SPLITTING THE COOKIE:
STRUCTURAL BIAS IN DIGITAL TECH
Meredith Broussard

Often, when people talk about making more equitable technology, they start with "fairness."

This is a step in the right direction. Unfortunately, it is not a big enough step in the right direction. Understanding why starts with a cookie. (A sweet and crunchy one, like you would eat—not like the cookie that you have to accept when visiting a web page.)

When I think of a cookie, I think of the jar my mother kept on our yellow Formica kitchen counter throughout my childhood. It was a large porcelain jar with a wide mouth, and often it was filled with homemade cookies. The porcelain lid clanked loudly every time a kid opened the jar for a snack. If I heard my little brother opening the jar, I wandered into the kitchen to get a cookie too. My brother did the same if he heard me. It was a mutually beneficial system—until we got to the last cookie.

When there was only one cookie left in the jar, my brother and I bickered about who got it. It was inevitable. My brother and I squabbled about everything as kids. (As adults, we work in adjacent fields, and there's still a fair bit of good-natured back and forth.) At the time, our cookie conflicts seemed high-stakes. There were often tears. Today, as a parent myself, I admire my mother for stepping in hundreds of times to resolve these kinds of kid disputes. I admire her more for the times she didn't step in and let us work out the problem on our own.

If this story were a word problem in an elementary school math workbook, the answer would be obvious. Each kid would get half of the cookie, or 50 percent. End of story. This mathematical solution is how a computer would solve the dispute as well. Computers are machines that do math. Everything computers do is quite literally a computation. Mathematically, giving each kid 50 percent is fair.

In the real world, when kids divide a cookie in half, there is normally a big half and a little half. Anyone who has had a kid or been a kid can tell you what happens next: there's a negotiation over who gets which half. Often, there is more arguing at this point, and tears. If I was in the mood for peaceful resolution when I was a kid, I would strike a deal with my little brother. "If you let me have the big half, I'll let you pick the TV show that we watch after dinner," I'd offer. He would think for a moment and decide that sounded fair. We would both walk away content. That's an example of a socially fair decision. My brother and I each got something we wanted, even though the division was not mathematically equal.

Social fairness and mathematical fairness are different. Computers can only calculate mathematical fairness. This difference explains why we have so many problems when we try to use computers to judge and mediate social decisions. Mathematical truth and social truth are fundamentally different systems of logic. Ultimately, it's impossible to use a computer to solve every social problem. So, why do so many people believe that using more technology, more computing power, will lead us to a better world?

The reason is technochauvinism.

WHAT IS TECHNOCHAUVINISM?

Technochauvinism is a kind of bias that considers computational solutions to be superior to all other solutions. Embedded in this bias is a priori assumption that computers are better than humans—which is actually a claim that the people who make and program computers are better than other humans.

Technochauvinism is what led to the thousands of abandoned apps and defunct websites and failed platforms that litter our collective digital history. Technochauvinist optimism led to companies spending millions of dollars on technology and platforms that marketers promised would "revolutionize" and digitize everything from rug-buying to interstellar travel. Rarely have those promises come to pass, and often the digital reality is not much better than the original. In many cases, it is worse.

Behind technochauvinism are very human factors like self-delusion, racism, bias, privilege, and greed. Many of the people who try to convince you that computers are superior are people trying to sell you a computer or a software package. Others are people who feel like they will gain some kind of status from getting you to adopt technology; people who are simply enthusiastic; or people who are in charge of IT or enterprise technology.

Mathematical truth and social truth are fundamentally different systems of logic.

Technochauvinism is usually accompanied by equally bogus notions like "algorithms are unbiased" or "computers make neutral decisions because their decisions are based on math." Computers are excellent at doing math, yes, but time and time again, we've seen algorithmic systems fail at making social decisions.

Algorithms can't sufficiently monitor or detect hate speech, can't replace social workers in public assistance programs, can't predict crime, can't determine which job applicants are more suited than others, can't do effective facial recognition, can't grade essays or replace teachers— and yet technochauvinists keep selling us snake oil and pretending that technology is the solution to every social problem.

The biases embedded in technology are more than mere glitches; they're baked in from the beginning.

The next time you run into a person who insists unnecessarily on using technology to solve a complex social problem, please tell them about cookie division and the difference between social and mathematical fairness. You may also want to throw in some comments about how equality is the not the same as equity or justice, how the power dynamics between me and my much-younger brother are a microcosm of power dynamics more broadly, how power in tech rests with those who write and own the code, and how privilege functions vis-à-vis differential access to technology (and cookies).

RACIST SOAP DISPENSERS

Digital technology is wonderful and world-changing; it is also racist, sexist, and ableist. For many years, we have focused on the positives about technology, pretending that the problems are only glitches. Calling something a glitch means it's a temporary blip, something unexpected but inconsequential. A glitch can be fixed. The biases embedded in technology are more than mere glitches; they're baked in from the beginning. They are structural biases, and they can't be addressed with a quick code update.

It's time to address this issue head-on, unflinchingly, taking advantage of everything we know about culture and how the biases of the real world take shape inside our computational systems. Only then can we begin the slow, painstaking process of accountability and remediation.

One of the examples I rely on to explain how tech is biased is the case of the racist soap dispenser. It is a good example of why tech is not neutral, and why the intersection of race and technology can reveal hidden truths.

The racist soap dispenser first bubbled up into public consciousness in a 2017 viral video. In it, a dark-skinned man and a light-skinned man try to use an automatic soap dispenser in a men's bathroom.

The light-skinned man goes first: he waves his hand under the soap dispenser and soap comes out. The dark- skinned man, Chukwuemeka Afigbo, goes next: he waves his hand under the soap dispenser and... nothing happens.

The viewer might think: It could have been a fluke, right? Maybe the soap dispenser broke or ran out of power, exactly at that moment?

Afigbo gets a white paper towel, shows it to the camera, and waves it under the soap dispenser. Soap comes out! Then he waves his own hand under the sensor again, and again nothing comes out. The soap dispenser only "sees" light colors, not dark. The soap dispenser is racist.

To a viewer with light skin, this video is shocking. To a viewer with dark skin, this video is confirmation of the tech bias they have struggled with for years. Every kind of sensor technology, from facial recognition to automatic faucets, tends to work better on light skin than on dark skin.

The problem is far more than just a glitch in a single soap dispenser. This problem has its historical roots in film technology, the old-school technology that computer vision is built on.

Computers are excellent at doing math, yes, but time and time again, we've seen algorithmic systems fail at making social decisions.

SHIRLEY CARDS

Up until the 1970s, dark skin looked muddy on film because Kodak, the dominant manufacturer of film-developing machines and chemicals, used pictures called "Shirley cards" to tune the film-processing machines in photo labs. The Shirley cards featured a light-skinned white woman surrounded with bright primary colors. Kodak didn't tune the photo lab equipment for people with darker skin, because its institutional racism ran so deep.

The company began including a wider range of skin tones on Shirley cards in the 1970s. While this was the decade in which Black stars like Sidney Poitier rose to greater prominence, the change wasn't the result of activism or a corporate diversity push. Kodak made the change in response to its customers in the furniture industry.

The furniture manufacturers complained that their walnut and mahogany furniture looked muddy in catalog photographs. They didn't want to print color catalogs, switching from their previous black and white catalogs, unless the brown tones looked better. Kodak's sense of corporate responsibility manifested only once it stood to lose money from its corporate clients.

Most people don't know the history of race in technology, and sometimes they blame themselves when technology doesn't work as expected. I think the blame lies elsewhere.

I would argue that we need to look deeper to understand how white supremacy and other dangerous ideas about race, gender, and ability are embedded in today's technology. Once we acknowledge this, we can reorient our production systems in order to design technology that truly gets us moving in the direction of a better world.

This process starts with recognizing the role that unconscious bias plays in the technological world.

I don't believe that the people who made soap dispensers intentionally set out to make their soap dispensers racist.

Nor do I think that most people who make technology or software get up in the morning and say, "I think I'll build something to oppress people today."

I think that what happens is that technology is often built by a small, homogenous group of people. In the case of the soap dispensers, the developers likely were a group of people with light skin who tested it on themselves, found that it worked, and assumed that it would therefore work for everyone else.

They probably thought, like many engineers, that because they were using sensors and math and electricity, they were making something "neutral."

They were wrong.

We need to audit all of our technology to find out how it is racist, gender-biased, or ableist.

MOVING BEYOND TECHNOCHAUVINISM

Many computer scientists have come around to the idea of making tech "more ethical" or "fairer," which is progress. Unfortunately, it doesn't go far enough. We need to audit all of our technology to find out how it is racist, gender-biased, or ableist.

Auditing doesn't have to be complicated. Download a screen reader and point it at your Twitter feed, and it's easy to understand why social media might exclude a Blind person from participating fully in online conversations. If a city moves its public alert system to social media, the city is cutting off access to a whole group of citizens—not just those who are Blind, but also those who lack technology access because of economics or a host of other reasons.

We should not cede control of essential civic functions to these tech systems, nor should we claim they are "better" or "more innovative" until and unless those technical systems work for every person regardless of skin color, class, age, gender, and ability.

Because the technological world is the same as the social world, policies in the tech world need to change too. By "policies," I mean the formal and informal rules of how technology works. Tech companies need better internal policies around content moderation. Local, state, and federal laws need to be updated to keep us all safer in the technical realm.

So, how do we get there? We have options.

In the case of the racist soap dispenser, the engineering team would have caught the problem during the most minimal testing if they had had engineers with a wide variety of skin tones on their team. The available evidence suggests that they did not.

Everyone reading this will be better able to anticipate how race, gender, ability, and technology intersect to potentially cause oppression.

Google's 2019 annual diversity report shows that the global tech giant has only 3 percent Black employees. Only 2 percent of its new hires that year were Black women. Black, Latinx, and Native American employees leave Google at the highest rates, suggesting that the company's climate and its pipeline need serious work.

This isn't a secret: Google made headlines in 2020 for firing Timnit Gebru, the world's most prominent AI ethics researcher and one of its few Black female employees. Google also fired Meg Mitchell, Gebru's ethical AI co-lead. There are other high-profile Silicon Valley whistleblower cases like Ellen Pao or Susan Fowler that are proxies for hundreds of similar but not publicized incidents of sexism, racism, and the intersection thereof.

None of the major tech companies is doing better. Google is typical of the "big nine" tech companies, and in fact Google's inclusion, diversity, and technology ethics efforts are among the industry's best. However, the problem is that their "best" is still pretty bad. Facebook launched a high-profile effort to hire more women that in 2021 actually resulted in a slight decline in the number of female workers at the company.

Building nonsexist, antiracist technology requires a different mindset and requires wrestling with complex social issues while also solving thorny technical challenges. The good news: everyone reading this will be better able to anticipate how race, gender, ability, and technology intersect to potentially cause oppression.

Instead of technochauvinism, I'm going to offer a different solution: using the right tool for the task. Sometimes the right tool is a computer. Sometimes it's not. One is not better than the other. Just like the 50 percent split of a cookie is unrealistic, and the bigger half of the cookie is not necessarily better than the smaller half of the cookie. **The better choice depends on your motivation and your needs.** ⏻

Info

Meredith Broussard

Ready to dig deeper into the book?
Buy a copy of More than a Glitch.

Want copies for your organization or for an event?
We can help: customerservice@porchlightbooks.com
800-236-7323

ABOUT THE AUTHOR

Meredith Broussard is Associate Professor at the Arthur L. Carter Journalism Institute of New York University and Research Director at the NYU Alliance for Public Interest Technology. She is the author of Artificial Unintelligence: How Computers Misunderstand the World (MIT Press). Her work has been featured in the New Yorker, the New York Times, the Atlantic, BBC, Wired, the Economist, and more. She appears in the 2020 documentary Coded Bias and serves on the advisory board for the Center for Critical Race & Digital Studies.

SHARE THIS
Pass along a copy of this manifesto to others.

SUBSCRIBE
Sign up for e-news to learn when our latest manifestos are available.

Porchlight