9 Popular Doesn't Mean Good

How can you take a "good" selfie? In 2015, several prominent American media outlets covered the results of an experiment that purported to answer this question using data science. The results were predictable to anyone familiar with photography basics: make sure your picture is in focus, don't cut off the subject's forehead, and so forth. The experiment used the same type of procedures we used to analyze the *Titanic* data in chapter 7.

What was notable about the experiment—but was not noted by the investigator, Andrej Karpathy, then a Stanford PhD student and now the head of AI at Tesla—was that almost all the "good" photos were of young white women, despite the fact that older women, men, and people of color were included in the original pool of selfies. Karpathy used a measure of *popularity*—the number of "likes" each photo garnered on social media—as the metric for what constituted *good*. This type of mistake is quite common among computational researchers who do not critically reflect on the social values and human behaviors that lead to statistics being produced. Karpathy assumed that the photos were popular, and therefore they must be good. By selecting for popularity, the data scientist created a model that had significant bias: it prioritized young, white, cisgender images of women that fit a narrow, heteronormative definition of attractiveness. Let's say that you are an older black man, and you give your selfie to Karpathy's model to be rated. The model will not label your photo as good, no matter what. You are not white and you are not a cisgender woman and you are not young; therefore you do not satisfy the model's criteria for "good." The social implication for a reader is that unless you look a certain way, your picture cannot possibly be good. This is not true. Also, no kind or reasonable person would say this to another person!

Copyright © 2018. MIT Press. All rights reserved.

Broussard, Meredith. Artificial Unintelligence : How Computers Misunderstand the World, MIT Press, 2018. ProQuest Ebook Central, http://ebookcentral.proquest.com/lib/uic/detail.action?docID=5355856.
Created from uic on 2020-10-02 05:15:41.

This conflation of *popular* and *good* has implications for all computational decision making that involves subjective judgments of quality. Namely: a human can perceive a difference between the concepts *popular* and *good*. A human can identify things that are popular but not good (like ramen burgers or racism) or good but not popular (like income taxes or speed limits) and rank them in a socially appropriate manner. (Of course, there are also things like exercise and babies that are both popular and good.) A machine, however, can only identify things that are popular using criteria specified in an algorithm. The machine cannot autonomously identify the quality of the popular items.

This brings us back to the fundamental problem: algorithms are designed by people, and people embed their unconscious biases in algorithms. It's rarely intentional—but this doesn't mean we should let data scientists off the hook. It means we should be critical about and vigilant for the things we know can go wrong. If we assume discrimination is the default, then we can design systems that work toward notions of equality.

One of the core values of the Internet is the idea that things can be ranked. Today's society is mad for measurement; it's unclear to me whether the mania for measurement arose from the mathematical frenzy for ranking, or if the mathematical frenzy is simply a response to a social incentive. In either case, ranking is king right now. We have college rankings, sports team rankings, and hackathon team rankings. Students jockey for class rank position. Schools are ranked. Employees are ranked.

Everybody wants to be at the top; nobody wants to be in the bottom, and nobody wants to hire (or select) from the bottom. However, in education, which is the area that I know best, there is a logical fallacy at work. If we look at a pool of one thousand students and their test scores, usually the scores will fit a bell curve. Half the students will be above average, half will be below, and there will be a small percentage who score really well or who score terribly. That's normal—but school districts and state officials insist that their goal is to have all students at a level of "competence." This is impossible unless you set the bar for competence at zero. It's quite popular for school districts to claim they want all of their students to be high-achieving, but it isn't necessarily good to strive toward an impossible ideal.

The idea that popular is more important than good is baked into the very DNA of Internet search. Consider the origin of search: Back in the 1990s, two computer science graduate students wondered what to read next. Their

Copyright © 2018. MIT Press. All rights reserved.

Broussard, Meredith. Artificial Unintelligence : How Computers Misunderstand the World, MIT Press, 2018. ProQuest Ebook Central, http://ebookcentral.proquest.com/lib/uic/detail.action?docID=5355856.
Created from uic on 2020-10-02 05:15:41.

discipline was only fifty years old (as opposed to the hundreds of years of history of their brethren in mathematics). It was difficult to figure out what to read outside of the syllabi handed out in class.

They had read some math about analyzing citations to get a citation index, and they decided to try to apply this math to web pages. (There weren't very many web pages at this point.) Their problem was how to identify "good" web pages, meaning web pages that they thought would be worth reading. The idea was that it would be just like academic citations: In computer science, the most-often-cited papers are the most important. By definition, the good papers become the most popular. Therefore, they built a search engine that would calculate how many incoming links pointed to a given web page, then they ran an equation to generate a ranking called PageRank based on the number of incoming links and the ranking of the outgoing links on a page. They reasoned that web users would act just like academics: each web user would create web pages that linked to other pages that each user considered good. A popular page, one with a large number of incoming links, was ranked higher than a page with fewer incoming links. PageRank was named after one of the grad students, Larry Page. Page and his partner, Sergei Brin, went on to commercialize their algorithm and create Google, one of the most influential companies in the world.

For a long time, PageRank worked beautifully. The popular web pages were the good ones—in part because there was so little content on the web that *good* was not a very high threshold. However, more and more people went online, content swelled, and Google began to make money based on selling advertising on web pages. The search-ranking model was taken from academic publishing; the advertising model was taken from print publishing.

As people learned how to game the PageRank algorithm to elevate their position in search results, popularity became a kind of currency on the web. Google engineers had to add factors to search so that spammers wouldn't game the system. They added multiple features, iterating and tweaking the algorithm. One interesting feature is the geolocation dimension they added to autocomplete what the user types in the search box. Search autocomplete is based on what's happening around you. If you type "ga" into the search box, it will autocomplete with "GA" if lots of people near you are searching for Georgia topics (or possibly UGA football) or with "Lady Gaga" if lots of people around you are searching for the musician. Now, there are

Copyright © 2018. MIT Press. All rights reserved.

Broussard, Meredith. Artificial Unintelligence : How Computers Misunderstand the World, MIT Press, 2018. ProQuest Ebook Central, http://ebookcentral.proquest.com/lib/uic/detail.action?docID=5355856.
Created from uic on 2020-10-02 05:15:41.

over two hundred factors that go into search, and PageRank has been augmented by many additional methods, including machine learning. It works beautifully—except when it doesn't.

A good example of how tech doesn't quite translate comes from how page designers create the front page of a newspaper. It's highly curated. The different areas have names: *above the fold* and *below the fold* are the most obvious. In the *Wall Street Journal* (WSJ), there's always a bright spot on the front page, called the *A-hed*. It brings levity. Longtime WSJ staffer Barry Newman writes:

> The "A-hed" started out as just another headline code. It soon became the code name for a story light enough to "float off the page." The A-hed is a headline that doesn't scream. It giggles.
>
> Great editors, it's been said, create vessels into which writers can pour their work. That's what Barney Kilgore did starting in 1941. The modern *Journal*'s first managing editor, he knew that into the world of business a little mirth must be poured.
>
> By putting the fun out front, wrapped around the day's woes, Kilgore sent a larger message: That anyone serious enough about life to read the *Wall Street Journal* should also be wise enough to step back and consider life's absurdities …
>
> Done well, an A-hed is more than a news feature. Ideas rise out of our personalities, our curiosities and our passions. A-heds aren't humor columns. They don't push opinions. We don't make stuff up. Sometimes, a touch of poignancy can set all joking aside. Yet two reporters reporting on the same oddity will always report its oddness in their own odd ways.[1]

This is substantially different than a scroll like the Facebook news feed because a page editor will consider a mix: something light, something dark, and a few mid-range stories to balance. A front page is a precise mix of elements. The *New York Times* has a team that curates its digital front page manually all day, every day. Few other news organizations can afford this staff effort, so at smaller outlets a homepage might be curated once a day or automatically populated based on the print front page. The page editor's curation adds value to the reading experience. This is good, but not popular: news homepage traffic has been steadily declining since social media began to eat the world.

It's popular to blame journalists, and journalism, for the decline of public discourse. I would argue that such blame is misplaced and not good for society, however. The switch from print to digital has had a dramatic impact on the quality of journalism produced in the United States. The Bureau of Labor Statistics reports that among information industries in

Copyright © 2018. MIT Press. All rights reserved.

Broussard, Meredith. Artificial Unintelligence : How Computers Misunderstand the World, MIT Press, 2018. ProQuest Ebook Central, http://ebookcentral.proquest.com/lib/uic/detail.action?docID=5355856.
Created from uic on 2020-10-02 05:15:41.

2015, average annual pay in the Internet publishing and web search portals industry was $197,549. Average annual pay was only $48,403 for those in newspaper publishing and $56,332 for radio broadcasting.[2] As newsrooms empty out because talented writers and investigative journalists are seeking higher-paying jobs, there are fewer people left to keep the foxes out of the henhouse.

This is a problem because cheating is baked into the DNA of modern computer technology and modern tech culture. Around 2002, when Illinois redesigned the image that would be imprinted on its quarter-dollar coins as part of the nationwide quarter redesign, state officials decided to hold a contest so that citizens could vote for the design they liked best. A programmer friend of mine had a clear favorite: the Land of Lincoln design, which showed a handsome young Abraham Lincoln holding a book inside an outline of the state of Illinois. To Lincoln's left was a silhouette of the Chicago skyline. To his right was a silhouette of a farm showing a house, barn, and silo. To my friend, this was the only design that ought to represent her state to the rest of the country.

So, she decided to commit a tiny bit of fraud to tip the balance in favor of Honest Abe.

Illinois officials were holding the voting online, hoping that using this then-new method of citizen engagement would allow them to reach new constituencies. My friend looked at the voting page and realized she could write a simple computer program that would repeatedly vote for Land of Lincoln. It took her all of a few minutes to write the program. She set it to run again and again, stuffing the ballot box in favor of Land of Lincoln. The design won by a landslide. In 2003, the design was launched to the rest of the country.

When my friend first told me this story in 2002, I thought it was funny. I still think about her every time I look at the spare change in my pocket and see an Illinois quarter. At first, I agreed with her that throwing a state quarter election was a harmless prank—but in the ensuing years, I came to see it as sad for the officials. The Illinois officials thought they were getting unprecedented response from the public about a civic issue. What they were really getting was the idle whim of a twenty-something who was bored at work one day. To the Illinois officials, it *looked* exactly like a lot of citizens weighing in on a civic matter. It probably made them happy to imagine that thousands of citizens really, really cared about graphic design

Copyright © 2018. MIT Press. All rights reserved.

Broussard, Meredith. Artificial Unintelligence : How Computers Misunderstand the World, MIT Press, 2018. ProQuest Ebook Central, http://ebookcentral.proquest.com/lib/uic/detail.action?docID=5355856.
Created from uic on 2020-10-02 05:15:41.

on currency. Dozens of other decisions must have been made based on the votes—people's careers, promotions, financial decisions inside the US Treasury.

This is the kind of fraudulent activity that happens every hour of every day on the Internet. The Internet is a magnificent invention, but it has also unleashed an unprecedented amount of fraud and a network of lies that move so fast that the rule of law can hardly keep up. After the 2016 US presidential election, there was a flood of interest in fake news. Nobody in tech was surprised that fake news was out there. They were surprised that people took it seriously. "Since when did people start believing everything on the Internet is true?" one programmer friend asked me. He honestly didn't realize that there are people who don't understand how web pages are made and how they get onto the Internet. Because he didn't realize this, he didn't realize that some people consider reading something on the Internet to be the same as reading something from a legitimate news outlet. It's not the same, but the two *look* so similar nowadays that it's easy to confuse legitimate and illegitimate information if you're not paying careful attention.

Few of us pay careful attention.

This willful blindness on the part of some technology creators is why we need inclusive technology, and we also need investigative journalism to keep the algorithms and their makers accountable. The foxes have been guarding the henhouse since the beginning of the Internet era. In December 2016, the Association for Computing Machinery (ACM), the main professional association for computer scientists, announced it was updating its code of ethics—for the first time since 1992. Many ethical issues have arisen since 1992, but the profession wasn't ready to confront the role that computers played in social justice issues. Fortunately, the new ethics code does recommend that ACM members should address issues of discrimination embedded in computational systems—prompted in part by the efforts of data journalists and academics who have taken on algorithmic accountability.[3]

Consider the case of eighteen-year-old Brisha Borden. She and a friend were goofing around on a suburban street in Florida. They spotted an unlocked Huffy bicycle and a Razor scooter. Both were child-sized. They picked them up and tried to ride. A neighbor called the police. "Borden and her friend were arrested and charged with burglary and petty theft for the items, which were valued at a total of $80," wrote ProPublica's Julia Angwin

Copyright © 2018. MIT Press. All rights reserved.

Broussard, Meredith. Artificial Unintelligence : How Computers Misunderstand the World, MIT Press, 2018. ProQuest Ebook Central, http://ebookcentral.proquest.com/lib/uic/detail.action?docID=5355856.
Created from uic on 2020-10-02 05:15:41.

in her coverage of the event.[4] Angwin then compared Borden's crime to another eighty-dollar infraction: the time that Vernon Prater, 41, shoplifted $86.35 in tools from a Florida Home Depot store. "He had already been convicted of armed robbery and attempted armed robbery, for which he served five years in prison, in addition to another armed robbery charge. Borden had a record, too, but it was for misdemeanors committed when she was a juvenile," Angwin wrote.

Each of these people was given a future risk rating when they were arrested—a move familiar from a movie. Borden, who is black, was rated high risk. Prater, who is white, was rated low risk. The risk algorithm, COMPAS, attempted to measure which detainees are at risk of recidivism, or reoffending. Northpointe, the company that developed COMPAS, is one of many such companies that are trying to use quantitative methods to enhance policing. It's not malicious; most of the companies hire well-intentioned criminologists who believe they are operating within the bounds of data-driven, scientific thinking on criminal behavior. The COMPAS designers and the criminologists who adopted the instrument truly thought they were being fairer by adopting a mathematical formula to evaluate whether someone was likely to commit another crime. "Objective, standardized instruments, rather than subjective judgments alone, are the most effective methods for determining the programming needs that should be targeted for each offender," reads a 2009 COMPAS fact sheet from the California Department of Rehabilitation and Correction.[5]

The problem is, the math doesn't work. "Black defendants were still 77 percent more likely to be pegged as at higher risk of committing a future violent crime and 45 percent more likely to be predicted to commit a future crime of any kind," Angwin writes. ProPublica released the data it used to perform the analysis. This was good because it enhanced transparency; other people could download the data, work with it, and validate ProPublica's results—and they did. This story unleashed a firestorm inside the AI and machine-learning communities. An absolute barrage of debate ensued—of the polite academic form, meaning that people wrote a lot of white papers and posted them online. One very important one was by Jon Kleinberg, a computer-science professor at Cornell University; Cornell graduate student Manish Raghavan; and Harvard economics professor Sendhil Mullainathan. In it, they proved that mathematically, it's impossible for COMPAS to treat white and black defendants fairly. Angwin writes: "A risk

Copyright © 2018. MIT Press. All rights reserved.

Broussard, Meredith. Artificial Unintelligence : How Computers Misunderstand the World, MIT Press, 2018. ProQuest Ebook Central, http://ebookcentral.proquest.com/lib/uic/detail.action?docID=5355856.
Created from uic on 2020-10-02 05:15:41.

score, they found, could either be equally predictive or equally wrong for all races—but not both. The reason was the difference in the frequency with which blacks and whites were charged with new crimes. 'If you have two populations that have unequal base rates,' Kleinberg said, 'then you can't satisfy both definitions of fairness at the same time.'"[6]

In short: algorithms don't work fairly because people embed their unconscious biases into algorithms. Technochauvinism leads people to assume that mathematical formulas embedded in code are somehow better or more just for solving social problems—but that isn't the case.

The COMPAS score is based on a 137-point questionnaire administered to people at the time of arrest. The answers to the questions are fed into a linear equation of the type you solved in high school. Seven *criminogenic needs*, or risk factors, are identified. These include "educational-vocational-financial deficits and achievement skills," "antisocial and procriminal associates," and "familial-marital-dysfunctional relationship." All of these measures are outcomes of poverty. It's positively Kafkaesque.

The fact that nobody at Northpointe thought that the questionnaire or its results might be biased has to do with technochauvinists' unique worldview. The people who believe that math and computation are "more objective" or "fairer" tend to be the kind of people who think that inequality and structural racism can be erased with a keystroke. They imagine that the digital world is different and better than the real world and that by reducing decisions to calculations, we can make the world more rational. When development teams are small, like-minded, and not diverse, this kind of thinking can come to seem normal. However, it doesn't move us toward a more just and equitable world.

I'm not confident that tech's utopians and libertarians are going to make a better world through using more technology. A world in which some things are more convenient, sure—but I don't trust the vision of the future in which everything is digital. It's not just about bias. It's also about breakage. Digital technology works poorly and doesn't last very long.[7] Phone batteries run down and stop holding a charge over time. Laptops stop working when their hard drives fill up after years of use. Automatic faucets don't recognize the motion of my hands. Even the elevator in my apartment building, which should be a simple algorithm based on workhorse technology that was invented decades ago, is flaky. I live in a high-rise building with a single bank of multiple elevators. Every hallway is identical. In one

Copyright © 2018. MIT Press. All rights reserved.

Broussard, Meredith. Artificial Unintelligence : How Computers Misunderstand the World, MIT Press, 2018. ProQuest Ebook Central, http://ebookcentral.proquest.com/lib/uic/detail.action?docID=5355856.
Created from uic on 2020-10-02 05:15:41.

of the elevators, there's something wrong with the wiring or the chips so that every few weeks, you press the button for your floor and it takes you to the floor above or below. It's unpredictable. A number of times, I've gotten off the elevator and walked to what I thought was my apartment door, and discovered that my key didn't work. I was on the wrong floor. The same thing has happened to everyone else who lives in my building. We chat about it in the elevators.

An elevator is a sophisticated machine with some programming embedded in it. There is an algorithm that determines which elevator goes to which floor and which one goes express down to the lobby and which one stops along the way. There are degrees of sophistication in the algorithms, too: newer elevators have programs that optimize the route for the people who push buttons at any given time. In the *New York Times* building, you push your destination floor on a central keypad at the elevator bank, and you're directed to an elevator optimized to get you to your destination as fast as possible, based on other people who also want to get to similar destinations at the same time. However, an elevator has one job. That one job is supported by highly qualified inventors, structural engineers, mechanical engineers, salespeople, marketing people, distributors, repair people, and inspectors. If all these people working together for decades can't make the elevator in my building do its one job, I don't have faith that a similar group of highly skilled people in a different supply chain will be able to make a self-driving car that will do multiple jobs simultaneously without killing me—or killing my kid or killing other people's kids who are riding a school bus or the kids who are innocent bystanders waiting at a bus stop.

The little things like elevators or automatic faucets matter because they are indicators of the functioning of a larger system. Unless the little things work, it's naive to assume the bigger issues will magically work.

Programmers' unconscious biases have been manifest for years. In 2009, Gizmodo reported that HP face-tracking webcams were not recognizing dark-skinned faces. In 2010, Microsoft's Kinect gaming system struggled to recognize dark-skinned users in low-light conditions. When the Apple Watch was released, it did not include a period tracker—the most obvious point of self-quantification for all women. Melinda Gates of the Bill & Melinda Gates Foundation commented on the omission: "I'm not picking on Apple at all, but just to come out with a health app that doesn't track

Copyright © 2018. MIT Press. All rights reserved.

Broussard, Meredith. Artificial Unintelligence : How Computers Misunderstand the World, MIT Press, 2018. ProQuest Ebook Central, http://ebookcentral.proquest.com/lib/uic/detail.action?docID=5355856.
Created from uic on 2020-10-02 05:15:41.

menstruation? I don't know about you, but I've been having menstruation for half my life, so far. It's just such a blatant error, and it's just an example of all the things we can leave out for women." Gates also commented on the lack of women in AI research: "When you look in the labs at who's working on AI, you can find one woman here, and one woman there. You're not even finding three or four in labs together."[8] Outside of the lab, the gender balance in leadership positions at tech firms is better—but not good. According to 2015 diversity figures compiled by the *Wall Street Journal*, LinkedIn is the big tech firm with the greatest percentage of women in leadership roles, with a measly 30 percent. Amazon, Facebook, and Google lag with 24, 23, and 22 percent, respectively. In general, statistics on leadership positions tend to be increased by women who rise to the top in marketing and human resources. These two departments tend to be more gender-balanced than engineering roles, as do social media teams. However, at tech firms the real power is held by the developers and engineers, not by the marketers or HR folks.

It's also worth considering the consequences of sudden, vast wealth on the community of programmers. Drugs play a large role in Silicon Valley and thus in the larger tech culture. Drugs were a major part of the 1960s counterculture, from LSD and marijuana to mushrooms and peyote and speed. In tech, drugs never became unpopular, but for years nobody really cared if developers were stoned so long as the code shipped on time. Now, with the opiate crisis reaching dramatic heights, it raises the question of how much technologists are facilitating the popularity and distribution of the ADD drugs and LSD and mushrooms and marijuana and nootropics and ayahuasca and DIY performance-enhancing drugs that are as popular in Silicon Valley as elsewhere. "With a booming startup culture cranked up by fiercely competitive VPs and adrenaline-driven coders, and a tendency for stressed-out managers to look the other way, illicit drugs and black-market painkillers have become part of the landscape here in the world's frothy fountain of tech," wrote Heather Somerville and Patrick May in the San Jose *Mercury-News* in 2014.[9]

In 2014, California ranked second among all states for highest rate of illicit-drug dependence and abuse among eighteen- to twenty-five-year-olds. That same year, the Bay Area had 1.4 million prescriptions for hydrocodone, a painkiller that is often taken recreationally. If you take speed to stay up, painkillers are useful to sleep. "There's this workaholism in the

Copyright © 2018. MIT Press. All rights reserved.

Broussard, Meredith. Artificial Unintelligence : How Computers Misunderstand the World, MIT Press, 2018. ProQuest Ebook Central, http://ebookcentral.proquest.com/lib/uic/detail.action?docID=5355856.
Created from uic on 2020-10-02 05:15:41.

valley, where the ability to work on crash projects at tremendous rates of speed is almost a badge of honor," Steve Albrecht, a San Diego substance-abuse consultant, told the *Mercury News*. "These workers stay up for days and days, and many of them gradually get into meth and coke to keep going. Red Bull and coffee only gets them so far." San Francisco, Marin, and San Mateo Counties had 159 visits per one hundred thousand people to hospital emergency rooms for stimulant abuse. This is triple the national average, which is thirty visits per one hundred thousand people.

Drug use is represented equally across all racial groups, as Michelle Alexander writes in the *New Jim Crow*.[10] However, while poor communities and communities of color are surveilled aggressively to enforce compliance with drug laws, the technology elites who build the surveillance systems seem to be free from scrutiny. Silk Road, an eBay-like marketplace for drugs, flourished openly online from 2011 to 2013. After its founder, Ross Ulbricht, was sentenced to prison, others stepped in to fill the gap. Alex Hern wrote in the *Guardian* in 2014: "DarkMarket, a system aiming to create a decentralised alternative to online drugs marketplace Silk Road, has rebranded as 'OpenBazaar' to improve its image online. OpenBazaar exists as little more than a proof of concept: the plan was sketched out by a group of hackers in Toronto in mid-April, where they won the $20,000 first prize for their idea."[11]

Two years later, an entrepreneur named Brian Hoffman took the OpenBazaar code, commercialized it, and got a $3 million investment from venture capital firms Union Square Ventures and Andreesen Horowitz to run the marketplace using Bitcoin, an alternative digital currency. In this, we can see the libertarian paradise that Thiel and others imagine: a new space, beyond the reach of government. It seems that their plan is working. Lend Edu, a fintech firm, surveyed millennials about their use of Venmo, a payments app owned by PayPal. Thirty-three percent of respondents said they had used Venmo to buy marijuana, Aderall, cocaine, or other illegal narcotics.[12] A site called Vicemo.com boasts the tagline "See who's buying drugs, booze, and sex on Venmo." It shows a constant livestream of people who publicly post their Venmo transactions. There is very little subtlety. A typical transaction reads like this post: "Kaden paid Cody/for my grub and ganja." Other users post emojis of pills or hypodermic needles. Trees or leaves or the phrase "cutting grass" typically signify marijuana transactions. Some of these are jokes, and some are actual horticultural payments,

Copyright © 2018. MIT Press. All rights reserved.

Broussard, Meredith. Artificial Unintelligence : How Computers Misunderstand the World, MIT Press, 2018. ProQuest Ebook Central, http://ebookcentral.proquest.com/lib/uic/detail.action?docID=5355856.
Created from uic on 2020-10-02 05:15:41.

of course. Either way, it's a little shocking to see the volume of transactions as the country struggles with an opioid crisis.[13]

Illegal drugs are popular. They've always been popular. Most people would argue that illegal drug use is not good, at least not for society as a whole—so when tech is being used to facilitate and distribute them, tech is being used in a way that's counterproductive for cultural good. Yet this is the logical outcome when tech is produced according to libertarian values with a willful disregard for application. If any of those people buying or selling drugs were apprehended and had their stats run through the COMPAS system, it would perpetuate yet another act of blatant discrimination. Therefore, it isn't enough to ask, of any new technical innovation, if it's good. Instead, we need to ask: Good for whom? We must investigate the wider application and implications of our technical choices and be prepared for the fact that we might not like what we find.

Copyright © 2018. MIT Press. All rights reserved.

Broussard, Meredith. Artificial Unintelligence : How Computers Misunderstand the World, MIT Press, 2018. ProQuest Ebook Central, http://ebookcentral.proquest.com/lib/uic/detail.action?docID=5355856.
Created from uic on 2020-10-02 05:15:41.