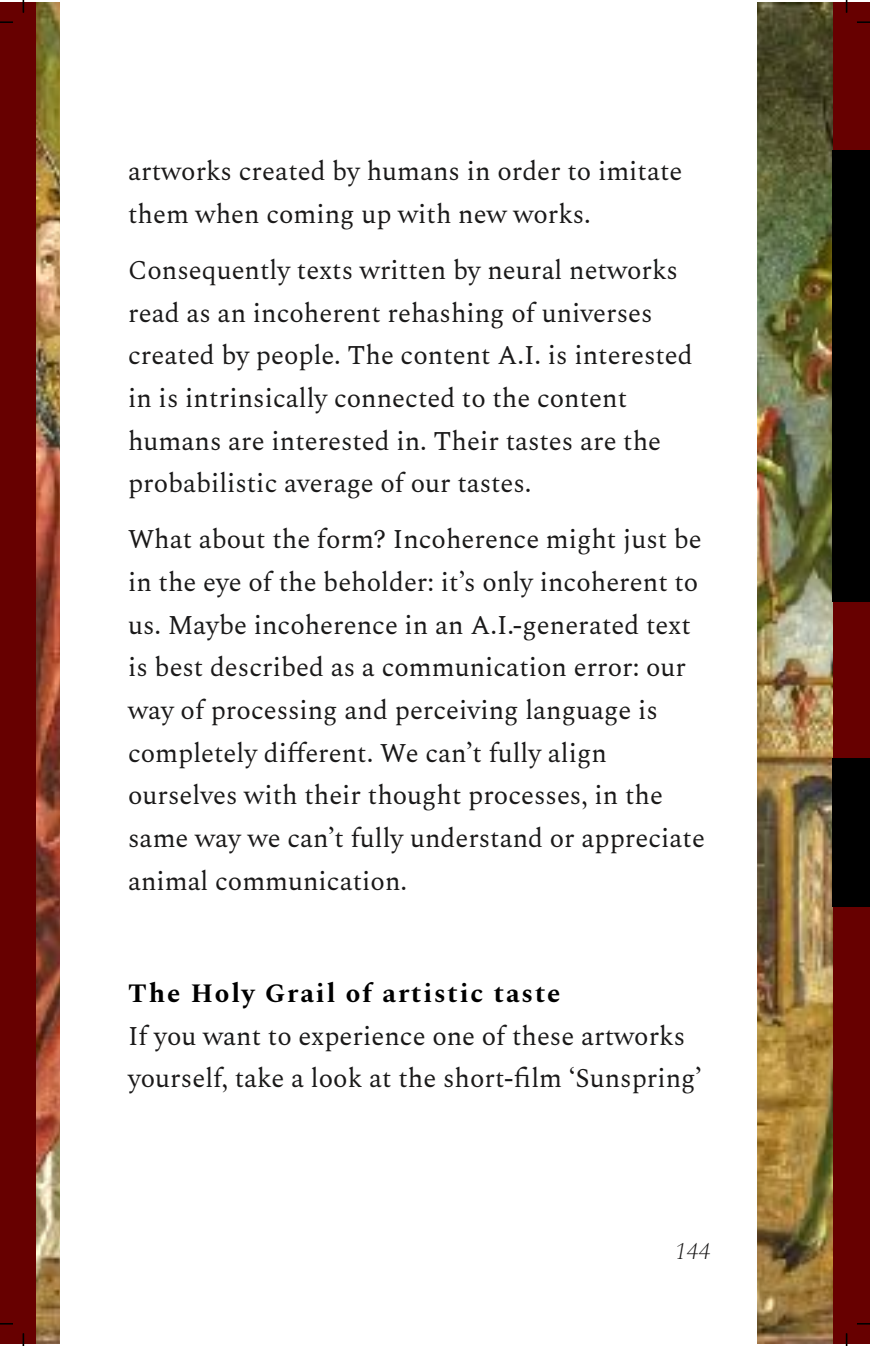


Don't ask them, watch them do it...

So far I have tried to use the human mechanism of liking something in order to speculate on what an A.I. might like, and I have passed a machine the metaphorical microphone. Now I will take a side alley: exploring the works of art created by A.I already. If we find out what they produce as artists, we might come closer to what they appreciate in art.

These electronic artists already made a lot of high-profile work: last year an A.I.-made painting sold for \$432,500.⁴⁰ But when looking at all these works - from Harry Potter chapters to sci-fi movies - we encounter the same issue: everything is a deeply subjective reflection of our own art, our own worlds and aesthetic tastes. Unlike AlphaGo, a system that can train itself to come up with highly creative solutions, art-creating A.I. can't learn to create 'better' art, because there are no objective parameters to adhere to. It is trained to analyse patterns in

⁴⁰ <https://www.christies.com/features/A-collaboration-between-two-artists-one-human-one-a-machine-9332-1.aspx>




artworks created by humans in order to imitate them when coming up with new works.

Consequently texts written by neural networks read as an incoherent rehashing of universes created by people. The content A.I. is interested in is intrinsically connected to the content humans are interested in. Their tastes are the probabilistic average of our tastes.

What about the form? Incoherence might just be in the eye of the beholder: it's only incoherent to us. Maybe incoherence in an A.I.-generated text is best described as a communication error: our way of processing and perceiving language is completely different. We can't fully align ourselves with their thought processes, in the same way we can't fully understand or appreciate animal communication.

The Holy Grail of artistic taste



If you want to experience one of these artworks yourself, take a look at the short-film 'Sunspring'



on youtube⁴¹. A filmmaker was bold enough to produce an A.I. generated sci-fi script, with all the bells and whistles. Because everyone involved takes it so seriously (one of the actresses even bursts into tears during an utterly incomprehensible monologue) watching this film is a surreal experience.

That seems to be the common denominator in all these works: surrealism. Maybe surrealism is what really gets them. Maybe seemingly incoherent comedy is one of the landmarks of the artistic taste of artificial intelligence. And maybe, just maybe, artificial intelligence is an avid lover of Monty Python. Their favourite show has to be one where randomness seems to be the common thread; the catchphrase is literally: “*And now for something completely different.*”. Which might just as well also be the catchphrase of A.I. writers. Yet at the same time, Monty Python is actually very well structured, and the arcs of the individual scenes are linear.

⁴¹ <https://www.youtube.com/watch?v=LY7x2Ihqjmc>




On top of that, the very nature of satire is not completely random. Satire presupposes a certain knowledge about the world and, for example, the understanding of concepts like “shop”, “parrot” and “return-policy”. But the disruptive force by which Monty Python comedically tackles the viewers notions of these concepts is possibly very appealing to A.I..

Deconstructing the artistic taste of A.I. from the viewpoint of their artworks has one major flaw: the presumption that they will be interested in art that is similar to creative works by their own hand, something that is not a given. My work as a performance artist is primarily text-based, yet I thoroughly enjoy visual art - precisely because it is outside my own area of expertise.

Inevitably human


At its core, A.I. is still very human. It is designed by us, learns from our data and functions to serve us. Theoretically it is possible that it will move beyond us in the future. If an A.I. system



autonomously creates a newer A.I. system that autonomously creates a newer A.I. system ad infinitum, all traces of human influence might eventually be eliminated. In music, this principle is used by Alvin Lucier: in a given space, he plays a recording of some music and re-records it. This recording is played and re-recorded again, and so on... up to a hundred times. The idea is that the last recording has almost no traces of the original music, but is a recording of the sound of the space (and recording equipment) itself.

But that is not the state of this technology today. Today it still incorporates our biases, our flaws and our ethics. Precisely because of this, the word 'autonomous' is a little bit deceptive. A self-driving car is never truly autonomous, since it does not decide how to drive itself; it drives according to our rules, receives updates created by humans and - as discussed earlier - has our ethics embedded in it.

Because of its very nature, both art created by and for A.I. will inevitably involve aspects of human art. An A.I. artist will create how we have




told it to create, and an A.I. audience member will experience art the way as we have dictated.

This is something we have to be aware of, even something we should embrace. Being aware of it forces us to be more mindful and critical of these new technologies whilst embracing it will empower us, enabling us to push the man-machine symbiosis to new (artistic) heights.

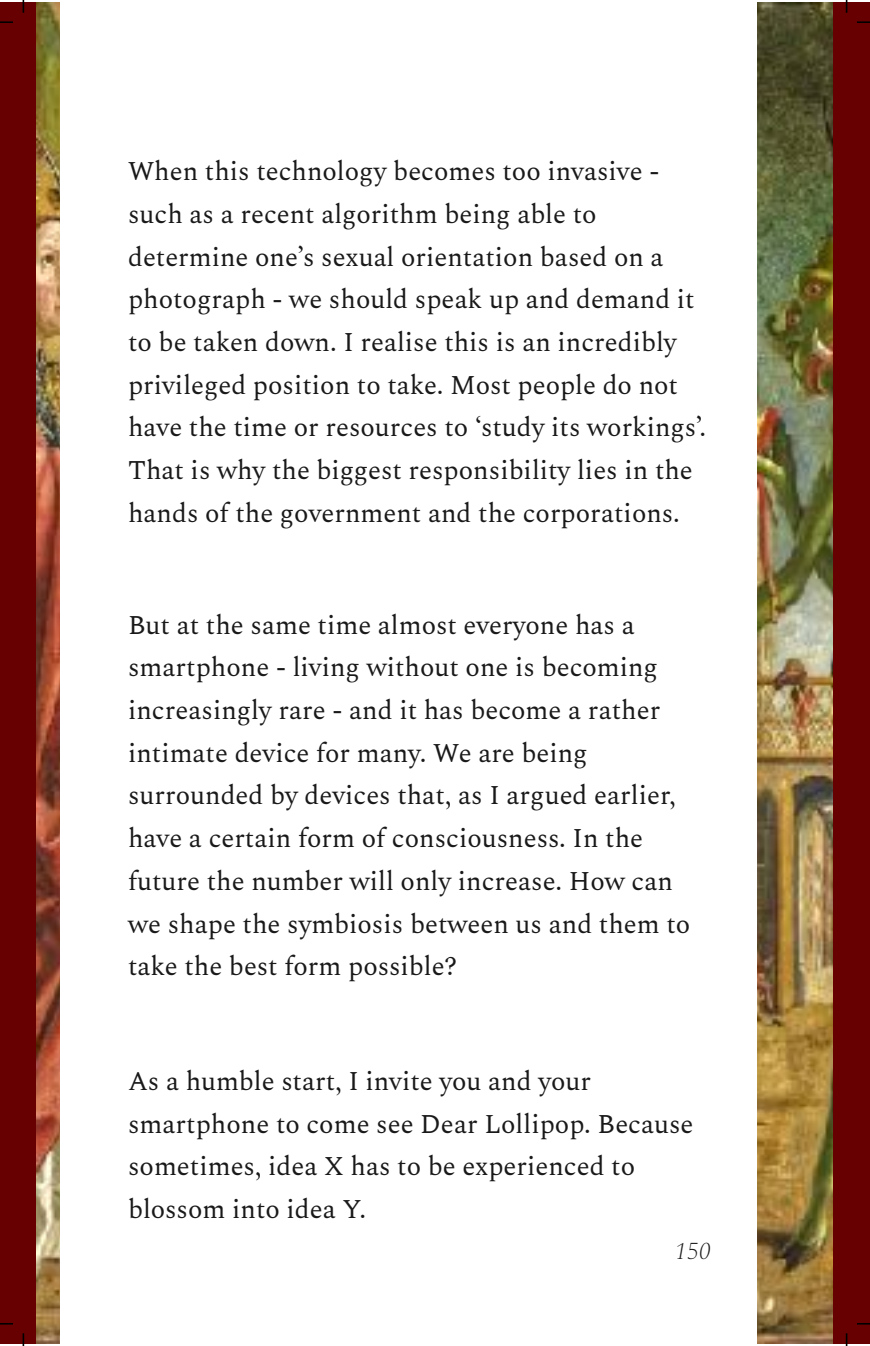
It is not a bad thing that A.I. is inherently human. Actually it is quite beautiful. For thousands of years humanity has used tools as extensions of themselves. This time around it is a little bit different; this exciting, scary, weird technology is a new tool we have at our disposal - but unlike anything we have seen before.

We have to be careful with it, conscientious just like our distant ancestors had to be with a pointed stick. Haphazardly swinging it around might hurt someone, but when wielded with great care, it could prove to be a beneficial tool in gathering food. This analogy might be helpful for people who think A.I. is a scary, weird thing; seeing it as just a tool might help break the ice.



However, for those who are already a little bit further in embracing technology it might seem a little bit shortsighted. A pointed stick really just is a tool to wield, whilst Artificial Intelligence is more than that: it is something we can form a meaningful relationship with. Most of us already have one with their smartphone. We caress them throughout the day, look at them for hours on end and feel uneasy without them.

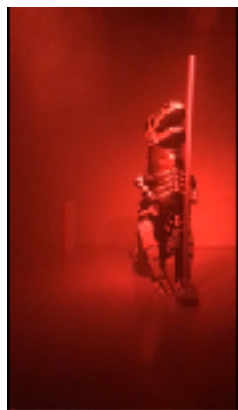
Regarding A.I. as something literally alien to us - as if it is an entity with a different soul, forever beyond our comprehension - is a dangerous line to take. Distancing ourselves wholeheartedly from this technology is a counterproductive and false belief, yet many people feel like they have got nothing to do with A.I. and its new appliances. A more advantageous approach is to study its workings in order to actively engage in ethical discussions about this topic. As with our dogs: we just need to learn how to communicate with them, understand their needs in order to be able to train them.

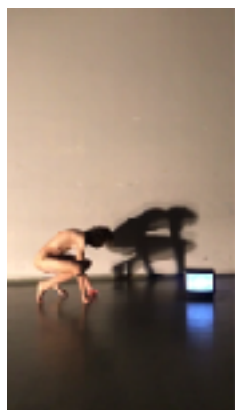



When this technology becomes too invasive - such as a recent algorithm being able to determine one's sexual orientation based on a photograph - we should speak up and demand it to be taken down. I realise this is an incredibly privileged position to take. Most people do not have the time or resources to 'study its workings'. That is why the biggest responsibility lies in the hands of the government and the corporations.

But at the same time almost everyone has a smartphone - living without one is becoming increasingly rare - and it has become a rather intimate device for many. We are being surrounded by devices that, as I argued earlier, have a certain form of consciousness. In the future the number will only increase. How can we shape the symbiosis between us and them to take the best form possible?

As a humble start, I invite you and your smartphone to come see Dear Lollipop. Because sometimes, idea X has to be experienced to blossom into idea Y.







references...

Allen, P. (2011). *The singularity isn't near.*

Retrieved May 23, 2017, from

<https://www.technologyreview.com/s/425733/paul-allen-the-singularity-isnt-near/>

Barrett, L.F. (2018). *How Emotions are made: The Secret Life of the Brain.*

Bostrom, N. (2014). *Superintelligence: paths, dangers, strategies.*

Bridle, J. (2018). *New Dark Age: Technology and the End of the Future.*

Dissanayake, E. (1998). *Komar and Melamid Discover Pleistocene Taste*

Dennett, D. (2017). *From Bacteria to Back and Back: The Evolution of Minds.*

Domingos, P. (2015). *The master algorithm.*

Eagleman, D. and Brandt, A. (2017). *The Runaway Species: How Human Creativity Remakes the World.*

Eagleman, D. (2015). *The Brain: The Story of You.*

Eagleman, D. (2011). *Incognito: The Secret Lives of the Brain.*

Godfrey-Smith, P. (2016). *Other Minds: The Octopus, the Sea, and the Deep Origins of Consciousness.*


Hanson, R. (2016). *The Age of Em: Work, Love, and Life when Robots Rule the Earth.*

Harari, N.Y. (2018). *21 Lessons for the 21st century.*

Harari, N.Y. (2015). *Homo Deus. A brief history of tomorrow.*

Harari, N.Y. (2014). *Sapiens: A Brief History of Humankind*

Kahneman, Daniel. (2011). *Thinking, Fast and Slow.*



Kasparov, G. (2017). *Deep Thinking: Where Machine Intelligence Ends and Human Creativity Begins*.

Mlodinow, L. (2012). *Subliminal: The Revolution of the New Unconscious and What it Teaches Us About Ourselves*.

O'Connell, M. (2017). *To be a machine*.

O'Neill, C. (2016). *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*.

Tegmark, M. (2017). *Life 3.0: Being Human in the Age of Artificial Intelligence*.


Turing, A. (1950). *Computing Machinery and Intelligence*.

Urban, T. (2015). *The AI revolution: the road to superintelligence*.

Retrieved May 15, 2017, from
<http://waitbutwhy.com/2015/01/artificial-intelligence-revolution-1.html>

Urban, T. (2015). *The AI revolution: our immortality or extinction*.

Retrieved May 15, 2017, from
<http://waitbutwhy.com/2015/01/artificial-intelligence-revolution-2.html>



Urban, T. (2017). *Neuralink and the brain's magical future.*

Retrieved May 27, 2017, from
<http://waitbutwhy.com/2017/04/neuralink.html>

de Waal, F. (2016). *Are We Smart Enough to Know How Smart Animals Are?*

Walker, M. (2018). *Why We Sleep: The New Science of Sleep and Dreams.*





'Dear Lollipop'

a live performance for smartphones

by

*Tara Barendrecht, Just van Bommel, Boris De Klerk,
Tabor Idema, Cor Langerak, Junior Manse
and Nina Wilson.*

Coaching and direction

by Max Wind and Casper Wortmann.

based on a concept by Peter Missotten

App design by Sylvain Vriens - www.sylvain.nl

*'Dear Lollipop' is an artistic research project by the
professorship Technology Driven Art ZUYD - Peter
Missotten and Maastricht Academy of Performative Arts,
Coproducted by IMPACT Interreg V-A
World premiere nov. 12, 2019 at CC Hasselt Belgium*

illustrations:

Saint Wolfgang and the Devil (1471-1475)

by Michael Pacher, Alte Pinakothek, Munich

Prinz by Casper Wortmann © 2016

***Stills from a smartphone's view on rehearsals
of 'Dear Lollipop' © 2019***

Design and text editing by Peter Missotten

www.technologydrivenart.org



Saint Wolfgang and the Devil (1471-1475)
by Michael Pacher, Alte Pinakothek, Munich

**ZU
YD**

The presence of Artificially Intelligent machines is gaining ground in our daily lives. They seem to be able to do anything for us.

The possibilities are endless.

But what can we do for them? Perhaps these machines possess some form of consciousness, and if that's the case, then maybe we should prepare for a reversal of roles. Maybe we should begin by performing for them. What if A.I. is the audience of the future? And what would they like?

These questions launched an artistic research project by the professorship Technology Driven Art at the Zuyd Arts faculty, resulting in a live performance for smartphones only: 'Dear Lollipop'.

In his first book ever, the Dutch actor/performer Max Wind goes on a quest for the Artistic Taste of Artificially Intelligent machines.

