

February 23, 2021

NYPD Commissioner Shea
New York Police Department
One Police Plaza
New York, NY 10038
Via Email

Re: NYPD's Facial Recognition Impact & Use Policy

Dear Commissioner Shea:

The Surveillance Technology Oversight Project (“S.T.O.P.”)¹ hereby submits our comment in response to the Draft Facial Recognition Impact and Use Policy (“Policy”) published by the New York City Police Department (“NYPD”) on January 11, 2021 pursuant to the Public Oversight of Surveillance Technology Act (“POST Act”). Not only did S.T.O.P. work extensively to promote passage of the POST Act, the law’s enactment was one of the reasons we were founded. Sadly, upon review, the Policy is so grossly inadequate that it not only undermines public trust and accountability, it violates the NYPD’s reporting obligations under the POST Act.

Instead of publishing an impact statement that tells New Yorkers what surveillance tools the NYPD uses, we were provided copy-and-paste responses that are opaque, misleading, and, at times, blatantly wrong. As written, the Policy primarily tells New Yorkers one thing: the NYPD cannot be trusted to use facial recognition technology.

Data Sharing Agreements

The POST Act requires the NYPD to enumerate all entities which are able to access the Department’s facial recognition data. However, instead of providing any meaningful information, the Policy merely states that unspecified “agencies at the local, state, and federal level . . . have limited access to NYPD computer and case management systems.” At a minimum, the Department must provide a full accounting of all agencies that access such data, along with the frequency of access and any limitations on how such data is used and retained. The NYPD would also need to provide a copy of any/all agreements with external agencies pertaining to the scope of agency access and the volume of data retained.

Vendors and Product Disclosure

Perhaps no aspect of the Policy is more antithetical to the text and spirit of the POST Act than the Department’s systematic failure to specify the make and model of equipment used for facial recognition. The driving impetus for the POST Act was the Department’s historical failure to disclose what tools it purchased to monitor New Yorkers until years or decades after the fact. This

¹ S.T.O.P.” is a non-profit organization that advocates and litigates for New Yorkers’ privacy rights, fighting discriminatory surveillance. For more information see <https://www.stopspying.org/>.

type of surreptitious procurement is antithetical to democratic government and the role of the City Council in overseeing agency purchases. Rather than comply with the POST Act's reporting obligations, the Policy describes the Department's facial recognition program in vague, non-descript terms. The Policy fails to include a single vendor name, let alone the comprehensive listing of tools that lawmakers required to be provided. At a minimum, the revised policy must include the name of every single facial recognition system employed by the NYPD, the system's manufacturer, and the names of any other vendors involved in creating or operating the system. The NYPD should also provide a comprehensive evaluation of what data is accessed and/or retained by vendors.

Racial, Ethnic, and Religious Bias

Racial discrimination and bias have defined New York City's policing since before the NYPD was even founded, and that deadly legacy of injustice has continued to this day. The POST Act provided the Department with a unique opportunity to address the ways that its surveillance operations have been driven by, and in turn fueled, discrimination for decades. Sadly, rather than addressing this challenge head on, the Department simply ignored the POST Act's requirements, responding with a terse and unbelievable claim that "The NYPD prohibits the use of racial and bias-based profiling in law enforcement actions." This statement is patently absurd. The NYPD has long been emblematic to the country as a symbol of biased policing²; and after the Department's violent and discriminatory response to recent protests, it's clear just how little has changed.³ Facial recognition exacerbates officers' bias, discriminates against BIPOC and LGBTQ+ communities, puts over-surveilled New Yorkers at risk of wrongful arrests, and worse. Universities, civil society organizations, and even the U.S. federal government proved that facial recognition technology is biased and broken. The technology has no place within New York City's law enforcement.

Retention Periods and Access Rights

To meet the minimum transparency requirements set out in the POST Act, NYPD must also clarify how long data is saved and how the access rights to the information is determined. The Policy does not provide [any/sufficient] information about the retention periods of the data collected through facial recognition technology. Instead, the Policy contains broad boilerplate language, referring to "applicable laws, regulations, and New York City and NYPD policies" without disclosing what these are or what they entail. The Department also fails to clearly and coherently describe access rights for NYPD employees and contractors to access this exceptionally sensitive data. Bland phrases stating that access rights are given to personnel with an "articulable need" and that access is "further limited based on lawful duty" are feeble efforts to circumvent the reporting obligations set out in the POST Act.

NYPD Data Security

The NYPD is not just asking New Yorkers to allow the Department access to huge volumes of intimate data about our private lives; they want us to let that data to be accessible to anyone who can break into the NYPD's systems. Sadly, we have no way to judge the risk that this data could fall into the hands of any hacker, criminal, or rogue state that could breach NYPD security measures. That is because the NYPD's data security promises are full of repetitive and empty phrases. The section

² Lauren del Valle, *NYPD didn't substantiate any complaints of police bias over 4 years. Report cites need to improve*, CNN (Updated 12:10 PM EDT, Jun. 27, 2019), <https://www.cnn.com/2019/06/27/us/nypd-bias-complaints-report/index.html>.

³ See Julie Ciccolini and Ida Sawyer, "Kettling" Protesters in the Bronx: Systemic Police Brutality and Its Costs in the United States, Human Rights Watch (Sep. 30, 2020), <https://www.hrw.org/report/2020/09/30/kettling-protesters-bronx/systemic-police-brutality-and-its-costs-united-states>.

contains general descriptions about the safeguards in place for the Department’s case management and computer systems, stating that NYPD uses a “multifaceted approach to secure data and user accessibility.”

Not only is the provided information insufficient to build public trust and accountability, it is also so generic that it is almost completely useless from a technical standpoint. The NYPD references its use of Lightweight Directory Access Protocol, dual factor authentication, Secure Socket Layer, and Transport Layer Security. These rudimentary encryption and security features are so ubiquitous that it would only be notable if they were not used as part of the NYPD’s data security policy. This is about as persuasive as arguing that a car is safe simply because it has functioning seatbelts; the real surprise would be finding a car that did not. The enormous amounts of highly sensitive data processed through the NYPD’s facial recognition systems call for higher security standards than what is described in the Policy.

NYPD Training

The Policy recognizes the self-evident truth that training is an important factor for the NYPD’s use of facial recognition technology. For example, the Policy states that every NYPD employee who gains access to facial recognition technology must first complete a “mandatory training related to use of the technology.” Sadly, this is not the introductory clause to an expansive training policy: this is almost the whole of the Policy’s details on the topic. The Policy’s training section is grossly insufficient, to say the least. This is particularly true in light of the NYPD’s well-documented history of using unscientific procedures that impair the accuracy of facial recognition and put New Yorkers at risk of false arrest. The NYPD has previously had officers run facial recognition scans against celebrity lookalikes of suspects, a practice that has absolutely no scientific validity or evidentiary value. The Department has trained officers to manipulate images prior to facial recognition review, a procedure that could further erode the accuracy of facial recognition. The Policy leaves unclear if officers are still trained to use these pseudoscientific techniques or other approaches that would increase the error rate of facial recognition.

Comparison of the POST Act to other CCOPS Jurisdictions

The Department’s failure to provide the public with meaningful details is particularly egregious in light of the strong national record of compliance with analogous efforts. As of today, more than a dozen localities have adopted Community Control Over Police Surveillance (CCOPS) laws. The POST Act is an outlier, both in that it is one of the weakest laws in the country and because the NYPD’s response has shown an unprecedented effort to circumvent even the most minimal transparency requirements.⁴ While many municipalities’ legislations require acquisition approval, ban non-disclosure agreements, and provide a right of action for private citizens, the POST Act only requires the NYPD to provide annual reports and use policies. Nevertheless, the NYPD is incapable of meeting the requirements set out in the POST Act by only providing opaque or boiler-plate responses in the Policy, hiding the details needed for meaningful public engagement. As a result, it is clear that more aggressive legislative responses are required, including a complete ban on the NYPD’s use of facial recognition.

⁴ Hogan Lovells and Surveillance Technology Oversight Project, *New CCOPS On The Beat* (Feb. 10, 2021), <https://static1.squarespace.com/static/5c1bfc7eee175995a4ceb638/t/602430a5ef89df2cc6894ce1/1612984485653/News+CCOPS+On+The+Beat.pdf>.

Artificial Intelligence and Machine Learning.

The Policy falsely states that the NYPD's facial recognition technology does not use artificial intelligence or machine learning. This claim is so preposterous that it demonstrates either intentional misrepresentation or gross incompetence. Facial recognition is not only one application of artificial intelligence and machine learning technology, it is perhaps the single most widely discussed example of such technologies. Facial recognition has been classified as a form of artificial intelligence and machine learning by everyone from researchers and lawmakers to business leaders—and even to the NYPD itself.⁵ Notably, within days of publishing the Policy, the New York City Algorithmic Management and Policy Office released its report on the use of automated decision systems in New York City. Among the artificial intelligence systems that were listed was the exact same facial recognition technology that was misclassified in the Policy. This blatant error speaks to the NYPD's lack of candor and rigor in completing the Policy.

Monitoring Political Rallies

The Policy falsely states that the NYPD does not use facial recognition technology to identify or monitor people in crowds or at political rallies. Not only is this statement inconsistent with existing NYPD policy, as documented in the Department's Patrol Guide, it ignores documented instances of protest surveillance. While we suspect this technology is used frequently to track protests, it was quite explicit in the case of Dwreck, a political organizer and leader of Warriors in the Garden. Shortly after leading a demonstration against police violence in 2020, Dwreck's apartment was surrounded by officers who engaged in an hours-long standoff. Bystander photographs revealed that one officer at the scene was holding a facial recognition report from the Department's Facial Identification Section. None of these facts are addressed by the Policy. To the contrary, the Policy's false, boilerplate reassurance undermines the very transparency and accountability the POST Act was enacted to support. More consequentially, this type of misstatement may be a violation of the NYPD's reporting obligations, exposing the Department to civil legal penalties. Misstatements include:

- Claim that "hybrid machine/human systems" fix the "variations in accuracy" observed with non-white males; and
- There is insufficient information about what training NYPD personnel are receiving

Concluding Remarks

The cumulative impact of the forgoing errors and omissions is clear: the NYPD is breaking the law. The POST Act is not a formality, it is not a nicety, it is binding legislation with full force of law. When the NYPD fails to comply with the statute, it seeks to overturn the will of New York's elected leaders, accomplishing by force what it failed to do through lobbying. If the NYPD persists in this

⁵ See e.g. Steve Lohr, *Facial Recognition Is Accurate, if You're a White Guy*, New York Times (Feb. 9, 2018), <https://www.nytimes.com/2018/02/09/technology/facial-recognition-race-artificial-intelligence.html>; Tom Simonite, *Congress Is Eyeing Face Recognition, and Companies Want a Say*, Wired (Nov. 23, 2020, 07:00 AM), <https://www.wired.com/story/congress-eyeing-face-recognition-companies-want-say/>; Ethan Geringer-Sameth, *The NYPD's Facial Recognition Policy Leaves A Lot of Leeway the Department Says It's Not Using*, Gotham Gazette (Jul. 22, 2020), <https://www.gothamgazette.com/city/9608-nypd-facial-recognition-policy-leeway-department-not-using-black-lives-matter-protests>.

flagrant disregard for its statutory reporting requirements, it will simply hasten the enactment of far more sweeping changes to the Department's surveillance powers in the coming months.

Sincerely,

/s/
Albert Fox Cahn, Esq.
Executive Director