



Markkula Center
for Applied Ethics
at Santa Clara University

Examining the Ethics of Facial Recognition

By: Xiomara Quinonez '24

**2023-24 Hackworth Fellow at the Markkula Center for
Applied Ethics at Santa Clara University**

Examining the Ethics of Facial Recognition

Facial recognition is a form of biometrics that allows users to determine the similarity between two images of separate faces, which is done using software. This software is then able to transform images into numerical expressions that are generated by neural networks. Through this, facial characterization also occurs which allows the software to classify a face into different categories such as age, gender, and emotion. “Templates” are the transformations created by transforming a facial image into a numerical expression (Crumpler & Lewis 2021). These templates are then compared to other faces to find matches.

However, before this occurs, there are steps taken to develop the system. First, the system is trained using machine learning and artificial intelligence to integrate a database of images that allows facial recognition to be performed. It is important to use images with different angles and lighting because this impacts how the system performs. An enrollment database is later used where images for comparison can be presented in the database (Partnership on AI 2020). After the development process, facial recognition can then be used by entities. To begin, an image is acquired digitally. Then, the software examines the image to find patterns based on different characteristics presented on the images. Patterns found can include similar eyes and/or similar mouths. Specific features are then extracted to generate a template that will also be used for comparison. The template generated is then compared with templates in the database to produce different results with different accuracy results. Lastly, the scores generated are used by entities to determine whether or not these scores are enough to produce a match (Li et al. 2016). Facial recognition has widely been used by law enforcement to aid in the identification and capturing of potential suspects; however, it has recently also become used as a form of surveillance and/or safety which has raised the question of ethics and privacy policies surrounding facial recognition.

In connection with all of this is the concern about the rise of racial biases becoming present in the facial recognition systems used for these reasons. As issues that put the general public at a disadvantage continue to arise, it is important to consider what can be done to address the development and use of facial recognition. This paper will examine potential causes for biases being carried over into facial recognition algorithms, their causes, and their effects.

Technological Lens:

The lack of technological development in facial recognition has become apparent in recent years. From a technological perspective, facial recognition is viewed as a technology that can improve. With the current artificial intelligence being implemented into the software, many believe that facial recognition shouldn't be used due to its performance errors. These doubts relate to the science behind AI. When facial recognition does not have access to enough pictures of people of color, AI in facial recognition becomes harmful. In this case, it has led to the rise of false positives being observed for facial recognition. Many reports of false positive identifications show that there may be a lack of fair implementation in the algorithms. Even though there is a high positive rate reported for facial recognition, this is not enough to make facial recognition trusted or accurate in general. More specifically, facial recognition is much more difficult to trust when presented with images of people of color, leading many to conclude that it may not be inclusive of minorities.

Facial recognition is just like any other AI system. The more trained it is, the better it becomes. However, when examining how little goes into training facial recognition before it is used in the field, it is evident that people of color are often excluded during the development and testing of this system. This occurs due to the datasets being composed for facial recognition, which mainly contain light-skinned subjects as their main subject (Leslie 2020). With

light-skinned subjects being the main subject, the error rate for these subjects can be observed to be far lower for misidentifying a potential criminal suspect, but much higher when it comes to misidentifying dark-skinner subjects. It is reported by *The Atlantic* that although companies market their facial recognition technologies as having an accuracy rate of 95%, these claims cannot be verified as companies do not have to go through any extensive testing before placing their products on the market (2016). Additionally, most of the time the groups accessing this underdeveloped technology are law enforcement, where again, little to no accuracy checks are performed before the systems are deployed. A common example comes from Microsoft's FaceDetect system model which had an error rate of 0% for light-skinned male suspects and a 20.8% error rate for dark-skinned female suspects (Leslie 2020). This shows how misrepresentation in the development process of AI systems causes racial biases to arise within the application of facial recognition in real-world scenarios. As the representation of subjects remains skewed to light-skinned subjects representation, this can be observed to increase error rates which then cause misidentifications and lead to harmful consequences. To examine the use of facial recognition by law enforcement, a notable example is the use of Amazon's Rekognition system. Rekognition was provided to the Orlando police department at no cost to pilot the software, along with a nondisclosure agreement that hid details about the agreement from the public. Through this use, there were reports of many misidentifications of people of color using Rekognition, leading to many wrongful arrests of the incorrect people. This led many to believe that Rekognition violated constitutional rights. These issues even led to extreme cases of misuse, including the misidentification of many members of Congress as criminals (Brandom 2018).

Although many examples show the possible need for a facial recognition ban in some places, corporations such as the New York Police Department (NYPD) are fighting it. On

NYPD's website, there is now a section that includes frequently asked questions and answers about the department's use of facial recognition. When asked about studies showing frequent incorrect analysis of minority groups caused by facial recognition, NYPD responds that as part of their policy, facial recognition-recommended matches are required to be analyzed by multiple human reviewers (2024). It is important to note that history has shown human statements are not to be trusted at times. This is comparable to eyewitness statements and the biases that these statements also hold. Studies have shown that humans are often susceptible to misinformation. When trying to identify perpetrators, many mock witnesses have been studied to not be able to correctly identify perpetrators, showcasing memory biases. With memory biases, predetermined beliefs about the world or groups of people contribute to the memories people create (Laney & Loftus 2024). In fact, at times, fake memories can even be created where people convince themselves of something to be true. This is comparable to the police use of facial recognition for verifying potential suspect matches with images. If many people are able to create false memories for identifying perpetrators, what is stopping law enforcement from doing the same to attempt to verify a potential match between images.

Upon conducting an informative interview with Irina Raicu, director of the Internet Ethics Program at the Markkula Center for Applied Ethics, one of the main points highlighted was that when negative implications like the ones mentioned occur, we should be questioning when we should use facial recognition, instead of when we should not use it. It is important to not accept harm to people who are already vulnerable and to not increase the vulnerability they face in new forms. Instead, we must think about the context where the benefit of facial recognition is so great, that we are willing to give up some of our privacy and accept various consequences in attempt to truly become safer in specific settings.

This level of benefit usually does not occur and instead more harm than good is done to vulnerable communities. Apart from the policing examples, another example of harmful context is the case of the use of facial recognition in Rite Aid stores. These stores used AI-based surveillance software to misidentify many people of color and accuse them of shoplifting, leading Rite Aid employees to follow many customers around the store and call the police (Federal Trade Commission 2023). This misuse led to Rite Aid stores being prohibited from using facial recognition and left people of color becoming targeted, something that they would have most likely not encountered if these stores were not using these systems.

This reckless use of facial recognition shows that oftentimes flawed technology does not help any group at all and will not make us safer, for as long as it continues to be a flawed technology. In order to improve the flawed technology, there needs to be a way to prove that these facial recognition technologies are truly working if we are going to use them in public spaces. This also depends on the context of when we should use it and what we are giving up as a society to get this benefit, if any benefit at all. For example, if we had evidence of more accuracy, the use of facial recognition for safety would then be very beneficial for security purposes, or similar contexts. Specific situations may include the use of facial recognition to prevent terrorist attacks in countries where these attacks are common. When considering situations like these, facial recognition would become immensely important in keeping communities safe. However, safety would only be possible if facial recognition was developed with the correct amount of accuracy and continued to be tested throughout time to ensure false positives were not an issue. More thorough testing must occur not only during the development phases but also once the product is already on the market. Technologies are changing systems and in order for them to continue developing effectively, we must perform testing that ensures

this development is happening correctly. In order to target one of the risks of public use of facial recognition, we must also target the issue of testing properly. Still, only because the technology can be improved, we should still examine where the improved technology should be used. For example, the ACLU reports that the National Institute of Standards and Technology (NIST) frequently tests facial recognition systems. NIST examines companies' algorithms by feeding images of a face to compare with potential matches in the database (2024). This is comparable to some real-world uses, such as the use of facial recognition to identify criminal suspects. However, the ACLU explains how no laboratory tests truly represent the reality of how facial recognition systems are used in real scenarios, specifically the use by law enforcement (2024). This is due to the lack of factors that contribute to the use of facial recognition in the field, such as varying qualities of images. This shows how even with continuous testing, even that amount of testing is not enough. It is important to keep in mind that external factors affect even the testing of facial recognition, which may make it difficult to truly target the issues presented by it.

Social Lens:

Although most people are aware of the concept of AI and most people know what facial recognition is often used for, many people do not realize the implications of both of these things, specifically the negative implications. Many also do not realize that improperly trained facial recognition systems become harmful largely due to the biases being carried over by the programmers developing these technologies. This also relates to the use of facial recognition by law enforcement. One of the biggest causes for these mismatches mentioned may be related to programmers carrying over their own discriminatory biases into the development of these algorithms. This has led to not only faulty algorithms being designed, but also introducing a new form of racial biases. The racial makeup of the team developing facial recognition systems and

their own personal biases play one of the biggest roles in how facial recognition systems perform. Often, the team of developers working on facial recognition systems are not minorities themselves, which means they often do not account for other racial groups when training these systems. This results in overlooking important features when training datasets, even if it is unintentional. Due to examples such as these, facial recognition has been described as widening pre-existing inequalities (Najibi 2020).

Along with the increase of racial biases being carried over by programmers into facial recognition, there is a lack of inclusivity in the teams developing these algorithms. While the racial makeup of groups developing these systems mainly includes white males, it results in the exclusion of different perspectives from being included in the development of algorithms that contribute to facial recognition – that is perspectives that would lead to the inclusion of more diverse faces and groups in training datasets, which could in turn improve facial recognition performance. Although oftentimes, programmers may carry over their predefined racial biases as mentioned above, it is also possible for it to not be intentional sometimes. Unfortunately, without being surrounded by different groups and perspectives, these programmers become limited to what they know and these racial biases may become simply all they know. This shows the need for a diverse set of data scientists, which includes minorities, to develop these products so that we are able to create more equitable facial recognition algorithms.

When considering other contributing factors to this issue, we must also consider that a lack of education regarding how to develop equitable forms of AI can lead to outcomes like these. By not having the availability or access to resources to be able to learn how to implement the algorithms developing facial recognition fairly and equitably, this may be affecting future developments of AI systems. The lack of availability to learn how to implement more fair

algorithms may be linked to the lack of learning ethics throughout students' education. Having the ability to explore ethical dilemmas during grade school has been shown to affect behavior changes and lead to more decisions without prejudice. Research shows that in order for people to become more caring and empathetic, they must receive help in all stages of life to become this way, with one way being the teaching of ethics (Shamout 2020). It is important to see that the lack of responsible AI being implemented now may be due to the lack of ethics that taught people how to think about ethically developing a product like this. By not having the ability to explore ethical dilemmas before entering the workforce, many then lack the thought of seeking to develop equitable systems during their careers, which may be because they were not previously able to explore these ideas. More specifically, in computer engineering college curriculums, there seems to be a lack of emphasis on ethical programming in major-related courses. Although students are able to partake in general ethics and philosophy courses, both related and unrelated to technology, there are not enough courses on ethics in actual computer engineering classes, that go into more detail on how to develop equitable and ethical algorithms. The lack of ethics teaching throughout the fundamental years of programmers' school lives ultimately contributes to the way these programmers are developing their computer programs. In the case of facial recognition, these factors may contribute to the way that algorithms are being developed racially biased, which may be unfortunately the only thing that programmers know.

Legal Lens:

The legal lens considers the abuse of technology and power that's happening with facial recognition, along with the lack of litigation to prevent this abuse. Although there are general laws that corporations and developers must follow, there seems to be a lack of specific laws related to the development of the technology. This means that there are very few AI-specific

regulatory rules that these corporations must follow before deploying facial recognition. However, there have been many lawsuits throughout the years against corporations that have proven the need for regulatory laws to prevent the unlawful use of facial recognition. An example is Amazon's unknown use of facial recognition in New York Go stores (Collier 2023). Due to the lack of regulation when the system was introduced, little was done to stop Amazon from using facial recognition in these stores, leading to their continued abuse of the system for years, Amazon was later discovered to have been using facial recognition technology in their physical stores in New York, without telling their customers. This led to the storage of the customers' biometric information and prompted a lawsuit against Amazon. It's important to note that New York requires businesses to make customers aware if they are tracking their biometric information, which again highlights the implementation of some laws. Even with laws such as these, it was not enough to require Amazon to develop and test their facial recognition systems in a better way. Amazon set up its Go stores, which allows customers to shop without having to go through the checkout process, leaving customers to be charged after they have left the store. Although this was possible due to the use of facial recognition, Amazon was urged to try and explore other implementations that could aid in their vision for the store. This along with other examples of misuse of facial recognition from a single corporation shows the need to develop more AI-specific laws to regulate the use of these systems and prevent the vulnerabilities we are introducing to communities.

While issues like these remain prominent, many are taking action to prevent more harm from being done. For example, cities such as San Francisco and Berkeley have banned the use of facial recognition by law enforcement (Nakar & D. Greenbaum 2022). Companies such as IBM have also agreed to step out of the business of facial recognition. However, the main focus to

create more ethical facial recognition systems is to expand the laws to regulate these systems. The push for laws has become more common in some states where state officials are urging for the regulation of facial recognition after witnessing riots in the last few years. An example is Illinois, which introduced the passage of the Biometric Information Privacy Act to create more transparency between private entities and consumers. Through this passage, businesses in Illinois are required to inform consumers that their facial profiles will be collected and may be used and/or shared with another entity (Hodge Jr 2022). If violated, this can result in a penalty charge. Another example included the biometric identifier law passed in Washington, which prohibits entities from enrolling consumers into a database without notice or consent. This not only includes facial images, but also fingerprints, voiceprints, and eye retinas. Laws such as these have regulated the use of facial recognition, but not the development of it.

The push continues for laws and policies to address more specific concerns — laws that target the problem while it is being developed, instead of when it is already in the market. Specifically, there is a push for right-to-be-forgotten laws to protect consumers from convictions. This is in an effort to also increase anonymity within facial recognition systems, specifically, in public. With anonymity in public spaces, citizens would be protected from government recordings through facial recognition, which would also protect their whereabouts (Nakar & D. Greenbaum 2022). While these laws have not been passed yet, they do address the main concerns of facial recognition. Still, there is the need for laws that target different applications of AI, to prevent loopholes and the misuse of different applications.

Upon conducting an interview with Linsey Krolik, Assistant Clinical Professor of Law at Santa Clara University, the lack of legal frameworks for responsible AI was reiterated in this conversation. Currently, there are not many frameworks that corporations are legally required to

follow to introduce new AI technologies to the market, resulting in not enough regulation for AI on a broad or global level. However, with the introduction of the mentioned European Union Artificial Intelligence Act, AI will be examined with risk management. With this act, AI will be looked at as a product and new measures such as physical and discrimination safety will be taken into account. More specifically, in regard to facial recognition, with this act, there is potential for facial recognition to be viewed as an unacceptable risk, which would result in potential banning in some instances. In addition, with acts such as these, companies will be legally required to let the public know if they are using AI. This will bring a lot more certainty to the public and provide a definite answer about whether or not AI is being used.

Exploring Recommendations

The injustices created by facial recognition can be explored through the different lenses examined in the paper. Each focus has its own strength and addresses different concepts of the issue. However, when examining from a holistic approach, it is clear that the social solution is more powerful than the rest as it targets different ideas from issues presented through each lens.

Technological Lens Recommendations

The technological lens discusses the AI that's being implemented into facial recognition. Based on this lens, the best solution is to improve the AI. The lack of testing both before and after the deployment process of facial recognition is one of the main contributors to the improper development and training of this system. In order to address this issue, the main solution would be to improve the technology behind this AI. This can be done by introducing new testing methods for facial recognition. Testing can occur by developing more accurate algorithms through the continued increased use of images of people of color for the training of the machine

learning systems behind facial recognition. This would create an improved accuracy for facial recognition, as well as be more consistent with the overall population of different communities. In addition, there is a need for continuous testing and testing that adapts to the way that technology is changing. Through the use of continuous testing, AI systems would be able to recognize images of people of color in a better way and reduce misidentifications.

This is important because there is a need to correct mistakes that are made on the technical side of facial recognition. More importantly, because these systems were developed and designed years ago, some developers might not even remember how they categorized images. By having continuous testing, there is a mandatory requirement for continuous acknowledgment of the way facial recognition continues to change, which leaves programmers to examine potential biases they might have carried over.

Additionally, there is a need to effectively introduce new testing constraints and methods. As mentioned, even though continuous testing helps improve the accuracy of facial recognition, it is only so helpful. There is a need to have more accurate laboratory testing that is more aligned with the way facial recognition is used in real-world scenarios. Through this method, there will be a lot more improvements in the accuracy, that work in favor of people of color.

Social Lens Recommendations

The social lens suggests that spreading awareness of racial biases would prevent further targeting of minorities. By educating people on the problems of facial recognition, people will learn how to use facial recognition more wisely. A key way that this education can be formed is by promoting it in college settings, specifically in engineering courses. By teaching computer engineering students more about ethics in their major-related courses, these students would

become better equipped on how to develop computer programs that are inclusive and ethical. This would target the problem early on.

One way teaching ethics can be introduced is through the teaching of what it means to have responsible AI. For example, teaching what are examples of responsible and nonresponsible AI. This goes hand in hand with teaching about the potential negative implications of developing systems that are biased. Most importantly, there is a need to teach what an example of an unethical computer program looks like. Through teachings like these, students would become more aware of the programs they are writing, which would make them feel more inclined to strive for creating ethical computer programs.

Legal Lens Recommendations

This solution deals with the entirety of the injustice and the privacy and policy concerns that harm minorities. A proposed solution would be to improve the litigation surrounding facial recognition before it is even introduced to markets. This would be especially useful in ensuring that we are only using facial recognition that we can prove works as intended.

In addition, there is a need for laws that can adapt. As described, continuous issues continue to arise, such as racial biases and the lack of freedom of speech it may create. Because of this, there is a need for laws that adapt to be able to target different implications imposed by facial recognition. This shows how new concepts for specific AI applications continue emerging and there is a need for laws that continue adapting with the way these technologies are changing to limit the risk they create.

Lastly, similarly to the social solution, there is a need for better education on the importance of ethical facial recognition and responsible AI for lawmakers. Because these people are ultimately the ones who are building laws that shape what is and is not allowed to be used, it

is important for these people to become educated on the implications of what facial recognition can create. These efforts will ultimately address concerns caused by facial recognition.

Examining the Need for a Better Ethics Curriculum

Integrating ethics is very important when introducing technologies such as facial recognition. It is important for students, programmers, and policymakers involved in the implementation and use of facial recognition to become educated on the ethical concerns of this technology. This can help create a better system for minorities. Better integration of an ethics education can be especially useful for addressing concerns examined by the technological, social, and legal lenses.

From a technological perspective, the introduction of ethics in education would allow people to better observe the need for more thoughtful, unbiased testing and development of facial recognition. This would also ensure that all ethical considerations are acknowledged and kept in mind during the creation of these facial recognition systems, ultimately ensuring more reliable systems are created.

From the social perspective, a better introduction of an ethics curriculum would allow engineering students to understand ethical issues in AI systems more easily, especially those issues presented in facial recognition. Creating this awareness would ultimately allow for programmers to better identify what can be improved with these systems. As well as, how to address and avoid unwanted issues such as racial biases from being carried over into the development of AI systems.

From a legal perspective, educating those in positions of authority about the ethical implications of AI and facial recognition would shape the policies around facial recognition in a

positive way. By introducing the importance of facial recognition ethics, these policymakers would promote the creation of more ethical facial recognition systems.

Addressing the negative implications of facial recognition through ethics education can be a way to address technical, social, and legal concerns surrounding this technology. This approach will ensure people develop a better understanding of what it means to have ethical and equitable facial recognition technology, making the possibility of a future with technologies centered around these values.

Works Cited

- Brandom, Russell. "Amazon's Facial Recognition Matched 28 Members of Congress to Criminal Mugshots." The Verge, theverge.com, 28 July 2018, <https://www.theverge.com/2018/7/26/17615634/amazon-rekognition-aclu-mug-shot-congress-facial-recognition>. Accessed 10 Apr. 2024.
- Collier, Kevin. "Amazon Sued for Not Telling New York Store Customers about Facial Recognition." NBC News, cnbc.com, 16 Mar. 2023, <https://www.cnbc.com/2023/03/16/amazon-sued-for-not-telling-new-york-store-customers-about-facial-recognition.html>. Accessed 15 May 2024.
- Crumpler, William, and James A. Lewis. "How Does Facial Recognition Work?" Center for Strategic and International Studies, June 2021, https://csis-website-prod.s3.amazonaws.com/s3fs-public/publication/210610_Crumpler_Lewis_FacialRecognition.pdf. Accessed 2 Feb. 2024.
- Flanagan, Linda. "What Students Gain From Learning Ethics in School." KQED, kqed.org, 24 May 2019, <https://www.kqed.org/mindshift/53701/what-students-gain-from-learning-ethics-in-school>. Accessed 12 Feb. 2024.
- Garvie, Clare, and Jonathan Frankle. "Facial-Recognition Software Might Have a Racial Bias Problem." The Atlantic, apexart.org, 7 Apr. 2016, <https://apexart.org/images/breiner/articles/FacialRecognitionSoftwareMight.pdf>. Accessed 7 Feb. 2024.
- Gerchick, Marissa, and Matt Cagle. "When It Comes to Facial Recognition, There Is No Such Thing as a Magic Number." NBC News, aclu.org, February 7, 2024,

<https://www.aclu.org/news/privacy-technology/when-it-comes-to-facial-recognition-there-is-no-such-thing-as-a-magic-number>. Accessed 2 Apr. 2024.

Hamilton, Masha. “The Coded Gaze: Unpacking Biases in Algorithms That Perpetuate Inequity.” The Rockefeller Foundation, rockefellerfoundation.org, 16 Dec. 2020, <https://www.rockefellerfoundation.org/case-study/unpacking-biases-in-algorithms-that-perpetuate-inequity/>. Accessed 2 Apr. 2024.

Hodge Jr., Samuel D. “The Legal and Ethical Considerations of Facial Recognition Technology in the Business Sector .” DePaul Law Review, vol. 71, no. 3, Summer 2022, <https://via.library.depaul.edu/cgi/viewcontent.cgi?article=4202&context=law-review>. Accessed 9 Nov. 2023.

Laney, Cara, and Elizabeth F. Loftus. “Eyewitness Testimony and Memory Biases.” NOBA, nobaproject.com, 2024, <https://nobaproject.com/modules/eyewitness-testimony-and-memory-biases>. Accessed 22 Apr. 2024.

Learned-Miller, Erik, et al. Facial Recognition Technologies in the Wild: A Call for Federal Office. Algorithmic Justice League, MacArthur Foundation, 2020, pp. 1–57, https://assets.website-files.com/5e027ca188c99e3515b404b7/5ed1145952bc185203f3d009_FRTsFederalOfficeMay2020.pdf. Accessed 22 Apr. 2024.

Leslie, David. Understanding Bias in Facial Recognition Technologies. The Alan Turing Institute, 2020, https://arxiv.org/pdf/2010.07023https://www.gse.harvard.edu/sites/default/files/parent_et_hical_kids_tips_.pdf. Accessed 24 May 2024.

Li, Lixiang, et al. “A Review of Face Recognition Technology.” IEEE Access , vol. 8, Aug.

2020,

<https://via.library.depaul.edu/cgi/viewcontent.cgi?article=4202&context=law-review>.

Accessed 9 Nov. 2023

Martinez-Martin, N. (2019, February). Neural Computing and Applications. *AMA Journal of Ethics*; journalofethics.ama-assn.org.

<https://journalofethics.ama-assn.org/article/what-are-important-ethical-implications-using-facial-recognition-technology-health-care/2019-02>. Accessed 10 Feb. 2024.

Najibi, Alex. "Racial Discrimination in Face Recognition Technology." Harvard Kenneth C.

Griffin Graduate School of Arts and Sciences, sitn.hms.harvard.edu, 18 Feb. 2020,

<https://sitn.hms.harvard.edu/flash/2020/racial-discrimination-in-face-recognition-technology/>. Accessed 10 Apr. 2024.

Nakar, Sharon, and Dov Greenbaum. "NOW YOU SEE ME. NOW YOU STILL DO: FACIAL RECOGNITION TECHNOLOGY AND THE GROWING LACK OF PRIVACY." *B.U. J. SCI. & TECH. L.*, vol. 23, no. 88:90,

<https://www.bu.edu/jostl/files/2017/04/Greenbaum-Online.pdf>. Accessed 11 May 2024.

Nawaz, N. (2022). Artificial intelligence applications for face recognition in recruitment process.

Neural Computing and Applications, 2. <https://doi.org/10.1007/s43681-021-00096-7>.

Accessed 10 Apr. 2024.

"NYPD Questions and Answers Facial Recognition." New York City Police Department,

nyc.gov, 2024,

<https://www.nyc.gov/site/nypd/about/about-nypd/equipment-tech/facial-recognition.page>.

Accessed 2 June 2024.

PAI Staff. "Bringing Facial Recognition Systems To Light." Partnership on AI,

partnershiponai.org, 18 Feb. 2020,

<https://partnershiponai.org/paper/facial-recognition-systems/>. Accessed 9 Nov. 2023.

Raising Caring, Respectful, Ethical Children. Harvard Graduate School of Education, 2024,

https://www.gse.harvard.edu/sites/default/files/parent_ethical_kids_tips_.pdf. Accessed

17 Feb. 2024.

"Rite Aid Banned from Using AI Facial Recognition After FTC Says Retailer Deployed

Technology without Reasonable Safeguards." Federal Trade Commission, [ftc.gov](https://www.ftc.gov), 19

Dec. 2023,

<https://www.ftc.gov/news-events/news/press-releases/2023/12/rite-aid-banned-using-ai-facial-recognition-after-ftc-says-retailer-deployed-technology-without>. Accessed 22 Feb.

2024.

Shamout, Omar. "Can Ethics Classes Actually Influence Students' Moral Behavior?" UC

Riverside News, news.ucr.edu, 20 Sept. 2020,

<https://news.ucr.edu/articles/2020/09/22/can-ethics-classes-actually-influence-students-moral-behavior>. Accessed 17 Feb. 2024.

"The EU Artificial Intelligence Act." EU Artificial Intelligence Act, artificialintelligenceact.eu,

2024, <https://artificialintelligenceact.eu/>. Accessed 3 June 2024.