

When Local Government Automates an Accusation

Heather Grizzle · August 25, 2026

A disputed red-light-camera citation reveals how automated enforcement, private vendors, public data, and weak accountability create local AI governance risks.

Contents

1. Observation became attribution
2. Human review may have been narrower than advertised
3. The burden shifted to the resident
4. Public enforcement was routed through private infrastructure
5. Accountability became distributed enough to disappear
6. Where detection becomes power
7. What procurement must govern

What a disputed red-light-camera citation reveals about AI governance, private vendors, public data, and institutional accountability

Novara Consulting Group | AI Governance × Local Government | August 2026

A resident opens an envelope from his city and finds an accusation inside: a photograph of an intersection, a timestamp, an enlarged image of a license plate, and a demand for payment. The vehicle is registered to him. He says he was miles away when the photograph was taken. The instinctive question, and the one around which most disputes of this kind are argued, is whether the camera worked. Was the plate read correctly, was the light truly red, was the equipment calibrated and certified. Those questions matter, and they are also the smallest questions available. Assume, for the sake of the institution's best case, that the answer to every one of them is yes. Assume the camera performed exactly as designed and the process operated exactly as written. The governance problem

survives intact, because the camera never observed the thing the citation asserts. It observed a vehicle. The citation accuses a person.

An automated enforcement program is not a camera. It is a chain of inferences. A sensor observes an event. Software associates the event with a record. An administrative process converts the association into an allegation. The allegation initiates enforcement against a named individual. Each link in that chain is a decision the government has made, whether or not anyone inside the institution experiences it as a decision, and a local government does not escape responsibility for those decisions by inserting an automated system and a private contractor between itself and the public. When technology observes an event, associates it with a record, generates an allegation, and initiates enforcement, every inference in that chain must be governed. The issue is not simply whether the camera accurately read a license plate. It is whether the government can justify how evidence about a vehicle became an accusation against a person.

A disputed traffic citation is a small case with small stakes, which is precisely what makes it useful. At this scale every component of the system can be examined, and the failure pattern the case exposes is not specific to traffic cameras. It is the standard failure pattern of automated decision systems deployed by public institutions. Five failures deserve attention, and none of them is about the camera.

1. Observation became attribution

The system reportedly established that a vehicle entered the intersection against the signal. If that is accurate, it is a genuine evidentiary accomplishment: a time-stamped photographic record of an event, produced without an officer present. But the event the system recorded is not the conduct the citation alleges. The system established what a vehicle did. It established nothing about who was driving. The enforcement process treated registered ownership as sufficient to initiate an accusation, and in doing so it converted a fact about an object into a claim about a person without ever acquiring the evidence that conversion requires.

That conversion is not a technical step. It is a legal and institutional one, and in most photo enforcement regimes it is accomplished by statute rather than by proof. Owner liability rules exist because identifying drivers is difficult and citing owners is administratively easy, and legislatures in many states have made that trade deliberately, usually by defining the violation as civil and attaching liability to the registration record. A government may be entitled to make that trade. It is not entitled to make it invisibly. When a citation arrives styled as a finding about the recipient's conduct while the underlying record supports only a finding about the recipient's property, the institution has overstated its own evidence. This is a recurring failure across automated governance generally: systems produce facts about objects, records, and proxies, and institutions quietly promote those facts into claims about people.

2. Human review may have been narrower than advertised

Programs of this kind are routinely defended with a sentence that sounds like governance: an officer reviews every citation before it is issued. The sentence deserves a follow-up question it rarely receives, which is what, exactly, the officer reviews. If the review consisted of confirming that the plate in the photograph matches the vehicle in the registration database, then the human validated the one link in the chain that the machine handles best and left untouched the one link the machine cannot address at all, namely the identity of the driver. The material inference, the inference on which liability actually turns, passed through the process unexamined by anyone.

"Human in the loop" is a governance claim, not a staffing description. It carries weight only when the human is positioned, authorized, and equipped to evaluate the inference that matters, and holds genuine discretion to stop the process when the evidence does not support it. A

reviewer whose function is confined to confirming what the machine has already established is not performing oversight; the reviewer is performing ratification.

Institutions that advertise human review owe the public a precise account of what the human examines, what the human is empowered to decline, and how often the reviewer actually declines. Where that account cannot be produced, the phrase is doing reassurance work that the process itself is not doing.

3. The burden shifted to the resident

The third failure follows from the first two. Because the system could not establish who was driving, and the review did not attempt to, the program resolved the identity question by transferring it. Rather than the government establishing who committed the violation before making an accusation, the registered owner was required to enter an administrative process, supply information, and redirect responsibility, in some regimes by identifying the actual driver and in others by documenting his own absence. Contestability existed, but it existed only after the system

had already acted, and the default consequence of doing nothing was liability. The sequence matters. An accusation that must be actively dislodged is a different exercise of power than an accusation that must first be proved.

Courts have recognized this structure before. In *State v. Kuhlman* (Minn. 2007), the Minnesota Supreme Court invalidated Minneapolis's photo enforcement ordinance, holding that it conflicted with the state's uniform traffic laws and reasoning that its presumption that the registered owner was the driver eliminated the presumption of innocence and shifted the burden of proof onto the accused. The lesson of that decision is not that automated enforcement is inherently unlawful. It is that burden allocation is a governance decision of the first order. Deciding who must prove what, at what cost, on what timeline, and with what consequence for silence is a decision about the relationship between the state and the people it governs. It cannot be inherited by default from a vendor's workflow, and residents should not discover how it was made for the first time when the envelope arrives.

4. Public enforcement was routed through private infrastructure

Every element of the accusation resides somewhere, and in a vendor-operated program much of it may reside outside the government. Images, plate reads, timestamps, registration lookups, violation records, payment data, and the contents of residents' challenges may pass through, or be stored on, systems the contractor builds and controls. The public is entitled to precise answers to unglamorous questions: who can access the information and under what authorization; where it is stored and under which jurisdiction's law; how long it is retained, and whether retention tracks the disposition of the citation or the convenience of the vendor; whether it may be reused, aggregated, analyzed, or sold in any form; and what becomes of it when the contract ends, when the vendor is acquired, or when the vendor fails. A resident cannot opt out of being photographed by an enforcement program at a public intersection. Data generated under that kind of compelled exposure is an artifact of public authority, and

its governance cannot default to the terms of a commercial agreement the public has never seen.

These are not hypothetical anxieties in this particular market. Photo enforcement has a documented history of misaligned incentives, including compensation structures that paid contractors per citation issued, and of procurement corruption, most prominently the bribery scheme surrounding Chicago's red-light program, which produced federal convictions. Texas responded to years of controversy by banning red-light cameras statewide in 2019. None of this establishes that vendors are bad actors. It establishes that vendors are vendors: private firms with commercial incentives, operating inside a function that is supposed to answer to entirely different incentives. That is a condition to be governed rather than an accusation to be litigated, and it is governable only in the contract.

5. Accountability became distributed enough to disappear

Ask who is responsible for the accusation and the program answers in fragments. The camera detected the event. The vendor operated the system. An officer checked the plate. The city authorized the program. State law structured the citation. Each participant can truthfully report that its own limited function performed as designed, and each is correct within its own boundary. Yet the outputs of those functions combined into a governmental accusation against a specific person who says he was miles away, and no single participant owns the composite act.

This is the quiet mechanism by which automation dissolves accountability, not through malice or malfunction but through decomposition. Responsibility is divided into components until it exists nowhere in particular.

Governance exists to prevent precisely this outcome.

Somewhere in the institution there must be an office, occupied by a person with a name, that is accountable for

the claims the system makes about members of the public, empowered to answer for those claims, and obligated to correct them. If no one inside the program can say that the accusation was theirs, the government has built a machine for producing accusations that no one is willing to own.

Where detection becomes power

The system may have accurately identified a vehicle while the institution inaccurately attributed conduct to a person. That is not merely a technical distinction. It is the point at which automated detection becomes governmental power. Nearly everything that procurement processes currently evaluate lives on the detection side of that line: equipment accuracy, vendor certifications, calibration schedules, uptime commitments, unit price. Nearly everything that matters institutionally lives on the attribution side: what the evidence can support, which inferences the government permits itself to draw from machine output, who bears the burden when an inference is contested, and who answers when it is wrong. A procurement that scrutinizes the camera and assumes the accusation has

evaluated the product and ignored the power.

None of this turns on whether the components involved qualify as artificial intelligence. Plate recognition is machine vision; much of the rest may be ordinary software and administrative workflow. The governance obligation attaches to the inference chain, not to the sophistication of the code, which is why the pattern generalizes. The same chain, observation to association to allegation to enforcement, sits inside gunshot detection, benefits fraud detection, automated license suspension, and every comparable system a local government buys. Michigan's automated unemployment insurance system, which falsely accused tens of thousands of residents of fraud before audits and litigation caught up with it, was this same architecture operating at catastrophic scale. The red-light camera is simply the version small enough to see whole.

What procurement must govern

Local governments are not merely purchasing cameras, algorithms, or administrative platforms. They are

purchasing systems that participate in public authority, and procurement is where that participation is defined. It is also the last moment at which the government holds full leverage. Before signature, every term is available. After deployment, governance becomes remediation.

Procurement must therefore govern more than performance and price. It must govern evidence, by requiring the contract and the citation itself to state precisely what the system can and cannot establish. It must govern inference, by specifying which conversions from machine output to legal claim are permitted and who authorizes each one. It must govern data, by settling custody, access, retention, reuse, security, and exit before the first image is captured. It must govern human review, by defining what the reviewer must independently determine and vesting the reviewer with real authority to stop issuance. It must govern contestability, by accounting honestly for the burden the program places on residents and defending that allocation in public. It must govern vendor incentives, by decoupling compensation from enforcement volume. And it must govern responsibility, by

naming the office that answers for what the system does to the people it touches.

None of this is utopian, because some of it is already law.

When Minnesota returned to automated enforcement in 2024, nearly two decades after Kuhlman, it did so through state legislation authorizing a narrowly conditioned four-year pilot rather than through local ordinance. The statute requires a warning for a first violation, mandates advance signage and published camera locations, requires independent evaluation and reports to the legislature, prohibits paying private contractors based on the number of violations or citations issued, and requires the implementing government to audit its vendor's data practices. Whatever one concludes about the wisdom of the program, the instrument is instructive. The difference between the ordinance the Minnesota Supreme Court struck down and the pilot the legislature authorized is not the camera. It is the governance written around the camera, before deployment, in enforceable terms. That is the discipline Novara Consulting Group practices as Operational AI Governance™: converting governance

principles into instruments a procurement office can actually execute while the government still holds the pen.

A government may outsource the technology. It cannot outsource its duty to justify what that technology does to the public.