

Context, Research, Refusal: Perspectives on Abstract Problem-Solving

Towards Trustworthy ML: Rethinking Security and Privacy for ML

International Conference on Learning Representations

April 26, 2020

Seeta Peña Gangadharan

London School of Economics and Political Science

I want to talk about what it means to be on the other side of the so-called machine learning revolution, give a sense of these experiences, and reflect broadly on their significance in relation to technological governance. My main argument revolves around the idea of context sensitization and the importance of moving beyond the abstractedness of computer science research and computational thinking. And by the end of this talk, I hope you might be willing to consider the role of context sensitization in the work that you do that could not only stimulate rethinking around trustworthy ML systems but also different ways of doing scholarship in this area.

This appeal to context sensitization derives from two main experiences, the first of which is the primary and most important. That first inspiration comes from co-leading *Our Data Bodies*, a multi-sited participatory study of privacy, data collection, data-driven systems, and marginalized communities. The second experience derives from my involvement in co-organizing an independent program on critique and reflexivity within the ACM Fairness, Accountability, and Transparency conference this past January. What I have observed a vast gap exists between different communities of practice—or epistemic communities, such as computer science researchers and marginalized communities.

My talk has three parts. In the first, I reflect on context. In the second, I introduce my research on data collection and data-driven systems in the lives of marginalized people and what they have said or done to challenge or refuse technological systems. And in the third, I share some perspectives on abstraction in computer science research and computational thinking.

On Context (Part 1)

Since this is a workshop that showcases privacy and security-related research in machine learning systems, and I am not a computer scientist, let me situate my terminology. To be clear, my use of the term context sensitization is inspired by the work by ethicist Clifford Christians and communication theorist James Carey. In their review of what constitutes the logic and aims of qualitative research, they cite to the importance of contextualization in investigation and analysis. You must become a master of context. You must be able to speak to the complex specificities of a situation or phenomenon.

[Slide 5]

You must be able to step out of these complex specificities in order to understand their relation to or connection with large systems and broader histories.

Doing so does not imply some sort of analytic movement from the particular to the general. It is not about focusing on an individual brick and bricklaying until you have an entire building. Instead, they argue, you must start from a broad level and from the perspective of interconnection or intersection between things, actions, and phenomena. Contextualization means having a fluent command of the histories, meanings, and relationships.

[Slide 6]

What I am describing refers a distinctive disciplinary trait. While some strands of social science make their mark by emulating the natural sciences, the qualitative tradition that Christians and Cliffords valorizes recognizes humans are living actors in an unfolding drama of life. Because of this, the qualitative tradition seeks to avoid reducing humans or, worse yet, dehumanizing them and their social practices. Rather it seeks to build understanding of the different processes by which humans “produce and maintain different forms of life and society and meaning and value” (Christians and Carey, 1989, p. 346).

Context is imperative to qualitative research.

Of course, the drive to humanize or to contextualize (as opposed to abstract or reduce) is not the sole domain of the social sciences. Humanistic motives in research are available to all, not just the humanistic social sciences or humanities, and computer science has claimed that mantle in small and unique ways. One example that comes immediately to mind is Computer Professionals for Social Responsibility and its opposition to computer defense systems and to computer science and military research collaborations.¹ The second—and more germane to research—is Phil Rogaway’s recent work and which shines a light on the material impacts of cryptographic research.² For Rogaway, math is never just numbers or just an abstract process. Math is political. Mathematical models are deeply connected to how we create and organize our lives together on this planet. So we should be asking important questions about when math is humanistic, and when it is not. Mathematics does not just comprise numbers and abstract computational processes.

Rogaway’s dictum—math is political—is critical, because it illuminates the broader context in which computer science research is done. It looks at the larger forces that connect mathematical models to social systems, political institutions, and economic activity. And in this manner, I see Rogaway’s work as an exercise in context

¹ <https://conservancy.umn.edu/handle/11299/107336>

² <https://web.cs.ucdavis.edu/~rogaway/papers/moral.pdf>

sensitization. He is sensitized to the ways in which human subjects are living actors in an unfolding drama of life and the ways in which cryptographic research intersects that drama. If we were to extend his provocation to research on machine learning systems, it would not be looking at how to bring complexity or diversity of the social life into mathematical models, but rather drawing connections between mathematical models and the complex world in which their problem-solving will meet. This includes the money needed to invest in them, the infrastructure that sustains ML systems, the institutions that steward these systems, the parts of the planet that are used to power these systems, the workers who keep them operational, the political processes that are transformed by them, and finally the people who must interact with ML systems.

I will come back to these points and reflect on the importance of context sensitization in the third part of my talk.

But for now, I want to move on to introduce Our Data Bodies, work that I co-lead and work that I believe can illuminate the lived contexts for members of marginalized communities dealing with data-driven systems. I want to focus particularly on the idea of refusal—refusal to accept the terms and conditions of data-driven systems in the way that institutions and the people who work for them offer. Refusal is significant to the argument about context sensitization—because in some ways, it reflects the sentiment that some members of marginalized communities simply do not want technical solutions or technological forms of governance for societal problems or systemic hardships they face.

[Slide 9]

Our Data Bodies (Part 2)

That drive for contextualization is one of the facets of the work we do at Our Data Bodies. Here is some of the background.

We are a research and organizing coalition, driven by justice, equity, and human rights, and we aim for our work to build with organizations involved in fights for racial justice, LGBTQI liberation, feminism, immigrant rights, economic justice, and other struggles to help us understand and address the impact of data-based technologies on social justice work.

We work in Charlotte, Detroit, and Los Angeles, where we are gathering and sharing stories of everyday surveillance and oppression which take place in people's encounters with data-driven systems and the institutions that manage them. By data-driven systems, we refer broadly to statistical and automated systems. (Note: we are not explicitly or expressly addressing machine learning systems, but they are relevant to our research.)

We're a five-person team. And we're community organizers, first, and researchers, second. Mariella's is connected to the Stop LAPD Spying Coalition. Tawana's with the

Detroit Community Technology Project. Tamika worked with Center Community Transitions. Kim is a writer, organizer, and artist based in Cape Town. Virginia Eubanks, who published the book *Automating Inequality*, was an original co-principal investigator and left the project in 2018.

Between late 2016 up until end of last year, we completed 140 in-depth interviews with individuals in these three cities' most marginalized neighborhoods. We also completed three reflective focus groups and conducted more than a dozen participatory workshops that allowed us to effectively understand how to best talk about all things data without stoking fear about data-driven systems. We actively tried to avoid creating paranoia or putting words in people's mouths.

In the spirit of contextualization, it is important to note that these cities differ in size and have unique histories. We entered communities in each from a different starting context to tap into these histories. In Charlotte, the country's second-largest financial capital and a city with high rates of imprisonment and recidivism of African-American populations, we concentrated on reentry (e.g., citizens returning from prison) and employment. In Detroit, a city which reeled from municipal bankruptcy and the Great Recession of 2008, we entered our research from the angle of utility shutoffs, evictions, and foreclosures. And in Los Angeles, a city with the third largest metropolitan economy in the world and home to an unhoused population of more than 35,000 people, our starting context was on housing at the intersection of criminal justice.

These starting contexts have provided essential ways in which to situate what we heard from people about the pain that they experienced when providing their data (or realizing data were being collected about them), and interacting with data-driven systems. Neither data nor system is an abstraction that can be used to create a mathematical model. By-and-large, systemic hardship weaves through the stories we have gathered.

For example, people talked of being caught in a cycle of disadvantage in which data and/or system was one part. After leaving jail, you can only find a temporary residence. When you go to apply for a job, you are denied, after listing your criminal background or the address of the temporary shelter where you live. So you get stuck in a rut. Or, if you get lucky and are able to land work, your luck goes only so far. As Jill, a Charlottean, explained: "I pled guilty to worthless checks in 2003... That's almost 15 years ago, but it's still being held against me... basically, all of my jobs have been temporary positions or contract positions."

[Slide 10]

Refusal in Context

I could go on and further elaborate similar kinds of experiences, but I want to draw your attention to the different ways in which members of marginalized communities

are responding. This also speaks to the aims and logic of context sensitization in data-driven research.

From a substantive point of view, the research of *Our Data Bodies* speaks to the profound need for democratic reimagining. This process involves a reimagining of political possibilities beyond the quote, unquote solution space that currently seems absent in the main conversations I have witnessed about fairness, accountability, transparency, or trust in statistical and automated systems. Let me unpack this below.

At the same time that we heard about systematic hardship, we also heard about refusal. People we talked to refused to settle for the data-driven systems or processes of data collection that were handed them. These forms of refusal, I contend, affirm the self (often in relation to a larger group, community, or collective). They assert one's agency. I want to review three forms of refusal here, because they directly implicate how, when, or why we prioritize technical forms of governance—say engineering trustworthiness, fairness, accountability, or transparency in systems—over other forms of governance.

[Slide 11]

Rectification

The first is refusal as rectification.

Allow me to introduce the story of Mellow, an older Black woman who is unhoused and living in the Skid Row area of Los Angeles. We interviewed her after a long struggle to get housed through the Coordinated Entry System, system that functions as a sort of Match.com for the unhoused population of Los Angeles. Coordinated Entry is powered by a scoring system called the Vulnerability Index, which is generated through a survey that welfare administrators pose to prospective housing recipients.

A resident of Los Angeles for ten years, and unhoused for four, Mellow fought tooth and nail to be placed in housing and refused to be defined by her data. And she credits her doggedness as the reason for success.

Despite being an at-risk individual, who should have easily scored “high” on the Vulnerability Index, she had been repeatedly denied housing. She could not identify the reason, but she felt she was being denied housing because she was outspoken. She discovered fraudulent shelter managers who demanded residents pay for shelter when shelters were already receiving benefits from the city. She saw security guards sexually harassing residents. She heard shelter workers tell her to eat rotting food and sleep in bed-bug infested beds. She had witnessed the death of a friend, whose diabetic condition went overlooked and who died while showering. Each time she could reapply for housing through the Coordinated Entry System, she was met with denial and told “your name has red tape on it.”

But just as she was dogged in documenting and reporting abuse at shelters, she was dogged in documenting her own needs for housing and appealing. She refused to accept the terms and conditions of a data-driven system that the city tried to present to her. Eventually, she challenged the data used to categorize her, rectified her record, and was granted housing.

Obfuscation

Mellow's story of refusal is one of individual resistance. In this next story, the refusal similarly chronicles what individuals are doing everyday to deny the people and institutions who manage and implement data-driven systems the ability to control and manipulate them. But unlike Mellow who sought to accurately represent herself, Ken, also based in Los Angeles, actively worked to misrepresent himself. Misrepresentation is, for Ken, a way to assert himself, to assert his personhood in a context of discriminatory state action.

When Ken spoke to Our Data Bodies, he explained his long history of struggle. A Native American man, he felt like most people around him targeted him just as the U.S. government had targeted indigenous people in 19th century: always wanted for dead. He struggled with mental health, drugs, and alcohol, was kicked out of shelters, and had to live on the streets, where law enforcement routinely harassed and abused him. In one instance, while Ken sitting down in front of some out-of-order public toilets, the local police tried to arrest Ken for trespassing. Ken countered back with the claim that the cops were harassing him for no good reason. When the police asked him his name, he gave them a false one, prompting the cop to respond, "Well, it's not in the computer." After several exchanges, the police issued him a ticket without a surname. As soon as they left, he tore it up. Clearly, Ken was practicing refusal (against a database-dependent police form). Like Mellow, he refused to accept the terms and conditions that data-driven systems presented to him. But in this case, he obfuscated his identity in order to remain excluded. And he did so in order to protect himself against a carceral system that again basically wanted him for dead.

Abolition

So far, I have described two individual-level examples of refusal. They show individuals working against and within data-driven systems to get their dignity and their due.

There is also refusal at the level of organized populations or communities. In Detroit, Tawana Petty, whose home organization is Detroit Community Technology Project, has been leading efforts to claw back the city's attempts to rapidly roll-out cameras with facial recognition capabilities. Called Project Greenlight, the program has promised "real-time crime fighting" and, last summer, boasted of partnerships with 550 institutions, including schools, churches, health centers, and gas stations that would stream images back to police headquarters.

For organizers in Detroit, the term “trustworthy” or “fair” real-time facial recognition is an oxymoron. There can be no trust or fairness in data-driven system within a context of broad disinvestment in social welfare, predation, gentrification, or trauma of demonizing media narratives about Detroiters.

Before the global pandemic, Tawana had been working with allied community organizations, tech privacy groups, and policymakers throughout the city and showing up to police commission meetings. She’s marshaled the stories and insights from Our Data Bodies, stressing the ways in which community members articulate their needs for safety and belonging, more so than privacy or security. The police commission has moved on some of Detroit Community Technology Project’s recommendations, though held back on others.

Since the global pandemic, the situation has become more complicated. The Michigan state senator, Senator Isaac Robinson, who had proposed a five-year moratorium on the use of facial recognition technology, lost his life to Covid-19. And Covid-19 has hit Detroit—a majority Black city—in profoundly disproportionate and devastating ways. It is hard to say whether banning facial recognition is a real possibility on the horizon. But it is without a doubt that Detroit Community Technology Project, in coalition with other local and state organizers, set the bar for collectively challenging the inevitability of data-driven systems advocated by state and private actors.

Before I get too lost in the current situation, I want to stress my larger point: Whether it’s refusal as rectification, refusal as obfuscation, or refusal as abolition, the people with whom we spoke want a decent life. People wanted access to resources to live. They wanted to be recognized as a human being, not quantified. People wanted respect from the individuals and institutions that they come into contact with. They desired meaningful human relationships and interactions. People wanted to be able to love themselves for who they are.

The Problem with Abstraction (Part 3)

What are we to make of these stories and strategies of refusal?

I see strategies of refusal as part of a larger demand for liberation from cycles of disadvantage and the dismantling of interlocking systems of oppression. Forget about human-centered machines or humans-in-the-loop. Forget about bad actors attacking systems from the outside. Forget about how ML systems interact with nature. Between the various acts of refusal, and the visions that we heard articulated in Charlotte, Detroit, and Los Angeles, members of marginalized communities call for, in many cases, radical transformation in the institutions that run the world, in the distribution of power in society, and in the dominant narratives that demonize marginalized people. Charlotteans, Detroiters, Angelinos want changes in the ways

that individuals understand, value, and relate to oneself, to one another, and to the planet. And they want changes that lead to their lives being filled with possibilities.

As I said at the beginning of my talk, I mentioned that I have witnessed a tremendous gap between epistemic communities or different communities of practice—between computer science research communities and the communities that Our Data Bodies is from and works with. I have heard computer scientists often talk and present their research in relation to real-world problems, as if researchers and their research are separate from the real world. I have also heard bald claims of machine-learning systems are here to stay, so we might as well do something about them. And I have read or listened to papers that disappear people into mathematical equations.

In contrast to research endeavors and disciplines that seem to abstractly relate to people and their lives, the research of Our Data Bodies is far more visceral. The practices of refusal don't dwell on the internal mechanics or constraints of systems, but technological systems in conjunction with social, political, and economic systems, and technological governance in conjunction with other forms of governance.

At a basic level, the problem which needs solving is that marginalized people are demonized, deprived, and lack effective opportunities to participate in the shaping of their destinies. At a more granular level, the problem is that designing and implementing data-driven technologies—even when well-intended—constrain people and their fields of action and legitimate the institutions who punish and demean them.

What is the point of tweaking data-driven systems to be fairer or more trustworthy when they make institutions even colder, more calculating, and more punitive than they already are for marginalized people who use their services? What is the point of tweaking data-driven systems to be more private and secure when the companies that control their production and diffusion siphon resources away from the social support and public infrastructure they need to live a decent life? As Petty recently wrote in a blog reflecting on our current times, why is their increased (supposedly limited-purpose) surveillance when it does not prevent Black people from dying? When some households still cannot get water? When air quality is still abysmal? When people cannot get the reliable healthcare they need? What is the point?

This is a profoundly different vantage from which conversations about technological systems begin.

And frankly the so-called more pragmatic approaches found in computer science feels like a taste of the kinds of impassive and profoundly undemocratic forms of technological governance that have yet to impact all populations. The processes of abstraction. The inability to humanize. The eventuality that technological systems make choices for humans in ways that we cannot easily track or intervene.

A culture of abstraction *is* a real-world problem, though with some effort, maybe this real-world problem is a tractable one. As I mentioned in the first section of my talk, I think there are examples to follow and a sensibility that can grow. At a minimum, the path towards context sensitization can involve framing questions in terms that different communities of practice might pose.

[Slide 16]

So, instead of health versus privacy and security, maybe the frame could be expanded so that it also acknowledges resource distribution problems.

Rather than building more technology, how might collective wellbeing be improved if tech companies were taxed at higher rates and that money went into public coffers that help governments respond to public health needs? Or more simply put, what would happen if instead of making machine learning systems more accurate, trustworthy, or robust, we focused on making companies contribute their fair share to society?

That's just one example of widening the lens and scoping societal problems. And it clashes with disciplinary traditions. But if discomfort is the price we must pay for producing new research, this cost pales in comparison to that which would arise from letting abstraction prevail.

Now before concluding, I want to share my minimum—that is, what I and researchers and organizers like me ought to commit to. (After all, we are in this together, I am a guest in your house, and I am grateful for the invitation.) At a minimum, I ought to commit to and do better at translating between computer science and social science. I know I could do more to identify when the experiences documented with Our Data Bodies do align with the problems that computer scientists can and do address. Discussions about adversarialism that are described in Kulynych et al. (2020) would appear to align with the disobedient practices of refusal that I have described.³ I have the sense that there is more. I know I connect more with computer science research on archiving or evaluating user feedback and embedding this feedback in social systems. But the communities that Our Data Bodies comes from and works with need you to help us, so we can get a better sense of when and where we are in alignment and when where we are not.

Between friction and alignment, I think change can happen.

I look forward to imagining those possibilities with you.

³ <https://arxiv.org/abs/1806.02711>