February ##, 2021

Commissioner Dermot Shea
New York City Police Department
One Police Plaza
New York, NY 10038

Re: Comments on Draft Surveillance Impact and Use Policies

Dear Commissioner Shea:

The New York Civil Liberties Union ("NYCLU") writes in response to the 36 draft surveillance technology impact and use policies issued by the New York Police Department ("NYPD") pursuant to the Public Oversight of Surveillance Technology ("POST") Act.

The POST Act was passed in response to the NYPD's long and troubling history of engaging in surveillance tactics that target political dissent, criminalize communities of color, and jeopardize all New Yorkers' privacy. Despite years of assurances from the NYPD to the contrary, the City Council recognized the obvious fact that the NYPD cannot be trusted to monitor its own use of surveillance technologies and to keep the full extent of its surveillance infrastructure secret from the public and policymakers alike.

The law's mandate is simple: the NYPD must disclose the technologies it currently possesses – and going forward, technologies that it plans to acquire – along with the policies that govern their use. The information required to be made public under this law is the baseline information needed to evaluate the ways in which NYPD surveillance practices target communities of color; magnify discrimination in areas like immigration, housing, and education; and contribute to our continued overinvestment in and militarization of law enforcement.

On January 12, 2021, the NYPD published draft surveillance technology impact and use policies purporting to cover all existing surveillance technologies in use by the Department. We include comments on specific policies throughout, but at the outset, it is apparent that the NYPD's overall approach to fulfilling its disclosure obligations was to do the absolute minimum. These policies lack serious consideration of the potential for biased and disparate enforcement, include overbroad groupings and generalizations, contain deeply inadequate provisions on data sharing and retention, and are replete with inaccuracies and misleading statements.

NYCLU

ACLU of New York

1 Whitehall Street, 3rd Fl.
New York, NY 10004
nyclu.org

Donna Lieberman
Executive Director

Olivier Sylvain
President

Moreover, the repeated use of stock language, apparently copied and pasted across policies without any meaningful attempts to more directly tailor that language to specific technologies, suggests that the NYPD, which was loudly on record in opposition to the POST Act, remains opposed to calls for greater transparency and self-reflection. This template-based response has led to a number of overarching issues with the POST Act disclosures writ large, which we address below.

No Serious Consideration of Bias and Disparities

Nowhere is the NYPD's lazy copy-and-paste approach more disturbing than in its assessment of potential disparate impacts. By and large, these sections repeat, verbatim, pronouncements that the NYPD is committed to impartial enforcement of the laws and that all of their policies, seemingly without fail, have sufficient "safeguards and audit protocols" to "mitigate the risk of impartial and biased law enforcement."

Only three among the 36 policies include much in the way of additional substantive consideration – and rejection – of the risk of racially disparate enforcement. Of these, the facial recognition policy downplays the well-documented failures of such technology to accurately identify women and people of color.[1] The Department's analysis similarly does not address the fact that the databases upon which the technology relies for comparators are, themselves, generated by law enforcement and likely to reflect the disproportionate rate at which communities of color are policed and their images entered into such databases. While the Department references a federal study finding that human observers can correct for racial inaccuracies, this ignores studies that have found a tendency for people to be more likely to trust and accept machine recommendations.[2]

The data analysis tools policy simply states that human oversight is integrated into the review process and that periodic assessments take place, without providing further detail on what that entails. The policy

---

[1] *See e.g.:* Cynthia M. Cook et al., *Demographic Effects in Facial Recognition and their Dependence on Image Acquisition: An Evaluation of Eleven Commercial Systems*, 1 IEEE Transactions on Biometrics, Behavior, and Identity Science 32–41 (2019); Joy Buolamwini & Timnit Gebru, *Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification*, Proceedings of Machine Learning Research (2018), http://proceedings.mlr.press/v81/buolamwini18a/buolamwini18a.pdf (last visited Feb 19, 2021).

[2] Lauren Chambers, *Bias All the Way Down: Research Shows Domino Effect When Humans Use Face Recognition Algorithms*, Privacy SOS (2020), https://privacysos.org/blog/bias-all-the-way-down-research-shows-domino-effect-when-humans-use-face-recognition-algorithms/.

on criminal group databases notes that the NYPD's gang database has been subject to criticism for its disparate impact but does not even acknowledge the glaring fact that nearly 99% of those in the database are people of color.[3] The policy claims that, since inclusion in the database is "only an investigative lead," it does not produce disparate, collateral consequences, but this fails to recognize the obvious fact that a database in which investigative leads are generated on thousands of people of color is a disparate impact in and of itself.

Deficient as these three policies are, the remaining 33 policies do not even attempt this level of analysis. The NYPD policy on ShotSpotter does not, for instance, consider how placement of the technology in particular neighborhoods will necessarily result in police responding to alerts from those neighborhoods. There is no discussion in this policy as to what this translates to in terms of the demographic makeup of communities with ShotSpotter sensors installed. The body-worn camera policy does not consider the fact that the cameras will undoubtedly generate more recordings of people of color based on the simple fact that NYPD enforcement activities overall involve disproportionately greater enforcement against communities of color. The NYPD claims that its use of License Plate Readers is not motivated by, among other factors, race, color, religion, or national origin, yet we know that the NYPD has previously been motivated by precisely these factors when they deployed license plate readers near mosques as part of its aggressive and unlawful surveillance of Muslims.[4]

The Department should consider not just the technical capabilities of a given piece of technology in terms of potential disparities, but also how those technologies are used in routine enforcement that already produces disparities. And the NYPD must do this work for each surveillance impact and use policy individually if the Department is to argue that it takes concerns about bias seriously.

Overbroad Categories and Obscuring of Information

While some policies are specific to a discrete technology or platform, like ShotSpotter, others are so broad that it is difficult to discern what technologies are actually being described. The policy for data analysis tools, for example, apparently covers a wide array of technologies, an undefined "some" of which utilize artificial intelligence

---

[3] Nick Pinto, *NYPD Added Nearly 2,500 New People to its Gang Database in the Last Year*, The Intercept, June 28, 2019, https://theintercept.com/2019/06/28/nypd-gang-database-additions/.

[4] *NYPD Defends Legality of Spying on Mosques*, CBS News, Feb. 24, 2012, https://www.cbsnews.com/news/nypd-defends-legality-of-spying-on-mosques/.

and machine learning, while others – again undefined – do not. It's unclear what information the NYPD hopes the public can glean from this entirely lacking description; indeed, this broad grouping together of technologies with apparently quite different capabilities seems intended to avoid more detailed reporting on individual tools and their more specific uses.

This aversion to nuance is also evident in the lack of detailed information on particular vendors for the technologies described in these policies, as well as the absence of any information on the number of specific devices in the NYPD's possession and their cost. Rather than provide the public with information on which specific products the NYPD uses and which vendors it contracts with, most policies simply state that the NYPD purchases technologies from "approved vendors," with whom the NYPD emphasizes the importance of confidentiality, and describe in general terms the legal requirements to which all vendors are subject.

The NYPD's policies should directly name which vendors it contracts with, the number and specific types of technologies it obtains, and the total cost to the city. The NYPD has – selectively – provided similar information in the past. For example, when the Department announced that it would begin using drones, the NYPD provided a detailed breakdown on the fourteen devices it had purchased, including the specific models and the cost of the overall program.[5] Yet the Department's surveillance impact and use policy for unmanned aircraft system contains none of that information. Now that there is a legislative mandate to be transparent about its surveillance practices, there is no excuse for the Department not providing similarly detailed information on its full surveillance arsenal.

Inadequate Provisions on Data Sharing and Retention

The NYPD sweeps up vast amounts of private data on New Yorkers through its surveillance practices. A key reason advocates demanded passage of the POST Act was to better understand how long that data is being kept and whether the NYPD is funneling information to other agencies, including federal law enforcement. Rather than provide this information in these policies, the NYPD has made rampant use of stock language to obscure such details.

---

[5] Ashley Southall & Ali Winston, *New York Police Say They Will Deploy 14 Drones*, N.Y. Times, Dec. 4, 2018, https://www.nytimes.com/2018/12/04/nyregion/nypd-drones.html.

Almost all of the policies note that "[o]ther law enforcement agencies may request" information from the relevant technology and that NYPD disclosure of such information "is governed by applicable laws and regulations, and NYPD policies." Similarly, most policies go on to say that "[g]overnment agencies at the local, state, and federal level, including law enforcement agencies other than the NYPD, have limited access to NYPD computer and case management systems." Neither of these copied-and-pasted statements provide the public with any useful information on the degree to which such external agencies actually do request and are permitted access to NYPD data.

Essentially, these provisions answer the question "Does the NYPD share surveillance data externally?" with: Maybe. That is simply not good enough and not in keeping with the letter and intent of the law to require that the NYPD disclose whether New Yorkers' private information is being shared. The Department must revise each of its policies to provide more detailed information on the actual extent to which such information sharing takes place.

Further, while a handful of policies provide at least some minimal information on how long types of data are retained, the majority of policies simply state that information is "retained in accordance with applicable laws, regulations, and New York City and NYPD policies." Language like this once again flies in the face of the POST Act's purpose of making information about NYPD practices easily accessible to the public. These policies must be amended to include more detailed information on the actual retention policies in place for each technology instead of just sending New Yorkers on a fishing expedition to look up other relevant laws, regulations, and policies.

Inaccurate, Misleading, and Ambiguous Provisions

Inaccuracies, misleading statements, and ambiguous provisions abound in the draft impact and use policies. These include statements that fly in the face of publicly available information and prior NYPD admissions. They also demonstrate the degree to which the NYPD does the public a disservice by adhering so closely to stock language throughout these policies in lieu of releasing more detailed and individualized policies tailored to specific technologies.

The NYPD makes a number of claims that certain technologies do not use artificial intelligence, machine learning, or video analytics. However, in many of these cases, there is direct evidence to the contrary. The closed-circuit television systems policy claims that it does not make use of such processes, but the NYPD has clearly done so in the past, and

it has partnered with vendors that provide such analysis.[6] The NYPD's ShotSpotter policy similarly claims that no artificial intelligence or machine learning is utilized, notwithstanding the fact that ShotSpotter's official website devotes a section to "Artificial Intelligence and Machine Learning" on its "Technology" landing page.[7] We know that the Department utilizes the machine learning tool Dataminr[8] – despite the fact that it is not specifically listed in any draft impact and use policy – but neither of the two policies where we might expect to find it (media aggregation services and social network analysis tools) admit to the use of machine learning tools.

Incredibly, the Department's facial recognition policy suggests that no artificial intelligence or machine learning are used, despite the fact that most facial recognition systems rely on exactly those mechanisms as a basic function. Since the policy does not provide more detailed information on exactly which vendor or product the NYPD uses, it is impossible to verify whether this improbable statement is accurate.

These same issues extend to the policies for the license plate readers and the Domain Awareness System. A 2017 article reviewed the NYPD's use of pattern recognition technology in conjunction with license plate readers, and the same article provided a comprehensive overview of the Domain Awareness System.[9] That article explicitly described the use of machine learning, pattern recognition, video analytics, and sensors,[10] all of which are disclaimed in the draft surveillance impact and use policy. The NYPD has also confirmed its use of Patternizr, yet another machine learning tool, as part of the Domain Awareness System.[11]

---

[6] *See* U.S. Dep't of Homeland Security, *Public Safety Analytics Terminal: Technology Scouting Resarch Summary*, Sept. 2019, https://www.dhs.gov/sites/default/files/publications/public_safety_analytics_terminal_updated_v3.pdf; IBM, *IBM SVS4.0 Research and Development Status Update 6 for NYPD*, Oct. 16, 2012, https://www.documentcloud.org/documents/4452844-IBM-SVS-Analytics-4-0-Plan-Update-for-NYPD-6.html.

[7] *See* ShotSpotter Technology, https://www.shotspotter.com/technology/ (last visited Feb. 22, 2021).

[8] Sam Biddle, *Police Surveilled George Floyd Protests with Help from Twitter-Affiliated Startup Dataminr*, The Intercept, July 9, 2020, https://theintercept.com/2020/07/09/twitter-dataminr-police-spy-surveillance-black-lives-matter-protests/.

[9] E. S. Levine et al., *The New York City Police Department's Domain Awareness System*, 47 INFORMS Journal on Applied Analytics 70–84 (2017).

[10] *Id.*

[11] Stephanie Kanowitz, *NYPD's Machine-Learning SoftwareSpots Crime Patterns,* GCN, Apr. 11, 2019, https://gcn.com/articles/2019/04/11/nypd-crime-patterns-ml.aspx;

While NYPD claims that its situational awareness cameras make no use of machine learning or video analytics, the "Digidog" described in that policy clearly possesses such capabilities.[12] The policy on digital forensic access tools states that the technology is not used "to engage in unauthorized access of 'hacking' of electronic devices," but that is precisely what tools meant to circumvent encryption do. In 2019, it was reported that the NYPD had been using a tool called Cellebrite to hack and decrypt phones.[13] And once again, contrary to the policy's assertions that no digital forensic access tools make use of artificial intelligence or machine learning, Cellebrite does just that.[14]

Beyond these outright falsehoods, the policies are awash in misleading statements, ambiguities, and omissions. The draft cell-site simulator policy states that the devices "are not used to collect the contents of any communications or any data contained on the device itself." Omitted from this statement is whether the devices in the NYPD's possession are *capable* of performing such functions, as some cell-site simulator configurations are.[15] The NYPD could resolve this ambiguity by more directly describing the full capabilities of the device in addition to its actual uses, as well as by providing specific information on the actual cell-site simulator models it uses. The license plate reader policy similarly omits some key details, including the fact that the NYPD has subscribed to the Vigilant Solutions commercial license plate database, which provides the Department with access to billions of license plate reads.[16]

If there are this many falsehoods, omissions, and ambiguities in disclosures related to the technologies that we already do know about, it raises serious questions about the degree to which the NYPD is still keeping critical information on its otherwise unknown surveillance

Alex Chohlas-Wood & E. S. Levine, *A Recommendation Engine to Aid in Identifying Crime Patterns*, 49 INFORMS Journal on Applied Analytics 154–166 (2019).

[12] James Carroll, *Artificial intelligence software expands Boston Dynamics' Spot robot capabilities*, Vision Systems Design (2020), https://www.vision-systems.com/embedded/article/14179537/artificial-intelligence-software-expands-boston-dynamics-spot-robot-capabilities (last visited Feb 19, 2021).

[13] Steven Melendez, *Oh, Great! NYC Law Enforcement can Probably Hack Your Phone Now*, Fast Company, Oct. 8, 2019, https://www.fastcompany.com/90414742/report-nyc-law-enforcement-can-hack-your-iphone-android.

[14] *Id.*

[15] Kim Zetter, *How Cops can Secretly Track Your Phone*, The Intercept, July 30, 2020, https://theintercept.com/2020/07/31/protests-surveillance-stingrays-dirtboxes-phone-tracking/.

[16] Nathan Tempey, *The NYPD is Tracking Drivers Across the Country Using License Plate Readers*, Gothamist, Jan. 26, 2016, https://gothamist.com/news/the-nypd-is-tracking-drivers-across-the-country-using-license-plate-readers.

capabilities secret from the public, despite the POST Act's requirements that demand such transparency.

***

The POST Act was intended to provide the public with enough information to actually understand the NYPD's use of surveillance technologies and to empower communities to engage in informed conversations about the ways in which they are policed. These draft policies suggest that the NYPD remains unwilling to engage in this effort. Instead of taking this opportunity to prove the Department's oft-stated commitment to transparency and oversight, these overbroad, inadequate, and inaccurate policies are the latest in a long series of examples of the NYPD rejecting calls for true democratic accountability. The NYPD must issue revised policies that account for and correct these deficiencies and that give full effect to the POST Act's legal mandates. And City lawmakers must seriously reconsider the degree to which the NYPD has been empowered to assemble its massive surveillance infrastructure in the first place and to deploy it without any serious regard for the impact on Black and Brown communities.

Sincerely,

Michael Sisitzky
Senior Policy Counsel

Daniel Schwarz
Privacy & Technology Strategist

Justin Harrison
Senior Policy Counsel

cc: Mayor Bill de Blasio
City Council Speaker Corey Johnson
City Council Committee on Public Safety Chair Adrienne Adams