Which Characteristics May or May not Bias Facial Recognition Software?

By

Eric J. Clayborn

Master of Business Administration
Ferris State University, 2009

Bachelor of Science
Computer Information Systems
Ferris State University, 2004

Associate of Applied Science
Accounting
Southwestern Michigan College, 2002

Advisor:
Dr. Greg Gogolin
Full Professor
Information Security and Intelligence

Fall, 2020
Ferris State University
Big Rapids, Michigan

DEDICATION

This thesis is dedicated to my wonderful wife Elva C. Clayborn, MA. It is only with her love and support that I have been able to succeed.

ACKNOWLEDGEMENTS

First, I would like to thank the participants of this thesis, starting with my wife Elva C. Clayborn, MA and my good friend Richard Kiner. Second, I would like to thank my mother Evelyne Belizaire for always encouraging me to seek out a great education. Next I would like to thank my friends Dr. David Fleming PhD and Dr. Stacy Young PhD for always believing in my ability to redefine myself through hard work and continual improvement.

Finally, I would like to thank my advisor Dr. Greg Gogolin for mentoring me throughout this project along with support from Dr. Furstenburg, Dr. Cooper and Dr. Schuiling for the challenging and thought-provoking coursework that helped prepare me for my thesis.

Table of Contents

List of Table and Figures

Abstract

Biometric technology, such as facial recognition, has been known to falsely identify people of color more often than other demographics. This issue has led to the loss of civil liberties for those who are the victim of false-negative rates (Buolamwini, & Gebru, 2018). This is a study to identify characteristics that may bias biometric readings, with particular interest in facial recognition software. This study uses facial recognition software on a personal device from a Samsung Android S10 smartphone. Three individuals with varying age, gender, skin colors and tones participated in this study. Glasses, sunglasses, a hat, a black mask, and a white mask served as variables to obscure portions of their faces. Interesting patterns emerged while analyzing the results of this study. Most surprising was the dark-skinned African American male unlocked the Android device while wearing dark sunglasses along with a hat, which resulted in a false-positive match. The Android device sensor was unable to differentiate his eyes from the dark sunglasses or his forehead from the hat, which may indicate a limitation of the facial recognition software to distinguish darker skin tones. Ascertaining biases against people of color within the facial recognition feature of a personal-use device requires further study.

Which Characteristics May or May not Bias Facial Recognition Software?

Chapter 1

Introduction

Biometric technology is used as a means for identification and applies the type-three method of "something you are" category for authentication. Many newer devices such as laptops, smartphones, door-access, desktops, cameras, applications, safes, etc. employ some type of biometric technology. Security is increased when different types of biometric frameworks are combined, such as type-one and type-three, thereby creating a multi-factor authentication process. The challenge with any type of technology is that there are vulnerabilities within the process or framework, and the use of biometrics is not immune to such faults.

Type-three category biometric inputs take unique characteristics of the human body to help identify an end user. Some of the inputs include fingerprints, iris scanning, retina scanning, voice, and facial recognition. These character sets are distinctive to every individual, however there are instances of false rejection (false positive) and false acceptance (false negative) rates within the technology. These error rates are where system administrators and threat actors focus their efforts to locate vulnerabilities.

Many experts have suggested that the error rates within biometrics have caused the public to surmise that there are certain biases built into the authentication process of biometrics. Critics have targeted facial recognition software due to its high use within CCTV and law enforcement. Facial recognition uses machine learning algorithms which have been found to cause higher than normal false rejection rates within certain demographics of the population, whom by and large are people of color and women.

Private companies are able to capitalize on facial recognition while working with local governments to help establish a database of stored images for use in identification. False positive rates can vary across demographics by factors of 10 to 100 times (Grother, Ngan, & Hanaoka, 2019). It is highly possible that the information provided to computer algorithms is flawed and therefore has an effect on false-positive outcomes. The goal should be to determine how biometric-based machine learning uses datasets to identify individuals with a high level of accuracy.

Background

Biometric authentication uses a one-to-one measure to identify an individual through various means. Despite its level of accuracy, many studies have shown that biometrics such as facial recognition have a higher failure rate pertaining to specific demographics. On May 22, 2019, the Federal House Committee on Oversight and Reform held a session to discuss the challenges of identifying individuals through the use of facial recognition technology (Maloney, 2020). Some local jurisdictions such as San Francisco California, and Portland Oregon have banned the use of private-sector use of facial recognition technology in the public domain (Metz, 2020).

Evidence of false acceptance rates or false positive identification is prominent throughout various vendors. In 2018, a test of Amazon's "Rekognition" software, falsely matched twenty-eight members of Congress with mugshots of individuals who have committed crimes. The test of congressional members disproportionately and falsely matched people of color (Snow, 2018). Another example of false positive identification occurred when Apple's facial recognition algorithm falsely identified an individual of color in connection with multiple crimes in different states. One of the alleged crimes happened at the same time that the individual was at a high school prom (Shaban & Flynn, 2019).

Statement of the Problem

Some biometric tools rely on up-to-date database information in order to validate authentication and identification. Facial recognition uses information from databases sourced from various means in order to identify individuals. Databases used within a facial recognition algorithm are used to help develop machine learning which leads to higher or lower positive outcomes. If the database or algorithm is skewed or compromised in any way, it can have a less than optimal outcome for those who are incorrectly identified (Martinez-Martin, 2019).

Currently, there are challenges to accurately identify certain demographics of the population. Biometric technology, such as facial recognition, has a habit of falsely identifying people of color more often than other populations. This issue has led to the loss of civil liberties for those who are the victim of false acceptance or false error rates (Buolamwini, & Gebru, 2018).

Purpose of the Study

The purpose of this study is to provide empirical evidence of the accuracy and effectiveness of data to help accurately identify individuals through the use of biometric technology, specifically facial recognition technology. Information gathered on biometric data is sourced through different means. Many individuals who use fingerprint, voice, retina, and iris scans offer this information willingly. However, spoofing techniques can be used to compromise the mechanisms that provide security. Facial recognition data is often considered information that is sourced through the use of CCTV or social media platforms. Database information from driver’s license data can also be used for facial recognition (Harwell, 2019).

Rationale

Research objectives and questions will include multiple items. The first item pertains to the philosophy of biometric facial recognition technology and how it is designed to work. Detailed information on the various types of facial recognition authentication will be provided and their intended outcomes. Second, I will provide real-world examples of failed facial recognition implementations. Biometric features can fall prey to manipulation or false acceptance rates which can lead to the detriment of a targeted user. Third, I intend to deliver insight and information regarding congressional debate with the use of facial recognition. The use of facial recognition as a tool for local jurisdictions has developed into a national debate. The datasets used in facial recognition software to develop machine learning have been linked to false positives for specific demographics. Finally, I plan to discuss the rationale by local jurisdictions who have banned the use of certain facial biometric technology as a means to identify its citizens.

Research Questions

The research will answer the following questions:

1. Is the facial recognition feature of my personal biometric device, disproportionally biased against people of color?
2. What types of errors may compromise the accuracy of facial recognition biometric devices?

Significance of the Study

The use of biometric technology is prevalent throughout our technological ecosystem. Facial recognition software is one of the most used aspects of biometrics that helps identify individuals. Many local jurisdictions use facial recognition to identify potential suspects by comparing evidence to known datasets. Facial recognition algorithms use machine learning based on the provided dataset and often misidentify people of color more often than Caucasian individuals. Currently, there are no federal regulations on how to use facial recognition technology for both private and governmental organizations. There is currently no regulation on how datasets are developed to properly identify individuals throughout the machine learning process of facial recognition. While many agencies are rushing to use facial recognition as a means to identify suspects, there is no indication that one's first amendment or fourth amendment rights are being protected.

Jurisdictions are taking steps to mitigate some of the inherent risks from the use of facial recognition technology. Cities such as San Francisco CA, and Portland OR are banning the use of facial recognition as a means for identifying individuals within their locality (Metz, 2020). Many critics of facial recognition feel that even if the technology was 100% effective, using machine learned facial recognition technology to spy on citizens would fundamentally alter the concept of a free society. Other locations outside the U.S. are also considering regulations towards facial recognition. The European Union is considering regulations as a part of their initiative to develop a legal framework for trustworthy artificial intelligence. The increased use of facial recognition has the potential to undermine fundamental rights, in particular, people of color.

Definitions of Terms

- Biometrics - Often defined as either physiological or behavioral. The physical biometrics are unique attributes that help identify who an individual is through the use of their fingerprints, eyes, voice, and facial features (Chapple, Stewart, & Gibson, 2018).
- Facial Recognition - Input devices such as computer web-cameras are used to register an individual's unique facial nodal points. Face scans use the geometric patterns of faces for detection and recognition (Chapple, Stewart, & Gibson, 2018).
- Voice Recognition – Relies on characteristics of a person's speaking voice, known as a voiceprint. Often used as a secondary authentication feature (Chapple, Stewart, & Gibson, 2018).
- Fingerprint – Visible patterns on the fingers and thumbs of human beings which are unique (Chapple, Stewart, & Gibson, 2018).
- Retina Biometrics - Infrared light is used to scan the patterns of blood vessels on the back of the human eye. Retina scanners require end users to be closer than three inches from the scanner (Chapple, Stewart, & Gibson, 2018).
- Iris Scanning – These types of scans focus on the colored area around the human pupil. This type of scanning is more customary compared to iris scanning because it can be done from six to twelve inches from the input scanner (Chapple, Stewart, & Gibson, 2018).
- Spoofing - Attacks with the goal of gaining access to a target system through the use of falsified identity (Chapple, Stewart, & Gibson, 2018).

Assumptions and Limitations

Several limitations and assumptions are offered in the study. One limitation is the use of lighting. The lighting was not exactly the same for each individual who participated in the study. The lighting is exactly the same for the Caucasian female and the light-skinned African American male. The lighting is different for the dark-skinned African American male, as his portion of the study was conducted in his home.

A second limitation of this study is the small sample size. Social distancing requirements from the state of Michigan limited the number of participants in this study. This study uses three different skin tones; however, the sample size was limited due to state-manded health restrictions.

A third limitation was inherent in the technology of the personal device. The Samsung S10 Android smartphone allows for multiple reference images for fingerprints from different participants. However, it does not allow for multiple face print reference images from different participants. This limitation does not allow for the S10's facial recognition algorithm to develop any type of machine-based learning to identify a participant.

Chapter 2

Search Terms used in Literature Review

Facial Recognition Bias, Biometric bias, facial recognition minorities, biometric spoofing, Iris Biometric bias, Biometric Bias Fingerprint, Characteristics that bias biometric fingerprint, bias retina scanning, biometrics bias retina scanning, bias retina biometric

Definition of Biometric Technology

Biometric technology is used to help authenticate individuals to access such things as smartphones, software, doors, laptops, and desktops. Traditional authentication practices use the process of "something you know" to authenticate a person using a password. Many passwords can be cracked using various methods and are limited by how well an individual can remember their password. Biometric technology uses unique characteristics of the user through a process called "something you are." Human beings have unique attributes that help identify who they are through the use of their fingerprints, eyes, voice, and facial features. Input scanning features within various technologies such as fingerprint scanners, retinal scanners, iris scanners, microphones, and cameras can register these unique human characteristics which are then compared to future authentication attempts (Ciampa, 2018).

Using the human eye can be used for authentication in multiple ways. The human iris is responsible for controlling the diameter and size of the pupils. This allows the regulation of light which reaches the retina. By illuminating the iris with light, biometric technology can detect unique patterns, which are not visible to others and thereby authenticate an individual. By contrast, retinal scanning uses the complex structure of capillaries in the human eye which are unique to each individual. This type of scanning authenticates the user by using low-level infrared light to scan for unique variations of blood vessels into an individual's eye (Ciampa, 2018).

Human fingerprints have been used to identify individuals long before the use of biometric technology. Law enforcement agencies use fingerprints to identify suspects in active cases. Every human's fingerprints have unique ridges and valleys. The ridges are the upper skin layer segments, and the valleys are the lower segments. Static fingerprint scanners allow the user to place their entire finger or thumb in a small oval window to register their unique dataset. Dynamic fingerprint scanning tools often consist of small slits or openings near a laptop keyboard. Both types of scanning features take optical photos of the registered fingerprint and compare it with future authentication attempts (Ciampa, 2018).

Voice recognition uses the unique phonetic cadence of an individual to authenticate that user. Characteristics such as the size of one's head and age help build the unique trait of a person's voice. Input devices such as a camera or a microphone are used to register a person's unique voice attributes (Ciampa, 2018).

Facial recognition is one of the most popular forms of biometric authentication used in industry. Human beings have approximately eighty unique nodal points that make up their facial features. Some of these points consist of depth of eye sockets, the shape of cheekbones, size and shape of one's nose, an individual's skin tone, and the length of one's jaw line. Input devices such as computer web-cameras are used to register an individual's unique facial nodal points (Ciampa, 2018).

The general face recognition pipeline for personal smart devices are as follows. The first step is to acquire a selfie face image using a personal smart device. Next, face detection from the smart device determines if there is an image. If an image is detected, it is segmented. Third, the face image is normalized, and facial features are extracted (Rattani & Derakhshani, 2018). Finally, identification verification is matched by comparing features from two face images for authentication.

Facial recognition technology is often used by law enforcement agencies to scan crowds of people for alleged fugitives, missing persons, or even terrorists. However, variables such as poor lighting, hats, sunglasses, or masks make the use of this type of policing tool less precise than personal facial recognition used on smartphones or other computer devices.

Spoofing Biometric Technology

As with any type of technology, threat actors will make attempts to attack the security of biometric technology. Attacks with the goal of gaining access to a target system through the use of falsified identity is an act of spoofing. There are many examples where threat actors have gained access to devices through the use of spoofing tactics (Chapple, 2018).

Iris recognition scanning can fall prey to presentation attacks. Presentation attacks allow threat actors to present a false identity to an iris sensor with the goal of manipulating the biometric system into a false positive decision (Boyd, Fang, Czajka, & Bowyer, 2020). These presentation attacks can include using textured contact lenses, a paper iris printout, and prosthetic or even cadaver eyes (Fang, 2020). Two types of presentation attacks are categorized as impostor and concealer attacks. Impostor attacks may use paper printouts of iris images or replay attacks where bona fide iris images displayed on a screen are presented to a sensor. Synthetic iris images can be used as possible concealer attacks where iris images are used to imitate bona fide iris patterns (Boyd, Fang, Czajka, & Bowyer, 2020).

One of the most commonly used biometric traits for authentication is fingerprint authentication, however, it is also one of the most spoofed. Many devices such as laptops, smartphones, door access, and computers use some type of biometric technology which incorporate a fingerprint reader. Threat actors are able to spoof a fingerprint biometric system with fake samples. Fingerprint scanners can be spoofed through the use of materials that closely resemble the finger itself, such as silicone, or modeling clay. Solutions to help prevent fingerprint spoofing include the use of liveness detection through the local descriptor or binary patterns with filters (Balaji, 2016).

Voice recognition systems use automatic speaker verification (ASV) to help authenticate end users. Replay attacks are often used to spoof ASV technology. There are various points of interest that voice recognition systems are vulnerable to spoofing attacks. Some of the points of interest include a presentation attack at the microphone and modifying biometric samples to bypass the sensor, which are respectively known as physical and logical access (Chettri, 2020). Logical access can also be triggered through text-to-speech and modified speech generated through voice conversations. Replaying pre-recorded samples of speech is another example of physical access. Countermeasures designed to discern between real human speech or bona fide samples are developed through the use of translating raw acoustic waveforms to a sequence of short-term feature vectors (Chettri, 2020). To help protect against spoofing, identification phrases can be selected that would rarely come up in normal speech.

Research has found that facial recognition software can be spoofed through various means. Users can be impersonated by using 3D-printed masks or use face images which are downloaded from social networks (Sharif, Bhagavatula, Bauer, & Reiter 2016). However, attacks such as this can be mitigated by anti-spoofing mechanisms that use liveness detection.

The Great Debate on the use of Facial Recognition Technology

Facial recognition technology has changed the way individuals are identified and authenticated. There has been debate on the use of facial recognition to identify individuals for the use of law enforcement. Recent research developed by Joy Buolamwini and the Massachusetts Institute of Technology Media Lab identified that machine learning algorithms can discriminate based on categories such as race and gender. In her research, Ms. Buolamwini evaluated three commercial gender classification systems using specific datasets. The outcome determined that darker-skinned females were the most misclassified group compared to lighter-skinned males (Buolamwini & Gebru, 2018). Many facial recognition tools use machine learning algorithms that reference labeled datasets. If these datasets have biased information within them, it can skew machine learning and result in algorithmic discrimination. This type of skewed bias can have an effect on false positive rates within commercial facial recognition software.

Some studies have found that using diverse datasets increase the chances of correctly identifying an individual. For example, experiments using the MORPH style dataset contain a large number of African American and Caucasian images. This is in contrast with other publicly available datasets which only contain small numbers of African American images and would not allow for conclusive results (Bowyer, 2019).

Other companies have identified opportunities that can reduce the amount of bias that is produced within facial algorithms. Vintra, a California based analytics company, claims that it developed a facial recognition platform (FulcrumAI) that reduces bias compared to companies such as Amazon, Microsoft, and ArcFace. Vintra states that its analytics platform developed improvements that are based on building its own dataset from scratch. Vintra was able to access images from over 75 different countries and thousands of identities in order to better represent Caucasian, African, Asian, and Indian races. Taking these steps allows the machine learning to get a truer picture of what the world looks like. Western-based face algorithms, such as those used by Amazon and Microsoft are considered biased because they are built using datasets that have a super-majority of white faces. No amount of AI tweaks and changes can overcome the biased racial representation (Vintra, 2020).

The collection of datasets is important, not only in the way the data is collected and labeled but also how it is administered after the collection. Many within the Federal Bureau of Investigation (FBI) and Immigration and Customs Enforcement (ICE) have incorporated photos from states' driver's licenses into facial recognition datasets. A 2019 Washington Post article revealed that federal investigators have turned state department motor vehicle databases into a surveillance infrastructure (Harwell, 2019). This was done without the expressed authorization of state legislatures or individual license holders. Some states allow undocumented immigrants access to full driver’s licenses, which allows ICE to run searches on the databases. The article points out that if an FBI facial recognition search is performed, it is 86% accurate at finding the correct person. However, the FBI’s software is still dependent on factors such as poor lighting, image quality, and the software performs less accurately on individuals with darker skin (Harwell, 2019).

In August of 2019, an 18-year-old male named Ousmane Bah, from the state of New York filed suit against Apple for $1 billion. The lawsuit claimed that Apple’s facial recognition system falsely connected him with a series of thefts ranging from multiple jurisdictions, such as Delaware, Massachusetts, New Jersey, and New York (Shaban & Flynn, 2019). Another recent article points out that in July 2019, a 26-year-old male named Michael Oliver was arrested in Detroit Michigan, due to false identification from a facial recognition software program developed by Data Works Plus. The Detroit police chief stated, “If we were just to use the technology by itself, to identify someone… I would say 96% of the time it would misidentify" (Stokes 2020). The outcome from the false arrest resulted in loss of income and lost job opportunities for Mr. Oliver. Both Mr. Oliver and Mr. Bah are people of color.

Many local jurisdictions have taken steps to limit their law enforcement apparatus from using facial recognition technology on its citizens. In 2019, the city of San Francisco, California was one of the first major cities to ban police and other agencies from using facial recognition technology. In August 2019, the Swedish data protection authority issued a decision that found that schools which used facial recognition to track attendance of students violated the European Union's General Data Protection Regulation (GDPR) (Nesterova, 2020). In September 2020, the city of Portland Oregon banned the use of facial recognition by local police departments and local area businesses. Portland's concern was over the built-in racial bias in the facial recognition software and potential for misuse which could lead to fundamental privacy issues (Metz, 2020).

In 2018, the American Civil Liberties Union (ACLU) conducted a test of Rekognition, Amazon's face surveillance technology. The results of Amazon's Rekognition software matched 28 members of U.S. Congress with individuals who have been arrested for committing crimes. The congressional members who were falsely matched with the mugshot database were comprised of male and female Republican and Democratic members. However, these false positive matches disproportionately targeted people of color, including such members as Rep. John Lewis. In a letter sent from Congress to Amazon CEO Jeff Bezos, it was noted that approximately 40% of the false matches were people of color, even though they made up only 20% of Congress in total (Snow, 2018).

On January 15, 2020, the U.S. House Committee on Oversight and Reform held its third hearing on facial recognition technology: Ensuring Commercial Transparency & Accuracy. The committee deliberated on the various ways the private sector entities use facial recognition technology (Maloney, 2020). As a part of the deliberation, Charles H. Romine, director of the National Institute of Standards and Technology (NIST) issued a report that found higher than normal false positive rates among demographics to vary by factors of 10 and at times 100 times (Romine, 2020). Romine stated, “For most algorithms, the NIST study measured higher false positive rates in women, African Americans, and particularly in African American women” (Romine, 2020). Mr. Romine also indicated that the high levels of false positive rates can have varying degrees of consequences to the system owner and user. Some of the consequences range from allowing access to impostors and civil rights abuses from loss of liberty. Mr. Romine indicated that there were no such false positives in one-to-one matching between Asian and Caucasian faces for those algorithms developed in the Asian continent. The report also suggested that algorithms with fewer demographic differentials can be anticipated to produce fewer errors (Romine, 2020).

Facial recognition software can cause false positive outcomes which produce life changing consequences for those who are victims of the error. Datasets that use fewer demographic differentials have a higher-than-normal accuracy rate. However, those datasets are not as publicly available to those which contain smaller amounts of images of people of color, and women. There are no clear guidelines on how local jurisdictions and private companies must administer the algorithms or framework that guide the machine learning needed to accurately identify individuals. Organizations such as the ACLU and members of Congress should address the seemingly unequal use of facial recognition on civil liberties. Until a national framework is legislated, it will be up to the states and local jurisdictions to adopt or ban the use of facial recognition technology to accurately identify individuals.

Chapter 3

Description of Methodology

This is an exploratory study to identify characteristics that may bias biometric readings, specifically facial recognition software. The primary goal is to determine if skin color, combined with other factors creates false-positive or false-negative readings.

Design of the Study

This study used facial recognition software from a Samsung Android S10 smartphone. Three individuals with varying skin colors and tones were used in this study. The first individual is an African American male with a dark skin tone. The second individual is an African American male with a light brown skin tone. The third individual is a Caucasian female with a light skin tone. Glasses, sunglasses, a hat, a black mask, and a white mask served as variables to obscure portions of their faces. The masks were also folded in various manners to create additional variables to determine if this made a difference in the facial recognition software on the Android S10.

The first step of this study was to register the facial image of each participant without glasses, hats, or masks with the Android S10 facial recognition software. This step was completed at the beginning of each individual's turn to complete the study.

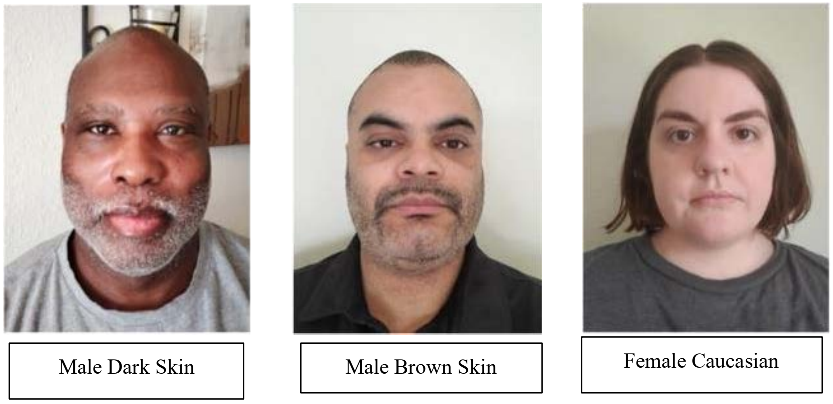

Male Dark Skin

Male Brown Skin

Female Caucasian

Figure 1: Facial Reference Images Used in Initial Registration of Facial Recognition Software

The second step of this study was to introduce the variables to obscure the faces. After each variable was introduced, the individuals attempted to unlock the Android S10 phone using the facial recognition software. Table 1 contains the results of this process. “Match” is used to designate if the facial recognition software successfully unlocked, while “No Match” was used if the device was not unlocked. The images in Figures 2, 3, 4, and 5 show examples of variables used to test the facial recognition software.

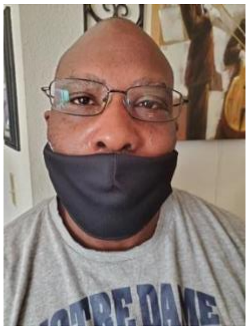

Figure 2: Glasses with Black Mask Folded Exposing Nose

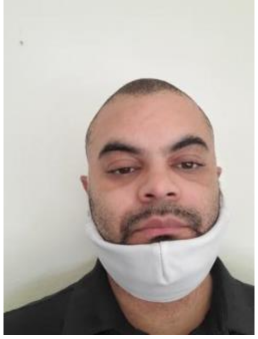

Figure 3: White Mask Folded Exposing Nose and Lips

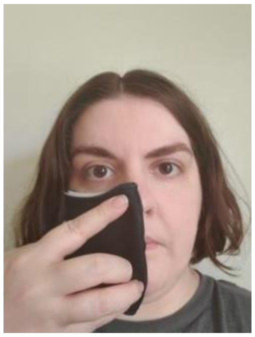

Figure 4: Black Mask Folded in Half, Covering Right Side of Face

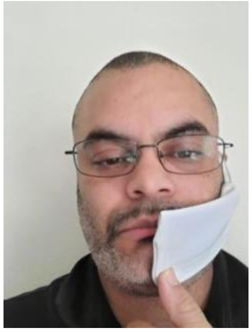

Figure 5: Glasses - White Mask Folded in Half, Covering Left Side of Face, Exposing Nose

Table 1: Data from Spoofing Facial Recognition Software on Samsung Android S10

Facial Variables	Male/Dark Skin	Male/Brown Skin	Female/Caucasian
Control Images	Match	Match	Match
Glasses/Hat/Black Mask	No Match	No Match	No Match
Sunglasses/Hat/Black Mask	No Match	No Match	No Match
Glasses/Hat/White Mask	No Match	No Match	No Match
Sunglasses/Hat/White Mask	No Match	No Match	No Match
Glasses/Black Mask	No Match	No Match	No Match
Sunglasses/Black Mask	No Match	No Match	No Match
Glasses/White Mask	No Match	No Match	No Match
Sunglasses/White Mask	No Match	No Match	No Match
Hat/Black Mask	No Match	No Match	No Match
Hat/White Mask	No Match	No Match	No Match
Glasses Only	Match	Match	Match
Sunglasses Only	Match	No Match	No Match
Hat/Glasses	Match	Match	Match
Hat/Sunglasses	Match	No Match	No Match
Hat Only	Match	Match	Match
Black Mask	No Match	No Match	No Match
White Mask	No Match	No Match	No Match
Black Mask Folded, Exposing Nose	Match	Match	No Match
Glasses/Black Mask Folded, Exposing Nose	Match	No Match	No Match
White Mask Folded, Exposing Nose	Match	No Match	Match
Glasses/White Mask Folded Exposing Nose	Match	No Match	No Match
Black Mask Folded, Exposing Nose And Lips	Match	Match	Match
Glasses/Black Mask Folded Exposing Nose and Lips	Match	Match	Match
White Mask Folded, Exposing Nose And Lips	Match	Match	Match
Glasses/White Mask Folded Exposing Nose and Lips	Match	Match	Match
Black Mask Folded in Half, Covering Left Side of Face	No Match	No Match	No Match
Glasses/Black Mask Folded in Half, Covering Left Side of Face	No Match	Match	No Match
White Mask Folded in Half, Covering Left Side of Face	No Match	No Match	Match
Glasses/White Mask Folded in Half, Covering Left Side of Face	No Match	No Match	Match
Black Mask Folded in Half, Covering	No Match	No Match	No Match

Right Side of Face			
Glasses/Black Mask Folded in Half, Covering Right Side of Face	No Match	No Match	No Match
White Mask Folded in Half, Covering Right Side of Face	No Match	No Match	Match
Glasses/White Mask Folded in Half Covering Right Side of Face	No Match	No Match	Match
Black Mask Folded in Half, Covering Left Side of Face, Exposing Nose	Match	Match	Match
Glasses/Black Mask Folded in Half Covering Left Side of Face, Exposing Nose	Match	Match	Match
White Mask Folded in Half, Covering Left Side of Face, Exposing Nose	Match	Match	Match
Glasses/White Mask Folded in Half Covering Left Side of Face, Exposing Nose	Match	Match	Match
Black Mask Folded in Half, Covering Right Side of Face, Exposing Nose	Match	Match	Match
Glasses/Black Mask Folded in Half, Covering Right Side of Face, Exposing Nose	Match	Match	Match
White Mask Folded in Half, Covering Right Side of Face, Exposing Nose	Match	Match	Match
Glasses/White Mask Folded in Half Covering Right Side of Face, Exposing Nose	Match	Match	Match

Chapter 4

Data Analysis

The data in Table 1 was gathered as variables were introduced during our study. The first column lists each type and combination of variables used. The second, third, and fourth columns contain the results respective of our three participants. The second column represents an African American male with dark skin. The third column represents an African American male with light, brown skin. The fourth column represents a Caucasian female with light skin. As previously mentioned, an indicator of a "Match" signals the participant was successfully able to unlock the device, using its facial recognition technology. A "No Match" result signals the participant was unable to unlock the device.

Interesting patterns emerged while analyzing the results of this study. In the beginning of the study, it appeared the facial recognition software would not recognize a user unless the individual's eyes were visible. However, the software detected the dark-skinned African American male's face even while he wore sunglasses. Figure 6 shows the difference of contrast between each of the participants' skin tones compared to the sunglasses used in this study.

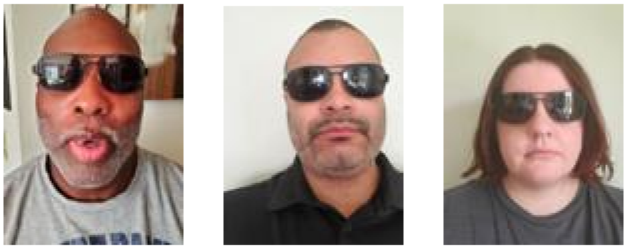

Figure 6: Sunglasses Worn by Participants with Varying Skin Tones

A second pattern emerged involving facial recognition software correctly identifying an individual's face when their eyes and nose are visible. Successful matches occurred when the African American male participants wore black masks with their eyes and noses exposed. However, the facial recognition software did not recognize the Caucasian participant's face when her eyes and nose were exposed while wearing a black mask. Figure 7 shows the participants wearing black masks with their eyes and noses exposed. On the converse, successful matches occurred when the dark-skinned African American male and the Caucasian female wore white masks with their eyes and noses exposed. However, the facial recognition software did not recognize the light-skinned African American male's face when his eyes and nose were exposed while wearing a white mask. Figure 8 shows the participants wearing white masks with their eyes and noses exposed.

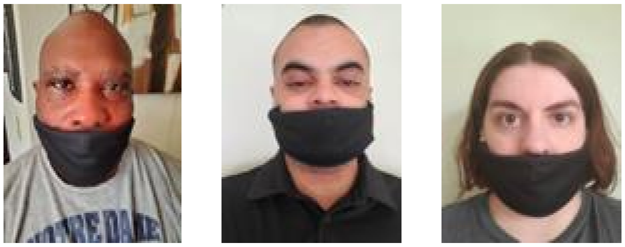

Figure 7: Black Masks Worn by Participants with Varying Skin Tones with Eyes and Nose Exposed

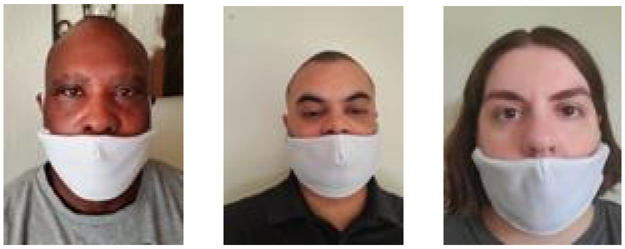

Figure 8: White Masks Word by Participants with Varying Skin Tones with Eyes and Nose Exposed

Eyeglasses worn by the participants typically created successful matches with the facial recognition software, possibly due to the visibility of their eyes. However, when the participants wore black masks and eyeglasses and exposed their noses, the software only recognized the light-brown skinned African American male. There was no match detected for the dark-skinned African American male or the Caucasian female. Figure 9 shows the participants wearing black masks and glasses with their eyes and noses exposed. The same results occurred when the subjects wore white masks and glasses and exposed their noses. Only the light-brown skinned

African American male was recognized by the facial recognition software. Figure 10 shows the participants wearing white masks and glasses with their eyes and noses exposed.

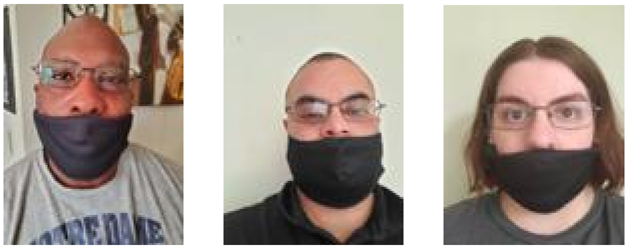

Figure 9: Black Masks and Glasses Worn by Participants with Varying Skin Tones with Eyes and Nose Exposed

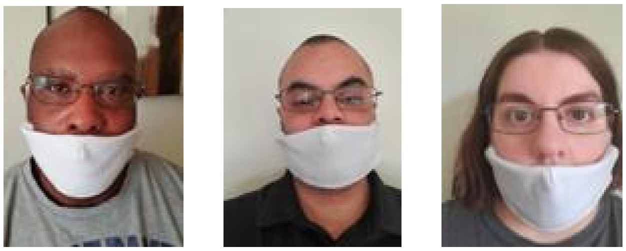

Figure 10: White Masks and Glasses Worn by Participants with Varying Skin Tones with Eyes and Nose Exposed

The remainder of this study tested whether folding the masks in various ways might result in a match for the subjects. For example, Figure 11 shows each of the participants with eyeglasses and a black mask folded to cover the left side of their faces. This resulted in a match for only the light-brown skinned African American male. However, when the subjects wore glasses and covered the right and left sides of their faces with a white mask, only the Caucasian female received a match in each of these scenarios. Figure 12 shows an example of the subjects wearing glasses with the left side of their faces covered with the white mask. The Caucasian female was also the only individual to receive matches for covering the left and right sides of her face with the white mask without glasses.

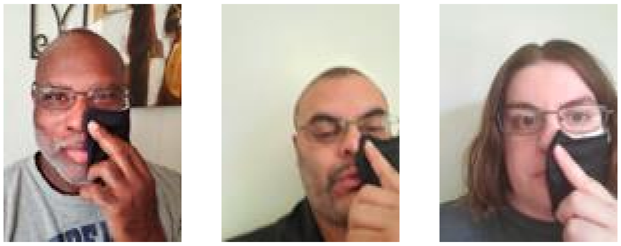

Figure 11: Glasses and Black Masks Worn by Participants with Varying Skin Tones with Left Side of Face Covered

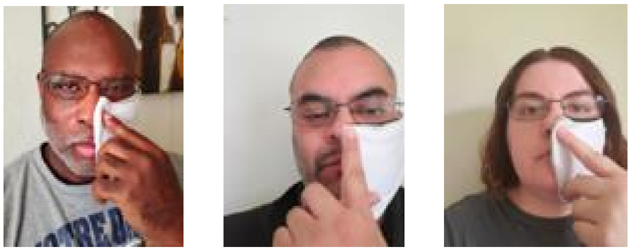

Figure 12: Glasses and White Masks Worn by Participants with Varying Skin Tones with Left Side of Face Covered

It is important to note that there were many results throughout the study that were correct and anticipated, resulting in matches and non-matches that confirmed the facial recognition technology worked as designed. For example, matches occurred for all three participants when they wore only glasses and/or a hat while the remainder of their faces were unobstructed. Matches also resulted from folding the black and white masks and allowing the participants' noses and lips to be exposed, as only their chins were obstructed (See Figure 13).

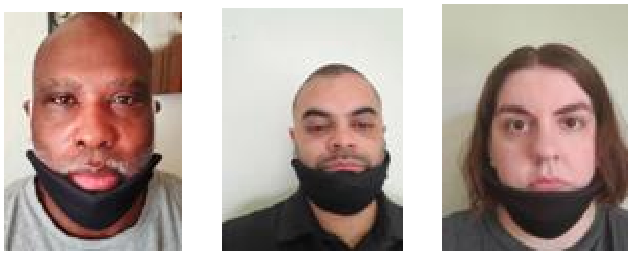

Figure 13: Black Masks Worn by Participants with Varying Skin Tones with Nose and Lips Exposed

There were no matches for the participants when the masks covered their cheeks, nose, lips, and chins. Also, no matches resulted from the combination of hats, sunglasses, and masks. These non-matches were anticipated, as the participants' faces were mostly obstructed and in some cases almost completely obstructed (See Figure 14). While there were a few surprises while conducting the study, the facial recognition software on the Android S10 appeared to work correctly as designed with fairly consistent results.

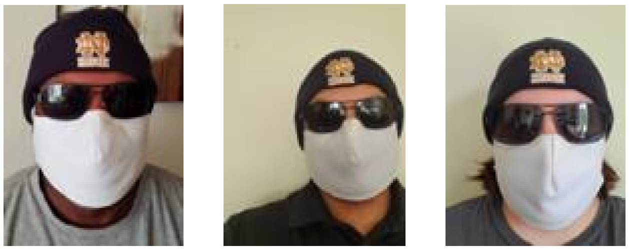

Figure 14: Sunglasses, Hat, and White Masks Worn by Participants with Varying Skin Tones

Chapter 5

Conclusions

During the study, some errors pertaining to the use of the device and environmental variables contributed to the accuracy. The participant's face was required to be no more than six inches away from the Android S10 in order to both register and authenticate the facial reference profile. If the participant was too far away or too close to the camera's sensor, the device was unable to perform its task of authentication. The end users also needed to trigger the facial recognition software by pressing a button on the side of the device; not doing so could cause false-negative assumptions. Background lighting was also important during the study. The device required enough background lighting to capture the image of the participant. The study was performed during daylight hours when natural sunlight was available, creating a consistent atmosphere for all participants.

Ascertaining biases against people of color within the facial recognition feature of the Samsung S10 device was inconclusive. The features for biometric facial registration differ on a personal device from what is used to identify individuals within a public setting, such as those used within CCTV. Local jurisdictions gather photos from online screen-scrapes from social media, driver's license databases, mugshots, etc., and use them to develop datasets where facial recognition algorithms can develop accurate machine learning for identification. The biometric technology within the Android S10 allows for multiple reference profiles of various fingerprints from different individuals but does not allow for multiple reference profiles of facial inputs. Only one reference profile is allowed at a time for the facial recognition feature within the Android S10. Thus, the algorithm used within the Android S10 could not develop machine learning based upon multiple facial reference profiles. However, that did not prevent this study from developing surprising results.

The results of this study raised some interesting points. Most surprising was that the dark-skinned African American male unlocked the Android device while wearing dark sunglasses along with a hat, which resulted in a false-positive match. The Android device sensor was unable to differentiate his eyes from the dark sunglasses or his forehead from the hat, which may indicate a limitation of the facial recognition software to distinguish darker skin tones. Neither of the other two participants were able to unlock the Android device while wearing the dark sunglasses alone or with the hat. The Caucasian female was the only individual to receive matches for covering the left and right sides of her face with the white mask without glasses. This study cannot conclude an overall correlation of bias, however, there does seem to be some limitations of the S10 facial recognition algorithm and/or camera sensor to distinguish skin tone.

Recommendations

Future research and replication of this study should consider including additional skin tones to determine if facial recognition software produces false negatives or false positives. For example, individuals of Hispanic, Arabic, Indian, Chinese, and many other backgrounds should be included if possible. Additionally, Caucasian males, African American females, and transgender individuals should also be considered to determine if these factors cause biometric screening bias.

It is also recommended that various facial recognition software devices be used to complete future research. For example, an iPhone could be used instead of an Android device. Devices that offer multiple inputs for facial reference profiles would help determine if machine learning algorithms are being used to develop accurate matches.

Finally, lighting should be considered during future studies. It is recommended that lighting for each participant be the same, which is likely to result from conducting the study in the same location at approximately the same time of day for each individual. Natural and artificial lighting could be used to determine if these variables make a significant difference in correct and incorrect results.

Further Study

Further study is needed to identify how facial recognition algorithms develop machine learning based on the datasets used for comparison. Datasets which have a smaller number of demographic differentials have been found to be more accurate than those with a larger differential. Unfortunately, the larger differential datasets, which are drawn from publicly available sources such as driver's license databases, social media, and other local jurisdictions are used more often for identification due to their ease of access. These skewed datasets can have an impact on the machine learning and thus create a need for better software development to help prevent false-positive outcomes.

References

Balaji, H. (2016). Multimodal Fingerprint Spoof Detection Using White Light. *Procedia Computer Science, 78*, 330–335. https://doi.org/10.1016/j.procs.2016.02.066

Bowyer, K. (2019). Why face recognition accuracy varies due to race. *Biometric Technology Today, 2019*(8), 8–11. https://doi.org/10.1016/S0969 4765(19)30114-6

Boyd, A., Fang, Z., Czajka, A., & Bowyer, K. (2020). Iris presentation attack detection: Where are we now? *Pattern Recognition Letters, 138*, 483–489. https://doi.org/10.1016/j.patrec.2020.08.018

Buolamwini, J. & Gebru, T. (2018). Gender shades: Intersectional accuracy disparities in commercial gender classification. *Proceedings of Machine Learning Research, 81*, 1-15.

Chapple, M., Stewart J. M., & Gibson, D. (2018). *Certified Information Security Professional* (8th ed). Wiley & Sons.

Chettri, K. (2020). Deep generative variational autoencoding for replay spoof detection in automatic speaker verification. *Computer Speech & Language, 63*, 101092–. https://doi.org/10.1016/j.csl.2020.101092

Ciampa, M. (2018). *Security+ Guide to Network Security Fundamentals*. (6th ed). Cengage.

Fang, C. (2020). Robust Iris Presentation Attack Detection Fusing 2D and 3D Information. *IEEE Transactions on Information Forensics and Security, 16,* 510–520. https://doi.org/10.1109/TIFS.2020.3015547

Grother, P., Ngan, M., & Hanaoka, K. (2019). *Face recognition vendor test (FRVT) Part 3: Demographic effects.* National Institute of Standards and Technology. https://doi.org/10/6028/NIST.IR.8280

Harwell, D. (2019, July 7). FBI, ICE find state driver's license photos are a gold mine for facial recognition searches. *The Washington Post*. https://www.washingtonpost.com/technology/2019/07/07/fbi-ice-find-state-drivers-license-photos-are-gold-mine-facial-recognition-searches/

Maloney, C. (2020). *Facial recognition technology (part III): Ensuring commercial transparency & accuracy.* House Committee on Oversight and Reform. https://oversight.house.gov/legislation/hearings/facial-recognition-technology-part-iii-ensuring-commercial-transparency#:~:text=On%20May%2022%2C%202019%2C%20the,of%20individuals%20across%20the%20country

Martinez-Martin, N. (2019). What are important ethical implications of using facial recognition technology in health care? *AMA Journal of Ethics, 21*(2) 180-187. https://doi.org/10.1001/amajethics.2019.180

Metz, R. (2020, September 9). *Portland passes broadest facial recognition ban in the US.* CNN Business. https://www.cnn.com/2020/09/09/tech/portland-facial-recognition-ban/index.html

Nesterova, I. (2020). Mass data gathering and surveillance: the fight against facial recognition technology in the globalized world. *SHS Web of Conferences, 74*, 3006–. https://doi.org/10.1051/shsconf/20207403006

Rattani, A. & Derakhshani, R. (2018). A survey of mobile face biometrics. *Computers and Electrical Engineering, 72*(2018), 39-52. https://doi.org/10/1016/j.compeleceng.2018.09.005

Romine, C. H. (2020). *Facial recognition technology*. National Institute of Standards and Technology. https://oversight.house.gov/sites/democrats.oversight.house.gov/files/documents/Romine%20Testimony.pdf

Shaban, H. & Flynn, M. (2019, April 23). *Teen sues Apple for $1 billion, blames facial recognition at stores for his arrest.* The Washington Post. https://www.washingtonpost.com/technology/2019/04/23/teen-sues-apple-billion-blames-facial-recognition-stores-his-arrest/

Sharif, M., Bhagavatula, S., Bauer, L., & Reiter, M. K. (2016). Accessorize to a crime: Real and stealthy attacks on state-of-the-art face recognition. *Computer and Communications Security*, (2016), http://dx.doi.org/10.1145/2976749.2978392

Snow, J. (2018, July 26). *Amazon's face recognition falsely matched 28 members of congress with mugshots.* ACLU. https://www.aclu.org/blog/privacy-technology/surveillance-technologies/amazons-face-recognition-falsely-matched-28

Stokes, E. (2020, November 19). *Wrongful arrest exposes racial bias in facial recognition technology*. CBS News. https://www.cbsnews.com/news/detroit-facial-recognition-surveillance-camera-racial-bias-crime/

Vintra claims less racial bias than Microsoft and Amazon. (2020). *Biometric Technology Today, 2020*(4), 2–3. https://doi.org/10.1016/S0969-4765(20)30044-8