Tech Fail: The Hidden Bias In Facial Recognition Systems

Sophia M. Harrison

Abstract:

In less than twenty years, artificial intelligence has evolved from a vague science fiction concept to an indispensable tool with broad real-world applicability. The field has enjoyed significant advances in recent years - some of the most promising of which have occurred in the area of facial recognition. As a result, the technology is now more accurate than ever before, making it an attractive resource for both the commercial and public sectors. However, the accuracy of facial recognition software can vary significantly, depending on the type of human subjects being analyzed. In fact, such systems frequently produce false positives and incorrect gender tags when presented with darker complexion subjects. An increasingly common part of the technology landscape, particularly in the area of law enforcement, the expanding prevalence of this new tool, necessitates the need to address, and remedy, its existing racial and gender biases.

Introduction

Conceived in the 1950s, facial recognition software is now used in everything from mobile apps to restaurants and even hospitals. Powerful though it may be, this technology is not without its flaws - the most concerning of which is racial and gender bias. However, despite this critical flaw, many software vendors have proclaimed the technology and the algorithms behind it, to be "highly reliable". When a facial recognition system fails in a commercial setting, consumers suffer a less-than-ideal user experience but are unlikely to endure any lasting negative consequences. But businesses aren't the only ones using this technology.

Dozens of federal, state and local law enforcement agencies throughout the U.S. rely on facial recognition software for the surveillance and apprehension of suspected criminals. A system error in such circumstances could threaten the rights and safety of innocent people, especially darker-complexion minorities, whom facial recognition programs all-too-often struggle to correctly gender tag or match with existing photos (i.e. produce a true positive). These harmful biases plus the overrepresentation of

African-Americans and Latinos in police encounters and, subsequently, facial recognition databases, necessitates swift and immediate action to prevent the violation of citizens rights and the rule of law. This paper looks at how and when facial recognition systems fail, the consequences of this failure and why it exists in the first place. It presents several actionable solutions (some of which are already being applied) for improving the accuracy of facial recognition systems and raising the standard of accountability for the vendors who create them.

How Facial Recognition Works

The term "facial recognition" refers to an algorithm or application designed to perform facial analysis and or facial recognition. Researchers and software vendors often utilize publicly available photo databases to train their algorithms. Generally speaking, the more images an algorithm has to train on the better it learns to correctly identify a face and its respective metrics. It's not uncommon for data scientists and computer programmers to use hundreds, or even thousands, of images to train an algorithm. Both the quantity and quality of data using during this training process play a critical role in the overall reliability of the software it will support and, in turn, whether or not the system will be biased, and if so, to what degree and under which circumstances.

Each human face is different. Some of these differences are easy to spot: skin color, head size, and eye shape being some of the most obvious. Others, however, are a lot less perceptible to the human eye. Known as a geometric approach, many facial recognition algorithms are designed to capture these differences, using attributes like the distance between someone's eyes, the length of their nose, or the width of their mouth to mark one face from another. A few programs are even sophisticated enough to differentiate between identical twins. Once all of the necessary metrics (as defined by the creator of the algorithm and its intended use) are identified, a "faceprint" or face template is created.

Some recognition algorithms use a photometric process that statistically normalizes facial images, distilling them down to only the most essential recognition data, which is then assigned a mathematical value and compared with facial templates to eliminate any statistical variance.

While not yet as common as other methods, 3D facial recognition technology is gaining traction. Unlike geometric or photometric methodologies, 3-D recognition is unaffected by changes in lighting and can more accurately capture facial metrics from different angles. However, it does challenges when confronted with non-neutral facial expressions - a shortcoming common in many facial recognition systems.

A program that is designed to recognize and catalog faces will process and store any newly detected facial data in a database or repository, to be accessed at a later date. Applications that perform facial matching take any identified markers and either compare them to a given photo of a known subject or run them against a database of other facial images, in search of a match. Several potential matches and their corresponding statistical likelihood of being a match, are presented for review. A fully functioning facial recognition system can be used for a diverse range of purposes including, but not limited to, emotion detection (recognizing a person's emotional state), age detection, gender detection, race/ethnicity detection and attention measurement (via eye movement tracking).

Who's Using Facial Recognition and Why

Mobile apps are the most visible application of facial recognition software today, but the technology can also be found in a variety of non-technical sectors, including HR, retail, and even the fast food industry. But no sector is more reliant on this tool than security and law enforcement. In fact, Market Watch analysts predict that the worldwide market for facial recognition technology will reach $9.6 billion by 2022, with North American security firms and government agencies the primary users.[1]

Law Enforcement

The FBI has used facial recognition for surveillance and tracking for decades and both the CIA and N.S.A. have been known to use the technology in the course of their daily operations - harvesting photos from social media platforms and intercepted online communication. A 2016 report by the Privacy & Technology Center at Georgetown Law, revealed that nearly half the US population, or approximately 150 million people, have been unknowingly cataloged in either the FBI's facial recognition database or one of the dozens of others managed by state, county or local law enforcement agencies.[2]

But the U.S isn't the only country where surveillance is becoming the norm. The U.K., France, China, the Australian Federal Police, and the Swedish Migration Agency also use facial recognition for everything from general surveillance to security

[1] Source: MarketWatch. (2016, June 28). *Facial Recognition Market Expected to Reach $9.6 Billion, Worldwide, by 2022 [*Press release*].* Retrieved from https://www.marketwatch.com/press-release/facial-recognition-market-is-expected-to-reach-96-billion-worldwide-by-2022-2016-06-29-8203

[2] Source: "The Perpetual Line-Up: Unregulated Police Face Recognition In America", Center on Privacy & Technology at Georgetown Law, October, 18, 2016.

at airports, sports stadiums, and public transit hubs. With more than 5.9 million surveillance cameras, the U.K is the most surveilled country in the world. And, with the use of facial recognition only governed by a "code of conduct", British law enforcement agencies have a width breadth when it comes to how it is applied.[3]

Amazon Rekognition

Amazon Rekognition is an image analysis service that is part of the company's larger suite of AI-based applications. With it, users can add visual analysis to nearly any application via APIs. Rekognition can perform real-time analysis, call up to tens of millions of requests with a consistent latency time and, as a pay-per-play service, requires no initial startup cost to get up and running. The application also offers cloud-based facial analysis and facial recognition.

In 2017, Novetta, a leading advanced analytics company, released a biometric study comparing Amazon Rekognition with four other cloud-based facial recognition (CBFR) algorithms.[4] Two photo databases were used for comparison: FRGC-1000, which contained 1000 well-lit conventional face images in front poses (with expressions); the other, FRAME ID, a database of 500 faces in various poses with expressions and good illumination. The FRGC -1000 database primarily consisted of white and Asian males, and white females - each with a significant amount of background noise.[5] In the FRAME ID database, 95 percent of the photos contained white subjects. Of the four facial recognition algorithms analyzed, Amazon Rekognition performed best against the FRGC-1000 database for both constrained and unconstrained images.

It's important to note that the sample size of both databases in the Novetta study was relatively small when compared to those typically used to train facial recognition algorithms. Likewise, the subjects in the databases were overwhelmingly homogeneous, a factor, which, in and of itself, can result in misleadingly positive results.

In addition, Rekognition's detection error tradeoff (DET) - a graphical plot of the false rejection rate (FNMR) vs. false acceptance rate (FMR) for input patterns (i.e. photos) showed that the algorithm "minimized FMR at the expense of FNMR". In short, the minimization of Rekognition's false acceptance rate corresponded with a higher false rejection rate. In a law enforcement scenario, this would be the equivalent of *correctly* identifying an innocent person, while simultaneously *failing*

[3] Source: Maxine Jacob, "Facial Recognition Gains Ground In Europe, Among Big-Brother Fears", Euractiv, October 20, 2017.
[4] Source: Novetta (Amazon Rekognition: Quick-Look Biometric Performance Assessment, 2017).
[5] Defined as the random variation in brightness or color in an image.

to correctly match the image of a criminal to their photo in a database of known criminals.

Where Recognition Fails

Despite using relatively small sample sizes to test the precision of facial recognition systems, startups and industry leaders alike often proclaim that their software to be "highly reliable". But a lack of larger sample sizes isn't the only area where facial recognition systems fail to meet the test for applicability in security and law enforcement.

Gender

Facial recognition software can be highly reliable - given specific circumstances. When an algorithm analyzes white male faces, facial identification and gender matching accuracy rates are consistently in the 90+ range. But analyses of minority facial images often produce less positive results.

Using a dataset made up of 1270 African and Nordic countries with a high percentage of female lawmakers, Joy Buolamwini, a computer science graduate researcher at the MIT Media Lab and Timnit Gebru, a graduate student at Stanford, performed a facial recognition experiment using algorithms that explicitly offer gender recognition.[6] The algorithms came from three companies: Microsoft, IBM, and Megvii. The maximum error rate for lighter-complexion females was 7 percent and only up to 1 percent for light-complexion males. The algorithms misidentified darker-complexion men up to 12 percent of the time.[7]

However, the results for darker-complexion women were notably less successful. Microsoft's algorithm incorrectly tagged darker-complexion women *as men* 21 percent of the time while IBM and Megvii's algorithms averaged a nearly 35 percent error rate. In collaboration with a dermatologic surgeon, each of the darker-complexion women was assigned a skin tone score of IV, V or VI, in accordance with the Fitzpatrick scale.[8] A further breakdown of the gender recognition algorithms' accuracy, with respect to this metric, showed that these groups of women were misidentified at an average error rate of 20.8, 34.5 and 34.7

[6] 385 photos of light-complexion men, 296 photos of light-complexion women, 318 of darker-complexion men and 271 photos of darker-complexion women.

[7] Source: Steve Lohr "Facial Recognition Is Accurate, If You're A White Guy", *The New York Times* February, 9, 2018

[8] Source: Retrieved 15 July 2018 from https://dermahealthinstitute.com/blog/the-fitzpatrick-scale/

percent of the time, respectively. Two of the algorithms had an error rate of 46.5 and 46.8 percent when analyzing women in group VI, the darkest complexion group.
In reference to the outcome of the study, Buolamwini stated "To fail on one in three, in a commercial system, on something that's been reduced to a binary classification (male vs. female) task, you have to ask, would that have been permitted if those failure rates were in a different subgroup?... our benchmarks, the standards by which we measure success, themselves, can give us a false sense of progress."[9]

Race

In 2015, an African-American developer discovered that the image recognition algorithm in Google's Photos service, designed to identify and auto-tag objects in photos, had labeled photos of him and an African-American friend as "gorillas". Instead of investigating the root cause of the problem, Google opted for a more short-term fix. In early 2018, Wired magazine decided to test the rigors of this solution. Using a database of 40,000 images of animals ranging from dogs and pandas to primates like apes and monkey, the magazine made a surprising discovery. While Google's algorithm was able to correctly label a wide range of animals, including baboons and orangutans, it failed to identify any primates. After reaching out to Google, the company conceded (two years after the initial incident) that it had removed all terms related to great apes and monkeys - citing the unreliability of nascent image recognition technology as the cause.[10]

In light of growing concerns about biases in existing facial recognition systems, the ACLU of Northern California decided to conduct its own analysis of the popular open source algorithm Amazon Rekognition. For less than $15, the legal nonprofit was able to access 25,000 publically available arrest photos. The ACLU then compared publicly available photos of all sitting members of Congress against the database of those with a criminal record.[11] In total, 28 members of Congress were falsely matched as persons with criminal records.[12]

But false positives disproportionately affected minority members of the House and Senate. In fact, 40 percent of Rekognition's false positives were of minorities, even though minorities only make up 20 percent of Congress. Among those falsely matched was Civil Rights legend John Lewis (D-Ga.). Likewise, of the 49 members of

[9] Source: Larry Hardesty, "Study Finds Gender and Skin-Type Bias in Commercial Artificial Intelligence Systems", MIT News Office, February 11, 2018.

[10] Source: Tom Simonite, "When It comes to Gorillas, Google Photos Remains Blind", *Wired*, January, 11, 2018.

[11] As of June 2018.

[12] Source: Jacob Snow, "Ämazon's Face Recognition Falsely Matched 28 Members of Congress with Mugshots", ACLU Northern California, July 26, 2018.

the Congressional Black Caucus, Rekognition falsely matched six as criminals (~12 percent).

Policing

According to a 2016 study by Roland G. Fryer, Jr., of Harvard University, African-Americans and Latinos are 1.5 times more likely to experience some form of non-lethal force at the hands of police.[13] Similarly, a study by the science, policy and public health journal *Injury Prevention* found that when African-Americans or Latinos were stopped and questioned by police, an arrest occurred 82-85% of the time.[14] False matches in recognition systems could further prejudice these interactions and result in illegal or even life-threatening violations of civil liberties.

In 2009, a 47-year-old African-American woman named Denise Green was pulled out of her vehicle and held at gunpoint by four San Francisco Police Department officers. A license plate image recognition system falsely matched her car with one that had been stolen. Police failed to perform a standard visual check of the license plate, which would have prevented the incident. Could an assumption of guilt have prevented this rudimentary check?

More than 20 minutes passed before police realized that the car did not match the one stolen, during which time, Ms. Green was held at gunpoint. In 2009, a district court rejected Ms. Green's claim that the SFPD violated her 4th Amendment rights.[15] However, in 2014, the 9th circuit court ruled in favor of her suit, awarding her $495,000.[16]

A year-long study on police use of facial recognition conducted by the Privacy & Technology Center at Georgetown Law reached out to more than 100 law enforcement agencies around the country.[17] The report found that 52 state and local law enforcement agencies are currently using or have recently used facial recognition software. This includes cities like Los Angeles, Chicago, and New York City, as well as, the state of Texas, among others.[18] In addition, while a handful of states do have

[13] Source: Fryer Jr., Roland G. (2017). *An Empirical Analysis of Racial Differences in Police Use of Force*.

[14] Source: Miller, Ted R. et al., (2017). Perils of Police action: A Cautionary Tale from US Datasets. *Injury Prevention, 23(1).*

[15] Source: Kade Crockford, "San Francisco Woman Pulled Out of Car at Gunpoint Because of License Plate Reader Error", ACLU, May 13, 2014.

[16] Source: Denise Green v. City and County of San Francisco et al. (2015)

[17] Source: "The Perpetual Line-Up: Unregulated Police Face Recognition In America", Center on Privacy & Technology at Georgetown Law, October, 18, 2016.

[18] The NYPD denied the center's Freedom of Information request for records on facial recognition use.

regulations in place to monitor facial recognition use and ensure the rights of citizens most usage is unregulated.

Amazon Rekognition, the same software algorithm that incorrectly identified several members of Congress as criminals, is being aggressively marketed to law enforcement agencies throughout the country. The software also being marketed for use in hiring, school surveillance, and border control enforcement.

Civil Liberty and Consent Issues

First Amendment rights like the right to protest may also be at stake as a result of biased facial recognition systems.

The death of Freddie Grey, who died of spinal cord injuries in 2015 sustained as a result of police negligence, as ruled by the Maryland State Medical Examiner, sparked protests throughout Baltimore - some of which turned violent. However, anyone who attended a protest, peaceful or otherwise, between January and October of 2016 was unknowingly added to a facial recognition database. This database is used by the Baltimore Police Department to monitor the city's residents, the vast majority of whom are African-American.

Facial data obtained during the Freddie Grey protests were obtained using a Cessna airplane outfitted with military-grade surveillance equipment. These planes were originally intended for use in foreign war zones.[19] A 2016 Department of Justice report cited the need for greater oversight and transparency of Baltimore Police Department surveillance techniques, which, in addition to aerial surveillance and facial recognition have also included cell phone tracking.

But protesting isn't the only way that American citizens are being monitored without their knowledge or consent. In Los Angeles, there are 16 undisclosed locations from which the LAPD monitors citizens.[20] Using facial recognition cameras, law enforcement is able to scan faces in real-time from a distance of up to 600 feet. The images are then compared to a database of those with arrest records, open warrants or anyone suspected of gang activity. However, since this surveillance is performed without public knowledge or consent there is unlikely to be any inherent oversight involved - giving the department the freedom to manage the images as it sees fit.

[19] Source: Benjamin Powers, "Eyes Over Baltimore: How Police Use Military Technology to Secretly Track You", Rolling Stone, January 6, 2017.

[20] Source: Claire Garvey and Jonathan Frankel, "Facial-Recognition Software Might Have a Racial Bias Problem", The Atlantic, April 7, 2016.

Pushback and Industry Opinion

The general public may not be well-informed about the issues surrounding facial recognition, but a growing number of concerned citizens are - and they're pushing back. A coalition of Amazon, Google and Microsoft employees, shareholders, civil rights groups and more than 400 members of the academic community are petitioning against the use of Amazon Rekognition as a government surveillance tool. More than 150,000 members of the general public have also joined the protest.[21]

Dr. Wu Shuang of YITU Technology Singapore, which took home the top spot in the National Institute of Standards and Technology's (NIST) facial recognition competition two years in a row (2017, 2018), believes that in light of the realities of facial recognition biases companies must be aware of these issues. He cites the need for software vendors to engage in "full disclosure...that means not just a single number for each metric."[22]

Brian Brackeen, founder and CEO of Kairos, a leading facial recognition company, said that he would not market his product to law enforcement agencies citing that existing bias could be used to "dehumanize entire groups". In a TechCrunch essay, Brackeen wrote: "Whether or not you believe government surveillance is okay, using commercial facial recognition in law enforcement is irresponsible and dangerous." [23]

In 2018, an email circulated by Amazon employees protesting the company's agreement to provide its facial recognition technology to Immigration and Customs Enforcement (ICE) read "Our company should not be in the surveillance business; we should not be in the policing business; we should not be in the business of supporting those who monitor and oppress marginalized populations."

Google faced a similar backlash when several employees resigned due to the company's Pentagon contract for militarized facial recognition software. The tech giant ultimately decided not to renew the contract.[24]

[21] Source: Jacob Snow, "Ämazon's Face Recognition Falsely Matched 28 Members of Congress with Mugshots", ACLU Northern California, July 26, 2018.

[22] Source: Chris Burt, "Better Understanding of Facial Recognition Needed for Reasonable Social Dialogue, Says YITU Technologies AI Research Scientist", *Biometric Update*, March, 27, 2018.

[23] Source: Brian Brackeen, "Facial Recognition Software Is Not Ready for Use By Law Enforcement", *TechCrunch*, June 25, 2018.

[24] Source: Drew Harwell, "Google to Drop Pentagon AI contract After Employees Called It the Business of War" *The Washington Post*, June 1, 2018.

Why Algorithmic Bias Exists

In general, the root cause of race and gender bias in facial recognition can be traced to one or more of the following three scenarios: insufficiently trained data, a lack of accounting for ethno traits (and the separate training that they necessitate) and simple human error.

Many publicly available databases are made up of photos that are overwhelmingly made up of white males. This is fine if an algorithm is testing accuracy for this group alone, but the results of this type of limited training would be insufficient and inaccurate as a representation of the overall accuracy of an algorithm/software. Obtaining more diverse training sets (i.e. those that have a healthy distribution of women and darker-complexion minorities) may require that companies built their own databases - a less convenient solution than uploading a database from an open source, but one that can significantly improve accuracy rates, and help reduce problematic race and gender bias.

However, using a more diverse training set alone is unlikely to reduce bias to levels that would make an algorithm/software acceptable for use in the context of law enforcement, where civil liberties and public safety must be top priorities. To achieve this, different algorithms that factor in phenotypic traits unique to each specific group, need to be constructed for each racial groups being analyzed. Gfycat, a video hosting service, has seen a notable improvement in its algorithm's identification and matching accuracy for Asian subjects by exclusively training it on this group (while adjusting for relevant ethno markers).[25]

AI has yet to progress to the point of sentience (i.e. independent decision making) so facial recognition systems need humans to select a final image from the number of potential matches that are presented once an algorithm has done its job. Human biases as this critical stage can lower the odds of selecting the right match down to the equivalent of random chance. As Joanna Bryson, a computer scientist at Princeton University and the University of Bath states, "AI is just an extension of our existing culture." [26] There's no quick or easy way to eliminate human biases, but implementing oversight procedures that include peer collaboration and review can certainly help.

How Researchers and Vendors Are Responding

[25] Source: Tom Simonite, "How Coders Are Fighting Bias in Facial Recognition Software", *Wired*, March, 29, 2018.

[26] Source: Matthew Hutson, "Even Artificial Intelligence Can Acquire Bias Against Race and Gender", *Science*, April, 13, 2017.

Some researchers and software vendors are actively working to combat facial recognition bias and, in turn, protect the rights of citizens.

University of Toronto

AI researchers at the University of Toronto have designed a privacy filter for photos which prevents recognition systems from reading/scanning facial data. They've achieved this privacy protection by removing marker features (e.g. eye spacing data, nose length data etc.) - making it challenging for facial recognition systems to identify a face.

IBM

IBM recently announced that it will release two databases this fall: one containing more than 1 million images (making it the world's largest facial data repository) and another, with a total of 36,000 photos with an equal representation of different races, genders, and ages. The company is making this data open source in an effort to facilitate deeper research into the bias that currently exists in facial recognition.

Accenture

At the 2018 AI Summit in London, Accenture showcased a new facial recognition tool that lets users define which fields they deem sensitive (e.g. race, gender, age) and how these factors correlate with other data. The result is a system that allows for de-biasing on a number of relevant data points. Referred to as "mutual information", the software also provides a visual demonstration of how accurate a model is when variable dependencies are decoupled. Lastly, the algorithm assesses for "predictive parity" by comparing, for instance, whether or not true positive and true negative rates are the same for differing groups. While these strides are significant and likely to play a vital role in the success of future recognition systems, Rumman Chowdhury, Global Lead for Responsible AI concedes that there is often a trade-off between algorithmic accuracy and fairness.[27]

Standards and Oversight

It's not unreasonable for consumers and software users to expect vendors to take steps to ensure the accuracy and unbiasedness of their systems - before they are

[27] Source: Jeremy Kahn, "Accenture Offers Way to Erase Gender, Racial, Ethnic Bias in Artificial Intelligence", *Insurance Journal, June, 13, 2018.*

brought to market. But these measures alone are not enough. Both industry and government oversight are needed.

Once every four years, the National Institute of Standards and Technology (NIST) holds a voluntary facial recognition testing competition. Making participation and subsequent certification mandatory for vendors bringing facial recognition products to market could go a long way towards increasingly software accuracy and reducing harmful bias. According to a representative for the NIST, Amazon Rekognition has never participated in this event, despite being the most prolific vendor of this technology.[28] Lawmakers can participate in this oversight by legally requiring that software vendors submit any existing and future algorithms to the NIST for accuracy testing - an oversight measure similar to existing consumer protection laws.

In cooperation with the MIT Media Lab, Joy Boulamwini, the leading expert on racial and gender bias in facial recognition systems, is currently working with the Institute of Electrical and Electronics Engineers to create an oversight committee specifically for accountability and transparency in facial recognition software.

Moving Forward

Facial recognition is becoming an integral part of the technology landscape. But like all new technologies, it's not without its flaws - the most concerning of which are racial and gender bias. These deficiencies, particularly in the context of law enforcement, are simply untenable, as they threaten the rights and safety of millions of Americans.

Improving the accuracy of facial recognition systems is central to protecting citizens. But a deliberate and cohesive effort on the part of researchers, vendors, politicians and the general public is needed in order to achieve this urgent goal. Practically, this means rigorous and diverse algorithm training, thoughtful consideration of ethno traits in data acquisition and analyses, and the consistent application of checks and balances to reduce human error. Equally important, is the need for full vendor transparency and accountability, ideally, in the form of legally mandated software testing. Technological progress can't be stopped but it can, and must, be held to the highest of standards. Society can't afford anything less.

References:

[28] This information was obtained via email correspondence with the NIST's Biometric and Standards Testing Lead, Patrick Grother.

Amazon’s Face Recognition Falsely Matched 28 Members of Congress With Mugshots | ACLU of Northern CA

Opinion | When the Robot Doesn’t See Dark Skin - The New York Times

Even artificial intelligence can acquire biases against race and gender | Science | AAAS

Study finds gender and skin-type bias in commercial artificial-intelligence systems | MIT News

Facial-Recognition Software Might Have a Racial Bias Problem - The Atlantic

Accenture Offers Way to Erase Gender, Racial, Ethnic Bias in Artificial Intelligence

NovettaBiometrics_AmazonRekognitionBiometricPerformanceAssessment_WP-8182017.pdf

YITU takes top spot on NIST Facial Recognition Vendor Test leaderboard | BiometricUpdate

One of the few police departments to use Amazon’s facial-recognition tech has stopped – for now - The Washington Post

Facial Recognition Is Accurate, if You’re a White Guy - The New York Times

Microsoft Improves Biased Facial Recognition Technology | Fortune

IBM to release world's largest facial analytics dataset

U of T Engineering AI researchers design ‘privacy filter’ for your photos that disables facial recognition systems

How Coders Are Fighting Bias in Facial Recognition Software | WIRED