Workshop Discussion Notes: Criminal Justice

Data & Civil Rights
October 30, 2014 – Washington, D.C.
http://www.datacivilrights.org/

This document was produced based on notes taken during the Criminal Justice workshop of the Data & Civil Rights conference. This document represents a general summary of the discussion that took place. Not all attendees were involved in every part of the conversation, nor does this document necessarily reflect the views and beliefs of individual attendees. All workshop participants received workshop materials prior to the event to spark discussion. The primer can be found at: http://www.datacivilrights.org/pubs/2014-1030/CriminalJustice.pdf

Overview

This workshop focused on improving efficiency, accountability, and transparency within criminal justice institutions while considering points of access for engagement with civil rights groups and communities. Participants sought to build a framework for moving forward in addressing both the embedded harms and the positive potential for more efficient and objective uses of data. Before turning to concrete steps, discussants attempted to lay the groundwork of key questions at stake. There was a general consensus on the need to ask what kinds of data are useful, as well as determining a set of consistently applicable standards for use and interpretation in identifying factors such as bias. If outside parties could have a say in common practices such as predictive policing or the use of risk assessment tools, then what kinds of methodologies for testing could be set in place? It also became clear that there was a range of limitations such as lack of uniformity and inadequate data management capabilities within law enforcement that would require collaboration across institutions to rectify.

Themes and Discussion Topics

Finding Points of Access

The discussion opened with a call for participants to identify some of the most pressing challenges or harms facing data collection and usage within the criminal justice system. What emerged were a range of intersecting issues that centered on finding points of access that would allow for transparency and collaboration with civil rights groups, without producing a chilling effect on the technology itself.

The workshop participants established first that there is a need to move beyond zero-sum understandings that pit public safety in conflict with privacy. In light of this, participants proffered ideas on what paths exist for encouraging transparency while protecting sensitive information. How can we ensure enough clarity on practices such as predictive policing, so that we can follow the decision-making process from prediction to action? If, for example, someone is placed on the Chicago "Heat List" because data analysis identifies them as a high criminal risk, then what actions or interventions follow as a result, and how can we test their efficacy? To this end, one

participant outlined the hype versus reality around predictive technologies, and what they can do (e.g. predictive policing tools cannot really predict future crime; risk assessment tools cannot actually tell you who will commit violent felonies upon being released). The need to ground civil rights conversations in a greater understanding of the comparative efficacy of the risk assessments being made with new tools, which, it was argued, are only marginally improved from traditional tools, was strongly emphasized.

One of the points of contention centered on the logistics and capacities of both police institutions and civil rights groups to consistently collect and interpret data. For example, one participant noted that police institutions lag far behind in terms of efficiency in amassing useful data on their own activities, often lacking basics such as computers, and relying on paper rather than electronic systems. On another level, the lack of uniformity in records has raised a challenge on how to effectively regulate cross-institutionally and between local and state levels. How can outside groups gain access to such data? One discussant contended that missing civil liberties infrastructure in law enforcement at the local and state level leaves a void of meaningful ways to engage, but it was countered that regional centers that incorporate privacy and civil liberties do exist, and could act as a launching pad for further involvement.

An ongoing challenge throughout the workshop was envisioning a viable role for grassroots institutions and communities. For this to be possible, it was agreed that data must be more widely available, while a certain level of technical standardization would need to be implemented too. Part of the obstacle is rooted in social attitudes towards big data and privacy, especially for communities concerned about surveillance and social control. If new narrative frames can be set in place, a space for dialogue with the criminal justice system can be opened. One example put forward was the analysis of available data on search and stop rates in North Carolina, which was successfully used by community organizers to identify disparities between police departments in arrest rates due to racial profiling, leading to a series of reforms. Likewise, this data could be used to dismiss unfair arrests in cases where a clear racial bias can be proven in an individual officer's arrest history. What this shows is that transparency is not simply about regulation, but can be used as a tool for collaboration between independent institutions and the police departments in mining databases, identifying biases, and suggesting better practices. The group also identified another positive outcome in the fact that civil rights groups will be able to make a stronger case for cultivating fairness when they are able to present concrete data points to support their claims.

While the data in such cases may be rich, participants raised concerns over the issue of capacity. One discussant made note of the fact that many civil rights groups, particularly smaller organizations, do not have the resources or analytic tools to begin delving into big data. Potentially, third parties and resources such as the Community Oriented Policing Services (COPS) of the Department of Justice could be brought in to help. Outside collaborations would be justified not only on the grounds of expertise, but also in promoting neutrality. If, for example, police departments themselves are collecting data on police misconduct, would this discourage people from submitting complaints, for fear of retaliation? One participant also raised the issue that law enforcement is already increasingly relying not only on in-house data but also information from public agencies and private companies: the question of how to regulate that trend was asked, but it the specifics of what is needed to do so were not addressed.

Another prominent thread in the discussion centered on the pitfalls of unintentionally producing a "chilling" effect that would deter law enforcement from taking advantage of new and

potentially useful technologies. This concern is evident in cases where the revelation of surveillance sources may cause people to move away from frequently used or visited platforms. Is there a way to reconcile the security and intelligence-gathering needs of law enforcement with the value of transparency?

What kinds of data do we want?

The latter half of the workshop session focused on posing answerable questions on how to develop better quantitative measures for assessing the fairness of practices such as predictive policing. The concept of efficacy was also addressed in more than one sense. The discussion encompassed not only data quality but also ways to define what counts as bad or unnecessary data, as well as what potential efficiency can be gained through solid implementation of risk assessment tools.

The benefits for civil rights issues were most evident in the discussion on detecting biases and approaching objectivity by looking at inputs to systems. Despite the myths of big data and the "implicit technophilic promise of accuracy and fairness", in practice, structural and individual racial biases will find their way into "neutral" systems. For instance, variables that are used in risk assessment algorithms, like neighborhood, family, education, etc. that are, either on their own or in combination with each other, effectively proxies for being poor or a person of color. Likewise, it was noted that raw data, as in the case of felon lists, often contain errors that produce a wave of consequences for individuals, such as obstructing their right to vote. Participants agreed that some form of independent evaluation of different risk assessment tools would be needed to encourage accountability and reduction of harms. One participant emphasized the need for more and better empirical studies on risk assessment and analysis of best practices.

If not all analytic tools are equally accurate or valuable, then how can we test and determine what works and what can be improved? The use of Automatic License Plate Recognition was presented as an example. After a certain point, there are diminishing returns to storing data after several months. Rather, holding onto data beyond the short term in this case becomes a liability for individuals' privacy rights and a drag on institutional resources. On this line of reasoning, there was some discussion about using risk analysis to implement time and energy reallocation or for identifying underperforming police officers, but one discussant countered that such measures may creep too close to becoming a job performance metric. Similarly, there were doubts aired over the efficiency benefits of improving risk assessment models. Does more complexity always lead to better results? It was noted that predictive efforts may not be sufficiently advanced for small percentage differences in probability risk to actually translate into tangible benefits. At the same time, others pointed out the clear incentives for moving towards collecting better data, including with regards to federal funding flows.

Finally, the discussion turned towards future-oriented concerns by asking, what kinds of data do we want? The group proposed a range of contexts in which data collection could lead to better practices, including data on traffic stops, sentencing, parole denials, use of force, and incarceration rates, among others. One consideration that the group kept in mind is that placing raw arrests, raw stops, and raw ticket issues as inputs measuring officers' performance specifically introduces negative incentives for officers to large numbers of arrests for petty misconduct, rather than to target activities that are better associated with proactively reducing future crime risk; as such, the group thought about the best ways to incentivize data collection, and about the kinds of data that

should be used as inputs in evaluating officer performances, such as quality performance measurements (like the number of arrests that lead to a prosecution). These law enforcement practices were coupled with a suggestion to gather data on more qualitative aspects such as use of narratives and success stories in instances that produced positive results, such as improved community safety.

The group concluded the session by brainstorming on desired practices in data collection. Participants suggested the need for more raw data, increased automation, and the decoupling of police practices from data recording, as in the case of license plates being photographed and run through Optical Character Recognition, without the need for self-reporting from officers. Ideas such as implementing implicit bias studies and conducting randomized control trials were additionally put forward. It was suggested that perhaps we can learn insights and methodologies from other domains where risk assessment and prediction are important factors, such as the field of medicine, to improve the comprehensiveness of data collection.

Areas for Further Exploration

One of the unifying themes to emerge from the workshop was how to make data useful in the face of broader limitations. How, for example, do we compensate for the disparity of resources and inconsistent modes of data collection across different states and police departments? Are there ways to educate on better and more objective methods to use such data, and who will track and analyze it? It became clear that the tasks at hand exceeded the capacities of any one institution. While, for example, a civil rights group may possess the knowledge of what problems face their communities and have an observational grasp on what law enforcement practices are discriminatory, they themselves may not have the resources and technological or statistical expertise to pull the necessary links from the data. On a social level, the goal of creating new narratives for speaking about law enforcement among civil rights groups was introduced, but there remains an uncertainty around what concrete steps can be taken to create new attitudes about the intersections of law enforcement, civil rights, and new technology, particularly when it comes to devising viable quantitative measures of what happens at each level of the criminal justice system. For example, if many of the potential gains of big data analysis are small due to lack of advanced tools or input, then how do we convince different institutions of the benefits, particularly when collection requires energy and resources? But as one discussant noted, are there in fact any alternatives but to turn to big data in this context?