Original Research Article

BIG DATA & SOCIETY

Big Data & Society
July–December 2019: 1–14

DOI: 10.1177/2053951719868492
journals.sagepub.com/home/bds
SAGE

Municipal surveillance regulation and algorithmic accountability

Meg Young[1], Michael Katell[1] and P. M. Krafft[2]

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, requirements for public engagement, and processes for establishing acceptable uses. Although the Seattle Surveillance Ordinance was 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 definitional gaps that could hinder future efforts toward algorithmic regulation. We argue that (i) finer-grained distinctions between types of information systems in the language of law and policy, and (ii) risk assessment tools integrated into their implementation would strengthen future regulatory efforts by rendering underlying algorithmic components more legible and accountable to political and community stakeholders.

Keywords

Accountability, algorithmic systems, ethnography, government, policy, surveillance

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 (Noble, 2018; Sweeney, 2013). 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 and Kim, 2017; Guidotti et al., 2018) or regulatory (Danaher et al., 2017; Jones, 2017; Zarsky, 2016; Ziewitz, 2016). Yet despite the groundswell of interest, persistent questions remain as to what interventions might be most effective.

As efforts around algorithmic regulation are in their early stages, we look to the related yet better-developed area of surveillance regulation, which shares common aims such as making public sector technologies legible to oversight. In recent years, surveillance regulations have been implemented in multiple municipalities across the United States in response to longstanding concerns about police surveillance practices, as well as to increasing public awareness about the use of data-intensive systems by government agencies. Surveillance ordinances have been passed in 13 US jurisdictions as of this writing.[1] Although surveillance laws are generally not intended as algorithmic accountability regulation, they provide a basis for responding to longstanding challenges in regulating algorithmic

[1]Information School, University of Washington, Seattle, WA, USA
[2]Oxford Internet Institute, University of Oxford, Oxford, UK

Corresponding author:
Meg Young, University of Washington, 1851 Grant Ln, Mary Gates Hall, Seattle, WA 98105, USA.
Email: megyoung@uw.edu

systems, in light of their progress toward subjecting surveillance technologies to public oversight.

Here, 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 (AI) techniques—to make or support decisions, judgments, or assessments, we also examine technologies disclosed under the Seattle Surveillance Ordinance that fit this definition 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.

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 functionality (Thomas et al., 2018), and intellectual property protections and "trade secret" laws are used to block efforts to make deeper assessments (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 US 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 marketing pitches promoting algorithmic 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. Some recent studies have begun to move towards this aim (Ananny and Crawford, 2016; Goodman, 2016; Yeung, 2017; Alkhatib and Bernstein, 2019; LaBrie and Steinke 2019; Morley et al. 2019; Reddy et al., 2019).

The push for algorithmic regulation is 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 (Eubanks, 2018; Madden et al., 2017), and in turn more likely to become the subjects of surveillance. The addition of data analytics transforms acquired data into raw material for predictions about human behavior that support automated decisions and interventions. Van Dijck (2014) analyses these practices as part of the logic of as "datafication," a dimension of surveillance in which machine learning and other algorithmic techniques surface patterns from prior events, 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, modern algorithmic systems introduce new risks as compared to 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. The resulting risks are particularly evident in predictive policing systems, which are susceptible for example to a type of risk known as "runaway feedback loops" (Ensign et al., 2017). At one popular system calculates a "risk score" based on a range of factors including the number of encounters the subject has had with police. Individuals assigned higher scores are more likely to be policed and each encounter further increases their score (Joh 2014). 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.

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 (American Civil Liberties Union (ACLU) Washington, 2017; Green, 2019). We adopted 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 first 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 second aim.

Our case study consists of 14 semi-structured and open-ended interviews (Weiss, 1995), participant observation (Atkinson and 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 US 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 its first and focused on administrative issues. Our observation also included online observation of news coverage, community blogs, and 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 (SIRs), 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 AI, we also sought input from experts in machine learning and AI via a survey of researchers in those fields (c.f. Krafft et al., 2019). 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, and informed our data analysis by attuning us to how functionalities 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 (ACLU).

Case study

Our case study traces the motivation and implementation of the current Seattle Surveillance Ordinance between 2017 and 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 (SPD) had acquired surveillance technologies without clear policy guidance or public disclosure; including a drone aircraft procured with federal funding, a closed circuit television (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 efforts ultimately led multiple cities in the US 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, some local observers criticized the 2013 effort, arguing it had fallen short of providing meaningful transparency and accountability. For instance, five months after the initial surveillance 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 up to that point (Mocek, 2014). Then, in 2016, a local newspaper reported that the 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. Council member Lorena González, a former civil rights lawyer and Chair of the Public Safety Committee, responded by leading 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 these concerns and others, 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 such 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, 2018a). Councilmember González recalled that similar concerns from the police department were a 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 crucial 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, 2018b). 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 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.

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 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.

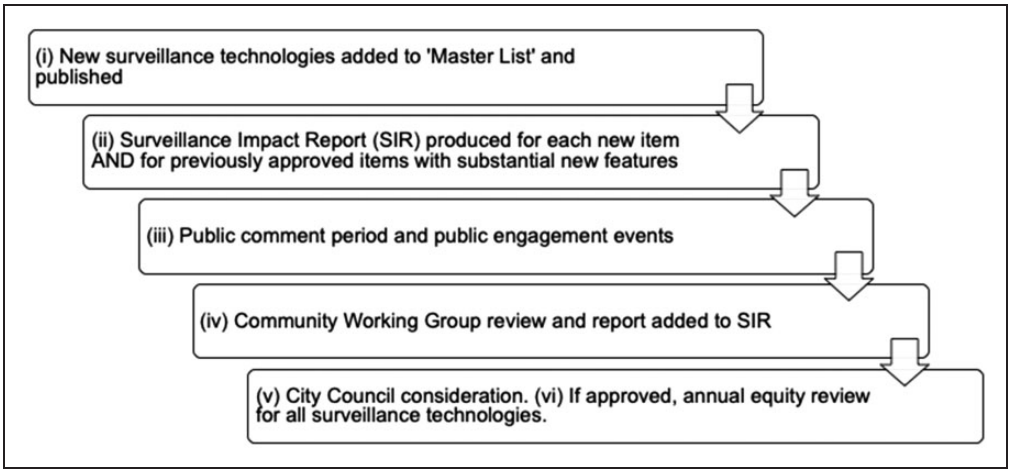

Figure 1. Reporting, review, and approval process in the Seattle Surveillance Ordinance.

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. As of our 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 ALPRs 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 surveillance technologies

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 AI components. Empty cells indicate information the authors could not obtain, judge, or infer with a reasonable degree of confidence. Column descriptions follow.

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

Table 1. Surveillance Technology Determination Criteria.

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.
	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.

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

SPD	Seattle Police Department
SCL	Seattle City Light
SDT	Seattle Department of Transportation
SFD	Seattle Fire Department
CD	Collects and stores digital data
AI	Uses 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

tools) already in use in other major municipalities. Many of the listed technologies are traditional surveillance hardware, such as cameras like 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 AI—two forms of automated license plate recognition, and a Booking Photo Comparison Software (BPCS). 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 us 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 BPCS used face recognition, its use may eventually be discontinued due

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

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	1	1	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	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)											
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	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						1		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	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			
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	1	0	0	1	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											

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.

Discussion

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 process (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 sounds and a notification appears 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 SPD 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 US 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, 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 and 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: 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, which leads to greater potential for software to identify—or misidentify—people of color. These considerations indicate a need to consider such risks in the existing evaluation procedures in pursuit of surfacing the disparate impacts of existing surveillance technologies.

Mental models of artificial intelligence and algorithmic systems

Several conversations in our field work indicated that broader definitions of algorithmic systems and AI 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 intel— . . . 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 interviews. 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 likely driven by two factors; first, an expectation that tasks which are easy for people—such as vision—should also be easy or engineers to program into computers; we call this the *Minsky Fallacy*. An apocryphal tale has it that Marvin Minsky, a 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 drastically Minsky underestimated the difficulty of these two computational 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, like Minsky, 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 AI 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 AI 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 AI.

Ordinance comparison

In addition to our in-depth case study in Seattle, we also considered regulatory efforts in other US 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 changes 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 other clauses (see Table 1 for the determination criteria in Seattle). Table 4 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 whether it includes 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 data analysis systems that act on passively collected data such as web or social media data to surface 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 ordinance definitions do not include either of those words. We find no explicit treatment of algorithmic systems in any of the other ordinances we reviewed.

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, 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).

Table 4. Comparison of municipal surveillance ordinances.

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.

Algorithmic oversight. We find that all of the recent surveillance regulations we reviewed, even those currently being drafted, have similar limitations with respect to 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. A primary motivation of ordinances in Seattle, Cambridge, and Oakland is to surface the civil rights implications of data intensive systems; here we identify opportunities to strengthen these existing reporting processes to surface algorithmic harms.

Easy wins for strengthening surveillance ordinances

The stated aim of the Seattle Surveillance Ordinance 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 with respect to the law's stated aims if these tools attended to the algorithmic details of these systems, for example, by including questions that would distinguish 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 ''Has any underlying data analysis software been independently evaluated for its potential for misclassification or other errors?''

Broader policy implications

Our 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 occur 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 progress in international discussions has been a lack of agreement over what counts as an autonomous weapon (c.f. Conn, 2016; Russian Federation, 2018). 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 SIRs used in implementing the Seattle Surveillance Ordinance.

Conclusion

We observe that several efforts to regulate surveillance technology in the US 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 data analysis processes. We also found that city policymakers and personnel drew distinctions between the surveillance technologies under review and algorithmic systems, even as disclosed technologies incorporated computer vision and automated inference techniques. Given that the Seattle law aims to surface the disparate impact of technologies adopted by city agencies, we argue that its implementation would be further strengthened with additional criteria that would address the data processing, analysis, and classification functions performed on surveillance data, with a focus on the disparate impacts that result from those functions.

Targeted changes to reporting procedures could surface computational functions and malfunctions, enabling more effective inquiry into potential disparate impacts of 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 could 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 and Ruhleder, 1996). To this end, we recognize an opportunity to extend lessons from the policy making of autonomous vehicles (Smith, 2017) and weapons (Lewis, 2014) to provide a conceptual schema characterizing degrees of automation, intended for a non-specialist, policy making audience. Such a schema could also promote evaluation against a list of risks particular to these technologies, such as false positives, misclassification, and feedback loops. This approach would better equip policymakers to locate uses of machine learning in otherwise quotidian technologies, and pull discussion of AI regulation toward closer consideration of systems currently in use.

Acknowledgements

We would like to thank the anonymous reviewers for their helpful suggestions and feedback. Great thanks to Anna Lauren Hoffmann, Jaime Snyder, Nic Weber, Quinn DuPont, Justin Petelka, Caitie Lustig, and the DataIRL Group at The University of Washington for their input, which greatly improved earlier drafts. We also thank Shankar Narayan for sharing the Master List with us in early conversations on the topic. All omissions and errors are our own.

Declaration of conflicting interests

The author(s) declared the following potential conflicts of interest with respect to the research, authorship, and/or publication of this article: Michael Katell has received funding from the American Civil Liberties Union in a subsequent related project, which was not offered, sought, or expected at the time the present work was conducted.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: This work was supported in part by the University of Washington's Information School through the Innovation Award titled, "Municipal Technology Policy: Algorithms and Accountability in Practice", by the Washington Research Foundation, and by a Data Science Environments project award from the Gordon and Betty Moore Foundation (Award #2013-10-29) and the Alfred P. Sloan Foundation (Award #3835) to the University of Washington eScience Institute.

Notes

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

ORCID iD

Meg Young https://orcid.org/0000-0002-9300-8575
Michael Katell https://orcid.org/0000-0003-2200-6246

References

Alkhatib A and Bernstein M (2019) Street-Level Algorithms: A Theory at the Gaps Between Policy and Decisions. In: *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, p. 530.

Ananny M and Crawford K (2018) Seeing without knowing: Limitations of the transparency ideal and its application to algorithmic accountability. *New Media & Society* 20(3), 973–989.

American Civil Liberties Union (ACLU) (n.d.) Community control over police surveillance campaign. Available at: www.aclu.org/issues/privacy-technology/surveillance-technologies/community-control-over-police-surveillance (accessed 10 July 2019).

American Civil Liberties Union (ACLU) Washington (2017) Seattle adopts nation's strongest regulations for surveillance technology, 8 August.

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, et al. (2016). Machine bias. *ProPublica*, 23 May.

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

Barocas S and Selbst AD (2016) Big data's disparate impact. *California Law Review* 104: 671.

Bowen GA (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*, 23-24 February 2018, New York University, NYC, pp. 77–91.

Caliskan A, Bryson JJ and Narayanan A (2017) Semantics derived automatically from language corpora contain human-like biases. *Science* 356(6334): 183–186.

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. In: *The Future of Life Institute*, 30 November.

Crawford K (2016) Artificial intelligence's white guy problem. *The New York Times*, 25 June.

Crump C (2016). Surveillance policy making by procurement. *Washington Law Review* 91: 1595.

Danaher J, Hogan MJ, Noone C, et al. (2017) Algorithmic governance: Developing a research agenda through the power of collective intelligence. *Big Data & Society*. Epub ahead of print 2017. DOI: 10.1177/2053951717726554.

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

Ensign D, Friedler SA, Neville S, et al. (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. New York, NY: St. Martin's Press.

Fellner J and Human Rights Watch (2009) *Decades of Disparity: Drug Arrests and Race in the United States*. New York, NY: Human Rights Watch.

Freelon D, McIlwain CD and Clark M (2016) Beyond the hashtags:# Ferguson,# Blacklivesmatter, and the online struggle for offline justice. American University, USA.

Goodman BW (2016) A step towards accountable algorithms? Algorithmic discrimination and the european union general data protection. Neural Information Processing Systems Symposium on Machine Learning and the Law.

González L (2018) Lorena González interview.

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

Guidotti R, Monreale A, Ruggieri S, et al. (2018) A survey of methods for explaining black box models. ACM Computing Surveys *(CSUR)* 51(5): 93.

Herz A (2016) How the Seattle Police secretly – and illegally – Purchased a tool for tracking your social media posts. *The Stranger*, 28 September.

Hirose M (2015) Newly obtained records reveal extensive monitoring of E-ZPass tags throughout New York. ACLU, 24 April.

Joh, E. E. (2014). Policing by Numbers: Big Data and the Fourth Amendment. *Washington Law Review* 89, 35.

Jones ML (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 MA (2018) Adverse detection: The promise and peril of body-worn cameras. In: *Surveillance, Privacy and Public Space*. Abingdon: Routledge, pp. 111–130.

Krafft PM, Young M, Katell M, et al. (2019) Policy versus Practice: Conceptions of Artificial Intelligence. *SSRN Electronic Journal*. DOI: 10.2139/ssrn.3431304.

LaBrie RC and Steinke G (2019) Towards a Framework for Ethical Audits of AI Algorithms. Twenty-Fifth Americas Conference on Information System.

Levine DS (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 Law Journal* 124: 1309.

Madden M (2014) Public Perceptions of Privacy and Security in the Post-Snowden Era. In: Pew Research Center's Internet & American Life Project website, 11 February 2015. Available at: http://www.pewinternet.org/2014/11/12/public-privacy-perceptions/.

Madden M, Gilman M, Levy K, et al. (2017) Privacy, poverty, and big data: A matrix of vulnerabilities for poor Americans. *Washington University Law Review* 9553.

Mittelstadt BD, Allo P, Taddeo M, et al. (2016) The ethics of algorithms: Mapping the debate. *Big Data & Society* 3(2): 2053951716679679.

Mocek P (2014) FOI gone afoul: Seattle City Council defer to Police Department on what is good for the public to know. In: Mocek.org, 17 March. Available at: https://mocek.org/blog/2014/03/17/foi-gone-afoul-seattle-city-council-defer-to-police/ (accessed 23 August 2018).

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

Morley J, Floridi L, Kinsey L and Elhalal A (2019) From What to How. An Overview of AI Ethics Tools, Methods and Research to Translate Principles into Practices. arXiv preprint arXiv:1905.06876.

Narayan S (2018) Shankar Narayan interview.

Noble SU (2018) *Algorithms of Oppression: How Search Engines Reinforce Racism*. New York: NYU Press.

Pasquale F (2015) *The Black Box Society: The Secret Algorithms that Control Money and Information*. Cambridge, MA: Harvard University Press.

Police Department Employee (2018a) Seattle Police Department interview.

Police Department Employee (2018b) Seattle Police Department interview.

Reddy E, Cakici B and Ballestero A (2019) Beyond mystery: Putting algorithmic accountability in context. *Big Data & Society* 6(1): 2053951719826856.

Russell SJ and Norvig P (2016) *Artificial Intelligence: A Modern Approach*. Kuala Lumpur: Pearson Education Limited.

Russian Federation (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 BW (2017) How governments can promote automated driving. *New Mexico Law Review* 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 SL 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): 44–54.

Thomas SL, Nafus D and Sherman J (2018) Algorithms as fetish: Faith and possibility in algorithmic work. *Big Data & Society* 5(1): 2053951717751552.

Van Dijck, J. (2014). Datafication, dataism and dataveillance: Big Data between scientific paradigm and ideology. *Surveillance & Society* 12(2), 197–208. https://doi.org/10.24908/ss.v12i2.4776.

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 RS (1995) *Learning from Strangers: The Art and Method of Qualitative Interview Studies*. New York, NY: Simon and Schuster.

Yeung K (2018) Algorithmic regulation: A critical interrogation. *Regulation & Governance* 12(4): 505–523.

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.