UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF COLUMBIA

UNITED STATES OF AMERICA, v. [redacted]., Defendant.	**CASE NO.** [redacted]

AFFIDAVIT OF CLARE GARVIE

Clare Garvie, having been duly sworn, deposes and says under penalty of perjury:

1. I am an attorney and senior associate at the Center on Privacy & Technology, a think tank based at the Georgetown University Law Center. I have studied and researched face recognition systems, their use by law enforcement agencies, and the federal and state laws that apply to these systems, since 2015. Through the course of my work, I have submitted more than 200 records requests to public agencies across the country regarding their use of face recognition technology, and have reviewed over 20,000 pages of records received in response to those requests. Based on this research I have gained a unique and thorough understanding of how law enforcement agencies use face recognition technology in conducting criminal investigations. My *curriculum vitae* is attached to this Affidavit.

2. I have co-authored or authored five academic reports on the subject.[1] I have trained over 2,000 attorneys on the use of face recognition in criminal cases, and have testified

[1] Clare Garvie, Alvaro Bedoya & Jonathan Frankle, The Perpetual Line-Up: Unregulated Police Face Recognition in America, Center on Privacy & Technology (Oct. 18, 2016), https://www.perpetuallineup.org/; Clare Garvie, Garbage In, Garbage Out: Face Recognition on Flawed Data, Center on Privacy & Technology (May 16, 2019), https://www.flawedfacedata.com/; Clare Garvie & Laura Moy, America Under Watch: Face Surveillance in the United States, Center on Privacy & Technology (May 16, 2019), https://www.americaunderwatch.com/; Jameson Spivack & Clare Garvie, A Taxonomy of Legislative Approaches to Face Recognition in the United States, AI Now

before Congress[2] and state and local legislatures on the subject. I serve as a technical expert to defense attorneys and journalists on police use of face recognition.

3. The following information contains matters of fact that are based on my aforementioned research and are true to the best of my knowledge. Any opinions stated in this affidavit reflect opinions based on this research.

Background

4. When used as an investigative tool, face recognition is a subjective feature comparison method. The overall foundational validity of face recognition as a forensic tool, as it is used in a typical U.S. law enforcement investigation (detailed below), has yet to be established through empirical, peer-reviewed study. In the absence of this, the reliability of identity evidence produced by face recognition investigative searches overall is not known, highly variable, and can only be understood on a search-by-search basis by examining the choices, biases, motivations, and degree of training of the human in the loop.

5. A single search will involve a number of both machine and human decision points, each of which introducing the possibility for faulty interpretation and cognitive bias, particularly in the absence of specialized forensic training or technical controls in place. Depending on the choices made by an agent at each step, the results of the face recognition system may vary widely, as will the possibility of misidentification. To understand the reliability of an identification produced by a face recognition search, therefore, it is vital to understand what happens during each of these steps.

Institute (Sept. 02, 2020), https://ainowinstitute.org/regulatingbiometrics.pdf; Clare Garvie, Face Recognition and the Right to Stay Anonymous in M. Ienca et. al. (eds), Cambridge Handbook of Information Technology, Life Sciences and Human Rights, Cambridge Univ. Press (forthcoming 2021).

[2] Clare Garvie, Testimony before the House Committee on Oversight and Reform on "Facial Recognition Technology: Its Impact on our Civil Rights and Liberties (May 22, 2019), https://oversight.house.gov/legislation/hearings/facial-recognition-technology-part-1-its-impact-on-our-civil-rights-and

6. Since the reliability of the identification of the defendant as the subject of the search speaks directly to the defendant's guilt or innocence, information about each of these steps must be disclosed during discovery to ensure the defendant is afforded due process. Depending on how they are conducted, each of these steps may also introduce the possibility that results of a search is unduly suggestive, biased, or overall unreliable as identity evidence.

7. The following steps take place during a law enforcement face recognition search at various points in time: (1) selecting which probe photo (the image of the unknown subject of the search) to use; (2) selecting the face recognition database to search against; (3) editing the probe photo prior to search; (4) selecting the algorithm to perform the search; (5) interpreting the results of the algorithm; and (6) confirming the identification made by the face recognition algorithm through further investigation. Not all steps will be present in every search. The remainder of the affidavit is organized around these steps.

Selecting the Probe Photo

8. The accuracy of face recognition systems is in large part determined by the quality and contents of the probe photo submitted to the algorithm. The probe photo is the photo of the unknown subject that an officer is seeking to identify. The less information the probe photo contains about what the subject looks like, the less information the algorithm has to process, and the less reliable the resulting identification will be. Low-quality probe photos may be blurry or pixelated; show a partial, obscured, or side view of the subject's face; be over- or under-exposed; have lens glare or distortions, or have another imperfection.

9. If an agent has multiple probe images to choose from, such as with surveillance camera footage with more than one frame containing the subject's face, that agent must decide which probe

photo(s) to run against the system. Information about why a certain image was selected, as well as whether other probe photos were run against the system and produced different, or no, matches will speak to the reliability of the identification evidence. If, for example, one frame of the subject's face produced no matches, or if different frames produced different confidence scores to the defendant's database photo, that may raise doubt about the identification of the defendant as the subject.

Selecting the Database

10. Law enforcement agents may have a choice about which face recognition database or databases to run a probe photo against. Most face photo databases on file with state and federal agencies are now face recognition databases and may be accessible to search by the investigating agency. Since a system can only identify someone enrolled in the database searched, information about what databases are searched, and the contents of those databases, speak to the reliability of the identification. For example, if the probe photo is searched against a database containing the defendant's photo and that photo is not returned, that may raise doubt about the identification of the defendant as the subject even if a search of a different database produced a possible match.

Editing the Probe Photo

11. It is not uncommon for the analyst or agent who is running the search to edit the probe photo or photos before submitting them to the face recognition algorithm, or for the system itself to make automated adjustments to a photo to correct for pose or other variations. What edits are made to a probe photo are inherently subjective, highly variable, and will not necessarily accurately represent identity information belonging to the subject in the probe photo. A non-exhaustive list of the types of edits law enforcement agents make to probe photos include:

a. Inserting open eyes from a photo of a different person in place of the subject's closed or averted eyes;

b. Inserting a mouth and chin into the subject's photo from a photo of a different person when the subject's mouth is open or obscured;

c. Using 3D modeling software or photo editing software to rotate the subject's face within the two-dimensional photo, filling in the missing information based on what an average face or the visible portion of the subject's face looks like;

d. Mirroring over a partial photo of the subject's face to create a complete face;

e. Using the "blur" tool in Photoshop or similar photo editing software to add in pixels to an otherwise blurry or low-quality photo of the subject;

f. Combining the subject's photo with a photo of a different person to create a less pixelated photo; and,

g. Replacing the subject's photo entirely with a "celebrity lookalike" when the subject's photo is of too poor quality to generate a match.

12. All these types of edits introduce new information to the face recognition algorithm that is not present in the original photo of the subject and thus does not reflect the subject's identity. This new information may be fabricated by a software program, such as using a blur tool to add pixels, or sourced from photos of people other than the subject of the search. The face recognition algorithm will not distinguish between what is added information and what is original evidence, giving the "noise" the same weight as the true identity evidence of the subject.

13. When sourced from photos of different people, these edits will add identity evidence of another person into the subject's biometric template. The algorithm has no way to know which

evidence belongs to the true subject of the search and which belongs to the person not being sought by the investigation. This practice effectively presents to the algorithm an intentionally mixed biometric sample.

Selecting the Algorithm

14. The face recognition algorithms used by law enforcement agencies are typically developed by private companies, each with its own team of designers and trained on different datasets. As a result, face recognition systems perform differently depending on the make and model of the algorithm used. Law enforcement agencies who run multiple algorithms simultaneously with each search have reported receiving different results from each algorithm, such as different confidence levels assigned to the matches returned or different matches returned altogether. This is also evidenced in the public testing conducted by the National Institute of Standards and Technology (NIST), which demonstrates that some algorithms perform more accurately than others. This means that the make and model of the algorithm used in a given investigation can directly influence the accuracy of the identification.

15. The same algorithm may also perform at different levels of accuracy depending on the age, race, and gender of the person being searched. Algorithms commercially available to law enforcement agencies may produce less reliable results on subjects with very dark skin, women, and young people, producing higher false non-match rates (missed identifications). The accuracy of many algorithms also declines when there is an age gap of multiple years between the subject in the probe photo and the corresponding database photo, which may lead the algorithm to miss a correct identification and instead return incorrect matches that are closer in age at the time the photographs were taken.

Interpreting the Results of the Algorithm

16. Face recognition systems used by U.S. law enforcement agencies typically produce a list of possible candidates, not just a single match result, to be reviewed by an agent or analyst running the search. These candidate lists vary in length depending on the presets chosen by a given agency, but may contain as many as a few hundred possible candidates. Candidate lists are typically presented in rank order, beginning with the candidate that the algorithm determines is the most likely match. The match candidates may or may not be presented with a corresponding confidence score produced by the algorithm, also depending on a given agency's presets. Confidence scores may be presented as a percentage (e.g., "Match 96.03%"), a whole number out of an unknown total (e.g. "535.000"), or some other metric or notation such as a decimal, a star ranking system, and/or a function of a logarithmic regression model.

17. The confidence score indicates the algorithm's certainty in the match, not the likelihood that the match is or is not correct. For example, a confidence score of 99% accompanying the defendant's photo does not mean there is a 99% chance the defendant is the subject and a 1% chance he or she is not. It merely means the algorithm has a 99% confidence in the similarities between the two photos, given the limitations of the algorithm's design and training, the evidence available for analysis in the probe photo, any information added to the probe photo during the editing process, and the contents of the database the algorithm runs against. Confidence scores presented in a format other than a percentage are likely the function of a logarithmic regression and must be interpreted as such.

18. The face recognition candidate list contains evidence that the algorithm may have determined that someone else looked similar to the subject of the search, or in fact more like the subject than the defendant. For example, documents from one police department indicate that a subject investigated

and ultimately charged was displayed at rank #319, meaning the algorithm produced 318 matches that it determined were more likely to match the subject of the search than the person ultimately charged.[3]

19. Deciding which candidate is a possible match is a decision made by an analyst or law enforcement agent. There is no certification requirement yet for forensic face analysts in the United States; however, the Facial Identification Scientific Working Group (FISWG), composed largely of law enforcement representatives, has developed best practice guidance for training. FISWG states that: "coursework alone is insufficient to establish expertise for facial comparison. In addition to coursework, on-the-job training with a mentor and on-going professional development are necessary to achieve and maintain expertise."[4] The standards body recommends that facial reviewers, responsible for generating investigative leads or law enforcement intelligence, pass proficiency examinations in a variety of skills and techniques encompassing the entire search process and receive a minimum of six months, full-time, on-the-job training under mentor supervision as well as ongoing technical training, research and literature reviews, and further instruction. Facial examiners, responsible for generating conclusions in forensic applications, are recommended to have similar training and a minimum of 12 months of full time, on-the-job, supervised training.

20. These training and supervision requirements are considered best practice to ensure reliable search results, but are not mandatory; thus, candidate list review may be performed by someone who has never received training in morphological comparison or any other technique to accurately and scientifically determine whether the algorithm made a correct identification. The analyst may

[3] NYPD Real Time Crime Center, Facial Identification Section Overview, FOIL Production No. NYPD_02779, on file with author.
[4] Guide for Mentorship of Facial Comparison Trainees in Role Based Facial Comparison, Facial Identification Scientific Working Group (FISWG), May 10, 2019, available at https://fiswg.org/documents html.

additionally have no background in how the algorithm works or how to interpret the confidence scores or other associated information produced by the system, which could lead to a misinterpretation of the results. Studies consistently show that untrained analysts, regardless of experience on the job, perform poorly at unfamiliar face comparison.

21. Cognitive biases that are present in other subjective feature comparison methods like latent fingerprint, tool mark, or fiber analysis likely also exist in face recognition searches, but no nationally recognized controls exist to protect face recognition search results from reflecting those biases and many systems lack controls. Biases in feature comparison methods include, but are not limited to, confirmation, context, circular reasoning, motivation, and illusory superiority. Below are two examples illustrating the potential for bias in a face recognition search.

22. Face recognition systems can only identify people in the database being searched, meaning that many searches may not yield the person being searched for but nonetheless produce a lengthy candidate list with high confidence scores. This may serve to bias the analyst in favor of agreeing with the algorithm and finding a match even when there isn't one.

23. The candidate list may also display or make available the arrest history of each possible match. This may lead to an identity determination that is context-dependent rather than solely based on the similarity of the two photos. An analyst may be inclined to choose the defendant over another more similar-looking candidate based on the similarity between the defendant's prior criminal history and the offense being investigated.

24. Since an analyst is aware a face recognition search has been conducted and is privy to the identity determinations made by the algorithm, the analyst may exhibit confirmation bias in favor of agreeing with the prior determination rather than conducting a thorough, independent, unbiased

review. A recent study of forensic face examination found this bias to be present both with prior human and machine determinations.[5]

Confirming the Face Recognition Results

25. There is a general recognition across most U.S. law enforcement agencies that face recognition is not reliable on its own to produce a positive identification. Most law enforcement agencies consider a face recognition match to be an investigative lead only, meaning that the identification produced by a face recognition search must be confirmed by additional investigation. What constitutes sufficient additional investigation is not defined by most agencies, however, meaning there is a high degree of variability in how much weight is placed on a face recognition match.

26. The way the results of a face recognition search are presented to the investigating officer or to a witness also may suffer from confirmation biases in favor of finding a match, skewing the investigation towards merely certifying, rather than independently corroborating, what the face recognition system proposed as a match. Real-world examples from law enforcement practice that would suffer from bias or otherwise not be a reliable way to independently verify a face recognition search result include:

a. Defendant's photo is presented on its own to a witness rather than in a photo lineup. This suggests to the witness that the defendant is the person being investigated by law enforcement, introducing confirmation bias in favor of finding a match even if the system identified the wrong person.

5 John J. Howard, Laura R. Rabbitt & Yevgeniy B. Sirotin, Human-algorithm teaming in face recognition: How algorithm outcomes cognitively bias human decision-making, PLoS One (Aug. 21, 2020), available at https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7444527/.

b. Defendant's photo is presented to an officer or other individual who was not a witness to the crime, who is nonetheless asked to determine identity. This process adds no additional evidence to the investigation, and may suffer from confirmation, motivation, illusory superiority, or other bias that increases the likelihood of misidentification.

c. Defendant's photo is presented along with associated information about the defendant's prior arrest history, often when the witness is a law enforcement officer. This adds a bias towards finding a match if charges in the defendant's criminal history are similar to the fact pattern being investigated. This means the identification is not made solely on the witness' recollection of what the subject looked like and rather whether he/she thinks the defendant is capable of committing the charged offense.

d. Defendant's photo is presented to the witness along with information indicating the photo was the result of a face recognition search. This may bias the witness towards agreeing with the algorithm and finding a positive identification, which is often perceived as having mathematical certainty.

e. Defendant's photo is presented along with the confidence score generated by the face recognition algorithm. Confidence scores may incorrectly suggest to the witness a probability that the defendant is a match. For example, a 99% confidence score may be interpreted as a 99% chance the defendant is the suspect, and a 1% chance that someone else is the suspect. This is an understandable, but an incorrect, interpretation of the confidence score, which merely indicates the degree to which the faces appeared similar to the algorithm.

Conclusion

27. For the foregoing reasons, it is my expert opinion that all available information pertaining to the use of face recognition by law enforcement during the course of an investigation speaks directly to the reliability of the ultimate identification, and therefore to the guilt or innocence of the defendant. As such, this information must be disclosed during discovery to ensure a defendant receives the due process afforded him or her under the Fifth and Fourteenth Amendments of the U.S. Constitution. Moreover, if it becomes apparent that a search was conducted without adequate training or controls in place to minimize subjective decision-making or cognitive biases, the search may be unduly suggestive or otherwise impermissible as reliable identification evidence, and should be suppressed.

28. As described in the criminal complaint, the face recognition search and human confirmation leading to the identification of Defendant [redacted] lacks clear indicia of reliability, and contains elements that suggest it was not adequately confirmed by additional, independent investigation, or protected from cognitive bias. In this search, while Lieutenant [redacted] ran the face recognition search, criminal complaint affiant Detective [redacted] acted as a forensic face analyst, visually comparing two images to determine whether they represent the same individual. Since Detective [redacted] was aware that a face recognition search took place which identified the Defendant as a possible match candidate, his review was not shielded from confirmation bias towards agreeing with the algorithm's determination rather than conducting an independent biometric review. Beyond the presence of a similar baseball cap in the backpack, the criminal complaint contains little information about independent investigative steps substantiating the identification, such as an eyewitness identification of the Defendant from a photo array, despite the presence of eyewitnesses.

29. No information is given about the quality of the probe image, any photo editing that may have taken place, the respective training or expertise of Lieutenant [redacted] or Detective [redacted] at performing face recognition searches or feature comparisons, the candidate list and Defendant's place within it, the confidence scores, or how Detective [redacted] arrived at his conclusion that there was a biometric match between the Defendant's identification card photo and the probe photo. This information all directly speaks to the reliability, or undue suggestibility, of the face recognition search process and resulting match.

30. No information was provided in the criminal complaint about the face recognition algorithm. The system used, titled the National Capital Region Facial Recognition Investigative Leads System (NCRFRILS), will be shut down no later than July 1, 2021, following the unanimous passage of a law in Virginia[6] that prohibits the deployment of face recognition system by local law enforcement agencies unless expressly authorized by statute.[7]

Clare Garvie

June 25, 2021
Date

[6] VA H.B. 2031 Enrolled, available at https://lis.virginia.gov/cgi-bin/legp604.exe?212+ful+HB2031ER+pdf.
[7] Letter from Chuck Bean, Executive Director of the Metropolitan Washington Council of Governments, to Jeramie Scott, Senior Counsel at the Electronic Privacy Information Center (May 14, 2021), available at https://epic.org/privacy/facerecognition/MWCOG-Letter-Ending-NCRFRILS.pdf.