Use of Facial Recognition by Groups Holding Positions of Power

A Research Paper submitted to the Department of Engineering and Society

Presented to the Faculty of the School of Engineering and Applied Science
University of Virginia • Charlottesville, Virginia

In Partial Fulfillment of the Requirements for the Degree
Bachelor of Science, School of Engineering

Edward Shen
Spring, 2021

On my honor as a University Student, I have neither given nor received unauthorized aid on this assignment as defined by the Honor Guidelines for Thesis-Related Assignments

Signature  Date 5/6/2021
Edward Shen

Approved  Date 5/4/2021

Hannah Rogers, Department of Engineering and Society

Abstract

Facial recognition is a relatively new technology that has found uses across many different parts of society. As with most new technologies integrating with society, its use raises some ethical concerns. Given enough time, these ethical issues eventually resolve, however, it does not without a significant effort to do so. This paper investigates ethical issues facial recognition faces with its integration with everyday society using a case study analysis. Facial recognition's use in healthcare, law enforcement, and education is analyzed with focus on the issues of bias, privacy, and trust. Through a case analysis of facial recognition technology, potential solutions targeting the primary ethical issues facial recognition faces are proposed. Findings show that current legislation inadequately addresses these issues, and further measures are needed to ensure facial recognition technology is as ethical as possible with its usage by those holding positions of power.

Use of Facial Recognition by Groups Holding Positions of Power

Introduction

Every high-end Apple or Samsung smartphone released after 2018 comes equipped with a feature called Face ID, a novel method that uses a user's facial features as a sign in tool for their smartphone. It is amazing to see a computer that fits into a pocket or a purse has the capabilities to recognize its owner based only on images of their face. But stop and consider the possibilities a computer a thousand times the size can do with the same technology. In the case of Robert Julian-Borchak Williams, he was unfortunate enough to experience the negative side of facial recognition technology. Williams is an African-American man wrongfully arrested based on a false positive from a facial recognition algorithm. (Hill, 2020, para 10) The only evidence used in his arrest was the camera footage of the crime and the false positive match that was determined by the police's facial recognition algorithm. The case against Williams was dropped without prejudice after his situation was clarified, however the time he spent in jail and the mental fatigue resulting from this situation stayed with him. It is a given that everyone who is an active member of society may fall victim to a faulty or mismatched facial recognition algorithm, and the likelihood of this happening increases as the use of facial recognition continues to increase. The use of facial recognition technology in society has raised the concerns of bias, privacy, and trust across multiple use cases, which need to be addressed if facial recognition is to have continued use. These concerns will be analyzed using actor network theory to determine potential solutions.

To define significant technical terms discussed in this paper, Computer Vision is defined as the field of study which focuses on the development of methods and techniques for computers to understand the content of digital media such as photos and videos. (Brownlee, 2019, para. 1)

Although humans can easily analyze and annotate an image, it is much more difficult for a computer to do the same. The general goal of computer vision is for a computer to extract useful information out of images/videos, which has been worked on by CV researchers for the past four decades (as cited in Brownlee, 2019, para. 30). Facial recognition is a subfield of computer vision, where a computer determines the identity of a person given an image of the person's face. Computer vision is not a new technology, with its origins dating back to more than forty years ago (as cited in Brownlee, 2019, para. 30). However, only in the past decade has developments in computing technology allowed its widespread usage in everyday life.

With many new technologies, people are wary of the consequences its introduction into society may have. In this case, those concerns may be warranted. Among the different applications computer vision has in society, one that has important and numerous concerns that needs to be addressed is facial recognition. The principle that a computer can identify a person is simple and straightforward, but the usage of facial recognition introduces concerns about its ethics. Because it's a relatively new and unique technology with rapid growth, it is hard to foresee backlash against its usage. For example, in an excerpt of Klosowski's (2020) paper:

> Facial recognition's first dramatic shift to the public stage in the US also brought on its first big controversy. In 2001, law enforcement officials used facial recognition on crowds at Super Bowl XXXV. Critics called it a violation of Fourth Amendment rights against unreasonable search and seizure. (para. 12)

It can be seen that the use of the technology in public was not well received. Because computer vision had minimal interaction with normal people in society, understanding the consequences is difficult without previous experiences. Now, there has been enough use that experts are aware of its usage and can analyze computer vision in context and predict ethical and societal concerns with its use.

Of course, computer vision also has its merits that that can be balanced with its consequences. In a scenario outlined by Kuflinski for catching criminals:

> Back in 2009, there were approximately 30 million surveillance cameras in use in the US, and now the number has grown exponentially. Imagine if these cameras, which reportedly captured an average of 4 billion hours of recording in a single week, were integrated with a facial recognition system. (para. 5)

Having a system that can pinpoint a criminal in a crowd of people anywhere in a city would greatly help in keeping a city safe. Violent offenders could be quickly identified and apprehended if a public facing facial recognition system detected them, making it more difficult for them to evade law enforcement. However, as seen with Williams, this technology does not always work as intended and can have significant consequences for those involved. There are many discussions where the government has the capability of tracking your every movement through technology, and these were generally seen in a negative light. Having cameras track your whereabouts wherever you go is what many consider a gross violation of personal privacy, an issue analyzed in this STS research.

Analysis of Facial Recognition Use Across Contemporary Society

Three primary issues with facial recognition use are examined in this paper: bias, privacy, and trust. These issues were chosen because they are not immediately apparent with facial recognition, yet these issues have significant impact on the user population facial recognition is targeted towards. Bias refers to how an algorithm may be more inclined to make decisions towards one social or ethnic group, privacy refers to breaches of our right to personal privacy, and trust is the relationship between the user of facial recognition and the system itself. These issues need to be studied in different use cases to understand the various impacts they have on user populations. Among the various different applications of facial recognition, the situations

that will be discussed focus on facial recognition use in healthcare, law enforcement, and education.

Healthcare

Begin by examining the application facial recognition technology has in healthcare. In recent years, facial recognition algorithms have been trained to be able to detect latent diseases based on a patient's facial features (Rigby, 2019, p121). By using a facial recognition algorithm, a medical practitioner can make better informed decisions to better aid their patient. However, this is strongly dependent on the accuracy of said facial recognition system. One fact about any machine learning system is they are never completely accurate, there will always be false positives and false negatives mixed in with accurate results. Then the issue arises when a medical practitioner makes a conflicting decision with the facial recognition algorithm. Would a patient trust the medical practitioner or the facial recognition algorithm? Even if they were able to choose, would they be able to completely trust either after they saw that they reached conflicting conclusions? Here the issue of trust is clearly outlined as a consequence of the application of facial recognition algorithms in healthcare.

Bias is also an important problem facing facial recognition in healthcare. With facial recognition, its ability to recognize certain features is not designed by engineers, rather engineers build the system to be able to train itself to recognize features on a large dataset. Consider a hypothetical scenario, the facial recognition system is provided with a dataset of one million images and corresponding features. After training the system, it is put to use in a hospital to use on real patients. The system works fine for some people, but for others it seems to make terrible decisions. Looking into the issue, you see that the system is extremely accurate on Caucasian males and significantly less accurate on other social groups. The dataset used to train the system

also contained mostly images of Caucasian males. The bias a facial recognition system has is not a result of any social construct, but a mere consequence of a biased dataset towards different social groups.

The issue of privacy needs to be considered in this use case as well. A person's face and the face template derived from it is considered biometric data and thus personally identifiable information (Martinez-Martin, 2019, p182). Currently, the Health Insurance Portability and Accountability Act (HIPPA) regulates patient health information and medical records, and includes privacy protections for personally identifiable information. (2019, p182). As a full facial photograph and facial template are considered biometric data, HIPPA would protect this data. However, not all existing statutes can be extended to protect privacy with facial recognition technology. From an excerpt of Martinez-Martin (2019):

> The Genetic Information Nondiscrimination Act (GINA) of 2008, for example, does not apply to FRT for genetic diagnosis, as FRT does not fit GINA's definition of genetic testing or genetic information. The Americans with Disabilities Act of 1990, which protects people with disabilities from discrimination in public life (e.g. schools or employment), would also likely not apply to FRT used for diagnostic purposes if the conditions diagnosed are currently unexpressed.

It is seen that existing legislation is insufficient to properly protect an individual's privacy, and further action is needed.

Law Enforcement

Looking back at the article by Hill (2020) and the case Williams faced when misidentified by a facial recognition algorithm, we see clearly how bias, privacy, and trust issues are present with facial recognition's application in law enforcement. Recent studies by the Massachusetts Institute of Technology and the National Institute of Standards and Technology have found that these facial recognition algorithms used in law enforcement perform relatively

well on white men, but are less accurate for other demographics. (Hill, 2020, para. 12) In research performed by Raji et.al., they found that facial recognition algorithms by Microsoft, Amazon, and Clarifai all performed the worst on darker and/or female subgroups. (2020, p146) The simple fact that it works better for one social demographic over other demographics makes it biased against certain groups, as they are more likely to be misidentified as evidenced by Williams. This inherent inaccuracy of the algorithm naturally discriminates against non-white male demographics, with errors of up to 35% (Kuflinski, 2019, para 20). Moreover, bias is present starting from the dataset curation process. When curating a dataset, data from a specific population may be supplemented or highlighted that was previously underrepresented in other datasets. "Efforts to increase representation of this group can lead to tokenism and exploitation, compromise privacy, and perpetuate marginalization through population monitoring" (Raji et. al., 2020, p148). Facial recognition technology has issues regarding demographic and social group bias from the first stage of collecting data, and pervades up to the application of a finished algorithm.

The next issue of privacy is a major concern regarding facial recognition's use in law enforcement. The most obvious privacy issue with facial recognition use in law enforcement is surveillance. When facial recognition algorithms are capable of identifying individuals in large crowds, the question arises of how far facial recognition should be allowed to go in surveillance. On one side, facial recognition could be used to identify people with warrants when they are spotted on a surveillance camera. On the other, the vast majority of people who are law abiding citizens will be monitored whenever they are seen on a surveillance camera. As stated by Andrejevic and Selwyn, knowing where people go provides intimate and wide ranging information about their personal, professional, and leisure lives. Keeping track of where a person

goes, such as a medical clinic or a casino, can reveal highly sensitive personal information. (2020, p117) This "big brother" watchdog is a usual criticism of facial recognition (Kuflinski, 2019, para 24). Just as the US population was outraged when Edward Snowden revealed the NSA was spying on its citizens through phone calls and telephone records, the same would apply through cameras collecting personal data without user consent.

Another privacy issue with facial recognition in this case is the collection of data. As mentioned earlier, the accuracy of a facial recognition algorithm is dependent on the size and quality of the dataset used to train it. For an algorithm to perform well, vast amounts of data are needed, in the case of facial recognition, pictures of people's faces. Thus, in training the algorithm alone, it may be possible that personal information is part of the dataset used to train the facial recognition algorithm. From Raji et.al., "while audit benchmark datasets should reflect the populations who will be impacted by the audited technology, collecting a sufficiently large and diverse dataset can present privacy risks for the individuals represented in the dataset" (2020, p148). Furthermore, depending on the data storage method and data dissemination policies, sensitive and biometric data could be made accessible outside of its original purpose in the facial recognition technology (2020, p148). However, even after an algorithm is trained, the privacy concerns do not stop. The concept of data storage can apply for a finished facial recognition algorithm. Once a person's image is used when suspicions of a crime are present, the image may be stored alongside other suspected individuals without consent, regardless if the individual was actually involved in any criminal activity or not (Kuflinski, 2019, para 27).

Examining trust in a facial recognition algorithm in the context of its use in law enforcement, it is seen in a survey performed by Smith that the majority (56%) of Americans trust law enforcement to use facial recognition technology. (2019) However, taking a simple

inverse of this statement, we see that 44% do not trust law enforcement to use facial recognition technology responsibly. Even more interesting is the responses to this survey based on demographic. After conducting this survey, Smith found that a substantially smaller share of young adults think it is acceptable for law enforcement to use facial recognition, likewise a smaller share of African American and Hispanic adults believe it is acceptable for law enforcement to use facial recognition when compared to white adults. As mentioned earlier in this paper, facial recognition algorithms tend to perform worse on women and people of color. There is indeed a positive correlation to the relative performance of facial recognition algorithms and the respective demographics, but there is not enough information or studies done that show a biased facial recognition algorithm lessens trust with it or vice versa. However, given cases such as the one documented by Hill with Williams falsely identified and arrested due to a false positive match from a facial recognition algorithm, it is reasonable to assume that these such cases are factors taken into consideration when marginalized demographics report their trust in law enforcement to use facial recognition responsibly.

Education

Facial recognition technology has pervaded many aspects of society, and schools were not spared from facial recognition's influence. There are now various applications of facial recognition, ranging from campus security systems, automated attendance, and even attention monitoring.

The more prominent application in schools of facial recognition technology is campus security. These systems are marketed to schools in the United States as an "all seeing shield against school shootings" (Andrejevic & Selwyn, 2020, p118). These systems are capable of identifying permitted or not permitted individuals on the premises, and can keep track of the

whereabouts of individuals detected. An extension of this use is attendance monitoring, where a facial recognition algorithm is simply used to verify attendance in a class. Another use that is more relevant given stay at home schooling is the use of facial recognition in virtual classrooms. Facial recognition is used in virtual classrooms to ensure only authorized individuals have access to online educational content as well as to proctor electronic assessments in place of a human proctor. The final use facial recognition sees in education as mentioned by Andrejevic and Selwyn is to monitor engagement levels among students. The primary use of this technology in monitoring engagement is to understand "the learners state of mind," and can highlight problems with knowledge, stimulation, anxiety, and/or frustration (Andrjevic & Selwyn, 2020, p118-119).

Looking into the technology behind facial recognition technology and considering the consequences of these applications in education, emerges the issue of bias once again. Facial recognition plays a role in foregrounding fixed attributions of a students race and gender in informing school decision making (Andrjevic & Selwyn, 2020, p119). When used this way, facial recognition arbitrarily divides social groups further. Furthermore, grouping students by their facial features is a discriminatory practice where a student is categorized by their biological characteristics and not their character, personality, or ability. Again, the problem of facial recognition technology performing better on certain demographics compared to others persists in education. When a facial recognition algorithm is attempting to detect engagement levels in a student, it may flag someone from a less predominant demographic in the dataset as inattentive or bored, when in reality they are and the facial recognition algorithm was just incapable of detecting it. Thus, if facial recognition was used inappropriately, certain demographics could be negatively impacted in an educational environment. Additionally, these facial recognition algorithms could be applied differently in different schools. Schools that have higher incidents

could use facial recognition more aggressively compared to schools that have fewer incidents. These circumstances could be magnified by other factors like funding, amplifying the bias involved with using facial recognition in schools.

With facial recognition in schools, student privacy is another topic that needs to be considered regarding facial recognition. Consider a facial recognition proctor when taking an online examination. Students may not have the option to opt out of such systems in fear of not passing the class, and are coerced into accepting terms of having a facial recognition proctor when taking an exam. Facial recognition algorithms may also incorporate eye tracking, where it can detect where a student is looking during an exam. Students who naturally look at other spaces when taking exams at school would no longer have that option and this would add extra burden on them in an already stressful environment of taking an exam. Additionally, facial recognition systems remove a student's ability to "lay low" or go under the radar at school. At a glance, this could be considered an undesirable behavior, but for some students it is a coping strategy or their way of going to school (Andrjevic & Selwyn, 2020, p-119). Being able to choose how one puts themselves out to society is a part of individual identity, and facial recognition algorithms take this ability away from students.

Trust in facial recognition use in schools draws very similar parallels to that of which was discussed in the previous healthcare and law enforcement sections. Once again, the largest concern with facial recognition is if it will work consistently for all groups subject to its use. If an engagement detector consistently performs worse on certain groups, then that group would not trust the algorithm to give their students an unbiased education. Likewise, with school surveillance systems, can groups trust the system to not infringe on their student's privacy? An issue with schools that is not present with healthcare or law enforcement is the primary group it

is used on is students. As most students are minors, they are not provided the right to vote, thus do not receive representation on issues they deem important. If the primary group facial recognition targets does not have the power to protect themselves, it is extremely difficult to ensure their privacy will be protected.

Legislation

Legislation regulating the use of facial recognition technologies are starting to emerge. For example, the Commercial Facial Recognition Privacy Act of 2019 attempts to resolve privacy issues with commercial facial recognition use. This bill prohibits commercial organizations from collecting or using user information without documentation of their facial recognition technology and the explicit consent of the user. However, the bill does not address the use of facial recognition in a non-commercial setting, i.e., in healthcare, law enforcement, and education. Another piece of legislation that begins to regulate facial recognition technologies is the Illinois Biometric Information Privacy Act. This act does not specifically target facial recognition, however, it classifies faces as biometric information, thus protecting an individuals privacy against facial recognition algorithms.

These pieces of legislation are making an attempt at protecting relevant social groups from the issues facial recognition has when implemented in society, but still lacks the coverage needed to address the issues of bias, privacy, and trust outlined in this paper. For bias, many of the issues in healthcare, law enforcement, and education stem from an algorithm not being able to be trained to have equal performance on different demographics. Thus, to address the root cause of bias, it would be best to remove bias from the technology itself. This would be a task for engineers to take on, as they are the one making progress in facial recognition technologies. For privacy, the Commercial Facial Recognition Privacy Act of 2019 and the Illinois Biometric

Information Privacy Act are steps in the correct direction to protect an individual's privacy. These pieces will need to be extended to broader use cases in order for facial recognition to integrate closer with society. Once both of these issues are acted upon the public's trust in facial recognition technology will follow as the issues that stem from a biased algorithm or a lack of personal privacy are addressed.

Conclusion

Facial recognition technology has expanded to wide varieties of uses recently, with analyses of its usage lagging behind. The issues of bias, privacy, and trust were identified across multiple use cases and studied to find overarching themes facial recognition has when applied in modern day society. These themes are used to provide insight into possible solutions to the ethical issues facial recognition faces with its real-world applications. Existing legislation inadequately addresses these issues, resulting in the need for new legislation to remedy these issues. Only when these issues are addressed can facial recognition technology integrate itself in society without facing significant ethical issues.

Further studies can be done as only three different use cases were focused on in this research paper. Perhaps new ideas could be revealed offering counterpoints to the ones made with healthcare, law enforcement, and education. Because the applications of facial recognition technology are vast, studying them all is not practical. These three cases were chosen due to their high impact of its use and its power to affect an individual's life. Furthermore, there may be issues aside from bias, privacy, and trust that facial recognition technology has in society, and were not addressed in this research paper. Facial recognition technology is seeing more widespread usage, and measures need to be taken to protect the people involved. Engineers can innovate new techniques to reduce the problem of bias build into machine learning technology

and lawmakers can pass legislation protecting individual privacy, and trust will follow once people see their rights are protected and the algorithms treat everyone fairly.

REFERENCES

Andrejevic, M., & Selwyn, N. (2020) Facial recognition technology in schools: critical questions and concerns. *Learning, Media and Technology, 45*(2), 115-128, doi:10.1080/17439884.2020.1686014

Brownlee, J. (2019, March) A gentle introduction to computer vision. In *Machine Learning Mastery.* Retrieved from: https://machinelearningmastery.com/what-is-computer-vision/

Fioravanti, C., Velho, L. (2010) Let's follow the actors! Does actor-network theory have anything to contribute to scientific journalism? *Journal of Science Communication, 9(*4). 1-8, doi:10.22323/2.09040202

Hill, K. (2020, June), Wrongfully accused by an algorithm. In *The New York Times.* Retrieved from: https://www.nytimes.com/2020/06/24/technology/facial-recognition-arrest.html

Klosowski, T. (2020, July) Facial recognition is everywhere. In *The New York Times.* Retrieved from: https://www.nytimes.com/wirecutter/blog/how-facial-recognition-works/

Kuflinski, Y. (2019, April) How ethical is facial recognition technology? In *Towards Data Science.* Retrieved from: https://towardsdatascience.com/how-ethical-is-facial-recognition-technology-8104db2cb81b

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, doi:10.1001/amajethics.2019.180

Mihajloic, L. (2019, April) Everything you ever wanted to know about computer vision. In *Towards Data Science.* Retrieved from: https://towardsdatascience.com/everything-you-ever-wanted-to-know-about-computer-vision-heres-a-look-why-it-s-so-awesome-e8a58dfb641e

Peregud, I. & Zharovskikh, A. (2020, August) Computer vision applications examples across different industries. In *In Data Labs.* Retrieved from: https://indatalabs.com/blog/applications-computer-vision-across-industries

Raji, I. D., Gebru, T., Mitchell, M., Buolamwini, J., Lee, J., & Denton, E. (2020) Saving face: investigating the ethical concerns of facial recognition auditing. In *AIES '20: Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society,* 145-151 doi:10.1145/3375627.3375820

Rigby, M. J. (2019) Ethical dimensions of using artificial intelligence in health care. *AMA Journal of Ethics, 21*(2). 121-124, doi:10.1001/amajethics.2019.121

Simonite, T., Barber, G. (2019, October) The delicate ethics of using facial recognition in schools. In *Wired.* Retrieved from: https://www.wired.com/story/delicate-ethics-facial-recognition-schools/

Singh, R. (2019, July) Recent advances in modern computer vision. In *Towards Data Science.* Retrieved from: https://towardsdatascience.com/recent-advances-in-modern-computer-vision-56801edab980

Smith, A. (2019, September) More than half of U.S. adults trust law enforcement to use facial recognition technology responsible. In *Pew Research Center.* Retrieved from: https://www.pewresearch.org/internet/2019/09/05/more-than-half-of-u-s-adults-trust-law-enforcement-to-use-facial-recognition-responsibly/