


Anatomy of an AI System

 anatomyof.ai




I

A cylinder sits in a room. It is impassive, smooth, simple and small. It stands 14.8cm high, with a single blue-green circular light that traces around its upper rim. It is silently attending. A woman walks into the room, carrying a sleeping child in her arms, and she addresses the cylinder.

‘Alexa, turn on the hall lights’

The cylinder springs into life. ‘OK.’ The room lights up. The woman makes a faint nodding gesture, and carries the child upstairs.

This is an interaction with Amazon’s Echo device. ³ A brief command and a response is the most common form of engagement with this consumer voice-enabled AI device. But in this fleeting moment of interaction, a vast matrix of capacities is invoked: interlaced chains of resource extraction, human labor and algorithmic processing across networks of mining, logistics, distribution, prediction and optimization. The scale of this system is almost beyond human imagining. How can we begin to see it, to grasp its immensity and complexity as a connected form? **We start with an outline: an exploded view of a planetary system across three stages of birth, life and death, accompanied by an essay in 21 parts. Together, this becomes an anatomical map of a single AI system.**



Amazon Echo Dot (schematics)

II

The scene of the woman talking to Alexa is drawn from a 2017 promotional video advertising the latest version of the Amazon Echo. The video begins, “Say hello to the all-new Echo” and explains that the Echo will connect to Alexa (the artificial intelligence agent) in order to “play music, call friends and family, control smart home devices, and more.” The device contains seven directional microphones, so the user can be heard at all times even when music is playing. The device comes in several styles, such as gunmetal grey or a basic beige, designed to either “blend in or stand out.” But even the shiny design options maintain a kind of blankness: nothing will alert the owner to the vast network that subtends and drives its interactive capacities. The promotional video simply states that the range of things you can ask Alexa to do is always expanding. “Because Alexa is in the cloud, she is always getting smarter and adding new features.”

How does this happen? Alexa is a disembodied voice that represents the human-AI interaction interface for an extraordinarily complex set of information processing layers. These layers are fed by constant tides: the flows of human voices being translated into text questions, which are used to query databases of potential answers, and the corresponding ebb of Alexa’s replies. For each response that Alexa gives, its effectiveness is inferred by what happens next:

Is the same question uttered again? (*Did the user feel heard?*)

Was the question reworded? (*Did the user feel the question was understood?*)


Was there an action following the question? (*Did the interaction result in a tracked response: a light turned on, a product purchased, a track played?*)

With each interaction, Alexa is training to hear better, to interpret more precisely, to trigger actions that map to the user's commands more accurately, and to build a more complete model of their preferences, habits and desires. What is required to make this possible? Put simply: each small moment of convenience – be it answering a question, turning on a light, or playing a song – requires a vast planetary network, fueled by the extraction of non-renewable materials, labor, and data. The scale of resources required is many magnitudes greater than the energy and labor it would take a human to operate a household appliance or flick a switch. A full accounting for these costs is almost impossible, but it is increasingly important that we grasp the scale and scope if we are to understand and govern the technical infrastructures that thread through our lives.

III


The Salar, the world's largest flat surface, is located in southwest Bolivia at an altitude of 3,656 meters above sea level. It is a high plateau, covered by a few meters of salt crust which are exceptionally rich in lithium, containing 50% to 70% of the world's lithium reserves. ⁴ The Salar, alongside the neighboring Atacama regions in Chile and Argentina, are major sites for lithium extraction. This soft, silvery metal is currently used to power mobile connected devices, as a crucial material used for the production of lithium-Ion batteries. It is known as 'grey gold.' Smartphone batteries, for example, usually have less than eight grams of this material. ⁵ Each Tesla car needs approximately seven kilograms of lithium for its battery pack. ⁶ All these batteries have a limited lifespan, and once consumed they are thrown away as waste. Amazon reminds users that they cannot open up and repair their Echo, because this will void the warranty. The Amazon Echo is wall-powered, and also has a mobile battery base. This also has a limited lifespan and then must be thrown away as waste.

According to the Aymara legends about the creation of Bolivia, the volcanic mountains of the Andean plateau were creations of tragedy. ⁷ Long ago, when the volcanos were alive and roaming the plains freely, Tunupa - the only female volcano – gave birth to a baby. Stricken by jealousy, the male volcanos stole her baby and banished it to a distant location. The gods punished the volcanos by pinning them all to the Earth. Grieving for the child that she could no longer reach, Tunupa wept deeply. Her tears and breast milk combined to create a giant salt lake: Salar de Uyuni. As Liam Young and Kate Davies observe, “your smart-phone runs on the tears and breast milk of a volcano. This landscape is connected to everywhere on the planet via the phones in our pockets; linked to each of us by invisible threads of commerce, science, politics and power.” ⁸




IV

Our exploded view diagram combines and visualizes three central, extractive processes that are required to run a large-scale artificial intelligence system: material resources, human labor, and data. We consider these three elements across time – represented as a visual description of the birth, life and death of a single Amazon Echo unit. It's necessary to move beyond a simple analysis of the relationship between an individual human, their data, and any single technology company in order to contend with the truly planetary scale of extraction. Vincent Mosco has shown how the ethereal metaphor of 'the cloud' for offsite data management and processing is in complete contradiction with the physical realities of the extraction of minerals from the Earth's crust and dispossession of human populations that sustain its existence.⁹ Sandro Mezzadra and Brett Nielson use the term 'extractivism' to name the relationship between different forms of extractive operations in contemporary capitalism, which we see repeated in the context of the AI industry.¹⁰ There are deep interconnections between the literal hollowing out of the materials of the earth and biosphere, and the data capture and monetization of human practices of communication and sociality in AI. Mezzadra and Nielson note that labor is central to this extractive relationship, which has repeated throughout history: from the way European imperialism used slave labor, to the forced work crews on rubber plantations in Malaya, to the Indigenous people of Bolivia being driven to extract the silver that was used in the first global currency. Thinking about extraction requires thinking about labor, resources, and data together. This presents a challenge to critical and popular understandings of artificial intelligence: it is hard to 'see' any of these processes individually, let alone collectively. Hence the need for a visualization that can bring these connected, but globally dispersed processes into a single map.



V

If you read our map from left to right, the story begins and ends with the Earth, and the geological processes of deep time. But read from top to bottom, we see the story as it begins and ends with a human. The top is the human agent, querying the Echo, and supplying Amazon with the valuable training data of verbal questions and responses that they can use to further refine their voice-enabled AI systems. At the bottom of the map is another kind of human resource: the history of human knowledge and capacity, which is also used to train and optimize artificial intelligence systems. This is a key difference between artificial intelligence systems and other forms of consumer technology: they rely on the ingestion, analysis and optimization of vast amounts of human generated images, texts and videos.



VI

When a human engages with an Echo, or another voice-enabled AI device, they are acting as much more than just an end-product consumer. It is difficult to place the human user of an AI system into a single category: rather, they deserve to be considered as a hybrid case. Just as the Greek *chimera* was a mythological animal that was part lion, goat, snake and monster, the Echo user is simultaneously a consumer, a resource, a worker, and a product. This multiple identity recurs for human users in many technological systems. In the specific case of the Amazon Echo, the user has purchased a consumer device for which they receive a set of convenient affordances. But they are also a resource, as their voice commands are collected, analyzed and retained for the purposes of building an ever-larger corpus of human voices and instructions. And they provide labor, as they continually perform the valuable service of contributing feedback mechanisms regarding the accuracy, usefulness, and overall quality of Alexa's replies. They are, in essence, helping to train the neural networks within Amazon's infrastructural stack.

VII


Anything beyond the limited physical and digital interfaces of the device itself is outside of the user's control. It presents a sleek surface with no ability to open it, repair it or change how it functions. The object itself is a very simple extrusion of plastic representing a collection of sensors – its real power and

complexity lies somewhere else, far out of sight. The Echo is but an ‘ear’ in the home: a disembodied listening agent that never shows its deep connections to remote systems.

In 1673, the Jesuit polymath, Athanasius Kircher, invented the *statua citofonica* – the ‘talking statue.’ Kircher was an extraordinary interdisciplinary scholar and inventor. In his lifetime he published forty major works across the fields of medicine, geology, comparative religion and music. He invented the first magnetic clock, many early automatons, and the megaphone. His talking statue was a very early listening system: essentially a microphone made from a huge spiral tube, which could convey the conversations from a public square and up through the tube, and then piped through the mouth of a statue kept within an aristocrat’s private chambers. As Kircher wrote:

*“This statue must be located in a given place, in order to allow the end section of the spiral-shaped tube to precisely correspond to the opening of the mouth. In this manner it will be perfect, and capable to emit clearly any kind of sound: in fact the statue will be able to speak continuously, uttering in either a human or animal voice: it will laugh or sneer; it will seem to really cry or moan; sometimes with great astonishment it will strongly blow. If the opening of the spiral shaped tube is located in correspondence to an open public space, all human words pronounced, focused in the conduit, would be replayed through the mouth of the statue.”*¹¹

The listening system could eavesdrop on everyday conversations in the piazza, and relay them to the 17th century Italian oligarchs. Kircher’s talking statue was an early form of information extraction for the elites – people talking in the street would have no indication that their conversations were being funneled to those who would instrument that knowledge for their own power, entertainment and wealth. People inside the homes of aristocrats would have no idea how a magical statue was speaking and conveying all manner of information. The aim was to obscure how the system worked: an elegant statue was all they could see. Listening systems, even at this early stage, were about power, class, and secrecy. But the infrastructure for Kircher’s system was prohibitively expensive – available only to the very few. And so the question remains, what are the full resource implications of building such systems? This brings us to the materiality of the infrastructure that lies beneath.



Statua citofonica by Athanasius Kircher (1673)

VIII


In his book *A Geology of Media*, Jussi Parikka suggests that we try to think of media not from Marshall McLuhan's point of view – in which media are extensions of human senses ¹² – but rather as an extension of Earth. ¹³ Media technologies should be understood in context of a geological process, from the creation and the transformation processes, to the movement of natural elements from which media are built. Reflecting upon media and technology as geological processes enables us to consider the profound depletion of non-renewable resources required to drive the technologies of the present moment. Each object in the extended network of an AI system, from network routers to batteries to microphones, is built using elements that required billions of years to be produced. Looking from the perspective of deep time, we are extracting Earth's history to serve a split second of technological time, in order to build devices than are often designed to be used for no more than a few years. For example, the Consumer Technology Association notes that the average smartphone lifespan is 4.7 years. ¹⁴ This obsolescence cycle fuels the purchase of more devices, drives up profits, and increases incentives for the use of unsustainable extraction practices. From a slow process of elemental

development, these elements and materials go through an extraordinarily rapid period of excavation, smelting, mixing, and logistical transport – crossing thousands of kilometers in their transformation. Geological processes mark both the beginning and the end of this period, from the mining of ore, to the deposition of material in an electronic waste dump. For that reason, our map starts and ends with the Earth's crust. However, all the transformations and movements we depict are only the barest anatomical outline: beneath these connections lie many more layers of fractal supply chains, and exploitation of human and natural resources, concentrations of corporate and geopolitical power, and continual energy consumption.

IX

Drawing out the connections between resources, labor and data extraction brings us inevitably back to traditional frameworks of exploitation. But how is value being generated through these systems? A useful conceptual tool can be found in the work of Christian Fuchs and other authors examining and defining digital labor. The notion of digital labor, which was initially linked with different forms of non-material labor, precedes the life of devices and complex systems such as artificial intelligence. Digital labor – the work of building and maintaining the stack of digital systems – is far from ephemeral or virtual, but is deeply embodied in different activities.¹⁵ The scope is overwhelming: from indentured labor in mines for extracting the minerals that form the physical basis of information technologies; to the work of strictly controlled and sometimes dangerous hardware manufacturing and assembly processes in Chinese factories; to exploited outsourced cognitive workers in developing countries labelling AI training data sets; to the informal physical workers cleaning up toxic waste dumps. These processes create new accumulations of wealth and power, which are concentrated in a very thin social layer.


Marx's dialectic of subject and object in economy



X

This triangle of value extraction and production represents one of the basic elements of our map, from birth in a geological process, through life as a consumer AI product, and ultimately to death in an electronics dump. Like in Fuchs' work, our triangles are not isolated, but linked to one another in the production process. They form a cyclic flow in which the product of work is transformed into a resource, which is transformed into a product, which is transformed into a resource and so on. Each triangle represents one phase in the production process. Although this appears on the map as a linear path of transformation, a different visual metaphor better represents the complexity of current extractivism: the fractal structure known as the Sierpinski triangle.

A linear display does not enable us to show that each next step of production and exploitation contains previous phases. If we look at the production and exploitation system through a fractal visual structure, the smallest triangle would represent natural resources and means of labor, i.e. the miner as labor and ore as product. The next larger triangle encompasses the processing of metals, and the next would represent the process of manufacturing components and so on. The ultimate triangle in our map, the production of the Amazon Echo unit itself, includes all of these levels of exploitation – from the bottom to the very top of Amazon Inc, a role inhabited by Jeff Bezos as CEO of Amazon. Like a pharaoh of ancient Egypt, he stands at the top of the largest pyramid of AI value extraction.



Sierpinski triangle or Sierpinski fractal

XI

To return to the basic element of this visualization – a variation of Marx’s triangle of production – each triangle creates a surplus of value for creating profits. If we look at the scale of average income for each activity in the production process of one device, which is shown on the left side of our map, we see the dramatic difference in income earned. According to research by Amnesty International, during the excavation of cobalt which is also used for lithium batteries of 16 multinational brands, workers are paid the equivalent of one US dollar per day for working in conditions hazardous to life and health, and were often subjected to violence, extortion and intimidation.¹⁶ Amnesty has documented children as young as 7 working in the mines. In contrast, Amazon CEO Jeff Bezos, at the top of our fractal pyramid, made an average of \$275 million a day during the first five months of 2018, according to the Bloomberg Billionaires Index.¹⁷ A child working in a mine in the Congo would need more than 700,000 years of non-stop work to earn the same amount as a single day of Bezos’ income.

Many of the triangles shown on this map hide different stories of labor exploitation and inhumane working conditions. The ecological price of transformation of elements and income disparities is just one of the possible ways of representing a deep systemic inequality. We have both researched

different forms of 'black boxes' understood as algorithmic processes, ¹⁸ but this map points to another form of opacity: the very processes of creating, training and operating a device like an Amazon Echo is itself a kind of black box, very hard to examine and track *in toto* given the multiple layers of contractors, distributors, and downstream logistical partners around the world. As Mark Graham writes, "contemporary capitalism conceals the histories and geographies of most commodities from consumers. Consumers are usually only able to see commodities in the here and now of time and space, and rarely have any opportunities to gaze backwards through the chains of production in order to gain knowledge about the sites of production, transformation, and distribution." ¹⁹

One illustration of the difficulty of investigating and tracking the contemporary production chain process is that it took Intel more than four years to understand its supply line well enough to ensure that no tantalum from the Congo was in its microprocessor products. As a semiconductor chip manufacturer, Intel supplies Apple with processors. In order to do so, Intel has its own multi-tiered supply chain of more than 19,000 suppliers in over 100 countries providing direct materials for their production processes, tools and machines for their factories, and logistics and packaging services. ²⁰ That it took over four years for a leading technology company *just to understand its own supply chain*, reveals just how hard this process can be to grasp from the inside, let alone for external researchers, journalists and academics. Dutch-based technology company Philips has also claimed that it was working to make its supply chain 'conflict-free'. Philips, for example, has tens of thousands of different suppliers, each of which provides different components for their manufacturing processes. ²¹ Those suppliers are themselves linked downstream to tens of thousands of component manufacturers that acquire materials from hundreds of refineries that buy ingredients from different smelters, which are supplied by unknown numbers of traders that deal directly with both legal and illegal mining operations. In *The Elements of Power*, David S. Abraham describes the invisible networks of rare metals traders in global electronics supply chains: "The network to get rare metals from the mine to your laptop travels through a murky network of traders, processors, and component manufacturers. Traders are the middlemen who do more than buy and sell rare metals: they help to regulate information and are the hidden link that helps in navigating the network between metals plants and the components in our laptops." ²² According to the computer manufacturing company Dell, complexities of the metal supply chain pose almost insurmountable challenges. ²³ The mining of these minerals takes place long before a final product is assembled, making it exceedingly difficult to trace the minerals' origin. In addition, many of the minerals are smelted together with recycled metals, by

which point it becomes all but impossible to trace the minerals to their source. So we see that the attempt to capture the full supply chain is a truly gargantuan task: revealing all the complexity of the 21st century global production of technology products.


XII

Supply chains are often layered on top of one another, in a sprawling network. Apple's supplier program reveals there are tens of thousands of individual components embedded in their devices, which are in turn supplied by hundreds of different companies. In order for each of those components to arrive on the final assembly line where it will be assembled by workers in Foxconn facilities, different components need to be physically transferred from more than 750 supplier sites across 30 different countries.²⁴ This becomes a complex structure of *supply chains within supply chains*, a zooming fractal of tens of thousands of suppliers, millions of kilometers of shipped materials and hundreds of thousands of workers included within the process even before the product is assembled on the line.

Visualizing this process as one global, pancontinental network through which materials, components and products flow, we see an analogy to the global information network. Where there is a single internet packet travelling to an Amazon Echo, here we can imagine a single cargo container.²⁵ The dizzying spectacle of global logistics and production will not be possible without the invention of this simple, standardized metal object. Standardized cargo containers allowed the explosion of modern shipping industry, which made it possible to model the planet as a massive, single factory. In 2017, the capacity of container ships in seaborne trade reached nearly 250,000,000 dead-weight tons of cargo, dominated by giant shipping companies like Maersk of Denmark, the Mediterranean Shipping Company of Switzerland, and France's CMA CGM Group, each owning hundred of container vessels.²⁶ For these commercial ventures, cargo shipping is a relatively cheap way to traverse the vascular system of the global factory, yet it disguises much larger external costs.

In recent years, shipping boats produce 3.1% of global yearly CO₂ emissions, more than the entire country of Germany.²⁷ In order to minimize their internal costs, most of the container shipping companies use very low grade fuel in enormous quantities, which leads to increased amounts of sulphur in the air, among other toxic substances. It has been estimated that one container ship can emit as much pollution as 50 million cars, and 60,000 deaths worldwide are attributed indirectly to cargo ship industry pollution related issues annually.²⁸ Even industry-friendly sources like the World Shipping Council admit that

thousands of containers are lost each year, on the ocean floor or drifting loose.²⁹ Some carry toxic substances which leak into the oceans. Typically, workers spend 9 to 10 months in the sea, often with long working shifts and without access to external communications. Workers from the Philippines represent more than a third of the global shipping workforce.³⁰ The most severe costs of global logistics are born by the atmosphere, the oceanic ecosystem and all it contains, and the lowest paid workers.



Cargo container

XIII

The increasing complexity and miniaturization of our technology depends on the process that strangely echoes the hopes of early medieval alchemy. Where medieval alchemists aimed to transform base metals into ‘noble’ ones, researchers today use rare earth metals to enhance the performance of other minerals. There are 17 rare earth elements, which are embedded in laptops and smartphones, making them smaller and lighter. They play a role in color displays, loudspeakers, camera lenses, GPS systems, rechargeable batteries, hard drives and many other components. They are key elements in communication systems from fiber optic cables, signal amplification in mobile communication towers to satellites and GPS technology. But the precise

configuration and use of these minerals is hard to ascertain. In the same way that medieval alchemists hid their research behind cyphers and cryptic symbolism, contemporary processes for using minerals in devices are protected behind NDAs and trade secrets.

The unique electronic, optical and magnetic characteristics of rare earth elements cannot be matched by any other metals or synthetic substitutes discovered to date. While they are called ‘rare earth metals’, some are relatively abundant in the Earth’s crust, but extraction is costly and highly polluting. David Abraham describes the mining of dysprosium and Terbium used in a variety of high-tech devices in Jianxi, China. He writes, “Only 0.2 percent of the mined clay contains the valuable rare earth elements. This means that 99.8 percent of earth removed in rare earth mining is discarded as waste called “tailings” that are dumped back into the hills and streams,” creating new pollutants like ammonium. ³¹ In order to refine one ton of rare earth elements, “the Chinese Society of Rare Earths estimates that the process produces 75,000 liters of acidic water and one ton of radioactive residue.” ³² Furthermore, mining and refining activities consume vast amount of water and generate large quantities of CO2 emissions. In 2009, China produced 95% of the world's supply of these elements, and it has been estimated that the single mine known as Bayan Obo contains 70% of the world's reserves. ³³

scandium 21 Sc 44.956	yttrium 39 Y 88.906	lanthanum 57 La 138.91	cerium 58 Ce 140.12
praseodymium 59 Pr 140.91	neodymium 60 Nd 144.24	promethium 61 Pm [145]	samarium 62 Sm 150.36
europium 63 Eu 151.96	gadolinium 64 Gd 157.25	terbium 65 Tb 158.93	dysprosium 66 Dy 162.50
holmium 67 Ho 164.93	erbium 68 Er 167.26	thulium 69 Tm 168.93	ytterbium 70 Yb 173.05
lutetium 71 Lu 174.97			

Rare earth elements

XIV

A satellite picture of the tiny Indonesian island of Bangka tells a story about human and environmental toll of the semiconductor production. On this tiny island, mostly ‘informal’ miners are on makeshift pontoons, using bamboo poles to scrape the seabed, and then diving underwater to suck tin from the surface through giant, vacuum-like tubes. As a *Guardian* investigation reports “tin mining is a lucrative but destructive trade that has scarred the island's landscape, bulldozed its farms and forests, killed off its fish stocks and coral reefs, and dented tourism to its pretty palm-lined beaches. The damage is best seen from the air, as pockets of lush forest huddle amid huge swaths of barren orange earth. Where not dominated by mines, this is pockmarked with graves, many holding the bodies of miners who have died over the centuries digging for tin.”³⁴ Two small islands, Bangka and Belitung, produce 90% of Indonesia's tin, and Indonesia is the world's second-largest exporter of the metal. Indonesia's



national tin corporation, PT Timah, supplies companies such as Samsung directly, as well as solder makers Chernan and Shenmao, which in turn supply Sony, LG and Foxconn. ³⁵

XV

At Amazon distribution centers, vast collections of products are arrayed in a computational order across millions of shelves. The position of every item in this space is precisely determined by complex mathematical functions that process information about orders and create relationships between products. The aim is to optimize the movements of the robots and humans that collaborate in these warehouses. With the help from an electronic bracelet, the human worker is directed through warehouses the size of airplane hangars, filled with objects arranged in an opaque algorithmic order. ³⁶

Hidden among the thousands of other publicly available patents owned by Amazon, U.S. patent number 9,280,157 represents an extraordinary illustration of worker alienation, a stark moment in the relationship between humans and machines. ³⁷ It depicts a metal cage intended for the worker, equipped with different cybernetic add-ons, that can be moved through a warehouse by the same motorized system that shifts shelves filled with merchandise. Here, the worker becomes a part of a machinic ballet, held upright in a cage which dictates and constrains their movement.

As we have seen time and time again in the research for our map, dystopian futures are built upon the unevenly distributed dystopian regimes of the past and present, scattered through an array of production chains for modern technical devices. The vanishingly few at the top of the fractal pyramid of value extraction live in extraordinary wealth and comfort. But the majority of the pyramids are made from the dark tunnels of mines, radioactive waste lakes, discarded shipping containers, and corporate factory dormitories.



Amazon patent number 20150066283 A1

XVI

At the end of 19th century, a particular Southeast Asian tree called *palaquium gutta* became the center of a technological boom. These trees, found mainly in Malaysia, produce a milky white natural latex called *gutta percha*. After English scientist Michael Faraday published a study in The Philosophical Magazine in 1848 about the use of this material as an electrical insulator, *gutta percha* rapidly became the darling of the engineering world. It was seen as the solution to the problem of insulating telegraphic cables in order that they could withstand the conditions of the ocean floor. As the global submarine business grew, so did demand for *palaquium gutta* tree trunks. The historian John Tully describes how local Malay, Chinese and Dayak workers were paid little for the dangerous works of felling the trees and slowly collecting the latex. ³⁸ The latex was processed then sold through Singapore's trade markets into the British market, where it was transformed into, among other things, lengths upon lengths of submarine cable sheaths.

A mature *palaquium gutta* could yield around 300 grams of latex. But in 1857, the first transatlantic cable was around 3000 km long and weighed 2000 tons – requiring around 250 tons of *gutta percha*. To produce just one ton of this material required around 900,000 tree trunks. The jungles of Malaysia and

Singapore were stripped, and by the early 1880s the *palaquium gutta* had vanished. In a last-ditch effort to save their supply chain, the British passed a ban in 1883 to halt harvesting the latex, but the tree was already extinct. ³⁹

The Victorian environmental disaster of *gutta percha*, from the early origins of the global information society, shows how the relationships between technology and its materiality, environments, and different forms of exploitation are imbricated. Just as Victorians precipitated ecological disaster for their early cables, so do rare earth mining and global supply chains further imperil the delicate ecological balance of our era. From the material used to build the technology enabling contemporary networked society, to the energy needed for transmitting, analyzing, and storing the data flowing through the massive infrastructure, to the materiality of infrastructure: these deep connections and costs are more significant, and have a far longer history, than is usually represented in the corporate imaginaries of AI. ⁴⁰



Palaquium gutta

Large-scale AI systems consume enormous amounts of energy. Yet the material details of those costs remain vague in the social imagination. It remains difficult to get precise details about the amount of energy consumed by cloud computing services. A Greenpeace report states: “One of the single biggest obstacles to sector transparency is Amazon Web Services (AWS). The world’s biggest cloud computer company remains almost completely non-transparent about the energy footprint of its massive operations. Among the global cloud providers, only AWS still refuses to make public basic details on the energy performance and environmental impact associated with its operations.” ⁴¹

As human agents, we are visible in almost every interaction with technological platforms. We are always being tracked, quantified, analyzed and commodified. But in contrast to user visibility, the precise details about the phases of birth, life and death of networked devices are obscured. With emerging devices like the Echo relying on a centralized AI infrastructure far from view, even more of the detail falls into the shadows.

While consumers become accustomed to a small hardware device in their living rooms, or a phone app, or a semi-autonomous car, the real work is being done within machine learning systems that are generally remote from the user and utterly invisible to her. In many cases, transparency wouldn’t help much – without forms of real choice, and corporate accountability, mere transparency won’t shift the weight of the current power asymmetries. ⁴²

The outputs of machine learning systems are predominantly unaccountable and ungoverned, while the inputs are enigmatic. To the casual observer, it looks like it has never been easier to build AI or machine learning-based systems than it is today. Availability of open-source tools for doing so in combination with rentable computation power through cloud superpowers such as Amazon (AWS), Microsoft (Azure), or Google (Google Cloud) is giving rise to a false idea of the ‘democratization’ of AI. While ‘off the shelf’ machine learning tools, like TensorFlow, are becoming more accessible from the point of view of setting up your own system, the underlying logics of those systems, and the datasets for training them are accessible to and controlled by very few entities. In the dynamic of dataset collection through platforms like Facebook, users are feeding and training the neural networks with behavioral data, voice, tagged pictures and videos or medical data. In an era of extractivism, the real value of that data is controlled and exploited by the very few at the top of the pyramid.

XVIII

When massive data sets are used to train AI systems, the individual images and videos involved are commonly tagged and labeled. ⁴³ There is much to be said about how this labelling process abrogates and crystallizes meaning, and further, how this process is driven by clickworkers being paid fractions of a cent for this digital piecework.

In 1770, Hungarian inventor Wolfgang von Kempelen constructed a chess-playing machine known as the Mechanical Turk. His goal, in part, was to impress Empress Maria Theresa of Austria. This device was capable of playing chess against a human opponent and had spectacular success winning most of the games played during its demonstrations around Europe and the Americas for almost nine decades. But the Mechanical Turk was an illusion that allowed a human chess master to hide inside the machine and operate it. Some 160 years later, Amazon.com branded its micropayment based crowdsourcing platform with the same name. According to Ayhan Aytes, Amazon's initial motivation to build Mechanical Turk emerged after the failure of its artificial intelligence programs in the task of finding duplicate product pages on its retail website. ⁴⁴ After a series of futile and expensive attempts, the project engineers turned to humans to work behind computers within a streamlined web-based system. ⁴⁵ Amazon Mechanical Turk digital workshop emulates artificial intelligence systems by checking, assessing and correcting machine learning processes with human brainpower. With Amazon Mechanical Turk, it may seem to users that an application is using advanced artificial intelligence to accomplish tasks. But it is closer to a form of 'artificial artificial intelligence', driven by a remote, dispersed and poorly paid clickworker workforce that helps a client achieve their business objectives. As observed by Aytes, "in both cases [both the Mechanical Turk from 1770 and the contemporary version of Amazon's service] the performance of the workers who animate the artifice is obscured by the spectacle of the machine." ⁴⁶

This kind of invisible, hidden labor, outsourced or crowdsourced, hidden behind interfaces and camouflaged within algorithmic processes is now commonplace, particularly in the process of tagging and labeling thousands of hours of digital archives for the sake of feeding the neural networks. Sometimes this labor is entirely unpaid, as in the case of the Google's reCAPTCHA. In a paradox that many of us have experienced, in order to prove that you are not artificial agent, you are forced to train Google's image recognition AI system for free, by selecting multiple boxes that contain street numbers, or cars, or houses.

As we see repeated throughout the system, contemporary forms of artificial intelligence are not so artificial after all. We can speak of the hard physical labor of mine workers, and the repetitive factory labor on the assembly line, of the cybernetic labor in distribution centers and the cognitive sweatshops full of

outsourced programmers around the world, of the low paid crowdsourced labor of Mechanical Turk workers, or the unpaid immaterial work of users. At every level contemporary technology is deeply rooted in and running on the exploitation of human bodies.

Mechanical Turk



XIX

In his one-paragraph short story "On Exactitude in Science", Jorge Luis Borges presents us with an imagined empire in which cartographic science became so developed and precise, that it needed a map on the same scale as the empire itself. ⁴⁷

"...In that Empire, the Art of Cartography attained such Perfection that the map of a single Province occupied the entirety of a City, and the map of the Empire, the entirety of a Province. In time, those Unconscionable Maps no longer satisfied, and the Cartographers Guilds struck a Map of the Empire whose size was that of the Empire, and which coincided point for point with it. The following Generations, who were not so fond of the Study of Cartography as their Forebears had been, saw that that vast map was Useless, and not without some Pitilessness was it, that they delivered it up to the Inclemencies of Sun and Winters. In the Deserts of the West, still today, there are Tattered Ruins of that Map, inhabited by Animals and Beggars; in all the Land there is no other Relic of the Disciplines of Geography."

Current machine learning approaches are characterized by an aspiration to map the world, a full quantification of visual, auditory, and recognition regimes of reality. From cosmological model for the universe to the world of human emotions as interpreted through the tiniest muscle movements in the human face, everything becomes an object of quantification. Jean-François Lyotard introduced the phrase “affinity to infinity” to describe how contemporary art, techno-science and capitalism share the same aspiration to push boundaries towards a potentially infinite horizon. ⁴⁸ The second half of the 19th century, with its focus on the construction of infrastructure and the uneven transition to industrialized society, generated enormous wealth for the small number of industrial magnates that monopolized exploitation of natural resources and production processes.

The new infinite horizon is data extraction, machine learning, and reorganizing information through artificial intelligence systems of combined human and machinic processing. The territories are dominated by a few global mega-companies, which are creating new infrastructures and mechanisms for the accumulation of capital and exploitation of human and planetary resources.

Such unrestrained thirst for new resources and fields of cognitive exploitation has driven a search for ever deeper layers of data that can be used to quantify the human psyche, conscious and unconscious, private and public, idiosyncratic and general. In this way, we have seen the emergence of multiple cognitive economies from the attention economy, ⁴⁹ the surveillance economy, the reputation economy, ⁵⁰ and the emotion economy, as well as the quantification and commodification of trust and evidence through cryptocurrencies.

Increasingly, the process of quantification is reaching into the human affective, cognitive, and physical worlds. Training sets exist for emotion detection, for family resemblance, for tracking an individual as they age, and for human actions like sitting down, waving, raising a glass, or crying. Every form of biodata – including forensic, biometric, sociometric, and psychometric – are being captured and logged into databases for AI training. That quantification often runs on very limited foundations: datasets like AVA which primarily shows women in the ‘playing with children’ action category, and men in the ‘kicking a person’ category. The training sets for AI systems claim to be reaching into the fine-grained nature of everyday life, but they repeat the most stereotypical and restricted social patterns, re-inscribing a normative vision of the human past and projecting it into the human future.



Quantification of Nature

XX

"The 'enclosure' of biodiversity and knowledge is the final step in a series of enclosures that began with the rise of colonialism. Land and forests were the first resources to be 'enclosed' and converted from commons to commodities. Later on, water resources were 'enclosed' through dams, groundwater mining and privatization schemes. Now it is the turn of biodiversity and knowledge to be 'enclosed' through intellectual property rights (IPRs)," Vandana Shiva explains. ⁵¹ In Shiva's words, "the destruction of commons was essential for the industrial revolution, to provide a supply of natural resources for raw material to industry. A life-support system can be shared, it cannot be owned as private property or exploited for private profit. The commons, therefore, had to be privatized, and people's sustenance base in these commons had to be appropriated, to feed the engine of industrial progress and capital accumulation." ⁵²

While Shiva is referring to enclosure of nature by intellectual property rights, the same process is now occurring with machine learning – an intensification of quantified nature. The new gold rush in the context of artificial intelligence is to

enclose different fields of human knowing, feeling, and action, in order to capture and privatize those fields. When in November 2015 DeepMind Technologies Ltd. got access to the health records of 1.6 million identifiable patients of Royal Free hospital, we witnessed a particular form of privatization: the extraction of knowledge value.⁵³ A dataset may still be publicly owned, but the meta-value of the data – the model created by it – is privately owned. While there are many good reasons to seek to improve public health, there is a real risk if it comes at the cost of a stealth privatization of public medical services. That is a future where expert local human labor in the public system is augmented and sometimes replaced with centralized, privately-owned corporate AI systems, that are using public data to generate enormous wealth for the very few.

Corporate border



XXI

At this moment in the 21st century, we see a new form of extractivism that is well underway: one that reaches into the furthest corners of the biosphere and the deepest layers of human cognitive and affective being. Many of the assumptions about human life made by machine learning systems are narrow, normative and laden with error. Yet they are inscribing and building those assumptions into a new world, and will increasingly play a role in how opportunities, wealth, and knowledge are distributed.

The stack that is required to interact with an Amazon Echo goes well beyond the multi-layered ‘technical stack’ of data modeling, hardware, servers and networks. The full stack reaches much further into capital, labor and nature, and demands an enormous amount of each. The true costs of these systems – social, environmental, economic, and political – remain hidden and may stay that way for some time.

We offer up this map and essay as a way to begin seeing across a wider range of system extractions. The scale required to build artificial intelligence systems is too complex, too obscured by intellectual property law, and too mired in logistical complexity to fully comprehend in the moment. Yet you draw on it every time you issue a simple voice command to a small cylinder in your living room: ‘Alexa, what time is it?’

And so the cycle continues.

Footnotes

Kate Crawford is a Distinguished Research Professor at New York University, a Principal Researcher at Microsoft Research New York, and the co-founder and co-director of the AI Now Institute at NYU.

Vladan Joler is a professor at the Academy of Arts at the University of Novi Sad and founder of SHARE Foundation. He is leading SHARE Lab, a research and data investigation lab for exploring different technical and social aspects of algorithmic transparency, digital labor exploitation, invisible infrastructures, and technological black boxes.

Amazon advertising campaign, “All-New Amazon Echo”, September 27, 2017, <https://www.youtube.com/watch?v=IyvZ41XjUjY>.

Emily Achtenberg, “Bolivia Bets on State-Run Lithium Industry,” NACLA, November 15, 2010, <https://nacla.org/news/bolivia-bets-state-run-lithium-industry>.

Christine Negroni, “How to Determine the Power Rating of Your Gadget’s Batteries,” *The New York Times*, December 22, 2017, <https://www.nytimes.com/2016/12/26/business/lithium-ion-battery-airline-safety.html>.

Jessica Shankleman et al., “We’re Going to Need More Lithium,” *Bloomberg*, September 7, 2017, <https://www.bloomberg.com/graphics/2017-lithium-battery-future/>.

Nicola Clark and Simon Wallis, “Flamingos, Salt Lakes and Volcanoes: Hunting for Evidence of Past Climate Change on the High Altiplano of Bolivia,” *Geology Today* 33, no. 3 (May 1, 2017): 104, <https://doi.org/10.1111/gto.12186>.

Kate Davies and Liam Young, *Tales from the Dark Side of the City: The Breastmilk of the Volcano Bolivia and the Atacama Desert Expedition* (London: Unknown Fields, 2016).

Vincent Mosco, *To the Cloud: Big Data in a Turbulent World* (Boulder: Paradigm, 2014).

Sandro Mezzadra and Brett Neilson, “On the Multiple Frontiers of Extraction: Excavating Contemporary Capitalism,” *Cultural Studies* 31, no. 2–3 (May 4, 2017): 185, <https://doi.org/10.1080/09502386.2017.1303425>.

Lamberto Tronchin, “The ‘Phonurgia Nova’ of Athanasius Kircher: The Marvellous Sound World of 17th Century,” *Proceedings of Meetings on Acoustics* 4, no. 1 (June 29, 2008), 4: 015002, <https://doi.org/10.1121/1.2992053>.

Marshall McLuhan, *Understanding Media: The Extensions of Man* (New York: Signet Books, 1964).

Jussi Parikka, *A Geology of Media* (Minneapolis: University Of Minnesota Press, 2015), vii–viii.

Chris Ely, “The Life Expectancy of Electronics,” Consumer Technology Association, September 16, 2014, <https://www.cta.tech/News/Blog/Articles/2014/September/The-Life-Expectancy-of-Electronics.aspx>.

Christian Fuchs, *Digital Labor and Karl Marx* (London: Routledge, 2014).

“This Is What We Die For: Human Rights Abuses in the Democratic Republic of the Congo Power the Global Trade in Cobalt” (London: Amnesty International, 2016), <https://www.amnesty.org/download/Documents/AFR6231832016ENGLISH.PDF>. For an anthropological description of these mining processes, see: Jeffrey W. Mantz, “Improvisational Economies: Coltan Production in the Eastern Congo,” *Social Anthropology* 16, no. 1 (February 1, 2008): 34–50, <https://doi.org/10.1111/j.1469-8676.2008.00035.x>.

Julia Glum, “The Median Amazon Employee’s Salary Is \$28,000. Jeff Bezos Makes More Than That in 10 Seconds,” *Time*, May 2, 2018, <http://time.com/money/5262923/amazon-employee-median-salary-jeff-bezos/>.

Frank Pasquale, *The Black Box Society: The Secret Algorithms That Control Money and Information* (Cambridge, MA: Harvard University Press, 2016).

Mark Graham and Håvard Haarstad, “Transparency and Development: Ethical Consumption through Web 2.0 and the Internet of Things,” *Information Technologies & International Development* 7, no. 1 (March 10, 2011): 1.

“Intel’s Efforts to Achieve a ‘Conflict Free’ Supply Chain” (Santa Clara, CA: Intel Corporation, May 2018), <https://www.intel.com/content/www/us/en/corporate-responsibility/conflict-minerals-white-paper.html>.

“We Are Working to Make Our Supply Chain ‘Conflict-Free,’” Philips, 2018, <https://www.philips.com/a-w/about/company/suppliers/supplier-sustainability/our-programs/conflict-minerals.html>.

David S. Abraham, *The Elements of Power: Gadgets, Guns, and the Struggle for a Sustainable Future in the Rare Metal Age*, Reprint edition (Yale University Press, 2017), 89.

“Responsible Minerals Sourcing,” Dell, 2018, <http://www.dell.com/learn/us/en/uscorp1/conflict-minerals?s=corp>.

“Apple Supplier Responsibility 2018 Progress Report” (Cupertino CA: Apple, 2018), https://www.apple.com/supplier-responsibility/pdf/Apple_SR_2018_Progress_Report.pdf.

Alexander Klose, *The Container Principle: How a Box Changes the Way We Think*, trans. Charles Marcum II (Cambridge, MA: The MIT Press, 2015).

“Review of Maritime Transport 2017” (New York and Geneva: United Nations, 2017), http://unctad.org/en/PublicationsLibrary/rmt2017_en.pdf.

Zoë Schlanger, “If Shipping Were a Country, It Would Be the Sixth-Biggest Greenhouse Gas Emitter,” *Quartz*, April 17, 2018.

John Vidal, “Health Risks of Shipping Pollution Have Been ‘Underestimated,’” *The Guardian*, April 9, 2009, sec. Environment, <http://www.theguardian.com/environment/2009/apr/09/shipping-pollution>.

“Containers Lost At Sea – 2017 Update” (World Shipping Council, July 10, 2017), http://www.worldshipping.org/industry-issues/safety/Containers_Lost_at_Sea_-_2017_Update_FINAL_July_10.pdf.

Rose George, *Ninety Percent of Everything: Inside Shipping, the Invisible Industry That Puts Clothes on Your Back, Gas in Your Car, and Food on Your Plate* (New York: Metropolitan Books, 2013), 22. Similar to our habit to neglect materiality of internet infrastructure and information technology, shipping industry is rarely represented in popular culture. Rose George calls this condition, “sea blindness” (2013, 4).

id., 175.

Ibid., 176.

Chris Lo, "The False Monopoly: China and the Rare Earths Trade," *Mining Technology*, August 19, 2015, <https://www.mining-technology.com/features/featurethe-false-monopoly-china-and-the-rare-earths-trade-4646712/>.

Kate Hodal, "Death Metal: Tin Mining in Indonesia," *The Guardian*, November 23, 2012, <http://www.theguardian.com/environment/2012/nov/23/tin-mining-indonesia-bangka>.

Cam Simpson, "The Deadly Tin Inside Your Smartphone," *Bloomberg*, August 24, 2012, <https://www.bloomberg.com/news/articles/2012-08-23/the-deadly-tin-inside-your-smartphone>.

Marcus Wohlsen, "A Rare Peek Inside Amazon's Massive Wish-Fulfilling Machine," *Wired*, June 16, 2014, <https://www.wired.com/2014/06/inside-amazon-warehouse/>.

Wurman, Peter R. et al., System and Method for Transporting Personnel Within an Active Workspace, US 9,280,157 B2 (Reno, NV, filed September 4, 2013, and issued March 8, 2016), <http://pdfpiw.uspto.gov/.piw?Docid=09280157>.

John Tully, "A Victorian Ecological Disaster: Imperialism, the Telegraph, and Gutta-Percha," *Journal of World History* 20, no. 4 (December 23, 2009): 559–79, <https://doi.org/10.1353/jwh.o.0088>.

Ibid., 574.

See Nicole Starosielski, *The Undersea Network* (Durham: Duke University Press Books, 2015).

Gary Cook, "Clicking Clean: Who Is Winning the Race to Build a Green Internet?" (Washington, DC: Greenpeace, January 2017), 30, <https://storage.googleapis.com/p4-production-content/international/wp-content/uploads/2017/01/35foac1a-clickclean2016-hires.pdf>.

Mike Ananny and Kate Crawford, "Seeing without knowing: Limitations of the transparency ideal and its application to algorithmic accountability," *New Media & Society* 20.3 (2018): 973–989.

Olga Russakovsky et al., "ImageNet Large Scale Visual Recognition Challenge," *International Journal of Computer Vision* 115, no. 3 (December 1, 2015): 216, <https://doi.org/10.1007/s11263-015-0816-y>.

Ayhan Aytes, "Return of the Crowds: Mechanical Turk and Neoliberal States of Exception," in *Digital Labor: The Internet as Playground and Factory*, ed. Trebor Scholz (London: Routledge, 2012), 80.

Jason Pontin, "Artificial Intelligence, With Help From the Humans," *The New York Times*, March 25, 2007, sec. Business Day, <https://www.nytimes.com/2007/03/25/business/yourmoney/25Stream.html>.

Aytes, "Return of the Crowds," 81.

Jorge Luis Borges, "On Exactitude in Science," in *Collected Fictions*, trans. Andrew Hurley (New York: Penguin, 1999), 325.

Jean Francois Lyotard, "Presenting the Unpresentable: The Sublime," *Artforum*, April 1982.

Yves Citton, *The Ecology of Attention* (Cambridge, UK: Polity, 2017).

Shoshana Zuboff, "Big Other: Surveillance Capitalism and the Prospects of an Information Civilization," *Journal of Information Technology* 30, no. 1 (March 1, 2015): 75–89, <https://doi.org/10.1057/jit.2015.5>.

Vandana Shiva, *The Enclosure and Recovery of The Commons: Biodiversity, Indigenous Knowledge, and Intellectual Property Rights* (Research Foundation for Science, Technology, and Ecology, 1997).

Vandana Shiva, *Protect or Plunder: Understanding Intellectual Property Rights* (New York: Zed Books, 2001).

Alex Hern, “Royal Free Breached UK Data Law in 1.6m Patient Deal with Google’s DeepMind,” *The Guardian*, July 3, 2017, sec. Technology, <http://www.theguardian.com/technology/2017/jul/03/google-deepmind-16m-patient-royal-free-deal-data-protection-act>.

Download

[PDF version of the map](#)

[Publication - A3 format \(Map + Essay \)](#)

Credits and acknowledgements

Authors: Kate Crawford and Vladan Joler

Maps and design: Vladan Joler and Kate Crawford

Published by: SHARE Lab, SHARE Foundation (<https://labs.rs>) and The AI Now Institute, NYU (<https://ainowinstitute.org/>)

Full citation: Kate Crawford and Vladan Joler, “Anatomy of an AI System: The Amazon Echo As An Anatomical Map of Human Labor, Data and Planetary Resources,” *AI Now Institute and Share Lab*, (September 7, 2018) <https://anatomyof.ai>

Acknowledgements: Our deep thanks go to Michelle Thorne and Jon Rogers at the Mozilla Foundation, who invited us to a retreat in summer 2017 where we first conceptualized this project. Thanks to Joana Moll and Meredith Whittaker for their inputs and inspirations on the first drafts of this text. Thanks also to all those who have given feedback, support and insights since, including Alex Campolo, Casey Gollan, Gretchen Krueger, Trevor Paglen, and Sarah Myers West at the AI Now Institute and Olivia Solis, Andrej Petrovski, and Milica Jovanovic at the SHARE Lab and all the wonderful folks from SHARE Foundation.

Finally, thanks to Irini Papadimitriou and all the curatorial staff at the V&A Museum. This map and essay will be on display there as part of the 'Artificially Intelligent' show from Sep 6 - Dec 31, 2018.