Article

Predictive Policing for Reform? Indeterminacy and Intervention in Big Data Policing

Aaron Shapiro

University of Pennsylvania, USA
ashapiro@asc.upenn.edu

Abstract

Predictive analytics and artificial intelligence are applied widely across law enforcement agencies and the criminal justice system. Despite criticism that such tools reinforce inequality and structural discrimination, proponents insist that they will nonetheless improve the equality and fairness of outcomes by countering humans' biased or capricious decision-making. How can predictive analytics be understood simultaneously as a source of, and solution to, discrimination and bias in criminal justice and law enforcement? The article provides a framework for understanding the techno-political gambit of predictive policing as a mechanism of police reform—a discourse that I call “predictive policing for reform.” Focusing specifically on geospatial predictive policing systems, I argue that “predictive policing for reform” should be seen as a flawed attempt to rationalize police patrols through an algorithmic remediation of patrol geographies. The attempt is flawed because predictive systems operate on the sociotechnical practices of police patrols, which are themselves contradictory enactments of the state's power to distribute safety and harm. The ambiguities and contradictions of the patrol are not resolved through algorithmic remediation. Instead, they lead to new indeterminacies, trade-offs, and experimentations based on unfalsifiable claims. I detail these through a discussion of predictive policing firm HunchLab's use of predictive analytics to rationalize patrols and mitigate bias. Understanding how the “predictive policing for reform” discourse is operationalized as a series of technical fixes that rely on the production of indeterminacies allows for a more nuanced critique of predictive policing.

A Techno-Political Gambit

Predictive analytics and artificial intelligence (AI) are now used regularly across criminal justice institutions and law enforcement agencies. Judges, parole boards, police commanders, and patrol officers make assessments, evaluations, and allocations every day that have profound impacts on the lives of those caught up in the system—suspects, defendants, parolees. Automated data analysis promises to make such institutional decision-making more objective, consistent, and rigorous—neutral and value-free.

Critical investigations of these systems, however, show that fairness claims do not hold up under scrutiny (Angwin et al. 2016; Lum and Isaac 2016). Julia Angwin et al. (2016) famously found that an algorithm used to determine bail was more likely to flag African American defendants as “high risk” than white defendants. The idea that machine learning algorithms and AI inherit the biases of their makers or even society at large now appears regularly in mainstream news (e.g., Crawford 2016; Lohr 2018; Wachter-Boettcher 2017). In the context of law enforcement and “predictive policing” applications, the focus has been on the data used to train predictive algorithms. Data that are limited, incomplete, inaccurate, or biased due to discriminatory policing practices stand to reinforce disparate treatments for already marginalized communities (Brayne, Rosenblat, and boyd 2015; Ensign et al. 2017; Jefferson 2017; Lum and Isaac 2016).

Shapiro, Aaron. 2019. Predictive Policing for Reform? Indeterminacy and Intervention in Big Data Policing. *Surveillance & Society* 17(3/4): 456-472.

<https://ojs.library.queensu.ca/index.php/surveillance-and-society/index> | ISSN: 1477-7487

© The author(s), 2019 | Licensed to the Surveillance Studies Network under a [Creative Commons Attribution Non-Commercial No Derivatives license](#)

In 2016, the ACLU (2016) issued a statement of concern, co-signed by several civil rights organizations, about applications of AI and predictive analytics in law enforcement. The statement argues that “[t]he data driving predictive enforcement activities—such as the location and timing of previously reported crimes, or patterns of community- and officer-initiated 911 calls—is profoundly limited and biased,” concentrating law enforcement in already over-policed communities while ignoring underlying causes of crime—structural racism, systemic disinvestment, and poverty.

Despite these powerful critiques, proponents—computer scientists, technology vendors, and policymakers—nonetheless insist that predictive systems not only improve efficiency in the criminal justice system but will ultimately “make the legal system fairer” (Siegel 2018). In an article for *Scientific American* titled “How to Fight Bias with Predictive Policing,” Eric Siegel (2018) describes predictive policing as “an unprecedented opportunity for racial justice” and “the ideal platform on which new practices for racial equity may be systematically and widely deployed.” A *New York Times* op-ed reaches similar conclusions. While the authors—computer scientists and computational social scientists—acknowledge the need for checks and balances to avoid disparate outcomes in algorithmic decision-making systems, they ultimately conclude that “well-designed algorithms can counter the biases and inconsistencies of unaided human judgments and help ensure equitable outcomes for all” (Corbett-Davies, Goel, and González-Bailon 2017).

How can predictive analytics be understood as both the source of and solution to discrimination and bias in criminal justice and law enforcement? This article provides a framework for understanding the technological gambit of predictive policing as a mechanism of police reform—a discourse that I call “predictive policing for reform.” Understanding this discourse allows for a more nuanced critique of predictive policing and of advocates’ claims that it leads to equitable criminal justice outcomes. I focus specifically on geospatial or location-based predictive systems (forecasts of where and when crimes will take place) and argue that “predictive policing for reform” should be understood as *a flawed attempt to rationalize police patrols through an algorithmic remediation of patrol geographies*. As I show throughout the article, the attempt is flawed because predictive systems operate on the sociotechnical practices of police patrols, which are themselves contradictory enactments of the state’s power to distribute safety and harm. As Lucas Introna (2016: 20) argues, “algorithms must be understood in situated practices—as part of the heterogeneous sociomaterial assemblages within which they are embedded.” The patrol is an assemblage of socio-technological mediations that enact urban geographies of authority and legitimacy, risk and danger (Merrill and Hoffman 2015; Nail 2015; Reeves and Packer 2013). The patrol also serves as a primary space for the managerial oversight of police officers in the field (De Lint 2000). Patrol technologies have thus historically been linked with efforts toward police reform and increased professionalism (see Manning 2008; Sklansky 2014). Predictive policing performatively embeds data-driven decision-making systems within these sociotechnical and institutional practices. It enrolls patrol officers into the predictive data structure as imperfect instruments of measurement and observation, through which reform efforts are imagined to take hold.

The argument builds on ethnographic research conducted with the product team at HunchLab, a geospatial predictive policing software suite used by several law enforcement agencies across the US.¹ The article details three incongruities that result from the complex and imperfect embeddedness of data-driven predictions with the sociotechnical practices of patrol: *indeterminacies*, *trade-offs*, and *unfalsifiability*. Understanding these incongruities and how they are strategically deployed by HunchLab illuminates fallacious assumptions baked into the logic of “predictive policing for reform.” The article concludes by discussing two incommensurable views of police patrol—as a distribution of public safety and a distribution of harm—to highlight the role of indeterminacy in legitimating intervention.

¹ Azavea sold HunchLab to ShotSpotter in October 2018, after the manuscript was accepted for publication. The findings reflect the version of HunchLab used up until that point.

The Patrol as Medium

The future-oriented risk management strategies of predictive policing are less novel than many observers would likely concede (e.g., Mantello 2016). Modern “liberal consent policing” has always involved geographies of preemption and anticipation (Anderson 2010; Massumi 2007) as a means to prevent crime and to govern racialized populations (De Lint 2000; Hartman 1997; Wood 2009). As Markus Dubber and Mariana Valverde (2006: 4–5) argue, the science of police, developed from the seventeenth through the nineteenth centuries, served a “hinge” function for the burgeoning liberal state, linking otherwise incommensurable and temporally disjointed logics of power: the backward-looking criminal law, which punishes crimes *ex post facto* and the future-oriented logics of prevention, which intervene before any infraction or crime has been committed.

The blurred line between prevention and punishment becomes especially problematic when the “power to police” is translated into actually existing police forces (Dubber and Valverde 2006). In the US and other colonial contexts, prevention has been intimately tied to violence and the racialization of enslaved bodies through the spectacle of state terror. The first town watch systems in the American colonies were organized in the seventeenth century to prevent slaves from running away or revolting (Beutin 2017; Williams 2015; Wood 2009). When the first professional municipal police forces were established, prevention was incorporated into founding charters (Dubber and Valverde 2006; De Lint 2000). The first of Robert Peel’s (1829) famous “Nine Principles of Policing,” issued at the establishment of the London Metropolitan Police in 1829, reads: “The basic mission for which the police exist is to prevent crime and disorder.” It is no coincidence that former Commissioner of the New York City Police Department (NYPD) and outspoken proponent of predictive policing, William Bratton (2014; *New York Times* 2014), refers nostalgically to predictive policing as a technologically motivated return to Peel’s basic tenets of policing. Situating algorithmic prediction within these historical legacies helps explain why advocates can argue for predictive policing’s capacity to improve fairness in law enforcement (by preventing crime and disorder) in the face of mounting critiques that the use of algorithmic systems in law enforcement will exacerbate racial and socioeconomic disparities.

The most elemental crime prevention technology is the police patrol beat. “The police beat conceives [of] policing as a question of the allocation of men to territorialized (or spatialized) jurisdiction” (De Lint 2000: 64). The patrol operationalizes the preventive function of the police. It creates spatial and temporal schema to distribute mobile police officers throughout uneven landscapes of anticipated risks, hazards, and environmental conditions, which are evaluated as more or less conducive to criminality. For example, early police experts (e.g., Colquhoun 1800) argued that patrol resources should be allocated “based upon the mobility demands of the fleeing criminal and the communication imperatives of a responsive/preventive police apparatus” (Reeves and Packer 2013: 363). The police patrol is thus an innately anticipatory spatial technology. It reconfigures urban landscapes in order to deter aberrant and deviant activity, preempting the flight paths of suspects and establishing the police’s logistical control over patrolled jurisdictions.

Following anthropologist William Mazzarella (2004) and media theorists Sarah Kember and Joanna Zylinska (2012), the patrol can be productively theorized as a social mediation—“a reflexive and reifying technology” that “makes society imaginable and intelligible to itself” (Mazzarella 2004: 346). Media are the material and technological frameworks that performatively enact societal relations; they enable and constrain social practices. Like other media technologies and techniques, the police patrol is “inseparable from the movement of social life and yet removed from it . . . at once obvious and strange, indispensable and uncanny, intimate and distant” (ibid.: 346). The patrol-as-medium is a concrete manifestation of the abstract state power to police (Dubber and Valverde 2006). It provides a pragmatic link between the power-knowledge of population statistics and territorial cartographies (Foucault 1980, 2007) and the sovereign authority to intervene—to stop, search, and seize; to make move. But the patrol also provokes extreme ambiguity. Relative to other social mediations (e.g., collective actions, social movements, religious or national ceremonies, riots and mobs; cf. Mazzarella 2009; Shapiro 2017), the patrol-as-medium is shaped by the tensions of “liberal consent policing” or “policing by consent”: a “delicate balancing of respectful

protection and intrusive penetration” (De Lint 2000: 55). The patrol embodies the state’s monopoly on the legitimate use of violence while also having to perform a *laissez faire* liberalism. According to Jacques Rancière (2010: 36), the power to police—to partition the social world²—is always an ambivalent intercession: it separates and excludes while enabling participation and inclusion. The patrol enacts legitimacy and authority by carving geographies of risk and danger (Lianos and Douglas 2000; Marcuse 1997): certain social behaviors and environmental conditions are normalized, while others are marked as suspicious, deserving of inspection, intervention, or violence (Garland 2001; Harcourt 1998, 2001). By its very design as a spatial economy, the patrol endorses a distribution of both public safety benefits *and* state violence.

According to Thomas Nail (2015), the patrol achieves these outcomes through a number of mediating functions. First, patrol is *preventive* and *circulatory*—it deters crimes before they can take place by “oscillating its presence to and fro” (ibid.: 121). Patrols do not police crime *per se*, but rather the potential for criminal activity (Garland 2001). They not only stop and inspect (and search and seize), they also *make move*—they conduct traffic to foster optimum conditions for transportation while providing “dromological support” for their own efficient circulation (ibid.: 128). Additionally, the police patrol is *kinoptic* and *kinographic*. The patrol’s kinoptic function describes a surveillant mobility that *sees* and *is seen* in the same instant, watchfully making its presence known. Its circulation is designed “as a moving image of perfection and order”—or what early police theorist Edwin Chadwick called an “ambulating lighthouse” (quoted in Nail 2015: 122). The image of the lighthouse is especially evocative of the kinoptic function. The patrol illuminates and makes visible, while simultaneously making itself visible. And this kinoptic apparatus is enacted according to rational schema, that is, the kinographic function. Kinography is the inscription of movement and geography, the patrol’s mapping of urban space. Patrols establish the most efficient routes and routines for their double optic, often by mapping criminal potentialities—the most efficient routes of flight (Reeves and Packer 2013). This rationalization requires a sprawling cartographic and documentary assemblage to identify patterns of movement and behavior, not only of the citizenry and potential offenders, but of the patrol officers themselves: commanders and district captains synchronize patrol routes and circuits, creating a mesh of ubiquitous presence, spatially and temporally.

As new technologies of mobility and communication become embedded within the sociotechnical assemblage of patrol, they also remediate it (Byrne and Marx 2011; Manning 2008). Joshua Reeves and Jeremy Packer’s (2013) concept of “police media” denotes the suite of communications and transportation technologies that police use to amplify their presence in urban environments and to maintain a logistical monopoly over circulation. The police, according to Reeves and Packer (ibid.: 378), maintain their authority “not simply through a monopoly on the use of violence, but by creating a monopoly on the use of logistical media.” Police media are logistical media insofar as they “create new capacities for manipulating the time/space axis”; used in practice, police media conceive of the city “as a technological and infrastructural problem dealing with how best to organize and regulate flows of people, commodities, and risks” (ibid.: 359–60). This has included technologies from police callboxes, introduced in the 1880s to establish lines of communication between patrols and district headquarters, to the police cruiser and two-way radio introduced in the 1910s and further propelling the mobility and communicative reach of police patrols, and on to mobile onboard computers, CAD (computer-aided dispatch), and MDTs (mobile data terminals), which connect officers in the field with real-time dispatch and departmental records on suspects and specific neighborhoods. We could add to this list reformist managerial programs like CompStat (“computerized statistics”), which leverage crime data to increase accountability in police command, as well as various “intellectual technologies” that sociologist David Garland (2001) identifies with a shift from rehabilitative goals to risk management in applied criminology, including environmental theories of criminogenesis—that is, the “Broken Windows” theory (Kelling and Wilson 1982)—that motivate geographic profiling and “hot spot policing” (cf. Byrne and Marx 2011; Ferguson 2011; Harcourt 1998, 2001; Jefferson 2017; Manning 2008).

² Rancière’s (2010) notion of *police* exceeds the narrow definition of law enforcement, but nonetheless articulates the ambiguities of liberal consent policing under consideration (cf. Nail 2015: 116).

What Reeves and Packer do not account for with the concept of “police media,” however, are the ways that communications and transportation technologies also expose officers in the field to their superiors’ managerial scrutiny—the institutional impacts of the kinographic function. Police historian Willem De Lint (2000) argues that new patrol technologies’ novel logistical affordances—their capacities for manipulating the time/space axis—always involve a double-edged outcome. New technologies may improve patrol mobility and surveillance capabilities, but they also invite new forms of “supervisory co-presence,” linking officers in the field with their commanders or staff sergeants stationed at the precinct or station. “With each new technology, a fuller and more penetrating gaze has been envisioned, both of the police into the polity and of police supervision on police officer mobilization. These technologies structure the decision-making of individual officers on patrol to organizationally vetted formats” (De Lint 2000: 70). MDTs and onboard computers may enable officers to run license plate checks, but they also create the possibility for superiors to monitor officers’ activity and penalize idleness. On one hand, this increased supervision responds to external community or political demands that officers be kept in check and the public protected from police abuses of power. On the other hand, such oversight also responds to the managerial problem of officer autonomy—for instance, by ensuring that officers in the field are positioned to respond most efficiently to situations demanding intervention or emergency response (De Lint 2000; Sherman 2013).

Predictive policing embeds algorithmic decision-making systems within the sociotechnical and institutional assemblage of the patrol. As I argue in the next section, predictive policing functions similarly to other police media: it rationalizes the surveillant and visible presence of the patrol in urban spaces while fostering an algorithmic “supervisory co-presence” that more tightly integrates managerial imperatives about where and when officers are to patrol. Through this embedding, however, predictive policing also introduces new incongruities that provide the basis for the “predictive policing for reform” discourse while simultaneously undermining foundational claims to objectivity and neutrality.

Predictive Policing for Reform

Predictive policing is somewhat unique among other contemporary technologies associated with police reform (e.g., body-worn cameras; see Beutin 2017). Predictive algorithms are not themselves overtly panoptic in the same way that prominently placed surveillance cameras might be (McGrath 2004).³ Rather, predictive policing coopts the patrol’s established surveillance mechanisms (e.g., the beat, the uniform, the prominently placed marked vehicle) while algorithmically remediating its geographies: data analytics determines optimal locales and routes for patrol circulations.

Law enforcement agencies have turned to predictive policing in the wake of major incidents of police violence or federal court-monitored consent decrees following civil rights lawsuits.⁴ As Andrew Guthrie Ferguson (2017: 29) writes, the adoption of Big Data policing strategies “grew out of crisis and a need to turn the page on scandals that revealed systemic problems with policing tactics.” For example, the \$6 million predictive policing experiment underway in Chicago is part of broader reform efforts resulting from a DOJ investigation triggered by widespread protests that followed the release of a video depicting a police officer shooting teenager Laquan McDonald (McLaughlin 2017). Or consider the St. Louis County Police Department’s adoption of predictive policing system HunchLab within a year of the massive protests in Ferguson, MO, which likewise followed the shooting death of black teenager Michael Brown by officer Darren Wilson (Chammah 2016).⁵ “In response to demonstrated human bias,” Ferguson (2017: 26) writes,

³ In fact, police departments tend to be extremely opaque about their use of predictive policing (Brayne, Rosenblat, and boyd 2015). For example, the NYPD was sued after they failed to respond to Freedom of Information Act requests for information on its testing and deployment of predictive systems (e.g., Levinson-Waldman and Posey 2018).

⁴ The consent decree is the de facto mechanism for federal enforcement of police reforms in the US (see Ross and Parke 2009).

⁵ Although Darren Wilson was an officer of the Ferguson Municipal Police Department, the St. Louis County Police Department managed the aftermath of Michael Brown’s murder, including the ensuing protests, and was largely criticized for its handling of the events (Serrano and Pearce 2015).

“it is not surprising that the lure of objective-seeming, data-driven policing might be tempting.” Consequently, the cycle of police violence—protests and unrest, failure to indict or convict, further protests, and DOJ investigation—may now culminate with the adoption of Big Data law enforcement solutions as a means to rein in discriminatory or biased patrol practices.⁶

Like CompStat, which incentivizes reductions in crime rates through a data-driven, chain-of-command accountability structure, predictive policing appeals to police departments interested in extending accountability to officers in the field. As a mechanism of reform, predictive policing is imagined to tighten managerial supervision over patrols. In this sense, algorithmic patrol allocation systems comport with trends in applied criminology and “evidence-based policing,” for which geographic information is used to track and manage “what police were or were not doing in relation to the dynamic patterns of crime and public safety problems” (Sherman 2013: 379). As Ingrid Burrington (2015) writes, while predictive policing may do little to transform what police officers do while out on patrol, “it does have the potential to increase the power that police management has over cops on the street.” The question is whether and how this managerial control will be used—will it rein in police abuses or will it simply create new perverse incentives, like quotas for officers to meet? As Burrington points out, “recent events in Ferguson, Missouri, . . . demonstrate [that] the tendency toward micromanagement too often leads to more petty arrests in pursuit of revenue and quotas.”

How do the producers of predictive policing systems conceptualize and operationalize reform imperatives within algorithmic systems? To respond to this question requires grappling not only with the technical details of the algorithms, but also with how system designers imagine predictive policing to remediate the police patrol. How do system designers imagine end-users—the patrol officers—engaging with predictive information? And how will officers’ use of this information affect the data gathered in the field and fed back into the system? Building on Sarah Brayne, Alex Rosenblat, and danah boyd’s (2015: 5) contention that binaries like “intuition-driven” and “data-driven” policing are deceptive, I ask: How do producers of predictive systems imagine the mutual imbrication of “intuition” and “data,” of human and machinic decisions?

HunchLab

Findings are based on ethnographic research conducted with the HunchLab product team. From October 2015 to May 2016, I participated in business meetings, sat in on planning sessions, gave feedback on webinars, met with potential clients, traveled on site visits, and attended all-staff events, such as visiting speakers and brown bag lunches at Azavea, the company that produces HunchLab. Azavea is a Philadelphia-based software company that focuses on web-based geographic data applications. Although the company employs over fifty people, the HunchLab product team itself is small, consisting of Jeremy Heffner, HunchLab’s product manager and senior data scientist, and two to three full-time product specialists. The product team collaborates with a team of programmers who work concurrently on a number of Azavea projects and products. In addition to extensive interviews with Heffner, I also conducted interviews with the product specialists; programmers; Azavea founder and president, Robert Cheetham; and with criminologists that HunchLab has partnered with or hired as consultants.

HunchLab is unique within the police technology sector. Most of its competitors are brought to market by large corporations (e.g., IBM, Microsoft, Motorola, Hitachi, LexisNexis) or by smaller companies that are backed by venture capital, such as PredPol, or with CIA seed funding, such as Palantir (Robinson and Koepke 2016; Winston 2018). Azavea, by contrast, is not beholden to shareholders, investors, or covert government funding schemes. Instead, it adheres to a strict set of criteria for corporate social responsibility,

⁶ While other Big Data applications have been considered—for example, pilot studies to use algorithms to flag officers before they use excessive force (Arthur 2016; Ferguson 2017)—departments are far more likely to adopt predictive policing systems. For example, a 2016 survey of the fifty largest police departments in the US found that twenty agencies had already adopted predictive analytics for crime forecasting and eleven were actively considering it (as of August 2016) (Robinson and Koepke 2016).

environmental sustainability, and transparency, which have earned the firm a certification as a “social benefit” company or B Corp. HunchLab is thus also unique within the company, as the product’s alignment with these values is questionable. Although Heffner and Cheetham maintain that their mission is to use the tool to reduce harm and improve fairness and accountability in policing, at least one member of the programmer team has opted to not work on anything HunchLab-related, a request that the company honors.

Like other predictive policing tools, the first HunchLab prototype was built with grant money awarded by the US government, in this case through the National Science Foundation.⁷ After founding Azavea, Cheetham, a former crime analyst, collaborated with a team of Temple University criminologists on a grant-funded project to evaluate applications of predictive analytics in law enforcement (see Taylor, Ratcliffe, and Perenzin 2015). To manage the project, Cheetham hired Heffner, a mathematician and statistician, who introduced techniques from AI and machine learning to the new HunchLab prototype. With these techniques, the team avoided having to commit to any specific crime-forecasting approach in their design. With enough processing power, several prediction methods could be incorporated into a single, theory-agnostic meta-model. The machine learning algorithm parses combinations of forecasting methods to determine the most predictively accurate model based on signals in local data. These combined methods include an early warning system that Cheetham developed for the Philadelphia Police Department; “near repeat” analysis, a forecasting technique based on the spatial and temporal distribution of crimes (e.g., Townsley, Homel, and Chaseling 2000, 2003);⁸ and “risk terrain modeling” (RTM), a system modeling the proximity of crimes to key urban features (bars, churches, transportation hubs, and so on) to create spatial risk profiles (Caplan, Kennedy, and Miller 2011).

The current version of HunchLab is sold as a subscription service. Pricing is determined by jurisdictional population size, but starts at about \$50,000 for the first year and \$35,000 for subsequent years (Chammah 2016).⁹ Subscribers receive access to several algorithmic features, but the core algorithm is called “Predictive Missions.” This models the different criminological approaches and generates geospatial risk scores. HunchLab trains its algorithm on a client department’s crime data from the previous five years and on several non-crime-related data sets: census data, weather patterns, moon cycles, school schedules, holidays, and concerts and events calendars, all of which are mapped onto a grid of five-hundred-square-foot cells overlaying a client’s jurisdiction. A series of thousands of decision trees recursively partitions the data set based on crime outcomes in each grid cell. If a crime occurred in a cell, then the regressions determine which variables influenced the occasion of that crime and to what extent; variables are then weighted accordingly, tailoring the model to the clients’ data.

The result is a hyper-localized and hyper-sensitive crime forecasting algorithm. Because it adjusts weights according to local crime data, models for the same crime will differ across jurisdictions. And because subscribing departments’ data are updated daily, the weighting for each crime-type may also change over time. For example, *location* could be the most predictive factor for *theft from automobiles* in Detroit, but in Philadelphia it might be *time of day*; both models would be automatically adjusted if the crime patterns changed. When the modeling is evaluated against ground truth data (where crimes actually occurred relative to the predictions), the results indicate high levels of accuracy—sometimes as high as ninety-two percent to ninety-seven percent, depending on the crime type.¹⁰

Indeterminacies

HunchLab boasts of these extraordinary performance rates to potential clients. If crime data are publicly available for a jurisdiction, a product specialist can create a mock-up model and present it during sales

⁷ <https://www.nij.gov/topics/law-enforcement/strategies/predictive-policing/Pages/welcome.aspx>

⁸ Competing system PredPol uses a parsimonious forecasting technique that approximates the “near repeat” method (Mohler et al. 2015; see Brayne, Rosenblat, and boyd 2015).

⁹ For a pricing perspective, HunchLab’s competitor PredPol charges about \$200,000 per year.

¹⁰ Author’s fieldnotes, St. Louis County, MO, December 2, 2015. HunchLab tested the accuracy of its St. Louis County Police Department model to be included in a presentation to departmental command staff during a site visit.

pitches. They simply withhold recent crime events from the training data and then juxtapose those with the model's predictions to illustrate where crimes were predicted relative to where they actually took place. But HunchLab is also quick to acknowledge that predictive accuracy can only be measured *prior* to a department's adoption and implementation of the system. Once predictions are put into use and officers start patrolling predicted crime locations, the algorithm's performance can no longer be evaluated: the ground truth data cease to be a controlled sample. When patrols are directed by predictive information, officers' visible presence in a predicted crime location affects the conditions being modeled, undermining any claims to the model's representational fidelity and predictive performance. By the same token, the officer patrolling the grid cell produces new data. If he or she makes an arrest in the area, for example, then that arrest will be incorporated into the model the next day. As in the Heisenberg uncertainty principle and posthumanist theories of performativity, HunchLab imagines the double optics of the patrol's kinoptic and kinographic functions—the observable and observing position of officers in the field—as introducing indeterminacy at the same time that it produces new data (Barad 2003).

HunchLab product manager Heffner conceptualizes these performative effects as a paradox, animated by competing probabilities: *detection* and *deterrence*. Detection refers to the increased likelihood that an officer observes crimes taking place by dint of his or her being in the predicted location, while deterrence refers to the increased likelihood that his or her visible presence will prevent crime from taking place there.¹¹ These countervailing forces defy accuracy measurements, illustrating how the double optics of patrol interact with, and potentially confound, predictive algorithms.

The result is an indeterminacy that is fundamental to prediction in general (Mackenzie 2015) but which has largely been unacknowledged in debates about big data policing. Engaging with indeterminacy betrays the extent to which predictive analytics is inaccurately theorized through an invisible and disembodied measurement apparatus (Haraway 1988). As soon as an officer steps foot into a predicted grid cell, he or she performatively shapes what takes place there. The task thus becomes capturing, measuring, and analyzing those performative effects. As when Google tracks users' responses to changes in the algorithm or design, HunchLab seeks to “fold the performativity of models back into the modeling process” (Mackenzie 2015: 443)—to statistically represent the performative effects of prediction within the predictive apparatus.

In predictive policing, “folding” the performativity back into the modeling can work if the desired outcomes are observable behaviors or actions. If clients want predictions to lead to higher arrest rates, then this outcome can be modeled because it is observable in the data. But if the desired outcome is *prevention*—as in Peel's (1829) final principle that “[t]he test of police efficiency is the absence of crime and disorder, not the visible evidence of police action in dealing with it”—then system managers are faced with a paradox: an event deterred is by definition unobservable (as in the truism that you can't prove a negative). Of course, prevention rates can be inferred by comparison between a treatment group and a control group (e.g., Hunt, Saunders, and Hollywood 2014; Ratcliffe, Taylor, and Askey 2017). But even this will always be an imperfect estimation, as no two jurisdictions, beats, or patrol shifts are identical. Further, maintaining a control group necessarily means only partially implementing predictions, a prospect that departmental clients may not be interested in, given that they are paying handsomely for the technology.

The result is an indeterminacy that cannot be avoided once predictive analytics are put into practice. The HunchLab team appears to be alone among police technology vendors in acknowledging this. And, crucially, such indeterminacy is central to HunchLab's claims that algorithmic mediation can serve police reform efforts. Indeterminacies provide an ambiguous opportunity for intervention. They open a space in which predicted outcomes can be thwarted and deterred, at the same time that they confound evaluative metrics like predictive accuracy: that which is statistically unobservable may actually be the most desirable.

¹¹ Author's fieldnotes, conversation with Jeremy Heffner, HunchLab offices, October 1, 2015.

The problem thus becomes how to incorporate predictions' performativity back into the modeling process to exploit this confounding (Mackenzie 2015). "Prescriptive analysis" (rather than *predictive*) is the name that HunchLab gives to its attempts to measure predictions' indeterminate performative effects and incorporate them within the modeling process. For example, in a webinar titled "Beyond the Box: Towards Prescriptive Analysis in Policing," Heffner and former product specialist Chip Koziara intimate that indeterminacies can be exploited to mitigate discriminatory patrol practices.¹² In business analytics, *prescription* occupies a more complex register than *prediction*: rather than merely predicting an outcome based on past events, prescription promises to account for how acting on predictions affects the conditions modeled, the goal being to optimize trade-offs and ensure desirable outcomes (Ransbotham, Kiron, and Prentice 2015). HunchLab seizes on the notion of prescription to position itself as disrupting entrenched patterns in law enforcement—"reducing harm associated with over-policing, and . . . helping officers find the best tactical solutions to improve their communities" (HunchLab, n.d.).

For example, working with Temple University criminologist Jerry Ratcliffe, HunchLab is developing a modeling system called HarmStat. HarmStat is based on the progressive notion that heightened police presence in low-income and minority neighborhoods is not perceived as a form of "protection" but as a source of harm that can be quantified and evaluated relative to crime harms through a cost-benefit analysis.¹³ The analysis is based on an estimate that the team calls "predictive efficacy," which describes the extent to which the most harmful crimes are able to be predicted and deterred. Violent crimes like assault and homicide may be more difficult to predict than property crimes like burglary or larceny (e.g., Ratcliffe, Taylor, and Askey 2017), but "communities" may find it more valuable to prevent violent crimes because they are more harmful; efforts to thwart more easily predictable crimes may have greater success rates, but the payoff from deterring less predictable and more harmful crimes may be higher.

This logic may make sense intuitively, but it glosses over the fact that indeterminacy remains baked into the very essence of HarmStat's modeling. Predictive efficacy is essentially an estimation of deterrence, which is unobservable. Further, HarmStat assumes an immeasurable precision and correctness to crime predictions, which cannot be validated after implementation due to the performative effects of the deterrence-detection tension detailed above. After predictions have been used to allocate patrols, neither HunchLab nor its clients can point to any evidence that the algorithm continues to accurately predict crimes. This is especially problematic given that government-funded research has failed to uncover statistical evidence that patrol predictions result in measurable decreases in crime (Hunt, Saunders, and Hollywood 2014; Saunders, Hunt, and Hollywood 2016). "Predictive efficacy," in other words, is conceptually fallacious. To believe that increased police patrols can serve as a source of public safety rather than harm in a cost-benefit sense, HarmStat users would need either to ignore or deny the fundamental indeterminacy of its own baseline metrics.

Trade-offs

While the Predictive Missions algorithm assigns risk scores to each grid cell in the jurisdiction, the "Allocation Engine" algorithm sorts through the risk-scored cells to select which should be patrolled during a given shift. The Allocation Engine's defining feature is its selective, strategic, and explicit insertion of *randomization* into the prediction process. Rather than directing patrol to the grid cells with the highest risk—how HunchLab was originally designed and how its competitors continue to operate—the algorithm directs officers instead to the second, third, fourth, or fifth riskiest places according to a probabilistic selection process based on randomization.¹⁴

¹² HunchLab webinar, "Beyond the Box: Towards Prescriptive Analysis in Policing." Available at <https://www.youtube.com/watch?v=NCXFDfQsYBE>

¹³ Author's fieldnotes, HunchLab offices, May 5, 2016.

¹⁴ The Allocation Engine involves a set of mathematical rules dictating the selection of mission areas. These rules can be tweaked by clients to prioritize strategic goals. Risk forecasts are transformed into z-scores, which are then used to filter out cells below a certain threshold, eliminating low-risk cells from being allocated as a Predictive Mission. Within

The benefits of randomization are matched by an acknowledged compromise in the perceived accuracy of deployments. That is, not selecting the highest risk cell every time means that were one to compare the Allocation Engine's grid cell outputs against (pre-deployment) ground truth, the system would be less accurate in predicting an actually occurring crime event. This makes sense given the indeterminacies discussed above. "If we're just trying to maximize our predictive accuracy," Heffner explained in a webinar introducing the Allocation Engine, "then, absolutely, selecting that [highest risk] cell every time would be what we'd do. But that's not the case here. You're going to act on [these predictions] and so you're going to start to skew things, displace crime, and so forth."¹⁵

HunchLab invoked the benefits of randomization strategically, allowing it to appeal to both potential clients (police department crime analysts or commanders) and to civil rights activists critical of predictive policing.¹⁶

When addressing potential clients, HunchLab referred to randomization as having a number of rationalizing benefits: it makes patrols less predictable to potential offenders and thus avoids crime displacement, and it keeps field officers invested in the usefulness of the system. The potential that location-based patrol strategies might displace crime is a general concern in experimental criminology (e.g., Bowers et al. 2011; Weisburd and Braga 2006) and has been especially problematic for "hot spot policing" (Braga 2005; COPS 2013). Hot spot policing is a rudimentary application of geospatial analytics to patrols, which are allocated based on thirty days' worth of historical crime data at the ZIP or neighborhood level. Crime displacement refers to the tendency for hot spots to become predictable to potential offenders, who simply move criminal activities into new areas. In a sense, predictive policing is simply a more granular version of hot spot policing and as such has raised new concerns about crime displacement. Randomization, HunchLab argues, mixes up which grid cells are selected for patrol during each shift, avoiding saturation and patrol predictability while simultaneously providing coverage for new areas.

Randomization is also invoked to tackle the rampant problem of officer boredom. Clients and potential clients often raised this issue during meetings with the product team. As two former police officers put it while consulting with HunchLab, "In Seattle, they had to shelve [PredPol] within a handful of months because the officers had no idea what to do and they were saying, 'I'm bored.'"¹⁷ If predictions send officers to the same high-risk grid cells for each shift, this may result in doubt and undermine officer buy-in. For example, a survey of the Burbank Police Department found that prediction-based deployments have resulted in officer malaise (Tchekmedyian 2016), exacerbating concerns about police officer deskilling. Here, predictive policing is but one among many "augmentation" and decision-support tools that experts like criminologists Dave Allen argue may result in a move away from experience and skill development and toward "management by remote control or blindly following information drawn from systems" (Allen, quoted in Wakefield 2013).

Crucially, the problem of deskilling presents an inverse to the solution of rationalized patrol allocations. If the "predictive policing for reform" discourse maintains that discriminatory practices are an aggregation of individual officers' misguided discretion and autonomy, then reining that autonomy in necessarily entails a "management by remote control"—or what De Lint called "supervisory co-presence." Randomization promises to do so in a way that mitigates rather than exacerbates officer boredom and malaise. Randomization, the HunchLab team argues, places patrol officers in areas that they may not have thought

the filtered collection of cells, weights are then used to differentiate between medium and high risk locations, and randomization is introduced in the selection within this narrowed set.

¹⁵ "Beyond the Box" webinar. See footnote 11.

¹⁶ Heffner and product team members often attend conferences to represent the police technology sector, often portraying HunchLab as an ethical alternative to other predictive policing products, for example, the 2015 Data and Civil Rights conference, sponsored by technology and social justice nonprofit Data & Society. <http://www.datacivilrights.org/>.

¹⁷ Author's fieldnotes, HunchLab offices, December 18, 2015.

to patrol previously because of the biases of their own on-the-job experience. Serendipitously, this also responds to the problem of boredom. The crime analyst for HunchLab client Greensboro Police Department has argued publicly that his department's officers enjoy the element of surprise that randomized predictions introduce.¹⁸ There is a craft to maintaining officer buy-in, Heffner explained to me—a delicate balancing act to ensure that officers believe the algorithm “knows what it’s doing,” sending them to known high crime-risk locations, and that it does “something new,” sending them to unexpected locales, and randomization is the technique of choice to accomplish the balance.¹⁹

When HunchLab team members addressed civil rights advocates critical of predictive policing, randomization was framed as a mechanism to *disrupt* the patrol's entrenched geographic biases. Speaking at a policing and civil rights conference at New York University's School of Law, Heffner explained,

If you're a police department and you're not using an algorithm, you're probably using a hot spot map . . . [which] has probably not changed very much for a long time. So what we do in HunchLab is, we sometimes *don't* send [patrols] to the highest risk places, because then we can see what happens when we don't send them there and we send them to a lower risk place . . . It's a bit of a randomization based upon the analysis to help us gain more insight into what it would look like when you don't saturate an area with police. Because maybe we don't have that in the training data, and we need to gain that knowledge.²⁰

The compromise between predictive accuracy—sending officers to the highest risk grid cells—and gaining knowledge by sending officers to new places echoes well-known trade-offs in computer science, between fairness and accuracy (e.g., Friedler et al. 2018; Zafar et al. 2017) and exploitation and exploration (Berger-Tal et al. 2014; Slivkins 2017). The fairness–accuracy trade-off refers to outcomes when algorithms operate on data about people. Fairness-aware algorithms seek to ensure that outcomes do not disproportionately impact members of a protected class (class, race, ethnicity, sexuality, gender, religion, and so on), but in doing so compromise predictive accuracy relative to the ground truth data (Friedler et al. 2018). In the exploitation–exploration trade-off, compromises must be made between obtaining new knowledge or information (exploration) and using the knowledge or information that one already has to improve performance (exploitation) (Berger-Tal et al. 2014). For HunchLab, randomization functions to introduce fairness through exploration, in the sense that randomizing allocations sends officers to areas that are underrepresented in the crime data because of uneven patterns of policing (that is, wealthier and whiter areas).

Though a technical intervention, randomization provides a flexible mechanism with which HunchLab links its rationalization of patrols with ethical concerns about biased geographic data and discriminatory patrol strategies. Although such interventions have largely been ignored in debates about algorithmic predictions in law enforcement, it is not enough to simply point to randomization as a universal corrective. We need to grapple with the trade-offs that randomization introduces in order to understand who is making such decisions, why, and with what consequences.

Unfalsifiability

Beyond the fairness–accuracy and exploration–exploitation trade-offs, compromises in scientific rigor are justified for the sake of pragmatic experimentalism. Beyond Predictive Missions and the Allocation Engine, HunchLab subscribers can also access a feature called “Advisor.” In line with the vision for “prescriptive policing,” Advisor provides a way for clients to experiment with different patrol tactics. On the surface,

¹⁸ HunchLab webinar, “HunchLab Predictive Missions at Greensboro PD: ‘Tell me what I don’t know!’” Available at <https://www.youtube.com/watch?v=E-QdYqZzQhY>

¹⁹ Author's fieldnotes, HunchLab offices, March 26, 2015.

²⁰ Heffner, presentation at the “Policing and Accountability in the Digital Age” symposium at the New York University School of Law, September 15, 2016 (emphasis reflects speech). <https://www.youtube.com/watch?v=M1saeirVqqU>

Advisor appears to merely automate the methodology of a randomized control trial (RCT): departments test different tactical responses to crime patterns and evaluate the outcomes relative to a control group. Yet, given the indeterminacies and trade-offs detailed above, Advisor also abandons some core tenets of the RCT, leading to a mode of experimentation that is ultimately untethered from ground truth.

Advisor consists of three distinct initiative types: Field Test, Experiment, and Adaptive Tactics. With Field Test, clients evaluate the effectiveness of different tactics in response to a specified crime-type. For example, following a wave of *home burglaries*, a department could use Field Test to study tactics like *high visibility patrol* and *canvassing homes and businesses* for reducing crime incidence. HunchLab offers suggestions for potential tactics to be tested (for example, “writ[ing] reports while parked in patrol cars at high risk locations”²¹), but these are also customizable fields that district commanders can update according to their own imperatives. After tactics are delineated, Field Test monitors the rate of *home burglaries* while a tactic is implemented and then compares the outcome to “what likely would have happened had you not been doing the field test.”²² The “what likely would have happened,” in turn, is determined by algorithmic prediction—but again, this can no longer be evaluated.

The second initiative type, Experiment, is similar but expands the process to the entire jurisdiction, randomly assigning beats, districts, or precincts as control or treatment groups (essentially replicating the RCT methodology). HunchLab maintains that Experiment promises certain advantages over traditional RCTs. For example, because it is designed for internal tests, departments may be less concerned with achieving proper statistical thresholds for determining significance. Experimentation can be implemented rapidly and, with software interfaces designed to automate methodology, the barrier to entry for officers and analysts without advanced degrees in statistics or experimental criminology is lowered.

The third of Advisor’s initiatives, Adaptive Tactics, is somewhat different, as it is not confined to a fixed experimental timeframe. Adaptive Tactics involves ongoing data collection on the effects that a selection of tactical responses will have on predicted crime rates. Like Field Test, Adaptive Tactics takes a list of tactical recommendations developed for a specific crime problem. Every time that crime-type is predicted, Adaptive Tactics makes a recommendation from the list and records its execution in relation to the risk profile for the grid cell. This begins as a randomized assignment, with zero confidence in the recommendations, but over time promises to accumulate enough data to make recommendations with higher levels of inferential certainty. As with the exploitation–exploration trade-off discussed above, Adaptive Tactics operates by a balancing act, between resources expended to acquire new information about various tactics’ effectiveness and simply exploiting what is already known to be effective. Because the fields are customizable, however, there are no guarantees that discriminatory tactics believed to be effective will not be imputed (for example, a department could easily impute unwarranted stops, searches, and seizures, “stop and frisk,” as a tactical recommendation).

The emphasis on pragmatic experimentation echoes calls from criminologists and law enforcement intellectuals for police departments to adopt a more iterative and experimental approach to patrol. Between 2014 and 2016, for example, the National Institute of Justice (NIJ) ran the Randomized Control Trial Challenge, promising grants of \$100,000 to five police departments to conduct research on police managerial strategies.²³ Jim Bueermann, president of the Police Foundation, predicted that by 2022, every police department would have a resident criminologist to test strategic efficacy (McCullough and Spence 2012). Similarly, criminologist Lawrence Sherman (2013) forecast that command staffs will be regularly deploying technologies to test patrol efficiencies by 2025. As HunchLab’s Heffner maintains, however, the kind of experimentation that criminologists call for can be burdensome and expensive for local departments,

²¹ HunchLab webinar, “HunchLab Advisor: Know What Works.” Available at <https://www.youtube.com/watch?v=hHDJfHPYTsU>

²² “HunchLab Advisor” webinar. See footnote 20.

²³ According to former NIJ executive George Ridgeway, the Challenge was canceled as of January 2016 because no departments applied for funding (Gregory Ridgeway, personal communication with the author, January 17, 2018).

especially absent a grant like the NIJ's.²⁴ Institutional pressures can also be prohibitive. In the era of intra-departmental accountability demands through programs like CompStat, pressures on district commanders to demonstrate strategic effectiveness through crime reductions “can lead to a sort of risk aversion,” with commanders becoming “less likely to experiment with things . . . because if we can just keep things generally as they are, they will likely turn out the same way at the next CompStat meeting—and so that’s a safe move.”²⁵ Advisor is marketed as capable of overcoming these barriers—through automation, easy-to-use interface design, and lower thresholds for significance.

The trouble is that Advisor is also built upon the indeterminacies discussed above. When Field Test or Adaptive Tactics evaluates the effectiveness of a specific tactic, they compare the crime outcome in that assigned cell against predictions of what likely would have occurred there based on the Predictive Missions algorithm. But given the performative effects of patrol on the crime, this scenario might best be understood as a fraction without a denominator, an RCT without a control. Inferences about a tactic’s effectiveness are admittedly unfalsifiable in the traditional sense of experimental and quasi-experimental designs. Scientific rigor is sacrificed for good-enough estimations. This could mean new, less discriminatory tactics, but it could also mean an optimized version of the same old practices; faith is simply placed in police departments’ willingness to test out less harmful tactics—canvassing rather than stopping and frisking, for instance. Algorithmic “rewards” are determined, and these inform the system how to sort outcomes as positive or negative. But if the only data legible to the system are arrests made based on predictive patrol (rather than the unobservable prevention or deterrence of crimes), then Advisor may simply learn how to reproduce the patterns most in need of change (Jefferson 2017; Robinson and Koepke 2016).

A Common Good or a Distribution of Harm?

Predictive policing has always been discursively supported by claims that it enhances control, certainty, and exactitude. As Los Angeles Chief of Police Charlie Beck puts it, algorithmic allocations offer “the ability to anticipate or predict crime provid[ing] unique opportunities to prevent, deter, thwart, mitigate, and respond to crime more effectively, ultimately changing public safety outcomes and the associated quality of life for many communities” (Beck and McCue 2009). In practice, however, it is indeterminacy, uncertainty, and a general “fudginess” that open a space for intervention. Patrol rationalizations operate on the scaffolding of “irrational” assumptions such as “predictive efficacy” or on unfalsifiable inferences as a criterion for evaluation. Most optimistically, these incongruities untether the predictive patrol from routinized, discriminatory practices and patterns in policing. As with Wendy Hui Kyong Chun’s (2018) recent work on the politics of proxies (cf. Wilk 2017), algorithmic prediction becomes an “ambivalent pharmacon”—a mode of intervening in patrol geographies through data proxies that challenge established epistemological premises for what constitutes evidence. Although prediction is nominally about representing and anticipating future events, it works by capturing distributions of the *present*. Untethering these from ground truth creates an opportunity to intervene, to divert futures from merely reproducing the uneven and inequitable patterns that shape policing today.

Critiques of predictive policing have largely ignored the fact that existing best practices, such as hot spot policing, are themselves a crude algorithmic remediation of patrol geographies²⁶ that facilitate the over-policing of poor and minority communities and expose them to police abuses. HunchLab team members are not wrong to critique hot spot policing; it is a blunt preventive instrument that allows for biased patrols, leaving officers plenty of space to abuse their power through an exercise of discretion. As Dubber and Valverde (2006) note, the police’s preventive power has always entailed a discretionary element, and that discretion reflects a form of localized sovereignty and power. “Predictive policing for reform” seeks to disrupt these entrenched patterns and introduce accountability for that discretion. The question that needs

²⁴ “HunchLab Advisor” webinar. See footnote 20.

²⁵ “HunchLab Advisor” webinar. See footnote 20.

²⁶ An exception is Brayne, Rosenblat, and boyd’s (2015) primer on predictive policing, which raises questions about the relationship between algorithmic predictions and current best practices in police patrol.

to be asked of algorithmic interventions, however, is in what ways that discretion is distributed between “data-driven” and “intuition-driven” patrols—between machinic and human controls—and how these distributions can be evaluated without recourse to unfalsifiable claims.

By the same token, indeterminacy’s destabilization of statistical representation should raise concerns—not simply about scientific validity or the merit of algorithmic prediction but about ethical validity and merit. As critical geographer Brian Jordan Jefferson (2017: 2) argues in his study of the Chicago Police Department’s use of geospatial analytics, the predictive data structure “ensures that negatively racialized fractions of surplus labor and the places they inhabit are only representable to state authorities and the public *as objects of policing and punishment*” (emphasis added). This narrow legibility owes to the fact that deterrence is central to the mediations of the police patrol but not representable by measurement: the sociotechnical and institutional assemblage of the police patrol is organized around “positive” (observable, quantifiable, and optimizable) outcomes—police use of force and punishment—to effect public safety benefits. Letting-alone, letting-be, respecting—these are statistically inscrutable.

Is it possible to expand this legibility and truly disrupt entrenched police abuses and violence? What would it look like to view HunchLab’s interventionary trade-offs—between fairness and accuracy, exploration and exploitation—with an eye toward remediating disinvestment in black communities, for example? Imagining such a scenario is difficult because the ambiguities, indeterminacies, and trade-offs that plague predictive policing are innate and defining features of the police patrol itself, not mere effects of algorithmic agency (Introna 2016).

Ultimately, “predictive policing for reform” is incapable of resolving two fundamentally incommensurate but concurrent functions of the police patrol. On one hand is a view of police patrols as distributing public safety as a common good—an idea that traces back to early modern theorists of the state (e.g., Locke [1698] 1988). On the other is the view from marginalized communities, who experience the patrol as an enactment of uneven geographies of legitimacy and authority, risk and danger, harm and abuse (e.g., Marcuse 1997; Merrill and Hoffman 2015). On the first view, patrols act to safeguard the public from criminality; accordingly, a more equitable distribution of protections can be optimized for. In the latter view, whole communities are criminalized through location-based patrols; police operate as an occupying force in neighborhoods isolated through histories of unjust policies, disinvestment, and urban renewal programs that isolate and racialize surplus populations (Jefferson 2017). For these communities, “optimization” would simply mean more “effective” sources and distributions of harm.

That these two functions are irreconcilable is not the fault of the algorithm per se, nor can an algorithm offer any sort of meaningful resolution. Improving public safety benefits for all communities—enacting more equitable geographies of risk and protection—will require grappling with, reorganizing, or even potentially dismantling the entire sociotechnical and institutional apparatus of the police patrol itself.

References

- ACLU (American Civil Liberties Union). 2016. *Predictive Policing Today: A Shared Statement of Civil Rights Concerns*. <https://www.aclu.org/other/statement-concern-about-predictive-policing-aclu-and-16-civil-rights-privacy-racial-justice> [accessed July 21, 2018].
- Anderson, Ben. 2010. Preemption, Precaution, Preparedness: Anticipatory Action and Future Geographies. *Progress in Human Geography* 34 (6): 777–98.
- Angwin, Julia, Jeff Larson, Surya Mattu, and Lauren Kirchner. 2016. Machine Bias. *ProPublica*, May 23, 2016. <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing> [accessed July 21, 2018].
- Arthur, Rob. 2016. We Now Have Algorithms to Predict Police Misconduct. *FiveThirtyEight*, March 9, 2016. <https://fivethirtyeight.com/features/we-now-have-algorithms-to-predict-police-misconduct/> [accessed July 21, 2018].
- Barad, Karen. 2003. Posthumanist Performativity: Toward an Understanding of How Matter Comes to Matter. *Signs* 28 (3): 801–31.
- Beck, Charlie, and Colleen McCue. 2009. Predictive Policing: What Can We Learn from Wal-Mart and Amazon about Fighting Crime in a Recession? *Police Chief* 76 (11). <http://acmcs373ethics.weebly.com/uploads/2/9/6/2/29626713/police-chief-magazine.pdf> [accessed July 21, 2018].

- Berger-Tal, Oded, Jonathan Nathan, Ehud Meron, and David Saltz. 2014. The Exploration-Exploitation Dilemma: A Multidisciplinary Framework. *PLOS ONE* 9 (4): e95693.
- Beutin, Lyndsey P. 2017. Racialization as a Way of Seeing: The Limits of Counter-Surveillance and Police Reform. *Surveillance & Society* 15 (1): 5–20.
- Bowers, Kate J., Shane D. Johnson, Rob T. Guerette, Lucia Summers, and Suzanne Poynton. 2011. Spatial Displacement and Diffusion of Benefits Among Geographically Focused Policing Initiatives: A Meta-Analytical Review. *Journal of Experimental Criminology* 7 (4): 347–74.
- Braga, Anthony A. 2005. Hot Spots Policing and Crime Prevention: A Systematic Review of Randomized Controlled Trials. *Journal of Experimental Criminology* 1 (3): 317–42.
- Bratton, William. 2014. Manhattan Institute: Bill Bratton on the Future of Policing. *YouTube*. https://www.youtube.com/watch?v=X_xjqOp7y0w [accessed July 21, 2018].
- Brayne, Sarah, Alex Rosenblat, and danah boyd. 2015. Predictive Policing. Workshop presented at annual conference of Data and Civil Rights: A New Era of Policing and Justice, Washington, DC, October 27, 2015. <https://datasociety.net/output/data-civil-rights-predictive-policing/> [accessed July 21, 2018].
- Burrington, Ingrid. 2015. What Amazon Taught the Cops. *Nation*, May 27, 2015. <https://www.thenation.com/article/what-amazon-taught-cops/> [accessed July 21, 2018].
- Byrne, James, and Gary Marx. 2011. Technological Innovations in Crime Prevention and Policing: A Review of the Research on Implementation and Impact. *Journal of Police Studies* 3 (20): 17–40.
- Caplan, Joel M., Leslie W. Kennedy, and Joel Miller. 2011. Risk Terrain Modeling: Brokering Criminological Theory and GIS Methods for Crime Forecasting. *Justice Quarterly* 28 (2): 360–81.
- Chammah, Maurice. 2016. Policing the Future. *Verge*, February 3, 2016. <https://www.theverge.com/2016/2/3/10895804/st-louis-police-hunchlab-predictive-policing-marshall-project> [accessed July 21, 2018].
- Chun, Wendy Hui Kyong. 2018. Ways of Knowing (Cities) Networks. Keynote lecture at Ways of Knowing Cities conference, Columbia University, New York, February 9, 2018. <http://c4sr.columbia.edu/knowning-cities/schedule.html> [accessed July 21, 2018].
- Colquhoun, Patrick. 1800. *A Treatise on the Commerce and Police of the River Thames*. London: J. Mawman. <http://archive.org/details/atreatiseoncomm00colqgoog> [accessed July 21, 2018].
- COPS (Community Oriented Policing Services). 2013. The Importance of Legitimacy in Hot Spots Policing. *Community Policing Dispatch* 6 (9): September, 2013. https://cops.usdoj.gov/html/dispatch/09-2013/the_importance_of_legitimacy_in_hot_spots_policing.asp [accessed July 21, 2018].
- Corbett-Davies, Sam, Sharad Goel, and Sandra González-Bailón. 2017. Even Imperfect Algorithms Can Improve the Criminal Justice System. *New York Times*, December 20, 2017. <https://www.nytimes.com/2017/12/20/upshot/algorithms-bail-criminal-justice-system.html> [accessed July 21, 2018].
- Crawford, Kate. 2016. Artificial Intelligence's White Guy Problem. *New York Times*, June 26, 2016. <https://www.nytimes.com/2016/06/26/opinion/sunday/artificial-intelligences-white-guy-problem.html> [accessed July 21, 2018].
- De Lint, Willem. 2000. Autonomy, Regulation and the Police Beat. *Social & Legal Studies* 9 (1): 55–83.
- Dubber, Markus D., and Mariana Valverde, eds. 2006. *The New Police Science: The Police Power in Domestic and International Governance*. Palo Alto, CA: Stanford University Press.
- Ensign, Danielle, Sorelle A. Friedler, Scott Neville, Carlos Scheidegger, and Suresh Venkatasubramanian. 2017. Runaway Feedback Loops in Predictive Policing. Paper prepared for the first conference on Fairness, Accountability, and Transparency in Machine Learning, New York University, New York, February 2018. <http://arxiv.org/abs/1706.09847> [accessed July 21, 2018].
- Ferguson, Andrew Guthrie. 2011. Crime Mapping and the Fourth Amendment: Redrawing 'High-Crime Areas.' *Hastings Law Journal* 63 (1): 179–232.
- Ferguson, Andrew Guthrie. 2017. *The Rise of Big Data Policing: Surveillance, Race, and the Future of Law Enforcement*. New York: New York University Press.
- Foucault, Michel. 1980. *Power/Knowledge: Selected Interviews and Other Writings*. Edited by Colin Gordon. New York: Vintage.
- Foucault, Michel. 2007. *Security, Territory, Population: Lectures at the College de France, 1977-78*. Translated by Graham Burchell. New York: Palgrave Macmillan.
- Friedler, Sorelle A., Carlos Scheidegger, Suresh Venkatasubramanian, Sonam Choudhary, Evan P. Hamilton, and Derek Roth. 2018. A Comparative Study of Fairness-Enhancing Interventions in Machine Learning. *FAT*19 Proceedings of the Conference on Fairness, Accountability and Transparency*, 329-338. Atlanta, GA. January 29-31, 2018. <http://arxiv.org/abs/1802.04422> [accessed July 21, 2018].
- Garland, David. 2001. *The Culture of Control: Crime and Social Order in Contemporary Society*. Chicago: University of Chicago Press.
- Haraway, Donna. 1988. Situated Knowledges: The Science Question in Feminism and the Privilege of Partial Perspective. *Feminist Studies* 14 (3): 575.
- Harcourt, Bernard E. 1998. Reflecting on the Subject: A Critique of the Social Influence Conception of Deterrence, the Broken Windows Theory, and Order-Maintenance Policing New York Style. *Michigan Law Review* 97 (2): 291.
- Harcourt, Bernard E. 2001. *Illusion of Order: The False Promise of Broken Windows Policing*. Cambridge, MA: Harvard University Press.
- Hartman, Saidiya V. 1997. *Scenes of Subjection: Terror, Slavery, and Self-Making in Nineteenth-Century America*. New York: Oxford University Press.

- HunchLab. n.d. Resources (webpage). <https://www.hunchlab.com/resources/> [accessed July 21, 2018].
- Hunt, Priscillia, Jessica Saunders, and John S. Hollywood. 2014. *Evaluation of the Shreveport Predictive Policing Experiment*. Santa Monica, CA: RAND. https://www.rand.org/pubs/research_reports/RR531.html [accessed July 21, 2018].
- Introna, Lucas D. 2016. Algorithms, Governance, and Governmentality: On Governing Academic Writing. *Science, Technology, & Human Values* 41 (1): 17–49.
- Jefferson, Brian Jordan. 2017. Digitize and Punish: Computerized Crime Mapping and Racialized Carceral Power in Chicago. *Environment and Planning D: Society and Space* 35 (5): 775–96.
- Kelling, George L., and James Q. Wilson. 1982. Broken Windows: The Police and Neighborhood Safety. *Atlantic Monthly*, March 1, 1982. <https://www.theatlantic.com/magazine/archive/1982/03/broken-windows/304465/> [accessed July 21, 2018].
- Kember, Sarah, and Joanna Zylińska. 2012. *Life After New Media: Mediation as a Vital Process*. Cambridge, MA: MIT Press.
- Levinson-Waldman, Rachel, and Erica Posey. 2018. Court Rejects NYPD Attempts to Shield Predictive Policing from Disclosure. *Brennan Center for Justice* (blog). <https://www.brennancenter.org/blog/court-rejects-nypd-attempts-shield-predictive-policing-disclosure> [accessed July 21, 2018].
- Lianos, Michaelis, and Mary Douglas. 2000. Dangerization and the End of Deviance: The Institutional Environment. *British Journal of Criminology* 40 (2): 261–78.
- Locke, John. (1698) 1988. *Locke: Two Treatises of Government*. Edited by Peter Laslett. 3rd edition. New York: Cambridge University Press.
- Lohr, Steve. 2018. Facial Recognition Is Accurate, If You're a White Guy. *New York Times*, February 9, 2018. <https://www.nytimes.com/2018/02/09/technology/facial-recognition-race-artificial-intelligence.html> [accessed July 21, 2018].
- Lum, Kristian, and William Isaac. 2016. To Predict and Serve? *Significance* 13 (5): 14–19.
- Mackenzie, Adrian. 2015. The Production of Prediction: What Does Machine Learning Want? *European Journal of Cultural Studies* 18 (4–5): 429–45.
- Manning, Peter K. 2008. *The Technology of Policing: Crime Mapping, Information Technology, and the Rationality of Crime Control*. New York: New York University Press.
- Mantello, Peter. 2016. The Machine That Ate Bad People: The Ontopolitics of the Precrime Assemblage. *Big Data & Society* 3 (2): 1–11.
- Marcuse, Peter. 1997. The Enclave, the Citadel, and the Ghetto: What Has Changed in the Post-Fordist US City. *Urban Affairs Review* 33 (2): 228–64.
- Massumi, Brian. 2007. Potential Politics and the Primacy of Preemption. *Theory & Event* 10 (2).
- Mazzarella, William. 2004. Culture, Globalization, Mediation. *Annual Review of Anthropology* 33: 345–67.
- Mazzarella, William. 2009. Affect: What Is It Good For? In *Enchantments of Modernity: Empire, Nation, Globalization*, edited by Saurabh Dube, 291–309. London: Routledge.
- McCullough, Debra Cohen, and Deborah L. Spence, eds. 2012. *American Policing in 2022: Essays on the Future of a Profession*. Washington, DC: Community Oriented Policing Services, National Institute of Justice.
- McLaughlin, Timothy. 2017. As Shootings Soar, Chicago Police Use Technology to Predict Crime. *Reuters*, August 5, 2017. <https://www.reuters.com/article/us-chicago-police-technology/as-shootings-soar-chicago-police-use-technology-to-predict-crime-idUSKBN1AL08P> [accessed July 21, 2018].
- Merrill, Heather, and Lisa Hoffman, eds. 2015. *Spaces of Danger: Culture and Power in the Everyday*. Athens: University of Georgia Press.
- Mohler, G. O., M. B. Short, Sean Malinowski, Mark Johnson, G. E. Tita, Andrea L. Bertozzi, and P. J. Brantingham. 2015. Randomized Controlled Field Trials of Predictive Policing. *Journal of the American Statistical Association* 110 (512): 1399–1411.
- Nail, Thomas. 2015. *Theory of the Border*. New York: Oxford University Press.
- New York Times. 2014. Sir Robert Peel's Nine Principles of Policing. April 16, 2014. <https://www.nytimes.com/2014/04/16/nyregion/sir-robert-peels-nine-principles-of-policing.html> [accessed July 21, 2018].
- Peel, Robert. 1829. *Sir Robert Peel's Principles of Law Enforcement 1829*. https://www.durham.police.uk/About-Us/Documents/Peels_Principles_Of_Law_Enforcement.pdf [accessed July 21, 2018].
- Rancière, Jacques. 2010. *Dissensus: On Politics and Aesthetics*. Translated by Steven Corcoran. New York: Continuum.
- Ransbotham, Sam, David Kiron, and Pamela Kirk Prentice. 2015. Minding the Analytics Gap. *MIT Sloan Management Review* Spring: 62–68.
- Ratcliffe, Jerry H., Ralph B. Taylor, and Amber P. Askey. 2017. *The Philadelphia Predictive Policing Experiment: Effectiveness of the Prediction Models*. Philadelphia, PA: Temple University. <http://www.cla.temple.edu/cj/center-for-security-and-crime-science/the-philadelphia-predictive-policing-experiment/> [accessed July 21, 2018].
- Reeves, Joshua, and Jeremy Packer. 2013. Police Media: The Governance of Territory, Speed, and Communication. *Communication and Critical/Cultural Studies* 10 (4): 359–84.
- Robinson, David, and Logan Koepke. 2016. *Stuck in a Pattern: Early Evidence on 'Predictive Policing' and Civil Rights..* Washington, DC: Upturn. <https://www.teamupturn.com/reports/2016/stuck-in-a-pattern> [accessed July 21, 2018].
- Ross, Darrell L., and Patricia A. Parke. 2009. Policing by Consent Decree: An Analysis of 42 U.S.C. § 14141 and the New Model for Police Accountability. *Police Practice and Research* 10 (3): 199–208.
- Saunders, Jessica, Priscillia Hunt, and John S. Hollywood. 2016. Predictions Put Into Practice: A Quasi-Experimental Evaluation of Chicago's Predictive Policing Pilot. *Journal of Experimental Criminology* 12 (3): 347–71.
- Serrano, Richard A., and Matt Pearce. 2015. Police Response to Ferguson Protests, in a Word, Failed, Federal Draft Report Says. *Los Angeles Times*, June 30, 2015. <http://www.latimes.com/nation/la-na-ferguson-draft-report-20150630-story.html> [accessed July 21, 2018].

- Shapiro, Aaron. 2017. The Medium is the Mob. *Media, Culture & Society* 39 (6): 930–41.
- Sherman, Lawrence W. 2013. The Rise of Evidence-Based Policing: Targeting, Testing, and Tracking. *Crime and Justice* 42 (1): 377–451.
- Siegel, Eric. 2018. How to Fight Bias with Predictive Policing. *Scientific American*, February 19, 2018. <https://blogs.scientificamerican.com/voices/how-to-fight-bias-with-predictive-policing/> [accessed July 21, 2018].
- Sklansky, David Alan. 2014. The Promise and Perils of Police Professionalism. In *The Future of Policing*, edited by Jennifer M. Brown, 343–54. New York: Routledge.
- Slivkins, Aleksandrs. 2017. *Introduction to Multi-Armed Bandits*. Unpublished manuscript. Microsoft Research, NYC. <https://arxiv.org/abs/1904.07272> [accessed July 21, 2018].
- Taylor, Ralph B., Jerry H. Ratcliffe, and Amber Perenzin. 2015. Can We Predict Long-Term Community Crime Problems? The Estimation of Ecological Continuity to Model Risk Heterogeneity. *Journal of Research in Crime and Delinquency* 52 (5): 635–57.
- Tchekmedyian, Alene. 2016. Police Push Back Against Using Crime-Prediction Technology to Deploy Officers. *Los Angeles Times*, October 2. <http://www.latimes.com/local/lanow/la-me-police-predict-crime-20161002-snap-story.html> [accessed July 21, 2018].
- Townsend, Michael, Ross Homel, and Janet Chaseling. 2000. Repeat Burglary Victimisation: Spatial and Temporal Patterns. *Australian & New Zealand Journal of Criminology* 33 (1): 37–63.
- Townsend, Michael, Ross Homel, and Janet Chaseling. 2003. Infectious Burglaries: A Test of the Near Repeat Hypothesis. *British Journal of Criminology* 43 (3): 615–33.
- Wachter-Boettcher, Sara. 2017. How Silicon Valley's Blind Spots and Biases Are Ruining Tech for the Rest of Us. *Washington Post*, December 13, 2017. <https://www.washingtonpost.com/news/posteverything/wp/2017/12/13/how-silicon-valleys-blind-spots-and-biases-are-ruining-tech-for-the-rest-of-us/> [accessed July 21, 2018].
- Wakefield, Jane. 2013. Future Policing Will Go Hi-Tech. *BBC News*, July 3, 2013. <https://www.bbc.com/news/technology-22954783> [accessed July 21, 2018].
- Weisburd, David, and Anthony A. Braga. 2006. Hot Spots Policing as a Model for Police Innovation. In *Police Innovation: Contrasting Perspectives*, edited by David Weisburd and Anthony A. Braga, 225–44. New York: Cambridge University Press.
- Wilk, Elvia. 2017. The Proxy and Its Politics: On Evasive Objects in a Networked Age. *Rhizome*, July 25, 2017. <http://rhizome.org/editorial/2017/jul/25/the-proxy-and-its-politics/> [accessed July 21, 2018].
- Williams, Kristian. 2015. *Our Enemies in Blue: Police and Power in America*. Oakland: AK Press.
- Winston, Ali. 2018. Palantir Has Secretly Been Using New Orleans to Test Its Predictive Policing Technology. *Verge*, February 27, 2018. <https://www.theverge.com/2018/2/27/17054740/palantir-predictive-policing-tool-new-orleans-nopd> [accessed July 21, 2018].
- Wood, Amy Louise. 2009. *Lynching and Spectacle: Witnessing Racial Violence in America, 1890-1940*. Chapel Hill: University of North Carolina Press.
- Zafar, Muhammad Bilal, Isabel Valera, Manuel Gomez Rodriguez, and Krishna P. Gummadi. 2017. Fairness Constraints: Mechanisms for Fair Classification, v5. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*. Fort Lauderdale, FL. <http://arxiv.org/abs/1507.05259> [accessed July 21, 2018].