COMMONWEALTH OF MASSACHUSETTS

SUFFOLK, ss

SUPERIOR COURT DEPT.
[redacted]

COMMONWEALTH

V.

[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 been studying 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 read 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.

2. Based on my research, I have co-authored or authored three reports on the subject, including The Perpetual Line-Up: Unregulated Police Face Recognition in America, Garbage In, Garbage Out: Face Recognition on Flawed Data, and America Under Watch: Face Surveillance in the United States. I have trained over 2,000 attorneys on the use of face recognition in criminal cases, and

recently testified before Congress 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.

4. Law enforcement face recognition searches involve a number of human decision points, each of which introduces the possibility for subjectivity and cognitive bias. Depending on the choices made by an agent at each step, the results of the face recognition system may vary, producing different identification evidence. In order 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. 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.

5. 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. The remainder of the affidavit is organized around these steps.

Selecting the probe photo

6. 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 law enforcement 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 any other imperfection.

7. 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

8. 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. 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

9. 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. What edits are made to a probe

photo are inherently subjective and highly variable, as they rely on a series of unconstrained choices by the agent making the edits and will vary from agent to agent and from search to search. 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 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;

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.

10. 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.

11. 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

12. The face recognition algorithms used by law enforcement agencies are typically developed by private companies, each with their 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.

13. 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 six or more years between the probe photo and the database photo.

Interpreting the results of the algorithm

14. 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. The New York Police Department (NYPD) face recognition system, for example, produces a list of 200 or more possible matches. 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, or a function of a logarithmic regression model.

15. 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.

16. 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.

17. Deciding which candidate is a possible match is a decision made by an analyst or law enforcement agent. There is no set degree or training required in the United States for this role, meaning that the degree of expertise in conducting facial comparisons is highly variable, so the accuracy of this process is highly variable as well. The task 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 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.

18. 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.

19. 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.

Confirming the face recognition results

20. 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.

21. 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. Examples from law enforcement practice include the following, sometimes in combination:

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.

b. 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.

c. 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.

d. 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

22. 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.

District of Columbia ss:
Sworn to before me
this 17th day of July 2019

(NOTARY PUBLIC)

Clare Garvie

JONATHAN ZIMMER NOTARY PUBLIC My Comm. Exp. Aug 14, 2022 DISTRICT OF COLUMBIA