

Prediction as Extraction of Discretion

Sun-Ha Hong
Simon Fraser University
sun_ha@sfu.ca

ABSTRACT

I argue that data-driven predictions work primarily as instruments for systematic extraction of discretionary power – the practical capacity to make everyday decisions and *define one’s situation*. This extractive relation reprises a long historical pattern, in which *new methods of producing knowledge generate a redistribution of epistemic power*: who declares what kind of truth about me, to count for what kinds of decisions? I argue that prediction as extraction of discretion is normal and fundamental to the technology, rather than isolated cases of bias or error. Synthesising critical observations across anthropology, history of technology and critical data studies, the paper demonstrates this dynamic in two contemporary domains: (1) crime and policing demonstrates how predictive systems are *extractive by design*. Rather than neutral models led astray by garbage data, pre-existing interests thoroughly shape how prediction conceives of its object, its measures, and most importantly, what it does *not* measure and in doing so devalues. (2) I then examine the prediction of productivity in the long tradition of *extracting discretion as a means to extract labour power*. Making human behaviour more predictable for the *client* of prediction (the manager, the corporation, the police officer) often means making life and work more unpredictable for the *target* of prediction (the employee, the applicant, the citizen).

CCS CONCEPTS

• **Social and professional topics** → User characteristics; Computing / technology policy; Computing / technology policy; Surveillance;

KEYWORDS

prediction, discretion, power, policing, labor

ACM Reference Format:

Sun-Ha Hong. 2022. Prediction as Extraction of Discretion. In *2022 ACM Conference on Fairness, Accountability, and Transparency (FAcT ’22)*, June 21–24, 2022, Seoul, Republic of Korea. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3531146.3533155>

1 PREDICTION AS EXTRACTION OF DISCRETION

The promise of prediction has become central to the real-world implications of AI and data-driven decision-making. Especially when applied to human behaviours and social outcomes, ‘prediction’ goes

far beyond narrow technical definitions like “a model’s output when provided with an input example” [107], and increasingly involves asserting those *probable* outcomes as bases for *definitive* judgments – despite mounting evidence that much of what we call AI is “fundamentally dubious” for predicting phenomena like employability or recidivism [69]. In this sense, prediction describes a social reality as well as technical process. It reshapes how we talk and think about what counts as true for the purposes of everyday decision-making. Specifically, this reshaping entails a *systematic extraction of discretionary power*: that is, people’s ordinary capacity to *define their situation*. Such extraction is at work when, for instance, automated exam proctoring technology imposes a criminalising gaze on every eye movement and browser click in a bid to catch cheaters. Elsewhere, human border officers insist on the discretionary power to set the terms of each interaction, from routine checks to litigious dissection – and resents any restriction of this freedom by the advent of automated systems [21]. I argue that the main societal effect of the spread of predictive systems is to redistribute this discretionary power, and often in a particular direction – extracting it from the target of prediction (the applicant, the suspect, the employee) and concentrating it in the client of prediction (the employer, the police officer, the manager).

Prediction as extraction of discretion extends the long history of how *new methods of producing knowledge generate a redistribution of epistemic power*: who declares what kind of truth about me, to count for what kinds of decisions? Critical histories of the modern fact amply demonstrate this pattern. New methods of knowledge production, such as the ‘elevation of particulars to the status of fact’ [78:92] or the experiment as a means of collective verification [87], not only generated alternative paths to credibility, but were themselves often products of contemporary political demands for such alternatives [27]. At the same time, the shape of this extractive relation immediately echoes Marx. As we will see in “Entrenching Discretion”, the technological extraction of discretion occurs most forcefully where it supports the extraction of labour power – a relationship that echoes across a century of anticipatory practices around automation. To extract from one is to concentrate in another. I argue that prediction typically entrenches already existing asymmetries of discretion, precisely because it is those who already have discretionary power that can afford to deploy these tools against those who cannot. At least in the US context, the emblematic site here is crime and policing. The police’s enthusiasm for expensive technologies with dubious results is consistent with the full history of the modern police as media institutions [82], and the very figure of the criminal is the product of an extractive principle where some people are defined as incapable of discretion, as dangerous forms of discretion. Below, the paper articulates the extractive process in precisely these efforts to predict (1) criminality and (2) productivity.

Across both, while the mechanisms of capture, analysis and circulation are often novel, the field of social struggles that they are

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

FAcT ’22, June 21–24, 2022, Seoul, Republic of Korea

© 2022 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-9352-2/22/06.

<https://doi.org/10.1145/3531146.3533155>

funded to serve tend to be marked by continuity. Surveillance capitalism, currently one of the most common critical frames around prediction, describes not a seismic revolution, but a variation upon the theme that is capitalism itself [68]. We might raise a similar point with the very idea of AI ethics. The question of how to pursue ethical values in AI has increasingly raised the problem of how to define those values, and *who* should get to do that defining to what ends [6, 38, 54, 77]. Today, these definitional questions are being settled not in the proverbial roundtable of independent thinkers engaged in good faith, but a quagmire of industry-funded lobby groups (such as the Information Technology Industry Council), corporate ethics teams that mistreat and fire their own ethics experts (like Google), and active co-option of critical vocabulary into forms of ethics washing [39]. In terms of funding, computational resources and prestige, Big Tech “make[s] the water in which AI research swims” [101:53] – as visible not only in expensive signature projects like the Chan Zuckerberg Initiative’s \$500 million donation to Harvard, but also the rapid growth of corporate affiliated authorship in AI research [8]. One result is that ethical concepts are narrowed and instrumentalised: bias becomes a form of numerical error to be corrected with better datasets, and ethics a bureaucratic checklist to be inserted into the production flowchart [8, 83] Often, paying for an AI prediction of productivity or criminality (which in many cases may or may not actually involve recognisably AI elements like machine learning) is less a way to establish a more objective foothold on future outcomes, than it is a way to reallocate discretionary power in one’s favour.

But why prediction in particular, amidst the wider umbrella of data-driven technologies? Because it is at the vanguard of both (allegedly) concrete, real world benefits and the enveloping fantasy of long-term progress. The claim to predict a tangible ‘problem’ like recidivism or productivity promises a powerful empirical basis for thinking about new data-driven technologies as desirable and inevitable. We are told that algorithms will allow us to, for example, automatically analyse U.S. prisoner phone calls to predict violent crime and suicide [91]; identify suspect individuals based on their head vibrations[102]; use facial features to predict prospective employees’ personality (HireVue); or conclude whether students are cheating or not in an online test (Mettl). This remains the case even when many such predictions remain unproven – or even specifically *disproven*. COMPAS, drawing from over 100 predictive features to predict recidivism, performs about as poorly as laypersons’ estimates [28]. New, AI-driven tools were celebrated as a potential difference-maker during the COVID-19 pandemic, but were found to be largely useless, plagued by poor input data and other fundamental problems [51]. Yet even as individual applications of prediction might be criticised and lampooned, such promises of real-world payoffs continue to generate emotional and cognitive investment in the idea of prediction (and through it, AI).

To better account for prediction’s social existence, I describe it as a *relational grammar*: a way of thinking and talking about how facts are made, and who tends to declare those facts about whom. The dissemination of predictions generates distinct expectations about what accuracy or objectivity looks like when it comes to human behaviour. It follows that such a grammar prioritises certain kinds of truthmaking over others. The paper first turns to crime and policing to explore how predictive systems are extractive by design. Choices

in what kinds of predictive models are created to pursue what kinds of objects generate systematic disparities between what tends to get measured and what does not. Pernicious myths, such as that prediction disasters are generally dataset problems (the ‘Garbage In, Garbage Out’ (GIGO) argument), continue to obscure the ways in which prediction manages these conditions of possibility. Turning to efforts to predict productivity, especially by Amazon, I argue that making human behaviour more predictable for the client of prediction often means making life and work *unpredictable* for the target of prediction. The advent of data-driven systems is often described as an inevitable hyper-rationality, “render[ing] all agonistic political difficulty as tractable and resolvable” [3]. But in this equation, it is *us* – and our own discretion – that is being defined as the risk and error.

2 PREDICTION AS RELATIONAL GRAMMAR

History of technology demonstrates that new technologies do not enter as blank slates, but that their meaning and usefulness are derived from preexisting social relations [65]. The social impact of technologies also tend to exceed their actual capabilities or implementation. The history of 20th century automation, for instance, was dominated by institutional and economic responses to what people *expected* automation to be and to do – responses which often persisted even when the technology never quite arrived in the expected way [53, 57]. Prediction is as much a way of thinking and talking about how we make facts, and who declares those facts about whom, as it is a set of calculative techniques. The struggle over these concepts – and the moral attitudes, affective orientations, and other pictures of the world embedded into them – shape our perceptions of what kind of technological arrangement is ‘inevitable’, or what kinds of reform, abolition, and alternatives are considered ‘plausible’. The frequent declarations that everything about human beings can be predicted with enough data and computational power, or that such predictive machines will ‘know you better than you know yourself’, reflect and perpetuate these kinds of shifting norms.

Social life is so often driven by the disjuncture between the substance of what passes as true in a given context, and the sense of what *kinds of things* are likely to look, feel, smell true. Foucault famously distinguished between what statements are treated as true, and what statements are considered *sayable* in the first place. The latter are less likely to become visible as formalised rules; in the context of Kuhnian paradigms, Ian Hacking suggests that scientific communities operate through ‘styles of reasoning’, basic tendencies for relating things and making connections that undergird and often outlast any particular theory [44]. From this perspective, descriptions of a facial recognition tool as ‘predicting’ criminality, or a generative adversarial network as ‘learning’, are not simply (misleading) descriptions of technical method; they perpetuate a more basic sentiment about how things relate to each other.

But what does it mean to *relate*? The anthropologist Marilyn Strathern writes that the English word ‘relation’, in its modern sense, serves as something of a master instrument for thinking the world in terms of discrete phenomena which may be joined, separated, sorted, mixed, in highly modular and indiscriminate ways [95]. To say something *relates* conveys a strange neutrality about

the kind of epistemic operation that I am undertaking: it relieves the pressure of identifying a more precise process (such as ‘extracts’, ‘subordinates’, or ‘causes’), and all the ethical and theoretical commitments that those designations would carry with them. (All this makes ‘relation’ an indispensable word for academics.) Similarly, in the context of big data and machine learning, prediction is funneled into a flexibly neutral space, where the scientist or engineer is required only to let ‘the data speak for itself’ and to document those relationships. And if the discovery – for instance, the use of facial recognition to predict sexual orientation – should facilitate clearly harmful uses of technology, it becomes easier to dismiss such consequences as things that simply could not be helped.

This very idea of modular, general relation is neither timeless nor universal. It is very much a modern way of seeing, and one which resonates with the long advent of quantification as a dominant method for understanding human and social phenomena [27, 79]. In this sense, the language of data-driven prediction imposes its own *relational grammar*: not so much the substantial claim that specific visual data of the face accurately identifies sexual orientation or criminality, for example, but the backgrounding tendencies in how we reason and model that makes it ‘plausible’ to try and predict one with the other in the first place. Dan McQuillan calls data science an ‘organising idea’ which imposes a very specific model of knowledge that he describes as neoplatonic: a “belief in a hidden mathematical order that is ontologically superior to the one available to our everyday senses” [67:254]. Others have described similar tendencies in terms of determinism [18] and positivism [103]. While such worldviews are often demurred in polite conversation, they tend to persist in partial and half-avowed ways. As David Golumbia notes, only some people will claim that human beings are just like computers (i.e. ‘the myth of computationalism’), but many more will flirt with the idea in practice [37]. Today, public polls continue to show widespread expectation that artificial general intelligence is imminent, or that automated prediction of phenomena like crime is inevitable [92, 105]. When Clearview AI’s notoriously unethical facial recognition technology became controversial through a *New York Times* report, it was telling to hear one investor fall back exactly on this kind of broad article of faith: “Laws have to determine what’s legal, but you can’t ban technology. Sure, that might lead to a dystopian future or something, but you can’t ban it.” [52]

Crucially, any statement of data-driven systems as objective and consistent describes an aspirational ideal rather than actual practice [56]. In their history of scientific objectivity, Lorraine Daston and Peter Galison write that “all epistemology begins in fear – fear that the world is too labyrinthine to be threaded by reason; fear that the senses are too feeble and the intellect too frail”. [23:372] Prediction as relational grammar does not simply dictate a smooth, totalising blanket of mathematical reason. Equally significant is what gets concealed or devalued in the process of imposing that grammar. Critical media scholar Tara McPherson calls it a lenticular logic, referencing the lenticular lens in old 3D postcards where one can flick between multiple images but never see them *together*. She argues that the history of computing – and its attendant lessons – are often sealed off from other domains [66]. Where McPherson takes up this point specifically around to race and computing, this also extends to how computing then imposes its own ways of seeing on those other domains. Bringing James Scott’s celebrated work

on ‘seeing like a state’ [85] to artificial intelligence, Ali Alkhatib argues that the indiscriminate proliferation of a data-driven way of seeing adds up to an asphyxiating utopia: an expectation of calculability which drives out everything that does not fit. “The models AI researchers train and deploy construct worlds without space for the dimension of our lives that make us who we are” [2:3]. How we attribute predictability – for instance, that a single statistical model can be ‘scaled’ to predict housing prices (Zillow) or ‘suspicious eye movements’ (Proctorio) across many different lives and locales – has direct repercussions for what kinds of actors are considered rational, and what kinds of behaviours are considered optimal. The grammar of prediction tends to prefer particular kinds of data and relations over others: a preference driven not purely on epistemic grounds, but by economic and institutional conditions that make some datasets and problems more available than others. In their critical history of AI’s pursuit of ‘generality’, Raji et al. note that datasets which become treated as general benchmarks, like ImageNet and GLUE, are built not through rigorous definition of what is truly representative of the target reality, but look “more like samples of convenience” [81:5]. Subsequently, any arbitrariness or partiality in such benchmarks tend to flow downstream into derivative applications in a ‘blood diamond effect’ [80]. These limitations are not purely technological in the sense that more diverse datasets and models are simply impossible; rather, they tend to replicate longstanding economic, political, and institutional asymmetries in what kinds of data and knowledge get produced, documented, reported, archived, in the first place. As Scott himself would later note: “legibility doesn’t come cheap.” [84:513]

3 EXTRACTIVE BY DESIGN

3.1 Bad questions

Prediction, like any knowledge regime, has a self-fulfilling element: it sees what it knows to see, and it measures what it can typically imagine measuring. These tendencies are shaped through longstanding economic and political asymmetries, whose influence is regularly written off as uncertain and uncontrollable ‘externalities’. When prediction troubles get compartmentalised as ‘just a dataset problem’, they obfuscate how patterns of extraction shape the research questions and the choice of what to measure (and what to dismiss without measuring). When predictive successes are defined as ‘beating’ humans in tests of accuracy, they reify the model’s own parameters as the world that matters, further burying these externalities beyond consideration.

The sociologist Luc Boltanski, with Laurent Thévenot then Eve Chiapello, has shown how capitalism constructs *regimes of justification*: the rules of the game by which we assess its own success and failure. Such regimes are also exercises in managing expectations [9, 10]. We learn to feel that it is surely unrealistic to expect that an employer will not squeeze its workers as hard as they can to maximise profit, or that if a technology is more ‘accurate’, then it must serve the common good. Some researchers have already adapted Boltanski’s notions to contemporary technoculture, showing how technologies like hiring systems provide justificatory scaffolding for firms [26], and help maintain collective belief in aspirational values like convenience and speed [59]. Prediction is so valuable precisely because it offers both the technical mechanisms and the

logic of justification through which pre-existing extraction of discretion can be replenished. This motivation is the starting point for the shape that predictive systems tend to take today – not the side effect.

Consider the concrete ways in which claims of data-driven prediction actually appear in social life today. Prediction regularly appears as an epithet of approval, in which the technique *du jour* (such as convolutional neural networks) is rapidly scaled onto an improbably wide range of complex social situations. Such claims often derive their attractiveness precisely by claiming to clarify human behaviours and propensities that have long remained contested and ambiguous. Crime and policing are emblematic, partly because the very history of crime as a measurable object, and of modern policing as an institution, is defined by this project of extracting discretion. In one notorious case, a 2020 Harrisburg University press release claimed that “with 80 percent accuracy and with no racial bias, [their] software can predict if someone is a criminal based solely on a picture of their face.”¹ Although this particular study was rejected by Springer, the publisher, following international criticism [50], efforts to graft machine learning tools onto physiognomic worldviews are increasingly common [1, 94]. The Harrisburg study was a norm, not an exception; it followed a now familiar grammar which effectively reduces any human condition into discrete empirical states. In this logic, ‘criminality’ appears as a state that inheres in the person, echoing centuries of imputation of bad blood or evil affliction. A similar exercise is performed with the human face, usually by relying on a variation of the facial action coding system (FACS) and its systematic reduction of facial expressions into objective, decontextualised data points [34:168]. Once frozen into such ossified forms, it only remains to demonstrate some statistical relation between the two artificially stable objects (‘face’ and ‘criminality’) to complete the equation.

On the surface, predictive systems rarely *admit* any commitment to a specific theory of criminality or the face. The grammar of data-driven prediction allows, even encourages, the researcher to avoid asking such questions in the first place – the argument being that it is unnecessary to understand what criminality *is* as long as we can produce actionable measurements bearing its name. Orit Halpern’s history of cybernetics recounts computation’s long rearticulation of objectivity away from the pursuit of ‘external truth or reality’, and towards a calculable grid of action, measurement, and replication [46:83]. This refusal to theorise is often claimed as an apolitical and amoral path; in reality, the very choice to predict the social enters into a politics of erasing ambiguity [7:10-11]. Articulating criminality or productivity as an object of prediction invariably leverages working definitions and heuristics already employed by the investors, engineers, and entrepreneurs involved in the production process – but rarely those of ‘targeted’ populations [20].

3.2 Bad measures

What kind of data tends to be easily available for what purposes? As critical researchers of race and technology have shown with particular clarity, pre-existing inequalities shape what kinds of predictions roam the world in the first place, endowed with what

kind of legitimacy [5, 16]. To suggest that the neutrality of data or algorithm cleanses predictions of their own historical provenance effectively describes a variation of money laundering: all of the decisions involved in any predictive system, from what kind of data is(n’t) gathered to who the ‘insights’ are sold to, are dismissed as somebody else’s problem. The model output, like ‘clean’ money, can then present itself as a neutral object of use. But far from wiping the slate clean when it comes to crime and prejudice, such practices must necessarily rely on existing records and practices to source the data, and are thoroughly shaped by longstanding institutional approaches to crime and patrolling. Ben Green describes how one University of Southern California project began by seeking to predict ‘adversarial groups’ like ISIS, but soon pivoted to targeting ‘criminal street gangs’. They ended up drawing data from the notoriously problematic LAPD gang data for the model, effectively amplifying and legitimating those biases [38]. All this happens despite the fact that “there is no agreement as to what predictive [policing] systems should accomplish [:::] nor as to which benchmarks should be used”, and even if there were, “like all evaluations of police technology, confounding factors make it impossible to measure directly its effectiveness at reducing crime.” [88:459]

At the same time, those who develop and sell these surveillance technologies are incentivised to constantly produce useful justifications around the dangerousness of crime and the necessity of advanced policing. Technologies like PredPol and ShotSpotter are directly sold to law enforcement clients, preying on their own anxieties about being left behind and their perpetual interest in expanding budgets [3:2-4]. Other data-driven tools lead with consumer-facing ‘safety’ services as a way to extract different kinds of urban data. The poster child here has been Ring, acquired by Amazon in 2018, and notorious for hawking its user data from individual door cameras to local law enforcement. Platforms like Ring’s Neighbours incite their userbase into forms of ‘digital vigilantism’ [96] in which ordinary people are encouraged into producing the neighbourhood alerts, video footage, and other ‘data’ that populate a picture of the world in which crime is always just around the corner – a picture which may then be sold back to the users and even law enforcement. Citizen, the most aggressive company in this space, repackages dramatic cases of crime thwarted into ‘magic moments’ for consumption, and even deploys ‘street teams’ into new areas to produce crime content [4]. With Amazon Ring signing up over 400 local police departments for two-way data sharing [49], or Nextdoor encouraging users to pass on their data to the police [55], these systems create new, ‘efficient’ pathways for connecting existing prejudices, perverse economic incentives with the aura of technological innovation [15].

So the choice of what to measure and how to measure it caters to these existing social scripts, often entrenching their distribution of differentially discretionary subjects (the suspect, the ‘bad’ neighbourhood, the police officer: ::). The corollary is that this is not simply a zero-sum game, where greater automation or rationalisation leads to less discretion. Rather, the deployment of predictive systems can serve to protect and relegitimise bureaucratic discretion by integrating those judgments into the system. The move towards data-driven policing often involves new systems of documentation and data-collection, as well as existing data taking on new uses and importance – following a longer pattern of modern

¹An archived version of this press release can be found at: <https://archive.is/N1HVe>

policing as *police media* [82]. The mundane work of documenting, photographing, and ‘writing up’ targets thus becomes a site where police workers’ discretion feeds into large data-driven systems. For instance, consider the LSI-R (Level of Services Inventory-Revised), a longstanding risk assessment tool for predicting individuals’ likelihood of recidivism. One version used by the Idaho Department of Correction [110] shows how details like prior convictions, “dissatisfaction with marital or equivalent situation”, or “unfavourable attitude towards convention” contribute to the quantification of individual risk. Each measure requires interpretive work, some more than others: a score of 1 out of 3 in ‘peer interactions’, the training documents explain, would describe an offender who “actively dislikes co-workers or has only limited contact with them [::and] often lets angry feelings toward others build up inside.” In practice, LSI-R has fared poorly in tests of inter-rater consistency, and evidence suggests its correlation with recidivism is low [48:82-3]. When predictive models are baked into decision-making processes, they do not simply shift the entire apparatus towards inhuman objectivity, but rather empower new norms around *who* gets to impose their discretion upon whom. Often, the promise of data as a universal illuminator conceals the reality in which it is data for me, and not for thee; more predictable production for the client of production, at the expense of less predictable conditions for the target.

3.3 The unmeasured

The common thread running throughout these different sites of data production is that where the police remains the primary purchaser of these predictions, the kind of data and models we “get” as a society will invariably reaffirm their definitions of crime and safety. Conspicuously missing in this thriving industry, for instance, is any effort to improve the disastrous undercollection and invisibility of data around police violence and misconduct. Kelly Gates writes that such asymmetry in the sheer availability of data entails a ‘symbolic annihilation’ [35:4], in which one’s experience is not simply underappreciated, but directly overwritten by police-oriented data production. From Michael Brown and Eric Garner in 2014 to George Floyd and Breonna Taylor in 2020, high-profile police killings of Black Americans precipitate repeated debate around just how prevalent such killings are. Yet each time, it is data on Black *crime* that tends to be readily available – while data on police misconduct and violence is underfunded, neglected, and sometimes actively suppressed. Even as protestors demanded widespread reform in the wake of George Floyd’s death, Chicago’s police union was working to destroy records of police misconduct – as they had been doing for years as part of their contract with the city [47]. Investment in predictive policing and other new technological systems also tend to come hand in hand with more funding for the police to produce the kind of data that they would prefer to exist – from scoring systems to identify the ‘worst of the worst’ offenders, such as LAPD’s Operation LASER [12], to algorithmically generated ‘heat lists’ that purport to predict ‘likely’ perpetrators *and* victims, such as those used by the Chicago Police Department [17]. Palantir’s success at winning contracts from law enforcement and intelligence agencies around the globe, including working with ICE to identify subjects for deportation [109], exemplifies this disparity in what kind of

data gets to exist in the first place. These technologies further generate a vicious feedback loop of legitimation by “assum[ing] the credibility of the underlying crime data—and the policing methods that generate that data in the first place.” [17]

As the Data 4 Black Lives movement has repeatedly shown, such asymmetries condemn certain kinds of suffering and lived experience as ‘merely anecdotal’, forced to an uphill struggle to count as ‘data’, while the police receive the numbers they need by default. Samuel Sinyangwe, co-founder of Mapping Police Violence, describes the recurrent barrier: ‘How often have you heard, “We don’t have the data. We don’t know if what you’re telling us is true.”’ [31] While more robust databases of police misconduct are becoming publicly available, they tend to come from small, independent alliances of researchers, journalists and activists working with limited resources. While the police are legally obligated to produce some kinds of misconduct data, they routinely engage in strategies of obfuscation and delay, turning theoretical availability into practical absence [22]. As a result, whatever becomes transparent about those killed by police – their misdeeds, their personal life – tends to be far more exhaustive than the data we can get about the police themselves [36]. Predictive systems are often used to delegitimise the kinds of lives and experiences that are already too disadvantaged to generate rich ‘data’ in the first place.

Again, these consequences don’t simply follow from who is targeted by predictive technologies and who isn’t. The experience of the predicted also varies depending on what kind of existing pool of discretionary power they can draw upon. Where many of the best known cases involve predictions of criminals, suspects and citizens, other tools target these instruments on police workers themselves (or, as in cases like PredPol, do both in tandem). Researchers have shown how the advent of predictive policing technologies provokes anger and resistance from the officers. As their managers hail a paradigmatic shift from ‘intuition-based’ to ‘data-driven’ policing [13:56], some police workers perceive this as threatening the spaces in which they could previously exercise their discretionary power, resulting in “experiential devaluation and increased managerial surveillance.” [14:2] We thus find them pushing back against third-party tech innovations, arguing that the data junkies do not understand the lived reality of policing [89:182-3]. Others engage in everyday acts of worker sabotage and resistance, such as by refusing to use the automatic vehicle locator and preferring to call in manually once at a location [14:9]. This, of course, is only possible due to the amount of discretionary power police officers retain around the use of predictive systems; the gig workers on DoorDash or students in a ‘smart’ school are systematically denied the possibility. For people who have traditionally been deprived of discretionary power, then, prediction arrives as a double blow. The technology is often designed and funded specifically to intensify this existing relation of extraction, *and* that extraction leaves them least able to contest, evade, or renegotiate the consequences of these predictive judgments.

4 ENTRENCHING DISCRETION

4.1 Defining the situation

Predictive technologies are typically wielded to *rationalise* judgment, in both common senses of the word: they establish rules-bound processes for decision-making, and further provide scripts for justifying that judgment to insiders and outsiders. This joint function indelibly links prediction to discretion. Predictive systems, by definition, reconfigure the existing distribution of decision-making relations: whose ‘gut feeling’ overrides whose? Which parts of the decision are considered sufficiently unknowable that they are open to (unequal) negotiation, and which parts become considered set in stone because the data has spoken?

It is not that prediction always enforces formalisation and Kafkaesque inevitability at the expense of discretion. Rather, prediction serve to *reallocate* discretionary power across different actors, and additionally to *obfuscate* the continuing role of discretionary power in decision-making. In an analysis of data-driven decision-making in European ‘smart’ borders, Alexandra Hall writes that discretion and rule ‘mutually constitute one another.’ [45:493]. The word ‘discretion’ has a complex etymological relationship with various meanings of separation or discrimination, shaping common use today as the ability to determine what is appropriate to a given situation. In particular, it describes the ability to determine how a rule might apply (or *not*) to the case at hand. Thus in legal scholarship, Ronald Dworkin memorably defined judicial discretion as the judge’s ability to ‘choose a solution’ [29], while in sociology, Michael Lipsky’s influential study of ‘street-level bureaucrats’ describes discretion as means of mitigating and coping with the tension between document and person, government agency and individual needs, that dominates the everyday work of bureaucracy [63].

The key point is that prediction and discretion can be used to regulate each other, rather than being locked into a static relation. Police departments might commission predictive systems to re-legitimise their existing discretionary practices upon the urban poor, while workplaces might see it more as a way to suffocate workers’ control over their work process. Discretion is not a relief from control any more than prediction renders judgment objective. Foucault makes this lesson clear in his analysis of *illegalisms*: the extralegal, informal, tacit agreements by which participants profit from or cope with the demands of regulated and official relations [32]. Yet illegalisms often do not represent an escape from the rule. Authorities might tacitly rely on black markets to supply goods when regular markets falter (as has often been the case in, say, North Korea in recent decades); police workers might tolerate some misdemeanours to focus on others. When Foucault says “respect for legality is [itself] only a strategy in this game of illegalism” [32:144], we might say similar of institutions’ respect for data or algorithmic judgment.

But it would be rather subjectless to say that discretion manages the rule and leave it at that. Discretion may allow the judge to choose a solution from the rule, but that also means the defendant is bound to obey both the rule *and* the discretion that went into its deployment. Kohler-Hausmann’s analysis of misdemeanour prosecutions in New York City demonstrates how defendants, once captured through a parking ticket or turnstile jumping, get sucked

into not simply the opprobrium of rule and all its painstaking requirements, but the prosecution’s ability to exercise discretion over that rule. Hence William Stuntz’ famous dictum that criminal law constitutes “items on a menu from which the prosecutor may order as she wishes” [60:12]. Street-level bureaucrats might employ discretion to relieve some of the stupidity of the rule, or to mitigate their own working circumstances, but they can also wrap their own discretionary judgment in the cloak of the rule to enforce it upon the citizen.

Discretion, in other words, describes the always unequal distribution of the power to *define the situation*: which bodies are immediately and tightly fastened to the regimes of visual surveillance and statistical correlation without appeal, and which bodies retain wriggle room, perhaps through the aid of a sympathetic human clerk and other forms of cultural capital? After all, the very justification for the deployment of predictive systems in the first place elevates data and interpretation above other kinds of evidence. This remains broadly true even where model outputs are subject to human review. A risk scoring system for child abuse and neglect may only issue non-binding recommendations, but they will tend to nudge and influence the judgment of overworked social workers, as well as how they justify their discretion to themselves and others [30:141-2]. And however probabilistic the insight, the decisions that result from them are singular and binding for each individual. Extraction never simply withdraws something from somewhere; it also means empowering some at the expense of others.

4.2 Predictable workers, unpredictable work

I now turn to efforts to predict productivity and worker behaviour. As I noted in the beginning, prediction’s *transfer* of discretion from the target of prediction to its client often dovetails with other historical patterns of extraction, and in particular, that of labour power. Consider the case of Amazon. As the fifth largest private employer in the world as of 2020 [90], Amazon is both test-bed and ground reality for prediction as the extraction of worker discretion. Its warehouses (‘fulfilment centres’) and delivery networks exemplify the synergy between intense, data-driven automation of labour on one hand, and the systematic deprivation of information, predictivity, and discretion for the workers on the other. For the former, the master metric is “the rate”: individual performances like boxes stowed, picked or packed per minute and hour. The rate is prominently cited as a reliable base measure for decisions including the hiring and firing of individual workers; it provides an always available source for justifying Amazon’s abnormally high rate of firings [61] - which themselves are increasingly automated [93]. Yet notably, the exact rate that a worker must hit at a given warehouse is never disclosed to the workers themselves. Sociologist Nantina Vgontzas, who has worked at Amazon warehouses as well as joined in workers’ organising efforts, recounts being trained against the rate of 400 picks per hour - and later being told by a human trainer that the actual expectation was around 360 [97]. Meanwhile, the demand for a machine-readable and predictable labour rate is further enacted through sensory opprobrium. Coloured graphs warn employees to work faster; ‘project loving energy’, counsels a workstation screen [24:56].

These quantified expectations governing the algorithmic workplace cater to managers and employers’ desire for a certain kind

of inhuman clarity, in which the many variations and ambiguities inherent in any act of labour are not actually eliminated, but erased, delegitimised and neglected. The consequence is that for the worker, their own work and life becomes both less predictable *and* less discretionary [33]. The Amazon ‘picker’ is constantly adapting to the algorithmic redistribution of boxes and goods, unable to accumulate their own rhythms for effective and safe work. Journalist Alec MacGillis relates the story of Hector Torrez, a San Francisco tech industry professional turned Amazon warehouse worker: “the challenge wasn’t so much the weight as that you couldn’t really tell, based on size, whether a box was going to be heavy or not when you went to pick it up. Your body and your mind never knew what to expect.” [64:4] This deprivation of someone’s ability to anticipate, plan, and adjust their own conditions of working and living, shows clearly that social realities don’t simply become more or less predictable, but that becoming predictable for some often comes at the expense of others.

This disparity in control results in all too familiar harms to working conditions. Amazon’s warehouse workers have long experienced significantly higher rates of severe injury compared to industry competitors [40]. Amazon’s delivery drivers are subject to similar factors: top-down dictations of punishing rates, sometimes of up to 400 deliveries per 10-hour shift, provide profits at the macro level, while workers pay in the form of extraordinarily high injury rates, and the now infamous practice of drivers having to pee in bottles or trash bags to meet the rate [43]. Amazon’s response to these problems tend to intensify this dynamic of profit-oriented datafication. In 2021, Jeff Bezos’ final letter to shareholders as CEO of Amazon prescribed “new automated staffing schedules that use sophisticated algorithms to rotate employees among jobs that use different muscle-tendon groups” [72] – a solution that would implement new forms of unpredictable, opaque decision-making on behalf of workers while retaining the non-negotiability of ‘the rate’ and the punitive overall level of productivity.

Amazon may be a pioneer in the algorithmic destruction of workers’ bodies, but the underlying economic and technical incentives are not at all unique. Coupang, South Korea’s answer to Amazon (though heavily dependent on SoftBank’s Saudi-backed funds, as is the case for many Silicon Valley companies), has also seen rocketing stock value and a reputation for superfast delivery on the back of algorithmic surveillance, high rates of workplace injury, and several workers that have died of overwork [62, 86]. In delivery platforms like DoorDash, the on-demand worker is shepherded into surge areas in one moment, then made to kick their heels curbside in another, never knowing how much they’ll make or how to pace their work [41:7,89:2]. Chinese food couriers, working for platforms like Ele.me and Meituan Waimai, describe sudden periods of intense pressure in which a delay of seconds might cost several orders’ worth of pay – the result of a business and technological model which profits from its ability to “stratify the value of people’s time” in rapid, finely-grained ways, extracting every possible sliver of ‘downtime’ [19:1563]. For the worker, it is their lived time chunked up into a pulsating mess of alarms and nudges, distractions and panic; for the manager and employer, it is a vast predictive matrix which externalises everything that the model would prefer not to predict (such as apartment complexes that refuse to let gig workers on their elevators, or inadequately represented crossings and road

conditions) and funnels the cost to the worker. What the model refuses to count cannot surprise it.

4.3 The history of the measured

This disparity between the predictor and the predicted reprises over a century of labour struggle – throughout which the extraction of discretion has served as a crucial instrument for the extraction of labour power. Alessandro Delfanti suggests that automation *separates* the worker from their skills and knowledge – a separation which then enables the factory to operate with deskilled workers, and to turn that information around to evaluate and control worker behaviour [25]. In Amazon warehouses, human ‘stowers’ handle incoming items of staggering variety, grouping them into bins in giant ‘pick towers’ – a task that involves a certain degree of discretionary judgment. When this process is laced with data-gathering systems, however, each worker action – such as the stower’s scanning of the item, their rate of pieces per hour, their bathroom breaks – provides data grist for improving data-driven predictions of the stowing process. Individual judgments and tacit knowledge are siphoned into a unique source of knowledge for managers, but not for the workers themselves; as we have seen above, the question of what the rate really is at any given time is for the manager to know and the worker to wonder. The key innovation is not merely to pack seven boxes instead of six, but ensuring that it is the manager who can *set* the ‘rate’ of seven, or eight, or two hundred, rather than the worker.

This dynamic runs right through the history of modern automation. Harry Braverman’s *Labour and Monopoly Capital* argues that this separation was one of the key legacies of Charles Babbage, and his argument that any labour process should be subdivided into ultra-fine fragments, such that each fragment could be regulated for minimal cost and maximum labour [11:55]. Not only does this presage contemporary microwork, Babbagification also requires ever more detailed apparatuses of measurement and justification through which those fragments can be priced and punished. For instance, David Noble shows how General Electric’s postwar investment into machine shop automation was an explicit response to the success of nationwide union strikes in 1946, and the fear that skilled workers are able to exercise discretion over the conditions of their workplace and to make demands of the employer [70]. GE Vice President Lemuel Boulware, whose name would become the label for its aggressive anti-union strategies of ‘Boulwarism’, argued that the company must “eradicate” the “fantasy” that “the employees . . . were in the driver’s seat.” [70:156] Such attitudes shaped what kinds of automation technologies would be accepted, and which functioning alternatives would be consigned to obscurity. GE’s favoured ‘record-playback’ approach to automation eventually lost out to numeric control in part because the managers and employers making the acquisition decisions wanted automation systems to transfer discretionary power away from the worker. In one instance, GE designed an automation system for a steel company, only for the client to complain that “operators were [still] controlling production, determining the output”, engaging in well-worn techniques of ‘stints’ and ‘pacing’ (worker-determined production quotas and rates) [70:164].

Here we return to the earlier point that there is always a choice to be made as to what kind of data is and is not collected, and what kind of predictions are made available for decision-making. Amazon exemplifies how it is not just possible, but profitable to produce cutting-edge data systems when they serve its goals for extracting worker control and discretion, and to actively pursue incompetence and negligence when it does not. During much of the ongoing COVID-19 pandemic, Amazon warehouses have refused to collect and disclose basic data about which of its workers had tested positive for the coronavirus. In one warehouse in Pennsylvania, managers initially informed workers of each new case, before stopping after around 60 – claiming that sharing the information ‘made no difference’. [99] One Amazon VP insisted that “it’s not a particularly useful number”, but that their warehouses’ rate of infections was “just under” community rates [71]. Such rhetoric draws from Big Tech’s well-established playbook of professing ignorance about inconvenient kinds of data, while jamming public debate with unverifiable statistical claims. In this same period, Amazon has moved quickly to expand its collection of data that it considers more ‘useful’ – including, for example, efforts to score and measure ‘risk of unionisation’ by employees in its Whole Foods subsidiary [76], and installing Pinkerton spies in warehouses around the world [42].

A company like Amazon – an emblem of our age for the promise of scaling data and technology to heretofore unimaginable wealth – thus exemplifies how the commitment to prediction and calculation is never universally consistent. When predictive systems are hailed as efficient or sophisticated, this describes not so much their ability to objectively measure the realities of production, but their ability to discount the less profitable aspects of that reality. In all this, Amazon faithfully follows the long tradition of Silicon Valley’s production industries, where low-paid, primarily migrant workers have long suffered from toxic exposures and resulting illnesses, and the data around those conditions actively delegitimised and suppressed by their tech employers [75:100-1].

These recurrent efforts to separate the worker from their discretion, and to wield prediction as an asymmetric instrument of control rather than collective transparency, come hand in hand with entrenching cultural tropes about the nature of the worker and what is *necessary* to make them work. As with crime and policing, such presumptions enter into the social life of prediction before and beyond any question of statistical bias in a dataset or the appropriateness of particular object labels. One persistent pattern across many seemingly disparate applications for predictive technologies today is that the promotion of ‘cutting-edge AI solutions’ leverages and amplifies a pre-existing fantasy that *criminalises* the targets of prediction. The worker is presumed to be, by default, a potential thief (of wages via low productivity, or directly of company property). Amazon has consistently cited package theft by delivery drivers as a key motivation for its Ring door cameras [108], pushing that fear onto consumers as a way to sell the devices [15:5]. This is now extending to driver surveillance. In March 2021, Amazon began forcing its drivers to sign away ‘biometric consent’ to keep their jobs, authorising a new suite of surveillance tools including claims of AI detection for ‘risky driving behaviours’, in addition to existing means for tracking and evaluating their rate of delivery

[98]. Perhaps the most strikingly explicit example of worker criminalisation as incentive for predictive technologies comes from an anonymous Microsoft engineer, who describes traveling to Kazakhstan to advise an oil production partnership between the Kazakh state and Chevron. After giving a presentation on possible uses of AI/ML technologies for mapping the ground to identify oil sources, the engineer finds that the managers in the room – mostly American expats – were actually far more interested in AI-driven worker surveillance. Their presumption was clear: “We have a lot of workers in the oil fields. It would be nice to know where they are and what they are doing [:::] If they are doing anything at all”, said one. [104] The worker is the suspect, and it is this *a priori* declaration that determines what role the data will play to begin with.

5 PREDICTING TO STAY THE SAME

Prediction grammatises – renders flexibly replicable, habituates, provides a template for – a widespread extraction of discretionary power: the spaces of practical ambiguity, the gap between rule and case, the moments of situational judgment, that were always unequally distributed across different subjects in the first place. From predictive approaches in policing and incarceration to their analogues in workplace surveillance, we find a consistent pattern. It is the *subjects* of measurement are preemptively defined as objects of suspicion and danger, whose exercise of discretionary power over their own circumstances is primarily seen as a source of unwelcome uncertainty. All this ‘seeing’ is done from the perspective of the managerial *clients* of prediction, for whom systems of datafication offer new ways to entrench their desired definitions of a productive worker or otherwise ‘good’ subject as neutral and objective facts.

When understood as a transfer and entrenchment of discretionary power, data-driven prediction appears more clearly as a literally *conservative* extension of longstanding strategies for social control, rather than any radically novel means for reducing human bias and irrationality. In *Forces of Automation*, David Noble relays a definition of workplace automation from none other than Peter Drucker, perhaps the closest available personification of corporate managerial logic: “what is today called ‘automation’ is conceptually a logical extension of Taylor’s scientific management [in which] productivity required that ‘doing’ be divorced from ‘planning’.” [70:231] Far from revolutionising this long project of worker deprivation, predictive technologies entrench them. In police surveillance and risk assessment, too, what matters is the choice of which numbers are collected in the first place: technology like body cameras can theoretically deliver impartial evidence around police conduct – but not when the police retain and exercise the practical power to turn cameras off citing ‘malfunction’, or refuse to release timely and unredacted footage [74, 106]. As Mimi Onuoha has observed, the abundance of video footage and other data around police violence against Black people in the United States has failed to lead to a ‘great awakening’ around institutional racism, because data can never speak truth to power on its own [58, 73].

To adapt Joseph Weizenbaum’s famous dictum on computing, not everything that can be predicted should be predicted [100:x]. The very act of declaring something predictable, of turning it into a prediction ‘problem’, already establishes the conditions foreextraction. Turning a remote education situation, or the allocation of police resources in a city, into a prediction problem means deploying a

well rehearsed template for decision-making in which the clients of prediction are empowered to consistently speak over and speak for the targets of prediction: the entrepreneur and venture capitalist over the students themselves, the police departments over inhabitants of overpoliced neighbourhoods, the managers and executives over warehouse workers. As technologies like facial recognition increasingly impose themselves over diverse social situations, they deliver not a hyper-rationalised environment in which the data speaks for itself, but a thin refurbishing of familiar asymmetries in which the same few might continue to judge the many with impunity.

ACKNOWLEDGMENTS

This paper is an accidental byproduct of research conducted through the Canadian Social Sciences & Humanities Research Council grant 430-2019-00078, “Personal Truthmaking in the Data-driven Society”.

REFERENCES

- [1] Blaise Agüera y Arcas, Margaret Mitchell, and Alexander Todorov. 2017. Physiognomy’s New Clothes. *Medium*. Retrieved from <https://medium.com/@blaisea/physiognomys-new-clothes-f2d4b59fdd6a>
- [2] Ali Alkhatib. 2021. To Live in Their Utopia: Why Algorithmic Systems Create Absurd Outcomes. In *CHI Conference of Human Factors in Computing*. ACM, Yokohama, Japan.
- [3] Louise Amoore. 2020. *Cloud Ethics: Algorithms and the Attributes of Ourselves and Others*. Duke University Press, Pittsburgh.
- [4] Boone Ashworth. 2020. Inside Citizen, the App That Asks You to Report on the Crime Next Door. *Wired*. Retrieved from <https://www.wired.com/story/citizen/>
- [5] Ruha Benjamin. 2019. *Race After Technology: Abolitionist Tools for the New Jim Code*. Polity Press, Cambridge.
- [6] Elettra Bietti. 2020. From Ethics Washing to Ethics Bashing. In *FAT* ’20: Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, Barcelona.
- [7] Abeba Birhane. 2021. The Impossibility of Automating Ambiguity. *Artif. Life* 27, (2021), 1–18.
- [8] Abeba Birhane, Pratyusha Kalluri, Dallas Card, William Agnew, Ravit Dotan, and Michelle Bao. 2021. The Values Encoded in Machine Learning Research. (2021). Retrieved from <http://arxiv.org/abs/2106.15590>
- [9] Luc Boltanski and Eve Chiapello. 2007. *The New Spirit of Capitalism*. Verso, London.
- [10] Luc Boltanski and Laurent Thevenot. 2006. *On Justification: Economies of Worth*. Princeton University Press, Princeton.
- [11] Harry Braverman. 1998. *Labor and Monopoly Capital: The Degradation of Work in the Twentieth Century* (25th Anniv ed.). Monthly Review Press, New York.
- [12] Sarah Brayne. 2017. Big Data Surveillance: The Case of Policing. *Am. Sociol. Rev.* 82, 5 (2017), 977–1008.
- [13] Sarah Brayne. 2021. *Predict and Surveil*. Oxford University Press, Oxford.
- [14] Sarah Brayne and Angele Christin. 2021. Technologies of Crime Prediction: The Reception of Algorithms in Policing and Criminal Courts. *Soc. Probl.* 68, 3 (2021), 608–624.
- [15] Lauren Bridges. 2021. Infrastructural obfuscation: unpacking the carceral logics of the Ring surveillant assemblage. *Information, Commun. Soc.* 24, 6 (2021), 830–849.
- [16] Simone Browne. 2015. *Dark Matters: On the Surveillance of Blackness*. Duke University Press, Pittsburgh. DOI:<https://doi.org/10.1215/9780822375302>
- [17] Ingrid Burrington. 2015. What Amazon Taught the Cops. *The Nation*. Retrieved from <https://www.thenation.com/article/archive/what-amazon-taught-cops/>
- [18] Alexander Campolo and Kate Crawford. 2020. Enchanted Determinism: Power without Responsibility in Artificial Intelligence. *Engag. Sci. Technol. Soc.* 6, (2020), 1.
- [19] Julie Yujie Chen and Sun Ping. 2020. Temporal arbitrage, fragmented rush, and opportunistic behaviors: The labor politics of time in the platform economy. *New Media Soc.* 22, 9 (2020), 1561–1579.
- [20] Sasha Constanza-Chock. 2020. *Design Justice: Community-led Practices to Build the Worlds We Need*. MIT Press, Cambridge.
- [21] Karine Côté-Boucher. 2016. The Paradox of Discretion: Customs and the Changing Occupational Identity of Canadian Border Officers. *Br. J. Criminol.* 56, (2016), 49–67. DOI:<https://doi.org/10.1093/bjc/azv023>
- [22] Kallie Cox and William H Freivogel. 2021. Police misconduct records secret, difficult to access. *AP News*. Retrieved from <https://apnews.com/article/us-news-police-reform-police-government-and-politics-fa6cbd7e017b85aa715e23465a90abbe>
- [23] Lorraine J Daston and Peter Galison. 2007. *Objectivity*. Zone Books, New York.
- [24] Alessandro Delfanti. 2021. *The Warehouse: Workers and Robots at Amazon*. Pluto Press, London.
- [25] Alessandro Delfanti. 2021. Machinic dispossession and augmented despotism: Digital work in an Amazon warehouse. *New Media Soc.* 23, 1 (2021), 39–55.
- [26] Lina Dencik and Sanne Stevens. 2021. Regimes of justification in the datafied workplace: The case of hiring. *New Media Soc.* (2021).
- [27] William Deringer. 2018. *Calculated Values: Finance, Politics, and the Quantitative Age*. Harvard University Press, Cambridge.
- [28] Julia Dressel and Hany Farid. 2018. The accuracy, fairness, and limits of predicting recidivism. *Sci. Adv.* 4, 1 (2018).
- [29] Ronald Dworkin. 1963. Judicial Discretion. *J. Philos.* 60, 21 (1963), 624–638.
- [30] Virginia Eubanks. 2018. *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor*. St. Martin’s Press, New York.
- [31] Adam Foss, Samuel Sinyangwe, Julia Angwin, Charmaine Arthur, and Kim Foxx. 2017. Automating (In)justice: Policing and Sentencing in the Algorithm Age. In *Data 4 Black Lives*, Cambridge.
- [32] Michel Foucault. 2015. *The Punitive Society: Lectures at the College de France 1972-1973*. Palgrave Macmillan, New York.
- [33] Maya Indra Ganesh. 2020. The ironies of autonomy. *Humanit. Soc. Sci. Commun.* 7, 1 (2020), 1–10.
- [34] Kelly A Gates. 2011. *Our Biometric Future: Facial Recognition Technology and the Culture of Surveillance*. New York University Press, New York.
- [35] Kelly A Gates. 2015. Big Data and State Transparency: What the Absence of Data on Police Killings Reveals. In *Democracy, Citizenship and Constitutionalism*, Philadelphia. Retrieved from <https://www.sas.upenn.edu/dcc/event/kelly-gates>
- [36] Kelly A Gates. 2019. Counting the Uncounted: What the Absence of Data on Police Killings Reveals. In *Digital Media and Democratic Futures*, Michael X Delli Carpini (ed.). University of Pennsylvania Press, Philadelphia, 121–142.
- [37] David Golumbia. 2009. *The Cultural Logic of Computation*. Harvard University Press, Cambridge.
- [38] Ben Green. 2019. “Good” isn’t good enough. In *NeurIPS 2019, AI for Social Good Workshop*, Vancouver.
- [39] Daniel Greene, Anna Lauren Hoffmann, and Luke Stark. 2019. Better, Nicer, Clearer, Fairer: A Critical Assessment of the Movement for Ethical Artificial Intelligence and Machine Learning. *Proc. 52nd Hawaii Int. Conf. Syst. Sci.* (2019). DOI:<https://doi.org/10.24251/hicss.2019.258>
- [40] Jay Greene and Chris Alcantara. 2021. Amazon warehouse workers suffer serious injuries at higher rates than other firms. *The Washington Post*. Retrieved from <https://www.washingtonpost.com/technology/2021/06/01/amazon-osh-injury-rate/>
- [41] Karen Gregory and Jathan Sadowski. 2021. Biopolitical platforms: the perverse virtues of digital labour. *J. Cult. Econ.* 14, 6 (2021), 1–13.
- [42] Lauren Kaori Gurley. 2020. Secret Amazon Reports Expose the Company’s Surveillance of Labor and Environmental Groups. *Motherboard*. Retrieved from <https://www.vice.com/en/article/5dp3yn/amazon-leaked-reports-expose-spying-warehouse-workers-labor-union-environmental-groups-social-movements>
- [43] Lauren Kaori Gurley. 2021. Amazon Furniture Assembly Service Adds New Chaos to Delivery Workers’ Lives. *Motherboard*. Retrieved from <https://www.vice.com/en/article/n7begd/amazon-furniture-assembly-service-adds-new-chaos-to-delivery-workers-lives>
- [44] Ian Hacking. 1992. Statistical Language, Statistical Truth and Statistical Reason: The Self-Authentification of a Style of Scientific Reasoning. In *Social Dimensions of Science*, Ernan McMullin (ed.). University of Notre Dame Press, Notre Dame, 130–157.
- [45] Alexandra Hall. 2017. Decisions at the data border: Discretion, discernment and security. *Secur. Dialogue* 48, 6 (2017), 488–504.
- [46] Orit Halpern. 2014. *Beautiful Data: A History of Vision and Reason since 1945*. Duke University Press, Durham.
- [47] Peter Hancock. 2020. State high court says police misconduct records cannot be destroyed. *Capitol News Illinois*. Retrieved from <https://capitolnewsillinois.com/NEWS/state-high-court-says-police-misconduct-records-cannot-be-destroyed>
- [48] Bernard Harcourt. 2007. *Against Prediction - Profiling, Policing, and Punishing in an Actuarial Age*. Chicago University Press, Chicago.
- [49] Drew Harwell. 2019. Doorbell-camera firm Ring has partnered with 400 police forces, extending surveillance concerns. *The Washington Post*. Retrieved from <https://www.washingtonpost.com/technology/2019/08/28/doorbell-camera-firm-ring-has-partnered-with-police-forces-extending-surveillance-reach/>
- [50] Taylor Hatmaker. 2020. AI researchers condemn predictive crime software, citing racial bias and flawed methods. *TechCrunch*. Retrieved from <https://techcrunch.com/2020/06/23/ai-crime-prediction-open-letter-springer/>
- [51] Will Douglas Heaven. 2021. Hundreds of AI tools have been built to catch covid. None of them helped. *MIT Technology Review*. Retrieved from <https://www.technologyreview.com/2021/07/30/1030329/machine-learning-ai-failed-covid-hospital-diagnosis-pandemic/>

- [52] Kashmir Hill. 2020. The Secretive Company That Might End Privacy as We Know It. *New York Times*. Retrieved January 18, 2020 from <https://www.nytimes.com/2020/01/18/technology/clearview-privacy-facial-recognition.html>
- [53] Sam Hind. 2021. Making decisions: the normal interventions of Nissan 'mobility managers'. *Mobilities* 00, 00 (2021), 1–17. DOI:<https://doi.org/10.1080/17450101.2021.1988682>
- [54] Anna Lauren Hoffmann. 2019. Where fairness fails: data, algorithms, and the limits of antidiscrimination discourse. *Inf. Commun. Soc.* 22, 7 (2019), 900–915.
- [55] Sarah Holder. 2020. How Nextdoor Courts Police and Public Officials. *Bloomberg*. Retrieved from <https://www.bloomberg.com/news/articles/2020-05-21/nextdoor-s-delicate-partnership-with-local-police>
- [56] Sun-ha Hong. 2020. *Technologies of Speculation: The limits of knowledge in a data-driven society*. New York University Press, New York.
- [57] Sun-ha Hong. 2021. Technofutures in Stasis: Smart Machines, Ubiquitous Computing, and the Future That Keeps Coming Back. *Int. J. Commun.* 15, (2021), 1940–1960.
- [58] Sun-ha Hong. 2021. Why Transparency Won't Save Us. *CIGI*. Retrieved from <https://www.cigionline.org/articles/why-transparency-wont-save-us/>
- [59] Jenny Huberman. 2021. Amazon Go, surveillance capitalism, and the ideology of convenience. *Econ. Anthropol.* (2021), 1–13.
- [60] Issa Kohler-Hausmann. 2018. *Misdemeanor Land: Criminal Courts and Social Control in an Age of Broken Windows Policing*. Princeton University Press, Princeton.
- [61] Colin Lecher. 2019. How Amazon automatically tracks and fires warehouse workers for 'productivity.' *The Verge*. Retrieved from <https://www.theverge.com/2019/4/25/18516004/amazon-warehouse-fulfillment-centers-productivity-firing-terminations>
- [62] Dong kyung Lee. 2021. [스트레이트] 얼마나 더... 멈추지 않는 쿠팡 노동자들의 죽음. *MBC News*. Retrieved from https://imnews.imbc.com/replay/straight/6139311_28993.html
- [63] Michael Lipsky. 2010. *Street-Level Bureaucracy: Dilemmas of the Individual in Public Services* (30th Anniv ed.). Russell Sage Foundation, New York.
- [64] Alec MacGillis. 2021. *Fulfillment: Winning and Losing in One-Click America*. Farrar, Straus and Giroux, New York.
- [65] Carolyn Marvin. 1988. *When Old Technologies Were New: Thinking About Electric Communication in the Late Nineteenth Century*. Oxford University Press, Oxford.
- [66] Tara McPherson. 2012. U.S. Operating Systems at Mid-Century: The Intertwining of Race and Unix. In *Race After the Internet*, Lisa Nakamura and Peter Chow-White (eds.). Routledge, New York, 21–37.
- [67] Dan McQuillan. 2018. Data Science as Machinic Neoplatonism. *Philos. Technol.* 31, 2 (2018), 253–272.
- [68] Evgeny Morozov. 2019. Capitalism's New Clothes. *The Baffler*. Retrieved from <https://thebaffler.com/latest/capitalisms-new-clothes-morozov>
- [69] Arvind Narayanan. 2019. How to recognize AI snake oil. Cambridge. Retrieved from <https://t.co/iCpyFw5uN?amp=1>
- [70] David F Noble. 2011. *Forces of Production: A Social History of Industrial Automation*. Transaction Publishers, New Brunswick.
- [71] Sara Ashley O'Brien. 2020. Amazon insists sharing data on coronavirus cases in its warehouses isn't useful. *CNN Business*. Retrieved from <https://www.cnn.com/2020/05/15/tech/amazon-warehouse-coronavirus-cases-data/index.html>
- [72] Edward Ongweso Jr. 2021. Amazon's New Algorithm Will Set Workers' Schedules According to Muscle Use. *Motherboard*. Retrieved from <https://www.vice.com/en/article/z3xeba/amazons-new-algorithm-will-set-workers-schedules-according-to-muscle-use>
- [73] Mimi Onuoha. 2020. When Proof Is Not Enough. *FiveThirtyEight*. Retrieved from <https://fivethirtyeight.com/features/when-proof-is-not-enough/>
- [74] Alex Pasternak. 2020. We spent a fortune on police body cameras. What have they done to policing? *Fast Company*. Retrieved from <https://www.fastcompany.com/90514774/we-spent-a-fortune-on-police-body-cams-why-havent-they-fixed-policing>
- [75] David Naguib Pellow and Lisa Sun-Hee Park. 2002. *The Silicon Valley of Dreams: Environmental Injustice, Immigrant Workers, and the High-Tech Global Economy*. New York University Press, New York.
- [76] Hayley Peterson. 2020. Amazon-owned Whole Foods is quietly tracking its employees with a heat map tool that ranks which stores are most at risk of unionizing. *Business Insider*. Retrieved from <https://www.businessinsider.com/whole-foods-tracks-unionization-risk-with-heat-map-2020-1>
- [77] Thao Phan, Jake Goldenfein, Monique Mann, and Declan Kuch. 2021. Economies of Virtue: The Circulation of 'Ethics' in Big Tech. *Sci. Cult. (Lond)*, 0, 0 (2021), 1–15. DOI:<https://doi.org/10.1080/09505431.2021.1990875>
- [78] Mary Poovey. 1998. *A History of the Modern Fact: Problems of Knowledge in the Sciences of Wealth and Society*. University of Chicago Press, Chicago.
- [79] Theodore M Porter. 1995. *Trust in Numbers: The Pursuit of Objectivity in Science and Public Life*. Princeton University Press, Princeton.
- [80] Vinay Uday Prabhu and Abeba Birhane. 2020. Large image datasets: A pyrrhic win for computer vision? (2020). Retrieved from <https://arxiv.org/abs/2006.16923>
- [81] Inioluwa Deborah Raji, Emily M. Bender, Amandalynne Paullada, Emily Denton, and Alex Hanna. 2021. AI and the Everything in the Whole Wide World Benchmark. In *NeurIPS 2021*.
- [82] Joshua Reeves and Jeremy Packer. 2013. Police Media: The Governance of Territory, Speed, and Communication. *Commun. Crit. Cult. Stud.* 10, 4 (2013), 359–384.
- [83] Morgan Klaus Scheuerman, Emily Denton, and Alex Hanna. 2021. Do Datasets Have Politics? Disciplinary Values in ComputerVision Dataset Development. In *Proc. ACM Hum.-Comput. Interact.*
- [84] James C. Scott. 2021. Further reflections on seeing like a state. *Polity* 53, 3 (2021), 507–514.
- [85] James C Scott. 1998. *Seeing like a State: How Certain Schemes to Improve the Human Condition Have Failed*. Yale University Press, New Haven.
- [86] Kap Seol. 2021. Coupang, the South Korean Amazon, Is Chewing Up and Spitting Out Its Workers. *Jacobin*. Retrieved from <https://jacobinmag.com/2021/04/coupang-south-korea-amazon-delivery-workers/>
- [87] Steven Shapin and Simon Schaeffer. 1985. *Leviathan and the Air-Pump: Hobbes, Boyle, and the Experimental Life*. Princeton University Press, Princeton.
- [88] Aaron Shapiro. 2017. Reform predictive policing. *Nature* 541, 7638 (2017), 458–460.
- [89] Aaron Shapiro. 2020. *Design, Control, Predict: Logistical Governance in the Smart City*. University of Minnesota Press, Minneapolis.
- [90] Amanda Shendruk. 2020. Amazon has become the world's fifth largest employer. *World Economic Forum*. Retrieved from <https://www.weforum.org/agenda/2020/12/amazon-employers-online-retail-covid-pandemic-coronavirus/>
- [91] David Sherfinski and Avi Asher-Schapiro. 2021. U.S. prisons mull AI to analyze inmate phone calls. *Reuters*. Retrieved from <https://mobile.reuters.com/article/amp/idUSKBN2FA000>
- [92] Aaron Smith and Monica Anderson. 2017. *Automation in Everyday Life*. Washington D.C. Retrieved from <https://www.pewresearch.org/internet/2017/10/04/automation-in-everyday-life/>
- [93] Spencer Soper. 2021. Fired by Bot at Amazon: 'It's You Against the Machine.' *Bloomberg*. Retrieved from <https://www.bloomberg.com/news/features/2021-06-28/fired-by-bot-amazon-turns-to-machine-managers-and-workers-are-losing-out>
- [94] Luke Stark and Jevan Hutson. Physiognomic Artificial Intelligence. *Fordham Intelect. Prop. Media Entertain. Law J.*
- [95] Marilyn Strathern. 2020. *Relations*. Duke University Press, Durham.
- [96] Daniel Trotter. 2017. Digital Vigilantism as Weaponisation of Visibility. *Philos. Technol.* 30, 1 (2017), 55–72.
- [97] Nantina Vgontzas. 2021. Toward Degrowth: Worker Power, Surveillance Abolition, and Climate Justice at Amazon. *New Glob. Stud.* December (2021).
- [98] James Vincent. 2021. Amazon delivery drivers have to consent to AI surveillance in their vans or lose their jobs. *The Verge*. Retrieved from <https://www.theverge.com/2021/3/24/22347945/amazon-delivery-drivers-ai-surveillance-cameras-vans-consent-form>
- [99] Karen Weise. 2020. 'Way Too Late': Inside Amazon's Biggest Outbreak. *The New York Times*. Retrieved from <https://www.nytimes.com/2020/05/19/technology/amazon-coronavirus-workers.html>
- [100] Joseph Weizenbaum. 1976. *Computer Power and Human Reason: From Judgment to Calculation*. W.H. Freeman and Company, New York.
- [101] Meredith Whittaker. 2021. The Steep Cost of Capture. *ACM Interact.* November-December (2021), 51–55.
- [102] James Wright. 2021. Suspect AI: Vibramage, Emotion Recognition Technology and Algorithmic Opacity. *Sci. Technol. Soc.* (2021).
- [103] Elvin Wyly. 2014. Automated (post)positivism. *Urban Geogr.* 35, 5 (2014), 669–690.
- [104] Zero Cool. 2019. Oil is the New Data. *Logic* 9 (2019). Retrieved from <https://logicmag.io/nature/oil-is-the-new-data/>
- [105] Baobao Zhang and Allan Dafoe. 2019. *Artificial Intelligence: American Attitudes and Trends*. Oxford. Retrieved from <https://governanceai.github.io/US-Public-Opinion-Report-Jan-2019/>
- [106] Ethan Zuckerman. 2020. Why filming police violence has done nothing to stop it. *MIT Technology Review*. Retrieved from <https://www.technologyreview.com/2020/06/03/1002587/sousveillance-george-floyd-police-body-cams/>
- [107] Machine Learning Glossary. *Google*. Retrieved from <https://developers.google.com/machine-learning/glossary#p>
- [108] 2018. Amazon and Ring Close Acquisition—Now Working Together to Empower Neighbors with Affordable Ways to Monitor their Homes and Reduce Crime in Neighborhoods. *Amazon*. Retrieved from <https://press.aboutamazon.com/news-releases/news-release-details/amazon-and-ring-close-acquisition-now-working-together-empower>
- [109] 2019. Palantir played key role in arresting families for deportation, document shows. *Mijente*. Retrieved from <https://mijente.net/2019/05/palantir-arresting-families/>
- [110] 2019. Documents Obtained by EPIC show Idaho's Use of Subjective Categories in Calculating Risk. *EPIC*. Retrieved from <https://epic.org/documents-obtained-by-epic-show-idahos-use-of-subjective-categories-in-calculating-risk/>