

[Search](#)

[Dashboard](#)

[Login or Signup](#)

[3.5 Resisting Reductio](#)

Published on Jul 16, 2018

DOI 10.21428/96c8d426

SHOW DETAILS



CITE



SOCIAL



DOWNLOAD



CONTENTS



Design Justice, A.I., and Escape from the Matrix of Domination

Essay Competition Winner

by [Sasha Costanza-Chock](#)



last released
7 years ago



Part 1: #TravelingWhileTrans

Millimeter Wave Scanning, the Sociotechnical Reproduction of the Gender Binary, and the Importance of Embodied Knowledge to the design of Artificial Intelligence




Image: 'Anomalies' highlighted in millimeter wave scanner interface, by Dr. Cary Gabriel Costello [Costello, Cary Gabriel, 2016. "Traveling While Trans: The False Promise of Better Treatment," in Trans Advocate. http://transadvocate.com/the-tsa-a-binary-body-system-in-practice_n_15540.htm]

It's June of 2017, and I'm standing in the security line at the Detroit Metro airport. I'm on my way back to Boston from the Allied Media Conference, a "collaborative laboratory of media-based organizing" that's been held every year in Detroit for the past two decades.¹ As a nonbinary, transgender, femme presenting person, my experience of the AMC was deeply liberating. It's a conference that strives harder than any that I know of to be inclusive of all kinds of people, including Queer, Trans, Intersex, and Gender Non-Conforming (QTI/GNC) folks. Although it's far from perfect and

GING) TOIKS. Although it's far from perfect, and every year inevitably brings new challenges and difficult conversations about what it means to construct a truly inclusive space, it's a powerful experience; a kind of temporary autonomous zone.² Emerging from nearly a week immersed in this parallel world, I'm tired, but on a deep level, refreshed; my reservoir of belief in the possibility of creating a better future has been replenished.

Yet as I stand in the security line and draw closer to the millimeter wave scanning machine, my stress levels begin to rise. On one hand, I know that my white skin, U.S. citizenship, and institutional affiliation with MIT place me in a position of relative privilege. I will certainly be spared the most disruptive and harmful possible outcomes of security screening. For example, I don't have to worry that this process will lead to my being placed in a detention center or in deportation proceedings; I won't be hooded and whisked away to Guantanamo Bay or to one of the many other secret prisons that form part of the global infrastructure of the so-called "War on Terror;"³ most likely, I won't even miss my flight while detained for what security expert Bruce Schneier describes as "security theater."⁴

On the other hand, my heartbeat speeds up slightly as I near the end of the line, because I know that I'm almost certainly about to be

subject to an embarrassing, uncomfortable, and perhaps even humiliating search by a TSA officer, after my body is flagged as anomalous by the millimeter wave scanner. I know that this is almost certainly about to happen because of the particular sociotechnical configuration of gender normativity (cis-normativity) that has been built into the scanner, through the combination of user interface design, scanning technology, binary gendered body-shape data constructs, and risk detection algorithms, as well as the socialization, training, and experience of the TSA agents.⁵

The TSA agent motions me to step into the millimeter wave scanner. I raise my arms and place my hands in a triangle shape, palms facing forward, above my head. The scanner spins around my body, and then the agent signals for me to step forward out of the machine and wait with my feet on the pad just past the scanner exit. I glance to the left, where a screen displays an abstracted outline of a human body. As I expected, bright fluorescent yellow blocks on the diagram highlight my chest and groin areas. You see, when I entered the scanner, the TSA operator on the other side was prompted by the UI to select ‘Male’ or ‘Female.’ Since my gender presentation is nonbinary femme, usually the operator selects ‘female.’ However, the three dimensional contours of my body, at millimeter

resolution, differ from the statistical norm of ‘female bodies’ as understood by the dataset and risk algorithm designed by the manufacturer of the millimeter wave scanner (and its subcontractors), and as trained by a small army of clickworkers tasked with labelling and classification (as scholars Lilly Irani and Nick Dyer-Witheford, among others, remind us⁶). If the agent selects ‘male,’ my breasts are large enough, statistically speaking, in comparison to the normative ‘male’ body-shape construct in the database, to trigger an anomalous warning and a highlight around my chest area. If they select ‘female,’ my groin area deviates enough from the statistical ‘female’ norm to trigger the risk alert, and bright yellow pixels highlight my groin, as visible on the flat panel display. In other words, I can’t win. I’m sure to be marked as ‘risky,’ and that will trigger an escalation to the next level in the TSA security protocol.

This is, in fact, what happens: I’ve been flagged, the screen shows a fluorescent yellow highlight around my groin. Next, the agent asks me to step aside, and (as usual) asks for my consent to a physical body search. Typically at this point, once I am close enough to the agent, they become confused about my gender. This presents a problem, because the next step in the security protocol is for either a male or female TSA agent to conduct a body search by running

their hands across my arms and armpits, chest, hips and legs, and inner thighs. The agent is supposed to be male or female, depending on whether I am male or female. As a nonbinary trans femme, I present a problem not easily resolved by the algorithm of the security protocol. Sometimes, the agent will assume I prefer to be searched by a female agent; sometimes, male. Occasionally, they ask whether I prefer a search by a male or female agent. Unfortunately, ‘neither’ is an honest but not an acceptable response. Today, I’m particularly unlucky: a nearby male agent, observing the interaction, loudly states “I’ll do it!” and strides over to me. I say “Aren’t you going to ask me what I prefer?” He pauses, seems angry, and begins to move towards me again, but the female agent stops him. She asks me what I would prefer. Now I’m standing in public, surrounded by two TSA agents, with a line of curious travelers watching the whole interaction. Ultimately, the aggressive male agent backs off and the female agent searches me, making a face as if she’s as uncomfortable as I am, and I’m cleared to continue on to my gate.

The point of this story is to provide a small but concrete example from my own daily lived experience of how larger systems - norms, values, assumptions - are encoded in and reproduced through the design of sociotechnical

data-driven systems, or in Langdon Winner's famous words, how "artifacts have politics." In this case, cisnormativity (the assumption that all people are cisgender, or in other words, have a gender identity and presentation that are consistent with the sex they were assigned at birth) is enforced at multiple levels of a traveler's interaction with airport security systems. The database, models, and algorithms that assess deviance and risk are all binary and cisnormative. The male/female gender selector UI is binary and cisnormative. The assignment of a male or female TSA agent to perform the additional, more invasive search is cis- and binary gender normative as well. At each stage of this interaction, airport security technology, databases, algorithms, risk assessment, and practices are all designed based on the assumption that there are only two genders, and that gender presentation will conform with so-called 'biological sex.' Anyone whose body doesn't fall within an acceptable range of 'deviance' from a normative binary body type is flagged as 'risky' and subject to a heightened and disproportionate burden of the harms (both small and, potentially, large) of airport security systems and the violence of empire they instantiate. QTI/GNC people are thus disproportionately burdened by the design of millimeter wave scanning technology and the way that technology is used. The system is

biased against us. Those who are also People of Color (PoC), Muslims, immigrants, and/or People with Disabilities (PwD) are doubly, triply, or multiply-burdened⁷ by, and face the highest risk of harms from, this system. Most cisgender people are unaware of the fact that the millimeter wave scanners operate according to a binary and cisnormative gender construct; most trans people know, because it directly affects our lives.

I share this experience here because I feel it to be an appropriate opening to my response to Joi Ito's call to "resist reduction," a timely intervention in the conversation about the limits and possibilities of Artificial Intelligence (AI). That call resonates very deeply with me, since as a nonbinary trans feminine person, I walk through a world that has in many ways been designed to deny the possibility of my existence. From my standpoint, I worry that the current path of A.I. development will produce systems that erase those of us on the margins, whether intentionally or not, whether in a spectacular moment of Singularity or (far more likely) through the mundane and relentless repetition of reduction in a thousand daily interactions with A.I. systems that, increasingly, will touch every domain of our lives.

In this response, I'd like to do three things: first, I've drawn from my own lived experience as a

I've drawn from my own lived experience as a gender nonconforming, nonbinary trans feminine person to illustrate how sociotechnical data-dependent systems reproduce various aspects of the matrix of domination (more on that below). Specifically, I've told a personal story that illustrates the reproduction of the binary gender system, and also hopefully demonstrates the importance of the intersectional feminist concepts of standpoint, embodied and situated knowledge, and nonbinary thought to A.I. systems design⁸. This first point, in a nutshell: different people experience algorithmic decision support systems differently, and we must redesign these systems based on the lived experience of those they harm. Second, in the next section I hope to extend Joi's critique of capitalist profitability as the key driver of A.I. by describing the paradigm shift wrought in many fields by the Black feminist concepts of intersectionality and the matrix of domination. Third, I'll briefly trace the encouraging contours of a growing community of designers, technologists, computer scientists, community organizers, and others who are already engaged in research, theory, and practices that take these ideas into account in the design and development of sociotechnical systems.

Part 2: A.I., Intersectionality, and the Matrix of Domination

Ito asks us to “examine the values and the currencies of the fitness functions and consider whether they are suitable and appropriate for the systems in which we participate.”⁹ He is primarily concerned with the reduction of fitness in A.I. systems to efficiency and capitalist profitability. I share this concern, but I would also argue that we must resist the urge to reduce the cause of the planetary ecological crisis to capitalism ‘alone.’ Instead, we’ll need to pay close attention to intersectionality and the matrix of domination, concepts developed by legal scholar Kimberlé Crenshaw and sociologist Patricia Hill Collins (the 100th president of the American Sociological Association), respectively. These concepts help us understand how capitalism, white supremacy, and heteropatriarchy (class, race, and gender) are interlocking systems: they are experienced simultaneously, by individuals who exist at their intersections. This has crucial implications for the design of A.I. systems.

Intersectionality was first proposed by legal scholar Kimberlé Crenshaw in her 1989 article “Demarginalizing the Intersection of Race and Sex: A Black Feminist Critique of Antidiscrimination Doctrine, Feminist Theory

and Antiracist Politics.” In the article, Crenshaw describes how existing antidiscrimination law (Title VII of the Civil Rights Act) repeatedly failed to protect Black women workers. First, she discusses an instance where Black women workers at General Motors (GM) were told they had no legal grounds for a discrimination case against their employer, because antidiscrimination law only protected single-identity categories. The Court found that GM did not systematically discriminate against all women, because the company hired white women, and that there was insufficient evidence of discrimination against Black people in general. Thus, Black women, who did in reality experience systematic employment discrimination as Black women, were not protected by existing law and had no actionable legal claim. In a second case described by Crenshaw, the court rejected the discrimination claims of a Black woman against Hugh Helicopters, Inc, because “her attempt to specify her race was seen as being at odds with the standard allegation that the employer simply discriminated ‘against females.’”¹⁰ In other words, the court could not accept that Black women might be able to represent all women, including white women, as a class. In a third case, the court did award discrimination damages to Black women workers at a pharmaceutical company, as women, but refused

to award the damages to all Black workers, under the rationale that Black women could not adequately represent the claims of Black people as a category.

Crenshaw notes the role of statistical analysis in each of these cases: sometimes, the courts required Black women to include broader statistics for all women that countered their claims of discrimination; in other cases, the courts limited the admissible data to that dealing with Black women only. In those cases, the low total number of Black women employees typically made statistically valid claims impossible, whereas strong claims could have been made if the plaintiffs were allowed to include data for all women, for all Black people, or both. Later, in her 1991 Stanford Law Review article “Mapping the Margins: Intersectionality, Identity Politics, and Violence Against Women of Color,” Crenshaw powerfully articulates the ways that women of color often experience male violence as a product of intersecting racism and sexism, but are then marginalized from both feminist and antiracist discourse and practice, and denied access to specific legal remedies.¹¹

The concept of intersectionality provided the grounds for a long, slow paradigm shift that is still unfolding in the social sciences, legal scholarship, and in other domains of research and practice. This paradigm shift is also

and practice. This paradigm shift is also beginning to transform the domain of technology design. What Crenshaw calls ‘single-axis analysis,’ where race or gender are considered as independent constructs, has wide-reaching consequences for A.I.

Universalist design principles and practices erase certain groups of people, specifically those who are intersectionally disadvantaged or multiply-burdened under capitalism, white supremacy, heteropatriarchy, and settler colonialism. What is more, when technologists do consider inequality in technology design (and most professional design processes do not consider inequality at all), they nearly always employ a single-axis framework. Most design processes today are therefore structured in ways that make it impossible to see, engage with, account for, or attempt to remedy the unequal distribution of benefits and burdens that they reproduce. As Crenshaw notes, feminist or antiracist theory or policy that is not grounded in intersectional understanding of gender and race cannot adequately address the experiences of Black women, or other multiply-burdened people, when it comes to the formulation of policy demands. The same must be true when it comes to our ‘design demands’ for A.I. systems, including technical standards, training data, benchmarks, bias audits, and so on.

Intersectionality is thus an absolutely crucial concept for the development of A.I. Most pragmatically, single-axis (in other words, non-intersectional) algorithmic bias audits are insufficient to ensure algorithmic fairness. While there is rapidly growing interest in algorithmic bias audits, especially in the Fairness, Accountability, and Transparency in Machine Learning (FAT*) community, most are single-axis: they look for a biased distribution of error rates only according to a single variable, such as race or gender. This is an important advance, but it is essential that we develop a new norm of intersectional bias audits for machine learning systems.

For example, Media Labber Joy Buolamwini and her project the Algorithmic Justice League have produced a growing body of work that demonstrates the ways that machine learning is intersectionally biased. In the project “Gender Shades,” they show how computer vision trained on ‘pale male’ data sets performs best on images of White men, and worst on images of Black women.¹² In order to demonstrate this, Buolamwini first had to create a new benchmark dataset of images of faces, both male and female, with a range of skin tones. This work not only demonstrates that facial recognition systems are biased, it also provides a concrete example of the need to develop intersectional training datasets,

how to create intersectional benchmarks, and the importance of intersectional audits for all machine learning systems. The urgency of doing so is directly proportional to the impacts (or potential impacts) of algorithmic decision systems on people's life-chances.

The matrix of domination

Closely linked to intersectionality, but less widely used today, the matrix of domination is a term developed by Black feminist scholar Patricia Hill Collins to refer to race, class, and gender as interlocking systems of oppression. It is a conceptual model that helps us think about how power, oppression, resistance, privilege, penalties, benefits, and harms are systematically distributed. When she introduces the term, in her book *Black Feminist Thought*, Collins emphasizes race, class, and gender as the three systems that historically have been most important in structuring most Black women's lives. She notes that additional systems of oppression structure the matrix of domination for other kinds of people. The term, for her, describes a mode of analysis that includes any and all systems of oppression that mutually constitute each other and shape people's lives.

Collins also notes that:

“People experience and resist oppression on three levels: the level of personal biography; the group or community level of the cultural context created by race, class, and gender; and the systemic level of social institutions. Black feminist thought emphasizes all three levels as sites of domination and as potential sites of resistance.”

We need to explore the ways that A.I. relates to domination and resistance at each of these three levels (personal, community, and institutional). For example, at the personal level, we might explore how interface design affirms or denies a person’s identity through features such as, say, a binary gender dropdown during account profile creation. We might consider how design decisions play out in the impacts they have on different individual’s biographies or life-chances.

At the the community level, we might explore how A.I. systems design fosters certain kinds of communities while suppressing others, through the automated enforcement of community guidelines, rules, and speech norms, instantiated through content moderation algorithms and decision support systems. For example, we know that Facebook’s internal content moderation guidelines explicitly mention that Black children are not a protected

mention that black children are not a protected category, while white men are; this inspires very little confidence in Zuckerberg's congressional testimony that FB is confident that they can deal with hate speech and trolls through the use of A.I. content moderation systems. Nor is Facebook's position improved by the recent leak of content moderation guidelines that note that 'White supremacist' posts should be banned, but that 'White nationalist' posts are within free speech bounds.

At the institutional level, we might consider how institutions that support the development of A.I. systems reproduce and/or challenge the matrix of domination in their practices. Institutions include various aspects of the State, especially funding agencies like NSF and DoD; large companies (Google, Microsoft, Apple); venture capital firms, standards-setting bodies (ISO, W3C, NIST), laws (such as the Americans with Disabilities Act), and universities and educational institutions that train computer scientists, developers, and designers. Intersectional theory compels us to consider how these and other institutions are involved in the design of A.I. systems that will shape the distribution of benefits and harms across society. For example, the ability to immigrate to the United States is unequally distributed among

different groups of people through a combination of laws passed by the U.S. Congress, software decision systems, executive orders that influence enforcement priorities, and so on. Recently, the Department of Homeland Security (DHS) had an open bid process to develop an automated ‘good immigrant/bad immigrant’ prediction system that would draw from people’s public social media profiles. After extensive pushback from civil liberties and immigrant rights advocates, DHS announced that the system was beyond ‘present day capabilities’. However, they also announced that they would instead hire 180 positions for people tasked to manually monitor immigrant social media profiles from a list of about 100,000 people. In other words, within the broader immigration system, visa allocation has always been an algorithm, and it is one that has been designed according to the political priorities of power holders. It is an algorithm that has long privileged whiteness, hetero- and cis-normativity, wealth, and higher socioeconomic status.

Finally, Black feminist thought emphasizes the value of situated knowledge over universalist knowledge. In other words, particular insights about the nature of power, oppression, and resistance come from those who occupy a subjugated standpoint, and knowledge

developed from any particular standpoint is always partial knowledge.

We have described the nearly overwhelming challenges presented by deeply rooted and interlocking systems of oppression. What paths, then, might lead us out of the matrix of domination?

Part 3: Building a world where many worlds fit

Against ontological reduction, towards design for the pluriverse, or, decolonizing AI

Ito ends “resisting reduction” on a hopeful note, with a nod towards the many people, organizations, and networks that are already working towards what he calls “a culture of flourishing”¹³. He mentions high school students and MIT Media Lab students; the IEEE working group on the design of A.I. around human wellbeing; the work of Conservation International to support indigenous peoples; and Shinto priests at Ise Shrine. I also believe that, despite the seemingly overwhelming power of the matrix of domination, it is important to center the real world practices of resistance and the construction of alternatives. Accordingly, I’ll end by describing a few more of the exciting emerging organizations and networks that are

-----o-o-o-----
already working to incorporate intersectional analysis into the design of A.I. systems.

The idea of intentionally building liberatory values into technological systems is not new. For example, the Appropriate Technology movement advocated for local, sustainable approaches to technological development in the countries of the Global South, rather than wholesale adoption of technology developed to serve the needs and interests of those in the wealthiest countries. In the 1980s, Computer Professionals for Social Responsibility emerged during the cold war to advocate that computer scientists resist the incorporation of their work into the nuclear arms race. In the 1990s, the values in design approach, developed by scientists like Batya Friedman, came to the fore.¹⁴ The past year has seen a wave of book-length critiques of the reproduction of race, class, and gender inequality through machine learning, algorithmic decision support systems, and AI, such as Virginia Eubanks' Automating Inequality, Cathy O'Neal's Weapons of Math Destruction, and Safiyah Noble's Algorithms of Oppression.

There is a growing community of computer scientists focused specifically on challenging algorithmic bias. As we touched on above, beginning in 2014, the FAT* community emerged

as a key hub for this strand of work. FAT* has rapidly become the most prominent space for computer scientists to advance research about algorithmic bias: what it means, how to measure it, and how to reduce it. This is such important work, with the caveat noted in the previous section (the current norm of single-axis fairness audits should be replaced by a new norm of intersectional analysis). This will require the development of new, more inclusive training and benchmarking datasets, as we saw with the work of the Algorithmic Justice League.

We need to also consider approaches that are beyond inclusion and fairness, and that center autonomy and sovereignty. For example, how do A.I. systems reproduce colonial ontology and epistemology? What would AI look like if it were designed to support, extend, and amplify indigenous knowledge and/or practices? In this direction, there is a growing set of scholars interested in decolonizing technology, including A.I. For example, Lilly Irani has argued for the development of postcolonial computing;¹⁵ Ramesh Srinivasan has asked us to consider indigenous database ontologies in his book *Whose Global Village*; and anthropologist and development theorist Arturo Escobar has just released a sweeping new book titled *Designs for the Pluriverse: Radical Interdependence, Autonomy, and the Making of Worlds*. In it,

— — — — —

Escobar draws from decades of work with social movements led by indigenous and Afro-descended peoples in Latin America and the Caribbean to argue for autonomous design. He traces the ways that most design processes today are oriented towards the reproduction of the 'One World' ontology. This means that technology is used to extend capitalist patriarchal modernity, the aims of the market and/or the state, and to erase indigenous ways of being, knowing, and doing (ontologies, epistemologies, practices, and life-worlds). Escobar argues for a decolonized approach to design that focuses on collaborative and place-based practices, and that acknowledges the interdependence of all people, beings, and the earth. He insists on attention to what he calls the ontological dimension of design: all design reproduces certain ways of being, knowing, and doing. He's interested in the Zapatista concept of creating "a world where many worlds fit," rather than the 'one-world' project of neoliberal globalization.

Happily, research centers, think tanks, and initiatives that focus on questions of justice, fairness, bias, discrimination, and even decolonization of data, algorithmic decision support systems, and computing systems are now popping up like mushrooms all around the world. These include Data & Society, the A.I. Next Institute, and the Digital Justice Lab in Miami.

now institute, and the Digital Equity Lab in New York City; the new Data Justice Lab in Cardiff, and the Public Data Lab. Coding Rights, led by hacker, lawyer, and feminist Joana Varon, works across Latin America to make complex issues around data and human rights much more accessible for broader publics, engage in policy debates, and help produce consent culture for the digital environment. They do this through projects like Chupadatos ('the data sucker'). Others groups include Fair Algorithms,¹⁶ the Data Active group,¹⁷ the Center for Civic Media at MIT; the Digital Justice Lab, recently launched by Nasma Ahmed in Toronto; Building Consentful Tech, by the design studio And Also Too in Toronto; the Our Data Bodies project, by Seeta Ganghadaran and Virginia Eubanks, and the FemTechNet network.

There are a growing number of conferences and convenings dedicated to related themes; besides FAT*, the past year has seen the Data4BlackLives conference, the 2018 Data Justice Conference in Cardiff, and the A.I. and Inclusion conference in Rio de Janeiro, organized by the Berkman-Klein Center for Internet & Society, ITS Rio, and the Network of Centers, as well as the third Design Justice Track at the Allied Media Conference in Detroit.

To end, it is worth quoting at length from the Design Justice Network Principles,¹⁸ first

developed by a group of 30 designers, artists, technologists, and community organizers at the Allied Media Conference in 2015:

Design Justice Network Principles

This is a living document.

Design mediates so much of our realities and has tremendous impact on our lives, yet very few of us participate in design processes. In particular, the people who are most adversely affected by design decisions — about visual culture, new technologies, the planning of our communities, or the structure of our political and economic systems — tend to have the least influence on those decisions and how they are made.

Design justice rethinks design processes, centers people who are normally marginalized by design, and uses collaborative, creative practices to address the deepest challenges our communities face.

1. We use design to **sustain, heal, and empower** our communities, as well as to seek liberation from exploitative and oppressive systems.


2. We **center the voices of those who are directly impacted** by the outcomes of the design process.
3. We **prioritize design's impact on the community** over the intentions of the designer.
4. We view **change as emergent from an accountable, accessible, and collaborative process**, rather than as a point at the end of a process.
5. We see the role of the **designer as a facilitator rather than an expert**.
6. We believe that **everyone is an expert based on their own lived experience**, and that we all have unique and brilliant contributions to bring to a design process.
7. We **share design knowledge and tools** with our communities.
8. We work towards **sustainable, community-led and -controlled** outcomes.
9. We work towards **non-exploitative solutions** that reconnect us to the earth and to each other
10. Before seeking new design solutions, **we look for what is already working** at the community level. We honor and uplift


traditional, indigenous, and local
knowledge and practices.

The Design Justice principles resonate closely with Ito's suggestion for "participant design".¹⁹ As we continue to race headlong towards the development of A.I. systems, we would do well to follow them.

In 1994, the Zapatistas appropriated the then-nascent 'Net to circulate a clarion call for "One No, Many Yeses." Fundamentally, it was a call to resist reduction. It is time to heed their words in our approach to the design of A.I. We need to listen to the voices of Indigenous peoples, Black people, Queer and Trans* folks, women and femmes, people with disabilities, immigrants and refugees, and all of those who are historically and currently the most marginalized, targeted, erased, under the matrix of domination. This is essential if we want to make space for many worlds, many ways of being, knowing, and doing, in our visions of A.I. and of planetary systems transformation.

FOOTNOTES • 19 

LICENSE •  [Creative Commons Attribution 4.0 International License \(CC-BY 4.0\)](#)

COMMENTS • 2    

? [Login to discuss](#)

? **Johnny Walker:** [arabuloku-okey oyna-agario modded-metin2 ...](#)

Sasha Costanza-Chock: Fixed spelling of Bruce Schneier's na...