


Abolish the #TechToPrisonPipeline
Crime prediction technology reproduces injustices and causes real harm.



Springer Publishing
Berlin, Germany
+49 (0) 6221 487 0
customerservice@springernature.com

RE: A Deep Neural Network Model to Predict Criminality Using Image Processing

June 22, 2020

Dear Springer Editorial Committee,

We write to you as expert researchers and practitioners across a variety of technical, scientific, and humanistic fields (including statistics, machine learning and artificial intelligence, law, sociology, history, communication studies and anthropology). Together, we share grave concerns regarding a forthcoming publication entitled “A Deep Neural Network Model to Predict Criminality Using Image Processing.” According to a recent press release, this article will be published in your book series, “Springer Nature — Research Book Series: Transactions on Computational Science and Computational Intelligence.”

We urge:

- The review committee to publicly rescind the offer for publication of this specific study, along with an explanation of the criteria used to evaluate it.
- Springer to issue a statement condemning the use of criminal justice statistics to predict criminality, and acknowledging their role in incentivizing such harmful scholarship in the past.
- All publishers to refrain from publishing similar studies in the future.

This upcoming publication warrants a collective response because it is emblematic of a larger body of computational research that claims to identify or predict “criminality” using biometric and/or criminal legal data.[1] Such claims are based on unsound scientific premises, research, and methods, which numerous studies spanning our respective disciplines have debunked over the years.[2] Nevertheless, these discredited claims continue to resurface, often under the veneer of new and purportedly neutral statistical methods such as machine learning, the primary method of the publication in question.[3] In the past decade, government officials have embraced machine learning and artificial intelligence (AI) as a means of depoliticizing state violence and reasserting the legitimacy of the carceral state, often amid significant social upheaval.[4] Community organizers and Black scholars have been at the forefront of the resistance against the use of AI technologies by law enforcement, with a particular focus on facial recognition.[5] Yet these voices continue to be marginalized, even as industry and the academy invests significant resources in building out “fair, accountable and transparent” practices for machine learning and AI.[6]

Part of the appeal of machine learning is that it is highly malleable — correlations useful for prediction or detection can be rationalized with any number of plausible causal mechanisms. Yet the way these studies are ultimately represented and interpreted is profoundly shaped by the political economy of data science[7] and their contexts of use.[8] Machine learning programs are not neutral; research agendas and the data sets they work with often inherit dominant cultural beliefs about the world. These research agendas reflect the incentives and perspectives of those in the privileged position of developing machine learning models, and the data on which they rely. The uncritical acceptance of default assumptions inevitably leads to discriminatory design in algorithmic systems, reproducing ideas which normalize social hierarchies and legitimize violence against marginalized groups.[9]

Such research does not require intentional malice or racial prejudice on the part of the researcher.[10] Rather, it is the expected by-product of any field which evaluates the quality of their research almost exclusively on the basis of “predictive performance.”[11] In the following sections, we outline the specific ways crime prediction technology reproduces, naturalizes and amplifies discriminatory outcomes, and why exclusively technical criteria are insufficient for evaluating their risks.

I. Data generated by the criminal justice system cannot be used to “identify criminals” or predict criminal behavior. Ever.

In the original press release published by Harrisburg University, researchers claimed to “predict if someone is a criminal based solely on a picture of their face,” with “80 percent accuracy and with no racial bias.” Let’s be clear: there is no way to develop a system that can predict or identify “criminality” that is not racially biased — because the category of “criminality” itself is racially biased.[12]

Research of this nature — and its accompanying claims to accuracy — rest on the assumption that data regarding criminal arrest and conviction can serve as reliable, neutral indicators of underlying criminal activity. Yet these records are far from neutral. As numerous scholars have demonstrated, historical court and arrest data reflect the policies and practices of the criminal justice system. These data reflect who police choose to arrest, how judges choose to rule, and which people are granted longer or more lenient sentences.[13] Countless studies have shown that people of color are treated more harshly than similarly situated white people at every stage of the legal system, which results in serious distortions in the data.[14] Thus, any software built within the existing criminal legal framework will inevitably echo those same prejudices and fundamental inaccuracies when it comes to determining if a person has the “face of a criminal.”

These fundamental issues of data validity cannot be solved with better data cleaning or more data collection.[15] Rather, any effort to identify “criminal faces” is an application of machine learning to a problem domain it is not suited to investigate, a domain in which context and causality are essential and also fundamentally misinterpreted. In other problem domains where machine learning has made great progress, such as common object classification or facial verification, there is a “ground truth” that will validate learned models.[16] The causality underlying how different people perceive the content of images is still important, but for many tasks, the ability to demonstrate face validity is sufficient.[17] As Narayanan (2019) notes, “the fundamental reason for progress [in these areas] is that there is no uncertainty or ambiguity in these tasks — given two images of faces, there’s ground truth about whether or not they represent the same person.”[18] However, no such pattern exists for facial features and criminality, because having a face that looks a certain way does not cause an individual to

commit a crime — there simply is no “physical features to criminality” function in nature.[19] Causality is tacitly implied by the language used to describe machine learning systems. An algorithm’s so-called “predictions” are often not actually demonstrated or investigated in out-of-sample settings (outside the context of training, validation, and testing on an inherently limited subset of real data), and so are more accurately characterized as “the strength of correlations, evaluated retrospectively,”[20] where real-world performance is almost always lower than advertised test performance for a variety of reasons.[21]

Because “criminality” operates as a proxy for race due to racially discriminatory practices in law enforcement and criminal justice, research of this nature creates dangerous feedback loops.[22] “Predictions” based on finding correlations between facial features and criminality are accepted as valid, interpreted as the product of intelligent and “objective” technical assessments.[23] In reality, these “predictions” materially conflate the shared, social circumstances of being unjustly overpoliced with criminality. Policing based on such algorithmic recommendations generates more data that is then fed back into the system, reproducing biased results.[24] Ultimately, any predictive algorithms that are based on these widespread mischaracterizations of criminal justice data justifies the exclusion and repression of marginalized populations through the construction of “risky” or “deviant” profiles.[25]

II. Technical measures of “fairness” distract from fundamental issues regarding an algorithm’s validity.

Studies like the aforementioned reflect a growing crisis of validity in AI and machine learning research that’s plagued the field for decades.[26] This crisis stems from the fact that machine learning scholars are rarely trained in the critical methods, frameworks, and language necessary to interrogate the cultural logics and implicit assumptions underlying their models. Nor are there ample incentives to conduct such interrogations, given the industrial incentives that are driving much machine learning research and development.[27] To date, many efforts to deal with the ethical stakes of algorithmic systems have centered mathematical definitions of fairness that are grounded in narrow notions of bias and accuracy.[28] These efforts give the appearance of rigor, while distracting from more fundamental epistemic problems.

Designers of algorithmic systems need to embrace a historically grounded, process-driven approach to algorithmic justice, one that explicitly recognizes the active and crucial role that the data scientist (and the institution they’re embedded in) plays in constructing meaning from data.[29] Computer scientists can benefit greatly from ongoing methodological debates and insights gleaned from fields such as anthropology, sociology, media and communication studies, and science and technology studies, disciplines in which scholars have been working for decades to

develop more robust frameworks for understanding their work as situated practice, embedded in uncountably infinite[30] social and cultural contexts.[31] While many groups have made efforts to translate these insights to the field of computer science, it remains to be seen whether these critical approaches will be widely adopted by the computing community.[32]

Machine learning practitioners must move beyond the dominant epistemology of computer science, in which the most important details of a model are considered those that survive abstraction to “pure” technical problems, relegating social issues to “implementation details.”[33] This way of regarding the world biases research outputs towards narrowly technical visions of progress: accuracy, precision and recall or sensitivity and specificity, F-score, Jaccard index, or other performance metric of choice, all applied to an ever-growing set of applications and domains. Machine learning does not have a built-in mechanism for investigating or discussing the social and political merits of its outputs. Nor does it have built-in mechanisms for critically exploring the relationship between the research they conduct and the researchers’ own subject positions, group memberships, or the funding sources that make their research possible. In other words, reflexivity is not a part of machine learning’s objective function.

If machine learning is to bring about the “social good” touted in grant proposals and press releases, researchers in this space must actively reflect on the power structures (and the attendant oppressions) that make their work possible. This self-critique must be integrated as a core design parameter, not a last-minute patch. The field of machine learning is in dire need of a critical reflexive practice.

III. Conclusion: Crime-prediction technology reproduces injustices and causes real harm

Recent instances of algorithmic bias across race, class, and gender have revealed a structural propensity of machine learning systems to amplify historic forms of discrimination, and have spawned renewed interest in the ethics of technology and its role in society. There are profound political implications when crime prediction technologies are integrated into real world applications, which go beyond the frame of “tech ethics” as currently defined.[34] At the forefront of this work are questions about power[35]: who will be adversely impacted by the integration of machine learning within existing institutions and processes?[36] How might the publication of this work and its potential uptake legitimize, incentivize, monetize, or otherwise enable discriminatory outcomes and real-world harm?[37] These questions aren’t abstract. The authors of the Harrisburg University study make explicit their desire to provide “a significant advantage for law enforcement agencies and other intelligence agencies to prevent crime” as a co-author and former NYPD police officer outlined in the original press release.[38]

At a time when the legitimacy of the carceral state, and policing in particular, is being challenged on fundamental grounds in the United States, there is high demand in law enforcement for research of this nature, research which erases historical violence and manufactures fear through the so-called prediction of criminality. Publishers and funding agencies serve a crucial role in feeding this ravenous maw by providing platforms and incentives for such research. The circulation of this work by a major publisher like Springer would represent a significant step towards the legitimization and application of repeatedly debunked, socially harmful research in the real world.

To reiterate our demands, the review committee must publicly rescind the offer for publication of this specific study, along with an explanation of the criteria used to evaluate it. Springer must issue a statement condemning the use of criminal justice statistics to predict criminality and acknowledging their role in incentivizing such harmful scholarship in the past. Finally, all publishers must refrain from publishing similar studies in the future.

Sincerely,

2425 professors, researchers, practitioners, and students spanning the fields of anthropology, sociology, computer science, law, science and technology studies, information science, mathematics, and more. See the full list [here](#).

Footnotes

¹ Scholars use a variety of terms in reference to the prediction of criminal outcomes. Some researchers claim to predict “anti-social” or “impulsive” behavior. Others model “future recidivism” or an individual’s “criminal tendencies.” All of these terms frame criminal outcomes as the byproduct of highly individualized and proximate risk factors. As Prins and Reich (2018) argue, these predictive models neglect population drivers of crime and criminal justice involvement (Seth J. Prins, and Adam Reich. 2018. “Can we avoid reductionism in risk reduction?” *Theoretical criminology* 22 (2): 258–278). The hyper-focus on individualized notions of crime leads to myopic social reforms that intervene exclusively on the supposed cultural, biological and cognitive deficiencies of criminalized populations. This scholarship not only provides a mechanism for the confinement and control of the “dangerous classes,” but also creates the very processes through which these populations are turned into deviants to be controlled and feared. As Robert Vargas (2020) argues, this type of scholarship “sees Black people and Black communities as in need of

being fixed. This approach is not new but is rather the latest iteration in a series of efforts to improve cities by managing Black individuals instead of ending the police violence Black communities endure.” Robert Vargas. 2020. “It’s Time to Think Critically about the UChicago Crime Lab.” *The Chicago Maroon* June 11. (Accessed June 17, 2020). For examples of this type of criminalizing language see generally: Mahdi Hashemi and Margeret Hall. 2020. “Criminal tendency detection from facial images and the gender bias effect.” *Journal of Big Data*. 7 (2) . Eyal Aharoni, et al. 2013. “Neuroprediction of future rearrest.” *Proceedings of the National Academy of Sciences* 110 (15): 6223–6228. Xiaolin Wu and Xi Zhang. 2016. “Automated inference on criminality using face images.” arXiv preprint arXiv:1611.04135: 4038–4052. Yaling Yang, Andrea L. Glenn, and Adrian Raine. 2008. “Brain abnormalities in antisocial individuals: implications for the law.” *Behavioral sciences & the law* 26 (1): 65–83. Adrian Raine. 2014. *The anatomy of violence: The biological roots of crime*. Visalia: Vintage Press.

2 AI applications that claim to predict criminality based on physical characteristics are a part of a legacy of long-discredited pseudosciences such as physiognomy and phrenology, which were and are used by academics, law enforcement specialists, and politicians to advocate for oppressive policing and prosecutorial tactics in poor and racialized communities. Indeed, in the opening pages of Hashemi and Hall (2020), the authors invoke the criminological studies of Cesare Lombroso, a dangerous proponent of social Darwinism whose studies the authors cited below overturn and debunk. In the late nineteenth and early twentieth century, police and other government officials relied on social scientists to create universalized measurements of who was “capable” of criminal behavior, based largely on a person’s physical characteristics. This system is rooted in scientific racism and ultimately served to legitimize a regime of preemptive repression, harassment, and forced sterilization in racialized communities. The connections between eighteenth and nineteenth century pseudoscience and facial recognition have been widely addressed. For examples of the historical linkage between physiognomy, phrenology, and automated facial recognition, see Blaise Agüera y Arcas, Margaret Mitchell, and Alexander Todorov. 2017. “Physiognomy’s New Clothes.” *Medium*, May 6.; on links between eugenics, race science, and facial recognition, see Sahil Chinoy. 2019. “The Racist History Behind Facial Recognition.” *New York Times*, July 10.; Stephanie Dick, “The Standard Head,” YouTube.

3 For example, Wu and Zhang (2016) bears a striking resemblance to the Harrisburg study and faced immense public and scientific critique, prompting the work to be rescinded from publication and the authors to issue a response (see Wu and Zhang, 2017.). Experts highlighted the utter lack of a causal relationship between visually observable identifiers on a face and the likelihood of a subject’s participation in criminal behavior. In the absence of a plausible causal mechanism

between the data and the target behavior, and indeed scientific rejection of a causal mechanism, the model is likely not doing what it claims to be doing. In this case, critics rightfully argued that the published model was not identifying criminality — it was identifying historically disadvantaged ethnic subgroups, who are more likely to be targeted by police and arrested. For a summary of the critique see here. The fact that the current study claims its results have “no racial bias” is highly questionable, addressed further below in Sections I (for whether such a thing is possible) and II (whether metrics for bias really capture bias).

4 As Jackie Wang (2018) argues, “‘police science’ is a way for police departments to rebrand themselves in the face of a crisis of legitimacy,” pointing to internally generated data about arrests and incarcerations to justify their racially discriminatory practices. While these types of “evidence based” claims have been problematized and debunked numerous times throughout history, they continue to resurface under the guise of cutting-edge techno-reforms, such as “artificial intelligence.” As Chelsea Barabas (2020, 41) points out, “the term ‘artificial intelligence’ has been deployed as a means of justifying and de-politicizing the expansion of state and private surveillance amidst a growing crisis of legitimacy for the U.S. prison industrial complex.” Sarah Brayne and Angèle Christin argue (2020, 1) that “predictive technologies do not replace, but rather displace discretion to less visible — and therefore less accountable — areas within organizations.” Jackie Wang. 2018. *Carceral capitalism* (Vol. 21). MIT Press. Chelsea Barabas. 2020. “Beyond Bias: Reimagining the Terms of ‘Ethical AI’ in Criminal Law.” 12 *Geo. J. L. Mod. Critical Race Persp.* 2 (forthcoming). Sarah Brayne and Angèle Christin. 2020. “Technologies of Crime Prediction: The Reception of Algorithms in Policing and Criminal Courts.” *Social Problems*.

5 The hard work of these organizers and scholars is beginning to gain public recognition. In recent weeks, major tech companies such as IBM and Amazon, and Microsoft have announced commitments to stop collaborating with law enforcement to deploy facial recognition technologies. These political gains are the result of years of hard work by community organizations such as the Stop LAPD Spying Coalition, Media Mobilizing Project (renamed Movement Alliance Project), Mijente, The Carceral Tech Resistance Network, Media Justice, and AI For the People. This on-the-ground work has been bolstered by research led by Black scholars, such as Joy Buolamwini, Timnit Gebru, Mutale Nkonde and Inioluwa Deborah Raji. See: Joy Buolamwini and Timnit Gebru. 2018. “Gender shades: Intersectional accuracy disparities in commercial gender classification.” Conference on fairness, accountability and transparency. Inioluwa Deborah Raji, Timnit Gebru, Margaret Mitchell, Joy Buolamwini, Joonseok Lee, and Emily Denton. “Saving Face: Investigating the Ethical Concerns of Facial Recognition Auditing.” ArXiv:2001.00964 [Cs], January 3, 2020. Mutale Nkonde. 2020. “Automated Anti-

Blackness: Facial Recognition in Brooklyn, New York.” Harvard Kennedy School Review: 30–36.

6 The Algorithmic Justice League has pointed out this blatant erasure of non-white and non-male voices in their public art project entitled, “Voicing Erasure” a project that was inspired in part by the work of Allison Koenecke, a woman researcher based at Stanford whose work uncovering biases in speech recognition software was recently covered in the New York Times. Koenecke was not cited in the original New York Times article, even though she was the lead author of the research. Instead, a number of her colleagues were named and given credit for the work, all of whom are men. In “Voicing Erasure” Joy Buolamwini pushes us to reflect on “Whose voice do you hear when you think of intelligence, innovation and ideas that shape our worlds?”

7 Timnit Gebru points out that “the dominance of those who are the most powerful race/ethnicity in their location...combined with the concentration of power in a few locations around the world, has resulted in a technology that can benefit humanity but also has been shown to (intentionally or unintentionally) systematically discriminate against those who are already marginalized.” Timnit Gebru. 2020. “Race and Gender.” In Oxford Handbook on AI Ethics. Oxford Handbooks. Oxford University Press. Facial recognition research (arguably a subset of AI) is no different — it has never been neutral nor unbiased. In addition to its deep connection with phrenology and physiognomy, it is entwined with the history of discriminatory police and surveillance programs. For example, Woody Bledsoe, the founder of computational facial recognition, was funded by the CIA to purportedly develop identify criminals and criminal behavior: Leon Harmon. 2020. “How LSD, Nuclear Weapons Led to the Development of Facial Recognition”. Observer. Jan 29. Shaun Raviv. 2020. “The Secret History of Facial Recognition.” Wired. Jan 21. See also: Inioluwa Deborah Raji and Genevieve Fried, “About Face: A Survey of Facial Recognition Datasets.” Accepted to Evaluation Evaluation of AI Systems (Meta-Eval 2020) workshop at AAAI Conference on Artificial Intelligence 2020. Likewise, FERET, the NIST initiative and first large scale face dataset that launched the field of facial recognition in the US was funded by intelligence agencies, for the express purpose of use in identifying criminals in the war on drugs. This objective of criminal identification is core to the history of what motivated the development of the technology. Phillips Jonathon, Harry Wechsler, Jeffery Huang, and Patrick J Rauss. “The FERET Database and Evaluation Procedure for Face-Recognition Algorithms.” Image and Vision Computing 16, no. 5 (1998): 295–306.

8As Safiya Umoja Noble (2018, 30) argues, the problems of data-driven technologies go beyond misrepresentation: “They include decision-making protocols that favor corporate elites and the powerful, and they are implicated in global economic and social inequality.” Safiya Umoja Noble, 2018. Algorithms of

Oppression. New York: New York University Press. D’Ignazio and Klein (2020) similarly argue that data collection environments for social issues such as femicide are often “characterized by extremely asymmetrical power relations, where those with power and privilege are the only ones who can actually collect the data but they have overwhelming incentives to ignore the problem, precisely because addressing it poses a threat to their dominance.” Catherine D’Ignazio and Lauren F. Klein. 2020. *Data feminism*. Cambridge: MIT Press. On the long history of algorithms and political decision-making, see: Theodora Dryer. 2019., *Designing Certainty: The Rise of Algorithmic Computing in an Age of Anxiety*. PhD diss. University of California, San Diego. For an ethnographic study that traces the embedding of power relations into algorithmic systems for healthcare-related decisions, see Beth Semel. 2019. *Speech, Signal, Symptom: Machine Listening and the Remaking of Psychiatric Assessment*. PhD diss., Massachusetts Institute of Technology, Cambridge.

9 As Roberts (2019, 1697) notes in her review of Eubanks (2018), “in the United States today, government digitization targets marginalized groups for tracking and containment in order to exclude them from full democratic participation. The key features of the technological transformation of government decision-making — big data, automation, and prediction — mark a new form of managing populations that reinforces existing social hierarchies. Without attending to the ways the new state technologies implement an unjust social order, proposed reforms that focus on making them more accurate, visible, or widespread will make oppression operate more efficiently and appear more benign.” Dorothy Roberts. 2019. “Digitizing the Carceral State.” *Harvard Law Review* 132: 1695–1728. Virginia Eubanks. 2018. *Automating inequality: How high-tech tools profile, police, and punish the poor*. St. Martin’s Press. Audrey Beard. 2020. “The Case for Care.” *Medium* May 27, 2020. Accessed June 11, 2020. See also: Ruha Benjamin, Troy Duster, Ron Eglash, Nettrice Gaskins, Anthony Ryan Hatch, Andrea Miller, Alondra Nelson, Tamara K. Nopper, Christopher Perreira, Winifred R Poster, et al. 2019. *Captivating Technology: Race, Carceral Techno-science, and Liberatory Imagination in Everyday Life*. Durham: Duke University Press. Chelsea Barabas et al. 2020. “Studying up: reorienting the study of algorithmic fairness around issues of power.” *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*. Anna Lauren Hoffmann. 2019. “Where fairness fails: data, algorithms, and the limits of antidiscrimination discourse”. *Information, Communication & Society* 22 (7): 900–915. Meredith Broussard. 2018. *Artificial Unintelligence: How Computers Misunderstand the World*. Cambridge, MA: MIT Press. Ben Green. 2019. “‘Good’ Isn’t Good Enough.” In *NeurIPS Joint Workshop on AI for Social Good*. AC. Sasha Costanza-Chock. 2020. *Design justice: Community-led practices to build the worlds we need*. Cambridge, MA: MIT Press.

10 As Ruha Benjamin (2016, 148) argues, “One need not harbor any racial animus to exercise racism in this and so many other contexts: rather, when the default settings have been stipulated simply doing one’s job — clocking in, punching out, turning the machine on and off — is enough to ensure the consistency of white domination over time.” Ruha Benjamin. 2016. “Catching our breath: critical race STS and the carceral imagination.” *Engaging Science, Technology, and Society* 2: 145–156.

11 By predictive performance we mean strength of correlations found, as measured by e.g. classification accuracy, metric space similarity, true and false positive rates, and derivative metrics like receiver operator characteristic curves. This is discussed by several researchers, most recently Rachel Thomas and David Uminsky. 2020. “The Problem with Metrics is a Fundamental Problem for AI.” arXiv preprint arXiv:2002.08512.

12 Scholars have long argued that crime statistics are partial and biased, and their incompleteness is delineated clearly along power lines. Arrest statistics are best understood as measurements of law enforcement practices. These practices tend to focus on “street crimes” carried out in low income communities of color while neglecting other illegal activities that are carried out in more affluent and white contexts (Tony Platt. 1978. “‘Street Crime’ — A View From the Left.” *Crime and Social Justice* 9: 26–34; Laura Nader. 2003. “Crime as a category — domestic and globalized.” In *Crime’s Power: Anthropologists and the Ethnography of Crime*, edited by Philip C. Parnell and Stephanie C. Kane, 55–76, London: Palgrave). Consider how loitering is treated compared to more socially harmful practices like wage theft and predatory lending. Similarly, conviction and incarceration data primarily reflect the decision-making habits of relevant actors, such as judges, prosecutors, and probation officers, rather than a defendant’s criminal proclivities or guilt. These decision-making habits are inseparable from histories of race and criminality in the United States. As Ralph (2020, xii) writes, with reference to Muhammad (2019), “since the 1600s, and the dawn of American slavery, Black people have been viewed as potential criminal threats to U.S. society. As enslaved people were considered legal property, to run away was, by definition, a criminal act...Unlike other racial, religious, or ethnic groups, whose crime rates were commonly attributed to social conditions and structures, Black people were (and are) considered inherently prone to criminality...Muhammad [thus] argues that equating Blackness and criminality is part of America’s cultural DNA.” Khalil Gibran Muhammad. 2011. *The Condemnation of Blackness: Race, Crime, and the Making of Modern Urban America*. Cambridge, MA: Harvard University Press; Laurence Ralph. 2020. *The Torture Letters: Reckoning with Police Violence*. Chicago: University of Chicago Press. See also: Victor M. Rios. 2011. *Punished: Policing the Lives of Black and Latino Boys*. New York: NYU Press.

13 Yet, criminal justice data is rarely used to model the behaviors of these powerful system actors. As Harris (2003) points out, it is far more common for law enforcement agencies to use their records to justify racially discriminatory policies, such as stop and frisk. David A. Harris. 2003. “The reality of racial disparity in criminal justice: The significance of data collection.” *Law and Contemporary Problems* 66 (3): 71–98. However, some data science projects have sought to reframe criminal legal data to center such powerful system actors. For example, the Judicial Risk Assessment project repurposes criminal court data to identify judges who are likely to use bail as a means of unlawfully detaining someone pretrial. Chelsea Barabas, Colin Doyle, JB Rubinovitz, and Karthik Dinakar. 2020. “Studying up: reorienting the study of algorithmic fairness around issues of power.” In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency (FAccT ’20)*. Association for Computing Machinery, New York, NY, USA, 167–176. Similarly, the White Collar Crime project is a satirical data science project that reveals the glaring absence of financial crimes from predictive policing models, which tend to focus on “street crimes” that occur in low income communities of color. Brian Clifton, Sam Lavigne, and Francis Tseng. 2017. “White Collar Crime Risk Zones.” *The New Inquiry* 59: ABOLISH (March).

14 Decades of research have shown that, for the same conduct, Black and Latinx people are more likely to be arrested, prosecuted, convicted and sentenced to harsher punishments than their white counterparts, even for crimes that these racial groups engage in at comparable rates. Megan Stevenson and Sandra G. Mayson. 2018. “The Scale of Misdemeanor Justice.” *Boston University Law Review* 98 (731): 769–770. For example, Black people are 83% more likely to be arrested for marijuana compared to whites at age 22 and 235% more likely to be arrested at age 27, in spite of similar marijuana usage rates across racial groups. (Ojmarrh Mitchell and Michael S. Caudy. 2013. “Examining Racial Disparities in Drug Arrests.” *Justice Quarterly* 2: 288–313.) Similarly, Black drivers are three times as likely as white drivers to be searched during routine traffic stops, even though police officers generally have a lower “hit rate” for contraband when they search drivers of color. “Ending Racial Profiling in America: Hearing Before the Subcomm.” 2012. on the Constitution, Civil Rights and Human Rights of the Comm. on the Judiciary, 112th Cong. 8 (statement of David A. Harris). In the educational sector, Nance (2017) found that schools with a student body made up of primarily of people of color were two to eighteen times more likely to use security measures (metal detectors, school and police security guards, locked gates, “random sweeps”) than schools with a majority (greater than 80%) white population. Jason P. Nance. 2017. “Student Surveillance, Racial Inequalities, and Implicit Racial Bias.” *Emory Law Journal* 66 (4): 765–837). Systematic, racial disparities in the U.S. criminal justice system run historically deep as well. In as early as 1922, white Chicagoans who testified on a report that city officials commissioned following uprisings after the murder of 17-year-old Eugene Williams asserted that “the police are systematically engaging in

racial bias when they're targeting Black suspects" (Khalil Gibran Muhammad, quoted in Anna North. 2020. "How racist policing took over American cities, explain by a historian." Vox, June 6. (Accessed June 18, 2020). These same inequities spurred William Patterson, then-president of the Civil Rights Congress, to testify to the United Nations in 1951 that "the killing of Negroes has become police policy in the United States." In addition, Benjamin (2018) notes that institutions in the U.S. tend toward the "wiping clean" of white criminal records, as in the case of a Tulsa, Oklahoma officer who had any evidence of her prosecution for the murder of Terrance Crutcher, a 43-year-old unarmed Black man, removed from her record altogether (Ruha Benjamin, 2018. "Black Afterlives Matter." Boston Review July 28. Accessed on Jun 1, 2020.) . All of these factors combined lead to an overrepresentation of people of color in arrest data.

15 On the topic of doing "ethical" computing work, Abeba Birhane (2019) avers "the fact that computer science is intersecting with various social, cultural and political spheres means leaving the realm of the 'purely technical' and dealing with human culture, values, meaning, and questions of morality; questions that need more than technical 'solutions', if they can be solved at all." Abeba Birhane. "Integrating (Not 'Adding') Ethics and Critical Thinking into Data Science." Abeba Birhane (blog), April 29, 2019. It is worth mentioning the large body of computer vision, machine learning, and data science research that acknowledges the gross ethical malfeasance of the work typified in the offending research, reveals the impotence of data "debiasing" efforts, and argues for deeper integration of critical and feminist theories in computer science. See, for instance: Michael Skirpan, and Tom Yeh. 2017. "Designing a Moral Compass for the Future of Computer Vision Using Speculative Analysis." In IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW), 1368–77. Honolulu, HI, USA: IEEE, 2017. Hila Gonen, and Yoav Goldberg. 2019. "Lipstick on a Pig: Debiasing Methods Cover up Systematic Gender Biases in Word Embeddings But Do Not Remove Them." ArXiv:1903.03862 [Cs], September 24. Green, Ben. 2019. "'Good' Isn't Good Enough." In NeurIPS Joint Workshop on AI for Social Good. ACM,. Rashida Richardson et. al.. 2019. "Dirty Data, Bad Predictions: How Civil Rights Violations Impact Police Data, Predictive Policing Systems, and Justice" NYU Law Review Online 192, February 13. Available at SSRN. Audrey Beard and James W. Malazita. "Greased Objects: How Concept Maintenance Undermines Feminist Pedagogy and Those Who Teach It in Computer Science." To be presented at the EASST/4S Panel on Teaching interdependent agency: Feminist STS approaches to STEM pedagogy, August 2020.

16 In these applications, both groupings of pixels and human-given labels are directly observable, making such domains suitable for machine learning-based approaches. Criminality detection or prediction, on the other hand, are not because

criminality has no stable empirical existence. See also: Momin M. Malik 2020. “A Hierarchy of Limitations in Machine Learning.” arXiv preprint arXiv:2002.05193.

17 Yarden Katz. 2017. “Manufacturing an Artificial Intelligence Revolution.” SSRN.

18 Arvind Narayanan. 2019. “How to Recognize AI Snake Oil.” Arthur Miller Lecture on Technology and Ethics, Massachusetts Institute of Technology, November 18, Cambridge, MA.

19 By insisting that signs of criminality can be located in biological material (in this case, features of the face), this research perpetuates the process of “racialization”, defined by Marta Maria Maldonado (2009: 1034) as “the production, reproduction of and contest over racial meanings and the social structures in which such meanings become embedded. Racial meanings involve essentializing on the basis of biology or culture.” Race is a highly contingent, unstable construct, the meaning of which shifts and changes over time with no coherent biological correlate. To imply that criminality is eminent in biology and that certain kinds of bodies are marked as inherently more criminal than others lays the groundwork for arguing that certain categories of people are more likely to commit crimes because of their embodied physicality, a clearly false conclusion. This has motivated leading scholars to move beyond analysis of race and technology to race as technology. In Wendy Hui Kyong Chun’s (2013, 7) words: “Could race be not simply an object of representation and portrayal, of knowledge or truth, but also a technique that one uses, even as one is used by it — a carefully crafted, historically inflected system of tools, mediation, or enframing that

builds history and identity?” See also; Simone Browne. 2010. “Digital Epidermalization: Race, Identity and Biometrics.” *Critical Sociology* 36 (1): 131–150; Simone Browne. 2015. *Dark Matters: On the Surveillance of Blackness*. Durham: Duke University Press; Alondra Nelson. 2016. *The Social Life of DNA: Race, Reparations, and Reconciliation After the Genome*. Boston, MA: Beacon Press; Amande M’Charek. 2020. “Tentacular Faces: Race and the Return of the Phenotype in Forensic Identification.” *American Anthropologist* doi:10.1111/aman.13385; Keith Wailoo, Alondra Nelson, and Catherine Lee, eds. 2012. *New Brunswick*: Rutgers University Press; Marta Maria Maldonado. 2009. “It is their nature to do menial labour’: the racialization of ‘Latino/a workers’ by agricultural employers.” *Ethnic and Racial Studies*, 32(6): 1017–1036; Wendy Hui Kyong Chun. 2013. “Race and/as Technology, or How to do Things to Race.” In *Race after the Internet*, 44–66. Routledge; Beth Coleman. 2009. “Race as technology.” *Camera Obscura: Feminism, Culture, and Media Studies*, 24 (70): 177–207. Fields across the natural sciences have long employed the construct of race to define and differentiate among groups and individuals. In 2018, a group of 67 scientists, geneticists, and researchers jointly dissented to the continuation of scientific discourse of race as a way to define

differences between humans, and called attention to the inherently political work of classification. As they wrote, “there is a difference between finding genetic differences between individuals and constructing genetic differences across groups by making conscious choices about which types of group matter for your purposes. These sorts of groups do not exist ‘in nature.’ They are made by human choice. This is not to say that such groups have no biological attributes in common. Rather, it is to say that the meaning and significance of the groups is produced through social interventions.” “How Not To Talk About Race And Genetics.” 2018. BuzzFeedNews March 30. (Accessed June 18, 2020).

20 For further reading on why “strength of correlations, evaluated retrospectively,” is a more accurate term for “prediction,” see Momin M. Malik. 2020. “A Hierarchy of Limitations in Machine Learning.” arXiv preprint arXiv:2002.05193; Daniel Gayo-Avello. 2012. “No, You Cannot Predict Elections with Twitter.” IEEE Internet Computing November/December 2012. Arvind Narayanan (2019)

21 These reasons for real-world performance being less than test set performance include overfitting to the test set, publication bias, and distribution shift.

22 Hinton (2016) follows the construction of Black criminality through the policies and biased statistical data that informed the Reagan administration’s War on Drugs and the Clinton administration’s War on Crime. She tracks how Black criminality, “when considered an objective truth and a statistically irrefutable fact...justified both structural and everyday racism. Taken to its extreme, these ideas sanctioned the lynching of black people in the southern states and the bombing of African American homes and institutions in the urban north before World War II...In the postwar period, social scientists increasingly rejected biological racism but created a new statistical discourse about black criminality that went on to have a far more direct impact on subsequent national policies and, eventually, serve as the intellectual foundation of mass incarceration” (19). Elizabeth Hinton, 2016. *From the War on Poverty to the War on Crime*. Cambridge, MA: Harvard University Press. See also: Charlton D. McIlwain. 2020. *Black Software: The Internet & Racial Justice, from the AfroNet to Black Lives Matter*. Oxford, UK: Oxford University Press. Data-gathering enterprises and research studies that uncritically incorporate criminal justice data into their analysis fuel stereotypes of African-Americans as “dangerous” or “risks to public safety,” the history (and violent consequences) of which is reviewed in footnotes 12 and 14. The continued propagation of these stereotypes via academic discourse continues to foment anti-Black violence at the hands of the police. It is within this historically embedded, sociocultural construction of Black criminality and Blackness as inherently threatening that police often find their justification for lethal uses of

force. Today, Black Americans are twice as likely as white Americans to be murdered at the hands of police. (Julie Tate, Jennifer Jenkins, and Steven Rich. 2020. “Fatal Force: Police Shootings Database.” Washington Post, May 13). As of June 9, 2020, the Mapping Police Violence project found that 24% of the 1,098 people killed by the police in 2019 were Black, despite the fact that Black people make up only 13% of the population in the U.S.

23 Wendy Hui Kyong Chun (2020) points to the performativity of predictive ML more broadly: “predictions are correct because they program the future [based on the past]. She offers a way to reimagine their use to work against an unwanted future: “In contrast, consider global climate-change models — they too make predictions. They offer the most probable outcome based on past interactions. The point, however, isn’t to accept their predictions as truth but rather to work to make sure their predictions don’t come true. The idea is to show us the most likely future so we will create a different future.” Wendy Hui Kyong Chun and Jorge Cottamay. 2020. “Reimagining Networks An interview with Wendy Hui Kyong Chun.” *The New Inquiry*.

24 Barocas et al. 2019. “A 2016 paper analyzed a predictive policing algorithm by PredPol, one of the few to be published in a peer-reviewed journal. By applying it to data derived from Oakland police records, they found that Black people would be targeted for predictive policing of drug crimes at roughly twice the rate of whites, even though the two groups have roughly equal rates of drug use (Lum and Isaac 2016). Their simulation showed that this initial bias would be amplified by a feedback loop, with policing increasingly concentrated on targeted areas. This is despite the fact that the PredPol algorithm does not explicitly take demographics into account.” (Solon Barocas, Moritz Hardt and Arvind Narayanan. 2019. *Fairness and Machine Learning*; Kristian Lum, and William Isaac. 2016. “To predict and serve?” *Royal Statistical Society* 13 (5): 14–19).

25 As Reginald Dwayne Betts (2015, 224) argues, “How does a system that critics, prisoners, and correctional officials all recognize as akin to torture remain intact today? The answer is simple: we justify prison policy based on our characterizations of those confined, not on any normative belief about what confinement in prison should look like.” Reginald Dwayne Betts. 2015. “Only Once I Thought About Suicide.” *Yale L&JF* 125: 222. For more on the construction of deviant profiles as a means of justifying social exclusion, see: Sharon Dolovich. 2011. “Exclusion and control in the carceral state.” *Berkeley Journal of Criminal Law* 16: 259. David A. Harris. 2003. “The reality of racial disparity in criminal justice: The significance of data collection.” *Law and Contemporary Problems* 66 (3): 71–98. Michael J. Lynch. 2000. “The power of oppression: Understanding the history of criminology as a science of oppression.” *Critical Criminology* 9: 144–152. Mitali Thakor. 2017. “How to Look: Apprehension, Forensic Craft, and the Classification of

Child Exploitation Images.” *IEEE Annals of the History of Computing* 39 (2): 6–8.

Mitali Thakor. 2018. “Digital Apprehensions: Policing, Child Pornography, and the Algorithmic Management of Innocence. *Catalyst* 4 (1): 1–16.

26 This crisis is nothing new: Weizenbaum noted some of the epistemic biases of AI in 1985 (ben Aaron 1985), and Agre discussed the limits of AI methods in 1997 (Agre 1997). More recently, Elish and boyd directly interrogated practices and heritage of AI. Diana ben-Aaron. 1985. “Weizenbaum Examines Computers and Society.” *The Tech*. April 9. Philip E. Agre. 1997. “Toward a Critical Technical Practice: Lessons Learned in Trying to Reform AI.” In *Bridging the Great Divide: Social Science, Technical Systems, and Cooperative Work*, edited by Geof Bowker, Les Gasser, Leigh Star, and Bill Turner. Hillsdale, NJ: Erlbaum. M.C Elish and danah boyd. 2018. “Situating Methods in the Magic of Big Data and AI.” *Communication Monographs* 85 (1): 57–80.

27 This is perhaps unsurprising, given the conditions of such interventions, as Audre Lorde (1984) points out: “What does it mean when the tools of a racist patriarchy are used to examine the fruits of that same patriarchy? It means that only the most narrow parameters of change are possible and allowable.” Audre Lorde. “The Master’s Tools Will Never Dismantle the Master’s House.” 1984. The specific challenges of “ethical” AI practice (due to a lack of operational infrastructure, poorly-defined and incomplete ethics codes, and no legal or business incentives, among others) have been well documented in the past several years. Michael A. Madaio, Luke Stark, Jennifer Wortman Vaughan, and Hanna Wallach. “Co-Designing Checklists to Understand Organizational Challenges and Opportunities around Fairness in AI.” In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 14, 2020. Stark, Luke, and Anna Lauren Hoffmann. 2019. “Data Is the New What? Popular Metaphors & Professional Ethics in Emerging Data Culture.” *Journal of Cultural Analytics*. 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,” In *Proceedings of the 52nd Hawaii International Conference on System Sciences*. Maui, HI. Hagendorff, Thilo. 2019. “The Ethics of AI Ethics — An Evaluation of Guidelines.” *ArXiv Preprint ArXiv:1903.03425*, 15. Jess Whittlestone, Rune Nyrop, Anna Alexandrova, and Stephen Cave. 2019. “The Role and Limits of Principles in AI Ethics: Towards a Focus on Tensions.” In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, 7,. Anna Jobin, Marcello Ienca, and Effy Vayena. “2019. The Global Landscape of AI Ethics Guidelines.” *Nature Machine Intelligence*, September 2.

28 Worthy of note are other discourses of “ethics” in AI, like transparency, accountability, ethics (with fairness, comprising the FATE framework), and trust. For discussion around fairness and bias, see Chelsea Barabas 2019. “Beyond Bias:

Reimagining the Terms of ‘Ethical AI’ in Criminal Law”. S.M. West, M. Whittaker, and K. Crawford 2019. “Discriminating Systems: Gender, Race and Power in AI. AI Now Institute”. However, many scholars have identified limitations of research and design within the Fairness, Accountability, Transparency and Ethics (FATE) streams of machine learning to over-simplify the “interlocking matrix” of data discrimination and algorithmic bias which are always differentially (and disproportionately) experienced (Costanza-Chock, 2018). Others have argued that the focus on fairness through antidiscrimination discourse from law, policy and cognate fields over-emphasizes a liberal framework of rights, opportunities and material resources (Hoffman, 2019: 908). Approaches which bring to bear the lived experience of those who stand to be most impacted into the design, development, audit, and oversight of such systems are urgently needed across tech ethics streams. As Joy Buolamwini notes, “Our individual encounters with bias embedded into coded systems — a phenomenon I call the ‘coded gaze’ — are only shadows of persistent problems with inclusion in tech and in machine learning.” Joy Buolamwini. 2016. “Unmasking Bias”. Medium. Dec 14. In order for “tech ethics” to move beyond simply mapping discrimination, it must contend with the power and politics of technological systems and institutions more broadly. Sonja, Solomun. 2021. “Toward an Infrastructural Approach to Algorithmic Power” in Elizabeth Judge, Sonja Solomun and Drew Bush, eds. [Forthcoming]. *Power and Justice: Cities, Citizens and Locational Technology*. UBC Press.

29 Greene, Hoffman, and Stark 2019. Chelsea Barabas, Colin Doyle, JB Rubinovitz, and Karthik Dinakar. 2020. “Studying up: reorienting the study of algorithmic fairness around issues of power”. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency (FAccT ’20)*. Association for Computing Machinery, New York, NY, USA, 167–176. Sasha Costanza-Chock. 2018. “Design Justice: towards an intersectional feminist framework for design theory and practice”. *Proceedings of the Design Research Society (2018)*.; Madeleine Clare Elish and danah boyd. 2018. “Situating methods in the magic of Big Data and AI”. *Communication monographs* 85, 1 (2018), 57–80.; Andrew D Selbst and Solon Barocas. 2018. “The intuitive appeal of explainable machines”. *Fordham Law Review* 87 (2018), 1085.

30 We borrow verbiage from set theory here to illustrate the deep complexity of such contexts, and to illustrate the peril of attempting to discretize this space.

31 In outlining parallels between archival work and data collection efforts for ML, Eun Seo Jo and Timnit Gebru (2020) bring forth a compelling interdisciplinary lens to the ML community, urging “that an interdisciplinary subfield should be formed focused on data gathering, sharing, annotation, ethics monitoring, and record-keeping processes.” Eun Seo Jo and Timnit Gebru. 2020. “Lessons from Archives: Strategies for Collecting Sociocultural Data in Machine Learning.”

Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency. For other great examples of this kind of interdisciplinary scholarship, see: Chelsea Barabas, Colin Doyle, JB Rubinovitz, and Karthik Dinakar. 2020. “Studying up: reorienting the study of algorithmic fairness around issues of power”. In Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency (FAccT ’20). Association for Computing Machinery, New York, NY, USA, 167–176.

32 Several key organizations are leading the charge in forwarding reflexive, critical, justice-focused, and anti-racist computing. Examples include Data 4 Black Lives, which is committed to “using the datafication of our society to make bold demands for racial justice” and “building the leadership of scientists and activists and empowering them with the skills, tools and empathy to create a new blueprint for the future” (Yeshimabeit Milner. 2020. “For Black people, Minneapolis is a metaphor for our world.” Medium May 29. Accessed June 4, 2020). Another example is Our Data Bodies, which is “based in marginalized neighborhoods in Charlotte, North Carolina, Detroit, Michigan, and Los Angeles, California,” and tracks “the ways [these] communities’ digital information is collected, stored, and shared by government and corporations...[working] with local communities, community organizations, and social support networks, [to] show how different data systems impact re-entry, fair housing, public assistance, and community development.” A third example is the Algorithmic Justice League, which combines “art, research, policy guidance and media advocacy” to build “a cultural movement towards equitable and accountable AI,” which includes examining “how AI systems are developed and to actively prevent the harmful use of AI systems” and “[preventing] prevent AI from being used by those with power to increase their absolute level of control, particularly where it would automate long-standing patterns of injustice.”

33 Abstraction as epistemology in computer science was independently developed by Malazita and Resetar (2019) and Selbst et al. (2019). James W. Malazita and Korryn Resetar. 2019. “Infrastructures of Abstraction: How Computer Science Education Produces Anti-Political Subjects.” *Digital Creativity* 30 (4): 300–312. Andrew D. Selbst, danah boyd, Sorelle A. Friedler, Suresh Venkatasubramanian, and Janet Vertesi. 2019. “Fairness and Abstraction in Sociotechnical Systems.” In Proceedings of the Conference on Fairness, Accountability, and Transparency — FAT* ’19, 59–68. Atlanta, GA, USA: ACM Press.

34 Sareeta Amrute (2019, 58) argues that the standard, procedural approach of conventional tech ethics “provide[s] little guidance on how to know what problems the technology embodies or how to imagine technologies that organise life otherwise, in part because it fails to address who should be asked when it comes to defining ethical dilemmas” and “sidesteps discussions about how such things as ‘worthy and practical knowledge’ are evaluated and who gets to make these

valuations.” In so doing, it risks reinforcing “narrow definitions of who gets to make decisions about technologies and what counts as a technological problem.” Alternatively, postcolonial and decolonising feminist theory offer a framework of ethics based on relationality rather than evaluative check-lists, in a way that can “move the discussion of ethics from establishing decontextualized rules to developing practices to train sociotechnical systems — algorithms and their human makers — to being with the material and embodied situations in which these systems are entangled, which include from the start histories of race, gender, and dehumanisation” (ibid. Sareeta Amrute. 2019. “Of Techno-Ethics and Techno-Affects.” *Feminist Review* 123 (1): 56–73). In other words, the conventional frame of “tech ethics” does not always acknowledge that the work of computer science is inherently political. As Ben Green (2019) states, “Whether or not the computer scientists behind [racist computational criminal prediction projects] recognize it, their decisions about what problems to work on, what data to use, and what solutions to propose involve normative stances that affect the distribution of power, status, and rights across society. They are, in other words, engaging in political activity. And although these efforts are intended to promote “social good,” that does not guarantee that everyone will consider such projects beneficial.” See also: Luke Stark. 2019. “Facial recognition is the plutonium of AI.” *XRDS: Crossroads, The ACM Magazine for Students* 25 (3): 50–55. For efforts that exemplify the relational approach to ethics that Amrute endorses and includes the people most marginalized by technological interventions into the design process, see Sasha Costanza-Chock. 2020. *Design Justice: Community-Led Practices to Build the Worlds We Need*. Cambridge, MA: MIT Press. For an example of an alternative ethics based around relationality, see Jason Edward Lewis, Noelani Arista, Archer Pechawis, and Suzanne Kite. 2018. “Making kin with the machines.” *Journal of Design and Science*.

35 Power is here defined as the broader social, economic, and geopolitical conditions of any given technology’s development and use. As Ruha Benjamin (2016; 2019), Safiya Umoja Noble (2018), Wendy Hui Kyong Chun (2013; 2019), Taina Bucher (2018), and others have argued, algorithmic power is productive; it maintains and participates in making certain forms of knowledge and identity more visible than others.

36 Crime prediction technology is not simply a tool—it can never be divorced from the political context of its use. In the U.S., this context includes the striking racial dimension of the country’s mass incarceration and criminalization of racial or ethnic minorities. Writing in 2020, acclaimed civil rights lawyer and legal scholar, Michelle Alexander (2020, 29) observes that “the United States imprisons a larger percentage of its black population than South Africa did at the height of apartheid”. Michelle Alexander. 2020. *The new Jim Crow: Mass incarceration in the age of colorblindness*. The New Press. For an ethnographic analysis of the stakes at play

when computer scientists and engineers partner with and expand the reach of policing networks, see Mitali Thakor. 2016. *Algorithmic Detectives Against Child Trafficking: Data, Entrapment, and the New Global Policing Network*. PhD diss., Massachusetts Institute of Technology, Cambridge.

37 The Carceral Tech Resistance Network (2020) provides a useful set of guiding questions to evaluate new projects, procurements and programs related to law enforcement reform. These questions are centered in an abolitionist understanding of the carceral state, which challenges the notion that researchers and private actors for profit can “fix” American policing through technocratic solutions that are largely motivated by profit and not community safety and reparations for historical harms. For a comprehensive record of the risks law enforcement face recognition poses to privacy, civil liberties, and civil rights, see Clare Garvie, A. M. Bedoya, and J. Frankle. 2016. “The perpetual line-up. Unregulated police face recognition in America”. Georgetown Law Center on Privacy & Technology.

38 2020. “HU facial recognition software predicts criminality.” Harrisburg University of Science and Technology, May 5. See also Rose Janus. 2020. “University Deletes Press Release Promising ‘Bias-Free’ Criminal Detecting Algorithm.” Vice, May 6