Big Data at the Margins

Digital Care and Cruelty: Social Provisioning and Deprivation in the Era of Big Data
Panel/Workshop Report
28-9 January 2021

This panel and workshop explored the question: What good can big data, automation and ratification intelligence do for individuals in need of social assistance and what harms can it perpetuate?

Panelist Dr. Virginia Eubanks describes the consequences of automated decision-making in public service programs and the advent of the "digital welfare state".

Eubanks approaches this tendency towards automated austerity by recounting experiences of individuals affected by these systems:

1. A little girl with important medical needs had her health insurance and care taken away due to an automated system error after the state of Indiana changed the eligibility process for Medicaid. The system sought to replace the work of local, hands-on case workers by automating the processing of claims via an online application. Applicants often faced a "failure to cooperate" error, without knowing what the error was, resulting in thousands of denied claims and restricted access to health services, with little avenue for recourse. The burden of finding and fixing these mistakes fell solely on the shoulders of those requesting these services, further exacerbating the struggles of some of the state's most vulnerable people.
2. A family was red flagged for child abuse because, facing ongoing poverty, they could not afford antibiotics for their child following an emergency room visit. A system implemented in Allegheny County, Pennsylvania used a statistical model to detect or predict child neglect by using data collected from various public agencies and institutions. The model's designers claimed that the system was programmed to eliminate bias. However, the bias moved from the bias of case workers to the bias of the economists and data scientists who designed and built the system. The question emerges: Who gets to have discretion in these systems?
3. A chronically unhoused man in Los Angeles, California's Skid Row was seeking housing through the county's coordinated entry program but was unable to achieve the necessary score on the VI-SPDAT (Vulnerability Index – Service Prioritization Decision Assistance Tool) questionnaire that LA County used to place individuals in housing. The questionnaire categorizes individuals based on comparative vulnerability but fails to address the overall lack of resources needed to remediate the issue of homelessness in the area.

Eubanks notes that, in all these cases, we can see that the use of these tools limit political vision through triage, by assuming a lack of resources, and depending on an ethical cover via automated rationing.

Panelist Dr. Sasha Costanza-Chock discusses the advent of algorithmic necropolitics at the crux of the first wave of the COVID-19 pandemic.

Necropolitics denotes the social and political power that dictates how and whether people live or die, and under what conditions their bodies will continue to exist. Necropolitics became manifest in the development of Massachusetts' Crisis Standards of Care Planning Guidance document, an algorithmic decision-making tool to triage patients for urgent medical care.

The document proposed a patient priority score using an allocation framework based on "saving the most lives" and "saving the most life-years", assigning patients to treatment priority groups accordingly, and determining which groups will have access to critical care interventions. However, socio-economic factors were not considered in the decision-making process in the goal of keeping things "objective"; this approach is at odds with the principle of social determinants of health, wherein social factors actually have an overwhelming impact on individual health outcomes and population-level health. Thus, the Crisis Standards of Care is a necropolitical document because it lays out the algorithm that will allocate scarce resources to some bodies and deny them to others.

This approach has serious implications for racialized and disabled individuals because the points system is racist (it pretends to be colourblind despite everything we know about the social determinants of health that structure risk for co-morbidities) and ableist because many impairments give disabled individuals a higher point score on prognosis for long-term survival (thus putting them in a lower priority group). Costanza-Chock notes that we should demand "just and ethical crisis standards that take into account political and social determinants of health to assure truly equitable allocation of resource", a demand that highlights the core problem of algorithmic resource provisioning.

Panelist Dr. Joanna Redden addresses the question: what does care and cruelty look like in datafied societies when we observe the project of automating (and eventually cancelling) public services systems?

Redden uses case studies from around the world, drawing from cases of cancelled fraud detection, justice/policing, and child protection systems to assess key implications. Investigating these systems can reveal the value sets that are classifying some groups as more worthy than others. Who is perceived to be a risk or a threat? Who is most affected (or stands to be affected) when these systems are introduced?

Redden notes that there is a type of cruelty inflicted upon those who are subject to the total surveillance innate in automated public services: Everyone in need of and applying for public benefits is screened for potential wrongdoing, but the lack of transparency about the workings of these automated systems makes it hard to contest or seek recourse for adverse or erroneous decisions. This is pointedly illustrated by Australia's Robodebt scheme and ongoing errors with the country's automated debt assessment and recovery system.

There is also a certain cruelty involved in the larger mission of overturning automated systems even after they are deemed detrimental to the public interest. Systemically, remediating the harm caused by faulty systems can require lengthy legal challenges and resource-intensive class-action lawsuits. Individuals affected by these systems may, in the interim, face personal bankruptcy, stress, family breakdown, and ongoing fights for restitution for errors made. Critical media investigations and civil society protests (outside parties or independent oversight prompting the cancellation of these services) play a key role in influencing the cancellation of harmful public service systems and are thus an important source of care in datafied societies.

Key themes

A number of key themes and questions emerged in the discussion:

The way we discuss technology use and social justice tends to exist in the abstract

Discussions about AI or big data systems are often difficult for lay people to understand. Obscure, technical language, and the rhetoric of "objectivity" and "universality" can result in what's been called "math-washing", deterring people (researchers, policy-makers, or community members) from intervening in the process of development and implementation.

In addition, the dynamics of research and decision-making related to these systems are subject to contested and conflicting notions of "expertise". The designation of what constitutes expertise becomes a political matter. For example, there are things that nurses know better than doctors, even though the medical system naturally privileges the expertise of doctors. In the case of digital care and big data systems, for example, machine-learning developers in the U.S. measure harm to children by monitoring repeat calls to Child Protective Services (CPS), while case workers know that racism and vendetta-calling can also lead to multiple calls to CPS, even when the child is not in a harmful situation. The conflicting politics of expertise can be mitigated by recognizing different types of expertise and making sure all parties are equitably involved.

Other ways to tackle the problem of abstraction in discussions about AI and big data are to focus on developing and sustaining a common and accessible language, one that can be used by technical experts and everyday community members alike and situating them in narratives and anecdotes that can make their effects/consequences more tangible and discernable. It must be emphasized that these systems, often positioned as highly technical and impenetrable, have very real impacts on individuals and groups. The idea of a unified, universal, and objective system is false; these systems have a variety of impacts on a wide variety of people (see: VI-SPDAT's classification of crisis/chronic homelessness).

Clarifying definitions make a big difference to how AI and data systems are designed and used

We can't talk about digital care without also examining how policy-makers and data scientists define "health". The definition of health by these parties has a primarily bio-medical inflection, namely a focus on physiological health or physical wellbeing. Health care is equated with medical treatment. It may be more effective for system developers to define health through a "social determinant of health" lens, which considers health in relation to broader personal, social and economic factors for individuals, groups and populations (e.g., employment and working conditions, education, access to social supports, income, etc.)

If we want to develop equitable data systems (specifically, systems that can adequately consider the intersections of race, gender identity, sexual orientation, socio-economic standing, etc.) we also have to re-evaluate the meaning of "diversity". Especially in commercial/institutional contexts, the neoliberal approach to diversity assumes that if you just diversify the people in the room, equity and liberation will follow. There needs to be a shift in the way we think about diversity and representation so that they align more with Sasha Constanza-Chock's design justice framework that considers Patricia Collins' "matrix of domination".

The dynamic that exists between academic researchers and community partners should be respectful, nurtured and mutually productive (conducting research with, not about, communities)

Some community-based research enforces an arbitrary divide between academics and community members. This can lead to uneven power dynamics, divergent goals, and misguided research approaches. Academics should be reflexive about their own social location and relationships to the research and consider what type of communities they identify with and belong to. This reflexivity would help create more equitable and productive relationships between academics and their communities.

While some research may focus on digital systems and fragmented, far flung groups, it is often more impactful to conduct research within local/geographic communities. This makes research more visible, tangible, and promotes accountability. It also allows researchers to better track the dynamics and histories that exist in the relationship between partners.

Academics should be asking community partners: "What can my research do for you?" Researchers may begin their research with preconceived or already established research questions/approaches that may not properly address the needs of community participants and partners. For example, a researcher may choose to examine access to internet resources in lower-income communities but define access as having internet in the home. But this choice overlooks the fact that access is also achieved through public libraries and schools. This preconceived notion of access as being in the home doesn't adequately encompass the experience of groups and communities. Researchers must remain open and flexible and be ready to reframe their questions or change their approaches.

Researchers must be mindful of the classist tendencies involved in (not) properly resourcing or compensating participants. We pay/compensate academic partners for the various types of expertise they provide; why shouldn't we pay community partners or groups in the same manner? This practice would demonstrate a certain sense of parity and respect for the time and effort that we draw from non-academic and community-based participants.

Data scientists and system developers operate within distinct institutional, educational, and epistemological contexts

There are many structural/hierarchical dynamics to deal with when conducting research of this nature - from both academic computer science departments and commercial/corporate entities. Effective, long-term change is difficult and requires educating leaders and upper management. Researchers must also contend with the demand for "efficiencies" at both the level of system development and government implementation.

In the context of STEM education, specifically engineering education, programs of study will often bracket out real world complexities (including experience-based or humanist elements), so as to simplify the process of system development. Obviously, this becomes problematic when students are employed in positions that develop tools for social provisioning or care management. It is very important to integrate humanities perspectives and data ethics when training data scientists or engineers.

Data scientists and system designers work with a different set of epistemological assumptions than community participants and academics in the humanities and social sciences. This exacerbates the problem of the privileging of technical expertise over social justice issues when discussing AI and big data. It also intensifies the rhetoric of "problem-solving" and "objective truth" that dominates the field (to the detriment of subjective and experience-based insights).

Empathy is important to social justice initiatives but tends to be dismissed given the nature of technical systems, to a detrimental effect

Most people fear the dehumanizing elements of big data and AI, so much so, that individuals interviewed by researchers say they "would rather deal with a flawed human than a computer". They would rather have a case worker who understands the nuances of their condition than be assigned resources by an algorithm.

These systems become a form of "empathy override" that tries to circumvent or dismiss the messy feelings involved in pressing social justice issues. Designers love to solve problems with technology because there is a sense that the concepts being used are clean and clear and easy to distinguish. This suits governments as well, since they are often impatient and in a hurry to address issues. This can lead to arrogance and a tendency to dismiss the very real, human concerns of community members. But value judgements and emotions are internalized in system design itself, and inevitably end up showing up in practice. This is something researchers need to be aware of when they approach these systems and work with community parties.