

Empire of AI is wildly misleading on AI water use

Andy Masley

Update 12/18/2025: The author has updated the book based on my two main criticisms in this post. [She explains the changes in her blog here](#). I think the chapter now leaves readers with a better understanding of how water's being used by AI. I have some remaining background disagreements with some of the presentations of the issue, but the author was super direct here and engaged a lot, gave me a lot of back and and corrected the two major issues I found, so I want to commend her for that. Really great to see.

Update 11/17/2025: the author [took the time to respond to me below](#). While I'm very grateful, the materials she sent actually seems to confirm my main criticism and I'm now very confident a key number in the book is 1000x too large and needs to be revised. [I summarize everything in my reply here](#).

I was taking a break from posting about AI and the environment, but after reading parts of Karen Hao's book [Empire of AI](#), I've stumbled on such wildly misleading claims that have so far gone unaddressed that I've felt the need to counter them here. Within 20 pages, Hao manages to:

Claim that a data center is using 1000x as much water as a city of 88,000 people, where it's actually using about 0.22x as much water as the city, and only 3% of the municipal water system the city relies on. She's off by a factor of 4500. This is the single largest error in any popular book that I've found on my own, and to my knowledge I'm the first person to notice it.

Imply that AI data centers will consume 1.7 trillion gallons of drinkable water by 2027, while the study she's pulling from says that only 3% of that will be drinkable water, and 90% will not be consumed, and instead returned to the source unaffected.

Paint a picture of AI data centers harming water access in America, where they don't seem to have caused any harm at all.

Frame Uruguay as using an unacceptable amount of water on industry and farming, where it actually seems to use the same ratio as any other country.

Frame the Uruguay proposed data center as using a huge portion of the region's water where it would actually use ~0.3% of the municipal water system, without providing any clear numbers.

These are all the significant mentions of data centers using water in the book. Read in this light, the chapter becomes somewhat ridiculous, because the rest includes descriptions of brutal acts of torture and plunder under colonialism, and then frames data center water use as a continuation of that same colonialism. If instead you see data centers using water in other countries as part of a simple trade the countries are making to get more taxable industry in the area, and that doesn't seem to harm water access, the central narrative thrust of the chapter becomes false.

This book has been very popular, and influenced the AI/environment conversation. One of the most common replies I've received to my water arguments is that I should read it. So I did, and I came away kind of shocked at how badly it covered the issue. Hao's points about water are all in the chapter "Plundered Earth."

My only ask for people writing about AI and the environment is that at the end, readers are left with a more accurate picture of how energy and water is being used overall in the regions covered, and not left with a contextless impression of AI as a huge environmental offender in places where it's not. Every single time water's mentioned in this book, the reader is left with a worse understanding than they came in with.

I make my case for each with easily accessible numbers. All quotes and page numbers are taken from the Kindle edition.

[A misleading projection of how much water AI is expected to use](#)

[A massive factual error: Hao claims a data center would use 4500x as much of a city's water as the actual value](#)

[A weird framing of Uruguay's water use](#)

[Misleading presentations of data center water issues in America](#)

[Conclusion](#)

The first time the book goes into detail on water is here:

The land and energy required to support these megacampuses are but two inputs in the global supply chain of data center expansion. So, too, is the extraordinary volume of minerals including copper and lithium needed to build the hardware—computers, cables, power lines, batteries, backup generators—and the extraordinary volume of potable—yes, potable—water often needed to cool the servers. (The water must be clean enough to avoid clogging pipes and bacterial growth; potable water meets that standard.) According to an estimate from researchers at the University of California, Riverside, surging AI demand could consume 1.1 trillion to 1.7 trillion gallons of fresh water globally a year by 2027, or half the water annually consumed in the UK.

(pp. 277-278)

The study mentioned is "[Making AI Less Thirsty](#)." The study does **not** say that AI demand will *consume* 1.1-1.7 trillion gallons of water annually. Hao seems to be getting this number from this part of the study:

Even considering the lower estimate, the combined scope-1 and scope-2 water **withdrawal** of global AI is projected to reach 4.2 – 6.6 billion cubic meters in 2027

4.2-6.6 billion cubic meters = 1.1-1.7 trillion gallons.

Withdrawal is very different from consumption. Withdrawal means the amount of water taken from a local source. Consumption is the amount of water taken *and not returned* to the local source. Many of the ways we use water (especially power plants) withdraw water temporarily from local sources, use it briefly, and then return it unaffected. This is called non-consumptive withdrawal. It's like diverting part of a river to run

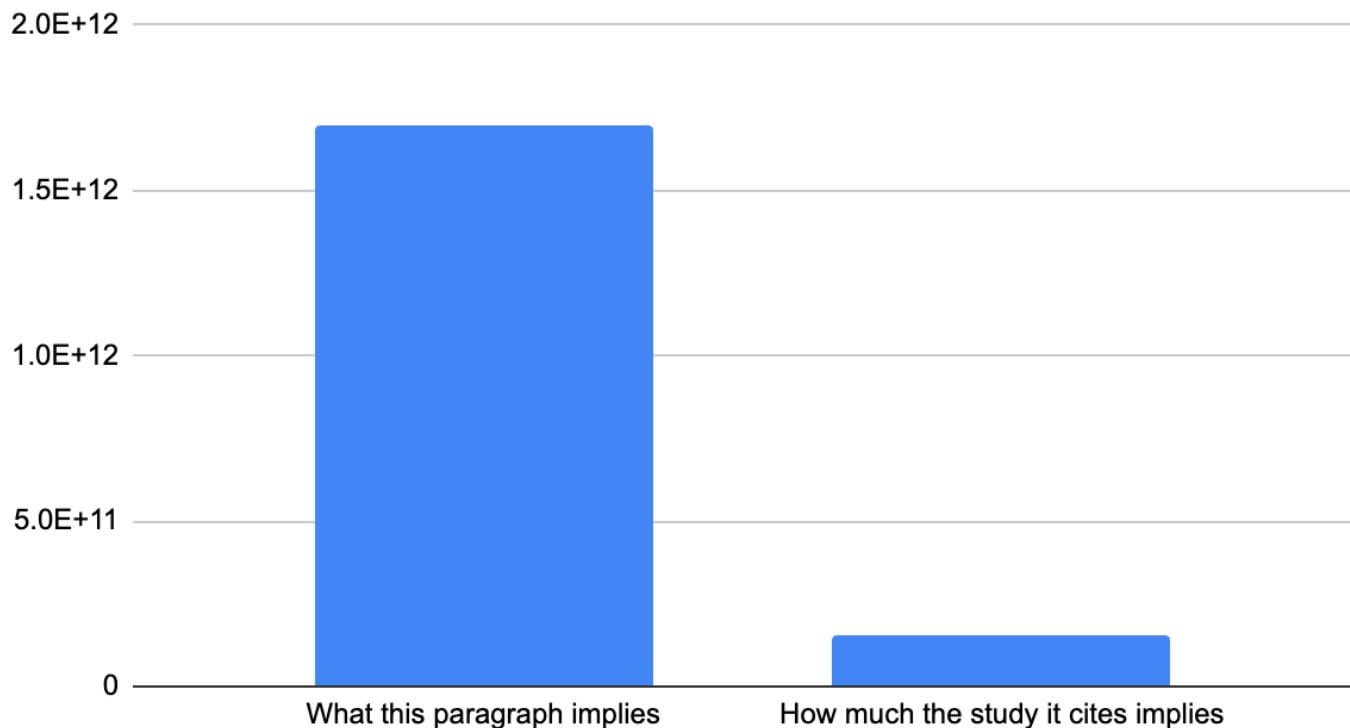
near a mill where you have some water wheels, and the diverted water then returning to the main flow after. Consumptive withdrawal is like sucking up and evaporating the water. The main water issue for most regions is water consumption, not just water withdrawal.

How different is consumption and withdrawal for AI? Very! The next sentence in the study Hao cites says:

Simultaneously, a total of 0.38 – 0.60 billion cubic meters of water will be evaporated and considered “consumption” due to the global AI demand in 2027.

This is equal to 100–158 billion gallons, only 10% of the number Hao reports. If someone’s writing a book that covers AI water use in depth, I would expect them to know the difference between consumption and withdrawal, especially when the consumption number is given in the very next sentence of the study. It’s very weird that Hao misidentifies withdrawals as consumption here.

Gallons of water consumed by AI in 2027



Why are withdrawals so much bigger than consumption for AI? It’s mainly because the authors of the study are measuring how much water offsite power plants use to generate electricity for AI as part of its total withdrawals number. The vast majority of water that power plants use is withdrawn and then returned to the local source (non-consumptive use on this graph).

So the withdrawal number is going to be way bigger than the consumption number. Withdrawal is just not nearly as much of an issue for local water access as water consumption.

But the way this is being measured gets weirder. Even just measuring consumptive water use, the amount of water actually removed permanently from a local source, the study authors still find that the vast majority of water consumed by AI is offsite in nearby power plants that data centers draw from. For the US average, only ~15% of the water AI is consuming is actually happening in data centers themselves (the on-site water

column).

Basically none of the water power plants use is potable. The only potable water used for AI is in data centers themselves. So the study Hao is citing is not only saying AI will consume just 10% as much water as she says, only 15% of that 10% will be drinking water.

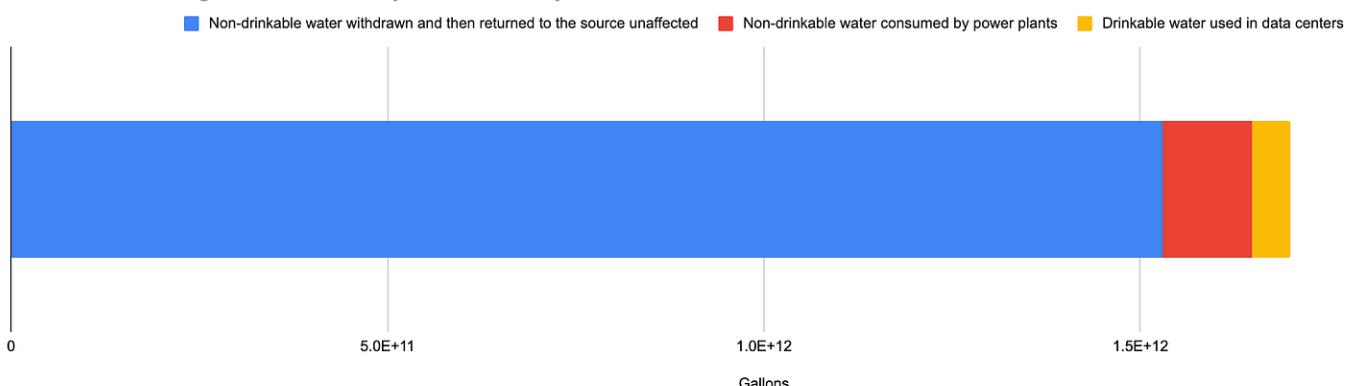
Read Hao's paragraph again. It strongly implies that AI will be using 1.1-1.7 trillion gallons of drinking water:

The land and energy required to support these megacampuses are but two inputs in the global supply chain of data center expansion. So, too, is the extraordinary volume of minerals including copper and lithium needed to build the hardware—computers, cables, power lines, batteries, backup generators—and the extraordinary volume of potable—yes, potable—water often needed to cool the servers. (The water must be clean enough to avoid clogging pipes and bacterial growth; potable water meets that standard.) According to an estimate from researchers at the University of California, Riverside, surging AI demand could consume 1.1 trillion to 1.7 trillion gallons of fresh water globally a year by 2027, or half the water annually consumed in the UK.

(pp. 277-278)

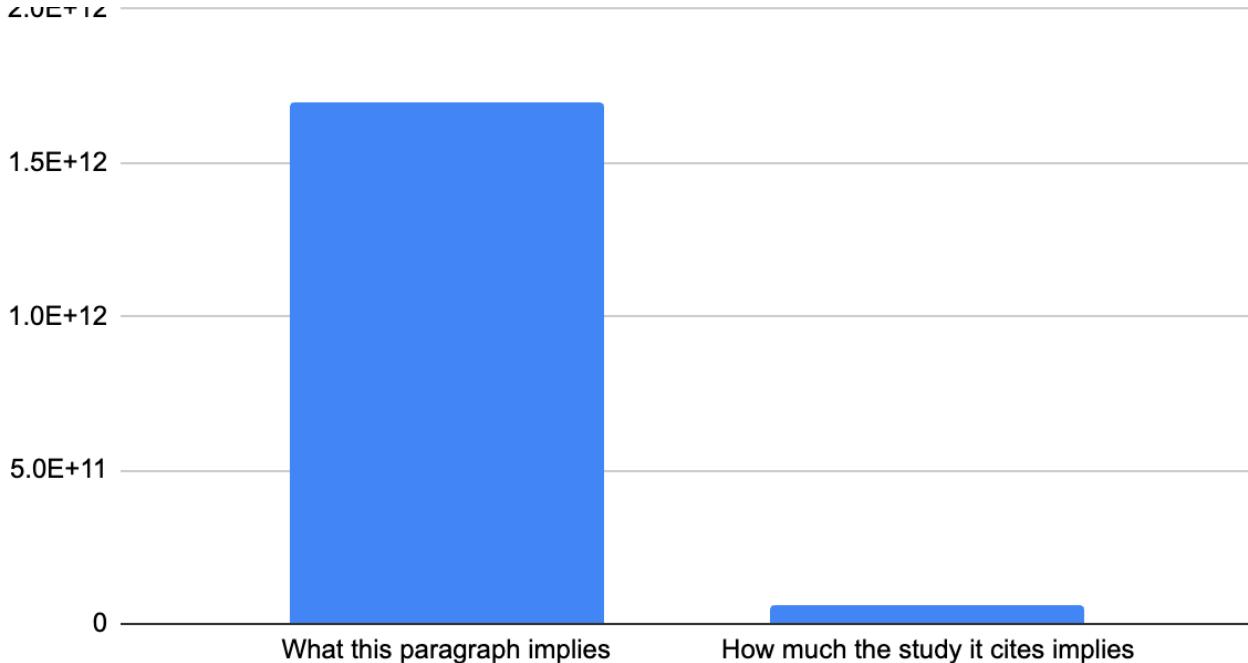
In reality, the study she's citing specifically says that the water used by AI data centers themselves (the only place potable water is used) could consume 150 – 280 billion liters. This is 40-74 billion gallons. This is just 3.6% of Hao's number. Further, the study notes that only 80% of the water used by Google's data centers is potable, bringing the number down 32-59 billion gallons, just 3% of Hao's number. Here's how much of Hao's number is non-drinkable water returned to the source unaffected (blue), non-drinkable water consumed by power plants and data centers (red), and drinkable water used in data centers themselves (yellow).

What "1.7 trillion gallons of water by 2027" actually means



The paragraph compares water to half the UK, and leaves the reader to infer that this is all potable drinking water. [London uses 2.6 billion liters of water every day](#), 690,000,00 gallons, which is 252 billion gallons per year. So the actual amount of potable drinking water this study Hao is citing is projecting that all global AI will use is at most 20% of just the London water system, and 1.5% of the UK's water use.

Drinking water used by AI in 2027



Hao repeats this misconception in her Atlantic article [AI is taking water from the desert](#):

Public data hint at the potential toll of this approach. Researchers at UC Riverside [estimated](#) last year, for example, that global AI demand could cause data centers to suck up 1.1 trillion to 1.7 trillion gallons of fresh water by 2027.

This leaves the reader with the very strong impression that the physical data center buildings themselves will “suck up” and not return up to 1.7 trillion gallons. In reality, 90% of this “sucked up” water will be returned to the source it was drawn from, and only 3% of it will be in physical data centers themselves.

I have some disagreements with the study itself. Among other things, a lot of the water it measures AI “using” offsite is actually water evaporated in lakes dammed by hydroelectric plants to generate electricity, but it ignores the water these lakes recapture from rain. [In similar studies this evaporated water ends up being 2/3rds of the total water use estimate of AI in general, so I think the actual offsite use should be significantly lower](#). The study is also two years old and I think we just have better estimates based on way more data now, specifically [this report](#). We also have [better numbers for chatbots](#) measured by companies themselves, which show them using significantly less water than the estimates we had available in 2023. [I also argue here](#) that whether water is potable doesn’t really matter, because what actually harms the availability of drinkable water is access to total freshwater, not potable water. It’s relatively easy to turn freshwater potable.

The chapter later turns to Chile, focusing a lot on the municipality of [Quilicura](#) on the outskirts of Santiago. The chapter covers ways that colonialism has seriously harmed the people and local environment and water there in the past, and then turns to the Google data center built there. The Google data center is framed as a continuation of the colonialism. It then goes on:

This unique combination—a history of neglect and a precious water source—created fertile ground for the blossoming of several environmental activist groups who were used to being watchdogs and were fiercely protective against the extraction of their resources. That summer, as Google filed a report with Chile’s

environmental agency for approval of its data center—a largely rubber stamp process—MOSACAT, a water activist group, began combing through all 347 pages of the filing. Buried in its depths, Google said that its data center planned to use an estimated 169 liters of fresh drinking water per second to cool its servers. In other words, the data center could use more than *one thousand times* the amount of water consumed by the entire population of Cerrillos, roughly eighty-eight thousand residents, over the course of a year.

MOSACAT found this unacceptable. Not only would the facility be taking that water directly from Cerrillos's public water source, it would do so at a time when the nation's entire drinking water supply was under threat. In 2019, as with Iowa and Arizona, Chile was already nine years and counting into a devastating and historically unprecedented megadrought.

(pp. 288-289).

Look at this line again:

In other words, the data center could use more than ***one thousand times*** the amount of water consumed by the entire population of Cerrillos, **roughly eighty-eight thousand residents**, over the course of a year.

Hao justifies this number in the notes section at the end of the book:

In other words, the data: The Google environmental impact report to SEA stated that the data center could use 169 liters of potable water a second, or 5,329,584,000 liters a year. According to the water service authority in Cerillos, the municipality consumed 5,097,946 liters in all of 2019, the year Google sought to come in; 5,329,584,000 liters a year divided by 5,097,946 liters a year equals 1,045.

(p. 454)

Something isn't adding up here. 5,097,946 liters a year is 13,966 liters per day. Dividing by 88,000 residents says that each resident is using 0.2 liters of water per day. That's about this much:





This implies that a city of 88,000 people is using as much water per day in total as [a single shower head left running](#). Adults need to drink ten times as much as this, [2-4 liters of water per day, to stay alive](#). The average citizen of Chile buys [180 liters per day](#) from their municipal water systems. I think the actual amount of municipal water supplied to people in Cerrillos is 900 times as much as Hao is claiming here.

I can't find the study Hao is referencing, but I did find [this document](#) from the local government saying that the potable water consumption for the region of Maipú, Cerrillos, part of Estación Central (which the water system seems to serve) in 2019 (the year the study was conducted) was 54,148,639 m³. That's 54,148,639,000 liters in a year, 148,352,000 liters per day. The total population the system is serving seems to be ~650,000 people.

So each person using this system is using ~230 liters of water per day. That's way more in line with Chile's average water use. The reason this is larger is probably that it also includes commercial buildings.

Hao got this number wrong by 3 orders of magnitude and reported the data center as using 1000 times as much water as a city of 88,000 people. I can't get over how crazy this is. She implied that a single building is using as much water as a city of 88 million people. That's over 4x the entire population of Chile! Readers are leaving this book assuming that Google built a single building in Chile that's using 20% of all the residential water on the whole continent. No wonder everyone's freaking out so much about AI and water use.

I think I know what happened: this municipality seems to report their water use in cubic meters. Each cubic meter is 1000 liters. Hao's number reported in liters is noticeably very close to the cubic meter value I see in my sources. I suspect she got her wires crossed and somehow recorded m³ as liters, making the city appear to use 1000x less water than it does. But she and others should have caught this in editing. I'm not aware of any building anywhere that uses 1000x as much water as a nearby city. That would be crazy! **Note:** Since writing this Hao [replied here](#), and the material she sent confirmed to me that this is actually what happened, [I explain here](#). I expect the book to be corrected.

But it gets worse! This is not the only misleading thing about how this number is presented.

This 169 L/second number [was reported by a local paper as the maximum permitted amount](#). Nowhere in Hao's writing does it make it clear that this number is the very maximum upper bound of how much water the data center will use, not the average normal water use. Reporting the maximum permit for water draw as the normal amount it will actually draw is very misleading. Here's my explanation for why [from my main post on AI and water](#):

Many articles about current or future AI data centers report the number in the water permit they apply for as the amount of water they actually use day to day. But this almost never happens.

When a data center is being built, the company needs to obtain water use permits from local authorities before construction. At this stage, they have to estimate their maximum possible water consumption under worst-case scenarios:

All cooling systems running at full capacity

Peak summer temperatures

Maximum IT load (every server rack filled and running)

Minimal efficiency from cooling systems

The permit needs to cover this theoretical maximum because regulators want to ensure the local water infrastructure can handle the demand and that there's enough water supply for everyone. It's easier to get a higher permit upfront than to come back later and request more, so data centers are incentivized to aim high.

Actual water usage is always significantly lower than what the permits allow, because they're designed with the absolute worst conditions in mind. But many popular articles about how much water data centers use give the number on the water permit, not how much the data center actually uses.

We don't have information on how much water the data center was actually expected to use, but we can at least compare it to other Google data centers permitted for the same amount of water to get a rough guess. Google's data centers in [The Dalles Oregon were permitted to draw the same maximum amount of water per day](#). The actual amount they used in a year was [275 million gallons](#), which is 0.75 million gallons per day. So another data center with this same permitted amount ended up using just 20% of its actual permit.

Maybe the data center would use more water because of different climates? The only times when data centers seem to use significantly more water is during heat spikes at the hottest times of the year, and The Dalles's hottest month is on average warmer ([88 F average in the Dalles, 71 F average in Cerrillos](#)).

This all points to the actual expected water usage of this data center being around 1 million, not 4 million, gallons per day. 1 million gallons per day is 3,800,000 liters per day, or 1,767,000,000 liters per year.

So in a place where the local water system sells 54,148,639,000 liters per year, this data center's normal operations would have raised the total water system's demand by 3%.

Compare Hao's sentence on this:

the data center could use more than **one thousand times** the amount of water consumed by the entire population of Cerrillos, **roughly eighty-eight thousand residents**, over the course of a year.

To what I think is the correct description:

A municipal water system serves 650,000 people. If a data center is built in the region, it would raise the system's annual demand by 3%, equivalent to the water used by 19,500 residents.

Still significant, but drastically different from the number Hao gives. This is the main story of a data center harming water access in the book, and the main number given for its water usage is 4500 times as high as the correct value. Ridiculous!

This book has been reviewed almost 1000 times on Amazon:

This also shows up in [an interview did with Tech Policy Press](#):

That's where Google wanted to add its data center and then tap into that freshwater resource to cool its data centers. The amount of water that they were proposing was to consume more than a thousand times the amount of water that residents in that community would typically consume. They went through, they fought tooth and nail just to get Google's attention because they not only had to make enough noise to pressure Google Chile, they had to make enough noise to then get that all the way back to headquarters in Mountain View.

And from what I can tell not a single person has noted "Oh hey the central story of a data center harming water access is assuming its water use is 4000x as big as it actually is." Way, way, way too many people are reading about AI water issues uncritically. Why is it so easy for me to so quickly stumble on these gigantic errors in popular writing, that it seems like no one else has found? This book has been really popular with environmentalist critics of AI. Surely someone would have noticed this mistake?

Hao [says here](#) that she spent time with activists in Chile focused on data centers. Surely she would have noticed somewhere that these numbers were off. She specifically says they were "fighting tooth and nail to stop these data centers from taking all their drinking water" so conversations must have focused on the amount of water the data centers would actually use. How did this go so wrong? Further, she [has a degree in mechanical engineering from MIT with a minor in energy studies and was a senior editor on AI at the MIT Tech Review](#). Surely she knows that there's just no way any single physical building we've ever built would use as much water as 1000 times a city of 88 thousand people. The data center with the most water used in the world as of writing seems to be Google's Iowa data center, [with a billion gallons used every year](#). Even this one is just using half as much water as Cerrillos. This new data center would have to be 2000 times as large as any previously existing data center to make her claim true. How did someone covering AI for a living for a decade, with an energy studies degree from MIT, who spent time in the region itself with water activists, not immediately clock that this is way too big?

The next section covers Uruguay. It opens with a discussion of drought in the country, and includes a strange observation:

The water crisis emerged from the compounding effects of climate change and a failure of the state's allocation of freshwater resources: **In Uruguay, more than 80 percent of the country's fresh water goes to industry instead of human consumption—most notably, cash crop agriculture.** These include industrial farms for soybeans and rice, and for trees that feed into paper production. Most such farms are run not by local companies but by multinationals that export what they grow and show little accountability for Uruguay's natural environment. Their activities deplete the nutrients in the soil, making it more difficult to grow actual food, and pollute the country's water streams with a volume of fertilizers that makes Uruguay one of the world's largest per capita fertilizer consumers and causes unusually high rates of cancer.

(p. 293)

The reason I think this is strange is that in basically every country, including the US, 80% of water is used on industry and commercial buildings rather than household consumption. I think a lot of readers don't know that. This framing implies Uruguay is holding back water from locals who need it more. In reality they

have the same split of water as every other country. Water's very important to agriculture, industry, and commercial buildings. Hao presents this as a sign that Uruguay's water is being used by big evil industries and being kept from the people. In reality it seems to have the same split of water as any other country.

The chapter then goes to a sociology researcher named Daniel Pena.

So when Google arrived, Pena was vigilant. During his regular scans of the Uruguayan environmental ministry's website, which lists major industrial projects, he came across the company's proposal for the data center. Pena had read about hyperscalers using potable water, even during major droughts, and the activism of communities like MOSACAT that had resisted the projects. But when he downloaded the details of the project, the water numbers were marked as confidential. After submitting a public information request, which he had successfully done around twenty times, the ministry continued to withhold the numbers, saying they were proprietary information. Pena wondered what they were hiding and worried about the precedent it would set for other cloud companies that would inevitably begin to eye Uruguay, following Google's lead, for their own expansion. So he evoked the water clause in the constitution. With the help of a lawyer friend who was willing to work pro bono, he sued the ministry.

In March 2023, four months later, Pena won the case in a surprising victory. The environmental ministry revealed that Google's data center planned to use two million gallons of water a day directly from the drinking water supply, equivalent to the daily water consumption of fifty-five thousand people. With much of Montevideo receiving salt water in their taps not long after, the revelations were explosive. Thousands of Uruguayans took to the streets to protest Google and all of the other industries that had led the government to squander the country's precious freshwater resources.

....

Near the end of 2023, Google silently updated its proposed data center in Uruguay to use a waterless cooling system and said it would reduce the facility to a third of its size.

(pp. 294-295)

I'm not going to lecture locals in a far away country what they should and should not build in their country, but I do want to know how nefarious Google was being here. Again, I suspect this 2 million gallons per day was an upper limit, and if it's anything like other Google data centers the number's probably closer to 400,000 gallons per day. The municipal water system serving the city of 1.7 million people there seems to use 500,000 m³ per day. This is 132 million gallons. So if the data center had been built, it would be using about 0.3% of the municipal water system. Again, not nothing, but the reader is left with zero sense of scale here. It should be pretty easy to provide that.

Hao repeatedly mentions data centers in America built in water stressed areas. Each mention I think is misleading. Take this example from Iowa:

Altman and other executives never brought up the data centers' environmental toll in company-wide meetings. As OpenAI trained GPT-4 in Iowa, the state was two years into a drought. **The Associated Press later reported that during a single month of the model's training, Microsoft's data centers had**

consumed around 11.5 million gallons, or 6 percent, of the district's water. GPT-4 had trained there for three months. (A Microsoft spokesperson said the company is working to increase its water efficiency by 40 percent above its 2022 baseline and to replenish more water than it consumes across its global operations by 2030, with a focus on the water-stressed regions where it works.)

A month of using 11.5 million gallons means each day OpenAI used 380,000 gallons of water. [Corn in Iowa uses between 0.1-0.2 inches of water to grow per day.](#) 0.1 inches of water over 1 acre is 27,154 gallons. So OpenAI was using as much water as 14 acres of an Iowa corn farm, or 0.02 square miles. [The average Iowa corn farm is 346 acres](#). This amount of water is equivalent to Sam Altman purchasing 4% of a single Iowa farm to grow corn for his employees. Here's that area on a map (the yellow box).



Does Hao's paragraph get this magnitude across? If you heard a company had bought 4% of a corn farm, how big of a problem would you assume this is for regional water access? What if the tech company were using this 4% to grow something that half a billion people would use every single week for a year?

The book turns next to Arizona.

Arizona, too, faces a severe water crisis. In 2022, as Microsoft laid the groundwork for Phase 3, a study in *Nature Climate Change* found that the Southwestern US had been facing the worst drought it had seen in

over a thousand years. That drought, combined with severe mismanagement, has drained the Colorado River, which Arizona and six other states rely on for fresh water, to dangerously low levels. Without drastic action, the river could cease to flow. The shortage compounds a power crisis, as climate change has slammed the region with relentless record-breaking temperatures and families have cranked up their air-conditioning. The region relies in part on hydropower from the Hoover Dam and water-cooled nuclear power plants. In other words, it needs water to produce more energy. In 2023, the Phoenix metro area hit multiple new heat records as well as the worst year for heat-related fatalities, which surged at least 30 percent from 2022 to over six hundred dead. “All things,” says Tom Buschatzke, the director of the Arizona Department of Water Resources, “are converging in a challenging direction.”

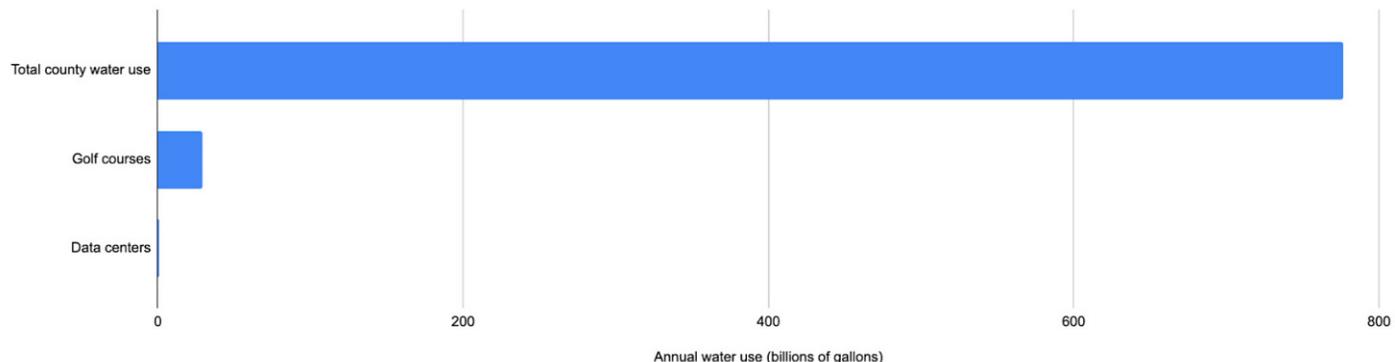
(pp. 280-281)

Notice that Hao does not give a number for how much water Arizona data centers are using. Here's a section of my main water post where I look at how much water data centers in Maricopa County (the main place they're being built there) are using compared to the total county water and golf courses, and how much tax revenue they're bringing in per unit water:

Take Maricopa County in Arizona. The county is home to Phoenix, and is in a desert where water is pumped in from elsewhere. It's also one of the places in the country where the most new data centers are being built.

[Circle of Blue](#), a nonprofit research organization that seems generally trusted, [estimates that data centers in Maricopa County will use 905 million gallons of water in 2025](#). For context, Maricopa County golf courses [use 29 billion gallons of water each year](#). In total, the county uses [2.13 billion gallons of water every day](#), or 777 billion gallons every year. Data centers make up 0.12% of the county's water use. Golf courses make up 3.8%.

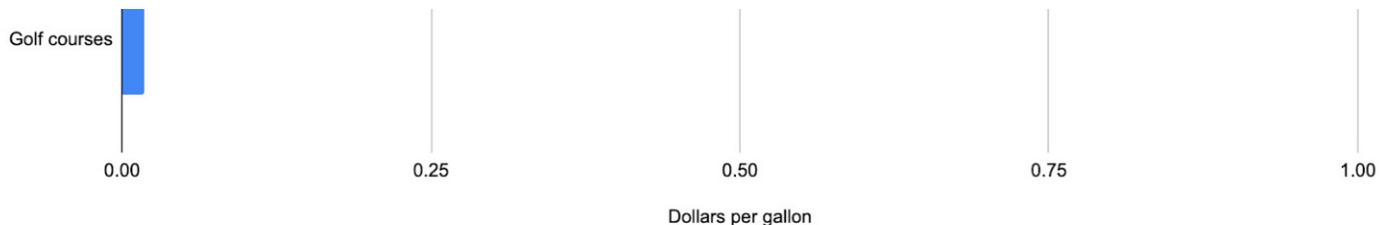
Water use in Maricopa County



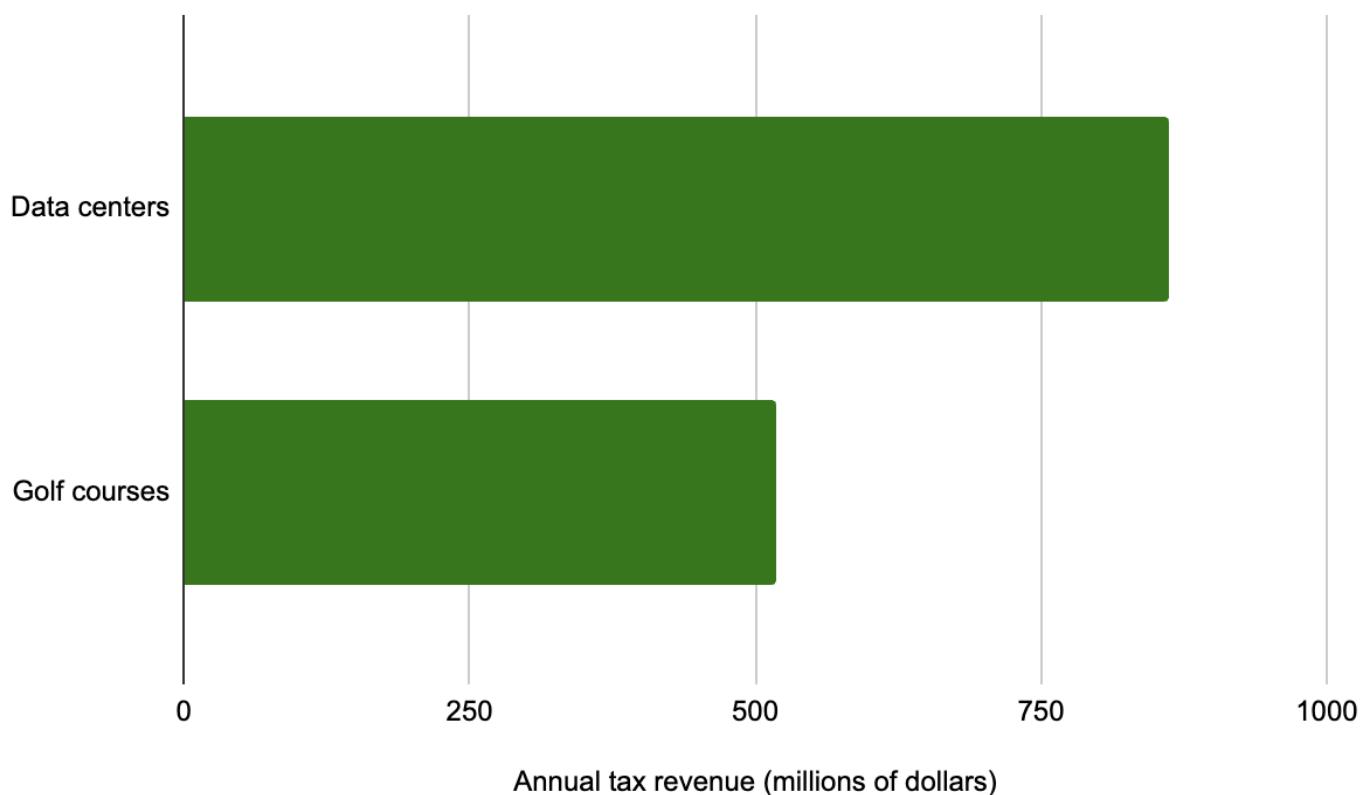
Data centers are so much more efficient with their water that they generate 50x as much tax revenue per unit of water used than golf courses in the county:

Tax revenue per gallon of water used





So even though data centers are using 30x less water than golf courses, they bring in more total tax revenue:³



Some people see this, and react with something like “Well I don’t think golf courses OR data centers should be built in the desert.” At some point this becomes an argument against anyone living in deserts in the first place. If you want to have a gigantic city in the desert, like Phoenix, that city needs some way of supporting itself with taxes, and giving jobs to the people who live there. Most industries use significant amounts of water. If Phoenix is going to exist, it’s going to need private industries built around it that are using some water. We have two options here:

Build industries that generate huge amounts of tax revenue relative to the water they use. Data centers fall into this category (though they don’t provide many jobs).

Do not build cities in the desert in the first place.

Arguments against data centers existing in the desert because they harm water systems there also often apply to building cities in the desert in the first place. It’s fine and consistent to say that Phoenix shouldn’t exist because it’s unnaturally pumping water from hundreds of miles away, but it’s inconsistent to say that Phoenix should exist, that its water bills should be kept as low as possible, but also that no industries that use any water should be built there.

Simply reporting data centers as “being built in water scarce areas like Arizona” I think leaves out way too much important context to leave readers more informed.

Hao opens the water section with this observation:

Another study found that in the US, one-fifth of data centers were already drawing that water before the generative AI boom from moderately or highly stressed watersheds due to drought or other factors.

(p. 278)

This is true, but data centers are subject to the laws of supply and demand like any other business. In places where water’s more scarce and expensive, data centers are more likely to use other methods of cooling. In Arizona, where lots of data centers are being built, they’re using very little water relative to many other industries, and the tax revenue they’re bringing in. Overall, [I haven’t found a single place where the normal operation of a data center has caused any issues for water access anywhere in America](#). We have plenty of other industry in medium and high water stress areas. Those communities also benefit from the tax revenue industry and commercial buildings bring in. But you wouldn’t know that from the book’s coverage.

Hao is smart. [She has a degree in mechanical engineering from MIT with a minor in energy studies and was a senior editor on AI at the MIT Tech Review](#). If she wanted to give readers a full complete picture of where AI data centers are fitting into the broader environmental picture, this would be extremely easy for her. But every individual mention of water in Empire of AI leaves the reader less informed.

More broadly, this is a terrible sign for just how bad the public understanding of AI water use is. This book has been out for 6 months, received 1000 reviews on Amazon, was recommended by Time, The New York Times, Vulture, The New Yorker, The Economist, Financial Times, and Kirkus Review. Hao also thanks her fact-checking team in the acknowledgements:

To my incredible fact-checking team: Lindsay Muscato, Matt Mahoney, Rima Parikh, and Muriel Alarcón. All four of them fastidiously combed through the draft, cross-checking the labyrinth of details against documents and sources, and stress-testing my word choices. Matt also supported early research in my book, and Lindsay fielded many calls from me to serve as the most patient sounding board, while Rima somehow turned her fact-checking notes into standup comedy. They are all lifesavers.

(p. 424)

And yet after all these people read the book, I’m somehow the first person to notice and comment on the fact that no building, anywhere on Earth, is using anywhere close to 1000 times as much water as a community of 88,000 people use. How did this happen? Why did no reader, anywhere, let the author know about this massive glaring issue?

I think this happened because **almost no one covering this issue is actually just looking at the numbers on AI and water use**. [The people who do are mostly experts less engaged in the public conversation](#).

Almost everyone covering this seems to just be going with vibes. The fact that so, so many journalists, self-identified environmentalists, and members of the public could just nod along to the claims here and let them slip by provides a useful intuition for how it could possibly be that after all the ink spilled on it by so many

professional people, the AI water issue is still somehow [fake](#).

[All my writing on AI and the environment](#)

[The AI water issue is fake](#)

[Data centers don't harm water access at all anywhere in America](#)

[More Perfect Union videos are wildly deceptive on data center water use](#)

[Contra the UK government, please don't delete your old photos and emails to save water](#)