A JOINT INFORMATIONAL HEARING OF THE ASSEMBLY COMMITTEE ON PRIVACY & CONSUMER PROTECTION AND THE ASSEMBLY SELECT COMMITTEE ON EMERGING TECHNOLOGIES & INNOVATION

SHAPING THE FUTURE OF FACIAL RECOGNITION TECHNOLOGY IN CALIFORNIA: IDENTIFYING ITS PROMISES AND CHALLENGES

March 10, 2020
3:00 p.m.
State Capitol, Room 127

BACKGROUND PAPER

I. Background

Facial recognition technology (FRT) refers to the use of automated devices to identify or verify a person from a digital image by determining whether two images of faces represent the same person. FRT consists of two component processes: face detection, or locating a face within a photo, and face identification, or the matching of facial information to an image or images in a specified database that link to identifying information. FRT relies on the use of biometrics, the statistical analysis of measurements of biological data, in order to compare these images, reducing complex images to numerical values that represent key facial measurements that distinguish individuals.

While the genesis of modern, automated FRT stemmed from the CIA-funded work of Woodrow Bledsoe in the 1960s, the concept of biometric identification using facial measurements traces back to the late 1800s and a technique referred to as "bertillonage."[1] Invented by French police officer Alphonse Bertillon, bertillonage involved measurement of 11 parts of the body of a suspect, criminal, or immigrant, including the length of the right ear and the length and breadth of the head. Pairing these measurements with photographs (e.g. "mug shots") would create a unique identifier that could be stored, retrieved, and cross-referenced in the event an apprehended individual had changed their appearance[2]. Ultimately bertillonage was supplanted by the more efficient and reliable process of fingerprinting, but it nonetheless generated controversy even then over its implications for privacy and the presumption of innocence, spurring several lawsuits attempting to compel destruction of these records for suspects who were detained but not convicted or whose convictions were expunged[3].

[1] Shaun Raviv, "The Secret History of Facial Recognition," *WIRED*, Jan. 21, 2020, https://www.wired.com/story/secret-history-facial-recognitions,

[2] José Regas, "What's in a Face ID?", *Slate,* Mar. 5, 2018, https://slate.com/technology/2018/03/with-apples-face-id-its-time-to-look-at-facial-recognition-techs-problematic-past.html.

[3] Matthew Guariglia, "Facial Recognition Technology Is the New Rogues' Gallery," *Slate*, Feb. 17, 2020, https://slate.com/technology/2020/02/rogues-gallery-facial-recognition-technology-history.html.

Early digital FRT was conceptually similar to bertillonage, relying on operators to manually identify and mark certain landmarks on the subject's face using a digital input tablet and a stylus. The distances between landmarks were then automatically measured, and compared to, the same metrics for other images[4]. Limitations on computational power at the time made comparisons with large databases virtually impossible – the database used by Bledsoe for this research consisted of only ten images – and limited accuracy based on the reliability of these measurements and the number of measurements that could be efficiently compared[5]. However, recent revolutions in computer science, namely in the fields of artificial intelligence and machine learning, have dramatically increased computational power, and have resulted in seismic improvements in the efficiency, accuracy, and scale with which FRT can be implemented[6]. The ability to train algorithms to identify and assign numerical values to facial variables that are far more complex and unique than those self-evident to a human observer (i.e. so-called "principal components") has simultaneously reduced the number of variables that must be analyzed to accurately identify a face and has increased the tolerance of FRT to poor-quality images[7,8]. According to a 2018 report from the National Institute of Standards and Technology, the most accurate commercial facial recognition algorithms in 2018 produced twenty times fewer errors than the most accurate algorithms tested just five years prior. This likely resulted from the widespread application of so-called "deep convolutional neural networks," a type of machine learning based loosely on the learning mechanisms of animal brains, and represents a nearly 300-fold increase in the accuracy of FRT since 1993[9]. These algorithms have improved the speed and ease with which facial information is collected, analyzed, and compared, and have reduced the sensitivity of FRT to confounding variables, such as differences in lighting, angle, age, and expression[10]. As a result, not only can modern FRT be applied on an unprecedented scale - the FBI face recognition unit's database consists of over 400 million photos of up to 125 million Americans - but it can also be used to identify individuals in real-time from surveillance video feeds[11].

These technological advances have revealed new applications for FRT in a variety of sectors and circumstances. Still, the ability to identify individuals on a large scale, potentially in real time, has significant implications for our fundamental rights to privacy and free expression. Extensive research has also indicated disparities in performance of FRT depending on characteristics of the subjects, with generally poorer performance when identifying people of

[4] Raviv, "The Secret History of Facial Recognition," *supra*.

[5] *Ibid.*

[6] *National Institute of Standards and Technology,* "Ongoing Face Recognition Vendor Test (FRVT) Part 2: Identification," P. Grother, M. Ngan, K. Hanaoka, *U.S. Department of Commerce,* Nov. 26, 2018, http://doi.org/10.6028/NIST.IR.8238.

[7] Turk & Pentland, "Eigenfaces for recognition," *Journal of Cognitive Neuroscience,* 3(1): 71-86.

[8] Navarrete & Ruiz-Del-Solar, "Analysis and Comparison of Eigenspace-Based Face Recognition Approaches," *International Journal of Pattern Recognition and Artificial Intelligence*, 16(7): 817-830.

[9] *NIST*, "Ongoing Face Recognition Vendor Test (FRVT) Part 2," *supra.*

[10] *Ibid.*

[11] *Center on Privacy & Technology*, "The Perpetual Line-Up: Unregulated Police Face Recognition in America," *Georgetown Law*, Oct. 18.2016, https://www.perpetuallineup.org.

color and women, as well as entrenchment of existing cultural biases based on the particulars of training data and algorithmic designs[12]. Though FRT retains ample promise as an emerging technology, it is imperative that California adopts a regulatory approach that is mindful of these shortcomings and prioritizes the maintenance of civil liberties guaranteed by the State and Federal Constitutions governing its residents. The purpose of this hearing is to foster discourse on the technical and practical complexities of FRT, and its impacts on civil society, as the Legislature considers approaches to regulating the use of this technology.

II. **Applications of Facial Recognition Technology**

FRT has several applications, both current and prospective, that offer utility for technological efficiency, safety, and security. Most frequently, discussions about the use of FRT take place in the context of law enforcement, where it can be used for surveillance and identification of perpetrators or suspects of criminal activity. Responses to public records requests by Georgetown Law's Center for Privacy & Technology suggest that at least 52 state and local law enforcement agencies surveyed are now using, or have previously used or obtained, FRT, indicating extensive adoption that precedes any substantive regulation of the manner in which it is used[13]. Specifically, use of FRT generally falls into one of two categories: face verification, i.e. the confirmation of one's claimed identity, or face identification, i.e. the determination of the identity corresponding to an unknown face.

Face verification, which performs a 1:1 match of an image of confirmed identity with a real-time facial scan, can be used, e.g., to identify fraudulent use of government-issued identifying documents, to authenticate identity for access to sensitive records online, or to confirm the identity of an individual boarding an airplane. Face identification, which is typically subject to greater public outcry, performs a so-called 1:N search that compares an image with an entire database of images of confirmed identity. Face identification in the law enforcement context has various applications. For example, it can assist in identifying an individual who either refuses or is unable to identify themselves, in determining whether an apprehended individual matches photos from unsolved crimes or has outstanding warrants, or in obtaining a list of candidates for further investigation based on a photo or video still of a suspect from a security camera, smartphone, or social media post. Most controversially, the use of face identification for real-time video surveillance allows law enforcement to extract faces from live video feeds to identify, passively, any individual within a given area in order to determine the locations of missing persons or suspects of interest. Currently, real-time face recognition is computationally expensive, but the rate of advancement of this technology suggests that it could become more pervasive in the coming years[14].

[12] *National Institute of Standards and Technology,* "Face Recognition Vendor Test (FRVT) Part 3: Demographic Effects," P. Grother, M. Ngan, & K. Hanaoka, *U.S. Department of Commerce,* Dec. 2019, https://doi.org/10.6028/NIST.IR.8280

[13] *Center on Privacy & Technology,* "The Perpetual Line-Up," *supra.*

[14] *Ibid.*

The applications of FRT extend beyond the law enforcement context, as well. For example, FRT can limit access to secure facilities in a manner similar to fingerprint scanners or iris scanners. Since FRT is passive, it does not require the authorized individual to engage directly with the scanner, boosting efficiency. Such technology is already in widespread use for authorizing access to smartphones, as Apple has provided for the use of highly sophisticated FRT to unlock their devices in the past several iterations of their operating system[15]. Beyond security uses, social media outlets, such as Facebook, have used FRT to alert users when photos of them are uploaded, whether or not the user is tagged in those images, and to suggest user tags for images[16]. In advertising, FRT can allow for recognition of an individual's age and gender to target advertisements. Tesco, the multinational grocery and merchandise retailer, has already announced plans to install screens at gas stations with built-in FRT for this purpose[17]. A research report conducted by *Component* assessing the market for FRT predicts that the commercial facial recognition industry in the United States alone will be worth over $7 billion by 2024[18].

One could envision several additional applications for this technology, including using one's facial information to verify purchases without the need for a credit or debit card, identifying who is viewing a streaming service to target suggested content, performing an instantaneous criminal background check, or confirming that one is of age to purchase alcohol or tobacco without the need for displaying identification. Despite the practical convenience of these applications, however, the use of FRT in any of these circumstances requires careful review, because it presents significant risks to individual privacy and free expression, and could exacerbate discrimination.

III. Privacy, Speech, and Bias Considerations

Privacy

While the United States Constitution does not explicitly guarantee a right to privacy, the Supreme Court has consistently ruled that the Constitution includes an implicit right to privacy granted by the First, Third, Fourth, and Fifth Amendments.[19] Emphasizing the fundamental importance of a right to privacy, the California Constitution made this right

[15] See, e.g., Yoni Heisler, "Infrared video shows off the iPhone X's new Face ID feature in action," *BGR,* Nov. 3, 2017, https://bgr.com/2017/11/03/iphone-x-face-id-video-infrared; Yoko Kubota, "Apple iPhone X Production Woe Sparked by Juliet and Her Romeo," *The Wall Street Journal*, Sep. 27, 2017, https://www.wsj.com/articles/apple-iphone-x-production-woe-sparked-by-juliet-and-her-romeo-1506510189.

[16] Tom Simonite, "Facebook Creates Software That Matches Faces Almost as Well as You Do," *MIT Technology Review*, Mar. 17, 2014, https://www.technologyreview.com/s/525586/facebook-creates-software-that-matches-faces-almost-as-well-as-you-do.

[17] Dean Nicolls, "What is Facial Recognition...and How Does It Differ from Facial Authentication?", *Jumio,* Jul. 9, 2019, https://www.jumio.com/facial-recognition-vs-facial-authentication.

[18] *Component,* "Facial Recognition Market," Jun. 2019, https://www.marketsandmarkets.com/Market-Reports/facial-recognition-market-995.html.

[19] See, e.g., *Griswold v. Connecticut* (1965), 381 U.S. 479; *Roe v. Wade* (1973), 410 U.S. 113; *Lawrence v. Texas* (2003), 539 U.S. 558.

explicit and inalienable under Article 1, Section 1, which states, "All people are by nature free and independent and have inalienable rights. Among these are enjoying and defending life and liberty, acquiring, possessing, and protecting property, and pursuing and obtaining safety, happiness, and privacy."[20] In addition, Sections 24 and 28 of the same Article reassert this right to privacy as it applies specifically to defendants and victims of criminal conduct, respectively, by preventing its universal suspension due to suspicion of criminal behavior[21]. Because FRT is capable of passively identifying virtually any individual in a recording or real-time feed, and without that individual's knowledge, the expanding adoption of FRT in both the public and private sectors has come under scrutiny for its implications for individual privacy.

FRT poses particularly serious threats to privacy because it is, at present, minimally regulated. In response to concerns relating to the lack of regulation of FRT, the California Legislature passed AB 1215 (Chap. 579, Stats. 2019), which placed a three year moratorium on the use of any biometric surveillance system, including FRT, in connection with police-worn body cameras. Apart from this measure, however, the use of FRT is generally unregulated, and it is unclear how existing privacy protections apply to the use of FRT. While law enforcement collection of biometric information (e.g. mouth swabs, fingerprinting, etc.) typically constitutes a "search" subject to certain protections under the Fourth Amendment, FRT, which can collect biometric information passively, does not require the physical seizure of that biometric information, and is thus categorically unique[22]. Federal district courts in California have held that an individual's reasonable expectation of privacy extends to records of their movements revealed by cell-site location information, and that a warrant must be approved for obtaining this information, indicating that a physical search is not necessary for the Fourth Amendment to apply[23]. However, no state or federal court has yet ruled on the application of Fourth Amendment protections to the use of FRT.

Other applications of FRT include tracking the behavior of an individual over time by identifying and compiling when an individual is in front of a given recording device/feed, or identifying individuals in public without their knowledge. These capacities raise concerns that adoption of FRT could spell the end of public anonymity. One can imagine the privacy and security threats in a world in which each person wears glasses with FRT that can identify any individual they encounter; a malicious person inadvertently bumped while passing another on the street could immediately identify their name, address, and any other publically available information to harass or harm that person. Already, a New York Times exposé revealed that a company specializing in FRT, *Clearview AI,* has aggregated over three billion images scraped from publically accessible media, including Facebook, YouTube, and Venmo, to create a database of online identities matched with images of those individuals

[20] Cal. Const. Art. 1, Sec. 1.

[21] Cal. Const. Art. 1. Sec. 24, 28.

[22] *Center on Privacy & Technology,* "The Perpetual Line-Up," *supra.*

[23] *Carpenter v. United States* (2018), 585 U.S. ___.

that can be used for facial recognition[24]. *Clearview AI* has allegedly provided this service to over 600 law enforcement agencies, allowing identification of virtually any individual in an image so long as that individual maintains an online presence[25].

One can also imagine highly invasive public uses of this technology by a regime that places surveillance cameras in all public areas to constantly aggregate information on the behavior of individuals, and sort that information by identity. In effect, application in this manner would create a database of where each person was, what their actions were, and who they were with any time they enter a public space. In conjunction with private technology in the home, e.g. one's smartphone, the same FRT could expand this database to include the behavior of that individual in private. This use of FRT for surveillance is already becoming commonplace in China, in which a vast network of over 300 million public-facing closed circuit surveillance cameras, coupled with advanced FRT, has been used to monitor its population for criminal conduct or dissident behavior[26]. China's government aims to couple the video data collected by these surveillance cameras with other personal data collected on citizens, including criminal and medical records, travel bookings, online purchases, and social media comments, to create a comprehensive government profile of each citizen[27]. Such invasive use of this technology highlights the potential for FRT to be used in manners that disregard personal privacy and suppress the exercise of expressed fundamental rights championed by the United States as a whole and California in particular.

Speech

As is the case with Fourth Amendment protections, it is currently unclear how the venerated First Amendment, which protects the freedoms of speech and assembly, applies to the use of FRT. In 1958, the Supreme Court held in *NAACP v. Alabama* that compelling the NAACP to disclose the identities of its members would likely hinder the ability of those members to advocate for their beliefs, and in 1960, the Supreme Court held in *Talley v. California* that a law prohibiting the anonymous distribution of pamphlets violated the First Amendment[28]. Taken together, these decisions indicate that the courts support an interpretation of the First Amendment that protects *anonymous* speech, which would presumably extend to the use of FRT to identify individuals exercising their freedoms of speech and assembly. Indeed, in its opinion in the case of *McIntyre v. Ohio Elections Commission* in 1995, the Supreme Court described anonymity as "the shield from the tyranny of the majority."[29] However, the unequivocal protection of anonymous speech under the First Amendment has not always

[24] Kashmir Hill, "The Secretive Company That Might End Privacy as We Know It," *The New York Times,* Jan. 18, 2020, https://www.nytimes.com/2020/01/18/technology/clearview-privacy-facial-recognition.html.

[25] *Ibid.*

[26] Simon Denyer, "China's watchful eye," *The Washington Post,* Jan. 7, 2018, https://www.washingtonpost.com/news/world/wp/2018/01/07/feature/in-china-facial-recognition-is-sharp-end-of-a-drive-for-total-surveillance.

[27] *Ibid.*

[28] *NAACP v. Alabama* (1958) 357 U.S. 449; *Talley v. California* (1960) 362 U.S. 60.

[29] *McIntyre v. Ohio Elections Commission* (1995) 514 U.S. 334, 357.

been the position of the Court. For instance, in *Laird v. Tatum* (1972), the Supreme Court held that military surveillance of public meetings did not have an "inhibiting effect" on the expression of First Amendment rights unless it created immediate danger of direct injury[30]. Several subsequent cases have used this decision to permit police photography of public demonstrations[31].

While such surveillance may be permissible, the courts have not weighed in on whether passive identification of the individuals surveilled during public demonstrations crosses the line into unconstitutional chilling of free speech and assembly. The use of FRT to disclose the identities of all individuals demonstrating for a given cause, as in *NAACP v. Alabama*, has the potential to hinder the ability to advocate for beliefs. In 2015, the FBI admitted to conducting surveillance flights over Ferguson and Baltimore during protests of police use of force, and that the Department of Homeland Security has reportedly surveilled protests by Black Lives Matter, an activist group focused on police brutality and discrimination[32]. This makes clear that the broad application of FRT to surveil political demonstrations could be particularly problematic. A 2011 Privacy Impact Assessment by the Department of Homeland Security, the FBI, and several state police agencies, in discussing the capacity for FRT to compromise anonymity in a manner inconsistent with the First Amendment, explicitly recognized that "surveillance has the potential to make people feel extremely uncomfortable, cause people to alter their behavior, and lead to self-censorship and inhibition."[33]

Bias

One of the major challenges raised by FRT concerns categorical disparities in performance, particularly in high-stakes situations in which errors can dramatically affect the life of an individual. The reliance of modern FRT on machine learning augments this concern, since training datasets and algorithmic parameters, designed or specified by humans, can reflect and magnify the existing biases of those individuals. Biases in automated mechanisms used to make critical decisions, such as whether an individual should be subject to invasive surveillance or whether someone's identification or citizenship documents are fraudulent, can result in errors leading to dire consequences for those individuals and to systemic discrimination against entire demographic groups.

[30] *Laird v. Tatum* (1972) 408 U.S. 1.

[31] E.g. *Donohoe v. Duling* (1972) 465 F.2d 196,202; *Phila. Yearly Meeting of Religious Soc'y of Friends v. Tate* (1975) 519 F.2d 1335, 1137-38.

[32] Eric Tucker, "Comey: FBI used aerial surveillance above Ferguson," *Associated Press,* Oct. 22, 2015, http://www.salon.com/2015/10/22/comey_fbi_used_aerial_surveillance_above_ferguson; George Joseph, "Exclusive: Feds Regularly Monitored Black Lives Matter Since Ferguson," *The Intercept*, Jun. 24, 2015, https://theintercept.com/2015/07/24/documents-show-department-homeland-security-monitoring-black-lives-matter-since-ferguson.

[33] The International Justice and Public Safety Network, "Privacy Impact Assessment: Report for the Utilization of Facial Recognition Technologies to Identify Subjects in the Field," Jun. 30, 2011, p. 016632.

Extensive research has determined the presence of biases in various forms of artificial intelligence, and recent reports indicate that FRT is no exception[34]. Studies have reported order-of-magnitude elevations in false positive rates (i.e. the number of images determined incorrectly to match an image in a database) for Asian vs. Caucasian faces and for African American vs. Caucasian faces. Additionally, those same studies reported lower false negative rates (i.e. the number of times an image did not matched to an image of the same individual that existed in a database) for African American vs. Caucasian faces[35]. These racial disparities in FRT performance could have resounding implications for racially biased law enforcement, particularly given the direction of the effects. For instance, if FRT is relied on for identifying criminal suspects from images, higher false positive rates and lower false negative rates for African American faces are likely to lead disproportionately to unwarranted investigation and arrest of African American individuals, who are already subjected to this form of discrimination. In other words, the biases existing in FRT have the potential to exacerbate and legitimize existing racial discrimination.

Despite substantial increases in overall accuracy in the past several years, a 2019 report published by the National Institute of Standards and Technology reaffirmed the prevalence of problematic biases in commercially available FRT algorithms[36]. The report detailed the performance of 126 FRT verification algorithms in matching 442,019 images from 24 countries with a database of 441,517 different individuals from the same countries. The report identified several problematic biases in commercially available FRT algorithms, including highest false positive rates for West and East African and East Asian faces, and lowest false positive rates for Eastern European faces. The report noted, however, that several of the algorithms developed in China reversed this effect, with East Asian faces showing the lowest false positive rates. The report also found higher false positive rates for women relative to men, and in the elderly and children compared with middle-aged adults. While the report encouragingly demonstrated that the most accurate algorithms were also the least biased between demographic groups, these results are nonetheless cause for concern, as consequences of their shortcomings seem to weigh most heavily on demographic groups already most vulnerable to discrimination. The perception that automated technology is entirely objective in its performance, even in the face of documented evidence to the contrary, makes these demographic disparities more concerning, as the inequities inherent in the technology can be easily overlooked.

IV. Creating a Regulatory Framework for Facial Recognition Technology

In their landmark treatise in the *Harvard Law Review* entitled "The Right to Privacy," Samuel Warren and Louis Brandeis mused on the difficulty of balancing the right to privacy with potential public utility, as follows:

[34] Center on Privacy & Technology, "The Perpetual Line-Up," *supra.*

[35] Phillips et al., *An Other-Race Effect for Face Recognition Algorithms*, 8 ACM Transactions on Applied Perception, 14:1, 14:5 (2011).

[36] *NIST,* "Face Recognition Vendor Test (FRVT) Part 3: Demographic Effects," *supra.*

> It remains to consider what are the limitations of this right to privacy, and what remedies may be granted for the enforcement of the right. To determine in advance of experience the exact line at which the dignity and convenience of the individual must yield to the demands of the public welfare or of private justice would be a difficult task.[37]

Developing a workable regulatory framework that acknowledges the utility of FRT to safety, security, and efficiency, while remaining conscientious of the potential for FRT to infringe on fundamental rights and civil liberties, such as the individual right to privacy and the freedom to express viewpoints anonymously, is indeed difficult, and requires consideration of several critical questions, both practical and conceptual. Further exacerbating the complexity of this task is the necessity in such a framework for sensitivity to the current technical shortcomings of FRT, including performance disparities between demographic groups and entrenchment of existing cultural biases by those designing and training the algorithms underlying these technologies. Questions requiring consideration as the Legislature contemplates confronting this challenge include the following:

- In what circumstances do the risks of FRT to privacy and civil liberties outweigh the utility? Should alternatives to FRT be preferentially employed whenever possible?

- How do the protections against unlawful search and seizure, guaranteed by the Fourth Amendment of the US Constitution, inter alia, apply to the use of FRT? Does the use of FRT to identify an individual constitute a "search"? Should the use of FRT to track an individual require a warrant?

- Does the use of FRT in public spaces inevitably threaten the presumption of innocence that is fundamental to the American justice system?

- Should an individual receive notification when entering their image into a facial recognition database? Should consent be required to enroll an image of an individual into a facial recognition database?

- Should an individual be informed when they may be subject to FRT in places open to the public? Should consent be required to subject an individual to FRT in a place open to the public?

- Should law enforcement facial recognition databases be limited to individuals who have been convicted of a criminal offense? Of a serious criminal offense?

- Should the State be permitted to provide drivers' license or other state ID photographs to law enforcement for enrollment in facial recognition databases?

[37] Warren & Brandeis, "The Right to Privacy," 4 Harvard L.R. 193 (1890).

- Should a law enforcement agency be permitted to purchase or otherwise enroll images of individuals that have been scraped from social media or other media accessible to the public?

- Should dragnet identification of any or all individuals in an image or real-time video be permissible, or should the use of FRT in real time be limited to targeted searches for specific individuals of interest?

- Should real-time, ongoing surveillance using FRT be permitted in any law enforcement context? In any commercial context? What burden of proof should be necessary to permit the use of ongoing surveillance using FRT by law enforcement? Reasonable suspicion? Probable cause?

- What types of cybersecurity safeguards should be required for databases containing FRT data?

- What types of regular audits of FRT should be required to identify misuse and ensure that performance does not result in discriminatory effects?

- Should meaningful human oversight of FRT be necessary for its use in making decisions of significant consequence?

The Committees recognize that the right solutions may not exist today, but are looking to identify the types of questions that the Legislature needs to consider to prepare adequately for the challenges that may arise from the adoption and implementation of FRT. It may very well be that future hearings are necessary, perhaps in conjunction with other policy committees, in further exploring topics raised at this informational hearing.