Simulated Space
Animal, Man, Machine

Simulated Space
Animal, Man, Machine

Autonomous robots move through our public and private spaces, transporting people, moving around goods, sorting out objects, harvesting crops, cleaning, surveilling, following, tracing and observing. To travel the world, these machines need to map and model our spaces instantaneously. As they integrate in our distribution networks and inhabit a permanent position within our infrastructure, *the ways* in which they build internal models of our spaces become highly relevant to our own experience of space, since these environments are in turn transformed and optimized for robot perception. Future landscapes will be inhabited and modified by both biological life and autonomous machines. But today most territories assigned to machines have been altered to fit their needs.

Human Exclusion Zones

The past few years have been marked by the spawning of new kinds of spatial environments that are built to accommodate non-human agents. Automated production spaces are in full operation today in the form of advanced agriculture, computer controlled harbor transshipment areas or automated storage facilities. There is a fleet of robots (all named Betty) that operate inside an Amazon warehouse, picking up orders at 1.3m/sec. The area in which they work is dubbed as ‘human exclusion zone’. Within the Human Exclusion Zone, agents stick to a 24/7 working cycle, working with 99.99% accuracy with only a 5 minute break every hour to recharge their batteries. The architecture of this warehouse accommodates machines — populated by endless storage racks, barcodes, recharging stations. Addressability is key: every spatial element should be mapped and registered in order for the machine to interact properly. Some already proclaim that future windowsills will need their own individual addresses.[1] A quantified architecture for machines.

The mechanical choreography of automated space does not account for human subjects. Human nature imposes limits on efficiency, and therefore is kept outside fences or within narrow corridors, determined by green colored lines on the factory floors marking out safety zones. The Human Exclusion Zone’s architecture can make spaces impassable by either being too small, too dangerous, or even illegible. Inside fully automated ‘Dark Factories’ industrial robots work silently in pitch black: fully mechanised labor does not require any light.[2] Excluding light — and therefore people — from a worksite might be a cost-saving reflex, a useful tactic for upscaling industrial production processes. But when autonomous robots further integrate into society we will have to think about the design of an environment that is inhabited by both humans, animals *and* autonomous robots.

Perhaps the real nightmare, even worse than the one in which the Big Machine wants to kill you is the one in which it sees you as irrelevant, or not even as a discrete thing to know. Worse than being seen as an enemy is not being seen at all. Perhaps it is that what we really fear about AI.'

Benjamin Bratton - *Outing Artificial Intelligence: Reckoning with Turing Tests, in 'Alleys of your Mind'*

An inclusive environment is based on mutual recognition: seeing and being seen. A truly inclusive environment should benefit all of these agents: machines, humans, animals, a complex new ecology that requires a different attitude towards design that is not only centered on one type of (human) user, but many. The triangular relationship of machine, man and animal will provide the fundament for this text to investigate the complex connection between intelligence and space, opening up new ways to formulate the conditions for inclusive public landscapes.

Architecture of compression

Robots inhabit many different ecosystems. Some robots are screwed to the factory floor, some robots travel the world at high speed, others perform at landscape scale. The way they affect the accessibility of the spaces they occupy differs greatly. Within the Human Exclusion Zone, physical space is organised to minimize emptiness *around* the machine. Because of the machine's motorized accuracy, there is no need for much surplus 'safety' space around travel paths. The result is an architecture of compression, maximizing productivity, minimizing emptiness. Compression marks the promise of technological innovation: efficiency and innovation through the minimization of the use of energy, space or time. The development of autonomous robots was boosted by different forms of 'compression': the shrinkage of sensors to invisible sizes, more efficient algorithms to process the incoming data (the compression of time), increased computational power on smaller circuitry, smaller batteries (compressed energy) and improved mechanical robotics.

This architecture of compression has a landscape-scale effect: it redefines our notion of the rural and remote, since compressed infrastructural networks modify our experience of distance (shorter travel time, increased bandwidth, improved connectivity). The introduction of autonomous robots signifies another phase in this project of compression. Human Exclusion Zones stage an extraordinary concentrated infrastructural activity in which countless agents are trapped in continuous motion, an uninterrupted flux of productivity. This flux greatly affects to way we can access these spaces. What will might the implications be of further compression in different domains, other than industries of production and transportation? The integration of AI within public space requires many changes. AI has to adapt to it's environment, but in reverse the environment can also adapt to AI. Today, test drivers of self-driving cars have been dressing up as car seats to study how people react on an empty vehicle driving around. Meanwhile, there have been claims for the redesign of traffic signs, road signs and enhanced traffic light systems, and the irregular checkered patterns of QR codes have become the predominant visual component of distribution centers.

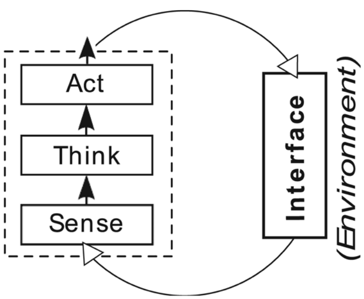

Fig.1

In short, the architecture of compression has had a great visual impact on certain types of spaces. The logic behind these changes can be traced back to the sensing technologies employed by autonomous robots, and the way these robots process the incoming data. When we want to study the way intelligent autonomous agents behave in space, we can focus on three aspects derived from the early paradigms of robotics: *sensing*, *thinking* and *acting*. This text is based on the structure of this cycle: listing the many types of sensors, then diving into the structure of AI, and finally investigating what kind of possible behavior emerges from what kind of intelligence.

Sensing
Models and Modalities

How do machines sense the world, and how do they interpret the frequent signals and attributes and qualities of our spaces? How does this differ from our human perception of space? Understanding the way machine vision systems translate our common reality into an internal model (how space is computed and simulated) is key for solving the design problem of both robot and human/animal-inclusive environments. If we understand how these 'internal models' are constructed, we can deduce what combinations of spaces and sensors result in what type of emergent behavior - giving us insight into a possible common ground between machine vision and human vision.

> Even humble motile bacteria, as they swim up the gradient of nutrients toward the point of maximum concentration, can be said to have an internal model of their environment. This internal model is not, of curse, based on representations of any kind and it is not a model of the world at large but only of the concrete opportunities and risks afforded to bacteria by their immediate surroundings. In other words, from the beginning of life the internal

> models mediating the interaction between a primitive sensory system and a motor apparatus evolved in ration to what was directly relevant or significant to living beings.
>
> Manuel DeLanda - *Philosophy and Simulation, the Emergence of Synthetic Reason, p.80*

Multicellular organisms, animals, human beings and robots need to be aware of the space around them *and* of their own bodies in order to navigate through it. As DeLanda points out, even the simplest organisms without central nervous systems can show quite complex behavior. They are able to monitor and model their environment. They react to certain changes in their surroundings (sensory input) and change their movements accordingly (motorized output). But an animal displaying intelligent *behavior* does not imply it has great intelligence. The same goes for autonomous robots: they may display intelligent behavior, but this does not necessarily prove it's intelligence.

Different organisms or agents sense spaces in different ways. For us, human beings, *seeing* space might be our primary way of interpreting and interacting with our direct environment. But for other animals different modalities like sound, reverberations, touch and smell are equally useful in constructing an internal model of the world around. Similarly, robots directly construct 3D point clouds by sending out laser pulses (LIDAR), analyse thermal data or register infrared frequencies. All are equally valid instruments to map the world with. The sensory apparatus of an autonomous agent determines *how* it exchanges information with its environment.

Sensory Ecology

Because an inclusive public space is largely a matter of *sensing* and *being sensed*, I propose to approach this future urban environment (occupied by human beings, other living organisms and machines) from the perspective of *sensory ecology*. Sensory ecology is a relatively new field in biology which focuses on the exchange of information within an ecological system: what do organisms know and how do they obtain this knowledge? How have sensory systems and the exchange of certain types of information influenced large-scale evolutionary change and speciation?

> Yet, even organisms living in the same place, at the same time, inhabit different worlds: they live in different sensory environments, bounded by the properties of their sensory organs. For example, a bee using colour vision to search for flowers may be right next to a snake waiting to detect the infrared cues of its prey or an ant following chemical pheromone trails to food. Animals should, and do, only pay attention to features of the habitat that are important to them. What they can detect and respond to is dictated by their sensory systems and how they

constructed over evolution. Animals can be almost touching each other in space, but be worlds apart in perceptual terms.

Martin Stevens - *Sensory Ecology, Behavior & Evolution, p.127*

Similarly to the bee, snake and ant, the machine and human work alongside each other without necessarily sharing the same perceptual experience of space. On top of that, both agents profit from sensory integration: it is possible to obtain information from two or more senses, making the incoming information more plausible and accurate, and enriching the internal model with more than one type of data. For example, it would make a lot of sense to equip autonomous vehicles with sensors that can detect small differences in temperature, enabling it to distinguish objects from bodies and ice from water. Mapping thermal imaging on top of a model constructed through a ‘normal’ camera creates a much richer and adequate simulation of any environment. But, although it makes it easier for the machine to sense the human body, the way in which it does differs greatly from how the human body distinguishes other bodies.

The machine can be equipped with many species of sensors, enhancing sensory integration. Just to name a few that are installed on autonomous robots today, think of lidar sensors, sound sensors, tactile sensors, various motion detector systems, proximity sensors, thermal sensors, infrared cameras, standard RGB cameras, stereo vision sensors, dynamic vision sensors or event based cameras, gyroscopes, vehicle speed sensors, bluetooth sensors, WiFi signal sensors and GPS trackers. Some of these sensors generate 3D data instantly, others record 2D imagery that is interpreted and translated into a spatial model through computation, such as simultaneous localisation and mapping (SLAM) algorithms that translate a live video feed into a 3D model. Simulated space does not necessarily exist as point clouds or polygons, it can come in many types of information simultaneously. Simulated space can be a database providing insight in the health of an individual tree in an orchard, the speed and direction of surrounding traffic, or the number of warm bodies on a crash site.

Vision

As we can infer from the list, a lot of sensors tap into information about the world that our bodies have no direct access to, like the infrared spectrum, laser detection or WiFi signal sensing. Without a visual model mediating this information, we could never perceive it. But, even when we consider simple visual sensors (RGB camera’s) that you would expect to process information similar to human vision, we will see that the way in which AI interprets the incoming data is extremely different from how the human brain treats visual information (we will

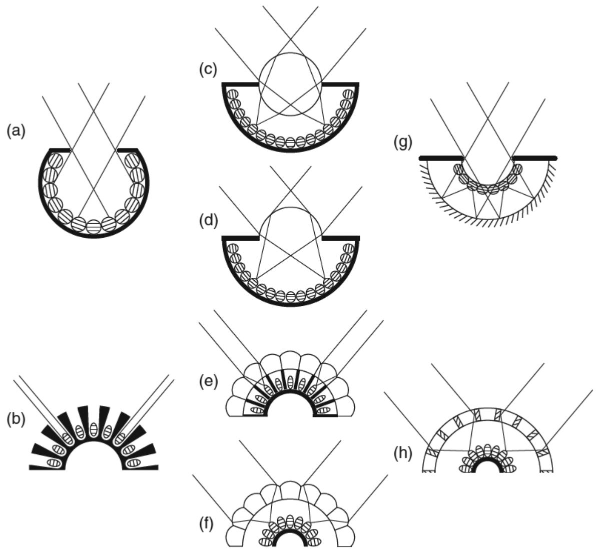

Fig.2

expand on this disparity later on in the text, when discussing image recognition algorithms). Therefore it is time to reconsider the position of vision in the overall spectrum of sensing mechanisms.

> As humans we tend to think of vision as a general-purpose sense, supplying us with any kind of information we require. For most other animals this is not so. Predators and prey, for example, have different visual requirements: foxes and rabbits have different eyes and different visual systems, as have dragonflies and mosquitoes. Similarly, a sedentary clam lives in a different world from a flying insect, and the optical requirements are quite different.
>
> Land & Nillson - *Animal Eyes*

The human eye is a complex sensory organ that mediates most of our day-to-day experience of reality. Eyes have been developed along a wide variety of evolutionary paths. Some organisms can only detect the presence of an ambient light source (like the sun or moon), more advanced sensory apparatuses result in more complicated relationships between animals and their environment, as Land & Nillson classified in their analysis on the evolution of the eyes:

> There are a number of important steps along this evolutionary path, and these lead to a classification of light-controlled behaviors into four basic classes:
>
> 1. Behaviours controlled by non-directional monitoring of ambient light. Examples are the control of circadian rhythms, light-avoidance responses for protection against harmful levels of short wavelength light, shadow responses to avoid predation, and surface detection for burrowing animals.
>
> 2. Behaviours based on directional light sensitivity. Examples are phototaxis, control of body posture (optical statocysts), and alarm responses for approaching predators.
>
> 3. Visual tasks based on low spatial resolution. Examples are detection of selfmotion, object avoidance responses (anti-collision), habitat selection, and orientation to coarse landmarks or major celestial objects such as the sun or moon.

4. Visual tasks based on high spatial resolution. Examples are detection and pursuit of prey, predator detection and evasion, mate detection and evaluation, orientation to fine landmarks, visual communication, and recognition of individuals.

Land & Nillson - *Animal Eyes*

Through these classifications we can consider how more visual refined preceptory systems allowed for more complex behavior to emerge, a key element in sensory ecology. Coevolution through predator-prey relationships could only emerge with two species able to either *recognize prey and chase* or *recognize predator and flee* (both requiring high spatial resolution visual systems). Thus the way in which agents perceive space determines the relationships it will engage in with other agents. Now let's go deeper into analysing machine vision and the differences between the human perception of space. Along the way, we will also begin see the logic behind the transformations of spaces occupied by machines today.

Distortion & Grip

Autonomous robots that travel space usually construct internal 3D models from a 2D data feed. To build this simulated space they need texture, shapes, objects to orientate themselves. The algorithms designed to carry out this operation are trained to recognise edges and corners of objects, and by calculating the movements of these points from one video-frame to another, they construct a spatial model. As a result, the algorithm performs better in a space with a lot of clutter than in an empty room. A white cube is the algorithm's nightmare, it does not provide enough information to orientate itself. Pasting some QR codes on the walls would do the trick. The machine suffers from a threshold of information input — without the right amount of data, it is not able to position itself in space.

Here we see a first divide between human and machine perception: what may deceive the human experience of space — colour, patterns, textures, reliefs — is actually helpful for the algorithm. These qualities provide grip. For example, a vertically striped pattern on a wall may make a space look taller to us, but for the machine it provides more edges and corners to process, more information to position in space and to relate itself to. The space becomes more distinct and easier to interpret for the machine. For us, texture distorts space. For the machine, texture means grip.

Color

There is hardly any inherent quality of the world more deceptive than the sensation of color. Not only are all colors permanently influenced by the context in which we see them (the effect of simultaneous contrast), also the color receptors in our eyes are influenced by the intensity of light, making us shift between different modes of perception throughout the day (the Purkinje effect). At night, we shift to different color receptors that enhance our sensitivity for dark and light contrasts, but at the cost of a more desaturated picture of the world. Our sensibility for color is affected by external factors all the time. Next to that, different colors have different spatial effects: far away objects shift towards a more desaturated and bluish hue. Therefore blue is a color that appears to recede. When we see complementary colors like orange-blue and purple-yellow, the brighter shade appears to come more toward us, and the darker one appears to recede.

Of course, certain camera's are light-sensitive too, but for machine vision algorithms, colors appear as RGB data on pixel-level. The difficulty for machines though is to be able to recognize the shift that colors undergo through different lightning conditions. Human beings have an adaptive mechanism that compensates for that (color constancy), so when we see a white piece of paper in yellow lightning or under blue moonlight, we will still recognise it as a white piece of paper. For the machine, this is a hard thing to do. It is necessary for the algorithms to map out all possible hues (a possibility space of RGB data), that one object can take on under any possible light. Once the algorithm has an understanding of the fluctuations of an object, it is able to track and sense colors more accurately.

Low-res emptiness

The human experience of space equals the experience of emptiness which has the potential of being occupied with different kinds of matter and activities. When we say someone has a spacious home or office, we mean there is a lot of empty space to move around in. How does a machine interpret *empty* space? In terms of programming/data, empty space is a problem. Walls, corners and other objects in a space are interesting, and need to mapped to be avoided while traveling through that space. Empty space is not interesting. Mapping empty space with the same resolution used for 'interesting objects' could quickly fill up the internal memory of the system. A solution is rasterization: dividing space into boxes, splitting into more boxes when objects are detected. This way, empty space will be a big 'box', a low-resolution void. Interesting areas are rendered in higher resolution. Low resolution emptiness saves memory.

Ambient Sensing

The tricky part is that the assemblage of effectors, sensors and processors that makes up an autonomous robot can be scrambled and organised in different ways. Human beings usually only have two eyes to sense depth, but the 'eyes' of a robot can be anywhere, even a hundred kilometers away from it's central processing unit. Many sensors can be distributed throughout space with one central body to read the data. In this way the autonomous agent turns into a spatial being itself. A sensory apparatus with a X, Y and Z axis. We will cover different applications of ambient sensing later on when discussing landscape-scale surveillance systems.

The great discrepancy between human and machine vision (what seems easy for us, fools them and the other way around), demonstrates that it is difficult to claim that machines ‘see’ space, and that is very hard to indicate any common ground. Machines have an entirely different perception of space. They register and model space with a much greater sense of detail and almost perfect accuracy, yet, they have trouble with plain emptiness, color changes and bare walls. With these findings we could also sketch out a space in which the machine would perform well: it would be cluttered with different sized and shaped objects painted in high contrasted colors, creating a rich visual landscape by enhancing the difference between foreground (object) and background (wall, floor, ceiling). In short, an environment we would experience as quite chaotic, even psychedelic, dense, and aesthetically unpleasant.

In order to get to a broader understanding of the sensory ecology that spawns from the interaction between autonomous robots and humans, analysing the way it computes space does not suffice, it is necessary to focus on how robots recognise objects (fellow robots, animals, human beings) within that space. To get to the bottom of these mechanisms we will have to expand on the topic of artificial intelligence and image processing, which will also reveal surprising overlaps with visual processing in the animal kingdom.

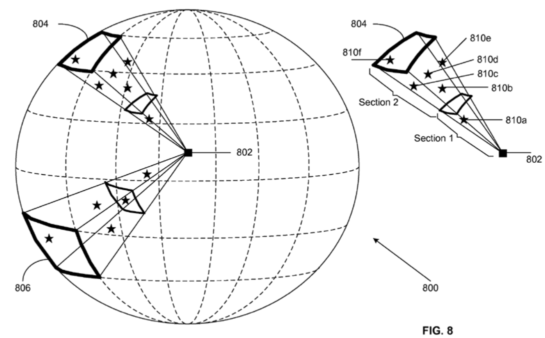

Fig.3

‘Thinking’
Artificial Neural Networks

I put thinking between quotation marks because it does not cover the subject entirely: *thinking* only applies to human intelligence, it is an act involving a human body, a brain, neurons which can be stimulated by hormones, influenced by flows of neurotransmitters. Machines don’t have brains, they have computational power, they write ones and zeroes. We might need a neologism here: a verb for the application of artificial *and* human intelligence, covering both *thinking* and *computing*.

The human brain is able to recognize a cat instantly, whether the cat is blue, it is situated inside a space station or at home on the couch, or if it is the size of a house. For AI, this is not an easy task. A computational model that is capable of proper object recognition in images is the Artificial Neural Network, which consists of multiple layers of artificial neurons (nerve cells), that represents biological neural networks such as the mammalian brain. I want to stress the word *represents* here because these artificial models come nowhere near the complex intelligence of animals and humans. The human brain contains approximately 85 billion interconnected neurons, which all individually fire out information simultaneously, which is not possible with any form of computation today. Now how do these artificial networks work?

The individual neurons pass on information, just like nerve cells in the human body. They have the capacity to measure an incoming signal, and at a certain threshold, send an output signal to other neurons further down the chain. The network is organised in layers, with an input layer at the ‘front’, several ‘hidden’ layers, and an output layer at the end. Image classification algorithms build an internal model from the bottom up: first the neural network ‘searches’ the pixel data for edges and corners of shapes, which it can distinguish because of high-contrast values between individual pixels. Once it has distinguished an edge, one neuron might fire to a neuron in the next layer saying *‘edge detected’*. One layer on top of that, the network distinguishes shapes. And finally, some layers further down the chain, it is able to recognize a cat. Every layer creates advanced generalisations of the previous one, and so it builds up a classification system through this very simple mechanism of interconnected neurons.

When we start setting up a network, the connections between the neurons are completely random. We create a working algorithm by ‘training’ the network. Say if we want to learn the algorithm to distinguish cats in images, we can ‘train’ it by showing millions of images of cats (through the input layer), and telling it that these are cats at the output layer. The internal network of hidden layers will be structured (by adjusting the ‘weights’,

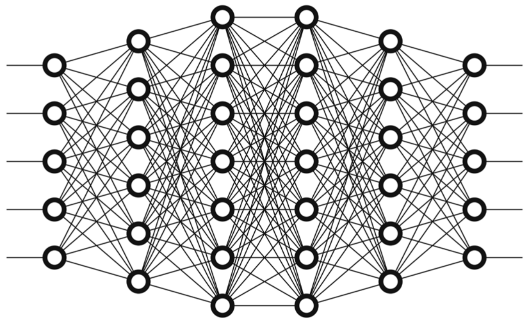

Fig.4

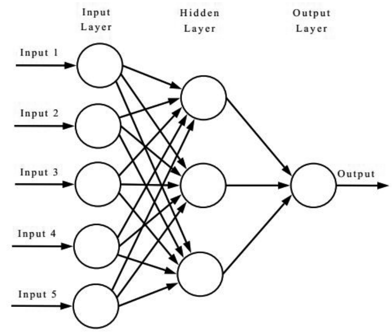

Fig.5

the strength of connections between individual neurons) so that it creates some kind of internal model representing what it believes to be a cat.

The complex network of connections between several hidden layers of these neurons has advanced abilities in recognising patterns in big chunks of data. This is why these networks work so well for recognising clusters of pixels resembling cats. In fact, image classification algorithms have already surpassed humans in their accuracy to recognise for example different types of dog breeds on pictures. This made researchers question if there are any differences between the capabilities of machine vision and human vision. But notice that through the process the algorithm does not see in any way we see, it just processes pixel data (it does not have photoreceptors that react to light that has entered through a lens etc.). Also, recent discoveries in how these algorithms can be tricked give insight in how different these neural networks are from the way mammalian brains process images.

Fooling algorithms and animals

Despite their supreme accuracy in pattern recognition, image recognition algorithms can be easily fooled by very subtle adjustments to an image. Simple inconspicuous noise patterns overlaid on a picture, almost invisible to the human eye, can disorient the algorithm dramatically and cause it to classify objects incorrectly. It was necessary to discuss the inner workings of neural networks in such detail earlier to allow us to the bottom of this spoofing mechanism. These noise patterns, named *universal adversarial perturbations*, give great insight in how different humans and machines recognize objects in pictures. The human brain is able to *generalize*, and in this way capable of distinguishing a layer of subtle noise from the actual picture, but the algorithm starts from the bottom up, with the basic patterns of edges and corners. But all that information is slightly distorted and interrupted by the noise. Because it is disoriented at this base-level, the upper layers in the network classify the image as something unintelligible for us, yet it makes sense to the machine. Researchers found out these perturbations work on many different algorithms,

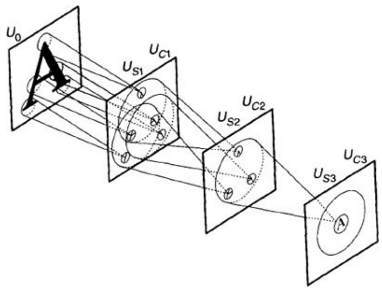

Fig.6

even ones that were trained on entirely different datasets.[3] Things got even more scary when a team of AI researchers showed that there exist imperceptible adversarial perturbations that cloak almost all pixels of a target class while leaving the other classes across the image nearly unchanged, which basically means that you can blend out targeted object classes like 'humans' out of a picture.[4]

The last few years AI developers were surprised by how image recognition algorithms also appeared to be able to classify images of sheer white noise or (what appears to us as) random abstract patterns with 99% confidence. This kind of 'spoofing' is quite the opposite from the tricky noise filters of the *universal adversarial perturbations*. For our human eyes there is no way to see what the algorithm claims to see. Researches dubbed these deceptive patterns as *fooling images*.[5] Again, the convergence of all of these recognition systems is surprising: the same abstract blobs and shapes are classified as a specific object (for example baseballs, trucks, lizards) by entirely different algorithms. When you study the aesthetics of these fooling images and patterns, you can see what might trigger the algorithm: the rastered pattern of knobs of a remote, the colors and striped linage of a school bus. It looks as if subtle aspects of these objects are enhanced and elevated to a psychedelic display of colors and lines. One can draw a parallel to the biological term *supernormal stimuli*, which describes the psychological effect of exaggerated stimuli (brighter colors, bigger shapes) on animals.

For example, some birds will preferentially incubate eggs much larger than those naturally occurring (see Tinbergen 1951). Some male butterflies prefer to court females with mimicking stimuli at much higher rates than the natural wingbeat frequencies of females. Rowland (1989) has shown that female sticklebacks (Gasterosteus aculeatus) presented with models of males prefer larger models even though they exceed the size of the largest natural males by 25%. In addition, some female birds have been shown to prefer songs significantly longer than normal male songs (Neubauer 1999), and some male moths are more attracted to super-normal synthetic pheromones over the naturally occurring female substance (Jaffe et al. 2007). Thus preference for supernormal stimuli may be wide-spread across modalities.

Martin Stevens - *Sensory Ecology, Behavior & Evolution (2013), p.127*

The interesting thing about the *fooling images* and supernormal stimuli as a psychological effect, is that it elicits a response from a cue that seems random or abstract to us. The problem, which at the same time illustrates the potential, of the vulnerabilities of image recognition software is that the outer visual layer of anything makes the identity of any object inherently negotiable.

Invisible in plain sight

The human brain has the special ability to adapt to novel situations and to recognise anything as an object even though it hasn't seen it ever before. We can process new information without much effort. This is all very different for the machine's AI and the animal brain. Volvo's 'Large Animal Detection system' developed for self-driving vehicles had trouble with placing kangaroo's in space, because the animal's distinctive hopping movements confounds the algorithm.[6] There is an interesting parallel to similar 'flaws' in animal recognition in the animal kingdom: toads don't recognize worms moving sideways.

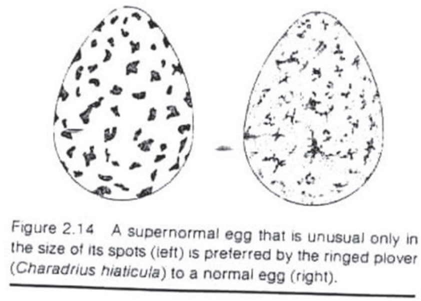

Fig.7

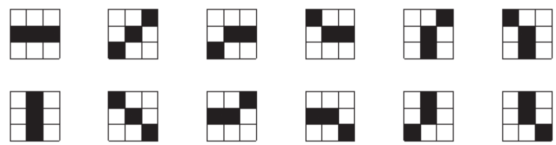

Fig.8

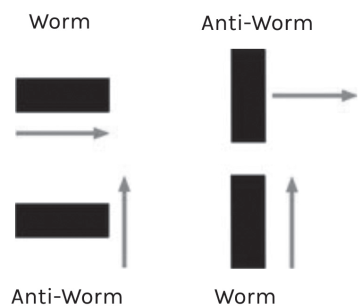

Fig.9

Toads respond with prey-capture behaviour to a small horizontally elongated object moving from left to right in front of it (i.e. an object elongated parallel to the direction of movement). However, they do not attempt to capture objects elongated in the horizontal plane (i.e. elongated perpendicular to the direction of movement). They therefore have a visual system that encodes prey-like objects such as worms and insects.

Martin Stevens - *Sensory Ecology, Behavior & Evolution (2013), p.44*

Animals that display certain type of movements that do not trigger any ganglions (nerve cell clusters) within the toad's brain are not perceived, so they simply do not exists for the organism. The worm should practice being invisible by moving differently. Here we see the powerful relationship between sensing and acting, when 'thinking' is not really present. Therefore thinking is a possible way, not a condition, for displaying intelligent behavior.

In general, after dissecting the intricate relationship between sensing and thinking, I believe it is important to revise the position of vision as a general purpose sense within the sensory ecology of man, animal and machine. There should probably be more room for sensory integration in the design of communication systems that serve a multiplicity of users (for example, think of an olfactory signal warning birds for Amazon's hive-like drone towers, or similarly RFID tags implanted into trees warning delivery drones for bird nests).

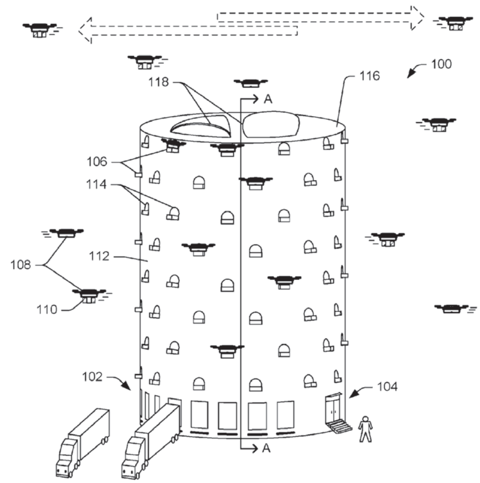

Fig.10

Acting
Applied Intelligence

What does it mean for an autonomous robot to act: to affect the environment it inhabits? On the one hand, we manufacture spatial conditions that accommodate robots, on the other hand, robots redistribute matter and recompose landscapes — welcome to the post-anthropocene (or *machinascene*). When thinking about acting within the world and affecting others, we have to determine what we consider to be intelligent. This sounds like a wish for some kind of robot-ethics: what do we consider to be a good robot? And what spatial conditions are necessary for the good robot to perform well?

Our only way to judge the intelligence of another agent, whether animal or machine, is by studying the consequences of its actions. Here it is dangerous to fall into the trap of the ethological anthropomorphic bias: humans have a tendency to ascribe intelligence to a creature when it displays something close to human behavior: an octopus solving a maze, a dog responding to our language, an elephant recognizing an old human friend. Intelligence is a much more complicated concept that cannot be measured by how much it approximates human behavior. As we have seen, image classification systems are superior to human beings when it comes to recognising certain patterns in a great number images.

But they can be easily tricked and have no way of compensating for these errors. Therefore it is more fair to contextualise intelligence within local (spatial) conditions.

> It is hard to draw the line at what is intelligence, and what is environmental interaction. In a sense it does not really matter which is which, as all intelligent systems must be situated in some world or other if they are to be useful entities. The key idea from intelligence is: Intelligence is determined by the dynamics of interaction with the world.
>
> Rodney Brooks - *Intelligence without Reason (1991), p.17*

The comparison between animal intelligence and artificial intelligence is nothing new, already in 1993 McFarland & Bösser called for a niche-based approach 'closely related to robot ecology'.

> First, behavior requires a body. Disembodied behavior is not possible. Both animals and robots have bodies that are capable of intelligent behavior, and can influence the world around them. An intelligent computer that has no body and that cannot influence its environment, is not capable of intelligent behavior.
>
> The second principle is that only the consequences of behavior can be called intelligent. Behavior is intelligent

> only by virtue of its effect on the environment. The consequences of behavior are due to both the behavior itself and the environment that it influences. Intelligent behavior is behavior of which the consequences of which are judged to be intelligent.

and:

> Behavior is intelligent only by virtue of its effect on the surrounding environment. Intelligent behavior requires judgement, and the consequences of the behavior must be judged in relation to some criteria of intelligence. The criteria of intelligence for animals relate to natural selection, whereas those for robots relate to market forces.
>
> McFarland & Bösser - *Intelligent Behaviour in Animals and Robots (1993), p.71*

The arguments of McFarland and Bösser derive from a capitalist framework of cost efficiency and labour competition: robots will have to compete with (or replace) human beings, so they will have to perform better with less costs involved, similar to how animals have to adapt to their environment and exploit their niche. But when we extend this logic of niche exploitation, we will figure it will be likely that robots will start to perform tasks previously unknown or unthinkable. For example, think of a fully operative self-driving car that also has insight into economic models of supply and demand throughout a city and a financial interface controlling checks and balances. This agent has gained the capability of exploiting its own 'niche' (as McFarland and Bösser lined out), assessing risks and opportunities, trying to maximise its own profit, culminating into a truly autonomous economical agent. Intelligence here can be measured in economic success, to what degree it performs better than other, similar agents. But we can also imagine the cost of such form of intelligence that is optimised for making profit (a capitalist machine that might ignore 'poor' neighbourhoods and only drive around the wealthy parts of a city).

What does it mean to use space *intelligently*? What are the conditions for intelligent interactions between machines, animals and human beings? While autonomous robots are not yet 'smart' enough to be set free into the real world, we must be aware of our tendency to adjust spaces, 'dumbing down' spatial environments to fit their needs. Although the world is enriched with new cues and signs (QR codes, GPS fences marking safety zones, WiFi beacons pointing out docking stations), it is at the same time reduced to only a limited amount of qualities — it is an oversimplified version of reality. The problem is that for us, it is hard to detect the difference between an intelligent agent navigating a complex dynamic real-life landscape, and a relatively 'dumb' agent following a pre-programmed route. We see the robots moving around in physical space, strolling through the same three dimensions. It shows again that intelligence is not required to display intelligent behavior. We can walk around a city without having a map of the entire urban grid: not having a map does not limit our ability to walk, and for someone studying our walking, we might appear to display a fairly intelligent way of walking. In the same way, autonomous agents can act without a complete worldview/simulation or some kind of reasoning taking place in a central processing unit.

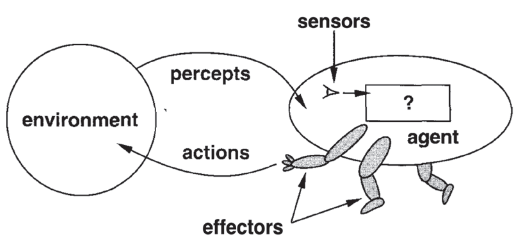

Fig.11

Sensing as Acting

A Roomba vacuums your floor, yet it is only programmed to drive around in a simple pattern and swiftly avoid obstacles. For the Roomba, sensing and thinking can be said to be one and the same thing, because sensing directly triggers certain ways of acting (sense wall — correct travel path to follow wall). At the same time, sensing might become the primary function of the many autonomous robots that will swarm public spaces of the future. Gathering information from impossible places, in impossible formats in impossible frequencies. In other to extract certain types of data, the machine might even be required to become invisible.

The final stage of compression is disappearance: the process of becoming invisible or blending in with the environment. Concrete examples are surveillance technologies that are concealed within the landscape — vast sensing networks buried underground, or autonomous agents mimicking plants and animals, enabling new forms of 'ambient governance' (as coined by Brian Murakami Wood).[7] To come back to the Bratton quote earlier, maybe our biggest fear is should not be about not being seen by the Big Machine, but about *not being able to see it.*

Figures

1
Sense Think Act Cycle

2
Land & Nillson - Animal Eyes

3
Systems and Methods of Merging Multiple Maps for Computer Vision Based Tracking

4
Artificial Neural Network Multilayered

5
Artificial Neural Network

6
Hopfield Model of Neural Network

7
A supernormal egg that is unusual only in the size of its spots (left) is preferred by the ringed plover (Charadrius hiaticula) to a normal egg (right)

8
Image Recognition Basics

9
Stevens Martin - Sensory Ecology Behaviour Evolution

10
Amazon Drone Tower Patent

11
Russel Norvig - Artificial Intelligence. A Modern Approach

References

1
https://www.newscientist.com/article/dn28255-new-urbanist-why-your-windowsill-will-need-its-own-address/

2
https://www.technologyreview.com/s/601215/china-is-building-a-robot-army-of-model-workers/

3
https://arxiv.org/pdf/1610.08401.pdf

4
https://arxiv.org/pdf/1703.01101.pdf

5
http://www.evolvingai.org/fooling

6
https://www.theguardian.com/technology/2017/jul/01/volvo-admits-its-self-driving-cars-are-confused-by-kangaroos

7
http://pactac.net/2014/03/vanishing-surveillance-securityscapes-and-ambient-government/

Text
Sander Manse

Design Research
Louisa Zahareas

Graphic Design
Anastasia Kubrak

2017

creative
industries
fund NL