




A Machine Learning Evaluation Framework for Place-based Algorithmic Patrol Management

Duncan Purves and Ryan Jenkins

August, 2023

A report supported by the
National Science Foundation and the
Center for Advancing Safety of Machine Intelligence

Electronic copy available at: <https://ssrn.com/abstract=4549861>

Primary Authors

Duncan Purves

dpurves@ufl.edu

Department of Philosophy

College of Liberal Arts and Sciences

University of Florida

Ryan Jenkins

ryjenkin@calpoly.edu

Ethics + Emerging Sciences Group

Philosophy Department

College of Liberal Arts

California Polytechnic State University, San Luis Obispo

Cover photo: Fred Moon, unsplash.com

Suggested citation: Purves, Duncan and Ryan Jenkins. "A Machine Learning Evaluation Framework for Place-based Algorithmic Patrol Management." Report funded by US National Science Foundation, awards #1917707 and #1917712 and the Center for Advancing Safety of Machine Intelligence. August, 2023.

Stable link: <https://ssrn.com/abstract=4549861>

Last major revision: August 3, 2023

Version: 1.0.1

Contents

Introduction	I
Normative Background	6
Background Goals, Values, and Norms	6
Legitimacy	7
Feature Assessment Fact Gathering	15
Data	15
Algorithmic Choices	24
User Interaction	25
Over-saturation and Feedback Loops	31
Organizational Ethics	44
Conclusion and Next Steps	48
New legal regimes for pass-through data collection	49
Enhanced community oversight and collaboration	50
Fairness in Machine Learning	51
Next Steps	53
About the Authors	54
Appendix: Condensed Recommendations	55
Legitimacy	55
Data	56
Algorithm Questions	57
User interaction	58
Over-saturation and Feedback Loops	60
Organizational ethics	63

Appendix: Workshop Participants and Consultants	64
Workshop Participants	64
Senior Personnel.	68
Project Consultants	69
Student Assistants	70
References.	71

Introduction

American law enforcement agencies continue to make substantial investments to develop data-driven technology to fight crime. Police departments are turning to data-driven analyses and forecasts to manage the exponential growth of intelligence data and digital assets they have access to (IDC 2021). As a result, the global market in digital policing technologies is expected to grow at a compounded annual growth rate of over 10% between now and 2027 (GME 2023).

The term “data-driven policing” encompasses a variety of big data applications. What characterizes big data systems, of which data-driven policing is one instance, is that they make use of vast quantities of data from a variety of data sources, and, using rapid computer processing and algorithms, identify correlations or patterns in the data (Brayne 2020 Paraphrasing Doug Lane (2001))—often patterns that humans would otherwise overlook—to classify new instances in the domain of interest. This process of pattern recognition and classification occurs through the use of computer algorithms, which are sometimes developed through the process of machine learning.

The landscape of data-driven policing methods is constantly evolving, and police departments offer limited publicity about the use of such technologies. It is therefore difficult to know precisely how many departments employ data-driven methods. Nonetheless, it is safe to say that most major cities across the U.S., as well as several smaller cities and towns, use one or more forms of data-driven policing technology. On its website, Geolitica (formerly PredPol) claims its software currently “help[s] protect roughly one out of every 30 people in the United States” ([Geolitica.com](https://www.geolitica.com) 2021). Furthermore, even some departments that have discontinued using certain methods still rely on others. For example, despite abandoning PredPol, the LAPD Strategic Plan for 2019–2021 lists “enhancing data-driven policing” as the #1 activity in its initiative to reduce crime and victimization (LAPD 2019).

One of the most widely used data-driven policing applications is place-based algorithmic patrol management (PAPM), more commonly referred to as place-based predictive policing—usually pejoratively. PAPM is the practice of collecting and analyzing data about previous crimes for identification and statistical prediction of geospatial areas with an increased

probability of criminal activity for the purpose of developing policing intervention and prevention strategies and tactics (Meijer and Wessels 2019).

Technology companies that develop and sell place-based algorithmic patrol management systems (PAPM systems) tout a variety of benefits (see, e.g., [Geolitics.com](https://www.geolitics.com) 2021; SoundThinking 2023). Most notably, PAPM systems are meant to provide a superior alternative to gut-based patrol allocation and hotspot analysis. Gut-based allocation can easily be swayed by conscious or unconscious human bias, and it offers little transparency or accountability for resource allocation decisions. Second, at a time when attrition in law enforcement is at an all time high, recruitment is more challenging than ever, and law enforcement budgets are shrinking, PAPM allows police departments to do more with less by outsourcing a significant part of crime analysis to the algorithmic system. PAPM therefore promises more efficient use of police patrol time, taskforce activities, and crime analyst resources for resource-strapped agencies. Third, a guiding aim of PAPM systems is to increase police visibility in a high risk area at high risk times, thereby deterring crime, without increasing the number of adverse police contacts with members of the public. Finally, some PAPM systems allow departments to closely monitor the movements of police on patrol and to analyze police officers' movements in relation to the predictions generated by the PAPM system. For example, the system can track how long police have spent in a specific area during a shift. This data offers greater transparency in police activity, and it can be used to address concerns raised by citizens that an agency is under- or over-policing in certain areas.

In spite of these advertised benefits, PAPM has been the subject of public scrutiny and suspicion because of concerns about privacy, bias, transparency, and community impacts, among others (Li 2022; Patel 2015; Meijer and Wessels 2019; Ferguson 2017a; Shapiro 2017). Yet, data-driven policing technologies are rapidly becoming police orthodoxy. In this document, we hope to anticipate the public's concerns and safeguard the benefits of these tools. In the absence of regulation targeted at the use of artificial intelligence in law enforcement, the responsibility falls to developers, law enforcement agencies, and community representatives to provide effective guidance and ensure the responsible development and deployment of these technologies.

While the social and ethical risks of PAPM have been widely discussed, little guidance has been provided to police departments, community advocates, or to developers of place-based algorithmic patrol management systems (PAPM systems) about how to mitigate those risks. The framework outlined in this report aims to fill that gap. This document proposes best practices for the development and deployment of PAPM systems that are

ethically informed and empirically grounded. Given that the use of place-based policing is here to stay, it is imperative to provide useful guidance to police departments, community advocates, and developers so that they can address the social risks associated with PAPM. We strive to develop recommendations that are concrete, practical, and forward-looking. Our goal is to parry critiques of PAPM into practical recommendations to guide the ethically sensitive design and use of data-driven policing technologies.

This document reflects a model of evaluation that has been developed based on many conversations and contributions. Ideally, the evaluation of a novel technology, such as a PAPM system driven by machine-learning, should begin before implementation. A natural first question is therefore: *Will this technology benefit society?* (The Policing Project 2020). In this case, however, place-based systems have been in use for over a decade. The task before us, then, is to reflect on the criticisms that have been raised against these systems by academics and communities and to parry those critiques into concrete recommendations to guide existing data-driven policing technology and subsequent generations of that technology.

Our approach addresses a variety of audiences. Some of our readers will be developers considering whether to begin a project in PAPM; others will be considering how to shape their product roadmap toward the next iteration of their product. Other readers will represent police departments, academics who study the intersection of law and technology, or policy makers interested in erecting safeguards around technologies that implicate important legal rights. Still others will be advocates for policed communities. **Our recommendations are directed primarily at developers, i.e., the frontline workers who create and refine the tools we are evaluating.** However, other recommendations are aimed at community advocates, policymakers, and law enforcement agencies. In many cases, the recommendations encourage closer collaboration between a variety of stakeholders and agencies. At any rate, we expect all of the audiences mentioned here will find something to gain from the recommendations we give below.¹

Because our recommendations are not directed at lawmakers but rather developers of PAPM systems, community advocates, and police depart-

1 The term “data-driven policing” is capacious and the term “predictive policing” has fallen out of favor. We are aware that the context of the implementation of PAPM we envision here is not the only one that police departments might undertake. They might, for example, enter into a partnership with a university to help them model and forecast crime. They might develop such tools internally—though few departments have the resources to spare for this. At any rate, we expect many of our recommendations will be portable to these other contexts. Whether the audience is composed of developers at a private firm, at a university, or within a law enforcement agency, our recommendations speak to the process of developing and deploying this technology.

ments, our recommendations do not directly address legislative action to regulate the use of PAMP technologies. Our recommendations address the current state of PAMP in use at the discretion of law enforcement agencies. We will not recommend legislative oversight or authorization of PAMP technologies, or a moratorium on the use of PAMP in this document. A moratorium has been endorsed by other scholars already (Robertson, Khoo, and Song 2020). Our recommendations are consistent with greater legislative oversight of PAMP, and they are consistent with a moratorium on the adoption of PAMP technology by law enforcement agencies, but we are interested in understanding what would be the ideal *implementation* of PAMP in law enforcement, on the assumption that it will continue to be used at the discretion of law enforcement agencies. However, if police departments and PAMP developers cannot meet the standards embodied in our recommendations below, then it may not be ethically acceptable for police agencies to adopt PAMP at all. Those who are fighting for greater oversight of PAMP technologies can also find something of value in our recommendations insofar as these recommendations could be used to inform legal requirements, should legislators and policymakers choose to take regulatory action.

Our focus in this document is PAMP applications such as ResourceR-outer (formerly ShotSpotter Connect) and Geolitica (formerly PredPol). Specifically, we focus on data collection and algorithmic choices in these applications, with other recommendations touching on the applications' user interface and interaction components and their integration into law enforcement workflows. We do not focus on person-based applications, such as Chicago's Strategic Subject List (the "Heat List") or London's Gangs Violence Matrix, or other applications of machine learning in criminal justice, though much of what we say here should port over to those applications as well. Other terms like "data-driven policing" should be understood to include applications involving large data sets and artificial intelligence, deployed elsewhere for law enforcement tasks besides patrol management, e.g. biometric and facial recognition for suspect identification, surveillance, and evidence-gathering and analysis.

This framework will be useful to: (a) system developers² as they design and build PAMP systems; (b) crime analysts and law enforcement offi-

2 In this document, we use the term "developer" when referring to those working to create machine learning applications, including data acquisition and ingestion, model training, and design of user interface and user interaction. We use the term "vendor" to refer to companies making larger-scale decisions of whether to develop PAMP products, how to market them and describe their functionality, and how to interact with customers during and after a sale. These terms cut across many different kinds of parties that might be interested in developing and employing these technologies,

cials responsible for technology acquisition who wish to make ethically responsible acquisition decisions; and (c) community advocates seeking guidelines by which to hold police departments accountable for their choices about whether and how they employ data-driven technologies.

This document has been supported by National Science Foundation awards #1917707 and #1917712, “Artificial Intelligence and Predictive Policing: An Ethical Analysis.” It has also been benefitted from a collaboration between Northwestern University, funded by a generous gift from UL (formerly Underwriters Laboratories). (This Northwestern–UL collaboration has since given rise to the Center for Advancing Safety of Machine Intelligence, or CASMI.) This document has benefited greatly from the input and feedback from many other contributors, consultants. This includes, especially, several workshop participants who convened in June 2022 in a two-day series of discussions, generously supported by the Northwestern Law and Technology Initiative, housed in Pritzker School of Law and McCormick School of Engineering. These workshop participants and other contributors are listed at the end of this document in an appendix.

The views represented in this report are the views of the authors alone, and not necessarily representative of or endorsed by any of the individuals, consultants, experts, or entities mentioned here or involved in its drafting.

e.g. for-profit vendors, or a law enforcement agency itself, or a non-profit community initiative developing a “home brew” application.

Normative Background

Before the facts related to the data, the algorithm, and the interaction are collected, it is important to understand what normative standards apply when evaluating a data-driven system. This requires understanding the goals, values, and constraints governing the domain in which the system will be used. Analyzing a system against the goals and values of a domain equips us with a set of analytical tools to judge the deployment of machine learning as appropriate or inappropriate in that domain and to articulate the tradeoffs, benefits and drawbacks of its human impact.^{3,4}

Background Goals, Values, and Norms

Background goals, values, and norms are the set of normative considerations that bear on the proper functioning of the entity (e.g., institution, agency, firm) that is using the ML application. Domain-specific background goals, values, and norms are important because they provide a standard against which to evaluate a specific technological system (Jenkins et al. 2022). Questions like these are helpful in identifying the normative considerations relevant in the domain:

- What are the goals, values, and norms of the domain in which the ML application will be deployed? What best practices do these goals, values, and norms suggest?

3 Note that this is not to suggest that the domain-specific normative standards cannot reflect more general normative standards that apply across domains. For example, as we see below, the commitment in policing to impartial enforcement of the law likely reflects a more general principle of anti-discrimination. Still, we think the domain of application is the place to begin the normative evaluation.

4 We reiterate that our recommendations here are meant to be agnostic in relation to the legislative and regulatory landscape governing law enforcement's use of technology. We are sympathetic to those who would suggest that the decision to deploy a technology like PAPM, which is one of the most momentous decisions contemplated in this document, should be regulated by governments, e.g. at the city or state level. Regulation could also address standards for maintaining the applications and making them available for audits or impact assessments. But taking up questions here, e.g. about the ideal shape of regulation, would take us outside the scope of our project.

- Is the use of the system in line with best practices within the domain? If not, what values might be violated if the best practices are not followed?
- Does the application impact any of the broader goals or best practices of the domain?
- Our first task, then, is identifying the goals and values of policing as a distinct domain of practice.

Legitimacy

It is not possible to provide a list of goals, values, or norms that have been adopted universally by law enforcement agencies across the United States. To understand the goals of policing, it is perhaps best to review some of the characteristic functions of police in a democratic society. Police are called upon to: identify criminal offenders and criminal activity and to apprehend those offenders where appropriate; reduce the opportunities for the commission of crime through preventative patrol or other means; assist individuals in danger of harm; protect constitutional guarantees; facilitate the movement of people and vehicles; assist those who cannot care for themselves; resolve conflicts (Perez and Moore 2013, 43; Klaver 2014, 9); create and maintain a feeling of security in the community; and maintain civil order (Perez and Moore 2013, 62,93; Caldero, Dailey, and Withrow 2018, 254, 259–60). Taken together, these activities suggest one core goal of policing:

Goal: The core goal of policing is to prevent crime and disorder. Catching and punishing criminals is one means of crime prevention, though ideally police would also prevent crime and disorder by deterring it before it happens.

This basic aim of policing provides a further standard by which to assess policing activities. We can ask: Do these activities prevent crime and disorder more effectively than other activities police might perform? The pursuit of crime prevention is constrained by various other considerations. The crisis of legitimacy faced by American police agencies in the past decade also provides a standard by which to evaluate policing activities and organizational arrangements. We can ask: Do these activities, practices, and organizational arrangements promote legitimate policing? Here “legitimacy” is understood as the property of being morally permitted to exercise political power (Monaghan 2021). Law enforcement agencies exercise the state’s political power when they stop, question, search, or detain citizens (B. Jones and Mendieta 2021, 6) Police legitimacy may seem

too nebulous a concept to be the basis of an evaluation of specific policing technologies, but below we spell out some widely accepted requirements for police legitimacy.

At a minimum, police agencies are not permitted to prevent crime and disorder if doing so violates citizens' constitutional rights. Legitimate policing is also sufficiently competent, proportionately enforces the laws, enforces laws equally regardless of a person's race or other protected characteristics, and promotes public trust in policing (Purves and Davis 2022; Monaghan 2021).

Sufficient competence demands that a police agency meets some minimum bar of effectiveness in achieving the aims that justify the agency's existence. We understand the basic aim of law enforcement to be the preservation of public safety, with core subsidiary aims of preventing crime and disorder. Ineptitude at achieving these aims threatens an agency's permission to exercise political power through law enforcement. To take an extreme example, a police agency that had a zero percent case closure rate, or that solely targeted innocent citizens but not offenders for arrest, or that eschewed the use of evidence when conducting investigations, would not meet a minimum level of competence and would thus be illegitimate (Lovell 2014, 411).

"Legitimate policing must be conducted impartially, with no favor given to some social groups over others."

Proportionality demands that the risk imposed on citizens by a means of crime prevention must be proportionate to the crime reduction goal to be achieved. There is often a tension between the measurable ends of policing and the means of achieving those ends. In a phenomenon referred to as 'noble cause corruption', police officers can justify extralegal or overly harsh tactics by imagining that they are reducing the rate of crime, which is what "really matters." Similarly, community groups can demand that police do whatever is necessary to bring crime under control, or a majority community can demand strict policing of a minority community in the service of overall public interests. (Caldero, Dailey, and Withrow 2018, 115,118,119; Perez and Moore 2013, 216; Elliott and Pollock 2014, 247; T. L. Meares 2021, 28). In each of these cases proportionality is threatened. At the very least, to be proportionate, the means used to achieve a legitimate policing aim must not cause more harm than they prevent. For example, using lethal force to enforce laws against operating unauthorized taxis would not be a permissible form of policing, because the means is not proportionate to the crime reduction goal to be achieved. A commitment to proportionate use of force entails a further commitment: the use of force is a **last resort** in crime prevention, and that public cooperation in crime prevention is always to be preferred. The reason for this is simple: use of force imposes the greatest risk of harm on citizens, and it therefore poses the

greatest threat to proportionality. Preventing crime without resort to force avoids risking harm to citizens. It should therefore always be preferred, other things being equal, for the sake of preserving police legitimacy.

Equality demands equal enforcement of the laws for all groups, and equal protection of the laws for all groups. Legitimate policing must be conducted impartially, with no favor given to some social groups over others. As philosopher Jake Monaghan puts the point, “A police department that tends to enforce laws against one race or ethnicity and not another is likely illegitimate; at that point the department stops looking like a police force [and more like a gang]” (Monaghan 2021, 44). If a policing practice creates a pattern of law enforcement that burdens one protected group while benefiting another, then the practice conflicts with the requirement of equality. In this sense, equality in policing requires, at a minimum, a commitment to anti-racist policing policies and practices.

Finally, **public trust** in and support of police are key to legitimate policing (Purves and Davis 2022). This is because trust is necessary for cooperation from the public, and cooperation from the public aids in crime prevention (Tyler and Huo 2002; Desmond, Papachristos, and Kirk 2016). Therefore, public trust supports competent policing. But public trust also matters in its own right. Police are servants of the public and the democratic state. Their exclusive authority to enforce the laws is therefore dependent on maintaining a relationship to the governed that citizens cannot reasonably reject (Purves and Davis 2022).

One of the ways to promote public trust is to give members of the public a say in policing tactics and priorities. Consider, for example, the emergence of community-led policing across the United States. According to some scholars, as many as 95% of police departments in cities of more than 250,000 people included a commitment to community policing in their mission statements (Skogan 2019). Community policing is characterized by efforts to include community members in law enforcement **priority setting and tactics development**. While the broad basic aim of policing is to prevent crime and disorder, police departments have limited resources, and the broad aim of preventing crime and disorder does not adjudicate between competing priorities in crime prevention. Officials must decide which laws to enforce and to what extent, in order to meet desirable metrics without violating the rights of citizens or allowing too many criminals deserving of punishment to go unpunished (Perez and Moore 2013, 68, 169; B. Jones and Mendieta 2021, 8; Caldero, Dailey, and Withrow 2018, 74). Should a department spend their limited resources targeting a recent rash of vehicular thefts, an alarming trend of sexual assaults on a college campus, or public intoxication and disorderly conduct in a city’s bar district? Community policing helps police to answer these ques-

“Public trust in and support of police are key to legitimate policing.”

tions, in part, by placing public safety priorities of community members front and center. Academic literature supports the conclusion that public trust and collaboration are key to both crime prevention and legitimate policing (Tyler 2004; Desmond, Papachristos, and Kirk 2016; Purves 2022; Purves and Davis 2022).

The problems of policing are not simply problems of finding “efficient” and “effective” means; they are problems of ends, of competing social values, interests, and priorities the resolution of which raise fundamental moral and political issues to be decided by an informed citizenry, not only scientific or technical issues to be decided by experts and technocrats. Hence, the most hopeful prospect of substantive police reform is the influence an informed public can exert on the direction of change in police agencies. (Rumbaut and Bittner 1979, 284)

Community policing demands greater transparency about police operations, including providing community members access to crime statistics. If police and citizens share the responsibility for crime reduction, then both police and citizens must have sufficient understanding of the tools of law enforcement being employed by police departments. This requires that technology companies build their products to be sufficiently transparent, with public-facing information portals, for instance, that cannot be hidden or disconnected. Police should communicate clearly and openly with their communities about the technologies and methods that they are considering adopting, and with a willingness to ‘pivot’ if told that those technologies are strongly opposed by the public or raise other concerns. Police departments must be receptive to skepticism from the public, because, again, the public should be seen as a full partner in crime reduction.

“Community policing demands greater transparency about police operations, including providing community members access to crime statistics.”

So-called “problem-oriented policing” also promotes public trust—and therefore legitimate policing—insofar as it emphasizes strategies to prevent crime that minimize adverse interactions between citizens and police. According to Herman Goldstein who created the approach (Goldstein 1979), enforcement of the laws is over-emphasized in traditional police work to the detriment of communities. Instead, problem-oriented policing proposes that problems—specifically problems brought forward by the public—are the starting point of police work. Once a problem is identified, the police should identify an array of potential solutions to those problems, only some of which will involve *enforcement*-oriented activities such as arresting people. As described by criminologist John Eck, problem-oriented policing is characterized by a *normative* commitment to the view that “police are supposed to reduce problems rather than simply respond to incidents and apply the relevant

law” (Eck 2019, 167). Recent scholarship supports the effectiveness of problem-oriented policing in preventing crime (Weisburd and Green 1995; Braga et al. 1999).

If we combine the basic aim of policing to prevent crime and disorder with the constraints on police activity imposed by the requirements of legitimacy, we can describe the core function of police, constrained by relevant rights, as follows:

Standard: A core goal of policing is to prevent crime and disorder, ideally by preventative means, while respecting citizens’ constitutional rights and preserving the conditions required for legitimate policing.

Standard provides a criterion against which we can assess the development and deployment of PAPM systems.

To meet Standard police must satisfy the requirements of legitimate policing. Among other things, police agencies must develop **evidence-based standards**. To *know* whether a policing practice is proportionate, for example, one must gather evidence of its effects, on crime to be sure, but also on innocent members of the community who may bear the burdens of the practice. Consider, for example, the New York Police Department’s stop, question, and frisk policy by which officers detain individuals suspected of being engaged in criminal activity for a short period of time and “pat down” individuals who are detained. Stop, question, and frisk imposes costs on detained individuals, ranging from embarrassment and a temporary infringement of their freedom of movement to serious physical harm. Furthermore, searches without reasonable suspicion threaten protections afforded by the Fourth Amendment to the US Constitution (Ferguson 2017b). Do the benefits in the form of crime reduction justify these costs? A number of scholars have argued that the answer to this question is “no,” because there is little evidence that stop, question, and frisk reduces crime. In fact, some evidence suggests that the drop in crime in New York City in the 90s and 2000s was due to other factors (Monaghan 2021; Fagan et al. 2020; Rosenfeld and Fornango 2014; Apel 2016). The failure of the NYPD to attend to this evidence caused them to adopt what might have been disproportionate enforcement practice, undermining the agency’s legitimacy. Call this further requirement “epistemic responsibility.” Evidence-based standards are also necessary to assess whether a policing practice threatens equality in policing. For example, evidence-based evaluations of stop, question, and

“If a policing practice creates a pattern of law enforcement that burdens one protected group while benefiting another, then the practice conflicts with the requirement of equality.”

frisk show that the majority of individuals stopped were people of color and a disproportionate number were black (ACLU of New York 2022; Speri 2021; Southall and Gold 2019). This calls into question whether the policy is enforced equally for all racial groups.

Because legitimate policing requires evidence-based standards, a crucial task for designers of PAPM systems is to validate those systems, ideally before deploying them in the field. Assessing the efficacy of a PAPM system ideally includes randomized controlled trials. These are resource-intensive research activities, requiring collaboration with academics (to ensure the soundness of the methods) and community members (to understand what the crime-reduction goals of the system should be). Randomized controlled trials require developing methods by which to compare the efficacy of using PAPM to forecast crime with a relevant control group. What the “control” group is will depend on what the status quo ante was prior to the implementation of the PAPM system. Few randomized controlled trials of PAPM technologies have been conducted to date (Hunt, Saunders, and Hollywood 2014; Mohler et al. 2015; Saunders, Hunt, and Hollywood 2016). The authors of the few controlled trials of a place-based algorithmic crime forecasting system described the control in their experiment as “hotspots maps produced each day and shift by dedicated crime analysts” (Mohler et al. 2015). The key to the control group is thus the distinctively human element involved in the crime analysis. The dilemma facing randomized controlled trials is that a randomized controlled trial is effectively an experiment on a community with PAPM as the “treatment.” But this would certainly not pass institutional review board (IRB) or bioethical standards.⁵ The only way to collect robust data demonstrating the efficacy of PAPM systems is therefore open to serious ethical challenges.

Short of a randomized controlled trial, PAPM developers can indirectly validate their work by using empirically validated theories and approaches to crime prevention. For example, one empirically validated practice is to observe the “Koper curve,” discussed above. The Koper curve is the product of hot spots policing experiments indicating that the positive deterrent effect of police patrols in an area is maximized when police spend about 10-15 minutes patrolling there. Returns begin to diminish greatly after 15 minutes (Koper 1995; Williams and Coupe 2017). Another example of an empirically validated principle from criminology is that the environmental features of places affect criminality at those places. For example, studies have shown that a variety of environmental features correlate with

5 For a related conversation about how to understand the deployment of emerging technologies in communities as a kind of experiment, and therefore requiring the kind of oversight traditionally provided by IRBs, see Van de Poel (2016; 2013).

increased crime risk (Barnum et al. 2017; Hart and Miethe 2014; Piza et al. 2016). A PAPM system that is appropriately informed by locations' environmental vulnerabilities to crime when generating predictions therefore stands on firm empirical footing. Reasonable parties might disagree about whether this form of validation is sufficient to justify deployment of PAPM in real-world settings without having first conducted a randomized controlled trial.

The requirements of police legitimacy, combined with some sensible lessons emerging from contemporary policing approaches such as community-led and problem-oriented policing, suggest the following recommendations for developers and police departments at the design and implementation phases of the PAPM system lifecycle:

- ▶ **RECOMMENDATION:** When implementing a new PAPM system, departments should, at the same time, establish standards and methods necessary to assess the efficacy of the system.
- ▶ **RECOMMENDATION:** Develop a feasible method by which to measure the system's effect on public approval of police actions as well as trust of police. If the system negatively affects public approval, modify the design or implementation of the system.
- ▶ **RECOMMENDATION:** When setting priorities, departments must seek input from community members about citizens' needs and concerns, and this often requires creating new channels for feedback from community members, including partnerships with civic organizations, town halls, regular consultation with neighborhood associations about safety concerns and priorities, and the creation of citizen-led advisory committees with real capabilities for oversight.⁶
- ▶ **RECOMMENDATION:** Law enforcement priorities, informed by community consultation, should guide technology adoption rather than technology adoption guiding law enforcement priorities. To do otherwise risks minimizing the role of community members in priority setting, in violation of the requirements of legitimate policing (Bennett Moses and Chan 2018). If a police department cannot employ a PAPM system without irreparably damaging public trust—because public concerns

6 The voices of those generally skeptical of law enforcement or the deployment of PAPM are especially important to consider. However, they may also be especially difficult to capture if they simply refuse to sit on or cooperate with a committee like the ones we suggest here. We do not see an obvious solution to this paradox except the gradual and effortful restoration of police-community relations.

about the technology cannot be satisfactorily addressed—the department should consider alternative means of crime analysis.

- ▶ **RECOMMENDATION:** Create mechanisms for public feedback about the use of PAPM systems by police.
- ▶ **RECOMMENDATION:** Ask whether the PAPM system encourages policing tactics that involve the use of physical force. If so, ask whether the technology can be used in combination with less aggressive tactics.
- ▶ **RECOMMENDATION:** Develop methods to determine whether the system encourages unequal enforcement of the laws (see “feedback loops” below).
- ▶ **RECOMMENDATION:** If, at present, use of the system encourages unequal enforcement of the laws, change the design or implementation of the system to achieve equal enforcement.

Some of these recommendations are very general, but in the next section we will offer specific steps that developers and police departments can take to address these general recommendations.

Feature Assessment Fact Gathering

As we consider a particular system, we need to identify both the features that are specific to it (data sourcing and quality, algorithms used, user interactions) and those that are linked to the more general features of the data, algorithms, and role of the system. For both system-specific and general features, we have posed key questions to ascertain the facts.

Data

Machine Learning systems are built on data and continually rely on data in their use. To evaluate a system's data foundation, considerations pertain to the intended use, quality, completeness and coverage of data.

The processes that define data flow also define the set of issues to evaluate in order to understand the system that results from them. Questions of how the data were gathered, cleaned, integrated with other sources, and augmented all have to be asked to develop a complete set of facts.

A Note on Privacy

We are aware of and sensitive to widespread concerns about the privacy implications of PAPM technologies. Communities of color, already disproportionately surveilled and anxious about government scrutiny, understandably bristle at the thought of supercharged data collection and analysis being “targeted” where they live. These concerns are important, but we ultimately have to leave those questions for another investigation, because this report emphasizes depth over breadth.⁷ Our analysis begins

⁷ Moreover, these concerns have already been the subject of excellent treatments elsewhere. See, especially, Brayne's book-length treatment (2020) and Ferguson's chapter on the widespread suspicion of data-driven policing (2017b, 54 ff.). See, for example, Ferguson's discussion of the 4th Amendment implications of PAPM and the concern that PAPM forecasts could constitute or buttress “reasonable suspicion” by officers:

Uncertainties about place-based predictive policing continue in the constitutional realm. Can location factor into a police officer's suspicion? Are there areas where police should be more suspicious?... The Supreme Court has held that police observation in a “high crime area” can be a factor in deciding whether the officer has reasonable suspicion or probable cause. Rather unhelpfully, the Court has

from the point at which developers are considering, of the data sets they have available, which they should use to train crime forecasting models. That is, our analysis begins after the moment that data are collected—it does not address concerns about the methods or propriety of data collection.

Framing Questions

- What is going to be learned from this data? What new information is going to be derived from this data?
- What is the data supposed to be representing? That is, what features are being captured by the data? Are they relevant to the task? Are other relevant features missing?
- Given what the data are supposed to represent, what features or attributes should be included in the data?
- What actions will this data be used to inform?

One of the first and most fundamental questions for designers of PAPM systems is to ask *which crimes* are going to be predicted—the answer to this first question will cascade throughout the design process. It is too simple to say that the system will “predict the timing and location of crimes,” because data on different crimes differ with regard to their completeness, reliability, bias, source, geographic distribution, and so on (Duffee et al. 2000).

Some crimes are highly geographically correlated. For example, there is often robust data about the location of Part II crimes such as “nuisance crimes.”⁸ These crimes are easier to model because the geographic data

never defined a “high crime area,” but as can be imagined, predictive policing technologies might be quite useful in mapping such areas... [T]he conclusion that a computer algorithm could alter Fourth Amendment freedoms in certain areas and especially in communities of color should be of great concern. Issues of disparate treatment, accuracy, transparency, and accountability all demand attention. If walking through a predicted red box changes my constitutional rights to be free from unreasonable searches and seizures, then a higher level of scrutiny might need to be brought to bear on the use of the technology. (Ferguson 2017b, 76–77)

- 8 The distinction between Part I and Part II crimes is due to the Uniform Crime Reporting program, where Part I crimes are designated as such specifically because of their greater seriousness or harm (FBI 2011): “The UCR Program collects data about Part I offenses in order to measure the level and scope of crime occurring throughout the nation. The program’s founders chose these offenses because they are serious crimes, they occur with regularity in all areas of the country, and they are likely to be reported to police... The Uniform Crime Reporting (UCR) Program divides offenses into two groups, Part I and Part II crimes. Each month, participating law enforcement

about these crimes is plentiful, and they tend to cluster in ways that are useful for police on patrol. Furthermore, nuisance crimes, because they occur in public view, can be effectively deterred by additional police presence in an area. It is therefore tempting to model nuisance crimes for the purpose of allocating police patrols.

But modeling nuisance crimes for the purposes of allocating targeted patrols can conflict with community-led policing, insofar as community members do not place a high priority on the enforcement of nuisance crimes. This raises a core issue in need of critical analysis: there are multiple, potentially conflicting justifications for policing tactics or strategies: What justifies a choice to model and patrol one crime rather than another? Because the community indicated that it was a priority? Because there would be disparate impact or unequal enforcement otherwise? Or because it was the costliest crime according to RAND's "cost of crime" calculator? These justifications could all align, but they could also conflict.

- ▶ **RECOMMENDATION:** By any reasonable measure of harm, not all crimes are equally serious (Heaton 2010). Developers should create PAPM models capable of prioritizing some crimes over others so that police departments can prioritize resource allocation accordingly. Police departments should ask developers for their rationale about which crimes to model as well as developing their own rationale.
- ▶ **RECOMMENDATION:** Developers and police alike should seek counsel from citizens, community organizations (e.g., neighborhood associations, citizen-led police advisory or oversight councils, chambers of commerce), advocacy and policy organizations (see, e.g., NYU's Policing Project), and other stakeholders throughout the development and deployment of PAPM systems. This helps to ensure that choices made reflect the priorities and concerns of community members. This includes, for example, choices about the crimes targeted for prediction and the interventions chosen in response to PAPM forecasts (more on this last point below). Due diligence could require consulting with groups opposed to the use of PAPM models in the first place, but even antagonistic perspectives are important.
- ▶ **RECOMMENDATION:** RMS developers and police departments should design technologies with an eye toward transparency, e.g. build in fea-

agencies submit information on the number of Part I offenses that become known to them; those offenses cleared by arrest or exceptional means; and the age, sex, and race of persons arrested for each of the offenses. Contributors provide only arrest data for Part II offenses."

tures that easily produce transparent reports accessible to community members.

Some crimes are not good targets for prediction by police departments because they are not deterred by the presence of police patrols. Examples of these include, for example, domestic violence, shoplifting, and “white collar” crimes like fraud or embezzlement. Because the outputs of the PAPM systems are used primarily to allocate police patrols, it makes little sense to model crimes that are unaffected by these interventions.

- ▶ **RECOMMENDATION:** Select crimes for modeling which are: geographically correlated; patrollable and deterrable through police presence, where “presence” includes situational crime prevention tasks; high priorities for community members, ideally, as evidenced by their explicit endorsement. Avoid modeling crimes if their associated training data is especially susceptible to enforcement bias (discussed further below).
- ▶ **RECOMMENDATION:** Police should seek explicit community endorsement and involvement in the selection of crimes to police. In the absence of such endorsement, police should target crimes for prediction only if the harm resulting from these crimes is serious enough that they are a reasonably high priority for interventions that might risk disparate impacts or the corrosion of police-community relations.
- ▶ **RECOMMENDATION:** Developers should avoid modeling crimes that are officer-discovered, as they might reflect differential selection on the part of police officers enforcing problematic policies such as “stop and frisk.”

Quality Questions

Once we have decided which crimes are going to be predicted, there is a related question of the source of the data about those crimes. We know of no PAPM systems that directly take protected characteristics such as race into account when generating forecasts— in fact, it might be illegal in the United States to do so (Hellman 2020; Corbett-Davies and Goel 2018a).

However, many sources of data that are useful for PAPM have potential to correlate with race or other protected factors, for example, the location of multi-family housing units, census tract information on median household income, the location of pawn shops, and so on. Developers face a tradeoff between (1) the utility of these factors in predicting crime and (2) their correlation with protected characteristics.

Most Part I crimes are called in by the community, so by using data only from community reports, a model will de facto exclude most Part II

crimes. Some crime data is community-generated, e.g. “calls for service” like 911 calls. Other crime data is generated by officers on patrol, i.e. “discovered” crime. Discovered crime data tends to have more pronounced correlations with race. Officer-discovered crime data therefore raises acute concerns about bias. The term “bias” is used to refer to a variety of phenomena by scholars writing about data-driven and AI systems. Here we define bias in a PAPM system as follows:

Bias. The tendency of a PAPM system to predict crime at some locations out of proportion to the actual crime rate at those locations.

If a PAPM system predicts crime at some locations out of proportion to the actual crime rate at those locations, it “overpredicts” crime at those locations. It is in that sense *biased with respect to* those locations. When a PAPM system overpredicts crime in locations, especially when racial minorities are concentrated in those locations, this raises concerns about disparate impact and unequal enforcement of the laws (Barocas and Selbst 2016). Bias in prediction is an especially acute worry for nuisance crimes. Because nuisance crimes, including vandalism, prostitution, and vagrancy tend to be “officer-discovered,” historical arrest data about those crimes will reflect where officers have decided to patrol (and, hence, find crime) in the past. If officers have targeted some places out of proportion to the actual crime rate at those locations, then models built on arrest data about nuisance crimes will demonstrate bias as defined above.

The records management systems (RMSs) used by police departments frequently, but do not always, include a simple data label for which crimes are community-reported or officer-discovered.

- **RECOMMENDATION:** RMS developers and police departments should include a data label to facilitate the exclusion of officer-discovered crimes in crime forecasting. In fact, if RMS systems universally included this label, and if the labels were filled out perfectly, this would significantly undermine the force of the “feedback loop” criticism, discussed further below, since developers could guarantee that data generated by officers on patrol was not creating a self-fulfilling prophecy.

Even when officer-discovered crimes are excluded from the set of training data, questions of racial correlations face developers throughout the development process: when and why are correlations between race and crime data problematic? When does a correlation become problematic between some predictively powerful feature and race? Second, we must investigate whether machine learning algorithms are able to overcome and counteract those biases. While humans are undoubtedly biased, and this is

likely reflected in historical crime data, some AI systems can be audited to excavate and extirpate the sources of that bias. Further, predictive models can incorporate additional data sources such as weather and geography, which may add predictive power without adding bias. To test a given feature's predictive power, developers can use feature importance methods which check the model's performance with and without the given feature (Zien et al. 2009). To ensure that given features don't add excessive bias to the model, developers can measure the given feature's association with protected variables (Datta et al. 2017). Both of these methods are algorithm agnostic, and developers can thus perform these tests even if their PAPM system utilizes black box algorithms.

Calls for service are a source of data that may be less influenced by police officer behavior. Using data about community calls for service to generate crime forecasts would impart a measure of democratic legitimacy to the use of AI models insofar as they reflect the perceived need for police assistance by community members. Some PAPM applications use community calls for service. For example, we know of one application whose forecasting models are initially trained from three to five years of community calls for service and historical crime data (excluding records of officer-discovered crimes). Six months of that data is reserved for the post-training testing phase. After that, data is collected every shift, and the models are retrained periodically roughly every month.

However, relying on calls for service instead of arrest data does not eliminate the possibility of bias. For example, residents of some neighborhoods call the police more than residents of other neighborhoods. This could be explained by the fact that there really is more crime in those places. But it could also be explained by the fact that residents of those neighborhoods are more willing to report the crimes that do occur, perhaps because residents of that neighborhood are more trusting of police than residents of other neighborhoods (Desmond, Papachristos, and Kirk 2016). This can lead to a paradoxical implication: if a policing method increases calls for service, this is consistent with two explanations: (1) crime in that area really is increasing, suggesting that the policing intervention has failed; or else (2) it is merely the community's willingness to report crime that has increased, which might in fact signal an *improvement* in police-community relations. Having access to reliable, longitudinal data on community sentiment would be necessary to disentangle these two competing explanations for increasing calls for service.

- **RECOMMENDATION:** Agencies employing PAPM should seek out methods to secure geographically granular, longitudinal data about community sentiment regarding crime, disorder, and the police. Because these data are expensive to generate, the most practical route would

likely involve partnerships with academia, nonprofits, or other NGOs. See, for example, the “Chicago Police Sentiment” survey conducted by the Chicago PD each month (Chicago PD, n.d.).

In order to generate reliable predictions for the timing and location of crimes, a minimum number of data points is required. Confidently predicting crimes that are rare is difficult simply because there are too few data points to infer a pattern. One solution to this problem is to lump different crimes together in the same “bucket” (Selbst 2018, 132).⁹ For example, a department might not have sufficient data on robberies, larceny, or grand theft auto, but they might reason that because these crimes are similar, they can be forecast together as a single group. Once combined, the department might have enough data points to generate sufficiently robust predictions about the future location of a crime that would be a robbery, larceny, or grand theft auto. However, the apparent similarity of these crimes could bely important differences between them. It might be the case, instead, that these crimes have significantly different drivers and aggravating factors. As a result, binning them together, and then responding to predictions with a single type of intervention, might fail to address the drivers of some of the binned crimes.

- **RECOMMENDATION:** While a technical solution to data sparsity, binning should be performed thoughtfully, if at all. In many cases it may turn out that binning ostensibly similar crimes together is unhelpful or even counterproductive.
- **RECOMMENDATION:** Officers on patrol should not be overlooked as a potential source of data collection to improve predictions moving forward. For example, if a recommended patrol was based on environmental factors such as abandoned buildings, bus stops, or liquor stores, officers can confirm the presence of those factors or correct the data.

Note that this recommendation comes with substantial caveats: First, the data collected by officers in the field would likely suffer from problems of consistency and completeness, unless officers are expected to understand and abide by best practices of data collection. This is unlikely, and so these data should be understood to be incomplete, and their utility under-

9 “For example, if the nuance between robberies and burglaries is missing because both are placed in the “property crime” bucket, the algorithm may not detect the difference between an area with high amounts of robberies and an area with a high number of burglaries, though the two crimes might be perpetrated by different people with different victims” (Selbst 2018, 132). Despite the initial attractiveness of binning crimes together, appropriate enforcement strategies for each crime could differ, i.e., burglaries and robberies might respond to different law enforcement practices.

mined accordingly. Unless these data are collected in ways that are tightly controlled and monitored, they may be worse than useless: they may end up exacerbating biases in future analyses.

This is crucial to appreciate when deciding on uses for these data collected in the field. For example, it may be acceptable to use these for informal, internal conversations, or to provide these data to community services such as urban renewal or social services via a report function in the user interface. But the data gathered this way are probably not reliable enough to use as the grounds for forecasting the location of crime or monitoring long-term trends in crime.

Second, directing officers to collect data while on patrol could erode community relations and exacerbate distrust or suspicion, as police incorporate surveillance responsibilities into their regular activities. (This is the case to the extent that it is not already obvious that police might take note of salient environmental factors or other observations while on patrol.) This is underscored by introducing police discretion into what kinds of observations are recorded, and also spurring resentment from officers (if, for example, they are asked to note every broken light they see).

Bias Questions

Perhaps the single most widespread and well known criticism of PAPM systems is that they enable racially biased policing. Because historical crime data often reflects a correlation between race and crime, police who act on the basis of these recommendations will be sent repeatedly to minority neighborhoods. Here, we explore methods for ameliorating this concern through deliberate choices at the data and model training phases of the ML pipeline.

**“Every church is perfect
until a human walks in.”**

The problem of racially-correlated crime forecasts can be ameliorated in part by incorporating other sources of data such as environmental factors that make a place vulnerable to crime. Other non-crime factors, such as census data or demographic information, which are superficially neutral, such as median household income, density of cars, or density of multi-family housing units, will often correlate with race. (Sometimes these correlations can be quite surprising, such as the correlations discovered by some of our workshop participants between race and the location of public schools, percentage of renters in an area, or the location of railroad bridges. Still, the strength of these correlations often differs by city or neighborhood, confounding this analysis further.)

In part, this reflects the sad truth that America’s urban design was for a long time informed by policies that were racist—overtly or covertly. The routes planned for highways often has a disparate impact on minority

communities (Mohl 2004); terrain elevation might correlate with socio-economic status because the poor disproportionately occupy flood-prone areas (Rentschler et al. 2022); and so on. As a result, modeling environmental features of a place that make it vulnerable to crime might reflect the racist policies from which those environmental features emerged. This presents an unavoidable challenge for developers.

On the other hand, removing from predictive models all features that correlate with race would undermine the models' predictive value (J. L. Skeem and Lowenkamp 2016).¹⁰ But some factors may contribute little, if any, predictive value. If these factors can be jettisoned without substantially undermining the accuracy and utility of predictions, then they should be. Otherwise, developers must be prepared to defend the inclusion of these factors to the communities subjected to the predictions.

- ▶ **RECOMMENDATION:** Developers should generate policies for navigating such trade offs including, if possible, explicit values for acceptable marginal increases in efficacy compared with increases in the correlation of between predictions and race. While this may be extremely time-consuming, it is also an indispensable step for guaranteeing that these technologies, already controversial, are deployed in ways that build community trust.
- ▶ **RECOMMENDATION:** Developers should check their model outputs for racial bias that cannot be justified by strong predictive power. Developers should also be able to justify the metrics used to measure model fairness, as fairness measures will often deliver conflicting results (Chouldechova 2017).

The concern that PAPM technologies are racially biased is a substantial driver of public skepticism and distrust. While there is no direct empirical evidence of this claim that we are aware of, there is good evidence for closely related claims concerning police-community relationships and trust of algorithms more broadly. Most relevant to this discussion is a Pew report citing widespread skepticism and anxiety about the use of machine learning algorithms in making predictions about the likelihood of a criminal recidivating—so-called “criminal risk assessments” or “risk scores.” These predictions are often used in judicial or parole hearings to determine the length and nature of a defendant’s sentence. The Pew Research Center found in 2018 that 50% of those surveyed thought that algorithms that gen-

¹⁰ One clever yet provocative recent suggestion is to develop auditing methods that explicitly take race into account in order to check for bias in place-based patrol recommendations (J. Skeem and Lowenkamp 2020).

erate predictive risk scores are unfair to people who are up for parole (Pew Research Center 2018, 1–4). Crucially, the percentage of people who found algorithmic risk scores to be unfair rose to 61% for black respondents (Pew Research Center 2018, 6). These results suggest that the public might find PAPM, like algorithmic risk scoring for parolees, unfair and that perceptions of unfairness might vary between communities.

Distrust undermines the effectiveness of these technologies. It even threatens to make their use counterproductive, if that distrust reduces calls for service or cooperation with police investigations.

- **RECOMMENDATION:** Departments should take steps to improve public understanding and reception of data-driven policing technologies, e.g. by disclosing methods of generating patrol recommendations to citizens via police advisory councils or discussing the results of third party bias audits with the public.

Algorithmic Choices

Once a data set has been gathered, the decision-making process shifts focus to the algorithms that are going to be applied to build the model. At this stage, developers must make decisions about which algorithms to use, which features to include, and how to segment the data set for training and testing. Given the choice of algorithm, specific questions related to the requirements, performance, explainability, and transparency are now at the fore. While there is some overlap between the analysis of both the data and the algorithms that use them, it is important to see each in its own light. As with the analysis of data, this begins with framing questions related to the goals of the domain where the system will ultimately be deployed.

In the case of forecasting crime, which would be worse: over-predicting crime in an area (false positives) or under-predicting it (false negatives)? If we knew our model had to miss in one of these directions, which one should we prefer? Over-representing crime in a community would mean sending police there more often than is actually appropriate. This could lead to negative police interactions with the communities they serve, including mistaken arrest, detention, search, or seizure. And this is in addition to the increased probability of “privacy harms” themselves (Calo 2011). And of course there are also serious risks of physical harm to innocent citizens. And patterns of over policing can devastate communities. As Tracy Meares observes, convicted felons are less likely to invest in their own ‘human capital’, perceiving such investment as a waste of time given their criminal record. This frays the relations between convicts and other members of their community because the person with a criminal record

is less capable of benefiting those members (Tracey L. Meares and Kahan 1998).

On the other hand, under-policing can also devastate communities, especially in cases where violent or property crime is rampant. A PAPM system that underrepresents crime in a place could mean failing to send police there when they could have deterred crime. This could constitute a dereliction of duty as well as a violation of the requirement that legitimate police institutions enforce the laws equally.

As we will see below in the “Interaction” section, programmers have access to some methods to ameliorate the risk of harm from police interactions with communities, namely, by suggesting non-enforcement activities to police who are sent to patrol specific areas. This can alter the cost-benefit analysis of over-representation by reducing the burdens of additional police presence on citizens.

- **RECOMMENDATION:** Police departments and developers should be prepared to justify cases in which their model over- or under-represents crime. Developers and departments should incorporate the priorities of community members when calculating costs and benefits of over- and under-representation.

User Interaction

Questions related to the model Interaction are focused on how systems that have been developed via machine learning are used in real context. Once a model has been produced, it is incorporated into a larger command system within a law enforcement agency. The development of the Human/Computer interaction with that model has tremendous impact on the ways in which the system guides human behavior. The PAPM system is going to be used by many different kinds of users, occupying different roles at the agency. If the PAPM system is designed with one set of users in mind, but another set of users ends up using it, then many of the design assumptions about choice structuring and user information needs might not apply. Effectively guiding user interaction requires uncovering the ways that a PAPM system’s predictions are interpreted by users.

The successful implementation of PAPM systems requires officer uptake. Anecdotally, many departments report encountering resistance from officers when officers are asked to rely on the predictions of PAPM systems (Brayne 2020). Whereas commanding officers might have a deep understanding of the PAPM system being used by the department, officers who have years of walking a beat distrust the predictions of computer

systems. Sarah Brayne quotes one LAPD officer as saying about PrePol, “I think that’s just witchcraft” (Brayne 2020, 87).

One challenge is to convince officers on a beat to internalize empirically validated practices like the “Koper curve,” discussed above. The Koper curve is the product of hot spots policing experiments indicating that the positive deterrent effect of police patrols in an area is maximized when police spend about 10-15 minutes patrolling there. Returns begin to diminish greatly after 15 minutes (Koper 1995; Williams and Coupe 2017). However, patrol officers, who are often very busy answering calls for service, and who are not intimately acquainted with criminological theory, might fail to appreciate the crime reduction benefits of observing the Koper Curve. Instead, they are presented with more algorithmically-managed patrols and told this will actually *reduce* their workload. Understandably, this can seem paradoxical. One challenge, then, is to make these crime reduction benefits of the PAPM system clear and tangible to officers using the PAPM system.

While detailed recommendations for such training programs fall outside the scope of this current project, we can offer a couple of modest recommendations.

- ▶ **RECOMMENDATION:** Consider surfacing information to officers to help explain why they are being directed to a particular location, e.g. the factors that contributed to this recommended patrol. A clear explanation makes it easier for officers to understand the reasoning behind the recommended patrols and increases trust in the system, especially in those cases where their “gut” might tell them something different.
- ▶ **RECOMMENDATION:** Consider what information is helpful and easily understandable to officers, and might plausibly nudge their behavior in ways that help serve the goals of the community (Jameson et al. 2014; Johnson et al. 2012). Consider surfacing only that information that could reasonably alter user behavior for the better, e.g. to frame their behavior to minimize the chance of unjustified harm or counter-productive interactions with community members.
- ▶ **RECOMMENDATION:** Identify and address barriers to sustained and continuous use of the application. One such barrier is relying solely on the clients to use the product without proper guidance and accountability from higher-ups, which can result in the product being abandoned as just “one more tool” that came and went. This involves explaining the benefits of the product and the rationale behind its features to all levels of the organization, as well as tracking engagement over time.

Maintain ongoing contact with agencies, including in-person contacts multiple times per year.

- **RECOMMENDATION:** To gain buy-in, understanding, and enthusiasm for the product throughout the organization police departments and developers should work together to collect data about patrol officers' attitudes toward the PAPM system. If the prevailing attitude of officers on patrol is that the PAPM system is a waste of time, it is vital that department leaders know this. This would serve the dual purpose of improving the algorithm and approach based on the 'on the ground' understanding and discoveries of police officers, while also learning how to build more trust in the system by police officers by better presenting information to them.

Failure to address informational framing effects can lead to unexpected and harmful consequences. The early implementation of CompStat by the New York Police Department is a cautionary tale. Because supervisors wanted to demonstrate productivity, they encouraged arrests. This led to the charge that police were harassing communities to drive up the arrest numbers. However, because supervisors also wanted crime to go down, this gave officers an incentive to miscategorize serious crimes as less serious crimes. The CompStat case demonstrates how perverse incentives can lead to data manipulation and aggressive policing, in violation of the legitimate policing requirements of proportionality and competence (Ferguson 2017b, 72).

Consider the contrast between these two cases: (1) telling an officer that an area was selected for patrol because it is near a football stadium, where a game is being played and where many cars have been left unattended; versus (2) telling an officer that the area has seen a rash of armed burglaries in the last week. These two different predictions are likely to "frame" or "prime" the officer's actions in the area toward very different kinds of interactions.

At the stage of police interaction, the expertise and experience of police who are familiar with a patrol must be balanced against validated, data-driven practices. Officers might bristle if they found the tool to be micro-managing or excessively prescriptive in its recommendations. Additionally, officers might be understandably uncomfortable being sent to areas they are less familiar with, i.e. outside their regular patrol.

Telling officers, for example, that the area is "forecasted to experience higher-than-average crime" is unhelpful. Telling them that there were three burglaries recently, at this time of day, in this neighborhood, is clearly

more actionable.¹¹ If the PAPM application considers environmental features of a place when generating crime forecasts, surfacing such information to officers on patrol could plausibly alter the way that an officer enters a location, e.g. rather than assuming that this area has suffered a rash of crimes recently, they may understand that this area contains “environmental aggravators” to crime, such as abandoned buildings or broken street lights.

- **RECOMMENDATION:** Vendors should carefully consider the extent to which the inner workings of their systems are made transparent—or malleable—to officers and analysts. Some explanations of the system’s functioning will be more useful than others.

For example, allowing departments to tweak settings concerning model training, probabilistic sampling, and so on, is likely unnecessary for effective use. Still, vendors must be mindful of the quantity and quality of pressure applied to departments, always maintaining a relationship with departments in the spirit of good faith collaboration—for the health of the community and for the department’s own goals. This task of shepherding begins before product delivery during the design and development phase, and extends after delivery through ongoing engagement.

The very nature of PAPM—that it generates a list of locations that are likely to see an increase in crime—naturally suggests certain interventions, at the very least, in-person patrols or surveillance of an area.¹² Still, officers could perform one of many different interventions.

- **RECOMMENDATION:** The choice of interventions suggested to officers on patrol should take into account the likely effects of different interventions on community relations. This generally counts in favor of non-enforcement activities¹³, such as community contacts, conversations with business owners, or identifying and cataloging features of the physical environment that drive crime. (Enforcement activities, in contrast, would include surveillance or high-visibility presence such

11 Note that not all factors that influence crime are actionable or meaningful. The fact that it is hot outside correlates closely with violent crime; the fact that it is a full moon means there is more light for criminals at night. But neither of these is useful to tell officers on patrol since neither is tied to a specific location.

12 Analogously, person-based systems naturally suggest targeted contacts with those people who are identified as being likely offenders or victims.

13 We characterize enforcement activities as immediately directed at enforcing the law or deterring crime through police presence. Non-enforcement activities would include e.g. conversations with business owners which, in contrast, directly aim at improving community trust.

as parking a patrol vehicle in the area for 10-15 minutes.) Interventions that are less intrusive and burdensome are easier to justify to those who are burdened.¹⁴

- **RECOMMENDATION:** Relatedly, police training should include communications training to enable productive conversations with community members about basic public safety improvements; don't overlook opportunities for impromptu positive interactions with community members.

Once again, the work of providing recommendations here is stymied by a lack of robust data on the success or efficacy of police interventions. It is unclear, for example, whether it is best for police to visit each designated hotspot once per patrol, or twice; once per week or more; how much time to spend in each patrol box; and so on.

- **RECOMMENDATION:** Departments should practice empirically validated interventions in combination with PAMP systems. For example, one empirically validated practice is to observe the “Koper curve,” according to which police presence in an area has a positive deterrent effect on crime up to around the 10-15 minute mark. Returns begin to diminish greatly after 15 minutes (Koper 1995).

Once again, assembling sufficient data and disentangling confounding factors is a vexing challenge. In an ideal world, analysts would need to have access to data on community sentiment¹⁵, the “ground truth” of the occurrence of crime, and other information that is not likely to be forthcoming.

- **RECOMMENDATION:** Technology developers and vendors have an important dual-role to play as both collaborators and shepherds of department operations.¹⁶ In order to facilitate uptake and use of a prod-

¹⁴ A full defense of an account of fairness or justification is outside the scope of this report. But it is worth considering the ability of officers to justify their use of place-based crime forecasting, including its attendant burdens, directly to the people who are burdened. The most promising way of doing this is by pointing to a benefit that *those community members themselves* receive and that outweighs the burden. The upshot of this is that the less severe the enforcement actions that are taken, the easier the justification is to get off the ground.

¹⁵ Even here, “community sentiment” is a capacious term. It could refer, for example, to how safe community members feel in their neighborhood; whether their neighborhood looks safe; how much they trust the police; how likely they are to call the police if they suspect or witness a crime taking place, etc.

¹⁶ One employee we spoke with, who works for a major manufacturer of place-based predictive systems, characterized this balance as between serving as a *disruptor*

uct, vendors obviously must accommodate the unique culture, practices, history, and local knowledge of a department. Early and ongoing conversations between departments and vendors are crucial to maintaining rapport and trust. This includes, for example, conversations early in the relationship about the nature of the tool and the ways it *can* and *should* be used.

For example, vendors might work with police departments to identify tactics that officers have found effective in reducing crime. If those tactics are not otherwise problematic, they could be loaded in “custom” for each department. Current data, unfortunately, is not robust enough to draw novel conclusions about correlations between tactics and crime, e.g. that embracing a certain tactic could be expected to lower crime in an area. These data would amount to a kind of Holy Grail: drawing conclusions, bespoke to a particular district or locale, about which tactics are most effective at deterring certain kinds of crime.

“Train, train, train—all up and down the command chain.”

- ▶ **RECOMMENDATION:** In a regulatory environment where vendors enjoy significant discretion about the design and application of PAPM technology, vendors have a duty to endorse responsible use and discourage irresponsible use of their technology, uses that might e.g. lead to over-policing, target crimes that are a poor fit for the technology, or otherwise erode relations with the community. This will require some basic education of police officers about the nature of machine learning, how the system works, and which theories of crime inform model training.
- ▶ **RECOMMENDATION:** After transferring the product, ongoing conversations should center around police use of the technology, e.g. whether police are faithful to the patrol recommendations of the model and whether there has been sufficient and appropriate uptake within the agency.
- ▶ **RECOMMENDATION:** Vendors should carefully consider, through internal dialogue and engagement with stakeholders, what uses of their technology would constitute “red lines,” e.g. exacerbating over-policing or policing nuisance crimes that are highly susceptible to biased enforcement.

versus as a *facilitator* of police practices.

- ▶ **RECOMMENDATION:** Other ways of discouraging certain uses could be “designed into” the product itself, such as limiting the choice of which tactics to surface to police on patrol.
- ▶ **RECOMMENDATION:** As part of the standard onboarding process for agency clients, include a set of general orders that agencies can easily incorporate. These orders should include information such as the division of responsibilities along the chain of command; the set of interventions that have been judged appropriate after consultation with the community; and a reminder to officers that not simply being in a high-risk zone is not sufficient to treat someone as suspicious.

Consider one particularly poignant and fruitful case: whether to allow officers or analysts to “override” the predictions of a PAPM system by placing their own “ad hoc box” to direct officers to a particular location.¹⁷ Departments may have good reasons, inaccessible to the model, to send officers to a location: perhaps there is a festival in town; perhaps there is intel about inter-gang retaliation. At the same time, this feature is vulnerable to abuse, for example, placing an ad hoc box permanently on a multi-family housing unit for the purposes of harassment or because of mere suspicion. Moreover, the underlying paradox is that place-based predictive systems are intended to replace the *status quo ante* of human crime analysts placing boxes on a map manually. What can be done?

- ▶ **RECOMMENDATION:** Vendors must not just accept but embrace their role as choice architects in these systems which have clear downstream effects on the rights and wellbeing of citizens. Explore options to facilitate reasonable customization of technology tools, while giving clear prescriptions when onboarding users to avoid problematic uses.

For example, consider: publicizing ad hoc boxes and rendering them not just transparent but highlighted on analyst reports; requiring additional “signatures” from those further up the command chain to place or renew an ad hoc box; and monitoring and reporting for abnormally high numbers of ad hoc boxes in an area or over time.

Over-saturation and Feedback Loops

A combination of system design and user interaction can lead to “over-saturation” as it is sometimes called. Because PAPM models are often

¹⁷ This issue surfaced during an interdisciplinary discussion at an NSF-sponsored workshop at Northwestern University in 2022.

used to predict the location of crimes which are themselves geographically correlated, it is possible that these models might direct police officers to patrol the same areas repeatedly. This leads to concerns about “over-policing,” or drastically and disproportionately concentrating police presence in small areas. Any form of over-saturation risks running afoul of the normative commitment to impartiality in policing, and it threatens proportionate policing by amplifying the police response to crime in an area beyond what is reasonable. Furthermore, insofar as over-saturation leads to unwanted and excessive police attention, it can also threaten the willing cooperation of the public.

“Vendors must embrace their role as choice architects in these systems which have clear downstream effects on the rights and wellbeing of citizens.”

Feedback Loops

One way that oversaturation can occur is through a “feedback loop” (Ensign et al. 2018; Lum and Isaac 2016).¹⁸ The prospect that PAPM systems will lead to feedback loops is perhaps the most widely-known criticism of PAPM systems. Thus it merits special attention by developers and users of PAPM systems, as well as community advocates. Consider a story of how PAPM could drive a feedback loop, even when the initial training data is not racially biased:

- (1) Imagine a non-racially biased model, trained on non-racially biased arrest data, which accurately shows that the majority of crimes of interest occur in neighborhood X.
- (2) As a result, the model recommends that police spend a great deal of time patrolling neighborhood X.
- (3) Because of these patrols, police encounter more suspicious behavior, perform more stops, have more interactions with community members, conduct more traffic stops, and so on, in that neighborhood. Because police spend more time in this neighborhood, they also effect more arrests in neighborhood X.
- (4) Those data on arrests, community interactions, traffic stops, etc., are then used to train the next iteration of the model.

¹⁸ Some feedback loops are good for crime prevention. For example, if a policing practice is effective at deterring crime, this can improve community relations, leading to more cooperation from citizens, leading to higher case closure rates, leading to a further improvement in community relations, and so on.

- (5) Because police spend most of their time in neighborhood X, that neighborhood is over-represented in the training data that are used to retrain the model.
- (6) This results in a new round of forecasts which recommend that police spend even more time in neighborhood X, and so on. If this process continues, police could be expected to spend more and more time patrolling neighborhood X. This “ratcheting up” effect risks negative community interactions, escalation, harms to innocent people, false arrests, and so on (Harcourt 2006).




Figure 1. The development of a feedback loop, even supposing that the crime forecasting model begins with a non-biased set of training data.

To be clear, not all feedback loops are problematic. For example, we can imagine a “virtuous” feedback loop by which police successfully deter crime in an area, thereby improving community relations, which leads to more cooperation by the public with police, which leads to greater deterrence, and so on. The public concern over feedback loops thus seems to be driven by a conjunction of two claims:

- (1) Any use of place-based PAPM technologies inevitably leads to feedback loops which lead to repeated visits by police to an area; and

- (2) Repeated visits by police to an area are per se problematic.

Neither of these claims, as stated above, is plausible. As we will see, there are multiple ways of diluting model training data and nudging police interventions that are much less likely to generate feedback loops, and police can have good reason to repeatedly visit the same areas of a city, namely, if those areas are, in fact, at the highest risk of crime. However, weakened versions of these claims are more plausible:

- (1*) The use of PAPM technologies *driven by certain data collection and modeling practices* risks the creation of feedback loops;
- (2*) Repeated visits by police to an area are problematic when conducted without good reason or when police employ problematic tactics while on patrol, because they risk harming members of policed communities and undermining community trust.

Consider in further detail the reasons why feedback loops might be problematic: First, by concentrating police attention in one area, they risk concentrating attendant harms from police-community interactions. Because demographics and geographics are often highly correlated, this also risks concentrating harms not just in particular areas but in particular demographic populations, oftentimes groups such as ethnic or racial minorities and the poor who have historically faced discrimination. Note that while this is happening, feedback loops will generate *inaccurate predictions* by trapping a model in a local maximum and leading police to overlook other areas of concern.

Lastly, feedback loops can have adverse effects on community relations, depending on the kind of crimes predicted by the algorithmic system in combination with tactics employed by police. Survey evidence suggests that formal contacts with police decrease community approval of police, whereas informal contacts increase community approval (Maxson, Henigan, and Sloane 2003). Formal contacts include residents' calls to police stations requesting service, police questioning of residents regarding possible crimes, as well as arrests. Informal contacts include police participating in community meetings, increasing officers' visibility in neighborhoods, and talking with citizens. Using PAPM systems to target crimes that tend to be discoverable by police on patrol (e.g., nuisance crimes like loitering or drug possession) is likely to increase formal contacts with police by members of a community when officers uncover these crimes in progress.

In this way, targeting such crimes could adversely affect police-community relations.¹⁹

To address concerns about oversaturation:

- ▶ **RECOMMENDATION:** Patrol recommendations should change often, e.g. at least every day. Still, this could result in officers spending time in a relatively small location over several days or weeks.
- ▶ **RECOMMENDATION:** Police departments should consider keeping a record of where models are directing police patrols and to attach a cost to sending police patrols to an area where they have recently been. This added cost could decay over time and with distance from the original patrol. This can be expected to reduce the amount of time that police spend in one small area and disperse patrols more evenly throughout the community. This cost function could be “tuned” to maintain the deterrent effect of police presence in the area so as to avoid underpolicing of an objectively high risk location.²⁰
- ▶ **RECOMMENDATION:** In addition, departments should consider implementing controls that discourage police from spending too much time in an area in any one shift. Research on the “Koper Curve” suggests that police efficacy declines after about 10–15 minutes spent in an area. Consider timing police presence in an area and encouraging officers to move along after 15 minutes have elapsed. Consider sharing this information with community stakeholders—as they might prefer that police spend more or less time in their area, too.
- ▶ **RECOMMENDATION:** If instead the use of the PAPM system increases *informal* contacts with police by members of the community, this could improve police-community relations.²¹ Whether this benefit can

¹⁹ The reader might be surprised that we do not include the discussion of over-saturation and feedback loops in the “Bias Questions” section. The reason for this is that while feedback loops can lead to racially unequal distributions of police patrols, these racially unequal distributions need not be a function of any racial bias in the data set on which the model is trained.

²⁰ Note the relevance here of the multi-armed bandit problem in probability theory. The multi-armed bandit is a problem of distributing finite resources among multiple courses of action whose expected utilities are unknown, i.e. until further resources are allocated to those courses of action. See (Vermorel and Mohri 2005) for a canonical, recent treatment of the problem. For an application of the problem to law enforcement, see (Akin 2017).

²¹ One example of this kind of informal contact that was suggested to us was for officers to be on the lookout for home- or business owners that are engaged in behavior

be expected depends heavily on enforcement priorities and existing community-police relations.

There are many ways to interrupt the so-called “feedback loop.” Consider interventions at some of the steps listed above.

- **RECOMMENDATION:** Data on interactions with community members—for example, traffic stops, officer-initiated contacts and arrests, and officer-discovered crimes—can produce a feedback loop. When deciding *which* crimes to predict in the first place, as noted above, data on some crimes are more likely to be infected by “enforcement bias” than others. Choosing crime data that are community-generated rather than officer-generated is perhaps the most significant single intervention that could be taken. However, community-generated data is not a panacea, because it can also be biased. For example, in communities with low levels of trust of police, crime reporting rates might not reflect actual crime rates. Therefore, the solution is not to prefer community-generated data *per se*, but instead to audit and examine all data used in generating crime forecasts.

- **RECOMMENDATION:** Developers should consider including additional sources of data that are not directly influenced by the behavior of police on patrol. For example, some PAPM systems consider environmental factors, such as proximity to and number of liquor stores, bus stops, graffiti, and broken windows. Data about environmental factors may have predictive value for forecasting crime, and because they are *not* influenced by the behavior of police on patrol, they can be used in model training to “dilute” the influence of bias in historical crime data. If data about environmental factors is given significant weight in generating forecasts, this can reduce the influence of bias in historical crime data.

“Data on interactions with community members—for example, traffic stops, officer-initiated contacts and arrests, and officer-discovered crimes—can produce a feedback loop.”

Though they are not susceptible to influence by officer behavior, PAPM systems that consider environmental features could also end up creating

that make themselves vulnerable to crime. Impromptu interactions, for example, an officer who notifies a homeowner that their garage door or side gate is open, can help build community trust. The IACP recommends other impromptu interactions such as joining pickup basketball or football games and visiting community events (IACP 2018).

feedback loops, albeit in a different way. This is because, while crime is dynamic, environmental features are static—they do not change over time (at least, not very quickly). For example, bus stops, liquor stores, and abandoned buildings might remain the same for years. A model that forecasts crime partly on the basis of these environmental features will tend to assign a higher risk score to locations that have these environmental features, independently of how much crime is being reported or recorded in those locations.²² The likelihood that patrols will be sent to a place then depends either on how common crime is in that (or surrounding) location and on how much the location’s environmental features resemble the features of other locations where crime is common. This means that if unchecked, and depending on the influence of the environmental features in the PAPM model, officers could be sent to a specific location repeatedly, out of proportion to the actual crime rate at that location.²³ When environmental factors that are predictive of crime are correlated with race—for example, if liquor stores are more densely concentrated in minority communities—this “environmental feedback loop” can be cause for concern about fair and equal enforcement of the laws.

The elasticity of crime

Bernard Harcourt has raised a concern for statistical profiling on the grounds that members of different groups have different elasticities when it comes to police deterrence of criminal activity (Harcourt 2014). The elasticity of a group with respect to criminal activity is a function of their sensitivity to disincentives. It is possible that the criminal behavior of people with low socioeconomic status (SES) will be less responsive to disincentives (e.g., stricter enforcement) than high SES groups, because their alternatives to crime are significantly less appealing than the alternatives for high-SES people. If this were true, then disproportionately using police resources

22 To invoke a distinction from philosophy, this kind of feedback loop occurs for *types* of neighborhoods, rather than *token* neighborhoods. Individual neighborhoods could be affected because they are tokens of a *type* that has been diagnosed as liable to experience higher rates of crime, even if the individual neighborhood itself defies the generalization. (For more on the type-token distinction, see (Wetzel 2018), especially some examples in section 2.3, “Science and Everyday Use.”) Note that this possibility is speculative and remains a topic of ongoing investigation by the authors.

23 Note that we are not claiming that crime remains static because environmental factors remain static. Criminological studies suggest that is not the case (Hatten and Piza 2021). Rather, the claim is that *predictions* about crime might remain static insofar as they are sensitive to static environmental factors, and thus fail to reflect the dynamicity of crime.

to target high-crime, low-SES groups with greater police attention might actually increase crime.

It is canon in criminology that some crimes track SES. When low-SES communities have lower elasticities but higher overall offense rates, singling them out for extra policing could actually decrease societal well-being. In other words, if the base rate of criminality is higher in a low-SES community, it still could be the case that focusing on that community enables members of other communities to engage in more crime than is prevented. High-SES people might feel immune to police intervention, which might increase their likelihood of engaging in criminal behavior (Harcourt 2014, 304). Harcourt argues that, depending on the elasticities of different groups, police patrols should be randomly distributed instead of being allocated on the basis of data about criminal activity by a person or at a place.

Harcourt's concern can be extrapolated to PAPM. PAPM systems are developed to forecast locations at high risk of crime, but they are not sensitive to elasticity in criminal behavior. It could therefore turn out that the crime behavior of populations at forecasted high-crime locations is less elastic than the crime behavior at forecasted low-crime locations. In this case, allocating more police resources in the forecasted high-crime locations, while allocating fewer resources in forecasted low-crime locations, may be counterproductive.

- **RECOMMENDATION:** In light of the dual phenomena of feedback loops and the elasticity of crime, police agencies should check recommendations for their geographic distribution, as mentioned above, to ensure that certain neighborhoods are not overrepresented in the recommended patrols. A best practice would be to “scramble” or inject a stochastic factor into the recommendations to diversify the location of patrols, displaying to police a selection of those locations forecast to be “high risk.” This will ensure that police do not end up visiting the same locations over and over again if those locations are always predicted to be at the highest risk. Because of elasticity, scrambling forecasts might achieve an optimal level of *efficacy* while simultaneously curbing *over-policing*. Conceivably, scrambling predictions could result in directing patrols away from the areas forecast as most likely to experience crime, but this may be worth the added value in police-community relations.
- **RECOMMENDATION:** A second strategy, complementary with the first, would be to “penalize oversaturation” of forecasted patrols. This would require tracking the frequency of visits to particular boxes, along with the intensity of intervention, and allowing areas to “cool off” before

recommending patrols to those same areas. This recommendation has limits; if the model is *very confident* about the likelihood of crime in an area, and especially if the predicted crime is *very serious*, then this confidence should arguably override the cost of oversaturating an area with patrols.

- **RECOMMENDATION:** Alternatively, developers can “gamify” the process of patrolling in order to incentivize officers to visit different areas, including areas they are unfamiliar with. This approach rewards officers for diversifying their patrol choices—for example, through internal “leaderboards,” praise, and recognition—as an alternative to filtering out over-saturated areas through model training and outputs. But beware: gamification can create perverse incentives that nurture harmful habits in the name of “pumping up” stats like arrests (Giacalone and Vitale 2017).

While officer compliance is largely up to the police departments’ organizational culture, gamification can encourage officer buy-in and to follow the model’s recommendations in a productive way. Unfortunately, place-based predictive policing may not be able to avoid the potential for over-saturation entirely. There are only so many areas in a city that officers can patrol, and as long as crime tends to be concentrated in certain areas, police are liable to return to those areas repeatedly. As a result, those areas might be subject to over-saturation even when the predictive model avoids a feedback loop.

- **RECOMMENDATION:** Include a layer of analysis on top of the geographic predictions of forecasting models that measures the concentration and frequency of police patrols in certain neighborhoods to protect against over-policing those areas.
- **RECOMMENDATION:** What police do (or don’t do) while on patrol can help to minimize potential feedback loops while also minimizing adverse interactions with the public. A PAPM system can suggest behaviors to officers on patrol that are less likely to lead to a feedback loop. A system trained on traffic stops, for example, which recommends that police head to a certain location to perform more traffic stops, is likely to create a feedback loop. Instead, consider recommending that police park in a visible location, which may have much the same deterrent effect on speeding as performing traffic stops, without generating more self-fulfilling data. Carefully consider which data are fed back to retrain the predictive model. Consider preventing data generated by police on recommended patrols from being used as training data in the next generation of predictions.

Notice, however, that using data about officer-discovered crime has at least two benefits. First, collecting data about the crimes police encounter while on recommended patrols could help determine whether the model is working. For example, is it really the case that, among areas where the model predicted that there was a 20% chance of a crime occurring, police witnessed a crime roughly 20% of the time? Second, if the police do encounter crimes while on patrol in those areas, then that historical crime data does, certainly, have some value in predicting where crime is likely to occur in a neighborhood. Excluding this data source could degrade the model's predictive utility.

- **RECOMMENDATION:** Monitor the recommended patrol locations over time so that the system does not end up recommending that police patrol the same locations repeatedly over days or weeks, and that they spend no more time than is necessary patrolling recommended areas.

Going beyond patrol allocation

PAPM systems that incorporate a variety of data sources, including crime data, census data, and environmental features, have potential utility beyond their use for patrol allocation. As Andrew Ferguson has argued, “A predictive policing algorithm can forecast a particular likelihood of crime in a particular place. But identifying risky places does not determine the appropriate remedy to fix the crime problem.” (Ferguson 2017b, 167). Once a risky place has been identified, a city should be open to a host of solutions, many of which do not involve enforcement-oriented policing. If a PAPM system predicts a rash of car thefts in a specific parking lot in a multi-family housing complex, sending police there on patrol is but one option. The complex could also install surveillance cameras to deter theft, improve lighting if thefts tend to occur in the dark, or hire a security guard. In reality, the best solutions to crime will be multi-faceted, involving a combination of policing and non-policing solutions.

Ameliorating the physical features of places that make them prone to crime can prevent crime before it occurs, reducing the need for enforcement-oriented tactics such as saturating an area with police patrols. Often, however, addressing the physical features of places that make them vulnerable to crime is outside the immediate purview of law enforcement. Community groups and non-law enforcement agencies often have resources the police department does not have that may well aid in crime prevention. Interventions should therefore be designed with groups/stakeholders and their respective resources and abilities in mind. Implementing non-enforcement solutions to crime therefore requires fluid collaboration between police, community groups, and other city agencies such as public works.

- **RECOMMENDATION:** Looking ahead, developers and police departments should remain open to the potential to assist with non-policing solutions to crime by sharing their data with non-law enforcement agencies and community groups. For example, modeling can be performed as a diagnostic—that is, it can help to identify physical features of places (e.g., dim street lighting) that make them prone to crime.
- **RECOMMENDATION:** When possible, technology developers and police departments should pursue alternative opportunities to put PAPM to use in the service of crime reduction. Taking advantage of these opportunities will require maintaining open channels of communication between police departments and a variety of non-law enforcement agencies.

Another positive development in law enforcement is the pairing of police intervention with mental health clinicians to better assist individuals suffering with mental illness. When officers respond to a call for service involving a person with mental illness, they can be ill-prepared to deescalate the situation or to find the help that the individual needs. This can be a drain on human resources and lead to adverse interactions between police and citizens. As a result, police departments around the country are teaming with mental health professionals to improve emergency response to calls for service involving individuals with mental illness. Any PAPM system should be compatible with the growing role of non-law enforcement professionals in aiding with emergency response to calls for service (Abramson 2021).

- **RECOMMENDATION:** Developers of PAPM systems should seek opportunities to assist with new approaches to law enforcement, such as growing collaborative partnerships between police and mental health professionals. For example, developers might explore whether it is possible to generate reliable recommendations about the type of emergency response a dispatcher should demand: police only, mental health professional only, or both.

Another way that PAPM systems can go beyond mere patrol allocation is to forecast expected safety risk for officers on patrol. In their 2019 report ‘Law Enforcement Officers Killed and Assaulted’ (FBI 2019) the Federal Bureau of Investigation noted that, in that same year, of the 475,848 law enforcement officers about whom data was collected 11.8%, or 56,034 officers, were assaulted (FBI, 2019, ‘Officers Assaulted’) and 48, of whom 8 were Black/African-American or Asian, were killed (‘Officers Feloniously Killed’) while on duty. Most relevantly for this paper of those who were killed, “30 officers were on assigned vehicle patrol when the felonious incidents occurred” (‘Officers Feloniously Killed’), 3 officers in “Two-officer

patrol” and 27 officers in “One-officer patrol” (Table 21); and of those who were assaulted 62.1% “were assigned to 1-officer vehicle patrols” and 17.3% to “2-officer vehicle patrols” (‘Officers Assaulted’), or a combined total of 44473 officers assaulted (‘Table 84’). Therefore the vast majority of law enforcement officers who were killed or assaulted, were killed or assaulted during some type of patrol. Within the framework of this paper, certain recommendations to PAPM system developers suggest themselves in response to these facts.

There is a difference between a location’s being predicted to be, on the one hand, either high-risk or low-risk for crime and, on the other hand, its being predicted to be high-risk or low-risk of *harm* to the patrolling officer. These two features of a place can come apart, because an area can have a high frequency for crime while being low-risk of harm to the patrolling officer. At the same time an area can have a low-frequency for crime but high-risk of harm to the patrolling officer. In the latter case, we have scenarios such as those resulting in increased police officer casualties, but also in community member casualties due to police overreaction to an anticipated ambush. Accurately predicting the safety risk for patrol officers promises a variety of benefits: first, frequent police casualties in the line of duty can exacerbate already dire recruitment problems facing American police departments. When recruitment suffers, communities suffer as well by becoming underserved by law enforcement. Second, accurately predicting safety risk avoids sending gratuitously armed officers to high-crime-risk but low-harm-to-officer-risk areas. This can help address community concerns that police patrols are ‘militarized’ or unnecessarily threatening in number or appearance.

A potential cost of estimating safety risk for patrolling officers is that it might lead to officers avoiding high-risk-to-officer areas, irrespective of crime frequency, leading to these communities becoming underserved. This cost must be balanced against the potential benefits of adding this product feature.

- **RECOMMENDATION:** PAPM system developers should explore crime models which account for data about law enforcement officer deaths and assaults, focusing on the locations, methods of assault, and the physical characteristics of the environment and the time of day, while avoiding data that relate to the protected characteristics or identities of the assailants.
- **RECOMMENDATION:** Developers should explore the potential of PAPM systems that can estimate the safety risk facing patrol officers during a recommended patrol alongside other metrics to ensure a proportionate response.

Data that has already been gathered about the circumstances surrounding these officer's deaths and assaults include, but are not limited to, the hour of the day (FBI, 2019, 'Table 21'), day of the week ('Table 6') and of the month ('Table 9'); "Lighting" and "Weather/environmental" conditions and location of the incident, such as whether it was inside of a structure or outside, and whether the structure was say commercial, residential or government ('Table 3').

Organizational Ethics

Many of the ethical concerns we have discussed here center on choices made within an organization, or in conversations between a vendor and a law enforcement agency customer. As a result, ethical foresight, and early, ongoing ethical deliberation within a technology vendor is likely to identify and head off many of these concerns. Developing a culture of seriousness and thoughtfulness about ethics within a technology company is one of the single most impactful changes that we can recommend. Developers should consider taking the following steps to develop a culture of seriousness and thoughtfulness about ethics.

- **RECOMMENDATION:** Formalize a position of Chief Ethics Officer with clearly stated responsibilities for forecasting and anticipating ethical concerns with their products. This position may be separate from or combined with a role charged with overseeing legal compliance and best practices for data handling. Employees should have multiple informal channels for surfacing ethical concerns confidentially, e.g. through an anonymous tip line, through dedicated internal message boards, or during open office hours with the Chief Ethics Officer.

Recognize that decisions made during the process of development carry the potential for ethically significant impacts. By establishing a Chief Ethics Officer, a developer demonstrates that they are treating ethics with the seriousness it deserves—e.g. on a par with legal compliance and other C-level responsibilities. This role could focus on forecasting potential ethical implications of the development and deployment of products and identifying ways to address them. Ensure that employees' ethical concerns are welcomed and valued.

While the establishment of a Chief Ethics Officer is a valuable step, it should not be the sole mechanism for ensuring ethical practices within an organization. The presence of “ethics champions” throughout the organization—individuals who are passionate about upholding ethical standards and can motivate others to do the same—can improve the effectiveness of ethical foresight and oversight.

- **RECOMMENDATION:** Develop statements of values or ethical commitments that the organization stands behind. Publicize these internally, at the very least, but ideally to the public, as far as is practical. Consider

adopting reasoned and consistent positions, for example, on which crimes will and will not be modeled and why. Reflecting on examples of tradeoffs or tensions that employees might be expected to navigate, or which they have navigated in the past can be helpful. Developing these statements is in itself a valuable “soul-searching” exercise likely to reveal tensions within the organization, but also illuminate firm consensus on other matters. Be prepared to defend these decisions to potential clients.

A clearly articulated statement of value commitments can provide anchor ethical decision-making within the organization. By developing and sharing these commitments, developers establish a set of ethical benchmarks to which to hold themselves accountable. This also helps employees understand the company’s stance on critical ethical matters, ensuring alignment in decisions, and facilitating decision making from a place of *principled consistency*. Furthermore, being transparent about the company’s ethical commitments can enhance trust with the public, showing that the organization is dedicated to upholding ethical standards.

“Be prepared to defend these decisions to potential clients.”

- **RECOMMENDATION:** Incorporate a session on ethical impacts and ethical product design into employee onboarding and orientation. This session should introduce the principles and practices of ethical decision-making that underpin the company’s work, the values statements or commitments, and channels for airing ethical concerns internally. Alongside more mundane information—e.g., where the bathrooms are located—these sessions would establish a firm foundation of ethical engagement from the start of an employee’s tenure within the organization.

Introducing ethics discussions from the outset signals the importance of ethical considerations within the organization and communicates that a serious attitude towards ethics is expected of everyone. These sessions can foster an understanding of the company’s ethical values and expectations, setting the tone early for employees’ attitudes towards ethical decision-making. Importantly, this ongoing engagement emphasizes that ethics isn’t a one-off task, but a continuous process and integral part of the employee’s work. Ensuring that this is understood from the outset can help to embed ethical thinking into the organizational culture and workflow, promoting ethical behavior and decision-making as the norm, rather than the exception.

- **RECOMMENDATION:** Normalize internal conversations about ethical concerns with technology development and use. Have frequent con-

versations, both in formal and informal settings, considering questions such as, “Is this the right thing to do?”, “Is there a less ethically risky way to do this?”, “How can we design this tool to discourage certain uses?”, “How might this choice impact the most vulnerable populations affected?” and “Is the ethical risk of this choice worth the social benefit?” Consider scheduling regular time for discussion and trainings on ethical reflection to communicate the seriousness of the responsibility—otherwise these discussions risk becoming perfunctory “box-ticking exercises.”

Regular, open discussions about ethics can help to promote a proactive culture of ethical awareness. Too often, employees feel that ethics is “not their job”—but responsible development requires that this commitment permeate the organization, making it clear that *ethics is everyone’s job*. When ethical concerns are seen as a normal part of conversations, they’re more likely to be thought about and addressed effectively. This also encourages the development of ethically sound solutions, as the team considers potential impacts on vulnerable populations and seeks to balance ethical risk with social benefit. In addition, by encouraging a dialogue about ethics, employees are more likely to feel comfortable raising their own ethical concerns.

“When ethical concerns are seen as a normal part of conversations, they’re more likely to be thought about and addressed effectively.”

- **RECOMMENDATION:** Include “ethics checks” in the product development process. For teams that use the Agile development process, consider, for example, including “consequence scanning” events in their workflow (DotEveryone n.d.).

Integrating ethical considerations into the routine practices of development ensures that potential issues are addressed proactively and systematically. Ethics checks can involve assessing potential negative consequences and considering whether there are ways to mitigate them. With Agile development’s iterative nature, “consequence scanning” at regular intervals provides opportunities to identify and address new ethical issues that emerge as the product evolves.

- **RECOMMENDATION:** Include ethical requirements in product specifications. For example, specify and continually revisit concrete requirements for acceptable levels of bias, transparency, and explainability. Consider ethical performance to be one more non-functional requirement of product design, alongside other metrics such as user experience, speed, computational cost, etc.

By treating ethical considerations as an integral part of product specifications—e.g. alongside the typical SWaP-C requirements or usability—technologists can establish a concrete framework to guide responsible development. (Moreover, this makes ethical requirements amenable to incorporation into multi-criteria decision analysis techniques, which is an active area of research in applied technology ethics (van de Poel 2015).) This strategy also ensures that ethical requirements such as bias minimization, transparency, and explainability are not afterthoughts but core considerations from the inception of the product throughout its development life-cycle.

“Develop statements of ethical commitments that the organization stands behind.”

Conclusion and Next Steps

This framework is founded on the conviction that a core goal of policing is to prevent crime and disorder, ideally by preventative means, while respecting citizens' constitutional rights and preserving the conditions required for legitimate policing. Legitimate policing requires sufficient competence on the part of police, proportionality in the response to crime, equality in the enforcement of the law, and public trust. We have demonstrated that the irresponsible development or deployment of PAPM systems can threaten this core goal of law enforcement, which in turn challenges police departments' institutional legitimacy to operate and undermines community trust. Managing these threats is key to enjoying the crime prevention benefits of PAPM while maintaining the integrity of our law enforcement agencies.

“Think of the scariest thing you’ve seen in a science fiction movie and that’s probably what’s coming.”

The challenges that we have emphasized in this report cut across three broad themes:

- how to mitigate bias and oversaturation of communities with police patrols;
- how to use PAPM systems in a way that cultivates community trust in law enforcement;
- and how to use PAPM systems in combination with policing tactics that minimize harm to citizens;

We have argued that addressing these challenges requires active attention to best practices in responsible AI development, including:

- close and ongoing collaboration between developers, law enforcement agencies, and policed communities throughout the technology life cycle;
- the careful integration of PAPM into police procedures; honoring community priorities in the choice of crimes to model;
- an active and mindful role for developers in designing the choice architecture of PAPM systems, paying attention to the nature of interventions that are surfaced to officers on patrol;

- thinking creatively about how to pair the forecasts of PAPM systems with non-enforcement oriented interventions, including problem-oriented, community-led crime solutions and involving mental health professionals;
- nurturing a culture of ethical sensitivity and seriousness within firms developing PAPM products.

There remain a number of important challenges that we have not addressed, including the development of new legal regimes for pass-through data collection, enhancing community oversight and collaboration, and achieving an industry standard for “fair” machine learning. We briefly describe these challenges below.

New legal regimes for pass-through data collection

According to case law on the 4th amendment, the US government is restricted from collecting certain types of information about its citizens. However, there do not seem to be legal restrictions on the *purchase* of individuals’ data, and many companies are legally allowed to collect data on citizens and sell it to the government (Cameron 2023). Furthermore, the law does not restrict the types of analysis that the government can use on whatever data it obtains. If the public has concerns about the government’s acquisition of data or its methods of analysis, laws would need to be put into place to restrict them (Solove 2001, 1137–40).

Increasingly, private companies are collecting, integrating, and selling far greater amounts of data, so much so that were the government to collect data in this fashion it would be acting illegally..

This raises important questions about the nature of this data collection and subsequent sale. Namely the question of whether it should be legal for these companies to collect this data with an eye toward selling it to a law enforcement agency? Moreover, should it be illegal for law enforcement agencies to purchase and use data that it would be illegal for them to collect on their own? Finally do the 4th amendment protections against search and seizure apply to the act of collecting the data?

This is the kind of outsourced collection and analysis, or “pass through” data collection and analysis wherein commercial firms analyze data and police ‘subscribe’ to the data.) This is akin to *laundering* of data analysis to get around 4th amendment protections, while also using public funding to collect data about the public. Further investigations in this topic will need to address the question of whether our commitment to privacy, which undergirds the fourth amendment, also suggests a restriction on police buying data that they would not be permitted to collect themselves.

Enhanced community oversight and collaboration

Our recommendations have urged greater citizen involvement in decisions about APM systems being deployed in their communities. There remains an important question about whether, and the extent to which, this involvement should take the form of *oversight*. Police departments and community members face many of the same challenges in exercising effective oversight of technology. Among other challenges, they lack the institutional resources, they lack the technical expertise in machine learning, and they are unable to easily interpret the data until developers facilitate transparent reporting and analysis. Simply put: it is expensive and difficult for communities and law enforcement to employ effective oversight mechanisms, and they cannot do it without the assistance of developers. With that said: What can technology companies do to facilitate oversight, accountability, analysis of the way the technology is being used and influencing policing operations? Using some of the technology discussed in this report and programs like SoundThinking could be one way to give communities more control, oversight, and knowledge of police priorities and behaviors.

Moreover, there are ways to build models which could involve using the same basic data about the timing and location of crime and other environmental factors, but which could diagnose underlying social issues plaguing a community as well as problematic or risky behavior by police on patrol (Carton et al. 2016)—Andrew Ferguson refers to these uses of data as “blue data” and “bright data.” Expanding the use of crime models to address the social drivers of crime as well as dangerous police behavior is a key step towards a holistic data-driven approach to crime prevention that promotes trust in policing (Ferguson 2017: 143 and 167).

Many of the recommendations contained in this document can facilitate increased community oversight and promote community approval. This can be facilitated by partnering with city agencies to increase social and other investments as well as by employing liaisons or customer success representatives to work with communities that can act as intermediaries and translators to explain the workings, impacts, and analyses of the technology.

These liaisons can further facilitate transparency through regular audits and reports. Regular audits would allow community members to inspect, correct, and contest the sources of data that go into model-building. The ability of community members to fully comprehend the inner workings of these programs is almost certainly impossible given the complexity of the model and the variety of the sources of data. Companies can, however, facilitate this and help to “make the ends meet” by bringing some

of the operations and outputs of the models down to a level a layperson could understand. This is closely related to recent work on reducing the “FATE” principles of AI—fairness, accountability, transparency, and ethics (Fjeld et al. 2020)—to a single overarching concern such as “contestability.” Contestability is treated as a gold standard in AI ethics, either as a *sui generis* requirement, as a fifth principle alongside FATE, or else as a broader overarching meta-principle that itself undergirds and generates the FATE principles (Lyons, Velloso, and Miller 2021b; 2021a; Henin and Le Métayer 2022)

Another option for front-end accountability is a certification process like that currently done by UL Solutions (formerly Underwriters Laboratories). Consider the process of a UL certification. In that context, brands take on the cost of safety testing in order to be licensed to apply the UL mark to their product. This in turn certifies them to sell that product in a certain market. Representatives of companies involved in the development of policing technologies agreed that the UL process is both quite expensive and too slow for the fast pace of development in the ML space. A major drawback is that this kind of auditing and external validation is “frickin’ expensive,” as one of our workshop participants put it, and does not move very fast. But if the market is there and the pressure from community groups is there, this could be an effective prototype to meet the demand for this kind of external validation. An independent body such as UL, because they are an independent body with less direct profit motives than developers of the technology themselves, could be better placed than companies themselves to serve community interests, and could be more efficient at assessment than government agencies.

Fairness in Machine Learning

There are multiple statistical measures of fairness in algorithmic classification systems (Berk et al. 2021; Verma and Rubin 2018). All of these measures are reasonable on their face—yet some of them are mutually inconsistent (Kleinberg, Mullainathan, and Raghavan 2016; Chouldechova 2017). Which measures of fairness seem most appropriate for place-based predictions?

A ‘biased algorithm’ is usually thought to be in some way unfair or discriminatory to the disadvantaged group. Bias in an algorithm can come from data sampling, data labeling, its decision structure, or its deployment structure, but statistical measures are usually going to focus on the first two of these sources of bias. The data that is used to train a model could produce bias in an algorithm because it is a misrepresentative sample of the facts on the ground. Both officer-discovered and community-reported crime data can produce bias in this way. Once data is being used by the

algorithm, bias can be exacerbated by classic ratcheting/looping/feedback effects, as, for instance, when police patrols are sent to areas with high reports and therefore discover more criminal activity in those areas, data about which goes back into the algorithm. We described this phenomenon in detail in this report.

Machine Learning (ML) literature has identified several standards that can be used to assess fairness in algorithms (Corbett-Davies and Goel 2018b). A ‘same treatment’ standard would require that different groups be treated equally in the sense that data about any protected features such as race be dropped from the dataset. According to this approach, race blindness means racial equality. A clear drawback of this measure is that it simply does away with bias by brute force and tends to undercut the usefulness of an algorithm’s results. This approach also ignores the influence that features that are proxies for protected features can have on algorithmic classification. Calibration is a standard requiring that for each group the percentage of individuals who receive a positive or negative classification matches that group’s base rate. Equalized odds is a standard requiring that false positive and/or false negative rates are equal for all protected groups. Satisfying both calibration and equalized odds is incompatible in many real-world cases (Chouldechova 2017, 153–63). Lastly, demographic parity is a standard that requires that equal percentages of individuals from all groups benefit from the algorithm’s results.

Only recently, however, has the ML literature on fairness come into conversation with political and moral philosophy (Binns 2018, 1-11,81). Those working in ML tend to see fairness as analogous to the procedural fairness of a coin in a coin toss, whereas political philosophers focus on substantive fairness, which would take into account both the process and its outcomes. One substantive standard of fairness would be egalitarian, which would require equal chances of good and bad outcomes, regardless of group membership. Another would be a desert-based standard, which would require equal chances for those exhibiting the same conduct, regardless of group membership (Binns 2018, 1-11,81). A further alternative is a justificatory standard of fairness, which would require that a process be justifiable to each person it affects because that standard is sensitive to the differential bargaining positions of the advantaged and disadvantaged.

The justificatory standard has the benefit of being disaggregated, since the analysis is not at the group level; the individuals made worse off would have to have only those burdens that were outweighed by the benefits of the system. In the case of an algorithm used for policing, the benefits would be crime reduction, the protection of property, and other similar goods, and the costs to those made worse off would be the direct costs of arrests and punishments, as well as the indirect costs of mistaken arrests (and the resultant loss of trust) and whatever knock-on effects there are of police

presence. Because this standard would need to point to community benefits that are sufficiently great that they outweigh the costs of certain community members being falsely arrested, it sets quite a high bar (Purves 2022). This high bar itself might be justified, however, by its potential to mitigate double victimization - being forced into crime and then being overly policed - resulting from historical injustice.

The debate about the extent to which competing fairness measures align with our best moral theories of fairness is still in its infancy, but momentum is building quickly (Eva 2022; Grant 2023; Hedden 2021; Hellman 2020; Huq 2019; Long 2021). An important next step in the scholarly literature on fairness in machine learning will be to understand how to apply these theoretical lessons to machine learning applications in the real world.

Next Steps

And here we return to a core commitment of the authors of this report: any useful evaluation of a machine learning technology will occur in light of the values and norms of the specific domain in which the technology is being applied. The evaluation that we have presented here is not only domain specific, but it is also multifaceted in that it has engaged with different stages in the development process. This report has targeted not only the way that the technology will be used, but also the processes of development and oversight that occur before and during implementation. Furthermore, addressing myriad risks of PAPM technologies requires a critical assessment of data processing and at times the restructuring or rethinking of the organization itself.

With this in mind, the scope of this framework is intentionally limited to PAPM systems. We have not addressed other sorts of predictive algorithms such as person-based systems, not to mention the multitude of data-driven policing technologies used in criminal investigations such as facial recognition. Each of these raises its own host of issues and deserves its own treatment.

As we flagged early in the report, the framework also contains only a brief discussion of the privacy concerns which undergird many of the issues that have been covered. Privacy too deserves its own full treatment. With that said, this framework begins from the assumption that there is already an existing data set at the disposal of developers. In that sense we are bracketing the question of: “Which data should we use?” and instead asking, “What should we do with the data we already have?” We hope this report serves as a useful guide for community advocates, police departments, and developers who are seeking an answer to that question.

About the Authors

Duncan Purves, PhD, is an associate professor in the Department of Philosophy at the University of Florida. His research focuses on the ethical dimensions of emerging technologies, especially decision-making aided by artificial intelligence. He has published his work in leading philosophy and ethics journals including *Philosophical Studies*, *Ethics & Information Technology*, and *Ethical Theory and Moral Practice*.

Ryan Jenkins, PhD, is an associate professor of philosophy and a senior fellow at the Ethics + Emerging Sciences Group at California Polytechnic State University in San Luis Obispo. He studies the ethics of emerging technologies, especially automation, cyber war, autonomous weapons, and driverless cars. His work has appeared in journals such as *Techné*, *Ethics & Information Technology*, *Ethical Theory and Moral Practice*, and the *Journal of Military Ethics*, as well as public fora including the *Washington Post*, *Slate* and *Forbes*.

Appendix: Condensed Recommendations

Legitimacy

Risk: A core goal of policing is to prevent crime and disorder, ideally by preventative means, while respecting citizens' constitutional rights and preserving the conditions required for legitimate policing. Failures of sufficient competence, proportionality, equality in enforcement, and public trust can lead to a loss of police legitimacy.

- **Recommendation:** When implementing a new PAPM system, departments should, at the same time, establish standards and methods necessary to assess the efficacy of the system.
- **Recommendation:** Develop a feasible method by which to measure the system's effect on public approval of police actions as well as trust of police. If the system negatively affects public approval, modify the design or implementation of the system.
- **Recommendation:** When setting priorities, departments must seek input from community members about citizens' needs and concerns.
- **Recommendation:** Law enforcement priorities, informed by community consultation, should guide the adoption of PAPM technology rather than technology adoption guiding law enforcement priorities.
- **Recommendation:** Create mechanisms for public feedback about the use of PAPM systems by police.
- **Recommendation:** Ask whether the PAPM system encourages policing tactics that involve the use of physical force, and consider less aggressive tactics where possible.
- **Recommendation:** Develop methods to determine whether the system encourages unequal enforcement of the laws (see "feedback loops" below).
- **Recommendation:** Ensure that the PAPM system encourages equal enforcement of laws.

Data

Framing Questions:

Risk: Developers of PAPM models will need to make choices between the prediction of different crimes that may sometimes conflict. How should these be prioritized?

- **Recommendation:** Developers should create PAPM models capable of prioritizing some crimes over others.
- **Recommendation:** Developers and police should ensure that choices about development and deployment of PAPM systems reflect the priorities and concerns of community members.
- **Recommendation:** RMS developers and police departments should prioritize transparency in the development process.
- **Risk:** PAPM systems should avoid including crimes that can not be effectively geographically targeted.
- **Recommendation:** Select models for crimes that are deterable using the method of police patrol. Avoid modeling crimes whose data are susceptible to enforcement bias.
- **Recommendation:** Crime prioritization choices should be endorsed by the community. Short of that, the crimes prioritized must have the potential to cause serious harm.
- **Recommendation:** Avoid modeling officer-discovered crimes.

Quality Questions:

Risk: The overprediction of crime in certain locations may result in biased prediction. Because nuisance crimes are officer-discovered crimes, they are especially prone to the risk of biased predictions.

- **Recommendation:** RMS developers should include a data label which facilitates the exclusion of officer discovered crimes.
- **Recommendation:** To disentangle competing hypotheses about the effects of interventions on crime and community sentiment, police departments should seek reliable, longitudinal data on community sentiment and the impacts of policing tactics on that sentiment.

Risk: The practice of "Bucketing," wherein similar crimes are grouped together to bolster the accuracy of predictions, could obscure important differences between crimes that respond differently to different interventions.

- **Recommendation:** Agencies should avoid “bucketing” crimes together, when possible, unless the crimes are amenable to the same kind of policing intervention.
- **Recommendation:** Excluding arrest data, agencies should take advantage of the ability of patrol officers to provide data for improved predictions, especially as concerns environmental factors.

Risk: Modeling environmental features of a place that make it vulnerable to crime risks being reflective of racially discriminatory policies from which those environmental features emerged. And yet, removing all features that correlate with race risks undermining the accuracy of predictions

- **Recommendation:** Developers should have an explicit policy concerning how to strike a balance between accuracy and controversial risk factors.
- **Recommendation:** Developers should audit their model outputs for unjustified correlations with race as well as develop sound rationales for the fairness metrics they use to audit their systems.

Risk: The concern that PAPM technologies are racially biased is a substantial driver of public skepticism and distrust. If this distrust manifests in refusal to cooperate with police, then it may undermine law enforcement effectiveness

- **Recommendation:** Departments should actively facilitate public understanding and endorsement of data-driven policing technologies by disclosing methods and publicizing third party audits.

Algorithm Questions

Risk: Many machine learning models are **opaque black box algorithms**, meaning that users have no way to interpret which features are driving their results (C).

- **Recommendation:** Interpretable models or models with accessible feature importance functionality should be preferred, for example (1) Linear models; (2) Tree-based models.

Risk: The algorithms may learn to use features as **proxy variables** (variables strongly associated with other variables not included in the data) **for protected features** (e.g. the model learns to “implicitly” consider the racial composition of a neighborhood, which deliberately was left out of the data, by considering other features associated with race, such as location, median income, and average level of education) (E).

- **Recommendation:** Check model output for racial bias by identifying and removing proxy variables, and striking a balance between ethical risk and efficacy.

Algorithmic Choices

Risk: The PAPM system might over-represent crime in some communities, leading to negative police interactions with the community. It might also under-representing crime, leading to the community being underserved.

- **Recommendation:** Police departments and developers should incorporate community priorities when calculating cost and benefits of over or under-representation, and they should reduce costs of overrepresentation by adopting non-enforcement-oriented interventions where possible.

User interaction

Risk: Even if the system both predicts risk zones and recommends low harm interventions, police officers may fail to follow the recommendations because of discomfort or system distrust.

- **Recommendation:** Departments should present information to officers that encourages faithful adherence to the PAPM system's recommendations including clear, actionable explanations for patrol recommendations.

Risk: Manually adding risk zones can import officer bias and negatively affect the functioning of the system in the long term.

- **Recommendation:** Developers should limit the ability of users to alter the PAPM system unless they are certain it can be done without risk of bias or misuse.

Risk: PAPM systems are often primarily used to allocate police patrols, rather than to inform problem-oriented policing approaches. This can lead to an increased police presence in an area, thereby increasing the risk of adverse or violent interactions with residents in that area.

- **Recommendation:** Data on the physical features of locations that make them prone to crime should be used proactively where possible to reduce reliance on enforcement-oriented tactics.

Risk: Failure to provide a sufficient explanation to patrol officers of a recommended patrol can lead to distrust in, and undue disillusionment with, the PAPM system.

- **Recommendation:** Information about why an officer is directed to a location should be made available to that officer.
- **Recommendation:** Surface only information that may alter user behavior in a way that better serves the community and minimizes the chance of harming an innocent citizen.
- **Recommendation:** Developers should maintain contact with agencies to encourage sustained and effective usage of the product, tracking the department's engagement with the product over time, and explaining the benefits of the product at all levels of the organization.
- **Recommendation:** Developers and police departments should work together to gather information on officer attitudes towards the PAPM system.

Risk: Officers may feel micromanaged by a PAPM system if the system does not effectively leverage officers' own expertise and experience or explain its forecasts in at a level of detail that is actionable by officers.

- **Recommendation:** Vendors should make a judicious selection about what elements of the system to make transparent to crime analysts and patrol officers.

Risk: Vendors, however, should be mindful of the pressure they apply and maintain a good faith relationship with departments. This is a delicate and resource-intensive task.

- **Recommendation:** Recommended interventions should take account of the burdens they place on the community being served as well as the effect of those interventions on community relations. Typically, non-enforcement activities should be preferred when possible.
- **Recommendation:** Police training should include communications training in order to facilitate positive interactions with community members.

Risk: It may initially be unclear which types of interventions are most effective in responding to recommendations of the PAPM system, for example, how frequently police should visit high risk areas to maximize deterrence.

- **Recommendation:** Departments should practice empirically validated interventions in combination with PAPM systems, observing the Koper curve when visiting high risk areas on patrol.
- **Recommendation:** Technology developers and vendors should communicate with police departments early and often, especially regarding the nature of the tool and ways it should be used.

Risk: Without robust data about correlations between certain tactics and reducing crime in a specific location, departments risk embracing bespoke solutions that may apply only to a particular district or locale.

- **Recommendation:** After transferring the product, ongoing conversations should center around police use of the technology, e.g. whether police are faithful to the patrol recommendations of the model.
- **Recommendation:** Vendors should carefully outline, through internal dialogue and engagement with stakeholders, what uses of their technology would constitute “red lines” that they are not permitted to consider.
- **Recommendation:** Other ways of discouraging certain uses could be “designed into” the product itself, such as limiting the choice of which tactics to surface to police on patrol.
- **Recommendation:** Vendors should define and distribute a set of general orders that agencies can easily incorporate when they adopt the technology.

Risk: Allowing officers and analysts to override the predictions of a PAPM system risks being abused and could result in oversaturation or harassment because of mere suspicion.

- **Recommendation:** Vendors must not just accept but embrace their role as choice architects, exploring options to facilitate reasonable customization of technology tools, while giving clear prescriptions when onboarding users to avoid problematic uses.

Over-saturation and Feedback Loops

Feedback Loops:

Risk: Due to the potential formation of feedback loops, there is an increased risk of concentrating harms from police-community interactions in communities who have historically faced discrimination as well as an increase in the ratio of formal to informal contacts between police and community members, which has been shown to decrease community approval of police.

- **Recommendation:** Patrol recommendations should change often, e.g. at least every day.
- **Recommendation:** To prevent patrols from spending too much time in one small area, departments should attach a cost to sending police patrols to an area where they have recently patrolled.

- **Recommendation:** Departments should enforce the “Koper Curve” by timing police presence in an area and encouraging officers to move along after 15 minutes have elapsed. Consider sharing this information with community stakeholders—as they might prefer that police spend more or less time in their area, too.
- **Recommendation:** Because of historical enforcement bias, all data should be examined and audited before being used to generate crime forecasts.
- **Recommendation:** Developers should include only data that are not influenced by the behavior of police on patrol. This includes but is not limited to such things as environmental factors.

The elasticity of crime:

Risk: If PAPM systems are not sensitive to elasticity in criminal behavior, allocating more police resources in the forecasted high-crime locations, while allocating fewer resources in forecasted low-crime locations, could be counterproductive

- **Recommendation:** Police departments should inject a stochastic factor to ensure that certain neighborhoods are not overrepresented in the recommended patrols.
- **Recommendation:** Developers should encourage officers to diversify their patrol choices. This can be achieved by gamification or other incentive mechanisms, including department “leaderboards” with corresponding rewards.
- **Recommendation:** Include a layer of analysis on top of the geographic predictions of forecasting models that measures the concentration and frequency of police patrols in certain neighborhoods to protect against over-policing those areas.
- **Recommendation:** PAPM systems should suggest behaviors to officers on patrol that are less likely to lead to a feedback loop such as parking in a visible location, which may have much the same deterrent effect on speeding as performing traffic stops, without generating more self-fulfilling data.

Risk: Concerns about feedback loops notwithstanding, using data about officer-discovered crime can have benefits. For example, collecting data about the crimes police encounter while on recommended patrols could help determine whether the model is working.

Going beyond patrol allocation:

Risk: Failure to address the physical features of places that make them prone to crime can leave a lot of crime prevention gains on the table. Addressing those features requires fluid collaboration between police, community groups, and other city agencies such as public works.

- **Recommendation:** Developers and police departments should be open to sharing certain types of data, such as physical features of places, with non-law enforcement agencies and community groups in order to assist with non-policing solutions to crime.
- **Recommendation:** Developers and police departments should maintain open channels of communications with non-law enforcement agencies to explore alternative uses of PAPM in effecting crime reduction.

Risk: In responding to an emergency response call for service concerning individuals suffering with mental illness, law enforcement officers are dispatched even when emergency mental health professionals are available, that are more prepared to deescalate the situation and find the help that the individual needs. This can lead to suboptimal care results.

- **Recommendation:** PAPM developers should facilitate novel approaches to law enforcement, such as police and mental health professional collaborations, by exploring the possibility of generating reliable recommendations about the type of emergency response a dispatcher should demand.

Risk: Conflating the frequency of crime in an area with the level of risk of harm to the patrolling officer can produce negative effects, ranging from a disproportional police response and a corresponding loss of public trust to otherwise preventable police and community casualties.

- **Recommendation:** PAPM system developers should explore crime models which account for data about law enforcement officer deaths and assaults, focusing on the locations, methods of assault, and the physical characteristics of the environment and the time of day, while avoiding data that relate to the protected characteristics or identities of the assailants.
- **Recommendation:** Developers should explore the potential of PAPM systems that can estimate the safety risk facing patrol officers during a recommended patrol alongside other metrics to ensure a proportionate response.

Organizational ethics

Risk: The risks identified above will be overlooked if ethics is not incorporated into the culture and workflow of technology companies. Ideally, these risks would be addressed in advance of product release and implementation.

- **Recommendation:** Formalize a position of chief ethics officer with clearly stated responsibilities for forecasting and anticipating ethical concerns with their products, who is accessible to employees throughout the organization.
- **Recommendation:** Develop statements of values or ethical commitments for internal distribution with concrete examples of tradeoffs or tensions that employees might be expected to navigate, or which they have navigated in the past. This statement can be used to justify real choices such as the crimes that will be modeled and why.
- **Recommendation:** Normalize internal conversations about ethical concerns with technology development and use.
- **Recommendation:** Include ethical requirements in product specifications. For example, specify and continually revisit concrete requirements for levels of bias, transparency, and explainability.
- **Recommendation:** Disclose methods of generating patrol recommendations to citizens or discuss the results of third-party bias audits with the public to improve public understanding and reception of data-driven policing technology.
- **Recommendation:** Incorporate a session on ethical impacts and ethical product design into employee onboarding and orientation, to establish a firm foundation of ethical engagement from the start of an employee's tenure.
- **Recommendation:** Include "ethics checks" in the product development process. For teams that use the Agile development process, consider, for example, including "consequence scanning" events in their workflow.
- **Recommendation:** Include ethical requirements in product specifications with clear ethical performance metrics; they should be treated as at least on par with other product metrics.

Appendix: Workshop Participants and Consultants

Workshop Participants

Paul Ames is ShotSpotter Senior Vice President, Products and Technology, leading the product development, software, hardware and operational engineering teams. Paul started his career in the UK and has held senior technology leaderships roles across a broad range of industries including communications, financial, professional information and media. He received his Computer Science education at the Polytechnic of South Wales and post graduate studies in Electronic Sound at University College, Cardiff.

Roseanna Ander serves as the Founding Executive Director of the University of Chicago Crime Lab (since 2008) and the University of Chicago Education Lab (since 2011), which work to design, test, and scale data-driven programs and practices that improve the public sector's approach to public safety and education. Since their inception, Ander has led the Crime Lab's and Education Lab's efforts on violence prevention, criminal justice reform, and improved educational outcomes in Chicago, New York, and around the nation. Ander also helped launch the University of Chicago Urban Labs network, with the creation of three independently run labs focused on poverty, health, and the environment. Prior to joining the University of Chicago, Ander worked at the Joyce Foundation, was a Soros Justice Fellow with the Massachusetts Attorney General's Office and worked for the Harvard Injury Control Center and the Harvard Project on Schooling and Children. Ander holds a M.S. from the Harvard School of Public Health.

Captain **Jonas Baughman** is a 19-year veteran of the Kansas City Police Department (KCPD) with experience in patrol, investigations, crime/intelligence analysis, and administration. His professional passions are crime analysis and data-driven public safety strategies, as well as data visualization. He has served as a sworn crime analyst, helped create and supervise the KCPD's Real-time Crime Center, and directed a gang intelligence Detective squad, among other assignments within the Fiscal Division and Information Services Division. Captain Baughman is currently assigned

to the Chief's Office where he provides strategic analysis and other performance metrics to executive command staff.

Clinton Castro is assistant professor of philosophy and director of the certificate in Ethics, Artificial Intelligence & Big Data at Florida International University (FIU) in Miami, Florida. His primary areas of study are information ethics, fair machine learning, and epistemology. His recently published book, *Algorithms and Autonomy* (co-authored with Adam Pham and Alan Rubel), examines how algorithms in criminal justice, education, housing, elections and beyond affect autonomy, freedom, and democracy. He is currently working on a book project with his FIU colleague Tim Aylsworth that argues for a moral duty to prevent digital distraction in ourselves and others.

Alexander Einarsson is a PhD candidate with the C3 Lab at Northwestern University under the supervision of Kristian Hammond. His research interests lie in the area of artificial intelligence for social good, where his main focus is to bridge the information gap between data and stakeholders in education. His work for the Predictive Policing Project has primarily been to lead the team in questions regarding data analytics and machine learning technologies, looking for potential flaws that may negatively impact the populations the PAPM systems are meant to help.

Kristian Hammond is the Bill and Cathy Osborn Professor of Computer Science at Northwestern University and the co-founder of the Artificial Intelligence company Narrative Science, recently acquired by Salesforce. He is the faculty lead of Northwestern's CS + X initiative, exploring how computational thinking can be used to transform fields such as the law, medicine, education, and business. He is director of Northwestern's Master of Science in Artificial Intelligence (MSAI) and the Engineering lead for Northwestern's Master of Business Administration and Artificial Intelligence (MBAi) that partners Northwestern's schools of business (Kellogg School of Management) and engineering (McCormick School of Engineering).

Yasser Ibrahim is responsible for leading Artificial Intelligence (AI), Imaging Systems and Connected Devices Software at Axon. In addition, Yasser works closely with Axon's independent AI Ethics Board to ensure our technologies are built and deployed with an ethical design framework. Prior to joining Axon in 2020, he served as Amazon's Head of Distributed Machine Learning at Alexa AI. Yasser led applied science and ML engineering teams to develop state-of-the-art deep learning frameworks and distributed algorithms powering Alexa's large scale ML experimentation and production platforms. He also led multidisciplinary teams building computer vision systems and computational imaging algorithms for Amazon Go, the world's first "Just Walk Out" store and technology platform. At Microsoft,

Yasser led computer vision system teams to deliver technologies such as Pixelsense and previously developed flight control and autopilot systems for flight simulators at CAE Inc., in Montreal. Yasser received his B.S. in Microelectronics and Telecommunication Systems from Ain Shams University and an M.B.A. from the University of Washington.

Mecole Jordan-McBride is Advocacy Director at the NYU Policing Project and the former Executive Director of the United Congress of Community and Religious Organizations in Chicago. She began her non-profit career in 2006, volunteering for a campaign to reduce the criminalization of individuals with drug addiction. She quickly became entrenched in the work, supervising reentry and violence prevention programs that provided comprehensive wraparound support for the formerly incarcerated and at-risk youth. She later began local and statewide organizing, conducting organizing and racial equity training for community organizations and parents engaged in education reform. She has gained significant experience working with and building broad-based coalitions for local and statewide reform efforts, particularly around Police Reform and Racial Equity.

Renée Jorgensen is an Assistant Professor of Philosophy at the University of Michigan and (for AY 21-22) a research fellow at the Safra Center for Ethics and Carr Center for Human Rights Policy at Harvard University. She works principally in social and political philosophy, specializing on moral rights and the ethics of risk, with particular focus on the use of predictive tools in criminal justice contexts.

Katie Kinsey is a Staff Attorney with the Policing Project at New York University School of Law where her work focuses on the ethical design and operation of policing technologies and the need to ensure that police use of emerging technologies is legal, ethical, and democratically accountable. Her primary area of research is police use of algorithmic tools such as face recognition. Prior to joining the Policing Project, Katie worked as a Managing Associate in the Complex Litigation Department at Orrick, Herrington & Sutcliffe and clerked for Magistrate Judge Mark P. Lane of the Western District of Texas. Katie graduated cum laude from the University of Texas at Austin School of Law where she served as a Notes Editor on the Law Review. She also holds an M.S. in Special Education from Brooklyn College and a B.A. in political science from Brown University.

Dan Linna has a joint appointment at Northwestern Pritzker School of Law and McCormick School of Engineering as a Senior Lecturer and the Director of Law and Technology Initiatives. Dan's teaching and research focus on innovation and technology, including computational law, artificial intelligence, data analytics, leadership, operations, and innovation frameworks. Dan is also an affiliated faculty member at CodeX—The Stanford Center for Legal Informatics. Previously, Dan was an equity partner

at Honigman Miller Schwartz and Cohn. Before beginning his legal career, Dan was an information technology manager, developer, and consultant.

Monica Mena is Director of Education and Outreach and leads Anti-counterfeiting initiatives for Underwriters Laboratories. The motivating force throughout her career has been to empower people through education. Monica leads a partnership with INTERPOL that has resulted in the development of an online, anti-counterfeiting learning platform. This online college educates over 35,000 law enforcement officers, in 6 languages, about how to combat the crime of counterfeiting. In 2019, Monica developed and launched the first consumer anti-counterfeiting education campaign, Be Safe Buy Real, that has reached over 10 million people with its safety messages.

Terrance Pitts joined the Center on Race, Inequality, and the Law at NYU School of Law as a Senior Research and Advocacy Fellow in 2021. Terrance's career has focused on disrupting mass criminalization and racial bias in the criminal legal system and transforming the conditions which allow violence to cause harm in communities of color. Terrance began his advocacy career addressing racial bias in the criminal legal system as a Project Director at the National Coalition to Abolish the Death Penalty. Since graduating from law school, he has worked at Open Society Foundations, the Ford Foundation, the Vera Institute of Justice, and in a number of consulting roles to support transformation of the U.S. criminal legal system. As a senior advisor at the Ford Foundation, Terrance managed a grantmaking portfolio focused on substantially reducing U.S. jail and prison populations. As a program officer at Open Society Foundations (OSF), Terrance's work focused on supporting advocacy and research to transform U.S. policing practices to make them more transparent and accountable. Between 2015 – 2021, he conducted numerous landscape analyses - interviewing more than 200 advocacy, government, and academic stakeholders to evaluate efforts in local jurisdictions focused on police reform.

Sarah Spurlock is the Associate Director of the Center for Advancing Safety of Machine Intelligence at Northwestern University. In this role, she is building and managing the operational capacity for this new research hub to execute on its mission of realizing the promise of machine intelligence that is safe, equitable and beneficial. Sarah has been with Northwestern since 2019. Prior to this she worked for the Institute of International Education and for Johnson & Johnson in a number of functions and leadership roles in operations, supply chain management, project management, business process improvement, and systems development and management.

Schuyler Sturm is a PhD candidate in Philosophy at the University of Florida. His research investigates the conditions under which it is permissible to use artificial intelligence to assist with decision making. He is interested especially in determining which standards of fairness and transparency apply to AI applications in different domains of use.

Dulani Woods is a data scientist and “data cowboy” at the RAND Corporation. He is adept at data acquisition, transformation, visualization, modeling, simulation, optimization, and more. Dulani has worked on policy related research efforts in multiple domains, with a primary focus on justice and homeland security including authorship of multiple peer-reviewed publications. Dulani’s RAND projects often employ crowdsourcing techniques including the Delphi method with expert focus groups criminal justice practitioners (law enforcement, courts, and corrections) as well as other homeland security and defense audiences. Dulani also specializes in building, operating, and maintaining large policy simulation models. He has developed or maintained models designed to estimate potential policy impacts on justice outcomes, health insurance markets, alcohol and public health policies, COVID and flu vaccination behavior, defense logistics, and Coast Guard mission execution. Dulani also has expertise at examining benefit, cost, performance, and risk tradeoffs for clients including the U.S. Coast Guard, Department of Defense, and the U.S. Forest Service. Dulani served for 10 years as a Coast Guard Officer on afloat and ashore assignments. He served for two years as a U.S. Peace Corps Volunteer in the Republic of Georgia. He currently holds an M.S. in Agricultural Economics (applied economics) from Purdue University, a B.S. in Mechanical Engineering, and a B.S. in Naval Architecture and Marine Engineering from the U.S. Coast Guard Academy.

Senior Personnel

Juan E. Gilbert is the Andrew Banks Family Preeminence Endowed Professor and Chair of the Computer & Information Science & Engineering Department at the University of Florida where he leads the Human Experience Research Lab. Dr. Gilbert has research interest in Human-Centered AI, machine learning, advanced learning technologies, usability and accessibility, and Ethnocomputing (Culturally Relevant Computing). He is the inventor of Applications Quest. Applications Quest is a patented AI used in admissions, scholarships and hiring decisions to select qualified applicants while maintaining diversity. He has published more than 180 articles, given more than 250 talks and obtained more than \$28 million dollars in research funding. He is an ACM Fellow, a Fellow of the American Association of the Advancement of Science and a Fellow of the National Academy of Inventors.

Keith Abney is a senior lecturer in the Philosophy Department and a senior fellow at the Ethics + Emerging Sciences Group at Cal Poly, San Luis Obispo. His work focuses on the ethics of emerging technologies, especially ethics and moral reasoning in robots and other autonomous systems in various domains, from space to autonomous vehicles. He has published in other areas of technology ethics, such as AI risk assessment, cyberwarfare, space colonization, and enhancement bioethics, and his work has appeared in public fora such as the Communications of the ACM, Slate, BBC World Radio, io9.com, and others. He co-edited both *Robot Ethics* (MIT Press, 2012) and *Robot Ethics 2.0* (OUP, 2017).

Patrick Lin is the director of the Ethics + Emerging Sciences Group at Cal Poly, where he is a full philosophy professor. He is currently affiliated with Stanford Law School, the 100 Year Study on AI, Czech Academy of Sciences, Center for a New American Security, and the World Economic Forum. Previous affiliations include: Stanford's School of Engineering, US Naval Academy, Univ. of Notre Dame, Dartmouth, UNIDIR, and the Fulbright specialist program (Univ. of Iceland). Prof. Lin is well published in technology ethics—especially related to security and defense—and is regularly invited to provide briefings on the subject to industry, media, and government.

Project Consultants

David Boonin, PhD, is Professor of Philosophy and Chair of the Philosophy Department at the University of Colorado Boulder. Dr. Boonin's research focuses on the areas of applied ethics, ethical theory, and the history of ethics.

Joel Caplan, PhD, is a professor at the Rutgers University School of Criminal Justice. Dr. Caplan's research focuses on geographic information systems (GIS), risk assessment, crime prevention, policing, and police-community relations.

Nick Evans, PhD, is a Professor of Philosophy at the University of Massachusetts Lowell. Dr. Evans' research focuses on national security, emerging technologies, and the ethics of infectious disease, with a focus on clinical and public health decision making during disease pandemics.

Andrew Guthrie Ferguson, JD, LL.M., is Professor of Law at the University of the District of Columbia. He is a national expert on juries, predictive policing, and the Fourth Amendment.

Katerina Hadjimatheou, PhD, is a criminologist and applied philosopher working on the ethical aspects of technological developments and the use of data in security, policing, and criminal justice. She chairs a police force

ethics committee in the UK and is a member of ethics committees of the National Crime Agency and Metropolitan Police.

Lyria Bennett Moses, JSD, LLM, LLB, BSc (Hons), is Professor of Law at UNSW Sydney and Director of the Allens Hub for Technology, Law and Innovation. She is an expert on law and technological change, legal and policy issues for artificial intelligence, and legal and policy issues for the use of data analytics for law enforcement and national security.

Andrew Selbst, JD, is an assistant professor at UCLA School of Law. His research examines the relationship between law, technology, and society.

Student Assistants

Several student assistants made significant contributions to this report, offering their skills and expertise. At Northwestern University, **Alex Einarsson** provided insightful technical advice, which was instrumental in enhancing the credibility and quality of the work. At Cal Poly, **Shelby Trudeau** contributed background research, laying a solid foundation for the report. At the University of Florida (UF), **Schuyler Sturm** diligently captured notes during our workshop, helping to consolidate and highlight insights and findings. The combined efforts of **Jack Madock** and **Panagiotis Saranteas** at UF ensured meticulous proofing and reference checking. Their attention to detail was invaluable.

References

- Abramson, Ashley. 2021. "Building Mental Health into Emergency Responses: More Cities Are Pairing Mental Health Professionals with Police to Better Help People in Crisis." *Monitor on Psychology, American Psychological Association* (blog). July 1, 2021. <https://www.apa.org/monitor/2021/07/emergency-responses>.
- ACLU of New York. 2022. "A Closer Look at Stop-and-Frisk in NYC | New York Civil Liberties Union | ACLU of New York." December 12, 2022. <https://www.nyclu.org/en/closer-look-stop-and-frisk-nyc>.
- Akin, Ezra. 2017. "A Multi-Armed Bandit Approach To Following A Markov Chain." Master's Thesis, Monterey, California: Naval Postgraduate School. <https://apps.dtic.mil/sti/pdfs/AD1046284.pdf>.
- Apel, Robert. 2016. "On the Deterrent Effect of Stop, Question, and Frisk: Stop, Question, and Frisk Practices." *Criminology & Public Policy* 15 (1): 57–66. <https://doi.org/10.1111/1745-9133.12175>.
- Barnum, Jeremy D., Joel M. Caplan, Leslie W. Kennedy, and Eric L. Piza. 2017. "The Crime Kaleidoscope: A Cross-Jurisdictional Analysis of Place Features and Crime in Three Urban Environments." *Applied Geography* 79: 203–11. <https://doi.org/10.1016/j.apgeog.2016.12.011>.
- Barocas, Solon, and Andrew D. Selbst. 2016. "Big Data's Disparate Impact." *California Law Review* 104 (3): 671–732.
- Bennett Moses, Lyria, and Janet Chan. 2018. "Algorithmic Prediction in Policing: Assumptions, Evaluation, and Accountability." *Policing and Society* 28 (7): 806–22. <https://doi.org/10.1080/10439463.2016.1253695>.
- Berk, Richard, Hoda Heidari, Shahin Jabbari, Michael Kearns, and Aaron Roth. 2021. "Fairness in Criminal Justice Risk Assessments: The State of the Art." *Sociological Methods & Research* 50 (1): 3–44. <https://doi.org/10.1177/0049124118782533>.
- Binns, Reuben. 2018. "Fairness in Machine Learning: Lessons from Political Philosophy." In *Conference on Fairness, Accountability and Transparency*, 149–59. PMLR. <http://proceedings.mlr.press/v81/binns18a.html>.
- Braga, Anthony A., David L. Weisburd, Elin J. Waring, Lorraine Green Mazerolle, William Spelman, and Francis Gajewski. 1999. "Problem-Oriented Policing In Violent Crime Places: A Randomized Con-

- trolled Experiment.” *Criminology* 37 (3): 541–80. <https://doi.org/10.1111/j.1745-9125.1999.tb00496.x>.
- Brayne, Sarah. 2020. *Predict and Surveil: Data, Discretion, and the Future of Policing*. 1st edition. New York, NY: Oxford University Press.
- Caldero, M. A., J. D. Dailey, and B. L. Withrow. 2018. *Police Ethics: The Corruption of Noble Cause*. 4th ed. New York: Routledge.
- Calo, Ryan. 2011. “The Boundaries of Privacy Harm.” *Indiana Law Journal* 86: 1131.
- Carton, Samuel, Jennifer Helsby, Kenneth Joseph, Ayesha Mahmud, Youngsoo Park, Joe Walsh, Crystal Cody, CPT Estella Patterson, Lauren Haynes, and Rayid Ghani. 2016. “Identifying Police Officers at Risk of Adverse Events.” In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 67–76. KDD ’16. New York, NY, USA: Association for Computing Machinery. <https://doi.org/10.1145/2939672.2939698>.
- Chicago PD. n.d. “Chicago Police Department Sentiment Dashboard.” Sentiment Dashboard. <https://home.chicagopolice.org/statistics-data/data-dashboards/sentiment-dashboard/>.
- Chouldechova, Alexandra. 2017. “Fair Prediction with Disparate Impact: A Study of Bias in Recidivism Prediction Instruments.” *Big Data* 5 (2): 153–63. <https://doi.org/10.1089/big.2016.0047>.
- Corbett-Davies, Sam, and Sharad Goel. 2018a. “The Measure and Mismeasure of Fairness: A Critical Review of Fair Machine Learning.” arXiv. <http://arxiv.org/abs/1808.00023>.
- . 2018b. “The Measure and Mismeasure of Fairness: A Critical Review of Fair Machine Learning.” arXiv. <https://doi.org/10.48550/arXiv.1808.00023>.
- Datta, Anupam, Matt Fredrikson, Gihyuk Ko, Piotr Mardziel, and Shayak Sen. 2017. “Proxy Non-Discrimination in Data-Driven Systems.” <https://doi.org/10.48550/ARXIV.1707.08120>.
- Desmond, Matthew, Andrew V. Papachristos, and David S. Kirk. 2016. “Police Violence and Citizen Crime Reporting in the Black Community.” *American Sociological Review* 81 (5): 857–76. <https://doi.org/10.1177/0003122416663494>.
- DotEveryone. n.d. “Consequence Scanning – an Agile Practice for Responsible Innovators.” Accessed June 17, 2023. <https://doteveryone.org.uk/project/consequence-scanning/>.
- Duffee, David, David McDowall, Lorraine Green Mazerolle, and Stephen D. Mastrofski. 2000. *Measurement and Analysis of Crime and Justice*. Vol. 4.

- Criminal Justice 2000. Washington, DC: United States Department of Justice.
- Eck, John E. 2019. "Advocate: Why Problem-Oriented Policing." In *Police Innovation: Contrasting Perspectives*, edited by David Weisburd and Anthony Allan Braga, Second edition. Cambridge New York, NY Port Melbourne, VIC New Delhi Singapore: Cambridge University Press.
- Elliott, K. A., and J. M. Pollock. 2014. "The Ethics of Force: Duty, Principle, and Morality." In *Law Enforcement Ethics: Classic and Contemporary Issues*, edited by B. D. Fitch, 231–56. Sage.
- Ensign, Danielle, Sorelle A. Friedler, Scott Neville, Carlos Scheidegger, and Suresh Venkatasubramanian. 2018. "Runaway Feedback Loops in Predictive Policing." In *Conference on Fairness, Accountability and Transparency*, 81:160–71. PMLR. <http://proceedings.mlr.press/v81/ensign18a.html>.
- Eva, Benjamin. 2022. "Algorithmic Fairness and Base Rate Tracking." *Philosophy & Public Affairs* 50 (2): 239–66. <https://doi.org/10.1111/papa.12211>.
- Fagan, Jeffrey A., Amanda Geller, Garth Davies, and Valerie West. 2020. "Street Stops and Broken Windows Revisited: The Demography and Logic of Proactive Policing in a Safe and Changing City." In *Race, Ethnicity, and Policing*, edited by Stephen K. Rice and Michael D. White, 309–48. New York University Press. <https://doi.org/10.18574/nyu/9780814776155.003.0013>.
- FBI. 2011. "Offense Definitions." Federal Bureau of Investigation. 2011. <https://ucr.fbi.gov/crime-in-the-u.s/2011/crime-in-the-u.s.-2011/offense-definitions>.
- . 2019. "Law Enforcement Officers Killed and Assaulted." Federal Bureau of Investigation. 2019. <https://ucr.fbi.gov/leoka/2019/home>.
- Ferguson, Andrew Guthrie. 2017a. "Policing Predictive Policing." *Washington University Law Review* 94 (5): 1108–89.
- . 2017b. *The Rise of Big Data Policing: Surveillance, Race, and the Future of Law Enforcement*. New York: NYU Press.
- Fjeld, Jessica, Nele Achten, Hannah Hilligoss, Adam Nagy, and Madhulika Srikumar. 2020. "Principled Artificial Intelligence: Mapping Consensus in Ethical and Rights-Based Approaches to Principles for AI." *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.3518482>.
- [Geolitica.com](https://geolitica.com). 2021. "Geolitica.Com | Trusted Services for Safer Communities." 2021. <https://geolitica.com/>.
- Giacalone, Joseph, and Alex Vitale. 2017. "When Policing Stats Do More Harm than Good: Column." USA TODAY. February 10, 2017. <https://www.usatoday.com/opinion/>.

- GME. 2023. "Global Digital Policing Market Analysis | Size & Forecasts." *Global Market Estimates Research & Consultants* (blog). March 29, 2023. <https://www.globalmarketestimates.com/market-report/digital-policing-market-3710>.
- Goldstein, Herman. 1979. "Improving Policing: A Problem-Oriented Approach." *Crime & Delinquency* 25 (2): 236–58. <https://doi.org/10.1177/001112877902500207>.
- Grant, David Gray. 2023. "Equalized Odds Is a Requirement of Algorithmic Fairness." *Synthese* 201 (3): 101. <https://doi.org/10.1007/s11229-023-04054-0>.
- Harcourt, Bernard E. 2006. *Against Prediction: Profiling, Policing, and Punishing in an Actuarial Age*. Chicago: University of Chicago Press.
- . 2014. "Henry Louis Gates and Racial Profiling: What's the Problem?" In *Law Enforcement Ethics: Classic and Contemporary Issues*, edited by Brian D. Fitch, 295–324. Los Angeles: SAGE.
- Hart, Timothy C., and Terance D. Miethe. 2014. "Street Robbery and Public Bus Stops: A Case Study of Activitynodes and Situational Risk." *Security Journal* 27 (2): 180–93. <https://doi.org/10.1057/sj.2014.5>.
- Hatten, David, and Eric L. Piza. 2021. "Measuring the Temporal Stability of Near-Repeat Crime Patterns: A Longitudinal Analysis." *Crime & Delinquency* 67 (4): 498–522. <https://doi.org/10.1177/0011128720922545>.
- Heaton, Paul. 2010. "Hidden in Plain Sight: What Cost-of-Crime Research Can Tell Us About Investing in Police." RAND Corporation. https://www.rand.org/pubs/occasional_papers/OP279.html.
- Hedden, Brian. 2021. "On Statistical Criteria of Algorithmic Fairness." *Philosophy & Public Affairs* 49 (2): 209–31. <https://doi.org/10.1111/papa.12189>.
- Hellman, Deborah. 2020. "Measuring Algorithmic Fairness." *Virginia Law Review* 106 (4): 811–66.
- Hunt, Priscillia, Jessica Saunders, and John S. Hollywood. 2014. *Evaluation of the Shreveport Predictive Policing Experiment*. Santa Monica, CA: Rand Corporation.
- Huq, Aziz Z. 2019. "Racial Equity in Algorithmic Criminal Justice." *Duke Law Journal* 68: 1043.
- IACP. 2018. "Steps to Building Trust." International Association of Chiefs of Police. August 15, 2018. <https://www.theiacp.org/resources/steps-to-building-trust>.
- IDC. 2021. "Data-Driven Policing: Everything You Wanted to Know but Suspected the FBI Already Knew." March 8, 2021. <https://blogs.idc>.

- [com/2021/03/08/data-driven-policing-everything-you-wanted-to-know-but-suspected-the-fbi-already-knew/](https://doi.org/10.1080/07418829500096231).
- Jameson, Anthony, Bettina Berendt, Silvia Gabrielli, Federica Cena, Cristina Gena, Fabiana Vernerio, and Katharina Reinecke. 2014. "Choice Architecture for Human-Computer Interaction." *Foundations and Trends® in Human-Computer Interaction* 7 (1-2): 1-235. <https://doi.org/10.1561/11000000028>.
- Jenkins, Ryan, Kristian Hammond, Sarah Spurlock, and Leilani Gilpin. 2022. "Separating Facts and Evaluation: Motivation, Account, and Learnings from a Novel Approach to Evaluating the Human Impacts of Machine Learning." *AI & SOCIETY*, March. <https://doi.org/10.1007/s00146-022-01417-y>.
- Johnson, Eric J., Suzanne B. Shu, Benedict G. C. Dellaert, Craig Fox, Daniel G. Goldstein, Gerald Häubl, Richard P. Larrick, et al. 2012. "Beyond Nudges: Tools of a Choice Architecture." *Marketing Letters* 23 (2): 487-504. <https://doi.org/10.1007/s11002-012-9186-1>.
- Jones, B., and E. Mendieta. 2021. "Introduction: Police Ethics After Ferguson." In *The Ethics of Policing*, edited by B. Jones and E. Mendieta, 1-22. New York University Press.
- Klaver, J. J. 2014. "Law Enforcement Ethics and Misconduct: An Introduction." In *Law Enforcement Ethics: Classic and Contemporary Issues*, edited by B. D. Fitch, 3-28. Sage.
- Kleinberg, Jon, Sendhil Mullainathan, and Manish Raghavan. 2016. "Inherent Trade-Offs in the Fair Determination of Risk Scores." arXiv. <https://doi.org/10.48550/arXiv.1609.05807>.
- Koper, Christopher S. 1995. "Just Enough Police Presence: Reducing Crime and Disorderly Behavior by Optimizing Patrol Time in Crime Hot Spots." *Justice Quarterly* 12 (4): 649-72. <https://doi.org/10.1080/07418829500096231>.
- LAPD. 2019. *The Los Angeles Police Department Strategic Plan 2019-2021*. <http://lapd-assets.lapdonline.org/assets/pdf/Strategic%20Plan%202019-2021.pdf>.
- Li, Jonathan. 2022. "Pitfalls of Predictive Policing: An Ethical Analysis." *Viterbi Conversations in Ethics* 5 (3). <https://vce.usc.edu/volume-5-issue-3/pitfalls-of-predictive-policing-an-ethical-analysis/>.
- Long, Robert. 2021. "Fairness in Machine Learning: Against False Positive Rate Equality as a Measure of Fairness." *Journal of Moral Philosophy* 19 (1): 49-78. <https://doi.org/10.1163/17455243-20213439>.

- Lovell, J. S. 2014. "Public Information in the Age of YouTube: Citizen Journalism and the Expanding Scope of Police Accountability." In *Law Enforcement Ethics: Classic and Contemporary Issues*, edited by B. D. Fitch, 405–21. Sage.
- Lum, Kristian, and William Isaac. 2016. "To Predict and Serve?" *Significance* 13 (5): 14–19. <https://doi.org/10.1111/j.1740-9713.2016.00960.x>.
- Maxson, Cheryl, Karen Hennigan, and David Sloane. 2003. "Factors That Influence Public Opinion of the Police." 197925. National Institute of Justice. <https://nij.ojp.gov/library/publications/factors-influence-public-opinion-police>.
- Meares, T. L. 2021. "The Quest for Lawful versus Effective Policing and the Possibility of Abolition as a Solution." In *The Ethics of Policing*, edited by B. Jones and E. Mendieta, 25–38. New York University Press.
- Meares, Tracey L., and Dan M. Kahan. 1998. "Law and (Norms of) Order in the Inner City." *Law & Society Review* 32 (4): 805–38. <https://doi.org/10.2307/827740>.
- Meijer, Albert, and Martijn Wessels. 2019. "Predictive Policing: Review of Benefits and Drawbacks." *International Journal of Public Administration* 42 (12): 1031–39. <https://doi.org/10.1080/01900692.2019.1575664>.
- Mohl, Raymond. 2004. "Stop the Road." *Journal of Urban History* 30 (5): 674–706. <https://doi.org/10.1177/0096144204265180>.
- 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. <https://doi.org/10.1080/01621459.2015.1077710>.
- Monaghan, Jake. 2021. "Legitimate Policing and Professional Norms." In *The Ethics of Policing: New Perspectives on Law Enforcement*, edited by Ben Jones and Eduardo Mendieta. New York: New York University Press.
- Patel, Faiza. 2015. "Can Predictive Policing Be Ethical and Effective?" *Brennan Center for Justice* (blog). November 18, 2015. <https://www.brennancenter.org/our-work/analysis-opinion/can-predictive-policing-be-ethical-and-effective>.
- Perez, D. W., and J. A. Moore. 2013. *Police Ethics: A Matter of Character*. 2nd ed. Delmar, Cengage Learning.
- Pew Research Center. 2018. "Public Attitudes Toward Computer Algorithms." November 16, 2018. https://www.pewinternet.org/wp-content/uploads/sites/9/2018/11/PI_2018.11.19_algorithms_FINAL.pdf.

- Piza, Eric, Shun Feng, Leslie Kennedy, and Joel Caplan. 2016. "Place-Based Correlates of Motor Vehicle Theft and Recovery: Measuring Spatial Influence across Neighborhood Context." *Urban Studies* 54 (13). <https://doi.org/10.1177/0042098016664299>.
- Poel, Ibo van de. 2013. "Why New Technologies Should Be Conceived as Social Experiments." *Ethics, Policy & Environment* 16 (3): 352–55. <https://doi.org/10.1080/21550085.2013.844575>.
- . 2015. "Conflicting ValuesValue Conflict in Design for Values." In *Handbook of Ethics, Values, and Technological Design*, edited by Jeroen van den Hoven, Pieter E. Vermaas, and Ibo van de Poel, 89–116. Dordrecht: Springer Netherlands. https://doi.org/10.1007/978-94-007-6970-0_5.
- . 2016. "An Ethical Framework for Evaluating Experimental Technology." *Science and Engineering Ethics* 22 (3): 667–86. <https://doi.org/10.1007/s11948-015-9724-3>.
- Purves, Duncan. 2022. "Fairness in Algorithmic Policing." *Journal of the American Philosophical Association* 8 (4): 741–61. <https://doi.org/10.1017/apa.2021.39>.
- Purves, Duncan, and Jeremy Davis. 2022. "Public Trust, Institutional Legitimacy, and the Use of Algorithms in Criminal Justice." *Public Affairs Quarterly* 36 (2): 136–62. <https://doi.org/10.5406/21520542.36.2.03>.
- Rentschler, Jun, Melda Salhab, and Bramka Arga Jafino. 2022. "Flood Exposure and Poverty in 188 Countries." *Nature Communications* 13 (1): 3527. <https://doi.org/10.1038/s41467-022-30727-4>.
- Robertson, Kate, Cynthia Khoo, and Yolanda Song. 2020. "To Surveil and Predict: A Human Rights Analysis of Algorithmic Policing in Canada." The Citizen Lab. <https://citizenlab.ca/2020/09/to-surveil-and-predict-a-human-rights-analysis-of-algorithmic-policing-in-canada/>.
- Rosenfeld, Richard, and Robert Fornango. 2014. "The Impact of Police Stops on Precinct Robbery and Burglary Rates in New York City, 2003–2010." *Justice Quarterly* 31 (1): 96–122. <https://doi.org/10.1080/07418825.2012.712152>.
- Rumbaut, Rubén G., and Egon Bittner. 1979. "Changing Conceptions of the Police Role: A Sociological Review." *Crime and Justice* 1 (January): 239–88. <https://doi.org/10.1086/449063>.
- Saunders, Jessica, Priscillia Hunt, and John S. Hollywood. 2016. "Predictions Put into Practice: A Quasi-Experimental Evaluation of Chicago's Predictive Policing Pilot." *J Exp Criminol* 12: 347–71. <https://doi.org/10.1007/s11292-016-9272-0>.

- Selbst, Andrew D. 2018. "Disparate Impact in Big Data Policing." *Georgia Law Review* 52 (1): 3373.
- Shapiro, Aaron. 2017. "Reform Predictive Policing." *Nature News* 541 (7638): 458. <https://doi.org/10.1038/541458a>.
- Skeem, Jennifer L., and Christopher T. Lowenkamp. 2016. "Risk, Race, And Recidivism: Predictive Bias And Disparate Impact." *Criminology* 54 (4): 680–712. <https://doi.org/10.1111/1745-9125.12123>.
- Skeem, Jennifer, and Christopher Lowenkamp. 2020. "Using Algorithms to Address Trade-offs Inherent in Predicting Recidivism." *Behavioral Sciences & the Law* 38 (3): 259–78. <https://doi.org/10.1002/bsl.2465>.
- Skogan, Wesley G. 2019. "Community Policing." In *Police Innovation: Contrasting Perspectives*, edited by David Weisburd and Anthony A. Braga, 2nd ed. Cambridge University Press. <https://doi.org/10.1017/9781108278423>.
- Solove, Daniel J. 2001. "Access and Aggregation: Public Records, Privacy and the Constitution Symposium: Modern Studies in Privacy Law." *Minnesota Law Review* 86 (6): 1137–1218.
- SoundThinking. 2023. "SoundThinking." SoundThinking. 2023. <https://www.soundthinking.com/>.
- Southall, Ashley, and Michael Gold. 2019. "Why 'Stop-and-Frisk' Inflamed Black and Hispanic Neighborhoods." *The New York Times*, November 17, 2019, sec. New York. <https://www.nytimes.com/2019/11/17/nyregion/bloomberg-stop-and-frisk-new-york.html>.
- Speri, Alice. 2021. "The NYPD Is Still Stopping and Frisking Black People at Disproportionate Rates." *The Intercept*. June 10, 2021. <https://theintercept.com/2021/06/10/stop-and-frisk-new-york-police-racial-disparity/>.
- The Policing Project. 2020. "Evaluative Framework for Responsible Tech." The Policing Project. 2020. <https://www.policingproject.org/tech-framework>.
- Tyler, Tom R. 2004. "Enhancing Police Legitimacy." *Annals of the American Academy of Political and Social Science* 593 (1): 84–99.
- Tyler, Tom R., and Yuen J. Huo. 2002. *Trust in the Law: Encouraging Public Cooperation with the Police and Courts*. The Russell Sage Foundation Series On Trust. New York - Russell Sage Foundation.
- Verma, Sahil, and Julia Rubin. 2018. "Fairness Definitions Explained." In *Proceedings of the International Workshop on Software Fairness*, 1–7. FairWare '18. New York, NY, USA: Association for Computing Machinery. <https://doi.org/10.1145/3194770.3194776>.

- Vermorel, Joannès, and Mehryar Mohri. 2005. "Multi-Armed Bandit Algorithms and Empirical Evaluation." In *Machine Learning: ECML 2005*, edited by João Gama, Rui Camacho, Pavel B. Brazdil, Alípio Mário Jorge, and Luís Torgo, 3720:437–48. Lecture Notes in Computer Science. Berlin, Heidelberg: Springer Berlin Heidelberg. https://doi.org/10.1007/11564096_42.
- Weisburd, David, and Lorraine Green. 1995. "Policing Drug Hot Spots: The Jersey City Drug Market Analysis Experiment." *Justice Quarterly* 12 (4): 711–35. <https://doi.org/10.1080/07418829500096261>.
- Wetzel, Linda. 2018. "Types and Tokens." In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Fall 2018. Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/fall2018/entries/types-tokens/>.
- Williams, Simon, and Timothy Coupe. 2017. "Frequency Vs. Length of Hot Spots Patrols: A Randomised Controlled Trial." *Cambridge Journal of Evidence-Based Policing* 1 (1): 5–21. <https://doi.org/10.1007/s41887-017-0003-1>.
- Zien, Alexander, Nicole Kraemer, Soeren Sonnenburg, and Gunnar Raetsch. 2009. "The Feature Importance Ranking Measure." <https://doi.org/10.48550/ARXIV.0906.4258>.