Title:

Municipal Surveillance Regulation and Algorithmic Accountability

Authors:

Meg Young, Michael A. Katell, P. M. Krafft

Forthcoming in Big Data & Society

Abstract:

A wave of recent scholarship has warned about the potential for discriminatory harms of algorithmic systems, spurring an interest in algorithmic accountability and regulation. Meanwhile, parallel concerns about surveillance practices have already led to multiple successful regulatory efforts of surveillance technologies—many of which have algorithmic components. Here, we examine municipal surveillance regulation as offering lessons for algorithmic oversight. Taking the 2017 Seattle Surveillance Ordinance as our primary case study and surveying efforts across five other cities, we describe the features of existing surveillance regulation; including procedures for describing surveillance technologies in detail, processes for public engagement, and processes for establishing acceptable uses. Although these surveillance-focused laws were not intended to address algorithmic accountability, we find these considerations to be relevant to the law’s aim of surfacing disparate impacts of systems in use. We also find that in notable cases, government employees did not identify regulated algorithmic surveillance technologies as reliant on algorithmic or machine learning systems, highlighting a definitional gap that could hinder future efforts toward algorithmic regulation. We argue that (i) finer-grained distinctions between types of analytic and information systems in the language of law and policy, and (ii) risk assessment tools integrated into their implementation would both strengthen future regulatory efforts by rendering underlying algorithmic components more legible and accountable to political and community stakeholders.

1. Introduction

Discriminatory harms in algorithmic systems have been examined in technologies as varied as facial recognition (Buolamwini and Gebru 2018), machine translation (Caliskan et al. 2017), and search engines (Sweeney 2013; Noble 2018). In response to these risks, recent work by scholars, activists, and policy-makers seeks to make algorithmic systems more accountable. Such interventions can be technical (c.f. Doshi-Velez & Kim 2017; Guidotti et al. 2018) or regulatory (Danaher et al. 2017; Jones 2017; Ziewitz 2016; Zarsky 2016). Yet despite the groundswell of interest, persistent questions remain as to how such interventions might be most effective.

As efforts around algorithmic regulation are in their early stages, we look to related yet better-developed area of surveillance regulation, which share common aims such as making public sector technologies legible to oversight. In recent years, surveillance regulations have been implemented in multiple municipalities throughout the United States in response to longstanding concerns about police surveillance practices, as well as increasing public awareness about the use of data-intensive systems by government agencies. Surveillance ordinances have been passed in thirteen jurisdictions as of this writing.

¹ While these laws are generally not intended as algorithmic accountability regulation, and because many surveillance technologies contain algorithmic components, these pieces of legislation contain salient aspects that provide a basis for responding to longstanding challenges in regulating algorithmic systems.

We present an ethnographic case study of the 2017 Seattle Surveillance Ordinance and survey more recent legislative efforts from five other cities. We provide a detailed account of how municipal surveillance technologies are regulated in practice. Defining algorithmic systems as technologies that use a computerized procedure—including machine learning or other artificial intelligence techniques—to

1

<https://www.aclu.org/issues/privacy-technology/surveillance-technologies/community-control-over-police-surveillance>

make or support decisions, judgments, or assessments, we also examine such systems listed as subject to the ordinance for their technical and social risk dimensions. To conclude, we provide recommendations for how the existing surveillance regulation could better describe and qualify functionality inside algorithmic systems and suggest that future work should include the development of interpretive tools for making algorithmic systems legible to community stakeholders and policy makers.

2. Background

Several attributes of algorithmic systems present unique challenges for regulatory oversight efforts. Marketing claims about the capabilities of various systems do not always match their actual utility (Thomas et al. 2018), and intellectual property protections and “trade secret” laws are used to block efforts to understand their functionality (Levine 2007; Pasquale 2015). Algorithmic systems adopted by government agencies are not just technical artifacts; they are situated in particular social contexts (Ananny 2016; Mittelstadt et al. 2016); their embeddedness in larger infrastructure systems makes them only situationally or relationally visible to their users (Star and Ruhleder 1996). Moreover, algorithmic systems are more emergent and polyvalent than previous media technologies (Ananny 2016) and limited lay understandings of algorithmic systems may inhibit effective policy making as a result (Danaher et al. 2017). Case studies of existing regulation of algorithmic systems affirm these underlying challenges. For example, digitized stock market trades have been monitored for regulatory compliance in the U.S. since the 1990s, but such oversight is inadequate in the face of ever-increasing complexity, such as in the case of high speed algorithmic trading (Snider 2014). Given the gap between the marketing pitches and discourses promoting such systems and the realities of use, Christin (2017) calls for these to be “decoupled” via ethnographic work on algorithmic systems’ use in practice to better understand their effects on the ground.

Interest in algorithmic regulation is yet more urgent in the context of considering discriminatory harms in surveillance and policing. People of color are disproportionately arrested and imprisoned (Warde 2013) leading them to be overrepresented in government databases (Madden et al. 2017; Eubanks 2018), and in turn more likely to become the subjects of surveillance. The addition of data analytics techniques transforms acquired data into raw material for predictions about human behavior and supports automated decisions and interventions. Mai (2016) describes these practices as “datafication,” a dimension of surveillance in which machine learning and other algorithmic techniques surface patterns from prior events that correlate to outcomes demonstrated by others, yielding new knowledge about a person, place, or group. Growing evidence indicates that decision-making based on datafied surveillance reproduces and amplifies the structural inequalities reflected in the underlying datasets employed (Barocas and Selbst 2016; Crawford 2016), which could have multiple adverse impacts for data subjects long term (Katell 2018).

In this way, algorithmic systems introduce particular risks over previous practices. Brayne (2017) distinguishes “big data policing” from prior surveillance practice by its use of quantified risk assessment applied systematically and categorically, as opposed to selectively. Brayne documents how big data shifts policing towards predictive analytics; alert-based (rather than query-based) information systems; lower thresholds for including information as compared with previous practices; and data integration across disparate sources. These risks are particularly evident in predictive policing systems, which are susceptible to a type of risk known as “runaway feedback loops” (Ensign et al. 2017). For instance, at least one popular system calculates a “risk score” based on a range of factors including the number of contacts the subject has with police; subjects assigned higher scores are more likely to be contacted by law enforcement and each contact further increases their score (Winston 2018). Similar problems are documented across government systems, notably in software used to assess a criminal defendant’s

recidivism risk (Angwin et al. 2016; Chouldechova, 2017). Although algorithmic harms were not in and of themselves the impetus for surveillance oversight in the campaign to which we turn, the convergence of surveillance practices and algorithmic harm informed our analysis of recent regulatory efforts.

3. Methods

In order to study the strengths, weaknesses, and lessons learned from municipal government efforts to regulate surveillance technology, we conduct an in-depth case study on the 2017 Seattle Surveillance Ordinance and survey laws from five other cities. The Seattle Surveillance Ordinance is recognized as landmark legislation because of its early adoption and the strength of its detailed reporting processes, public engagement mechanisms, and direct political oversight functions (ACLU Washington 2017; Green 2019). We engaged in an ethnographic approach in order to understand the social and political factors that led to the design of the current regulation, to describe its content and implementation, and to characterize the perspectives of key stakeholders on its achievements and opportunities for improvement. Our primary aim was to describe the Seattle Surveillance Ordinance on its own terms: what motivated policy-makers in Seattle to pursue this effort? What rationales led to these specific definitions and procedures? And, how are these definitions and procedures being enacted in practice? These “emic” (i.e., grounded in participant views and epistemologies) perspectives provide the necessary foundation to understand the ordinance on its own terms, at which point we go on to apply an etic (i.e., external to the participants’ own perspective) reading of the law as informing algorithmic accountability.

Our case study consists of 14 semi-structured and open-ended interviews (Weiss 1995), participant observation (Atkinson & Hammersley 1998) of two public meetings, and document analysis (Bowen 2009) of artifacts related to the law and its implementation. Eleven of our interviews were with Seattle policymakers, city employees, local activists, and advocates involved in the development and

implementation of the Seattle Surveillance Ordinance, and two interviews were with vendor personnel for technologies purchased by the city. Interviewees were selected for their centrality to the development of the ordinance, or for their familiarity with technologies subject to City Council oversight. For insights into these ongoing regulatory efforts in areas where they were still emergent, we also conducted a phone interview with an academic focused on U.S. municipal technology policy.

Our field observations consist of one public engagement meeting in 2018 and one Surveillance Community Working Group meeting in 2019. The public engagement event we attended sought public input on Automated License Plate Reader (ALPR) technology and occurred in a public library in South Seattle in October, 2018. The working group meeting was the first meeting of the group and focused on administrative issues. Our participant observation included field notes and reviewing coverage in news stories, community blogs, and feedback delivered via public listservs. Documents were chosen based on their ability to provide insight into the motivation for the law and its implementation; these included the text of the law, publicly posted information about its implementation such as surveillance impact reports, reporting templates for the ordinance and related city IT governance processes, emails from public listservs, and news articles. This analysis also included the 2018 and successive versions of the “Master List” of technologies determined to be surveillance, as disclosed by municipal departments for review by the City Council.

To better characterize the functionality of the technologies on the Master List, we conducted an iterative, inductive qualitative coding procedure of the disclosed technologies. The coding scheme shifted over these iterations from a detailed typology to a simplification oriented towards exposing dimensions of the technologies relevant to analyzing their algorithmic functions. To validate the finding that at least three technologies in current use were likely to be using machine learning or artificial intelligence, we also

sought input from experts in machine learning and artificial intelligence. The results of our review of the 2018 Master List informed our interview data collection by prompting us to ask personnel what systems under review used algorithms and machine learning; as well as our data analysis by attuning us to how analytic capabilities of existing systems are surfaced in the language of reporting templates.

In addition to our in-depth case study of Seattle, we also reviewed the policy language of five other surveillance ordinance efforts from Oakland, California; Berkeley, California; Davis, California; Nashville, Tennessee; and Cambridge, Massachusetts. These efforts varied in their strengths and maturity. Our review was aimed at providing a basis for us to assess what elements in Seattle's ordinance are common to laws in other cities. Our data collection for the comparative study consisted of documents on (i) the ordinances themselves, (ii) local and national press coverage of each ordinance, and (iii) model language advanced in a national campaign by the American Civil Liberties Union. We conducted web searches using the keyword "surveillance" joined with "ordinance," "regulation," "municipal," and "local." We also added specific jurisdiction names to these searches (e.g. cities, transit authorities) based on information about the status of surveillance regulation efforts reported by the ACLU (ACLU n.d.).

4. Case Study

Our case study traces the motivation and implementation of the current Seattle Surveillance Ordinance between 2017-2019. Seattle has long been home to an active and engaged public on issues related to government technology use. In 2012, local activists were catalyzed to act by news reports that the Seattle Police Department had acquired surveillance technologies without clear policy guidance or public disclosure; including a drone aircraft procured with federal funding, a CCTV network, and a mesh network capable of tracking WiFi-enabled devices (Crump 2016). Around the same time in 2013, the Black Lives Matter movement elevated systemic racism and police brutality to a wider audience (Freelon

et al. 2016). Edward Snowden's 2013 revelations about the extent of National Security Agency spying practices also raised national awareness of the technological frontiers of surveillance generally (Madden 2014). Advocacy and activism associated with these developments ultimately led multiple cities in the U.S. to enact laws subjecting surveillance systems to oversight.

Responding to these developments, in 2013 the Seattle City Council passed an ordinance mandating that city departments provide the Council with detailed descriptions of how new surveillance equipment would be used---requiring Council approval prior to procurement. However, successive developments led some local observers to feel that the 2013 effort had fallen short of providing meaningful transparency and accountability. For instance, five months after these reports were to have been submitted to the Council, a local privacy and police accountability activist requested their disclosure and learned that no technologies had been submitted for Council review to that point (Mocek 2014). Then, in 2016, a local newspaper reported that the Seattle Police Department (SPD) was using software called Geofeedia to scrape, collate, and store social media posts associated with geolocations (Herz 2016). As the existing 2013 ordinance only subjected new *surveillance equipment* (i.e. hardware) to public approval, SPD did not seek Council review for their use of Geofeedia as a *software* application. The news report was met with a public outcry. Councilmember Lorena González, a former civil rights lawyer and Chair of the Public Safety Committee, championed the effort to revise the existing ordinance to include a new focus on the racial and social justice dimensions of surveillance. Describing her concern with new data collection, González (2018) explained, "Once we collect data, using any type of technology, that becomes part of our public records. And then it's susceptible to being requested by ICE [Immigration and Customs Enforcement], by DHS [Department of Homeland Security], or other federal agencies who do not have—at this point in time—the best interests of communities of color in mind" (González 2018)

In recognition of the salience of such concerns to the policy making effort, the language of the 2017 ordinance highlights the need to assess the “disparate impact” of surveillance on marginalized groups (City of Seattle Surveillance Ordinance 125376 2017). Specifically, the 2017 ordinance (amended in 2018) contains a series of declarations that signals these risks as directly inspiring the law:

“WHEREAS, Seattle residents can significantly benefit from carefully considered deployments of surveillance technologies that support the City’s responsibility to provide public safety and other services to the public, but such technologies also create risks to civil liberties related to privacy, freedom of speech or association, or disparate impact on groups through over surveillance, and....

“WHEREAS, protocols proposed by City departments for the use of surveillance technologies should include specific steps to mitigate civil liberties concerns and the risks of information sharing with entities such as the federal government, and should incorporate racial equity principles into such protocols to ensure that surveillance technologies do not perpetuate institutionalized racism or race-based disparities...” (Seattle Surveillance Ordinance 118930 2017).

As the law was being constructed to apply to a broader array of technologies, some expressed concern that a broad mandate for increasing public awareness about surveillance technologies might negatively impact the department’s ability to protect public safety. As one department employee told us, “With the surveillance ordinance there still is a concern that it basically is going to shut down SPD's ability to do any sort of surveillance on known criminals....we will fight crime with one hand tied behind our backs.” (Police department employee 2018b). Councilmember González recalled that similar concerns from the police department were the primary source of tension in the policymaking process, particularly the degree to which inter-departmental intelligence gathering is seen as essential to cross-jurisdictional efforts to fight major crimes, such as international terrorism, large-scale drug sales, sex trafficking, and human trafficking. SPD was concerned that expanding the mandate of the surveillance ordinance beyond equipment could include inter-agency surveillance and data sharing practices which are seen to be fundamental to public safety efforts.

Ultimately this tension resulted in a compromise over the scope of how surveillance technologies would be defined, and thus which activities would be subject to disclosure and approval. Whereas local advocates hoped the law would extend beyond City of Seattle-procured technologies to include data flowing via SPD's inter-jurisdictional intelligence sharing relationships (González 2018; Narayan 2018), SPD stated this approach would be infeasible. One respondent said: "The local chapter of the ACLU wanted to take [the scope of the ordinance] further, they wanted to address every piece of data that might qualify as surveillance data. That would be overwhelming considering we have literally millions of records... Maybe someday, but we are not there yet" (Police department employee 2018c). The final language of the amendment achieved a compromise with SPD to apply to hardware and software "designed, intended to be used, or primarily used" for the purpose of surveillance—but not datasets collected by hardware and software owned by other agencies.

A final key debate in the drafting of the ordinance was with respect to the form of public engagement, given the law's stated intent to surface the race and social justice impacts of each technology. Community advocates and the bill's sponsor desired the direct involvement of historically targeted communities in the vetting process. Within the city, some expressed reservations about whether laypersons could potentially veto a surveillance technology after considerable resources had already been dedicated to its procurement and reporting requirements, preferring community engagement to occur at the beginning of the process. Despite these reservations about late-stage community input, the law was updated in 2018 to include a Community Surveillance Working Group, with members appointed by the Mayor and City Council, to evaluate draft reports on each technology and prepare privacy and civil liberties impact assessments (Seattle Municipal Code § 14.18.080).

The ordinance includes procedures for disclosing the city’s surveillance technologies, reporting and departmental commitments to particular uses of each technology, and community input. The process consists of six phases (see Figure 1): (i) the creation of a standing “Master List” of technologies in use or in procurement by the city to be filed with the City Clerk and made publicly available; (ii) the compilation of a detailed Surveillance Impact Report (SIR) for each technology and its use; (iii) a public comment period and public engagement events specific to each technology; (iv) review of the SIR and public comments by the Community Surveillance Working Group which adds a privacy and civil liberties impact assessment to each SIR; (v) City Council approval; and (vi) an annual equity impact assessment of acquired technologies. Each SIR must include descriptive details about the functionality of the technology, its planned or reasonably foreseen uses, the types of identifiable information collected, a detailed data management plan, potential civil rights and disparate impacts, potential benefits to the public agency adopting the technology, and a description of public engagement efforts along with any feedback received. Surveillance oversight and reporting joins existing city-wide processes for data governance in city-procured technology, including privacy and cybersecurity compliance processes. The context of this integration affords additional oversight mechanisms; for instance, when system capabilities or uses advance beyond those approved by City Council, city personnel determine whether those new capabilities are sufficient to require a renewed SIR and political approval process. The city also has a default clause in its vendor agreements called the “Right of First Refusal,” under which it can decline vendor-initiated changes to existing systems.

Figure 1. Reporting, review, and approval process in the Seattle Surveillance Ordinance

Surveillance Technology Determination Criteria	
Criteria	Detail
Does the technology meet the definition a Surveillance Technology?	Technology whose primary purpose is to observe or analyze the movements, behavior, or actions of identifiable individuals in a manner that is reasonably likely to raise concerns about civil liberties, freedom of speech or association, racial equity or social justice. Identifiable individuals also include individuals whose identity can be revealed by license plate data when combined with any other record.
Do any of the following exclusion criteria apply?	Technology that is used to collect data where an individual knowingly and voluntarily provides the data.
	Technology that is used to collect data where an individual knowingly and voluntarily provides the data.
	Technologies used for everyday office use.
	Body-worn cameras.
	Cameras installed in or on a police vehicle.
	Cameras installed pursuant to state law authorization in or on any vehicle or along a public right-of- way solely to record traffic violations.
	Technology that monitors only City employees in the performance of their City functions.

Do any of the inclusion criteria apply?	The technology disparately impacts disadvantaged groups.
	There is a high likelihood that personally identifiable information will be shared with non-City entities that will use the data for a purpose other than providing the City with a contractually agreed-upon service.
	The technology collects data that is personally identifiable even if obscured, de-identified, or anonymized after collection.
	The technology collects data that is personally identifiable even if obscured, de-identified, or anonymized after collection.

Table 1. Surveillance Technology Determination Criteria

We attended a public comment event in October 2018 that presented ALPR technology in parking enforcement and patrol uses. The event featured presentations, semi-structured discussion, and both verbal and written opportunities to provide public comment. The room setup featured whiteboards, large adhesive posters with discussion questions posted on the wall, and chairs clustered around each whiteboard. The first 30 minutes featured presentations from city employees including police officers. An SPD employee responsible for police technology privacy and transparency gave a short presentation on how the department uses ALPR in patrol vehicles and parking enforcement, explaining how scans of license plates in the field are checked against a “hot list” of plate numbers associated with persons involved with violent crimes, missing children and senior citizens, or stolen cars. This presentation was followed by planned remarks from SPD leadership and field officers as to the value and use of ALPR in the field. Throughout the presentations, members of the public asked clarifying questions. After the presentation, employees from the city led breakout sessions around respondents’ hopes, concerns, and suggestions for the city, including ideas they have for how else department aims could be accomplished in the absence of ALPR.

In posts on public email listservs after the event, members in a local privacy activist group expressed their skepticism that the benefits of ALPR use outweighed its risks. One person wrote, “In the context of SPD's

use of ALPR, [a presenter] showed a slide with a map which had stolen vehicle recoveries plotted on it. When someone in the audience asked her if those were incidents in which ALPR was used to recover the vehicles, her response ...strongly implied that while [the map] implied such, they were, in fact, not” (Mocek, 2018, listserv post). Here, the speaker expresses concern about a map presented by SPD to indicate stolen vehicle recoveries, while SPD responses to audience questioning suggested that the vehicle recoveries were not attributable to the use of ALPR. This gap between how ALPR is portrayed and its use in practice raised the activist’s critique that technologies were being framed so as to incline public approval.

This episode also illustrates the multiple purposes that the law pursues to increase transparency, accountability, and the public trust. Notably, the most salient aspects of the law vary with respect to different vantage points; to civil rights advocates, the ordinance is intended to address the disparate impacts of surveillance; to employees working on data privacy, the law is a further expression of the city’s commitment to data governance; to the police department, the law can build trust within the community---without which the police cannot work effectively.

In practice, while thorough documentation and publication of reports is a primary tool for accountability and transparency in the surveillance ordinance, reporting processes are time-intensive. At the time of this writing in early 2019, SIRs for five technologies have been published and led through the public comment process, and are still under review by the Community Surveillance Working Group. A first cluster of technologies, grouped by theme, will be presented to the City Council for approval beginning in Fall 2019. Respondents highlighted the page length of completed SIRs as a challenge to public engagement processes. For example, the report on SPD patrol car use of Automated License Plate Readers is 349 pages long prior to additional input from the Community Working Group; a camera used by the Fire

Department to remediate hazardous materials spills (“hazmat camera”) is 118 pages; another technology used by the Department of Transportation to collect anonymized hardware ID pings from mobile devices to aggregate data on point-to-point travel time (“Acyclica”) is 38 pages.

Qualitative Coding of Algorithmic Systems

As another dimension of our analysis, we conduct a review of the 2018 Master List of 28 surveillance technologies that were disclosed by municipal departments in accordance with Seattle's ordinance. Technologies were included on the Master List if they met the ordinance's determination criteria for surveillance technology (see Table 1), namely, “any electronic device, software program, or hosted software solution that is designed or primarily intended to be used for the purpose of surveillance.” Surveillance is defined as “means to observe or analyze the movements, behavior, or actions of identifiable individuals in a manner that is reasonably likely to raise concerns about civil liberties, freedom of speech or association, racial equity or social justice. Identifiable individuals also include individuals whose identity can be revealed by license plate data when combined with any other record. Under the determination criteria, a technology is not surveillance if an individual “knowingly and voluntarily consented to provide the information, or had a clear and conspicuous opportunity to opt out of providing the information” (Seattle Municipal Code § 14.18.010). Our review of the Master List allows us to ground our analysis of the surveillance ordinance in the technologies to which it is being applied.

To be systematic and to facilitate comparison in our review of these technologies, we inductively coded the features of the disclosed technologies to create a typology. Our aim was for the typology to capture (i) functionality, (ii) characteristics, and (iii) potential for misuse. The dimensions of our final typology indicated which municipal department is using the technology, and other attributes of the technologies described in Table 2. Table 3 shows the codes of the technologies on the Master List. Our coding process

was based on information about the technologies available from public documents, public hearings, and our interviews, but should be interpreted as our best informed assessments given that we did not have access to audit the physical technologies or any software specifications. Grey shaded rows indicate technologies that likely employ algorithmic or artificial intelligence (AI) components. Empty cells indicate information the authors could not obtain, judge, or infer with a reasonable degree of confidence. Column descriptions follow.

SPD	Seattle Police Department
SCL	Seattle City Light
SDT	Seattle Department of Transportation
SFD	Seattle Fire Department
CD	Collects and stores digital data
AI	Use a computerized procedure, including machine learning or other artificial intelligence (AI) techniques, to make or support decisions, judgments, or assessments
ID	Involves image data
VD	Involves video data
AD	Involves audio data
RD	Involves relational, i.e. network, data
CS	Potentially collected collateral data, i.e. non-targeted data as a result of targeted data
MS	Is a mobile system
TS	Is a tracking system
CM	Engages in continuous monitoring
P	Has the pervasive deployment that causes panoptic harms

Table 2. Departmental abbreviations and typology for qualitative codes used in Table 3.

Dept	Tech	CD	AI	ID	VD	AD	RD	CS	MS	TS	CM	P
SPD	Automated License Plate Recognition (ALPR)	1	1	1	0	0	0		1	0	1	1
SPD	Booking Photo Comparison Software (BPCS)	0	1	1	0	0	0	0	0	0	0	0
SPD	Forward Looking Infrared Real-time video (FLIR)		0	0	1	0	0	1	1	0	0	0

SPD	Undercover/ Technologies - Audio recording devices	1	0	0	0	1	0	1	1	0	0	0
SPD	Undercover/ Technologies - Camera systems	1	0	0	0	1	0	1	1	0	0	0
SPD	Undercover/ Technologies - Tracking devices	1	0	0	0	0	0	0	1	1	0	0
SPD	Computer-Aided Dispatch (CAD)			0	0	0				0		0
SPD	CopLogic	1		1	0	0	0	1	0	0	0	1
SPD	Hostage Negotiation Throw Phone	1	0	0	0	1	0	1	1	1	0	0
SPD	Remotely Operated Vehicles (ROVs)	0						0	1	0	0	0
SPD	911 Logging Recorder	1		0	0	1		0	0	0	0	0
SPD	Computer, cellphone and mobile device extraction tools	1	0					1	0	0	0	0
SPD	Video Recording Systems	1	0	0	1	1	0	0	0	0	0	0
SPD	Washington State Patrol (WSP) Aircraft		0				0	1	1	0	0	0
SPD	Washington State Patrol (WSP) Drones						0	1	1			1
SPD	Callyo	1	0	0	0	1	0	1	0	0	0	0
SPD	I2 iBase		1						0	0		
SPD	Parking Enforcement Systems											
SPD	Situational Awareness Cameras Without Recording	0	0	0	1	0	0	1	1	0	0	0
SPD	Crash Data Retrieval	1	0	1	0	0	0	0	0	0	1	0
SPD	Maltego	1	1	1	1	1	1	1	0	1	1	1
SCL	Binoculars/Spotting Scope	0	0	0	0	0	0	1	0	0	0	0
SCL	SensorLink Amp Fork	1	0	0	0	0	0	0	0	0	0	0
SCL	Check Meter Device	1	0	0	0	0	0	0	0	0	0	0
SDT	License Plate Readers	1		0	1	0	0	1	0	1	1	1
SDT	Closed Circuit Television Equipment	0	0	0	1	0	0	1	0	0	1	1
SDT	Acyclica	1		0	0	0	0	1	0	1	1	0
SFD	Emergency Scene Cameras	1	0	1	0	0	0	1	1	0	0	0
SFD	Hazmat Camera	1	0	1	0	0	0	1	1	0	0	0
SFD	Computer-Aided Dispatch			0	0	0				0		0

Table 3. Metadata, names, and qualitative codes of the 28 surveillance technologies disclosed in the 2018 Master List of the Seattle Surveillance Ordinance.

In part due to years of active privacy and security advocacy in the community, Seattle does not use several data-intensive technologies (such as predictive policing tools) already in use in other major municipalities. Many of the listed technologies are traditional surveillance hardware, such as cameras like

closed circuit television (CCTV) and emergency scene cameras. Alongside these technologies we noted some systems whose software presents the potential for algorithmic harms by engaging in datafication or utilizing algorithmic functions. At least three disclosed technologies rely in some capacity on machine learning or artificial intelligence (AI)—two forms of automated license plate recognition, and a booking photo comparison software. One software tool used for data analysis in police intelligence work, i2 iBase, can be enhanced with plugins that use inferential data analysis, but these are not currently included in the software in use by the department.

At the conclusion of this coding midway through our interview data collection, we decided to ask city employees about what algorithmic systems are in use in the city. Of our six total interviews with city employees whose roles relate to the ordinance, we asked three what algorithmic systems were in use at the city without defining algorithmic systems in asking this question. Two employees said that the city was not using algorithmic systems; another said that while Booking Photo Comparison Software used face recognition, its use may eventually be discontinued due to civil liberties concerns. This same respondent considered that algorithmic bias regulation would have to be taken up by an entirely separate regulatory effort, and that bias was not in scope of the intended goals for the ordinance. Other city employees framed the value of the law as an exercise in data governance, and focused on the data collection dimension of surveillance devices---rather than their analytic or datafication properties. In our discussion, we reflect on the lessons of these findings for future regulatory efforts.

5. Discussion

5.1 Unaccounted Algorithmic Harms

Those closely involved with the Surveillance Ordinance considered algorithmic bias considerations to be out of scope, requiring separate legislation. However, we find that at least two systems listed on the Master List merit additional assessment as to their potential for algorithmic harm. For example, ALPR employs an algorithmic processes (optical character recognition, or OCR) to recognize the characters on license plates in view and then compares this data to a “hot list” of cars associated with theft, missing persons, criminal investigations, and substantial unpaid parking tickets. When a scanned plate is a match for the hot list, an alert sound and notification alarm on the device interface in the police patrol vehicle. Here, a key algorithmic risk is “false positive” identifications resulting in unwarranted police engagements. The computer vision technology of OCR that underlies ALPRs is considered to be generally reliable in certain contexts. However, ALPRs still yield many false positives. A field officer we interviewed indicated that the rate of false positives is so high, “hits” or positive identifications are often ignored, at least by experienced field officers.

False positives may arise not from errors in the OCR alone, but from limits of system functionality. According to police employees, ALPRs in use in Seattle do not identify the state on a license plate, so a common form of false positive is for a hit matching the plate number alphanumerically, but for a different state. Although Seattle Police Department policy requires officers to verify the identified plate with a dispatcher, one officer interviewed in our field work indicated that mistakes are common, especially when new officers are learning to use the system. The same officer noted that these mistakes have led to unwarranted police engagements with drivers in Seattle. Considering that traffic stops have the potential to escalate (especially when someone is suspected of a crime that contributes to a plate’s inclusion on the “hot list”), this failure mode warrants serious attention. Notably, Seattle is one of many U.S. cities whose racially disparate application of policing has led to federal investigations and court-mandated supervision (Moy 2019). Another technology included in the 2018 Master List, Booking Photo Comparison Software

(BPCS), compares security camera stills with King County Jail booking photos and returns matches. As with OCR, face recognition is susceptible to false positives, which could misidentify innocent individuals. Compounding this problem, current face recognition software has been shown to have higher false positive rates for darker skin tones (Buolamwini and Gebru 2018).

Data-driven technologies pose risks for marginalized groups not only when their outputs are inaccurate, but also when their prompts reflect an accurate synthesis of data that is used in a racially inequitable social context. Scholars have recognized the disproportionality of arrests and convictions for people of color for decades (cf. Fellner & Human Rights Watch 2009). In algorithmic systems, these distortions have the potential to create feedback loops, as previously mentioned. Historic patterns of racialized policing produce data that amplify and distort a link between race and unlawful behavior: “For example, if police officers have been more likely to stop black drivers than white drivers, police data may encode a statistically significant link between race and traffic violations” (Moy 2019, p.15). ALPR and BPCS are both potentially capable of feedback loops of this nature. For instance, ALPR technology is used to identify drivers with unpaid parking tickets; to the extent that people experiencing poverty are less likely to be able to pay tickets on time, a positive ALPR match for unpaid citations could lead to additional fines. Similarly, BPCS utilizes a database of booking photos; disproportionate levels of policing leads to disproportionate representation of people of color in police databases. This leads to greater potential for software to identify---or misidentify---people of color. These considerations argue for the need to consider such risks in the existing evaluation procedures in pursuit of surfacing the disparate impacts of existing surveillance technologies.

5.2 Mental Models of Artificial Intelligence and Algorithmic Systems

Several conversations in our field work indicated that broader definitions of algorithmic systems and artificial intelligence could empower regulatory approaches to algorithmic accountability and transparency for technologies in current use. When we asked a city employee who works on police technologies, “Do the technologies that the Police Department use have an algorithmic or machine-learning component?”, the employee replied “Well ALPR is not, doesn't have that algorithmic intell-- ...I mean it's just.. we go in and we look at reads. What it can do is it can match up optical characters against a list of numbers, or license plates. But it doesn't, it's not like an iterative, AI type of process that-- you know-- ‘it learns from its mistakes’ and it ‘becomes increasingly invasive--’ there's nothing like that.” Others mentioned the difficulties of understanding the underlying technologies or asserted machine learning or AI were not being used within the city. Respondents did not provide detailed information as to their own definitions of AI and algorithmic systems, and our team did not provide definitions of algorithmic systems as part of our questioning. However, these responses point toward a sense among some city employees that algorithmic regulation is more appropriate for technologies on the horizon, rather than the quotidian systems in current use.

Through our discussions with city employees and data analysis, we came to understand this sense as driven by two factors. First, an expectation that tasks which are easy for people---such as vision---should also be easy for us to program into computers; we call this the *Minsky Fallacy*. An apocryphal tale has it that Marvin Minsky, a legendary forebearer of the field of AI, provided a student mentee a summer project to build a computer vision system that could perform image segmentation and recognition. The story is well known because of how causally Minsky treated problems that have since fostered decades of work dedicated to solving them. One respondent’s description was evocative of Minsky in claiming, “What it can do is it can match up--you know--optical characters against a list of numbers, or license

plates...” Here, the response neglects the underlying difficulty of this computational task and overlooks the presence of algorithmic processing.

We also identify a second related definitional disconnect at play; a mental image of artificial intelligence as involving complex forms of agency. In this conception, if algorithms or AI are not necessary for “easy tasks” such as computer vision, then the scope of AI reaches towards greater capabilities. Our respondent describes an AI process as something that “learns from its mistakes” and “becomes increasingly invasive,” a conception of AI we refer to as the *Terminator Fallacy*. Here we draw a parallel to the dystopian science fiction movie, *The Terminator*. The eponymous character of this movie is an AI-powered robot who acts in intelligent, recognizably human ways, contributing to its complex (and deadly) capabilities. While there is no consensus definition of artificial intelligence among experts working in the field, we note that the definitive landmark textbook Russell and Norvig’s (2016) *Artificial Intelligence: A Modern Approach* includes sections on handwritten character recognition and on image recognition in its advanced chapters, suggesting that these capabilities found in ALPR systems are included within conventional definitions of artificial intelligence.

5.3 Ordinance Comparison

In addition to our in-depth case study in Seattle, we also considered regulatory efforts in other U.S. jurisdictions at the time of writing. Table 4 summarizes these comparisons. Overall, we concluded that the Seattle ordinance was among the strongest and most comprehensive. Opportunities for public engagement are more robust in the Seattle ordinance than recent efforts in other areas, such as Nashville, where there appears to be no public engagement component. The Seattle ordinance is also the most thorough with respect to providing mechanisms for further review of a surveillance system’s functionality if it expands after initial approval.

Defining Surveillance Technology

In all of the ordinances we examined, devices and systems are regulated only if they meet the ordinance definition of “surveillance technology” and are not excluded by the ordinance by name (see Table 1 for determination criteria in Seattle). Table 4 (below) summarizes the definition of surveillance technology used in each ordinance, indicates whether the law includes references to algorithmic systems, whether an intent to minimize harm to marginalized groups is mentioned, and summarizes community participation in evaluating the technologies. We evaluated which definitions of surveillance technology prompt the evaluation of algorithmic features that may be present in a technology, such as the image identification capabilities of a camera system, or analytic systems that act on non-surveillance data to surface surveillant insights about subjects or events. We note that the surveillance technology definitions employed in Cambridge, Davis, Nashville, and Oakland all contain some reference to systems that “process” information. Of those four, all but Davis also refer to systems that “analyze.” The Seattle and Berkeley ordinances do not include either of those words. In each ordinance, we find no explicit treatment of algorithmic systems in any of the other ordinances we reviewed.

Jurisdiction	Surveillance Technology Definition	Mentions algorithms	Mentions disparate impact, marginalization, or racism	Community Oversight Process
Seattle	“designed or primarily intended to be used for the purpose of surveillance.”	No	Yes	Surveillance Impact Reports must include public comments from community meetings, city website. Civil Liberties Impact Assessment to be provided by an appointed Community Surveillance Working Group. Annual review includes complaints or concerns received by departments from the public.
Berkeley	“Designed, or primarily intended to remotely and surreptitiously collect...”	No	No	Surveillance Technology Reports to include complaints received by the City from the public.
Cambridge	“capable of collecting, capturing, recording, retaining, processing, intercepting, analyzing, monitoring, or sharing...”	No	Yes	Community concerns or complaints to be included in Annual Surveillance Report.

Davis	"used, designed, or primarily intended to collect, retain, process, or share..."	No	No	Community concerns or complaints to be included in Annual Surveillance Report.
Nashville	"capable of collecting, capturing, recording, retaining, processing, intercepting, analyzing, monitoring, or sharing..."	No	No	None specified.
Oakland	"used, designed, or primarily intended to collect, retain, analyze, process, or share..."	No	Yes	Privacy Advisory Commission to review SIR, Surveillance Use Policy, and Annual Surveillance Reports prior to Council approval. Annual Surveillance Reports to include complaints received by the City from the public.

Table 4. Comparison of municipal surveillance ordinances.

The Seattle definition is the most narrow in limiting the definition to technologies “designed or primarily intended to be used for the purpose of surveillance” (Seattle Municipal Code § 14.18.010). While many surveillance technologies and data sets would likely fall under this definition, other relevant technologies would not. For instance, in New York City, E-ZPass readers, which detect RFID identifiers carried in cars and trucks intended for highway and bridge tolling, had been installed and re-purposed throughout Manhattan to measure traffic patterns without relevant data governance policy limiting how resulting data traces were used (Hirose 2015).

Community Oversight

Only the Seattle and Oakland ordinances include concrete community oversight mechanisms. In Seattle, civil rights advocates and Councilmember González were motivated to give a voice to members of communities historically targeted by government surveillance in the city’s adoption of surveillance technologies. As Councilmember González, said, “There is no doubt that when surveillance technology is used in the law enforcement context, it will be used primarily against black and brown communities” (González 2018). Public engagement in Seattle includes public presentations and structured discussion, a

public comment period, and privacy and civil liberties impact assessments of the Community Surveillance Working Group. The Oakland Surveillance Ordinance requires oversight by a “Privacy Advisory Commission,” who is provided an opportunity to make recommendations to the City Council for every proposed acquisition or change of previously approved uses of a surveillance technology (Oakland Municipal Code §9.64.020). Privacy Advisory Commission members are appointed by the Mayor and subject to approval by the City Council. While membership on the Oakland Privacy Advisory Commission is not as narrowly targeted as Seattle’s Community Surveillance Working Group (see Section 3), it is designed to represent a range of interests including privacy and technology activists, legal scholars, financial professionals, and technology experts (Oakland City Council Ordinance 13349 C.M.S., 2016).

Algorithmic Oversight

We find that all of the recent surveillance regulations we reviewed, even those currently being drafted, have similar limitations in service of explicit processes for algorithmic accountability and transparency. While four of the six ordinances we analyzed include references to the processing of information in their technology definitions, and three of them also mention data analysis, general language such as “analyze” and data “processing” may not highlight many inferential, machine learning, or automated processes with their attendant concerns. Whereas a primary motivation of the Seattle Ordinance (and those of Cambridge and Oakland) is to surface to the civil rights and discrimination dimensions of data intensive systems, adding additional criteria to existing reporting processes could further surface algorithmic risks and potential algorithmic bias.

5.4 Easy Wins for Strengthening Surveillance Ordinances

The stated aim of the law is to surface the civil liberties and disparate impacts of surveillance technologies, primarily via public engagement and detailed reporting processes that seek to surface the functionality and use of each system. Given these factors, we find that the reporting tools used in the ordinance’s implementation would be improved against the law’s stated aims if they attend to the analytic capabilities of these systems, for example, by including questions that would distinguish in detail between different data processing methods. Adding these questions to existing reporting tools and processes would increase the effectiveness of the SIR reporting process as an exercise in surfacing the disparate impact of such systems, by increasing the ability of existing reporting processes to surface risks particular to algorithmic systems, such as classification errors, false positives, and feedback loops.

The existing reporting template calls for general descriptions of how a technology works, what data it collects, and how it is used; these existing processes would be strengthened by adding heuristics to existing reporting processes that help identify whether a technology incorporates predictive or other algorithmic features, and what failure modes to anticipate for each of those features. For example, an impact report template like the one used in Seattle could include prompts such as “Does the technology rely on a software algorithm for analysis?”, and “What harms are possible if this software makes a mistake?” or “What information is a human operator being given about the reliability of the data?” or “Have the underlying analytic systems been independently evaluated for their potential for misclassification or other errors?”

5.5 Broader Policy Implications

This work also has broader implications in the area of tech regulation. Our findings suggest differences in how government employees versus the academic community define algorithmic systems. Although the Seattle Surveillance Ordinance was not explicitly billed as algorithmic accountability and transparency

legislation, the potential to overlook algorithmic harms could obtain in other efforts towards tech fairness legislation (such as two recent draft bills in Washington State in 2019), autonomous vehicle laws, or regulation of algorithmic warfare and autonomous weapons. In the case of autonomous weapons, for instance, one of the major barriers to consensus in international discussions has been a lack of agreement over what counts as an autonomous weapon (c.f. Russian Federation, The 2018; Conn 2016). Lessons may be drawn from other communities of practice for next steps in defining and classifying levels of automation and autonomy. For example, in the domain of automated driving vehicles, the Society for Automotive Engineers International defines five levels of vehicle automation: no automation, driving assistance, partial automation, conditional automation, high automation, and full automation; these levels include detailed definitions for making assessments that have been useful in ongoing regulatory efforts (Smith 2017). Our work affirms a need for both definitional and translational work to bring clearer and more consistent understandings of algorithmic systems to a wider array of stakeholder groups. To remediate this gap, future work could develop reporting tools and heuristics that might bolster existing oversight processes like those found in the Surveillance Impact Reports used in implementing the Seattle Surveillance Ordinance.

6. Conclusion

We observe that several efforts to regulate surveillance technology in the U.S. are motivated by a goal to surface disparate impact to social groups who have been historically targeted by surveillance practices. On its own terms, the Seattle law is addressed to “surveillance technology” with an interest in the racial and social justice impact of surveillance; we find its reporting documents are primarily concerned with the data collection function of surveillance systems, as opposed to analytic processes. Second, we found that city policymakers and personnel drew distinctions between the surveillance technologies under review and algorithmic systems, even as several technologies incorporated computer vision and automated

inference techniques. Given that the aim of the law is to surface the disparate impact of technologies adopted by city agencies, we argue that its implementation of the law may be further strengthened with additional criteria that would address the data processing, analysis, and classification processes performed on surveillance data.

Targeted changes to reporting procedures could surface analytic capabilities and enable more effective inquiry into potential disparate impacts in surveillance systems. Integrating finer-grained guidance on relevant distinctions between types of algorithmic and information systems would strengthen regulatory efforts by making the potential harms of underlying algorithmic components more legible to political and community stakeholders. Future work will explore more scalable heuristics and definitional distinctions by which algorithmic systems could become legible to lawmakers as machine learning and AI, considering the degree to which their presence in the system is “relational,” that is, shifting with respect to the viewer (Star & Ruhleder, 1996). To this end, we recognize an opportunity to extend lessons from the policymaking of autonomous vehicles (Smith, 2017) and weapons (Lewis, 2014) to provide a conceptual schema characterizing degrees of automation, intended for a non-specialist, policymaking audience, as well as evaluating them against a list of risks particular to these technologies, such as false positives, misclassification, and feedback loops. Such an approach could better equip policymakers to locate uses of machine learning in otherwise quotidian technologies, and pull discussion of artificial intelligence regulation toward closer consideration of systems in current use.

References

ACLU Washington (2017). Seattle Adopts Nation’s Strongest Regulations for Surveillance Technology. Press release, August 8, 2017.

ACLU (n.d.). Community Control Over Police Surveillance Campaign.

<<https://www.aclu.org/issues/privacy-technology/surveillance-technologies/community-control-over-police-surveillance>>. Accessed July 10, 2019.

Ananny, M. (2016). Toward an ethics of algorithms: Convening, observation, probability, and timeliness. *Science, Technology, & Human Values*, 41(1):93–117.

Angwin, J., Larson, J., Mattu, S., and Kirchner, L. (2016). Machine bias. *ProPublica*, May, 23.

Atkinson, P., & Hammersley, M. (1998). Ethnography and participant observation. *Strategies of Qualitative Inquiry*. Thousand Oaks: Sage, 248-261.

Barocas, S. and Selbst, A. D. (2016). Big data's disparate impact. *Cal. L. Rev.*, 104:671.

Bowen, G. A. (2009). Document analysis as a qualitative research method. *Qualitative research journal*, 9(2), 27-40.

Brayne, S. (2017). Big data surveillance: The case of policing. *American Sociological Review*, 82(5):977–1008.

Buolamwini, J. and Gebru, T. (2018). Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on Fairness, Accountability and Transparency*, pages 77–91.

Caliskan, A., Bryson, J. J., and Narayanan, A. (2017). Semantics derived automatically from language corpora contain human-like biases. *Science*.

Chouldechova, A. (2017). Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. *Big data*, 5(2):153–163.

Christin, A. (2017). Algorithms in practice: Comparing web journalism and criminal justice. *Big Data & Society*, 4(2):2053951717718855.

Conn, A. (2016) The problem of defining autonomous weapons. The Future of Life Institute. Blog post, Nov. 30, 2016.

Crawford, K. (2016). Artificial intelligence's white guy problem. *The New York Times*.

Crump, C. (2016). Surveillance policy making by procurement. *Wash. L. Rev.*, 91:1595.

Danaher, J., Hogan, M. J., Noone, C., Kennedy, R., Behan, A., De Paor, A., Felzmann, H., Haklay, M., Khoo, S.-M., Morison, J., et al. (2017). Algorithmic governance: Developing a research agenda through the power of collective intelligence. *Big Data & Society*.

Doshi-Velez, F., & Kim, B. (2017). Towards a rigorous science of interpretable machine learning. *arXiv preprint arXiv:1702.08608*.

Employee, S. P. D. (2018b). Seattle Police Department interview.

Employee, S. P. D. (2018c). Seattle Police Department interview.

Ensign, D., Friedler, S. A., Neville, S., Scheidegger, C., and Venkatasubramanian, S. (2017). Runaway feedback loops in predictive policing. *arXiv preprint, arXiv:1706.09847*.

Eubanks, V. (2018). *Automating inequality: How high-tech tools profile, police, and punish the poor*. St. Martin's Press.

Fellner, J., & Human Rights Watch. (2009). *Decades of disparity: drug arrests and race in the United States*. New York, NY: Human Rights Watch.

Freelon, D., McIlwain, C. D., & Clark, M. (2016). Beyond the hashtags:# Ferguson,# Blacklivesmatter, and the online struggle for offline justice. *Center for Media & Social Impact, American University*.

Garvie, C. (2016). *The perpetual line-up: Unregulated police face recognition in America*. Georgetown Law, Center on Privacy & Technology.

Guidotti, R., Monreale, A., Ruggieri, S., Turini, F., Giannotti, F., & Pedreschi, D. (2018). A survey of methods for explaining black box models. *ACM Computing Surveys (CSUR)*, 51(5), 93.

González, L. (2018). Lorena González interview.

Green, Ben. (2019). *The Smart Enough City: Putting Technology in Its Place to Reclaim Our Urban Future*. MIT Press.

Herz, A. (2016). How the Seattle Police Secretly—and Illegally—Purchased a Tool for Tracking Your Social Media Posts. *The Stranger*.

Hirose, M. (2015) Newly Obtained Records Reveal Extensive Monitoring of E-ZPass Tags Throughout New York. ACLU, April 24, 2015.

Jones, M. L. (2017). The right to a human in the loop: Political constructions of computer automation and personhood. *Social studies of science*, 47(2), 216-239.

Katell, M. A. (2018). Adverse detection: The promise and peril of body-worn cameras. In *Surveillance, Privacy and Public Space*, pages 111–130. Routledge.

Levine, D. S. (2007). Secrecy and Unaccountability: Trade Secrets in Our Public Infrastructure. *Florida Law Review*, 59, 135–194.

Lewis, J. (2014). The case for regulating fully autonomous weapons. *Yale LJ*, 124:1309.

Madden, M. (2014). Public Perceptions of Privacy and Security in the Post-Snowden Era.

Madden, M., Gilman, M., Levy, K., and Marwick, A. (2017). Privacy, poverty, and big data: A matrix of vulnerabilities for poor americans. *Wash. UL Rev.*, 95:53.

Mai, J.-E. (2016). Big data privacy: The datafication of personal information. *The Information Society*, 32(3):192–199.

Mittelstadt, B. D., Allo, P., Taddeo, M., Wachter, S., and Floridi, L. (2016). The ethics of algorithms: Mapping the debate. *Big Data & Society*, 3(2):2053951716679679.

Mocek, P. “FOI gone afoul: Seattle City Council defer to Police Department on what is good for the public to know” Mocek.org, March 17, 2014.

<<https://mocek.org/blog/2014/03/17/foi-gone-afoul-seattle-city-council-defer-to-police/>>.

Moy, L. (2019). How Police Technology Aggravates Racial Inequity: A Taxonomy of Problems and a Path Forward. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.3340898>

Narayan, S. (2018). Shankar Narayan interview.

Noble, Safiya Umoja. (2018). *Algorithms of Oppression: How Search Engines Reinforce Racism*. NYU Press.

Pasquale, F. (2015). *The black box society: The secret algorithms that control money and information*. Harvard University Press.

Russell, S. J., & Norvig, P. (2016). *Artificial Intelligence: A Modern Approach*. Malaysia; Pearson Education Limited.

Russian Federation, The. (2018). *Russia's Approaches to the Elaboration of a Working Definition and Basic Functions of Lethal Autonomous Weapons Systems in the Context of the Purposes and Objectives of the Convention*. Convention on Certain Conventional Weapons. Group of Governmental Experts on Lethal Autonomous Weapons, WP.6.

Smith, B. W. (2017). How governments can promote automated driving. *NML Rev.*, 47:99.

Snider, L. (2014). Interrogating the algorithm: Debt, derivatives and the social reconstruction of stock market trading. *Critical Sociology*, 40(5):747–761.

Star, S. L. and Ruhleder, K. (1996). Steps toward an ecology of infrastructure: Design and access for large information spaces. *Information systems research*, 7(1):111–134.

Sweeney, L. (2013). Discrimination in online ad delivery. *Communications of the ACM*, 56(5):44. 00114.

Thomas, S. L., Nafus, D., and Sherman, J. (2018). Algorithms as fetish: Faith and possibility in algorithmic work. *Big Data & Society*.

Warde, B. (2013). Black Male Disproportionality in the Criminal Justice Systems of the USA, Canada, and England: a Comparative Analysis of Incarceration. *Journal of African American Studies*, 17(4):461–479.

Weiss, Robert S. (1995) *Learning from strangers: The art and method of qualitative interview studies*. Simon and Schuster.

Winston, A. (2018). Palantir has secretly been using New Orleans to test its predictive policing technology.

Zarsky, T. (2016). The trouble with algorithmic decisions: An analytic road map to examine efficiency and fairness in automated and opaque decision making. *Science, Technology, & Human Values*, 41(1), 118-132.

Ziewitz, M. (2016). Governing algorithms: Myth, mess, and methods. *Science, Technology, & Human Values*, 41(1), 3-16.