An Analysis of Facial Recognition in Black Lives Matter Protests

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

Nikash Sethi
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 April 20, 2021
Nikash Sethi


Approved Date April 19, 2021
Richard Jacques, Department of Engineering and Soc

INTRODUCTION

Perhaps the most promising and fascinating innovation within artificial intelligence in the past few decades is computer vision. State-of-the-art models and algorithms can classify images they have never seen before with incredibly high accuracy, and researchers continue to implement and deploy interesting applications of the technology into society. A particularly interesting problem within computer vision is zero-shot learning, in which models classify categories of images that have not been trained at all (Wang et al., 2019, p. 1). Few-shot learning is a technique closely related to zero-shot learning and involves scenarios in which only a few image features are available as labeled examples during training (Rahman et al., 2018, p. 5655). Dickson (2020) argues that the most influential application of few-shot learning is facial recognition (para. 4). In this application, algorithms learn representations of faces with few data points and can classify faces based on these embeddings.

Facial recognition has come under heavy scrutiny over the past year, as police departments around the country started using the technology to track down activists who participated in Black Lives Matter protests following the death of George Floyd last summer (Rector & Winton, 2020, para. 1). The relationships between police departments, activists, legislators, and technology companies that develop facial recognition algorithms have become increasingly complicated, and tensions surrounding whether facial recognition should be used in policing have grown considerably. Therefore, it is important to objectively analyze the parties involved and how they interact with each other to determine potential solutions to the controversy.

My STS research tightly couples with the technical research, as both look to advance applications of computer vision and specifically zero-shot learning. I focus on analyzing the

various actors surrounding facial recognition in policing through an Actor-Network Theory framework that Callon (1984) and Latour (1996) pioneered. In this paper, I present an objective overview of the controversy of deploying facial recognition in police departments, which points toward potential solutions or regulations we can apply on the technology that benefit society as a whole.

AN OVERVIEW OF FACIAL RECOGNITION IN POLICING

Advancements in computer vision technologies, especially those in zero-shot and few-shot learning, have resulted in highly accurate facial recognition technologies that have several industrial applications. One of the most popular and recently most controversial applications of this technology is in policing. During Black Lives Matter protests, police departments across the country used facial recognition technologies to track potential criminals (Rector & Winter, 2020, para. 1). This practice has angered activists and propelled legislators to advocate for stricter regulation surrounding the industrial use of facial recognition (Vincent, 2020, para. 5). My STS research project analyzes the use of facial recognition in policing through the Actor-Network Theory framework that Latour (1996) helped define, in order to explore the complex and intricate relationships between various actors and entities involved in this controversy (p. 380).

Industrial applications of facial recognition first became prevalent when American military and intelligence used facial recognition in Iraq and Afghanistan to identify potential terrorists (Williams, 2015, para. 1). Williams (2015) found that police departments across America began adopting and deploying facial recognition around 2010 to identify and track criminals (para. 10). At the time, many lawmakers argued that tracking criminals and identifying suspects using facial recognition technologies were far more efficient than using traditional techniques like fingerprint matching (Williams, 2015, para. 2).

More recently, though, these facial recognition technologies have disproportionately harmed and targeted minorities (Rector & Winton, 2020, para. 20). Fitch (2019) reported on research done by the National Institute of Standards and Technology that found that industry-leading facial recognition technologies misidentify Asian and African Americans far more than Caucasians (para. 3). This report examined 189 facial recognition algorithms made by 99 companies, which is considered to be most of the industry (Fitch, 2019, para. 5). Fowler (2020) worries that while the technology may help against criminals escaping arrest, it could also "open a slippery slope to a world of supercharged policing that's likely to disproportionately impact people of color" (para. 7). Given recent events surrounding Black Lives Matter movements, many activists are lashing out against unfair policing tactics, which has only compounded opposition against the use of facial recognition technologies in policing.

UNRAVELING THE CONTROVERSY

A deeper look into this controversial application of facial recognition reveals several actors and entities with different perspectives and motivations toward the use of the technology. Parties involved in this intricate network demonstrate varying levels of support for facial recognition applications in policing. Further, some actors publicly display perspectives contradictory to their practices and actions. The complexity of relationships between different groups that are involved in the controversy of deploying facial recognition for policing demands the need for an in-depth analysis of the network that these actors form.

My STS research analyzes this intricately related network and the relationships between each party involved. Beyond understanding the perspectives, attitudes, and goals of each actor, I identify how the actors in this space interact and lay out public and hidden relationships that they have with each other. Ultimately, my work unravels these complicated relationships and presents

a clear view on the environment surrounding the application of facial recognition technologies in policing. To do so, I rely heavily on applying an Actor-Network Theory framework to the controversy.

The Actor-Network Theory framework was pioneered by Latour (1996) and Callon (1984), and has been applied on various societal problems and dilemmas related to advancements in technology. Actor-Network Theory is used to understand the role of science and technology in structuring power relationships (Callon, 1984, p. 197). The framework describes societies in a structured and systematic way, by delimiting the identity of actors and the interactions they have (Latour, 1996, p. 370). By interpreting facial recognition technology and its applications in policing through an Actor-Network Theory perspective, we can begin identifying potential solutions and regulations that mutually benefit all involved parties. Latour (1996) argues that applications of Actor-Network Theory offer an objective, black-and-white approach to analyzing complex and intricate networks, which makes it perfectly applicable to study the use of facial recognition in policing (p. 380).

Once we apply the Actor-Network Theory framework to the application of facial recognition in policing, we can begin understanding the perspectives and relationships that each entity maintains in the larger networks they are a part of. Figure 1 overviews four of the major actors in this network, their primary goals, and their support for or opposition to the use of facial recognition technologies in policing. Police departments are the primary users of facial recognition technologies in this context and have therefore come under heavy scrutiny from activists and lawmakers. Activists advocate for stricter regulation against facial recognition from legislators. These legislators seek to regulate the use and misuse of facial recognition technologies in applications both inside and outside policing. Finally, technology companies who

produce and sell facial recognition algorithms to police departments influence and react to the actions of legislation around facial recognition. The following section of this paper will dive deeper into each of these actors to clarify this complex and intricate network.

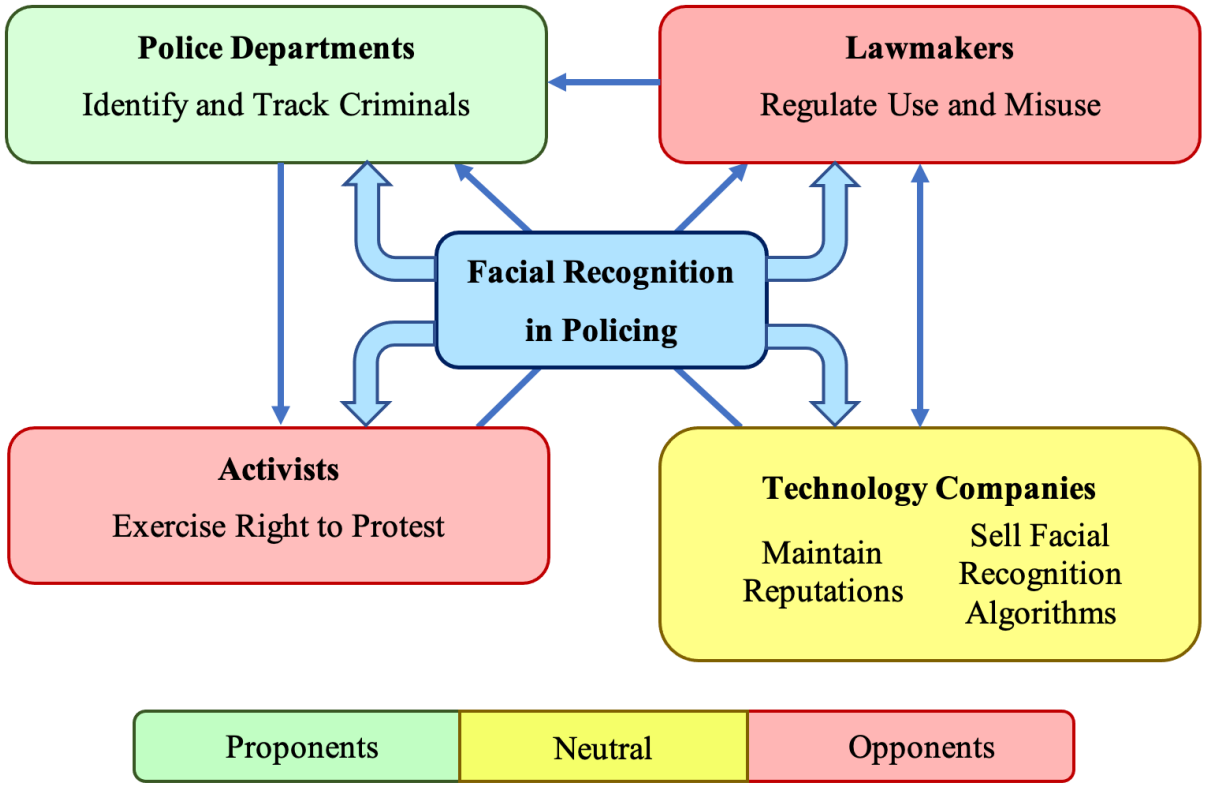

Figure 1: Actor Perspectives on Facial Recognition. Various actors have conflicting motivations and main goals related to facial recognition technologies in policing, which has deterred regulation on this practice from being widely accepted. (Sethi, 2020).

DEFINING THE ACTORS INVOLVED

Police departments across the country rely on facial recognition for identifying and tracking suspects. The Los Angeles Police Department has used facial recognition software almost 30,000 times since November 2009 (Rector & Winton, 2020, para. 1). The New York Police Department processed 9,850 requests to use the software in 2019 alone (Vincent, 2020, para. 6). This reliance continued into the Black Lives Matter protests during the summer of 2020. In one incident, Vincent (2020) reported that the New York Police Department used facial

recognition to track down the address of Derrick Ingram, an activist who was accused of assault during a Black Lives Matter protest (para. 2).

When using facial recognition, police departments typically run images of suspects from surveillance cameras against a database of legally obtained mugshots (Rector & Winton, 2020, para. 1). That said, some algorithms commonly used by police departments around the country collect data by scraping images off of online social media pages (Council, 2021, para. 7). Many believe police use of such algorithms crosses ethical boundaries, and some police departments such as the Boston Police Department decided not to deploy facial recognition practices altogether (Williams, 2015, para. 21).

Police departments obtain these facial recognition products from technology companies. Amazon, Microsoft, and IBM, each large-scale technology companies that research cutting-edge facial recognition algorithms, have all publicly stated they will stop selling facial recognition technologies to police departments (Fowler, 2020, para. 11). According to Fitch (2019), IBM even released a collection of images of people of different races and genders to help train algorithms in less biased ways (para. 9). Technology companies that release public datasets and advocate for transparency in facial recognition methodology believe that their action will decrease bias in the algorithms that police departments use.

That said, legislators argue that these public statements hold little promise and are merely tactics to maintain a good reputation. While these technology companies appear to stand against facial recognition, Fowler (2020) believes they are hesitant to deny the business opportunities of selling the technologies to police departments (para. 22). After investigating companies like Microsoft, Fowler found that technology companies do not actually regulate the technology as it is used and instead only provide vague guidelines and safety concerns when selling their

algorithms (para. 30). Despite their voluntary commitments, large technology companies may actually slow legislative efforts to ban the use of facial recognition in policing (Fowler, 2020, para. 4).

It is important to note that large-scale, household name technology companies such as Amazon and Microsoft only account for a minor portion of the facial recognition products used by police departments (Fowler, 2020, para. 11). The companies that develop and deploy most of the facial recognition technologies in police departments in America, including NEC Corporation and Clearview AI, have not publicly stated that they will make the same commitments as Amazon and Microsoft (Fowler, 2020, para. 11). Clearview AI, responsible for selling their algorithms to about 2,400 law enforcement agencies in America, claim that their models are trained on data that is lawfully obtained, and have no plans to stop providing their technologies to police departments (Council, 2021, para. 7). Recently, activist groups have sued Clearview AI for violating privacy rights of California citizens, arguing their technology makes individuals vulnerable to being tracked or targeted for their activism (Council, 2021, para. 1)

The contrast between "Big Tech" companies like Microsoft and the lesser-known main players in the facial recognition space has led media and society to place incorrect focus and awareness on actors that are not directly involved in selling facial recognition algorithms. Fowler (2020) argues that if NEC Corporation and Clearview AI continue providing the technology to police departments and public attention does not shift to these companies, new legislation will be the only way to stop facial recognition in policing (para. 14).

Advocators and legislators generally push for broad federal regulations to govern applications of facial recognition. Figure 2 outlines the current state of legislation on facial recognition technologies and the level of regulation that lawmakers advocate for. Recent

legislative efforts are limited to local and state legislation and do not cover the full scope of facial recognition technologies (Learned-Miller et al., 2020, p. 3). In a white paper directed to legislators, Learned-Miller et al. (2020) call for the creation of a new federal office to regulate facial recognition (p. 3). Modeled after the Food and Drug Administration, this federal office would standardize the use of facial recognition and hold developers and deployers accountable (Learned-Miller, 2020, p. 4). Still, some lawmakers believe facial recognition is far more effective than fingerprinting at identifying criminal suspects and are therefore hesitant to restrict or ban its use (Williams, 2015, para. 2). With a better understanding of the controversy, we can start to hypothesize regulations and solutions that best benefit all parties involved.

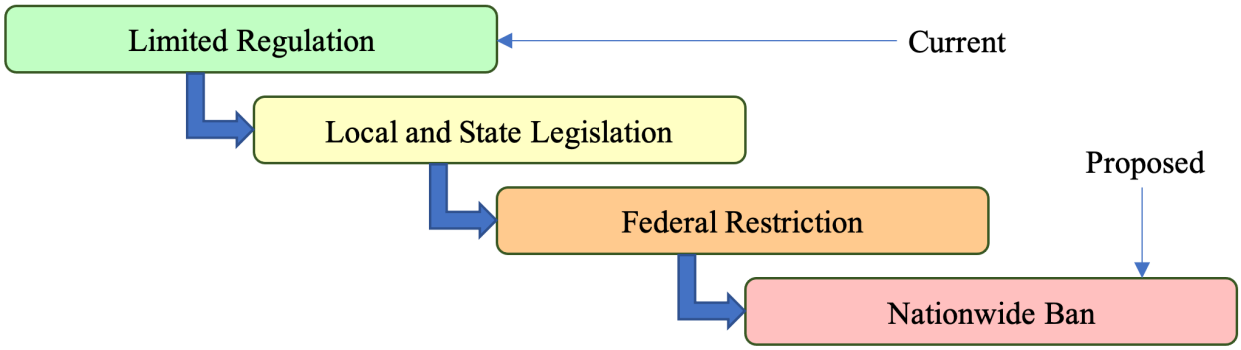

Figure 2: Levels of Facial Recognition Regulation. While the current regulation on facial recognition is limited and varies across different regions, activists and lawmakers advocate for stricter laws that ban the technology from being used in policing altogether. (Sethi, 2020).

REGULATING FACIAL RECOGNITION: A CASE STUDY

Across the country, cities including San Francisco, Portland, and Minneapolis have enacted bans on police use of facial recognition, largely due to the bias that such algorithms may introduce (Hill, 2021, para. 2). This dramatic action has received mixed reactions from activists advocating for regulation against facial recognition and police departments who use the technology regularly. Recently, Massachusetts became the first state to pass legislation regulating facial recognition; legislation that does not take an all-or-nothing approach to

regulating the technology (Hill, 2021, para. 2). Analyzing this police reform bill as a case study for broader state-wide legislation provides us with valuable insight into a novel approach into regulating facial recognition that I propose would extend well nation-wide.

Politicians and police departments in the state of Massachusetts were largely against banning police use of facial recognition altogether because of the proven efficiency and effectiveness of the technology. As Hill (2021) notes, Massachusetts police departments used facial recognition to solve two cases of homicide and child sexual abuse, which made the idea of banning the technology "politically impossible" (para. 23). Therefore, the new police reform bill, going into effect in July, takes a much more marginal approach. While still allowed to use facial recognition, police officers must now get a judge's permission before running a search, and then have someone from state or federal police perform the search for them (Hill, 2021, para. 5). This process contradicts the previously relaxed approach to running facial recognition searches, in which officers could simply download a facial recognition model and perform the search themselves without external approval (Vincent, 2020, para. 6).

I propose that federal lawmakers should install a system similar to the one recently enacted in Massachusetts. Police officers across the country should be required to get approval from external, experienced entities before performing their facial recognition searches. As Learned-Miller et al. (2020) propose in their white paper, a heavily standardized federal office created to carry through facial recognition searches would prevent misuse and abuse of the technology while maintaining the known benefits that it provides. Importantly, widescale efforts by technology companies providing facial recognition algorithms will still have to be completed in order to mitigate bias in these models, but this federal legislation would be a great unified first step toward ensuring the technology is being used cautiously and properly.

ETHICALLY ADVANCING APPLICATIONS OF COMPUTER VISION

Advances in computer vision techniques have led to large-scale social controversies in the application of facial recognition in policing. Various actors and relationships are at play in this dilemma, which emphasizes the need for an objective analysis of the goals and perspectives of police departments, technology corporations, activists, and legislators. Through my analysis, I propose a legislative approach to regulating facial recognition that is not all-or-nothing, mirroring a case study of Massachusetts's new police reform bills. Although such an approach still allows police departments to use potentially biased facial recognition algorithms, it is an important step toward mitigating the disproportionate targeting of minority groups currently impacted by facial recognition in policing.

With the development of any technology, we must consider social and ethical implications in parallel with innovation. As an engineer, I cannot disregard the non-technical aspects and implications of my technical research in zero-shot learning. Ultimately, the tight coupling between my computer vision research in zero-shot learning and my STS research in facial recognition applications in policing serves as a testament to conscious engineering and allows me to innovate in the right way.

References

Callon, M. (1984). Some elements of a sociology of translation: domestication of the scallops and the fisherman of St Brieuc Bay. *The Sociological Review*, *32*(1), 196-233. doi:10.1111/j.1467-954X.1984.tb00113.x

Council, J. (2021, March 10). Facial recognition tool used by police faces civil lawsuit in California. *The Wall Street Journal*. Retrieved from https://www.wsj.com/

Dickson, B. (2020, August 12). What is one-shot learning? *TechTalks*. Retrieved from https://bdtechtalks.com/

Hill, K. (2021, February 27). How one state managed to actually write rules on facial recognition. *The New York Times*. Retrieved from https://www.nytimes.com/

Fitch, A. (2019, December 19). Facial-recognition software suffers from racial bias. *The Wall Street Journal*. Retrieved from https://www.wsj.com/

Fowler, G. A. (2020, June 12). Black Lives Matter could change facial recognition forever – if Big Tech doesn't stand in the way. *The Washington Post*. Retrieved from https://www.washingtonpost.com/

Latour, B. (1996). On actor-network theory: a few clarifications. *Soziale Welt, 47*(4), 396-381. https://www.jstor.org/stable/40878163

Learned-Miller, E., Ordóñez, V., Morgenstern, J., & Buolamwini, J. (2020). *Facial recognition technologies in the wild: a call for a federal office.* Retrieved from The Algorithmic Justice League website: https://bit.ly/34GOzou

Rahman, S., Khan, S., & Porikli, F. (2018). A unified approach for conventional zero-shot, generalized zero-shot, and few-shot learning. *IEEE Transactions on Image Processing, 27*(11), 5652-5667. doi:10.1109/TIP.2018.2861573

Rector, K. & Winton, R. (2020, September 21). Despite past denials, LAPD has used facial recognition software 30,000 times in the last decade, records show. *The Los Angeles Times*. Retrieved from https://www.latimes.com/

Sethi, N. (2020). *Actor perspectives on facial recognition.* [Figure 1]. *Prospectus* (Unpublished undergraduate thesis). School of Engineering and Applied Science, University of Virginia. Charlottesville, VA.

Sethi, N. (2020). *Levels of facial recognition regulation.* [Figure 2]. *Prospectus* (Unpublished undergraduate thesis). School of Engineering and Applied Science, University of Virginia. Charlottesville, VA.

Vincent, J. (2020, August 18). NYPD used facial recognition to track down Black Lives Matter activist. *The Verge*. Retrieved from https://www.theverge.com/

Wang, W., Zheng, V., Yu, H., & Miao, C. (2019). A survey of zero-shot learning: settings methods, and applications. *ACM Transactions on Intelligent Systems and Technology, 10*(2), 1-37. doi:10.1145/3293318

Williams, T. (2015, August 12). Facial recognition software moves from overseas wars to local police. *The New York Times*. Retrieved from https://www.nytimes.com/