AI Facial Recognition in Law Enforcement: Benefits, Risks, and the Need for Regulation

By Harrison Gothelf

Artificial intelligence facial recognition technology has become one of the most consequential and contested tools adopted by law enforcement agencies in the modern era. By employing machine learning algorithms to analyze facial features and compare them against images stored in large databases, facial recognition systems enable investigators to identify individuals with a speed and efficiency that no prior investigative technology could match. Law enforcement agencies have adopted the technology for a range of purposes: identifying suspects captured on surveillance cameras, locating missing persons, and solving crimes that might otherwise remain unsolved. Yet despite these capabilities, the expansion of facial recognition technology in the law enforcement context has produced substantial legal and ethical controversy. Scholars and policymakers have documented that facial recognition systems produce biased results, particularly when identifying individuals from certain demographic groups. The ability of government agencies to identify and track individuals in public spaces raises important questions about privacy and civil liberties under the Fourth Amendment. Recent empirical research further suggests that the adoption of facial recognition technology by police departments may exacerbate racial disparities in arrest patterns. Documented wrongful arrests arising from algorithmic misidentification have already occurred in the United States, underscoring that these are not speculative harms.

How Facial Recognition Technology Works

Facial recognition technology refers to computer systems that identify individuals by analyzing unique facial characteristics. These systems function by converting an image of a person's face into a mathematical representation, commonly referred to as a facial template, that measures distinctive features such as the distance between the eyes, the shape of the nose, and the structure of the jawline. The system then compares that template against templates stored in a database to determine whether a match exists. Modern facial recognition systems rely heavily on artificial intelligence and machine learning. Algorithms are trained on large datasets containing thousands or millions of facial images, enabling them to learn patterns that distinguish one face from another. As algorithms are exposed to additional data, their capacity to recognize subtle variations in facial features improves. In recent years, this technology has become increasingly integrated with digital surveillance infrastructure. Cameras located in public areas, transportation systems, and commercial establishments generate vast quantities of visual data that facial recognition software can analyze in real time. Law enforcement agencies may submit images from surveillance footage into these systems, enabling investigators to compare the image against databases containing mugshots, driver's license photographs, or other government records.

Public Safety and Investigative Use Cases

The most prominent law enforcement application of facial recognition technology is the identification of suspects captured on surveillance cameras. In many criminal investigations,

investigators possess imagery of a suspect but lack information about the individual's identity. Facial recognition systems enable investigators to compare such images against large databases and generate potential matches, pursuing leads that might otherwise remain undiscovered. Supporters of the technology argue that these capabilities improve public safety by helping police identify suspects involved in serious crimes. Beyond suspect identification, facial recognition has been used to locate missing persons and identify victims of human trafficking. Because the systems operate quickly and can analyze vast quantities of imagery simultaneously, they allow investigators to cover investigative ground at a scale that would be impossible through manual review alone. These are genuine benefits that any regulatory framework must take into account.

Algorithmic Bias and Discriminatory Impact

One of the most serious criticisms of facial recognition technology involves algorithmic bias. Researchers have found that many facial recognition systems perform differently when analyzing faces from different demographic groups. Multiple studies have demonstrated that these algorithms exhibit higher error rates when identifying women and individuals with darker skin tones. The landmark Gender Shades study by Buolamwini and Gebru (2018) documented substantial disparities in commercial facial recognition systems' accuracy across intersectional race and gender categories, providing foundational empirical support for this concern. The National Institute of Standards and Technology's Face Recognition Vendor Testing reports have since corroborated these findings across a broader set of commercially deployed systems. These

disparities frequently arise from the composition of datasets used to train machine learning models. Where training datasets contain a disproportionate number of images from certain demographic groups, the resulting algorithm may achieve greater accuracy for those groups while performing less reliably for others. The consequences in the criminal justice context are severe. A misidentification may cause investigators to focus on innocent individuals, potentially resulting in wrongful arrests or prosecutions. Even where facial recognition outputs are nominally treated as investigative leads, investigators in practice may accord them substantial evidentiary weight. A study analyzing data from more than one thousand U.S. cities found that jurisdictions adopting facial recognition technology experienced increased racial disparities in arrest patterns following adoption. Arrest rates for Black individuals rose relative to arrest rates for white individuals in cities that implemented these systems. While the study does not establish that facial recognition technology directly causes discriminatory policing, its results suggest that the technology interacts with existing policing practices in ways that produce unequal outcomes. Several structural factors likely contribute. Law enforcement databases frequently contain a disproportionate number of images from communities that have historically been subjected to elevated levels of police surveillance. When facial recognition systems query these databases, they may generate matches more often for individuals from those communities, compounding rather than correcting existing racial disparities in the criminal justice system. The equal protection implications of these disparities warrant careful doctrinal analysis. Under Washington v. Davis, a facially neutral governmental action does not violate equal protection absent proof of

discriminatory intent. A policy producing racially disparate effects alone is constitutionally insufficient. However, the systematic nature of algorithmic bias, operating predictably and repeatedly to generate misidentifications at higher rates for Black individuals, raises distinct constitutional questions that existing doctrine has not yet fully resolved. Where policymakers have actual knowledge of documented disparate error rates and continue to deploy these systems without remediation, the case for discriminatory intent becomes substantially stronger.

Fourth Amendment and Privacy Concerns

The widespread deployment of facial recognition technology raises foundational Fourth Amendment questions about government surveillance in public spaces. Under Katz v. United States, the Fourth Amendment protects individuals' reasonable expectations of privacy. Although individuals generally have no reasonable expectation of privacy in their physical appearance while present in public spaces, the systematic, automated identification and tracking of individuals through facial recognition technology presents a qualitatively different form of government surveillance. The Supreme Court's decision in Kyllo v. United States is instructive. In Kyllo, the Court held that the use of thermal imaging technology to detect heat from inside a private residence constituted a search under the Fourth Amendment, reasoning that the government may not use sense-enhancing technology not in general public use to gather information from a constitutionally protected area that it could not otherwise obtain without physical intrusion. The Kyllo rationale, that technology enabling the government to obtain information it could not otherwise acquire implicates Fourth Amendment protections, provides a

foundation for arguing that automated facial recognition in public spaces is similarly subject to constitutional scrutiny. More directly, Carpenter v. United States substantially limits the third-party doctrine's application to comprehensive digital surveillance. In Carpenter, the Court held that the government's warrantless acquisition of cell-site location information constituted a Fourth Amendment search, reasoning that the comprehensiveness of digital location tracking, its ability to chronicle an individual's movements over time, implicated a reasonable expectation of privacy that the third-party doctrine's mechanical application would not protect. The Carpenter majority reasoned that when the government assembles detailed, long-term records of a person's activities, the intrusion on privacy is categorically different from isolated observations. Facial recognition technology deployed through public camera networks has the potential to enable precisely the kind of comprehensive tracking the Carpenter Court found constitutionally suspect. Unlike traditional surveillance, which requires human officers physically to follow an individual, an integrated facial recognition system could potentially track a person's movements across an entire city, generating a comprehensive record of their associations, political activities, and daily patterns. Legal scholars have warned that such capabilities could produce a chilling effect on free expression and political participation. Where individuals reasonably believe that their identities will be recorded during public demonstrations or political events, they may be deterred from exercising constitutional rights. The doctrinal gaps in this area are substantial. Lower courts have not yet established clear standards governing when facial recognition searches require warrants, what level of suspicion justifies their use, or whether their use in real-time tracking contexts

differs constitutionally from their use in retrospective investigations. These gaps create precisely the conditions under which congressional legislation is most needed.

Toward a Federal Regulatory Framework - Proposals

A comprehensive federal regulatory framework governing law enforcement use of facial recognition technology should incorporate several interlocking elements. Each listed addresses a distinct dimension of the technology's potential for harm. In order to find more success, some of these proposals should be in place.

Warrant Requirements and Judicial Oversight

Congress should require law enforcement agencies to obtain a judicial warrant before conducting a facial recognition search, except in narrowly defined exigent circumstances. Consistent with Carpenter, the warrant requirement should apply both to prospective monitoring of individuals through facial recognition systems and to retrospective searches comparing surveillance footage against law enforcement databases. Judicial oversight would ensure that facial recognition is employed only where investigators possess a legitimate, legally cognizable basis, providing a structural check against pretextual or speculative searches. The warrant application should require investigators to specify the crime under investigation, the basis for believing the facial recognition search will yield relevant evidence, and the database against which the search will be conducted.

Mandatory Algorithmic Auditing and Accuracy Certification

Before any law enforcement agency may deploy a facial recognition system, that system should be required to undergo independent algorithmic auditing to assess accuracy rates disaggregated by race, sex, age, and skin tone. Vendors should be required to certify that their systems meet minimum accuracy thresholds across all demographic categories, not merely in aggregate. Systems that produce error rates above specified thresholds for any demographic group should not be eligible for law enforcement use. Existing proposals that require only warrant procedures, without addressing the algorithmic accuracy dimension, are insufficient. A warrant requirement protects against unreasonable use of a system; it does not ensure the system is reliable. Both protections are necessary.

Human Review and Corroboration Requirements

A federal statute should prohibit arrests, detentions, or prosecutions based solely on a facial recognition match. Because facial recognition algorithms produce probabilistic outputs, not identifications, treating a match as conclusive proof of identity creates an unacceptable risk of wrongful arrest. Documented cases of wrongful arrest arising from facial recognition misidentification have already occurred, underscoring that this is not a theoretical concern. Federal law should require that any facial recognition match be reviewed by a trained human examiner before it is used as a basis for investigative action, and that the match be corroborated by independent evidence before an arrest may be made. Officers who receive facial recognition

results should be trained to understand the technology's limitations and to treat its outputs as investigative starting points rather than evidentiary conclusions.

Transparency and Public Reporting Obligations

Many law enforcement agencies have adopted facial recognition systems without meaningful public disclosure or democratic deliberation. A federal statute should impose mandatory transparency requirements to enable democratic oversight of this technology. Agencies should be required to publicly disclose the specific facial recognition systems they use, the frequency of deployment, the categories of crimes for which the technology is employed, and the databases queried. Annual public audit reports should document aggregate accuracy metrics, the number of searches conducted, the number of matches generated, the number of arrests following a match, and the outcomes of subsequent prosecutions. Prior to adopting a new facial recognition system, agencies should be required to submit to a public notice-and-comment process that allows affected communities and elected officials to weigh in on the decision.

Data Governance and Privacy Protections

The sources of facial image data used to construct law enforcement databases warrant independent regulatory attention. Some commercial facial recognition vendors have assembled databases containing billions of images scraped from social media and other online platforms without individuals' knowledge or consent. Clearview AI's operation exemplifies this practice.

Federal law should prohibit vendors providing services to law enforcement from assembling databases through the unauthorized scraping of publicly posted images. Law enforcement databases should be limited to government-held records, mugshots and driver's license photographs, for which there is a documented legal basis for collection. Federal law should additionally require that facial recognition search data, including the query image, the database queried, the results generated, and the disposition of any investigation, be deleted after a defined retention period unless the data is needed for an active criminal prosecution. Permanent face-tracking databases, in which individuals' movements are indefinitely logged, should be prohibited. Individuals whose biometric data is stored should be entitled to notice and provided with a legal mechanism for challenging the retention of their information.

Due Process and the Right to Challenge

Federal law should establish procedural rights for individuals subjected to facial recognition identification in the context of criminal proceedings. Defendants should have the right to know whether facial recognition was used at any point in the investigation leading to their arrest or prosecution. They should have the right to obtain the underlying facial recognition evidence, including the query image, the database, and the algorithmic output, and to challenge its reliability in court through expert testimony and cross-examination. Courts should develop standards governing the admissibility of facial recognition evidence that account for the documented limitations of specific systems and the demographic-specific error rates established

through independent auditing. These procedural protections are necessary to ensure that the use of probabilistic algorithmic outputs in criminal proceedings satisfies basic due process requirements.

Independent Oversight Boards

The foregoing requirements will be effective only if they are subject to meaningful enforcement. Existing internal oversight mechanisms within law enforcement agencies have proved insufficient in other technology contexts, and there is no reason to expect they would perform differently here. Congress should therefore establish an independent oversight body, structured along the lines of the Privacy and Civil Liberties Oversight Board, with authority to audit agency compliance with facial recognition regulations, investigate complaints of misuse, and approve or deny proposed deployments by federal agencies. The board should include technical experts, civil liberties advocates, and community representatives, and should be empowered to issue binding directives, impose civil penalties for violations, and publish its findings publicly. State and local law enforcement agencies receiving federal funding should be required to comply with the board's standards as a condition of funding.

Artificial intelligence facial recognition technology presents a paradigmatic case of innovation outpacing governance. Its investigative benefits are real, but so are its documented harms: algorithmic misidentification that falls disproportionately on Black individuals and women,

deployment without judicial oversight, secret adoption without democratic deliberation, and a structural capacity for comprehensive public tracking that threatens both privacy and expressive freedom. The doctrinal frameworks currently available, Fourth Amendment reasonableness analysis and equal protection doctrine, are ill-equipped to address these harms comprehensively and in real time. Congress should enact legislation establishing a comprehensive regulatory framework: warrant requirements and judicial oversight; mandatory algorithmic auditing and accuracy certification; human review and corroboration requirements; transparency and public reporting obligations; data governance protections; due process rights for defendants; and independent oversight with enforcement authority. This framework would not prohibit law enforcement use of facial recognition technology. Rather, it would ensure that the technology is deployed in a manner consistent with the constitutional guarantees that define the relationship between the government and the people it serves. The question is not whether facial recognition technology is useful. It plainly is. The question is whether its use is governed by rules adequate to the constitutional stakes. At present, it is not. This framework proposed would bring law enforcement use of facial recognition into alignment with the principles of judicial oversight, equal treatment, and democratic accountability that constitutional governance requires.