CALIFORNIA GOLDEN FLEECE® AWARDS

# The Pitfalls of Law Enforcement License Plate Readers in California and Safeguards to Protect the Public

Jonathan Hofer

INDEPENDENT INSTITUTE

INDEPENDENT INSTITUTE

Independent Institute, 100 Swan Way, Oakland, CA 94621-1428
T: 510-632-1366 • F: 510-568-6040 • info@independent.org • www.independent.org

The *California Golden Fleece® Awards* shine a spotlight on waste, fraud, and abuse in California government to give valuable information to the public, enabling them to provide needed oversight and demand meaningful change.

## About the Independent Institute

The Independent Institute is a nonprofit, nonpartisan, public-policy research and educational organization that shapes ideas into profound and lasting impact through publications, conferences, and effective multimedia programs. The mission of the Independent Institute is to boldly advance peaceful, prosperous, and free societies grounded in a commitment to human worth and dignity.

INDEPENDENT INSTITUTE

*Briefing*
August 16, 2022

CALIFORNIA GOLDEN FLEECE® AWARDS

# THE PITFALLS OF LAW ENFORCEMENT LICENSE PLATE READERS IN CALIFORNIA AND SAFEGUARDS TO PROTECT THE PUBLIC

By Jonathan Hofer

## Contents

## Overview

In 2009, a 47-year-old Black woman named Denise Green was forced to the ground at gunpoint by several San Francisco police officers during her car ride home from work. During the lengthy hold up, the officers searched Green's vehicle, while other officers had their guns pointed at her while she was handcuffed. Green never had a criminal record.

Her crime? The police alleged she was a car thief, but after an extensive detention, police acknowledged that Green's burgundy Lexus was, in fact, not the gray GMC truck they were looking for. An automated license plate reader, or ALPR, notified police that Green's car was stolen after misreading her license plate. The lesson of her story is that this could happen to anyone on the road.

California law enforcement agencies have come to embrace ALPRs enthusiastically. ALPR systems gather information from passing cars faster than police officers can visually confirm license plates, and the systems compare the plate numbers against a registry or they relay the plate numbers to dispatchers. ALPRs are high-speed cameras that can rapidly scan numerous computer-readable images, eliminating the need for law enforcement personnel to do manual checks.

Despite their increasing prevalence, local governments have paid little attention to their departments' sensitive technologies. Some municipalities have failed to

| Image credit: LukaTDB / iStock

adopt measures to prevent abuse before purchasing the equipment, and the few ALPR laws on the books are often ignored or are not comprehensive enough to prevent misuse. A lack of an overarching governance framework is to blame. Now out in the wild, ALPRs represent a significant risk to civil liberties.

Californians would benefit greatly from ALPR data-collection limits, regular data cleaning, and transparency. Until a structure is in place that protects individuals' privacy and provides law enforcement with a template to ensure accountability, no ALPR network is satisfactory.

Therefore, the Independent Institute awards its thirteenth *California Golden Fleece® Award* to agencies that have adopted ALPR technology. The award is granted for the circumvention of state laws and ordinances, the failure to implement policies before the use of automated license plate readers, the lack of adequate safeguards ensuring the safety and civil liberties of individuals, misleading the public, and essentially testing surveillance systems on Californians. Honorable mentions are given to the California Legislature for failing to address the problems associated with state departments installing ALPR networks, such as the California Highway Patrol, while having negligent policies or procedures, and the Northern California Regional Intelligence Center (NCRIC) for hosting and sharing vast quantities of ALPR data with poor safeguards.

## Background

Beginning in the early 2000s, ALPRs quickly exploded onto the law enforcement scene. An estimate by Axon, one of the largest suppliers of police body cameras and cruiser dash cams, notes that by 2019 ALPRs were "one of the most widely used surveillance systems in existence." California cities have been especially eager to adopt ALPR cameras, including those within the largest urbanized areas of Los Angeles-Long Beach-Anaheim, San Francisco-Oakland, San Diego, Riverside-San Bernardino, Sacramento, San Jose, and Fresno.

ALPRs are high-speed cameras that capture video or images of passing vehicles. What distinguishes ALPR cameras from other high-speed cameras is their ability to read the alphanumeric characters on license plates. The software technology enables the cameras to recognize and record optical characters quickly. It is similar to the ways in which some companies perform data entry from documents, or how Google Books allows people to search printed publications electronically. Some ALPRs used to utilize infrared imaging, although the current trend is to use standard imaging techniques like those in consumer digital cameras.

As vehicles are detected by an ALPR, or the camera takes several pictures of each vehicle, the camera scans for the license plate and applies an algorithm. The optical character recognition software then reads and deciphers the license plate.

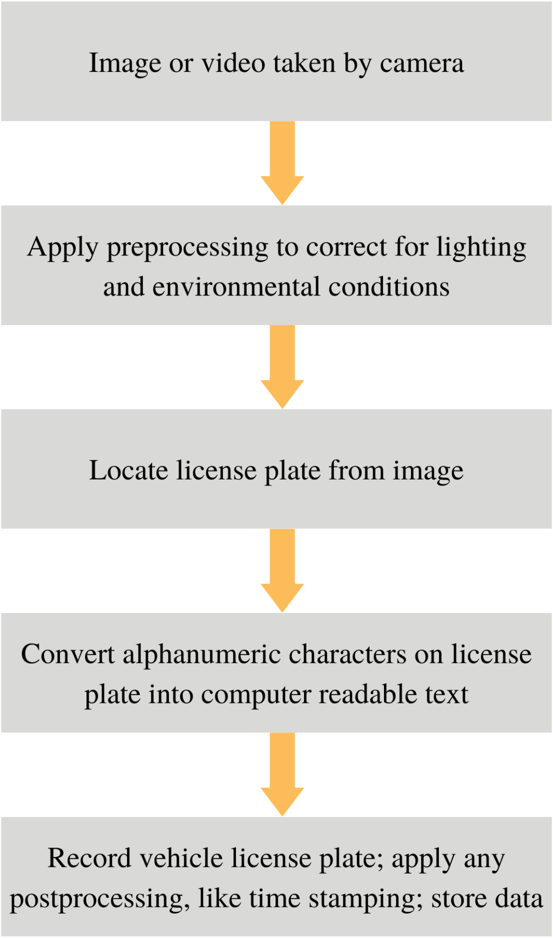
Image or video taken by camera
Apply preprocessing to correct for lighting and environmental conditions
Locate license plate from image
Convert alphanumeric characters on license plate into computer readable text
Record vehicle license plate; apply any postprocessing, like time stamping; store data

After reading the plate, the camera system typically logs the license plate's number along with the time and coordinates of the scan. The record is stored in a database. Police then receive alerts when a vehicle of interest in the database has

*(Example of stationary ALPRs over a roadway. Photo credit: LudvikaSweden Photography | Wikimedia Commons CC BY-SA 4.0)*

been spotted by a camera. Vehicles of interest generally are said to be on a "hot list," especially those that are reported stolen.

Broadly, ALPRs come in two varieties, either stationary or mobile. Stationary ALPRs are fixed in place, such as on a traffic light or over a freeway, and their placement usually is intended to scan large numbers of vehicles.

Mobile mounted variants are attached to vehicles, typically police cruisers. Mobile ALPRs also are sometimes attached to more ordinary city or county vehicles. San Francisco, for example, uses ALPRs on city buses to identify cars blocking bus stops. And San Jose outfits garbage trucks with ALPRs to send data directly to the police.

Although stationary ALPRs and mobile ALPRs are functionally equivalent in their technologies, their data collection tactics differ slightly. If stationary ALPRs are installed throughout a given road or highway, they can deduce the direction and speed of the passing cars. Combined with historical data, individuals' travel patterns could be determined. Mobile ALPRs allow officers to direct their cameras at specific areas, such as particular parts of a city, or to monitor the traffic coming to and departing from certain businesses. Mobile ALPRs likewise can be deployed to fill in gaps left by the absence of installed stationary cameras and to narrow down flagged vehicles of interest while patrolling.

## Stops Gone Wrong

Proponents of ALPRs argue that the readers help police "identify stolen vehicles, people wanted for a crime and missing persons." Despite proponents' claims that they aid in catching criminals, ALPRs suffer from several substantial drawbacks. For instance, the readers, which cost as much as $20,000 each, have troubling error rates. According to an estimate by an ALPR data aggregator, the cameras misread one out of 10 license plates—rather poor accuracy, which is especially concerning, considering that the cameras can scan 2,000 plates per minute. Given the error rate of ALPR systems, mistakes and misidentifications are frequent. Since vehicle stops by law enforcement sometimes are based on ALPR "hits," mistakes are magnified.

*(Example of a mobile ALPR. Photo credit: Mbrickn | Wikimedia Commons, CC BY 4.0)*

ALPR-related stops generally are not routine traffic stops, whereby an officer might pull over a car and issue a warning or a citation for speeding. Because ALPRs are intended to be a part of a department's strategy of fighting car theft, kidnappings, and so on, the stops that ALPR hits trigger frequently are considered felony or high risk traffic stops. In these situations, standard police procedures direct officers to respond with "guns at the ready" because the presumption is that the officers are dealing with known or suspected felons who may be armed and dangerous.

This can lead to frightening circumstances for victims of ALPR errors, as Mark Molner discovered in 2014. Molner was driving home from a sonogram appointment with his pregnant wife when a police vehicle aggressively darted in front of his BMW and blocked his path in a Kansas City suburb. Puzzled as to what he could have done to cause such a stop, Molner was even more shocked—as was his wife, who was witnessing the scene—when the officer unholstered his gun, though he never pointed the weapon at Molner.

As with the case of San Francisco's Denise Green, the stop was triggered by an ALPR misread. The ALPR had read a "7" as a "2" on Molner's license plate and erroneously alerted police that the plate belonged to a stolen vehicle. The responding officers neglected to manually check Molner's license plate after the ALPR's first scan. The police officer eventually verified that Molner's BMW was not the stolen Oldsmobile sought by the police.

Mistakes like this are not always caused by a camera misreading a license plate number, however, and may also occur when ALPR systems rely on unclear or faulty information. In a 2020 viral incident, police from Aurora, Colorado, pulled over Brittney Gilliam, who was taking her younger sister, daughter, and nieces to get their nails done. After realizing that the salon was closed, the family members returned to their car and were quickly surrounded by police officers with their guns drawn. The officers separated Gilliam from the children, who ranged from 6 to 17 years old, and detained them at gunpoint. Only the 6-year-old was not handcuffed.

The officer said the police had received an alert by an ALPR notifying them that Gilliam's car, a minivan, was stolen, prompting the felony stop. The ALPR, however, seemingly had confused Gilliam's Colorado license plate with a motorcycle's plate carrying an identical number from Montana.

Gilliam's case also stresses a vital point about ALPR data. Gilliam's car was stolen earlier in the year. She reported it stolen, and the car quickly was recovered. At the time of Gilliam's stop, police raised the issue of her car being reported stolen, and the officers admitted that the report could have been the reason the ALPR flagged her vehicle. Relying on the technology can be dangerous, even when the technology works. The ALPR did not mistranslate Gilliam's license plate numbers. In fact, the camera read the characters correctly. The issue may have been duplicated plate numbers from different states—license plate numbers are not unique—or the stolen vehicle report could have triggered the alert. Errors of this kind can be database issues. If the police never cleared her stolen vehicle report from the ALPR's database, the system would still flag her vehicle as stolen.

Sometimes vehicles should not even be on a stolen car list. In a gruesome story, Ali Badr, a man from Oakland, California, was on his way to work in

December 2020 when San Ramon police officers stopped his car after being tipped off by an ALPR hit. In a police video obtained by the *San Francisco Chronicle*, Badr is seen being mauled by a police dog despite no apparent provocation. Although the non-resisting Badr complied with officers' directions, the K-9 violently bit Bahr's arm, ripping it apart. The K-9's handler, instead of calling the dog off, walked up to Badr and pointed his gun at Badr's head. According to reporting on the resulting lawsuit,

> [The arresting officer] then grabbed Badr's left arm, while the dog's teeth were still sunken into Badr's other arm, and threw Badr to the ground. The officer then knelt on Badr's back, and grabbed his neck and forced him face-down on the pavement while two other officers also knelt on Badr's back as he was handcuffed. …
>
> The police K-9 was allowed to continue biting the Plaintiff for over 50 seconds. … During the traffic stop, all officers named as defendants pointed their guns at Badr. No officers intervened during the dog attack.

Badr lost the use of his arm because of the mauling. In the video, Badr is heard exclaiming in pain, "What I did? What I did?"

Badr was driving a Toyota Camry that he was renting from CarMommy, a rental service catering to delivery drivers and gig workers. Badr previously worked as an Uber and Lyft driver and started delivering food during the COVID-19 pandemic. Unbeknownst to Badr, CarMommy had reported the car stolen to the San Jose Police Department, placing the Camry's license plate on shared databases of stolen cars, which the San Ramon Police Department's ALPR identified.

Badr had fallen behind on his rental car payment by a couple of days, although he had been in contact with the rental company, telling the company he would pay them soon, as he had done previously. The rental agreement's language allowed only for the car to be reported as stolen if specific criteria were met, meaning that CarMommy's report of the stolen vehicle may have been submitted in bad faith. The ALPR system, however, is unable to distinguish between good faith and bad faith reports.

Rental cars have been a particular source of inappropriate stops in California. Dozens of customers renting cars from Hertz have been falsely arrested and jailed in recent years because Hertz reported the cars as stolen. ALPRs could have alerted police to those false reports. California Gov. Gavin Newsom (D), in 2019, signed Assembly Bill 391, which reduced the time window from five days to three days that rental car companies must wait to report a car as stolen if the rental contract had expired, possibly encouraging more situations like Badr's in the future.

Those stories highlight the dangers of automated policing: One inaccurate piece of information or one computer error can lead to a serious confrontation. Even if all the technology is working correctly, a car is not a person. An officer may be stopping a vehicle because they think they are pursuing the driver, but automated technology does not have the capacity for flexibility, understanding nuanced situations, or investigating erroneous or misleading ALPR hits.

## What Does an ALPR Capture?

Not only can ALPRs affect policing outcomes, but they also represent a genuine privacy concern. It is not solely an issue of a camera taking a license plate's picture. Because they record the times and places of vehicle movements, they can provide an intimate picture of people's lives. When many points of data are aggregated, ALPRs become powerful surveillance tools.

In addition to recording a license plate number, ALPR images include the car itself being photographed—and enough of the car usually is visible to identify the vehicle from the image. A US Department of Homeland Security memo about its recognition software claims that the agency is testing the software to make it capable of identifying the make and model of the car. A car with bumper stickers or a conspicuous paint job may be more likely to be identified from those pictures. Even the vehicle's occupants may be photographed.

After requesting that the City of San Leandro in Alameda County, California, send him a record of every time his car had been photographed, Mike Katz-Lacabe was shocked to learn that pictures of his car were taken more than 100 times in a year just by the San Leandro Police Department's ALPRs. One of the pictures was of him with his daughters exiting their car in their home's driveway.

Because of the magnitude and depth of the information that ALPRs gather, such license plate databases are ripe for abuse. In February 2022, a former Everett, Washington, police officer stood trial for using police resources, including a license plate database, to stalk a woman and frame her boyfriend for drug crimes and theft in order to break them up. The following are some other recent examples:

- The background of a 2020 US Supreme Court case, *Van Buren v. United States*, involved a successful sting operation conducted by the FBI that included an informant bribing a Georgia police officer to look up a license plate for a woman the informant said he met at a strip club, in exchange for $5,000.
- An audit of Minnesota's Dakota County Sheriff's Office revealed that 104 different police officers were looking up the driver's license records—a type of database that typically is linked to ALPR databases—of a particular female police officer they were stalking online. Her records were accessed 425 times.
- A Pennsylvania officer has been accused of tracking his estranged wife's movements using his department's ALPRs. "The printout that we received regarding his use of the license plate readers included over 100 pages of entries as far as the positions, locations, and times of family members," the chief of police said.
- A Massachusetts officer ran the plates of his ex-wife's friends while stalking her.
- A 2021 lawsuit claimed that New Jersey officers accessed a man's license plate information to harass him for "befriending the ex-girlfriend of one of the officers."

In addition to the stalking risk, constant monitoring is likely to create a "chilling effect," whereby lawful activity is suppressed out of fear and social pressure. In 1998, a Washington, DC, police officer admitted that he had used a license plate database to extort people whose cars were parked outside of a gay nightclub. Virginia State Police officers used ALPRs to scan the license plates of vehicles going to rallies for Barack Obama and Sarah Palin in 2008. US Immigration and Customs Enforcement (ICE) used local law enforcement scans to record cars going to a 2010 gun show. The use of government power to surveil, threaten, and extort or otherwise punish people for holding certain political beliefs or engaging in other perfectly legal activities is a frightening prospect that has only been made easier and more enticing by ALPR technology.

## The Dangers of Data Sharing

Since many of the ALPR databases are interconnected, a significant concern is the proliferation of data-sharing agreements among public-sector institutions. Substantial information sharing is expected as vehicles travel into and out of different jurisdictions. For example, a local police department could process a stolen vehicle report and place the vehicle on hot lists to which the California Highway Patrol has access. A CHP ALPR could then register a hit on the vehicle if the car travels from the local surface streets onto an interstate highway.

State and local agencies are not the only ones that collect and utilize ALPR data. In addition to municipal governments frequently housing their own ALPR record depositories, large numbers of ALPR scans in California are relayed to the Northern California Regional Intelligence Center (NCRIC). NCRIC is a "fusion center," a body that is meant to be a point of intelligence sharing, usually for counterterrorism purposes. NCRIC was established by the High Intensity Drug Trafficking Area Program (HIDTA) in the Office of National Drug Control Policy, operating under the White House. It is a participant in the Department of Homeland Security's National Network of Fusion Centers.

Agencies without their own ALPR systems can access such records from other departments. Private companies, such as Motorola Solutions and Leonardo, also may collect their own license plate records and sell them to law enforcement departments or hand them over to NCRIC.

Fusion centers such as NCRIC raise concerns for civil liberties since their bulk data collection efforts can create detailed portraits of individuals' behavior and travel patterns. Fusion centers exist in a "no-man's land" between the federal government and the states, and they operate with questionable authority and limited oversight. Some states and localities may have promulgated stronger regulations on data collection and privacy standards. But by working with a fusion center—effectively, a third party—the intentions of privacy laws can be skirted as fusion centers obfuscate who "owns" the data. Fusion centers may operate with weaker regulations on data collections and privacy standards than many local government agencies. Operating outside the purview of public scrutiny is worrisome, as some fusion centers rely on other third parties to collect or host data, increasing the risk that unauthorized persons may have access to sensitive information.

With loads of data in the hands of many parties, it is essential that databases are accurate and up to date. As police can contribute to a common pool of data, any changes in the status of vehicles, such as the recovery of a stolen car, must be reflected in the database in a timely manner. If one department reports a car as stolen, but never updates the database to reflect that the car was later found, other departments' ALPRs will send out an alert, possibly leading to a false arrest.

Given the sensitivity of ALPR data, reasonable care of the stored records is expected, and agencies should demonstrate a legitimate need for access to the information. Law enforcement agencies, however, often open their ALPR databases to others. California Civil Code sections 1798.29 and 1798.90.5 permit individual departments to share ALPR data with other public agencies, provided that they do so with due consideration to individuals' privacy. But police departments appear to have few, if any, standards concerning another agency's request; access usually is granted, no questions asked.

In 2020, the California State Auditor's Office issued a report on its investigation of the ALPR practices of several local California police departments. They found that Sacramento's police department, for example, shares data with more than a thousand different agencies across the country.

Among the more peculiar image-sharing arrangements, the police departments of Fresno, Marin, and Sacramento all share ALPR data with the Honolulu Police Department. This practice is curious, to say the least, since few cars are likely driven in both Hawaii and California.

Data sharing also may thwart state and city sanctuary laws. In 2020, the City of Pasadena purchased $80,000 worth of ALPR equipment while promising that none of its license plate logs would be furnished to ICE after concerns were raised that ALPR technology "fuels ICE's deportation machine." Months later, documents showed that Pasadena police were passing license plate data to ICE through a Homeland Security investigations team.

In 2018, the City of Long Beach issued a sanctuary memorandum. The Long Beach Values Act, as it was called, expanded the limits on data sharing with federal immigration agencies already put in place by California Senate Bill 54. Since 2019, a revision to the Long Beach Police Department Policy Manual prohibits employees from assisting "in the enforcement of federal immigration law" barring special circumstances. Then, through a 2020 public records request, a Long Beach resident found that the department had been sending data directly to ICE.

In 2018, it was discovered that the Bay Area Rapid Transit (BART) system had shared information with NCRIC that was accessible to ICE, despite BART's internal sanctuary policies. BART had installed

ALPRs in its MacArthur Station parking lot. Oakland resident Tracy Rosenberg said of the news, "I don't think how many times I park at MacArthur BART is any of the business of the Department of Homeland Security, and I'm a citizen."

Sanctuary cities demonstrate the difficulty of regulating the public deployment of ALPRs without broader state laws and strong local practices. The federal system of the United States creates important checks and balances between the national government and state governments. The Jeffersonian principle that state laws may differ from federal practices is ingrained in the text of the US Constitution. As the language of SB 54 says, "Entangling state and local agencies with federal immigration enforcement programs diverts already limited resources and blurs the lines of accountability between local, state, and federal governments." Data sharing circumvents the spirit of the separation of powers. It gives federal agencies access to California resources in order to enforce federal laws that the state explicitly rejects.

## Law Enforcement's Poor Record of Data Protection

Despite the real problems of shared databases, even the locally hosted databases of individual police departments may suffer from negligent oversight. In the 2020 California State Auditor's report, the ALPR data processing procedures were reviewed for the Fresno Police Department, Los Angeles Police Department (LAPD), Marin County Sheriff's Office, and Sacramento County Sheriff's Office. Of the four agencies reviewed, the auditor found that none of the departments implemented all of the practices required by Senate Bill 34—one of the few California laws that apply to ALPR systems. SB 34 requires training for personnel on how to use the system, permits only authorized personnel to access it, and places restrictions on the transfer of ALPR data.

The state auditor's report noted that the LAPD did not even have usage or privacy policies at all. The LAPD and the Sacramento County Sheriff's Office would add names, addresses, dates of birth, and criminal charges to their ALPR record systems, sometimes including information from the California Law Enforcement Telecommunications System (CLETS) maintained by the state's Department of Justice. State law requires special protection for such data, including encryption requirements, background checks, and employee training. Yet, the Sacramento County Sheriff's Office could not demonstrate to auditors that they did any vetting of their information storage and retrieval systems. The Fresno Police Department, Marin County Sheriff's Office, and Sacramento County Sheriff's Office all were unable to confirm who has access to the system, who is responsible for oversight, or how to delete ALPR data.

Those three departments also contracted out the storage of photos and ALPR data to a third-party cloud storage vendor. No department could confirm that the vendor met CLETS standards. The Fresno, Marin County, and Sacramento County agencies outsource ALPR data to a cloud database operated by Vigilant Solutions (a subsidiary of Motorola Solutions since 2019), but auditors found that because the Vigilant software is by default accessible via the Internet, officers could access the data on their personal devices, bypassing the agencies' network security safeguards.

The three agencies storing ALPR data in Vigilant's cloud never spelled out or enforced all the required security precautions in their contracts with Vigilant. The auditor found that "[t]he agencies' contracts, for example, do not stipulate that Vigilant would store its data in the US or Canada. Marin's contract is vague about who owns the data it uploads to the ALPR system."

The entire ALPR system of the Marin County Sheriff's Office was not even operated or maintained by an ALPR administrator, but instead was run by a deputy in the auto theft department who had no background in network security. One former employee retained access to Marin's ALPR database despite resigning the previous year.

In 2022, the Marin County Sheriff settled a lawsuit after three Marin residents alleged the Sheriff was sharing unprotected ALPR data with federal, state, and local agencies in violation of SB 34 and California sanctuary laws. Marin has offered data access to more than 400 out-of-state agencies, including federal immigration enforcement agencies. In the settlement agreement, the Sheriff's office conceded and agreed to stop sharing license plate information with agencies outside of California in order to comply with state law.

Other cities, which were not subject to the audit, have poor data protection records as well. Berkeley did not even have a retention policy before installing ALPRs and, according to a response to the ACLU, may have never had a "firm discussion" about ALPR retention before the system's implementation. In a 2021 Oakland city commission meeting, the city's police department acknowledged that it gave "the FBI 'unfettered access' to the license plate data in violation of the city's policy."

## The Questionable Legality of ALPRs

At present, no US Supreme Court rulings directly address the general use of ALPRs. In Fourth Amendment jurisprudence, a well-established concept is the plain view doctrine. This legal doctrine permits police to perform a search or a seizure if the officer:

1. Is in a place they have a right to be, e.g., a public road,
2. Does not have to "enter" something to search (such as a car or house), and
3. Makes an inadvertent discovery.

If these criteria are met, the officer does not need to obtain a warrant.

License plates are visible to the general public. Therefore, taking a picture of a car's license plate on a public roadway would not ordinarily violate an individual's reasonable expectation of privacy. As such, ALPRs, by simply taking pictures of cars on public thoroughfares, are not violating the Fourth Amendment. In *United States v. Knotts* (1983), the US Supreme Court concluded that visual surveillance did not constitute a Fourth Amendment search because a "person traveling in an automobile on public thoroughfares has no reasonable expectation of privacy in his movements from one place to another."

Recording and retaining the geolocation of a car's whereabouts, however, which ALPR data amounts to, pose serious constitutional questions. The government cannot track a person or vehicle with GPS without a warrant, an issue the Supreme Court affirmed unanimously in *United States v. Jones* (2012). Justice Sonia Sotomayor's concurrence stressed that the government may violate an individual's expectation of privacy with many forms of surveillance, even when there is no physical intrusion.

ALPR systems are, thus, legally questionable because, while they are not technically the same as a GPS device, ALPR records effectively substitute for GPS by consistently recording a vehicle's location at various points in time. As a vehicle travels about and passes ALPRs, the ALPR system notes the place and time and effectively tracks the vehicle. Legal scholars have suggested that the practice falls under a Fourth Amendment legal doctrine called "mosaic theory." The mosaic theory holds that even if collecting an individual data point on a person's location is not unconstitutional per se, when many data points are connected over time, that surveillance amounts to a Fourth Amendment search, which would require a warrant.

In a landmark 2018 case that dealt directly with the 1983 *Knotts* decision, the US Supreme Court in *Carpenter v. United States* took a step going beyond merely prohibiting real-time tracking as it did in *Jones*. In *Carpenter*, the Court held that the government may not even use historical information from a cell phone tower to "retrace the steps" of a person without first obtaining a warrant. The Court belabored the point that "a person does not surrender all Fourth Amendment [privacy] protection by venturing into the public sphere." The geolocation of a cell phone, much like an ALPR image, is an extensive log that an observer could use to determine the habits and patterns of travel. The only true cure would be to limit the time that ALPR data can be retained to prevent retracing a person's steps to recreate a pseudo-GPS map.

In *Commonwealth v. McCarthy* (2020), the Massachusetts Supreme Court endorsed extending the mosaic theory to ALPRs. The court declined, however, to rule in favor of Jason McCarthy, holding that not enough data had been collected by the ALPRs on either side of a single bridge on which McCarthy was traveling to create a detailed enough picture of his movements, and, thus, that the ALPRs did not constitute an unlawful search. Similarly, in *United States v. Yang* (2020), heard by the Ninth Circuit Court of Appeals, and *Uhunmwangho v. State* (2020) from the Texas Ninth District Court of Appeals, judges ruled in favor of the police because the cars in those cases were scanned only once—not sufficient for a mosaic. In each of those cases, the opinions of the court warned that if more cameras were used then a Fourth Amendment search could be triggered because this would allow law enforcement to reconstruct past movements.

In many California cities where ALPRs saturate the roads, ALPRs are scanning cars continuously and retaining the data. In Piedmont, California, a city of only 1.7 square miles, the thirty-nine ALPRs cover virtually every street. A heat map of Oakland's ALPR use, obtained by the Electronic Frontier Foundation, shows that nearly the entire city is covered.

With that many data points, combined with the long data retention typical of California police departments, the ALPRs deployed in California seemingly would amount to a Fourth Amendment search under mosaic theory. Accordingly, a warrant should be required to access the historical records of a vehicle, or vehicles, especially for purposes of any data analysis or mapping.

## The Undetermined Effectiveness of ALPRs

Despite the civil liberties concerns and a high potential for mistakes, local governments and police departments across California continue to purchase and operate ALPRs, claiming that they are invaluable crime-fighting tools. In February 2022, the California city of Saratoga purchased seven cameras for $20,000 for a pilot ALPR program. Acknowledging the privacy concerns, but nonetheless advocating for the ALPRs, Saratoga Mayor Tina Walia said, "In my mind there has to be a balance keeping the public safe as well as protecting their privacy."

But do ALPRs increase public safety? Does credible research support the claim of improved safety?

The supporting literature is sparse. In a short 2011 experiment by the Center for Evidence-Based Crime Policy at George Mason University, researchers concluded that "[ALPRs do] not achieve a prevention or deterrent effect" on crime.

A 2012 study conducted with a vehicle theft unit in Mesa, Arizona, looked at ALPRs' effectiveness in "recovering stolen automobiles, apprehending auto thieves, and reducing auto theft." While remaining optimistic about the future of ALPRs, the researchers concluded,

> While we do not find that the [ALPR] was able to reduce auto theft we did find that another hot spots policing approach (the same auto theft unit doing focused police work but doing manual checking of license plates) was able to reduce auto theft. …
>
> We found no vehicle theft crime displacement or diffusion of benefits from our targeted routes to areas adjacent or near these routes related to any of our models. Our results suggest that a specialized vehicle theft unit can have an effect on reducing vehicle theft compared to the control group, but only when this group does manual checking of plates as opposed to using the [ALPR] equipment.

After an Atlanta suburb in Cobb County installed thirteen ALPRs in 2019, even police had difficulty attributing anecdotal drops in local nonviolent crimes to the cameras. Stuart VanHoozer, Cobb County's deputy chief of police, said, "To make it very clear, we are not 100 percent positive that Flock cameras [ALPRs] were the difference."

In my 2021 study of the Police Department of Piedmont, California, I reviewed vehicle theft and ALPR hit data from 2004 to 2021. For years, Piedmont tracked how many investigative leads were generated by ALPR hits on cars from shared databases. Investigative leads could count discoveries such as witness or suspect identification, or the locations of stolen vehicles. The dataset was one of the most comprehensive of its kind, and no other department in the nation is known to have kept such complete historical records.

After performing statistical analysis, a weak positive correlation was found between an ALPR hit on a flagged car and an investigative lead. Moreover, additional analysis also demonstrated a weak correlation between hits and recovered stolen vehicles. The statistical evidence suggests that ALPRs do not provide strong benefits.

Proponents of ALPRs claim many benefits, but little evidence of such benefits has been reported. More scholarly analysis must be undertaken before we can know if ALPRs produce any benefits and, if so, what those benefits might be. (Appendix A also includes a more technical discussion on ALPRs, algorithmic bias, and predictive policing.)

## Key Recommendations on How to Protect Public Privacy and Safety When Jurisdictions Use ALPRs

The problems associated with ALPRs discussed above point to several key recommendations that must be adopted to protect both public privacy and public safety.

### 1. Limit database access and use.

ALPR data should be accessed and shared only on a need-to-know basis. The standard adopted by the California Law Enforcement Telecommunications System (CLETS) is that only employees with appropriate training should be entrusted with access to these databases. That standard should be universal for all agencies that use ALPRs. Personnel no longer employed by their respective departments should have access terminated when they leave their jobs. All systems should use two-factor authentication to prevent inappropriate access. Supervisors should be able to view data usage reports to see the activities, such as search queries, of ALPR system administrators and individual officers.

Police departments should also adopt language similar to that found in Oakland's Police Department Policy Manual, which identifies the administrator of the ALPR program and requires personnel to be trained in the equipment they operate and trained in all applicable laws, how to safeguard and appropriately access data, and what to do in the event of a data breach.

Police departments should, additionally, add procedures to their police manuals that require department officials to analyze and evaluate their sharing arrangements with other agencies on an annual basis, at minimum. Departments should consider whether requesting agencies have a demonstrable need for the images and whether requesting agencies have appropriate safeguards to store and access the shared data.

### 2. Adopt short retention periods.

In California, retention periods vary by city and county, though they generally range from 60 days to upwards of five years. San Francisco, for example, stores data for one year. The Los Angeles County Sheriff Department, Los Angeles Police Department, and the Los Angeles Port Police all store downloaded data for a minimum of five years. The retention period of ALPR data and images should be as brief as possible. The state of New Hampshire, by contrast, imposes a retention limit of three minutes from the time of capture, except for arrests, or ALPR-identified vehicles that were subject to a missing or wanted person broadcast. This is a reasonable standard.

Thus, the California Legislature should adopt that limit and codify a three-minute retention standard statewide. Long-term historical data pose privacy concerns, and data that are not up to date may produce inappropriate stops. Three minutes is enough time to allow for the spotting of stolen vehicles. A short retention period also would safeguard against an aggregation of data points that would create such a detailed "mosaic" of one's movements that it would constitute a violation of a person's rights to privacy and freedom from unreasonable searches under the Fourth Amendment. Moreover, no one should rely on agencies deleting obsolete or irrelevant records manually. All data that has been downloaded or collected should be automatically purged after three minutes unless there is an articulable reason to believe the data will become legal evidence in a criminal or civil case. In such circumstances, the data should be removed from any server or hard drive that is used for storage and downloaded to a portable device and booked into evidence.

### 3. Clean ALPR databases.

Police departments should practice "data hygiene" proactively by taking steps to maintain accurate records. Responding officers should manually confirm ALPR hits, which is generally required but may not be universally mandatory for all agencies. "Dirty" data inevitably leads to misidentification and potentially false arrests, sometimes at gunpoint. When one agency fails to update its information, other agencies may undertake high-risk felony stops based on that faulty data. All ALPR databases should, thus, be regularly scrubbed to ensure information quality. This includes adopting the following best practices:

- Purging existing data already held by agencies that are not accessed for an investigation
- Immediately updating stolen car registries and ensuring that shared databases contain up-to-date information
- Detecting and removing duplicate records
- Removing typographical errors
- Ensuring that entered data are consistent across department systems and shared databases, and that the data are accurate and complete

### 4. Require periodic impact and assessment reports.

The California Legislature should require agencies that operate ALPRs to perform annual audited reporting to identify their ALPR networks, how the agency uses the data, the costs and demonstrated benefits of using the equipment and the information stored in it, and make the findings available on a public and searchable webpage. Those reports should include the types of data the ALPRs capture; how many cameras are in use; which agencies have access to the data; hardware costs; maintenance expenses; and effectiveness measures, such as stolen vehicles recovered, financial savings, and the number of privacy complaints received and their outcomes. Agencies should ensure that all appropriate safeguards are in place and best practices are followed.

Several California cities already have passed ordinances that require similar reporting. San Francisco Administrative Code, Section 19B, passed in 2019, requires that departments submit to the Board of Supervisors an impact report for each surveillance technology in use. The city already has produced a report for its ALPRs since the code's passage. The Police Department of Piedmont, California, has been a national leader in reporting ALPR data since the implementation of its camera system and now features a transparency portal on its website. After past criticisms of its privacy policies, BART (Bay Area Rapid Transit), a special district of California, promulgated a requirement to produce annual reports on its surveillance technology and request that the supervisory board approve its continued use of the technology. All jurisdictions installing ALPRs should adopt such transparency safeguards.

### 5. Create enforcement mechanisms for bad actors.

To ensure that ALPR practices and policies are conducted in good faith and provide adequate safeguards, Californians should have a right to legal remedies if violations occur. Civilian oversight at the city level, such as a privacy commission, may help to offer guidance and recommendations to city officials before procuring cameras or approving practices. The cities of Oakland and San Diego are notable examples of such commissions. Civilian oversight, however, can be hindered by having poor access to ALPR data, which often is shielded from review for privacy reasons (even from the commission). In addition, the members of the commission may have limited knowledge of proprietary ALPR technology. To ensure the accountability of law enforcement agencies and compensate for what civilian oversight bodies lack, the active involvement of inspectors general and the California State Auditor's Office is crucial. State reviewers should conduct frequent audits of departmental ALPR practices, just as the State Auditor did in its 2020 report.

Furthermore, local governments and law enforcement agencies also should be civilly liable for failure to comply with local ordinances and other applicable laws, with persons who are truly harmed bringing the agency to court under a private right of action. The surveillance ordinances of several California cities already provide for such rights, and their protocols should be adopted elsewhere. In San Francisco,

> Any alleged violation of [the city's surveillance ordinances]…that is not corrected by the Department within 30 days of receipt of the notice, constitutes a legally cognizable basis for relief, and any person affected thereby may institute proceedings for injunctive relief, declaratory relief, or writ of mandate to remedy the violation, in any court of competent jurisdiction to enforce this [ordinance]. An action instituted under this subsection (b) shall be brought against the City.

### 6. Create a guidance document to assist governments and law enforcement agencies.

The ALPR policies of local governments currently are bad mixes of piecemeal ordinances that may conflict with those of other jurisdictions, while some municipalities do not have any policies in place at all. For example, Oakland's mobile ALPR units regularly travel through the nearby cities of Emeryville and Alameda while recording data. Emeryville has adopted its own ALPR policy, while Alameda, until recently, had not. To address the conflict between individual jurisdictions, the California Legislature should require the Department of Justice to create and circulate a memo and accompanying template for all jurisdictions addressing how to deal with conflicting policies. Agencies would be able to copy or adapt it for their unique circumstances.

While it would not be legally binding, such a template could, nevertheless, avoid confusion and encourage responsible ALPR use by outlining the proper administration of an ALPR program with due consideration to the sensitivity of the technology and its data. The memo would explain how California interprets the relevant statutes and compile the related regulations and standards in a single place, thereby enhancing transparency. Although California has few statewide ALPR laws, city police departments should not be recording invasive data from nearby cities. Moreover, the template should incorporate the policy recommendations spelled out above so that responsible use can further be ensured.

After a series of data-security scandals, BART, which uses ALPRs in its parking areas, leads all of California's jurisdictions with respect to ALPR policies, although its thirty-day data retention period is excessively long. On the other side of the spectrum, the residents of Los Angeles suffer perhaps the worst ALPR practices in the nation, with its lack of appropriate use or privacy policies, combined with its practice of retaining data for a minimum of five years.

## Conclusion

Public sector automated license plate readers should not be installed unless all proper safeguards are implemented. Jurisdictions should either do it right or not do it at all. Therefore, because of their failure to meet basic standards for protecting civil liberties and public safety, California cities that have adopted ALPRs, the California Legislature, and the Northern California Regional Intelligence Center have earned the thirteenth *California*

*Golden Fleece® Award* for failing to implement and maintain proper safeguards to protect California residents and visitors.

Allowing drivers and owners of vehicles in California to be subjected to extensive surveillance without proper protections amounts to a breach of the public's trust and the public's right to know what the government is doing.

That is not to say that ALPRs have never been successful in recovering stolen vehicles, generating investigative leads, or assisting law enforcement to solve heinous crimes. It is incumbent on public servants, however, to put proper safeguards and policies in place before adopting ALPRs or any new technology. For ALPRs, key safeguards include (1) limiting database access and use, (2) adopting short retention periods, (3) cleaning ALPR databases, (4) requiring periodic impact and assessment reports, and (5) creating enforcement mechanisms for bad actors.

If jurisdictions install ALPRs with those safeguards in place, the public will have a greater expectation that civil liberties will be protected and public safety enhanced. If, however, future scholarly research demonstrates conclusively that ALPRs yield net negative benefits, i.e., benefits minus costs are negative, even with proper safeguards in place, then taxpayer money and police resources should not be invested in ALPR technology.

## Appendix A: ALPRs, Algorithmic Bias, and Predictive Policing

ALPR cameras that feed or interoperate with criminal forecasting databases pose an especially high risk of reinforcing existing policing biases—and the issue is nearly impossible to rectify. In 2009, the Santa Cruz Police Department kicked off what has been likened to a "data revolution" in policing. Eight years of local crime reports were analyzed to predict the times and locations of future crimes. The math behind the predictions was adapted from models predicting earthquake aftershocks. The underlying assumption was that crimes can be approximate repeats of other crimes nearby in space and time—similar to how aftershocks follow earthquakes. Santa Cruz's "predictive policing" model was touted as the first of its kind.

The Los Angeles Police Department quickly followed suit. Popularizing the practice, the LAPD used predictive analytics to comb large datasets in an attempt to detect trends in crime. The LAPD's interest in the data-driven strategy led them to contract with PredPol, which was founded by the same researcher involved with Santa Cruz's predictive policing program. PredPol promises a "machine-learning algorithm to calculate predictions" about future crimes. The service boasts that it can assist police in reducing crime rates and victimization by providing the who, what, and where of hotspots to direct patrol patterns. A LAPD audit in 2019 admitted they could not determine that the program had "helped reduce crime." Despite this, predictive policing continues to grow rapidly across the nation. A *USA Today* article describes how "... a 2012 survey by the Police Executive Research Forum found that 70% of roughly 200 police agencies planned to implement the use of predictive policing technology in the next two to five years."

The danger of combining ALPR data and predictive policing models is that predictive policing relies on data fed into its algorithm by police officers or police equipment. Any predictive model has a vulnerability to "feedback loops." A feedback loop arises when the output of a model, in this case, a predicted criminal event, gets fed back into the model as training data. A feedback loop amplifies any existing errors in the data. As a popular adage in data science goes, "garbage in, garbage out."

For example, if crimes are reported more frequently in a particular location, police could be dispatched to that area more often. Even if no crime has occurred, a misleading data point created by the ALPR could find its way into the database, thereby training the predictive model to spit out bad results. Over time, the ALPR data will assign too much weight to the location; police would then patrol that area more—picking up more data and perpetuating the cycle.

This risk makes policing less efficient by suggesting that a particular location's crime rate is higher than its actual crime rate, thereby diverting police resources from other areas of need. It also risks propagating biased policing. The resource-allocation problem has become an area of increasing concern. Police have been known to dispatch license plate cameras in religious minority and low-income neighborhoods. If ALPRs are deployed more frequently in certain areas, ALPRs could overweight the data collected from those communities. An algorithm trained on these data may be skewed toward inaccurate and discriminatory predictions.

A common observation in the scholarly literature on policing is racial discrepancies in traffic stops. The Stanford Open Policing Project, an aggregator of millions of traffic stops by law enforcement, finds that "data show that officers generally stop black drivers at higher rates than white drivers, and stop Hispanic drivers at similar or lower rates than white drivers. These broad patterns persist after controlling for the drivers' age and gender."

Some researchers dispute that the discrepancies are indicative of police racial bias and instead suggest that the disproportionate stops could be driven by minorities committing more crimes, or confounding variables such as racial driving patterns. Perhaps Blacks frequently drive when more police are on patrol, leading to more traffic stops. The tendency of Black drivers to be stopped more frequently, however, remains relevant, even in the absence of a demonstrable causal link between bias and police stops, because the actual causal relationship between race and traffic stops could be unknown. It is plausible that feedback loops would lock the disparity in and push police toward racially biased practices.

The same effect of feedback loops has been shown in algorithms designed to assist judges with rating a defendant's propensity to reoffend, inform parole decisions, and assign bond amounts. A ProPublica investigative report discovered that "risk assessment scores" resulted in Black defendants being "77 percent more likely to be pegged as at higher risk of committing a future violent crime and 45 percent more likely to be predicted to commit a future crime of any kind." The algorithm that was responsible for the risk assessment scores, however, failed to accurately classify the defendant's risk of reoffending because the algorithm was trained on biased data. Creating better datasets as a fix is not always possible, especially regarding policing based on historical reports of crimes or records of infractions. Not only can a dataset be too homogeneous, meaning that the training population—which lacks demographic diversity—differs from the general population, but a problem also occurs whenever available historical data are corrupted or biased before being fed into a predictive model. A statistical model could adhere to the highest standards of scientific scrutiny, but that might not matter when individuals on the ground can introduce biased data.

In a much-publicized finding, a team of Stanford University researchers analyzed bodycam footage from Oakland police officers during traffic stops. The team relied on computational linguistics of the transcripts of the bodycam videos, and at a confidence level of 95 percent, the model was able to predict the driver's race accurately. The researchers concluded, "Police officers speak significantly less respectfully to black than to white community members in everyday traffic stops, even after controlling for officer race, infraction severity, stop location, and stop outcome."

Data scientists have suggested that if studies like the Stanford study could identify a particular rate of bias, the weights of the variables could be changed to fix the algorithmic bias. But that solution would not be enough to overcome the potentially erroneous data issues. Not only would such a strategy seemingly accept at face value a degree of officer bias, but law enforcement personnel also might feel the need to overcompensate to correct the shortcomings of a predictive model. Conversely, officers may police in a more lackadaisical manner, believing that the predictive model is doing the heavy lifting and officers need only follow what the model is telling them.

An inherent problem with predictive models is that they can be slow to account for changes in the behavioral patterns of police or the public. Changes can skew the algorithm's output as it needs to reconcile new data with the data on which the algorithm has been trained.

To the extent that police bias exists and is known, it would be implausible to compensate for police bias in a predictive model accurately. If officers are more biased than the algorithm's compensator, worse policing outcomes could result because the extra amount of bias would be discounted. Officers can be less biased than the compensator too. Therefore, if a variable for a particular demographic group is underweighted because of perceived discrimination, then the algorithm's compensator becomes an inequitable advantage for that demographic.

Even if reformers correctly calculated the bias rate, the rate would not be static. Factors such as the hiring of individual officers, the culture of the department, and the socioeconomic status of the city's residents could affect police bias. It is highly unlikely that the bias rate would be the same for police departments in the State of Maine and the State of California, let alone two officers within the same department. As time progresses, this rate would be expected to fluctuate.

Despite the foregoing concerns, police continue to press forward. Lexipol is a private-sector law enforcement consultancy that writes policies for many of California's central police departments. It provides services to some 3,400 agencies, "saturating California, where its clients include more than 90 percent of law enforcement agencies." A post on Lexipol's blog touts ALPRs' part in driving "intelligence-led policing," saying that "ALPR results, triangulated with other statistical data, can reveal patterns of activity associated with criminal events that analysts can use to establish probabilities of crimes and their locations." Its sales pitch continues confidently, "ALPR data [are] used to identify areas where stolen vehicles frequent often uncovering where known drug dealers, parolees and criminal actors live."

## About the Author

**Jonathan Hofer** is a policy research associate with the *California Golden Fleece® Awards* and a policy analyst at the Independent Institute's Center on Entrepreneurial Innovation in Oakland, California. He has written extensively on both California and national public policy issues. He holds a BA in political science from the University of California, Berkeley. His research interests include privacy law, student privacy, local surveillance, and the impact of emerging technologies on civil liberties. He is author of the policy report *Automated License Plate Readers: A Study in Failure*, and his articles on surveillance and policing have appeared in such publications as *Human Events*, *Orange County Register*, *The Hill*, *Towards Data Science*, and *The Daily Californian*.

Each quarter, the Independent Institute highlights a California state or local government spending program, tax, or regulation that fleeces taxpayers, consumers, or businesses. The *California Golden Fleece® Awards* shine a spotlight on waste, fraud, and abuse in California government to give valuable information to the public, enabling them to provide needed oversight and demand meaningful change.

*Fleece Award* winners are announced quarterly on Independent.org and posted on Independent's Twitter, Facebook, LinkedIn, and Instagram accounts. We encourage people—both inside and outside of government—to send us *Fleece* candidates. To learn more and to submit your candidates, go to www.independent.org/cagoldenfleece.