



40 Rector Street, 9th Floor
New York, New York 10006

www.StopSpying.org | (646) 602-5600

STATEMENT OF
ALBERT FOX CAHN, ESQ.
EXECUTIVE DIRECTOR
SURVEILLANCE TECHNOLOGY OVERSIGHT PROJECT, INC.

&

LIZ O'SULLIVAN
RESIDENT TECHNOLOGIST
SURVEILLANCE TECHNOLOGY OVERSIGHT PROJECT, INC.

BEFORE THE
COMMITTEE ON TECHNOLOGY
NEW YORK CITY COUNCIL

FOR A HEARING CONCERNING,
OVERSIGHT - UPDATE ON LOCAL LAW 49 OF 2018 IN RELATION TO
AUTOMATED DECISION SYSTEMS USED BY AGENCIES.

PRESENTED
April 4, 2019

Good afternoon, my name is Albert Fox Cahn, and I serve as the Executive Director for the Surveillance Technology Oversight Project (“STOP”). STOP advocates and litigates for New Yorkers’ privacy rights, fighting discriminatory surveillance. I commend the committee and Chairman Koo for today’s hearing; for bringing much-needed attention to the stagnating work of the New York City Automated Decision System (“ADS”) Task Force.

In 2018, I was proud to partner with the City as part of its ADS Task Force, meeting with City leaders, academics, and advocates to shape recommendations for the future role of artificial intelligence in New York City Government. In my initial discussions with the Task Force, I was hopeful that we would leverage this unique opportunity, setting a national model of how to address the role of artificial intelligence and other ADS in government. Sadly, more than halfway through the Task Force’s lifecycle, those hopes are almost completely gone.

Outreach

For an entity tasked with promoting transparency, the Task Force itself is alarmingly opaque. Rather than engaging with the public, soliciting input on the ADS questions that impact every single New Yorker, the Task Force hosted a single web page. In my time working with the Task Force, I was never notified about a single public comment. To the extent that public comments were made, whether through 311 or the NYC.gov website, they were never presented to Task Force members during my tenure. More recently, the Task Force has expanded public outreach, but I fear it’s too little, too late.

Given the complexity of ADS, a city-wide outreach and engagement campaign would have been difficult during the full eighteen-month lifecycle of the Task Force, but it’s quite impossible to shoe-horn such efforts into the short remaining time. Additionally, it’s concerning to see the lack of clarity about the purpose of the Task Force’s for a. We still don’t know the extent to which Task Force members will be available for public questions or the extent to which this will be a truly democratic process, versus engagement from hand-picked participants.

Role of city agencies

Local Law 49 of 2018 envisioned the Task Force as being independent.¹ That’s why the Council required it to include experts on ADS “fairness, accountability and transparency” and representatives of nonprofits that “represent persons in the city affected by” ADS.² The goal was to create a body of experts who could hold the city’s feet to the fire and challenged agency norms around ADS.

Unfortunately, that does not reflect the composition of the Task Force that we’ve seen in practice. Officially, one-third of the Task Force members are city employees, including representatives from ACS, the NYPD, and DOE. This number does not include unofficial Task Force members, the numerous city staffers who have effectively driven this process from the start, despite having no official position on the Task Force. Rather than being a member-driven undertaking, the Task Force has been repeatedly asked to sign-off on plans and language that were already set by city officials, with limited input from the experts who are meant to control this process.

¹ Local Law No. 49 (2018) of City of New York § 1(B)(2), Available at <https://legistar.council.nyc.gov/View.ashx?M=F&ID=5828207&GUID=DDC22B3B-B5F8-41D7-AE14-2A7BC5838233>.

² *i.d.*

Task Force Management

One example of Task Force members being structurally disempowered is the role of the Jane Family Foundation. The foundation effectively drove the process for much of the first year, despite never being an official, publicly recognized part of the Task Force. The foundation's role grew from initially providing background briefings to providing proposed language and policy documents for the Task Force to ratify. Increasingly, the foundation was writing a first draft of the Task Force's report. The foundation's role drew complaints from numerous Task Force members, so it was eventually phased out, but it's a telling example of how the role of Task Force members themselves was circumscribed as part of this process.

This organizational structure led to a lack of Task Force member buy-in on key questions that go to the heart of automated decision systems. We should have asked the nationally-recognized experts who sat on the task force to decide how to shape and engage with this process. Instead, they were asked to sign off on proposals put forward by outside groups. Rather than being asked to set the agenda for meetings, Task Force members were often presented with the City's pre-planned list of agenda items. While Task Force members could always provide feedback, it felt like commentators on the process we were supposed to lead. Unfortunately, this internal dysfunction has come at a cost.

Lack of access to info about relevant city systems

Sadly, the lack of transparency goes even farther. Not only has the city failed to make the Task Force accessible to the public, but it has also failed to make city records accessible to Task Force members themselves. Throughout this process, Task Force members have pushed to learn more about the ADS that New York uses already, but the Mayor's office has refused. It's hard to overstate how large an obstacle this presents for the Task Force.

How can we ask Task Force members to build a map to the future if we don't know where we are today? It's challenging enough to imagine all the ways that ADS might become a part of city life in the future, but it's impossible to do so if we don't know how those systems are already being deployed. Going forward, the Task Force must have a comprehensive listing of ADS deployed around the city, along with the background and context information needed to understand how they impact New Yorkers. Without this information, the Task Force becomes nothing more than an academic exercise.

Losing our lead

At the time I ended my work with the Task Force, nearly a year after its launch, we had yet to reach agreement on the most basic questions about what an ADS. This includes a lack of consensus about what even constitutes an ADS. We had no clear pathway to a consensus on the countless derivative questions that also needed to be addressed, such as the basic regulatory framework for ADS, an effective model of dispute resolution, or way in which our findings would be updated going forward

Sadly, while our city was initially a leader in the national movement for ADS transparency, we have quickly fallen behind other localities that have engaged in the type of public outreach that is indispensable for the Task Force's work. Vermont, which created a statewide Artificial Intelligence Task Force five months after New York City, has already engaged in the public outreach that the

city is only now contemplating.³ The Vermont Task Force has had experts testify at publicly broadcasted hearings, it publicly posts its meeting agendas and minutes, and its Task Force is expected to release a report by this summer, far ahead of our own.⁴

In another case, when the Pennsylvania Sentencing Commission engaged in broad public engagement around the role of ADS in risk assessment. It held public hearings throughout the state, including advocates, people affected by the criminal justice system, researchers, lawyers, and lawmakers.⁵ Additionally, when California evaluated its state law, replacing cash bail with risk assessment ADA, advocates provided prolific comments on every aspect of the policy.⁶

Best Practices

In contemplating the Task Force's recommendations, we must remember that algorithms are not magic. Like all statistical models, they have limitations. New York City should look to some of the best practices found in the private sector, such as “model cards” that explain a model's methodology and limits.⁷ Additionally human training practices are needed to teach decisionmakers how bias (conscious and unconscious) impacts ADS outputs, along with the impact of “automation bias.”⁸

Moreover, we must make sure that our city-wide community engagement primarily targets directly impacted communities. For these New Yorkers, ADS is not an abstraction, but a potentially insurmountable barrier to their next job, their preferred school placement, or even maintaining their freedom. These communities' voices must be heard if we're to ensure that ADS are not merely

³ H.B. 378, 91 Leg., Reg. Sess. (Vt. 2018), <https://legislature.vermont.gov/Documents/2018/Docs/ACTS/ACT137/ACT137%20As%20Enacted.pdf>; H.B. 2701, 191 Leg., Reg. Sess.(Ma.2019), <https://malegislature.gov/Bills/191/HD951>; H.B.1655, 66 Leg., Reg. Sess. (Wa. 2019), <https://app.leg.wa.gov/billssummary?BillNumber=1655&Initiative=false&Year=2019>.

⁴ Vermont Agency of Commerce and Community Development, Artificial Intelligence Task Force, <https://accd.vermont.gov/economic-development/artificial-intelligence-task-force>.

⁵ See e.g, Samantha Melamed, Move afoot to dismantle work of Sentencing Commission, Pitts. Post-Gazette (Dec. 12, 2018),<https://www.post-gazette.com/news/politics-state/2018/12/12/Move-afoot-to-dismantle-work-of-Sentencing-Commission/stories/201812120200>

⁶ Upturn, Comments on Proposed California Rules of Court 4.10 and 4.40 (2018), https://www.upturn.org/static/files/2018-12-14_Final-Coalition-Comment-on-SB10-Proposed-Rules.pdf; Human Rights Watch, Comments on Proposed California Judicial Council Rules 4.10 and 4.40 (2018), <https://www.hrw.org/news/2018/12/10/human-rights-watch-comments-california-judicial-council-bail-reform-rules>; Electronic Frontier Foundation, Written Comments on Proposed Rules 4.10 and 4.40 (2018), <https://www EFF.org/document/written-comments-eff-proposed-california-rules-court-410-and-440>; JusticeLA S.B.10 Comments (2018), <http://justicelanow.org/sb10comments/>.

⁷ Model cards explain the training materials, methodology, limitations, known biases, and unknown or untested capacities that the models might harbor. Understanding how narrow the focus of a model is, or whether it includes racial features or racial proxy features like zip Code can impact the decision-making of a human agent involved in reading the output of the algorithm.

⁸ “Automation bias” is the phenomenon that people presented with an algorithmic prediction will confirm its truth rather than deny it. As we saw in the Boeing 737 case, poor training can result in catastrophic outcomes, especially when the machines and the humans disagree.

transparent, but also just. ADS justice also requires explainability tools with which New Yorkers can understand how they have been impacted by ADS.⁹

In addition to explainability tools, we must deploy fairness testing that examine ADS for potential bias.¹⁰ We must agree-upon a metric with which to audit ADS and ensure that biases are not coded into our government.¹¹ But algorithmic fairness is not enough, since even a “fair” algorithm will produce discriminatory results when applied to a discriminatory dataset.¹² While there is no such thing as an “unbiased algorithm”, ongoing and rigorous evaluation can reduce the biases that are unavoidable in any human-designed systems.

Next Steps

To make the most of the remaining time, it’s essential that the Task Force give a broad cross-section of the public a meaningful opportunity to voice their concerns. The planned fora should include a chance for New Yorkers to question and speak to members directly. We must better publicize the work done by the Task Force to date, publicizing meeting agendas and summaries. We should empower Task Force members to understand the full range of ADS operating in New York City and how the data from those systems is used. Lastly, the Task Force’s report must memorialize the full range of opinions expressed by members. A document that effectively preserves members’ competing positions is far more helpful than a document that provides us with artificial consensus.

⁹ Any particular transaction should be made available through a searchable database that provides explainability metrics for why that decision was the output. This is the only way that humans will be able to appeal machine-made decisions.

¹⁰ There are more than 50 definitions of fairness, but one of the most prominent is “equality of odds.” This test asks if the algorithm makes equally many errors (false positives and false negatives) on people of one protected group compared to another.

¹¹ For example, reducing gender categories to “Male” or “Female” on a driver’s license excludes non-binary individuals from representation. This flattens/simplifies the data, representing individuals more crudely, and thus with less predictive power, something known as the “undersampled majority.”

¹² An instructive example is predictive policing, which often relies on evidence about past crimes to determine where future crimes will likely occur. Since historic policing data reflects bias against communities of color, which face higher rates of policing and arrests (including false arrests), the product of such predictive policing algorithms is skewed against those same communities. Data collection and labeling must examine and remedy such biases before this data can be used for ADS.