FRIDAY, NOVEMBER 1, 2019

CONVERSATIONS ON POLICING TECH:

CONVENING ON REGULATING FACIAL RECOGNITION

FALL 2019 SALON SERIES

THE POLICING PROJECT

About the Policing Project Salon Series

As part of our ongoing work to help promote the ethical use of policing technology, the Policing Project is hosting a series of closed-door salons to work through some of the most difficult questions we face. Made possible through support from Microsoft, the salons enable us to vet our projects and discuss pressing issues around law enforcement's use of technologies with a diverse set of experts, including privacy advocates, technology vendors, police chiefs, academics, legal experts, community leaders, and government officials.

About the Policing Project

We partner with communities and law police to promote public safety through transparency, equity and democratic engagement.

Our work focuses on front-end, or democratic, accountability — meaning the public has a voice in setting transparent, ethical, and effective policing policies and practices before the police or government act. The goal is to achieve public safety in a manner that is equitable, non-discriminatory, and respectful of public values.

For more information, visit **www.PolicingProject.org**

The report was written by Policing Project Technology Fellow Emmanuel Mauleon.

EVENT DESCRIPTION AND REVIEW

Several jurisdictions are using facial recognition to aid law enforcement—from helping identify suspects and witnesses, to locating missing children. At the same time, other jurisdictions have banned facial recognition outright, fearing the potential for abuse of such technologies, and warning of the severe invasion of privacy posed by such systems.

At the Policing Project, we favor an approach to emerging technologies that depends on public authorization and careful regulation to maximize any benefits while minimizing, or eliminating entirely, anticipated harms. However, these issues are exceedingly complicated around facial recognition. There are sharp societal disagreements on where to draw the lines between acceptable and unacceptable uses, and many pressing questions that remain, including:

- How do we measure and ensure that facial recognition is reliable and unbiased?
- What are acceptable law enforcement uses of this technology, and what uses go too far?
- How do we regulate the data and databases used to train algorithms, and the databases against which faces are matched, to ensure fairness and equal justice?

The Policing Project and the Information Law Institute at NYU Law recently hosted a diverse group of stakeholders—including privacy advocates, technology vendors, policing leaders, academics, legal experts, and government officials—to discuss these questions. The convening was held under a modified Chatham House Rule that permits us to capture the nature of the conversation and identify participants, without attributing a particular viewpoint to any individual.

The day was broken into three sessions, focusing on accuracy and bias, acceptable use-cases of facial recognition, and databases.

Pictured clockwise from top: Hanna Wallach from Microsoft Research New York and Professor Jason Schultz from NYU School of Law; Detroit Police Chief James Craig; and Jumana Musa, director of the Fourth Amendment Center of the NACDL.

SESSION 1: ACCURACY & BIAS

The first session, led by **Policing Project Executive Director Farhang Heydari**, began with a seemingly simple question—what does it mean for a facial recognition system to be accurate? As it turns out, this question proves much harder to answer than to ask.

It quickly became apparent that the advertised rates of accuracy of facial recognition systems (e.g., "99.9% accurate") can be quite misleading. For example, a participant shared that the accuracy figures of their system were determined after a human picked through dozens of false-matches; another system only accounted for searches that returned positive matches. Often, advertised accuracy results reflect a controlled testing environment, not real-world conditions. Vendor or law enforcement representations as to the accuracy of these systems could present a difficult evidentiary presumption to overcome, in which a jury (or even simply an investigator) may latch onto a number—such as "99.9% accurate"—without understanding the context in which that number was derived.

Participants also spoke to the biases and disparities—particularly across different races and genders—that have been shown in different facial recognition systems. The discussion covered considerable ground, with some clear takeaways: Although algorithmic disparities are often deemed "technical errors," they can have real life costs, including stops, arrests, and imprisonment. As such, deploying facial recognition technology prior to achieving parity in accuracy across different groups is a deeply concerning proposition. Although some dangers can be mitigated by ensuring that there is human-in-the-loop review, even such a safeguard cannot ensure that the technology is used equitably—there are those who question whether instead it will simply mirror existing disparities in policing.

Although both accuracy and bias are critically important issues to address, participants agreed that there is a paucity of concrete information on these aspects of facial recognition systems. There remains a need for more information, transparency, and research into facial recognition systems before a rigorous evaluation of the accuracy and disparities in such systems can take place, and thoughtful regulation can be crafted to address any issues.

SESSION 2: REGULATING SPECIFIC USES AND LIMITING MISSION CREEP

The second session, led by the **Policing Project Founder Barry Friedman**, asked whether effective regulation could distinguish between acceptable and unacceptable uses of facial recognition. This transformed into a broader discussion about whether it necessarily was possible to distinguish among uses, or to prevent mission creep once certain practices were allowed.

Some law enforcement agencies already employ facial recognition for one "use": to identify suspects, witnesses, and even victims. Yet despite employing facial recognition for the same use case, participants described widely varying methods to accomplish the same goal.

Alvaro Bedoya, director of the Center on Privacy and Technology.

Deputy Chief Sabih Khan, Chicago Police Department, and Assistant Chief James Wilcox, New York City Police Department.

Information Law Institute Faculty Director Katherine Strandburg.

Professor Solon Barocas from Cornell University and Microsoft Research New York.

Jeramie Scott from the Electronic Privacy Information Center.

For example, some agencies use a system that provides just one possible match, while others provide up to fifty. Several had dedicated officers trained to sift through possible matches, and return results to detectives only after an additional investigation to verify the identity of a particular suspect. Others allowed any officer to run a search. Every department, however, utilized a "human-in-the-loop" screening for identification purposes—none simply took the best algorithmic match as definitive identification of a suspect.

Although legislation could be crafted that delineates acceptable uses, many participants relied upon history to argue that mission creep is inevitable. They also made the point that historically abuses of technology by police routinely accrue against vulnerable communities, frequently poorer communities, communities of color, and immigrant communities.

Participants acknowledged the need for additional guardrails to constrain law enforcement from particular uses, particularly persistent surveillance. However, they recognized the difficulty of this if and when facial recognition is used widely in the private sector. For example, a law that banned law enforcement surveillance but permitted private sector use may lead police to solicit facial recognition information from the private users, much like how they may ask individuals for access to past security footage.

Some felt that even for location tracking, facial recognition could be used as a limited tool of last-resort, perhaps only for certain serious crimes, or when a suspect in a serious crime could not be located. Some also argued that judicial oversight might provide an effective means of addressing unique situations with appropriate care.

Although others remained skeptical of the effectiveness of such limits, there was acknowledgement that any eventual regulation likely would need to be tweaked in different contexts. There was greater consensus still that, to the extent facial recognition was used, notice should be provided to the persons who were being charged with a crime in which facial recognition had played any part, and potentially those whose faces were used in comparison sets.

SESSION 3: DATA & DATABASES

The final session, moderated by **Information Law Institute Director Katherine Strandburg**, centered on the data and databases that are employed in facial recognition or result from its use. She posed several broad questions at the start of the session to guide the conversation:

- What data sources should we allow as comparison sets, and which should we constrain?
- Should we be as comprehensive as possible, or as limited as possible?
- Should we allow commercial access to public databases, and vice versa?
- How do we demarcate the lines of retention, and what information is retained, searchable, and available for analysis?

When discussing training data—collections of unique images of faces used to improve the underlying algorithms of facial recognition software—some quickly drew a distinction between scraped-data (images pulled from public repositories or from across the Internet) and consented data (images for which companies received individualized permission to use). One recommendation was that facial recognition legislation may need to assert a right to control one's face and data, mirroring European protections and giving individuals more control whether to opt in or out of particular databases.

Concerns over bias and skewed data were voiced about both types of datasets. Participants noted parallels in medical research based upon coercive incentives that affect who agrees to be a test subject and who might not. Moreover, participants noted that often the person who consents to terms of use and uploads a photo to a web service used to train algorithms may not be the only one featured in a photo, muddying the question of consent. There was sharp disagreement in the room about which comparison or "target" photo sets would yield the most equitable and effective facial recognition results, while producing minimal harms.

Some posited that the only equitable solution is a dataset which contains as many faces as possible, such as a national photo book. They argued that making universal inclusion the standard could force a more robust conversation on the use of facial recognition systems among people who otherwise would feel comfortable using the technology on others, but not themselves. And they made the point that if the goal of facial recognition is to identify suspects, witnesses, and victims, the widest possible dataset made sense.

Others cautioned, however, that universal inclusion would not necessarily result in equal enforcement. Instead, larger datasets likely could be used to target minorities and vulnerable communities in even greater numbers and with more efficiency. Many pointed to China's tracking and detention of the Uighurs—aided through facial recognition—as a chilling example. Regulation would be needed to propel concurrent changes to how policing is performed, including guidelines for which neighborhoods would be targeted, and where cameras would be placed and deployed.

Limited comparison sets, such as arrest photos (which tend to over represent over-policed segments of the public), presented other hard regulatory questions. Limiting algorithmic comparisons to mugshots likely would compound previous biases and prove ineffective at catching first-time offenders. One participant pointed out that even if limited to those convicted of crimes, most cases are plea-bargained, which could result in disparities between those with the means to hire an attorney and those without. Participants agreed that any eventual regulation would need to ensure that facial recognition did not simply serve as a means of perpetuating biases already extant in the criminal justice system.

CONCLUSION

Overall, the discussion was incredibly robust. We learned a great amount about how current facial recognition systems work, how accuracy is determined by end-users and vendors, and what red flags should be addressed prior to adoption of facial recognition by any jurisdiction. Although disagreements remained among participants as to the best means of regulating facial recognition, the convening provided a serious venue for addressing concerns and presenting possible solutions for a technology that is not likely to disappear into the background any time soon. We hope to build on the lessons we've learned as we consider the future of facial recognition in the coming years, and as we proceed with a planned convening on biometrics. ∎

Pictured clockwise from top: Roxane Panarella, assistant general counsel for the FBI, Matthew Cagle from the ACLU of Northern California, and Lee Tien from the Electronic Frontier Foundation; Assistant U.S. Attorney Monica Ryan; and Megan Quattlebaum, director of The Council of State Governments Justice Center.

ATTENDEES

FALL SALON SERIES

Solon Barocas, *Principal Researcher*, Microsoft Research New York, and *Assistant Professor*, Cornell University

Alvaro Bedoya, *Founding Director*, Center on Privacy and Technology, Georgetown Law

Steve Block, *Director*, AWS Public Policy, U.S. Federal, Amazon

Matthew Cagle, *Technology and Civil Liberties Attorney*, ACLU of Northern California

Albert Cahn, *Executive Director*, Surveillance Technology Oversight Project (STOP)

Theodore Christakis, *Professor of International Law*, Université Grenoble Alpes and *Co-director*, Grenoble Alpes Data Institute

James Craig, *Chief*, Detroit Police Department

Brandon del Pozo, *Former Chief*, Burlington Police Department

Emiliano Falcon, *Technology for Liberty Policy Counsel*, ACLU of Massachusetts

Andrew Ferguson, *Professor of Law*, University of the District of Columbia, David A. Clarke School of Law

Barry Friedman, *Jacob D. Fuchsberg Professor of Law and Affiliated Professor of Politics*, and *Faculty Director*, Policing Project, NYU School of Law

Farhang Heydari, *Executive Director*, Policing Project, NYU School of Law

Brian Hofer, *Chair*, Oakland Privacy Advisory Commission, and *Chair and Executive Director*, Secure Justice

Pamela Hrick, *Associate*, Stockwoods LLP

Sabih Khan, *Deputy Chief*, Bureau of Technical Services, Chicago Police Department

Jumana Musa, *Director*, Fourth Amendment Center, and *Staff Attorney*, National Association of Criminal Defense Lawyers (NACDL)

Roxane Panarella, *Assistant General Counsel*, Criminal Justice Information Services, FBI

Megan Quattlebaum, *Director*, The Council of State Governments Justice Center

Monica Ryan, *Assistant U.S. Attorney*, District of Arizona

Jason Schultz, *Professor* of Clinical Law, *Director*, Technology Law & Policy Clinic, and *Area Lead in Law & Policy*, AI Now Institute, NYU School of Law

Jeramie Scott, *Senior Counsel* and *Director*, Domestic Surveillance Project, Electronic Privacy Information Center (EPIC)

Katherine Strandburg, *Alfred B. Engelberg Professor of Law*, and *Faculty Director*, Information Law Institute, NYU School of Law

Lee Tien, *Senior Staff Attorney* and *the Adams Chair for Internet Rights*, Electronic Frontier Foundation (EFF)

Hanna Wallach, *Senior Principal Researcher*, Microsoft Research New York

Andrew Weissmann, *Distinguished Senior Fellow*, Center on the Administration of Criminal Law, New York University School of Law

Jason Wilcox, *Assistant Chief*, New York City Police Department