SUPERIOR COURT OF NEW JERSEY
APPELLATE DIVISION
DOCKET NO. A-3078-21T1
INDICTMENT NO. 21-01-0035-I

	:	CRIMINAL ACTION
STATE OF NEW JERSEY,	:	On Leave to Appeal Granted
Plaintiff-Respondent,	:	From an Interlocutory Order of the Superior Court of
v.	:	New Jersey, Law Division, Hudson County.
FRANCISCO ARTEAGA,	:	
Defendant-Appellant.	:	Sat Below:
	:	Hon. Mitzy Galis-Menendez, J.S.C.

BRIEF OF AMICI CURIAE ELECTRONIC PRIVACY INFORMATION CENTER, ELECTRONIC FRONTIER FOUNDATION, AND NATIONAL ASSOCIATION OF CRIMINAL DEFENSE LAWYERS

Christopher J. Frascella
(SBN 410442022)
ELECTRONIC PRIVACY INFORMATION CENTER
1519 New Hampshire Ave NW
Washington, DC 20036
(202) 483-1140
frascella@epic.org

On the brief:
Christopher J. Frascella
Jacob Wiener
Jennifer Lynch
Hannah Zhao
Alan Silber
Clare Garvie

TABLE OF CONTENTS

TABLE OF AUTHORITIES

CASES

OTHER AUTHORITIES

INTEREST OF *AMICI CURIAE*

The Electronic Privacy Information Center (EPIC) is a public interest research center in Washington, D.C., established in 1994 to focus public attention on emerging civil liberties issues and to protect privacy, the First Amendment, and other constitutional values. EPIC routinely participates as amicus curiae in privacy cases throughout the country, including in New Jersey. See, e.g., Brief of Amicus Curiae EPIC Supporting Appellant, Bozzi v. Jersey City, 434 N.J. 326 (2021) (No. 84392) (arguing that disclosure of personal information held in a government record presents a colorable privacy claim that is not outweighed when that record is requested for commercial purposes). Brief of Amicus Curiae EPIC Supporting Appellant, State v. Andrews, 243 N.J. 447 (2020) (arguing that the Fifth Amendment protects privacy interests in cellphone passcodes); Brief for EPIC as Amicus Curiae EPIC Supporting Appellant, State v. Earls, 214 N.J. 564 (2013) (No. 68765) (arguing that individuals have a reasonable expectation of privacy in the current location of their cell phones). EPIC has subject-area expertise in government use of facial recognition technology. EPIC has testified on law enforcement use of facial recognition technology in Congress and state legislatures.

The Electronic Frontier Foundation (EFF) is a member-supported, non-profit civil liberties organization that has worked to protect free speech and privacy rights in the online and digital world for 30 years. With over 30,000 active donors, EFF represents the interests of people impacted by new technologies in court cases and broader policy debates surrounding the application of law in the digital age. EFF has special familiarity with and interest in constitutional issues that arise with new forensic technologies and the use of algorithms in criminal investigations and specifically with facial recognition. See State v. Pickett, 246 A.3d 279 (App. Div. 2021); Lynch v. State, 260 So.3d 1166 260 (Fla. Dist. Ct. App. 2018). EFF also participated in the GAO's inquiry regarding forensic technology, which was prompted by concerns from elected officials about the use of these technologies in criminal proceedings Forensic Technology: Algorithms Used in Federal Law Enforcement, U.S. Government Accountability Office (2020), https://www.gao.gov/products/GAO-20-479SP. And EFF has testified on law enforcement use of facial recognition in both the U.S. Senate and the House of Representatives.

The National Association of Criminal Defense Lawyers (NACDL) is a nonprofit voluntary professional bar association that works on behalf of criminal defense attorneys to ensure justice and due process for those accused

of crime or misconduct. NACDL was founded in 1958. It has a nationwide membership of many thousands of direct members, and up to 40,000 with affiliates. NACDL's members include private criminal defense lawyers, public defenders, military defense counsel, law professors, and judges.

NACDL is the only nationwide professional bar association for public defenders and private criminal defense lawyers. NACDL is dedicated to advancing the proper, efficient, and just administration of justice. NACDL files numerous amicus briefs each year in the U.S. Supreme Court and other federal and state courts, seeking to provide amicus assistance in cases that present issues of broad importance to criminal defendants, criminal defense lawyers, and the criminal justice system as a whole.

NACDL has a particular interest in cases that involve surveillance technologies and programs that pose new challenges to personal privacy. The NACDL Fourth Amendment Center offers training and direct assistance to defense lawyers handling such cases in order to help safeguard privacy rights in the digital age. NACDL has also filed numerous amicus briefs in this Court and the Supreme Court on issues involving digital privacy rights, including: Carpenter v. United States, 138 S. Ct. 2206 (2018); Riley v. California, 134 S. Ct. 2473 (2014); United States v. Jones, 132 S. Ct. 945 (2012).

SUMMARY OF THE ARGUMENT

As a police identification technique, facial recognition searches involve numerous components and steps that each introduce the possibility of misidentification. In many cases, facial recognition searches have been the sole, or primary, means of identification of a suspect and, in some of these cases, have led to the misidentification and wrongful arrest of innocent individuals. Even if a facial recognition search is used as an investigative lead that is corroborated by additional steps, such as a photo array, it will influence all subsequent steps in the identification process and the risk of error will carry forward. In other words, additional investigative steps do not automatically cure facial recognition accuracy problems. The only way to cure this risk of error and protect a defendant's constitutional right to disclosure anchored in Brady v. Maryland, 373 U.S. 83 (1963), is to require the discovery of information about the facial recognition search process and how it may have influenced the identification.

BACKGROUND

I. A facial recognition search is composed of several steps all of which can impact the accuracy of a search.

Facial recognition is a tool for identifying an unknown person in a photograph or video by comparing that image against a database of images whose identities are known.[1] Law enforcement use of facial recognition search as an investigative technique typically relies on algorithms and subjective human judgment to compare facial features and generate identification leads. This use is based on the presumption that as measured by both algorithms and humans, faces are biometrics, unique to each individual and stable across time. A search process will include most or all of the following five steps: 1) probe photo selection; 2) database selection; 3) photo editing; 4) algorithmic search; and 5) human review.

New York Police Department (NYPD) searches are run by analysts in the Facial Identification Section (FIS) of the Real Time Crime Center (RTCC). The facial recognition program runs images against NYPD mugshots collected since 1996, desk appearance tickets, and pistol license applications. The NYPD may also have access to other facial recognition databases, including

[1] This brief focuses on the definition and use of face recognition for identification, also known as 1:many face recognition. Other applications of face recognition include verification, or the comparison of two photos to determine whether they are the same individual, and face analysis, an attempt to label or classify individuals based on facial characteristics such as age, race, sex, emotional state, and more.

through its involvement in the High Intensity Drug Trafficking Area (HIDTA) with New Jersey. See Clare Garvie, Garbage In, Garbage Out: Face Recognition on Flawed Data (2019).[2]

A. Probe Photo Selection

During the first step, an officer or analyst chooses the photo, video still, or other image to run through the face recognition system, called the "probe photo." The characteristics of the chosen image impact the accuracy of the subsequent algorithmic search. A face captured on grainy, low-resolution video taken at night will be much more difficult to match than a high-quality image of someone facing and looking into the camera. NYPD's own documents suggest that it has no minimum quality requirement for probe photos; that this is a case-by-case inquiry.

B. Database and System Selection

The database that is searched is also material to the reliability of the identification. An analyst may have the option to select a county or state mugshot database, a DMV database, or even a database from another state or the federal government. If the database does not contain a photo of the subject in the probe image, the search results will only be incorrect matches.

[2] Information about the NYPD's program comes primarily from public records disclosed to the Center on Privacy & Technology at Georgetown Law. https://www.flawedfacedata.com/.

Detectives in this case originally opted to request a search of the databases accessible to the New Jersey State Regional Operations Intelligence Center (ROIC), which returned no matches, prior to requesting a search of the NYPD's database. Da 40.[3] Whether or not Mr. Arteaga is present in any of the databases searched by ROIC therefore directly speaks both to the accuracy of the search and Mr. Arteaga's culpability.

C. Probe Photo Editing

The analyst may edit the probe photo before running a search, casting further doubt on the assertion that facial recognition searches are based on biometrics unique to the individual. Specifically, NYPD FIS detectives are permitted to make significant edits. The facial recognition program used by the NYPD has options closely resembling Photoshop's editing tools, which the NYPD has used for 1) performing the "removal of facial expression" or "insertion of eyes" which amounts to cutting and pasting a different person's facial features into the probe photo; 2) "creating a virtual probe," or combining face photographs of two different people to identify one of them; 3) using the "blur effect" to add pixels into a low quality image; 4) using the "clone stamp tool" to "create a left cheek and the entire chin area" of a subject who's face wasn't completely visible; and 5) using 3D modeling software to generate

[3] Da – Appendix to Defendant-Appellant's Brief.

missing parts of a face turned away from the camera in a probe photo. Garvie, Garbage In, Garbage Out. These edits contaminate the supposed biometric sample in a way that is undetectable by the matching algorithm.

D. Algorithmic Search

The fourth step involves running an algorithm that compares the probe photo to the images in the databases selected in step two. The algorithms used by law enforcement are typically developed by private companies, each with its own team of designers, and each algorithm is trained using different datasets. In this case, the NYPD's system ran algorithms from two companies: NEC and Rank One Computing. Da 10. These programs are "black box" technology — it is impossible to know exactly how the algorithms reach their conclusions without looking at their source code. But each algorithm will and does produce different results.

Facial recognition algorithms diverge on their conclusions due to how the matching is conducted and how the algorithm is trained. Algorithms create templates, also known as "facial vectors," of the probe photograph and the photographs in the database and compare the templates to find a match, but different algorithms will focus on different points of a face in creating those templates.

Additionally, each algorithm's result is specific to the design and training of that algorithm. Training entails having the algorithm examine pairs of photographs chosen by humans and initially being told which are the same person and which are not. Using this data, the machine then "learns" how to identify similarities and differences in future pairs of photos. Because the training is individual to each algorithm, especially from different companies, the results of two facial recognition algorithms may vary significantly.

After an algorithm runs a search, it returns a "candidate list" of possible matches of photographs ordered according to a confidence score (produced by the algorithm) in descending order. The NYPD's system is programmed to return a list of 200 or more possible matches across multiple pages of results. The confidence scores appear as a three- or four-digit whole number with a three-point decimal (e.g. 586.000). Da 8. Whether these numbers relate to a percentage or are based on another metric such as a logarithmic scale is unclear.

E. Human Review

Although human analysts will review the probe photo and candidate list for matches, numerous studies show that overall facial recognition search accuracy is highly dependent on the training the analyst receives for this task. See, e.g., Alice Towler et al., Evaluating the feature comparison strategy for

forensic face identification; 23 J. of Experimental Psychol.: Applied 1, 47 (2017), see Alice Towler et al., Do professional facial image comparison training courses work? PLoS One 14(d): e0211037 (2019); Clare Garvie, A Forensic Without the Science: Face Recognition in U.S. Criminal Investigations (forthcoming 2022). Yet, there is no nationally recognized competence metric or certification requirement. The NYPD does have a generic facial recognition system user guide, and analysts do receive some form of training, but specifically targeted public record requests have revealed no information about what that training entails. See Garbage In, Garbage Out, supra.

Human review is also impacted by many other factors, including the analyst's innate ability to analyze faces, motivation to find a match, fatigue from performing a repetitive task, time limitations, and cognitive and contextual biases that all humans have. See A Forensic Without the Science, supra. Bias factors within the NYPD system include knowledge of the theory of the case and the ability to review prior criminal arrest histories of matches, both of which may influence the analyst regardless of whether the subjects are similar visually.

At the end of the facial recognition search process, the analyst ultimately decides whether or not one of the candidates identified by the program is a

likely match. In the NYPD, these are sent back to the investigating officer as a "Possible Match Report." Da 6.

II. Current facial recognition accuracy testing does not account for all types and sources of error.

Facial recognition algorithms can produce two types of errors in their outputs: false positives in which a match result is not of the same person as the one in the probe image (a misidentification); and false negatives in which the algorithm fails to output an accurate match that exists in the database (a missed identification). When a false negative error happens, every person the algorithm returns as a candidate has been misidentified. In law enforcement systems there is a trade-off between these types of errors; the lower the likelihood for one type of error, the more likely the other type of error is to occur.

Understanding how likely a facial recognition system is to produce either type of error requires extensive testing. The National Institute of Standards and Technology (NIST) conducts ongoing Face Recognition Vendor Tests (FRVTs) which evaluate algorithm performance in a variety of different conditions. Since participation in these tests is voluntary, not all systems available for purchase have been tested. The longest running tests are based on ID documents and other clear, high-quality frontal images, while newer testing includes mugshots and images from border crossings and ATM kiosks to

introduce a broader range of image variability. However, NIST does not yet regularly test algorithms against the types of photos that police are likely to encounter in investigations, such as surveillance camera images where the subject is blurry, looking away from the camera, in poor light, partially obscured, or edited. As such, the results are useful for comparisons between systems, but don't offer a good picture of how accurate facial recognition algorithms are in real-world police investigations.

Human errors that can occur throughout the process are also not accounted for in NIST's results. Research has shown that humans struggle to identify and distinguish unfamiliar faces, often exhibiting error rates near or worse than random chance when asked to properly identify a person they had seen earlier in a photograph. See Ahmed Megreya et al., Matching Face Images Taken On the Same Day or Months Apart: The Limitations of Photo ID, Applied Cognitive Psychol. 700–706 (2013) (finding 21 percent error rate on unfamiliar face images taken same-day); Vicki Bruce et al., Verification of Face Identities From Images Captured on Video, 5 J. of Experimental Psychol.: Applied 4, 349 (1999). Thus, human review introduces errors that are not accounted for in available testing as well.

ARGUMENT

I. Each facial recognition system presents a unique risk of error that requires robust discovery to be assessed.

The chance that any single facial recognition search produces an error depends on the system used and its specific version, the quality of the probe photograph, photo editing, analyst competence and training, and more. Poor-quality data and systems produce poor-quality results.

Moreover, facial recognition searches do not expose all individuals to the same risk of misidentification. People of color, women, the old, and the young are more likely to be misidentified by some facial recognition systems than adult, lighter-skinned men. The choice of database can amplify the risk of bias as well; many mugshot databases are historically biased and over-include minorities.

The combination of algorithm, photograph, database, and individual creates a unique risk of bias and misidentification for each facial recognition search.

A. Facial recognition systems utilized by law enforcement differ substantially in accuracy.

Facial recognition systems — and the specific versions — used by law enforcement vary greatly in accuracy, as shown in NIST testing. In the latest round of NIST testing, false negative rates ranged from 0.12 to 50 percent of

searches against a mugshot database. Patrick Grother et al., Face Recognition Vendor Test (FRVT) Part 2: Identification, NISTIR 8271 Draft Supplement (2022) at 42.[4] An algorithm with a 50 percent false negative error rate will return all wrong results in an astounding half of its searches. Of the algorithms tested, the worst were roughly 1,000 times more likely to make a mistake than the most accurate algorithms, leading NIST to urge a "buyer beware maxim" attitude. Ibid.[5]

Different versions of a company's algorithm can also yield different results. Ibid. For example, Rank One Computing, one of the algorithms that the NYPD uses, has submitted 14 different versions of its facial recognition algorithm for NIST testing. When searching against the same database, an older version of the algorithm was roughly 200 times more likely to make a mistake than the newest version. See NIST FVRT 1:N Identification: Identification Performance (2022).[6]

B. Probe photo quality directly impacts the accuracy of a search, and additional errors depend on the algorithm that was used.

It is no surprise that using a lower-quality photograph will produce more inaccurate results. But using photographs taken from different angles can also

[4] https://pages.nist.gov/frvt/reports/1N/frvt_1N_report.pdf.
[5] Version 3 from 2018 missed a correct identification 14.3 percent of the time, while version 13 from 2022 missed a correct identification in 0.7 percent of searches.
[6] https://pages.nist.gov/frvt/html/frvt1N.html.

substantially decrease accuracy. In police investigations, photographs of the perpetrator are of highly variable quality, pose, and lighting. The less information these photos contain — for example, if the image is pixelated or blurry or the individual's face is turned away from the camera — the less reliable the resulting search results will be.

While most algorithms perform relatively well when given a front-facing mugshot as the probe image, the false negative error rate for even the best algorithms jumps to more than 20 percent when applied to images where the sides of a person's face are cropped or the face is tilted down, like a still image from an ATM kiosk or a video camera mounted high in a store. NISTIR 8271 Draft Supplement, supra at 6. Most algorithms NIST tested were 6 to 10 times worse at identifying faces in a lower-quality webcam photo as opposed to a front-facing mugshot. Ibid. at 55-61. Algorithms were totally flummoxed by the side-view images, with most missing a match in 30 to 90 percent of searches. Ibid.

Probe photos in police investigations are often low quality, poorly lit, and non-frontal, e.g., stills from high-mounted security cameras taken at night. Moreover, police at times run images unrelated to the investigation without good justification. The NYPD Facial Identification Section, after running a search on a probe photo that was of too poor quality to return any usable

results, ran a photo of Woody Harrelson because an analyst thought that the suspect looked similar to the actor. Garbage In, Garbage Out (2019), supra. This illustrates the ability — and willingness — of agencies to submit garbage data into their systems, including information like someone else's face, which should necessarily fail to return any reliable results.

Because each algorithm performs differently under different conditions, defendants need detailed discovery on how the search was performed to establish the likelihood that they were misidentified by a facial recognition system.

C. The choice of algorithm substantially impacts the risk of a racially-biased wrongful identification.

Groundbreaking research in 2018 on facial classification algorithms — a set of systems closely related to facial recognition algorithms — showed that people of color, and in particular Black women, were far more likely to be incorrectly classified; the algorithms couldn't identify their gender. Joy Buolamwini & Timnit Gebru, Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification, 81 Proc. Machine Learning R. 1-15 (2018).[7] In 2019, NIST extended those findings to facial recognition algorithms, finding "empirical evidence for the existence of demographic

[7] https://proceedings.mlr.press/v81/buolamwini18a/buolamwini18a.pdf.

differentials in the majority of contemporary face recognition algorithms" it evaluated. Patrick Grother et al., FRVT Part 3: Demographic Effects, NISTIR 8280 (2019) at 6.[8] A number of studies concur with, and build on, NIST's findings. See Cynthia Cook, et al., Demographic effects in facial recognition and their dependence on image acquisition: An evaluation of eleven commercial systems, IEEE Transactions on Biometrics, Behavior, and Identity Sci. (2019); K. S. Krishnapriya et al., Characterizing the variability in face recognition accuracy relative to race, CoRR (2019).

Using country of origin as a proxy for race, NIST's report found that while not all algorithms performed the same, false positive rates (misidentifications) were generally higher for subjects from West and East Africa and East Asia than for Eastern European subjects; for women; and for the elderly and children. NISTIR 8280, supra. False negatives (missed identifications) were higher for Asian and American Indian individuals than for white or African American faces. Ibid. An individual's precise demographics including race, age, and gender thus may have a direct impact on the likelihood that individual will be misidentified.

Disparate impacts across race and gender remains even as the technology advances, and new sources of bias are regularly discovered. For example, a

[8] https://nvlpubs.nist.gov/nistpubs/ir/2019/NIST.IR.8280.pdf.

study from August 2022 found that facial recognition systems are more prone to misidentify Black faces when the probe image is converted to a common but lower-quality file format. Seyma Yucer et al., Does lossy image compression affect racial bias within face recognition?, Durham Univ. (2022).[9] One of the key takeaways from the recent history of facial recognition research is that bias will go unnoticed unless it is specifically looked for. Preserving the details of a facial recognition search and making them available to defendants is one way to mitigate the risk of currently unknown biases in these systems.

D. Database selection also produces disparate racial impact.

Even if the algorithm in a facial recognition system does not have differential accuracy rates, the database used can still inject bias and amplify the risk of a wrongful identification for minorities.

Facial recognition systems that search mugshot databases will recreate the biases in those databases. These databases disproportionately contain people of color, reflecting the history of over-policing poor and minority communities. In New Jersey, for example, Black people are 12 times more likely than whites to be incarcerated — the highest disparity of any state. Ashley Nellis, The Color of Justice (2021).[10] Given this, using facial

[9] https://www.researchgate.net/publication/362728786_Does_lossy_image_compression_affect_racial_bias_within_face_recognition.

[10] https://www.sentencingproject.org/wp-content/uploads/2016/06/The-Color-of-Justice-Racial-and-Ethnic-Disparity-in-State-Prisons.pdf.

recognition on mugshot databases will over-expose minorities to the risk of wrongful identification. And when a biased algorithm is used on a historically skewed database, the risk of error increases even more dramatically.

Because demographic risks compound the baseline risks of misidentification from inaccurate systems and poor-quality probe images, discovery should allow defendants to meaningfully assess the risks unique to their demographics in relation to the algorithm and database used.

II. Human review cannot cure algorithmic errors.

Any assumption that a "human in the loop" will correct and compensate for errors by a facial recognition algorithm is erroneous. Humans are prone to misidentifying unfamiliar faces and are subject to the same biases present in facial recognition systems.

A. Humans are prone to face identification errors.

People are substantially worse at correctly identifying or distinguishing between strangers' faces than faces of those they know. Vicki Bruce, Zoë Henderson, Craig Newman, and A. Mike Burton, Matching Identities of Familiar and Unfamiliar Faces Caught on CCTV Images, 7 J. of Experimental Psychol.: Applied 3, 207 (2001). These errors are magnified when variations like image quality, pose, age between photographs, or similar-looking imposters are introduced, just like facial recognition systems. See Garvie, A

Forensic Without the Science, supra; Vicki Bruce, et al., Verification of Face Identities From Images Captured on Video, 5 J. Experimental Psychol. Applied 4, 349 (1999). In one study testing individuals' ability to identify subjects in low-quality surveillance images, participants made correct identifications at a rate only marginally better than chance. See Forensic Without the Science, supra.

Experience in performing identifications does not improve performance. Separate studies on law enforcement agents and on passport officers found that individuals with years on the job performed just as poorly as non-professional participants. Ibid. And while training may help performance, forensic facial identification training remains inconsistent and not systematic performance evaluation scheme exists. Rather than providing a valuable check against errors, therefore, the human in the loop introduces additional risk of error into the facial recognition identification process.

B. Humans and algorithms make similar types of errors.

Facial recognition algorithms tend to choose possible matches from within the same demographic groups. See John J. Howard et al., Quantifying the Extent to Which Race and Gender Features Determine Identity in Commercial Face Recognition Algorithms (2021); see NISTIR 8280, supra.[11]

[11] https://www.dhs.gov/sites/default/files/publications/21_0922_st_quantifying-commercial-face-recognition-gender-and-race_updated.pdf; https://doi.org/10.6028/NIST.IR.8280.

If the probe photograph is of a Hispanic-looking male in his mid-thirties, for example, the resulting candidate list photographs will likely also be primarily Hispanic males in their mid-thirties. This may be because algorithms rely in part on demographic facial markers to make identifications rather than markers that only correspond to identity.

Humans make the same types of errors. We are more likely to mistake two people of the same ethnic origin, race, sex, and age than people that are of different demographics.[12] Thus, if the algorithm produces a misidentification, the analyst is likely to agree with that misidentification and pass that misidentification forward to the investigating officer.

C. Humans tend to agree with prior decisions rather than independently evaluate them.

Humans are prone to confirmation bias, that is, focusing on or interpreting new information in a manner consistent with existing expectations or beliefs. Consequently, analysts will be biased towards agreeing with an algorithm's conclusion rather than independently reviewing biometric similarities and differences between faces. See Garvie, A Forensic Without the Science, supra. A 2020 study of facial recognition systems sponsored by the Department of Homeland Security Science & Technology Directorate

[12] This is especially true if the analyst running the search is of a different demographic than the search subject, a phenomenon of misidentification referred to as the "cross-race bias" effect.

demonstrated this and cautioned that the human in the loop may be biased towards agreeing with an algorithm's false positive determination. John J. Howard et al., Human-algorithm teaming in face recognition: How algorithm outcomes cognitively bias human decision-making, PLoS One 15(8): e0237855 (2020).

Coupled with the fact that humans and machines make similar types of errors, these errors risk being compounded, and passed forward to later stages in the identification process, rather than mitigated.

D. Later investigative steps regarding identification, such as photo arrays, cannot cure algorithmic errors.

When the investigative follow-up is not sufficiently blinded, or protected from the influence and errors of the facial recognition search, misidentifications by the search algorithm will not be cured. A failure to sufficiently scrutinize results can happen if the officer performing the investigative follow-up is aware a facial recognition search took place, which is almost always the case. Knowing about the search may also impact eyewitnesses in the same way especially if they believe algorithms are inherently reliable.

The investigating officer receives a single printout with two images on it: the probe image and the "possible match." See, e.g., Da 6. Because these two images are likely to be of people with matching demographics, the officer

is likely to agree with, rather than independently evaluate, the possible match — particularly because the officer is aware a facial recognition search took place. That officer is then likely to focus the investigation on that person.

The defendant now becomes part of the photo array as a result of both the algorithmic search and the subsequent analyst's and officer's agreement with that result, regardless of its accuracy.

Given the above, any error the algorithm makes is likely to carry forward and affect subsequent human determinations of identity, despite human review.

III. Facial recognition searches routinely determine the course of investigation, and errors have resulted in numerous wrongful arrests.

As noted in Part I.E supra, the result of facial recognition searches will guide the course of a case, especially where there is a dearth of other evidentiary leads. Errors from these facial recognition searches have already resulted in numerous wrongful arrests, including in New Jersey. And the risk is especially high as matches are sometimes treated as a de-facto identification of the perpetrator even when officers are instructed that leads are investigative in nature only. See Garvie, A Forensic Without the Science, supra. Even where there are additional investigative steps, they often fail to cure the wrongful

identification. These errors derail human lives and waste law enforcement and judicial resources.

Michael Oliver was wrongfully arrested by Detroit police after a facial recognition search in circumstances similar to the identification of Mr. Areaga. Police investigating the alleged grabbing and smashing of a smartphone ran an image of the suspect through a facial recognition system that returned Mr. Oliver as a match. See Tate Ryan-Mosley, The new lawsuit that shows facial recognition is officially a civil rights issue, MIT Tech. Rev. (2021).[13] Officers then presented an eyewitness with a photo array containing Mr. Oliver's photo with five fillers, and the eyewitness confirmed, rather than corrected for, the mistake. Mr. Oliver was arrested and had charges pending for 4 months even though Mr. Oliver has face and arm tattoos while the suspect in the photo does not. See Natalie O'Neill, Faulty Facial Recognition Led to His Arrest—Now He's Suing, Vice (2020).[14] Oliver's case shows that an eyewitness identification from a photo array is not sufficient to correct for facial recognition errors and prevent the wrong man from being arrested.

Detroit police wrongfully arrested another innocent man, Robert Williams, after facial recognition erroneously identified him as the subject.

[13] https://www.technologyreview.com/2021/04/14/1022676/robert-williams-facial-recognition-lawsuit-aclu-detroit-police/.
[14] https://www.vice.com/en/article/bv8k8a/faulty-facial-recognition-led-to-his-arrestnow-hes-suing.

The only corroboration sought was a statement from a non-witness to the crime who nevertheless confirmed the identification based on her review of security footage and lineup photos. See Kashmir Hill, Wrongfully Accused by an Algorithm, New York Times (2020).[15] A photo array did not cure the facial recognition system's error in Mr. Williams case either, which led to his wrongful arrest in front of his wife and young daughters.

New Jersey is no stranger to facial recognition misidentifications. Police arrested Nijeer Parks for a crime he did not commit after a facial recognition system incorrectly flagged him as the perpetrator. Mr. Parks was identified from a photo on a fake ID left at the scene by a suspect in a shoplifting case. Wrongfully Accused by an Algorithm, supra. A detective compared Mr. Parks' photo with the fake ID and believed that they were a match. Ibid; Affidavit of Probable Cause, State of New Jersey v. Nijeer Parks, Police Case No. 19010123 (Woodbridge Mun. Ct 2019); Complaint and Demand for Trial by Jury, Nijeer Parks v. John E. McCormack et al, No. 2:21-CV-03021 (Sup. Ct. NJ 2020).[16] Mr. Parks was jailed for 11 days and experienced harsh treatment in custody; his charges were pending for nearly a year before they were dropped. Mr. Parks seriously considered taking a plea deal despite knowing he

[15] https://www.nytimes.com/2020/06/24/technology/facial-recognition-arrest.html.
[16] https://int.nyt.com/data/documenttools/new-jersey-facial-recognition-lawsuit-nijeer-parks-v/38ff3e74088a95a9/full.pdf.

was innocent. Khari Johnson, How Wrongful Arrests Based on AI Derailed 3 Men's Lives, Wired (2022).[17]

Parks, Williams, and Oliver are all Black men. In the thousands of cases over the past 20 years where facial recognition was used, and cases where innocent defendants took plea deals, many more facial recognition errors may exist that have not been discovered and remain un-redressed.

Discovery provides a last chance to identify and mitigate the worst harms of misidentifications by facial recognition systems. Such misidentifications are documented, shown to cause harm, and likely more common than is currently known.

IV. Discovery is necessary in this case to allow the defendant to understand the evidence against him.

The sections above demonstrate why Mr. Arteaga must have access to discovery of the facial recognition process used to identify him. Any identification procedure based on a facial recognition search process contains substantial risk of error, which is potentially exculpatory. Because the identification of the defendant as the suspect is material to his guilt or innocence, information about this identification should be considered Brady

[17] https://www.wired.com/story/wrongful-arrests-ai-derailed-3-mens-lives/.

material. See Supp. Brief of Defendant-Appellant at 32–45, New Jersey v. Arteaga, No. A-003078-21 (2022).

A. There is significant potential for error introduced at numerous stages in the facial recognition search process in this case, which may be cured by discovery.

First, the probe photographs used by the NYPD are of poor quality, with the potentially identifying features of the subject's face obscured from view due to the low quality of the image and the fact that the subject's face is turned away from the camera. Da 6-37. This makes the face comparison process for both the algorithm and the "human in the loop" harder and raises the risk of misidentification. In addition, analysts ran multiple different probe photos, which returned different results and suggest variability in the accuracy of each search. Ibid.

Second, the officers involved in this case had the opportunity to run the probe photograph(s) against numerous databases. But no information was provided to the defendant regarding the specific database searched. The fact that searches were run against multiple databases with different results — the New Jersey search returning no results — is material. (Da 62). If Mr. Arteaga was in the database used by New Jersey, the failure to return matches suggests that Mr. Arteaga is not the subject of the probe photo (and thereby innocent) or that the quality of image was of too low quality to ever obtain reliable results.

Information regarding the database used is necessary to properly understand what occurred in this case.

Third, the analyst running the facial recognition search could have edited the photo prior to running the algorithm as per NYPD internal training and guidance documents. See Garvie, Garbage In, Garbage Out, supra. Editing increases the chance of error and may amount to contaminating a biometric sample. It is impossible to know how the edits affected the search without knowing what edits were made with what tool and the source code of the search algorithm used.

Fourth, multiple searches were run, each returning different results. Da 8-26. Related, each search was run on two different algorithms which each produced different results. The resulting "possible match" photo is only present in some, not all, of the many candidate lists generated. This suggests that the NYPD analyst decided that some searches were more accurate than others, but the reasons were not disclosed. The analyst's decision to disregard some results of algorithms that the NYPD presumably trusts is information the defendant is entitled to discover.

Fifth, the analyst's process to narrow down these images to a single "possible match" is of critical importance in determining the reliability of the identification and arrest. This analyst had to narrow down the results of the

various algorithmic searches from hundreds if not thousands of possible matches to a single photograph. There is immense room for error and cognitive bias. Supra Parts III and IV. Yet none of the training, expertise, nots about thought processes and decisions, or other tools used to make this critical determination have been made available for review.

Sixth, the NYPD returned the results to the New Jersey authorities as a search result report with the following words: "Possible match. This is not a positive identification and is not probable cause to arrest, merely a lead. Further investigation is needed to develop probable cause to arrest." The document does not instruct the investigating officer what additional steps to take, or what additional information is needed, before the probable cause threshold is met. This introduces wide variability in the degree to which the face recognition search is relied upon in order to make an arrest. In addition, this process also guarantees that the investigating officer is aware that a facial recognition search took place. This knowledge will impact what additional steps are taken and how thoroughly the investigative lead is verified.

Disclosure of information about the facial recognition search process will substantially reduce the risk of misidentification. Since risk of error is present at all stages of the search, it is important for all aspects of the search process — the photographs used, all searches run, information about the

algorithms' reliability and the analyst's competence, to name a few — to be made available to the defendant. Protective orders enable this information to be shared in a way that does not pose a risk to private entities that may be implicated in a discovery request.

B. Information about a facial recognition search should be considered Brady material regardless of the State's intent to introduce it.

For all the reasons face recognition 1) risks creating mistaken identity and 2) impacts the resulting identification that is introduced in court, information about a facial recognition search process must be disclosed in order to comply with the requirements of Brady. Supp. Brief p. 24–45. The state has the responsibility to disclose material information that tends to exculpate the defendant and/or undermine the credibility of its witnesses. Brady v. Maryland, 37 U.S. 83 (1963); Kyles v. Whitley, 514 U.S. 419 (1995). Information is material if it tends to undermine confidence in the result of the criminal case. United States v. Bagley, 473 U.S. 667 (1985).

Information about a facial recognition search may negate guilt by suggesting someone other than the defendant, such as someone else in the candidate list, committed the offense in question. Facial recognition additionally produces information that may negate guilt or undermine the confidence in the result of a case as it suggests the State's reliance on an

investigative process that has not been thoroughly tested and determined to be reliable.

Information about a facial recognition search additionally may be Brady material because the system acts as an impeachable witness. The algorithm performs the task of selecting what it calculates to be the most likely matches, out of a much larger pool of individuals and ranks these matches. The analyst then also performs a task not unlike that of an eyewitness by reviewing the candidate list, likely of similar-looking individuals, and selecting the most likely match. The analyst's competence, the suggestiveness of the candidate list, or other potentially biasing factors impact the reliability of this identification as it would that of an eyewitness reviewing a photo array.

Although neither the algorithm nor the analyst was an eyewitness to the crime, the system components may still be considered impeachable under the expert witness theory — the algorithm and analyst perform a biometric forensic search process. Like any other forensic expert, however, the analyst and algorithm must still be made available to the defense for review for impeachment purposes through cross-examination to ensure that the search process was, in fact, scientifically sound.

Disclosure of information about the facial recognition search process will not just substantially reduce the risk of misidentification but will additionally protect the defendant's constitutional right to due process.

CONCLUSION

Amici respectfully request that the court reverse the Superior Court's ruling and find that Defendant is entitled to discovery on the details of how he was identified using a facial recognition system because the likelihood of a misidentification is a fact-specific determination that can only be made with discovery.

Dated: 09/26/2022

Christopher J. Frascella
(SBN 410442022)

ELECTRONIC PRIVACY
INFORMATION CENTER
1519 New Hampshire Ave NW
Washington, DC 20036
(202) 483-1140
frascella@epic.org

Jacob Wiener (*PHV Pending)
ELECTRONIC PRIVACY
INFORMATION CENTER
1519 New Hampshire Ave NW
Washington, DC 20036
(202) 483-1140

Jennifer Lynch (*PHV Pending)
Hannah Zhao (*PHV Pending)
ELECTRONIC FRONTIER FOUNDATION
815 Eddy St.
San Francisco, CA 94109
jlynch@eff.org
zhao@eff.org
(415) 436-8333

Alan Silber
(SBN 208431965)
PASHMAN STEIN WALDER HAYDEN
Counsel for *Amicus Curiae,* THE NATIONAL ASSOCIATION OF CRIMINAL DEFENSE LAWYERS
Court Plaza South
21 Main Street, Suite 200
Hackensack, NJ 07602
asilber@pashmanstein.com
(973) 610-8405

Clare Garvie (*PHV Pending)
NATIONAL ASSOCIATION OF CRIMINAL DEFENSE LAWYERS
1660 L St NW #12
Washington, DC 20036
cgarvie@nacdl.org
(202) 465-7657