

This beautiful map shows everything that powers an Amazon Echo, from data mines to lakes of lithium

James Vincent

That the modern world is a complex place will not have escaped your notice.

We are all dimly, unsettlingly aware that our lives are enmeshed in systems we can't fully comprehend. The last meal you ate probably contained produce grown in another country that was harvested, processed, packaged, shipped, then sold to you. The phone in your hand is the end-product of an even more convoluted chain; one that relies on human labor from mines in Africa, assembly lines in China, and standing desks in San Francisco.

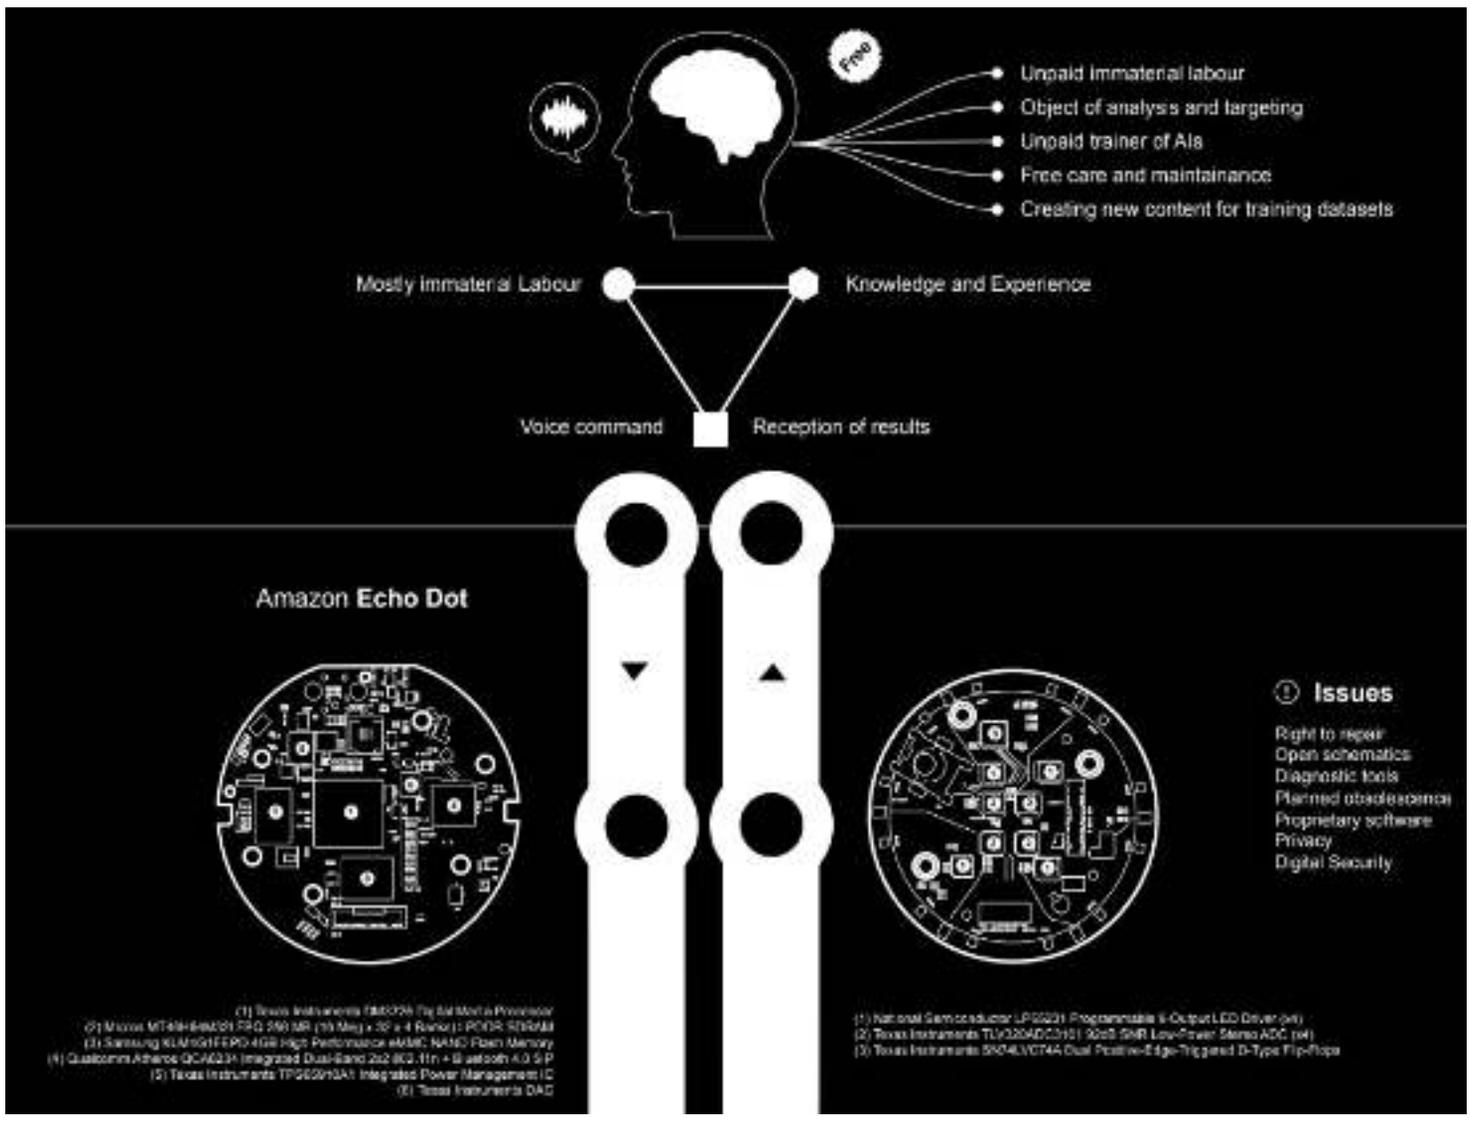
Explaining how these systems connect and the effect they have on the world is not an easy task. But it's what professors Kate Crawford and Vladan Joler have attempted to do in a [new artwork and essay](#), unveiled last Friday at the Victoria & Albert Museum in London.

The main artwork is a huge map, two meters high and five meters across, which traces the systems used to power one of the most complex products of the modern day: an AI-powered gadget, specifically, an Amazon Echo. It's a mess of branching lines in stark black and white, and looks more like the schematics of a nuclear reactor than an everyday gadget. The print is called *Anatomy of an AI system*, but its subtitle explains its scope: "The Amazon Echo as an anatomical map of human labor, data and planetary resources."

Ahead of the unveiling *The Verge* spoke to Crawford, a professor at New York University and co-founder of the [AI Now Institute](#), an organization that examines the social implications of developing artificial intelligence. Crawford and her collaborator Joler, a professor at the Academy of Arts at the University of Novi Sad, say they created *Anatomy* because of the lack of awareness of the structures that support modern gadgets, particularly AI.

"We need to have a deep and more complex conversation about the implications of building artificial intelligence at scale," says Crawford. "And with [*Anatomy*] it's really something you can look at and start to understand as part of a much bigger picture."

This interview has been lightly edited for clarity and brevity. [Click here](#) to see the whole Anatomy of an AI System print and read the accompanying essay.



The first part of Anatomy shows how an Amazon Echo collects data and feedback from human users.

So, first of all, why did you choose an Amazon Echo as the focus of this project?

I was really interested in the very simple voice-based interactions we have with these systems. The Echo sits in your house, looks very simple and small, but has these big roots that connect to huge systems of production: logistics, mining, data capture, and the training of AI networks. It's an entire infrastructural stack you never see. You just give a simple voice command — "Alexa, turn on the lights" — and it feels like magic.

But trying to really investigate and almost do archeology of *how* that magic is working is what this project is about. The Echo is powerful because of this sense of convenience, but when you open up the hood you can see the full cost of it.

Some would say that this has always been the case with technology. And in your essay you show this too, when you talk about how the need to harvest natural latex to insulate undersea cables in the 19th century led to huge deforestation. What's different now?

We've been through many technology booms before that have the concomitant extraction of resources to make it possible. This is certainly a trend. But I will say that the turn to AI is different for two really, really big reasons. First, it's operating at a level which starts to change the way society itself works, because AI systems are being built into the institutions that are most important to us, from health care to criminal justice. These systems change how you interact with the world on each level, so there's this difference in scale.

And second, I think there's the way in which the *consumer* is different which is very particular to AI. With AI devices, consumers exist in a hybrid state where they are someone who buys a product but also a resource, in that their voice commands are retained and analyzed as part of a corpus of training data.

They're also a worker in that they're providing unpaid labor by giving feedback to the system. Their responses help assess the accuracy and usefulness and quality of the AI. And they're also a product in that all of their interests [captured via these interactions] become a profile that you can sell to advertisers. This combination of being a consumer, a worker, a resource and a product, is something that's very new.

If consumers are being used in this way, providing unpaid labor, do you think there needs to be a movement to recapture this value? Should we be able to [sell their data](#) somehow, for example?

First, people need to grasp what's going on. People are just beginning to understand that social media sites, for example, are not just there to share your photos and connect with your friends. They're large systems that are extracting completely different forms of value than you think you're giving by saying hi to your mum or liking a picture of a cat. There's this conceptual shift that's needed to understand that the industry itself has changed, and that there are forms of value being extracted that didn't even exist seven or eight years ago.

The second thing is reform our idea social accountability so it matches our needs. And in terms of what *that* looks like, it's a big question. It's one of the things AI Now has been focusing on, and we have a very big civil law and policy contingent [...] working with organizations like the ACLU to work out how we can draw clear lines around what is acceptable use [of data], what sorts of technologies have serious downsides, and how to we have accountability. And that's something we see as a very big research project.



The Echo is built using a number of minerals, including lithium harvested from the Uyuni Salt Flats in Bolivia.

Photo by Dean Mouhtaropoulos/Getty Images

Anatomy itself is divided into three broad systems, each of which you refer to as an “extractive process.” There’s an extractive process for material resources, one for data, and one for human labor. Why do you think it’s useful to frame these systems in this way, as “extractive”?

All those processes extract value in different ways. When you think of coal mining, to take one example,

you might think of an industry that drove rampant growth, high profits, but that also produced costs that were initially overlooked and uncounted within the economic system. The true picture of resource mining can take decades to emerge. Does data mining have similarly unknown costs that exceed our current economic frame?

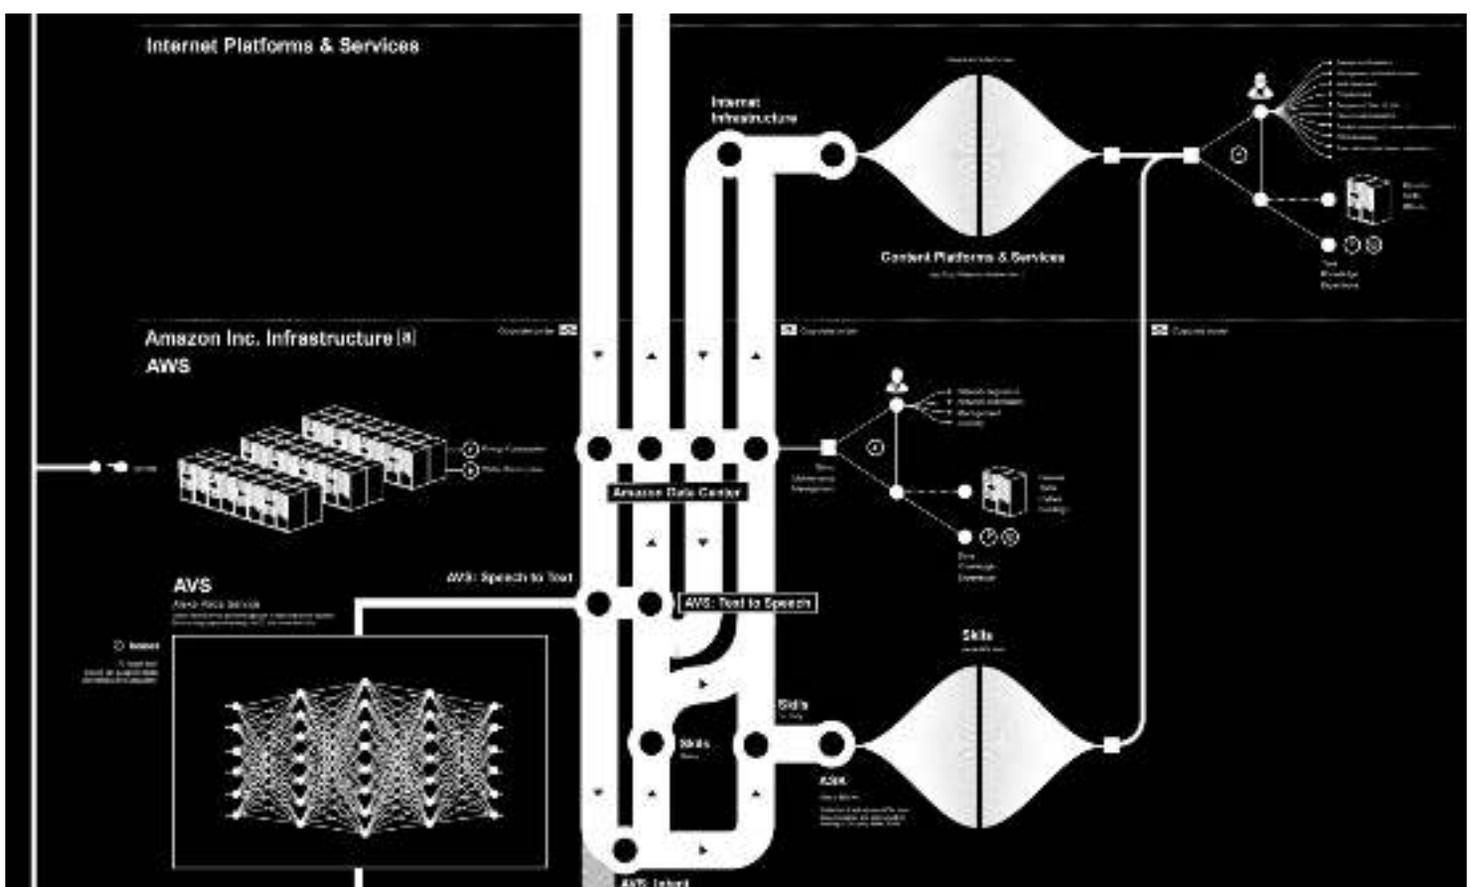
The [Cambridge Analytica scandal](#) is just one of many examples of costs to political systems and civil society that weren't being accounted for. You can see that pattern repeating at many levels: from the labor practices of clickwork, to the mass harvesting of user data, to the rare earth minerals needed to build consumer tech devices. AI systems are extracting surplus value from all kinds of human activities — right down to human emotions and facial expressions — and the costs are often obscured from the end-user and take years to be fully understood.

You present the Echo as the epitome of a certain type of gadget — one that we can't open (because we'll void the warranty) and can't fully control (because the software lives in the cloud and is updated without our permission). How does this paradigm affect our interactions with technology? Or with society?

Exactly. The concept of algorithmic black boxes is now well-known, thanks to the important work of academics like Frank Pasquale. Our project was interested in how that connects with other kinds of black boxes. The Echo itself is a type of box that is extremely hard to examine: a user can't see how it works, how it records data, or how its algorithms are trained. Then there's the hidden logistics around how the simple components inside it are harvested and smelted and assembled, through multiple layers of contractors, distributors, and downstream component manufacturers.

In the essay we write about the example of how it took Intel several years just to understand its own supply chain well enough to be able to ensure no tantalum from the Congo was contained in its microprocessors. Imagine a company that well-resourced, and with highly skilled employees, and a well-established set of record keeping and databases, and it took years to understand its own purchasing patterns!

That shows how hard these processes can be to investigate and analyze from the inside of a company, let alone for the researchers and journalists working on the outside. But that process of telling the stories of production is so important, and needed: it's how we can begin to see into the dizzying complexity of the global production of technology products.



AI systems use interactions with customers to get smarter.

In the essay you talk about the accumulation and concentration of wealth, and the terrible working conditions for those lower down the chain. Do you think AI naturally exacerbates this sort of inequality?

(And, as an aside, I should say I love the way there are multiple “thin crusts” in your map, each of which accumulates value of a different sort. There’s the lithium on the salt lakes in Bolivia; mineral has accrued over millions of years. And there’s the tech elite in Silicon Valley, collecting the value of all that unpaid clickwork by their customers.)

Exactly — those layers are thin. There are a few billionaires at the top of the system, who extract the maximum value, and the further you go down the chains of logistics and production, closer to the raw materials, the more extreme the disparity becomes.

For example, Amazon’s CEO Jeff Bezos made an average of \$275 million a day during the first five months of 2018, according to the *Bloomberg’s* Billionaires’ Index. That’s obviously a large number, but it’s hard to truly grasp until you put it in contrast with other workers much further downstream. Amnesty published a report tracking child labor in the Congolese mines where cobalt is tracked for lithium-ion batteries. In context, a child in one of those cobalt mines would need to work for approximately 700,000 years non-stop to earn the same amount as a single day of Bezos’ income.

These kinds of inequality repeat throughout industrial history, and are not particular to AI. But large-scale AI systems require so much data, infrastructure and maintenance that there are very few companies who are able to build and operate them. Thinking about what that means over time is important if we are going to be able to govern those systems well.

Anatomy of an AI System by Kate Crawford and Vladan Joler can be seen at the Victoria and Albert Museum in London as part of the *Artificially Intelligent* exhibition. You can also explore the map and read the accompanying essay [at the Anatomy website](#).