

Article

Automated License Plate Readers and Data Centers as Networked Mass Surveillance Infrastructure: The Systemic Erosion of Privacy and Free Expression

Haris Alibašić 

Department of Business Administration, Lewis Bear Jr. College of Business, University of West Florida,
Pensacola, FL 32514, USA; halibasic@uwf.edu

Highlights

Please indicate how your work links to systems science via your contributions to systems practice, theory, and/or methodology.

- Conceptualizes automated license plate reader networks as cross-vendor socio-technical systems linking sensors, cloud data centers, databases, interagency access, police action, and democratic oversight.
- Introduces the Rights Control Deficit and rights-preserving hybrid intelligence as systems approaches for diagnosing unmatched surveillance capabilities and designing institutional counter-controls.

What are the main findings and/or the implications of the main findings?

- Infrastructure aggregation, authority diffusion, asymmetric feedback, and rights invisibility create conditional risks of privacy loss and expressive chill; they do not establish a universal ALPR-specific causal effect.
- Meaningful human oversight is an institutional control property requiring authorization, minimization, audit, contestability, continuity, sanctions, and termination authority—not merely an officer’s presence at an interface.

Abstract

Automated license plate readers (ALPRs) are often evaluated as discrete police tools, although their public power arises from cross-vendor socio-technical infrastructure. This article examines roadside and mobile sensors, vehicle-attribute classification, cloud archives, commercial databases, real-time crime center integration, interagency access, automated alerts, and police action. Flock Safety supplies the principal documentary case because unusually extensive public records permit system-level tracing; Axon/Fusus, Motorola Vigilant/VehicleManager, and federal access to commercial ALPR data establish the wider vendor-independent boundary. A structured documentary analysis of 59 sources triangulates official records, peer-reviewed research, vendor materials used only for stated functions, and record-based investigations. It integrates boundary critique, control-structure mapping, feedback analysis, constitutional doctrine, a STRIDE-informed threat model, and empirical research on policing effectiveness and surveillance effects through 3 August 2026. The analysis identifies four conditional mechanisms: infrastructure aggregation, authority diffusion, asymmetric feedback, and rights invisibility. The article reformulates the Rights Control Deficit (RCD) as a non-arithmetic profile relation between operational demands and effective governance capacity and applies it to three documented configurations and a clearly labeled normative benchmark. Seven falsifiable propositions specify variables, indicators, suitable methods, and disconfirming conditions for later empirical study. A



Academic Editor: Sitalakshmi Venkatraman

Received: 18 July 2026

Revised: 8 August 2026

Accepted: 14 August 2026

Published: 18 August 2026

Copyright: © 2026 by the author.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

rights-preserving hybrid-intelligence architecture combines bounded automation with judicial authorization, short retention, sensitive-location protections, immutable audit, availability safeguards, independent review, contestability, sanctions, and credible termination authority. The evidence identifies capabilities, activated pathways, and conditional risks; it does not estimate population prevalence or a universal ALPR-specific causal effect. Meaningful human oversight is an institutional control property, not merely an officer's presence at an interface.

Keywords: automated license plate readers (ALPR); Axon Fusus; Flock Safety; Motorola Vigilant; mass surveillance; hybrid intelligence; privacy; freedom of expression; socio-technical systems; public administration

1. Introduction

A roadside camera appears modest: it photographs a passing vehicle, translates a plate, records time and location, and transmits the event to a server. That description is technically accurate and systemically incomplete. Contemporary ALPR infrastructure joins fixed and mobile sensors with machine-learning classification, communications networks, cloud storage, commercial repositories, searchable histories, hotlists, user permissions, cross-agency sharing, real-time crime centers, police workflows, procurement contracts, and legal authority [1–6]. The relevant object is therefore an information-and-action infrastructure rather than a camera.

The vendor landscape makes this distinction important. Flock Safety (Atlanta, GA, USA) markets fixed ALPR devices and searchable vehicle attributes [2,3]. Axon Enterprise, Inc. (Scottsdale, AZ, USA) combines fixed ALPR, livestreaming, and vehicle-attribute recognition with its Fusus real-time intelligence platform, which supports search, alert, export, hotlist, and audit functions across supported data sources [4,5]. Motorola Solutions, Inc. (Chicago, IL, USA) markets Vigilant VehicleManager with partial-plate, location, association, and predictive vehicle-location analytics across local and shared detections [6]. These products differ in architecture, contracting, database reach, and controls. Each nevertheless lowers the cost of converting transient public-road observations into persistent and actionable vehicle intelligence.

Flock remains the principal documentary case, not because it exhausts the field or presents the most intrusive imaginable configuration, but because its rapid municipal diffusion has generated an unusually dense public record: vendor policies, local contracts, audit findings, litigation, public-record investigations, statutory disputes, and termination decisions. The cross-vendor comparison prevents a product-specific critique. It also shows why reform limited to one company would leave the underlying governance problem intact.

The public-safety case is material but empirically bounded. Experimental and quasi-experimental research has found limited or inconsistent deterrent effects from ALPR patrols [7,8]. Large-scale investigative deployments may improve selected outcomes, yet causal attribution and cost-effectiveness remain uncertain [9,10]. Organizational research further shows that investigators reinvent broadly available and interoperable technology for uses beyond its original patrol function [11,12]. A credible rights analysis should acknowledge validated recoveries and investigative assistance while asking whether population-scale capture, expansive historical search, and cross-jurisdictional access are necessary for those benefits.

Socio-technical systems research locates outcomes in relations among people, organizations, technologies, rules, incentives, and environments [13]. Hybrid intelligence

similarly concerns the structured integration of computational speed with contextual human judgment, public values, institutional responsibility, and learning [14]. Under this view, the central governance question is not whether a person clicks before enforcement. It is whether the entire control structure preserves authority over purpose, collection, access, interpretation, coercive action, review, correction, continuity, and termination.

Privacy and free expression meet at movement data. Repeated sightings may expose visits to clinics, houses of worship, political organizations, protests, union offices, newsrooms, shelters, and law offices. Contextual integrity asks whether information flows remain appropriate to their actors, purposes, and transmission principles [15]. Surveillance and big-data policing scholarship explains how aggregation sorts populations, reallocates police attention, and converts routine collection into intelligence infrastructure [16–18].

Surveillance effects need careful causal language. Experimental and observational research links perceived monitoring to self-censorship and avoidance in other digital settings [19–22]. An ALPR survey experiment also reported consequences for trust in law enforcement [23]. These studies do not establish the magnitude of an ALPR-specific participation effect in every jurisdiction. They support a plausible mechanism and a research agenda, not a universal causal estimate.

Constitutional doctrine likewise turns on aggregation and system boundary. *United States v. Jones* recognized concerns arising from technologically assisted movement aggregation [24]. *Carpenter v. United States* treated access to historical cell-site location information as a search because aggregated location records reveal intimate patterns [25]. *Chatrie v. United States* held in June 2026 that police acquisition of Google Location History information was a Fourth Amendment search, while leaving warrant sufficiency to the lower court [26]. The decision did not decide ALPR legality, but it rejects the assumption that third-party cloud custody or a short query window automatically removes constitutional protection.

Lower courts have treated current ALPR configurations more narrowly. A federal district court upheld Norfolk's present Flock deployment in *Schmidt* [27], while amici on appeal urge an ecosystem-level analysis [28]. Leaders of a Beautiful Struggle treated persistent aerial reconstruction of movement as constitutionally significant [29]. Virginia's Court of Appeals held in *Robinson* that Norfolk's then-current 172-camera, 30-day configuration did not constitute an unreasonable search, emphasizing its incomplete coverage [30]. California's *Bartholomew* decision separately held that collection and use without the privacy policy required by state law constituted actionable harm [31]. These authorities show that legality depends on architecture, facts, jurisdiction, and statutory duties rather than the label attached to a camera.

Three research questions guide the article: (1) How do sensors, cloud platforms, commercial repositories, interoperability, organizational permissions, and human operators couple to create capability beyond each component? (2) Through which mechanisms does that infrastructure create risks to privacy, expression, association, and participation? (3) What rights-preserving hybrid-intelligence architecture would constrain drift while retaining proportionate public-safety uses?

The article makes four contributions. It specifies a vendor-independent ALPR system boundary; reformulates the RCD as a non-arithmetic profile diagnostic grounded in requisite variety; translates seven mechanisms into falsifiable propositions; and develops a control architecture that integrates authorization, minimization, verification, security, availability, audit, contestability, and exit. Claims are calibrated to a theory-building documentary design. Observed incidents establish activated pathways. They do not establish prevalence or a causal population effect.

2. Materials and Methods

2.1. Research Design and Unit of Analysis

This is a theory-building systems study supported by structured documentary analysis [32–37]. It does not estimate a treatment effect or validate a psychometric scale. General systems thinking supplies boundaries, interdependence, emergence, feedback, and delay [32,33]. Cybernetics supplies requisite variety: a controller requires sufficient response variety for the consequential states of the process it seeks to regulate [34]. System-theoretic safety analysis redirects attention from isolated component failure to unsafe interactions and inadequate control [35]. Critical systems heuristics asks who defines purpose, who benefits, who bears costs, and who lacks voice [36]. Document analysis supplies a transparent method for assessing legal, administrative, technical, scholarly, and public records [37].

Ashby's term variety refers to the range of states or responses available to a system or controller [34]. It is not statistical variability, dispersion, variance, or sampling error. The manuscript uses variability only when discussing empirical variation across observations; it uses variety only in the cybernetic sense.

The unit of analysis is the networked ALPR governance system: procurement; placement; observation; automated classification; transmission; cloud storage; commercial database access; matching and historical search; hotlists; interagency sharing; real-time crime center integration; human interpretation; police action; retention; audit; complaint; judicial review; legislative oversight; continuity; and termination. Flock supplies the primary case. Axon/Fusus, Motorola Vigilant/VehicleManager, and federal commercial-database access test whether the identified mechanisms depend on a single vendor [1–6].

The analytic boundary is functional rather than contractual, jurisdictional, or vendor-specific. A component is inside the boundary when it materially changes capture, persistence, inference, access, action, feedback, or the authority to constrain or terminate those functions, regardless of ownership. The boundary therefore includes edge devices, networks, vendor and commercial repositories, integrated platforms, partner agencies, human decision points, and oversight institutions. Unconnected closed-circuit television, one-time visual plate checks that create no retained searchable record, and general infrastructure with no ALPR-specific control role remain outside.

Four key constructs follow this boundary rule. Operational capability means consequential states a documented configuration can reach. An activated pathway is a documented use or failure and does not establish frequency. A rights effect is an observed change in privacy, expression, association, participation, or legal position supported by evidence appropriate to that claim. Effective governance is a legally authorized, implemented, inspectable control able to prevent, detect, correct, or terminate the corresponding system state.

2.2. Documentary Review Protocol

The revision converted the original purposive corpus into an auditable structured documentary review. The audit covered materials published through 3 August 2026. Discovery used the U.S. Department of Justice and National Institute of Justice libraries; DHS, NIST, Supreme Court, state legislative, attorney-general, and municipal repositories; publisher platforms including SpringerLink (Springer Nature, Cham, Switzerland), Emerald Insight (Emerald Publishing, Leeds, UK), SAGE Journals (SAGE Publications, Thousand Oaks, CA, USA), Taylor & Francis Online (Taylor & Francis, London, UK), and MDPI (Basel, Switzerland); court-document mirrors where official docket text was not readily indexed; vendor documentation; and targeted web discovery using exact titles, docket numbers, DOI strings, and domain-restricted queries. Supplementary Material S1 reports the repository families, reconstructed search families, dates, source classifications, source inventory, central claim crosswalk, domain-level sensitivity display, and case-comparability record.

Search strings combined an ALPR term—"automated license plate readers," "automatic license plate reader," "ALPR," "Flock Safety," "Axon Fusus," "Motorola Vigilant," or "VehicleManager"—with one or more governance terms: cloud, commercial database, retention, sharing, audit log, hotlist, protest, abortion, immigration, privacy, First Amendment, Fourth Amendment, chilling effect, public trust, system, requisite variety, STRIDE, threat model, continuity, or availability. Legal searches used party names, docket numbers, statute numbers, and official-report identifiers.

The starting manuscript contained 44 cited records. Corrective searches identified 23 additional candidate records. Two duplicates were removed, three records were excluded at title or summary screening, and three were excluded after full-text review because they were superseded, outside the system boundary, or insufficiently traceable. The final corpus contained 59 records. Figure 1 documents this flow. This accounting describes a corrective documentary audit; it is not presented as PRISMA-compliant exhaustive evidence synthesis.




Figure 1. Corrective documentary review and reference-audit flow. Counts describe the revision audit rather than an exhaustive systematic review.

Inclusion required evidence relevant to at least one of six functions: technical architecture and reach; effectiveness or public-safety outcomes; access, retention, sharing, and audit; privacy and expression pathways; legal authority; or governance and systems design. Exclusions covered duplicate reporting, generic commentary without inspectable support, obsolete product claims displaced by current documentation, and materials unrelated to networked vehicle-location data.

Each source received a role and quality classification. Tier A comprised enacted law, official opinions, government audits, and official technical standards; Tier B comprised peer-reviewed research and academic books; Tier C comprised vendor documentation used only for product functions and stated policies; Tier D comprised investigative journalism and civil-society investigations with disclosed records or methods. Independence, methodological transparency, claim proximity, corroboration, and currency were recorded separately. Vendor claims never supplied independent effectiveness evidence. Tier D

sources established documented allegations, search-log findings, or public controversy and were triangulated where official records existed.

The author performed the coding. No independent coder was available, and the study does not claim inter-coder reliability. Within that constraint, the revision used source-role triangulation: central claims were checked, where functionally distinct evidence was available, against official law, court, audit, or technical records; peer-reviewed research; vendor documentation limited to stated functions; and record-based investigations. Claims resting on one source role were worded narrowly, and contradictory evidence remained visible. Auditability is further supported by explicit anchors, a central claim-to-source crosswalk, retained uncertainty labels, and a domain-level sensitivity display for contested RCD judgments. This procedure assesses corroboration and inference discipline, not coder agreement.

The protocol is retrospective and corrective rather than preregistered. Repository interfaces, indexing, and result counts change over time, and contemporaneous database hit counts and timestamps were not retained for the original drafting searches. They are therefore not reconstructed as if they were contemporaneous records. The review is designed to make source selection and analytical use inspectable; it does not claim the exhaustiveness or search replicability of a systematic review.

2.3. Analytical Procedure

Analysis proceeded in six passes. Boundary mapping identified components, actors, data flows, trust boundaries, authority flows, and affected populations. Control-structure mapping identified controllers, controlled processes, control actions, feedback, and responsibility gaps. Feedback analysis distinguished reinforcing signals favoring expansion from balancing signals supporting correction or restraint. Legal and empirical evidence was then mapped to privacy, expression, public trust, and enforcement pathways. A STRIDE-informed security analysis assessed spoofing, tampering, repudiation, information disclosure, denial of service, and elevation of privilege [38]. Finally, hybrid-intelligence principles were translated into institutional and technical controls and compared with documented failure modes.

The analysis applied an explicit causal-evidence rule. Architecture establishes what a system permits; a documented incident establishes that a pathway has been activated; only a suitable empirical design can estimate frequency, magnitude, or causation. Accordingly, permits, enables, creates a pathway, and presents a risk are used for system-level claims, while causes, reduces, or deters are reserved for findings supported by an appropriate empirical design. This rule also prevents vendor descriptions of capability from serving as evidence of effectiveness or behavioral effect.

The seven propositions are outputs of theory building rather than tested hypotheses. They specify constructs, units, indicators, suitable designs, expected relationships, and disconfirming evidence. Expert interviews, Delphi panels, multi-site audits, and community studies are appropriate next-stage validation methods; adding them after peer review would constitute a new empirical study rather than a corrective clarification of this design.

The propositions are specified at the construct level appropriate to theory building. Study-specific confounders, measurement windows, comparison groups, and minimum evidence thresholds depend on the jurisdiction, platform configuration, data-generating process, and outcome selected for a later empirical design; those elements should be preregistered when such a study is undertaken rather than invented without a dataset or sampling frame.

2.4. Rights Control Deficit: Definition and Coding

The RCD asks a practical question: in which life-cycle domains can an ALPR system reach consequential states that its effective controls cannot match? It is a profile, not a sub-

traction score. Let D contain four paired domains. For each domain d , s_d is the documented operational demand and g_d is the effective governance response, each described as low, moderate, or high. The unmatched-domain set is:

$$D_{\text{RCD}}(S, G) = \{d \in D \mid s_d \succ_{(d)} g_d\}. \quad (1)$$

The relation $\succ_{(d)}$ means that documented operational demand in domain d exceeds the effective response range of the corresponding governance controls. Because the output is a set, domains are not added, averaged, weighted, or offset. A strong audit program cannot compensate for absent lawful authorization; short retention cannot compensate for uncontrolled external access.

The four domains follow the ALPR life cycle and its control hierarchy: (D1) capture matched to authorization and minimization; (D2) persistence and inference matched to retention enforcement and elevated authorization; (D3) access, sharing, and integration matched to identity governance, audit, and jurisdictional control; and (D4) action and error propagation matched to verification, remedy, continuity, sanctions, and exit. Pairing demand with response makes the control relationship visible and avoids treating twelve unpaired dimensions as a validated scale.

Applied use requires four steps: (1) map documented reachable states; (2) rate operational demand in each domain; (3) rate the implemented and inspectable governance response in the paired domain; and (4) flag any domain in which demand exceeds response. Low, moderate, and high limit false precision because the public record does not support interval measurement. Borderline judgments remain ranges, such as moderate–high, and the sensitivity check moves each uncertain classification one level toward the least-deficit interpretation before asking whether the mismatch remains.

Table 1 states the anchors. Table 2 applies them to Los Angeles, Norfolk, and Austin as documented configurations and to a bounded-local normative benchmark. The benchmark is not an empirical case. The comparison tests descriptive discrimination, not external validity.

Table 1. Paired ordinal anchors for the Rights Control Deficit.

Paired Domain	Operational-Demand Anchor, s_d	Governance-Response Anchor, g_d
D1 Capture ↔ authorization/minimization	Low: targeted sites or short event windows; non-hit capture minimized. Moderate: broader coverage with documented necessity review. High: dense, mobile, private, or multi-source coverage.	Low: procurement-only approval. Moderate: public policy and periodic approval. High: impact assessment, alternatives analysis, mapped placement, sensitive-location rules, and sunset.
D2 Persistence/inference ↔ enforceable limits	Low: transient known-plate matching. Moderate: short archive and limited attribute search. High: extended retention, reverse/pattern search, association, or predictive analytics.	Low: policy text or user attestation. Moderate: supervisory approval and partial automation. High: technically enforced deletion, case preservation, judicial/elevated review, deletion proof.
D3 Access/sharing/integration ↔ identity/audit control	Low: named local users, sharing off. Moderate: bounded partners and limited integrations. High: cross-jurisdictional, commercial, or multi-platform access.	Low: logs unreviewed or stale accounts. Moderate: internal sampling and periodic recertification. High: immutable logs, independent risk-based audit, rapid revocation, restrictions traveling with data.
D4 Action/error propagation ↔ verification/remedy/continuity/exit	Low: lead only; no coercive action from unverified alert. Moderate: documented verification with exceptions. High: automated escalation, stale hotlists, downstream reuse, or dependency on remote service.	Low: informal correction and weak exit. Moderate: supervisory correction and feasible continuity. High: stop rules, notice/appeal, sanctions, fail-safe degradation, tested recovery, pause and termination authority.

Table 2. RCD profiles for three documented configurations and one normative benchmark.

Configuration	Operational Profile	Governance Profile	RCD Reading and Sensitivity	Status/Evidence
Los Angeles, 2025–2026	D1 high; D2 high; D3 high; D4 high	D1 low–moderate; D2 low; D3 low; D4 moderate–high because termination was exercised	D1–D3 remain unmatched; D4 narrows only at exit. Least-deficit recoding leaves D1–D3 unmatched.	Documented audit [39,40]
Norfolk, 2024–2026	D1 moderate; D2 moderate; D3 moderate–high; D4 moderate	D1 moderate; D2 low–moderate; D3 uncertain; D4 moderate	D2 and potentially D3 unmatched; classification remains legally contested.	Documented litigation [27–30]
Austin, through June 2025	D1–D3 moderate–high before cancellation; D4 moderate	D1–D3 low–moderate; D4 high at termination	Pre-cancellation mismatch narrowed through non-renewal rather than demonstrated fine-grained control.	Documented exit episode [41]
Bounded-local benchmark	D1 low; D2 low; D3 low; D4 low	All domains high	No unmatched domain by design; benchmark is normative, not empirical.	Normative synthesis [42–45]

The documented configurations were selected purposively for maximum variation in scale, legal posture, oversight evidence, and termination status: Los Angeles supplies a large-deployment inspector-general audit; Norfolk supplies active constitutional litigation and appellate analysis; and Austin supplies a publicly documented non-renewal episode. They do not form a probability sample, share a uniform observation window, or support controlled causal comparison. Their function is to test whether the same paired-domain coding rules discriminate among materially different governance configurations while preserving uncertainty where the public record is incomplete.

The present RCD application is illustrative and theory building. Repeatability refers to transparent application of stated anchors and dominance rules, not to demonstrated inter-coder reliability, external validity, or predictive validation. Those forms of validation require a new multi-coder, multi-jurisdiction empirical study.

2.5. Scope, Limitations, and Research Integrity

Public records do not expose every query, account, integration, backup, contract amendment, or enforcement outcome. Cross-jurisdictional configurations evolve rapidly. The case profiles therefore include uncertainty rather than inferred precision. The study does not estimate misuse prevalence, ALPR accuracy, crime effects, or an ALPR-specific participation effect. It also does not claim that Flock, Axon, or Motorola systems are identical.

The original reference list was audited in full during revision. The placeholder in the original reference list was replaced with the published Bartholomew citation [31]. Citation numbering was rebuilt in order of first appearance; URLs, dates, decisions, DOI strings, publication status, figure placement, table cross-references, and the author’s 2026 Springer book record were checked against the identified sources. Supplementary Material S1 preserves the claim-level audit trail.


Generative artificial intelligence disclosure: The original submitted manuscript was authored by the author. During revision, an AI-assisted research and document-production environment was used to help locate possible citation inconsistencies, organize the response matrix, support language editing, and format files. It was not treated as an evidentiary source, author, or autonomous verifier. The author checked each retained source against

the linked publication or official record, decided every substantive revision, and accepts full responsibility for the analysis, citations, figures, and text.

3. Analytical Findings

3.1. Cross-Vendor ALPR Infrastructure and System Boundary

Figure 2 maps the boundary recovered from the corpus. Edge devices capture vehicles indiscriminately within their field of view. Classification converts images into plates and attributes. Networks transfer event records to cloud services or repositories. Platforms index histories, execute queries, distribute alerts, and connect partner data. Users interpret outputs and initiate stops, searches, arrests, reports, or intelligence products. Oversight institutions receive partial feedback unless contracts, law, audit access, and public reporting open the system.



Balancing feedback requires audit visibility, corrective authority, sanctions, and exit

Figure 2. Conceptual vendor-independent ALPR system boundary, principal data flows, and control loops. Arrows from governance controllers represent authorization, constraint, audit, correction, and termination; the upper return path represents expansion pressure generated by visible operational outputs. The diagram is a systems abstraction, not an empirically validated implementation map of every vendor configuration.

Cross-vendor evidence differentiates four capability patterns. Flock emphasizes distributed fixed capture, cloud retention and search, customer-managed alerts, and sharing [2,3]. Axon Outpost adds livestreaming and vehicle-attribute recognition, while Fusus can unify supported sensor feeds in a real-time operational view [4,5]. Motorola VehicleManager emphasizes access to billions of detections, partial-plate and location search, association analysis, predictive vehicle-location analytics, and agency-defined retention and sharing [6]. DHS documentation shows that federal users can query vendor-owned commercial LPR data without owning the originating cameras [46]. These are documented product or access functions, not independent effectiveness findings. Together they show that capture, fusion, inference, and repository reach vary by platform while the governance problem crosses brands.

Flock is analytically useful because cloud retention, customer-managed sharing, alerts, audit logs, and public controversies are unusually visible [2,3]. It is not the upper bound of surveillance capability. Systems combining ALPR with live video, real-time crime centers, commercial repositories, or association analytics potentially occupy a larger state

space. The RCD therefore evaluates reachable functions and controls rather than vendor reputation. Table 3 summarizes the system layers, operational functions, emergent risks, and corresponding control questions.

Table 3. System layers, operational functions, rights risks, and control questions.

Layer	Operational Function	Emergent Rights/Security Risk	Control Question
Edge capture	Observe plates, vehicles, place, and time	Indiscriminate collection; unequal placement; sensitive-location exposure	Is coverage necessary, proportionate, mapped, public, and time-bounded?
Classification and hotlists	Translate plates/attributes and generate matches	Misread plate; stale or overbroad list; automation bias	What confidence, source-age, verification, expiration, and stop rules apply?
Cloud and commercial repositories	Index, retain, join, search, back up, and delete	Retrospective tracking; breach; secondary use; deletion failure	Who controls keys, retention, backups, deletion proof, and model-training use?
Integration and sharing	Connect agencies, vendors, video, cases, and real-time centers	Authority diffusion; policy shopping; purpose drift; national reach	Which users, jurisdictions, purposes, and datasets receive technically bounded access?
Human workflow and police action	Interpret leads and exercise public authority	Confirmation bias; erroneous stop; informal workaround; downstream persistence	Who verifies source, context, lawful basis, and correction before coercive action?
Governance and continuity	Authorize, inspect, correct, recover, sanction, and end	Symbolic oversight; unavailable service; unsafe fallback; contractual lock-in	Are independent audit, contestation, recovery, safe degradation, and shutdown effective?

3.2. Four System Mechanisms

Infrastructure aggregation. ALPR capability depends on complementarities among coverage, retention, attribute resolution, query breadth, database size, and interoperability. The term nonlinear does not claim an estimated functional form. It denotes a positive interaction: the marginal significance of additional coverage rises when records persist and remain searchable; longer retention matters more when coverage is dense; attribute search and association analysis become more consequential when repositories cross jurisdictions. This is a testable interaction hypothesis rather than a measured exponent.

Authority diffusion. Ownership is only one control dimension. Defaults, user provisioning, partner access, backup custody, analytics, contract interpretation, and sanction authority may be divided among vendors, local agencies, private camera owners, state institutions, and federal users. A local vote does not govern every later query. Restrictions that do not travel with data become weaker as origin and use separate.

Asymmetric feedback. Recoveries, arrests, and investigative assists are vivid and rapidly counted. Uninvolved exposure, false-alert burden, avoided participation, erroneous inference, and lost trust are delayed, distributed, or absent from operational dashboards. This information asymmetry reinforces expansion unless independent institutions produce balancing signals with authority to alter configuration.

Rights invisibility. Platforms represent plates, vehicle attributes, time, location, users, alerts, and searches. They do not inherently represent why a person entered a clinic, met a journalist, attended worship, joined a protest, or visited counsel. Governance has to supply context omitted by the data model. Without that context, even accurate technical processing can create constitutionally and democratically significant risks.

3.3. Democratic Chill: Operational Definition and Evidence Boundary

Democratic chill is defined here as an observable reduction, avoidance, or alteration in lawful speech, association, assembly, worship, press contact, health-care access, or civic participation linked to awareness or reasonable belief that participation generates durable and shareable movement records. Indicators include attendance change, route or transportation substitution, willingness to express minority views, organizational membership or contact, avoidance of sensitive locations, perceived safety in participation, and trust in public institutions. For empirical use, the construct requires both observed behavioral change and evidence linking that change to perceived surveillance; awareness or coverage alone is not sufficient.

This definition separates mechanism from measurement. Existing surveillance studies support self-censorship and avoidance pathways [19–22], while ALPR research addresses trust and legitimacy [23]. NAACP v. Alabama establishes the constitutional importance of associational exposure [47]. Yet no included study estimates a Flock-, Axon-, or Motorola-specific causal effect on protest attendance. The article therefore uses potential democratic chill when discussing ALPR outcomes and reserves democratic chill without qualification for the defined construct.

Direct evidence shows that the pathway is operationally available. A civil-society analysis of disclosed Flock search logs reported protest-related searches by more than 50 agencies [48]. That finding does not establish unlawfulness for every query or behavioral response by every observed person. It establishes that protected-activity queries are not merely hypothetical and supports heightened authorization plus direct empirical study.

Distributional risks also matter. Records produced through unequal enforcement environments may carry those patterns into later analytical use [49]. Socio-technical fairness research likewise warns that abstraction from organizational and social context hides how formally uniform data practices produce unequal burdens [50]. These sources address related policing and analytical systems rather than supplying a direct ALPR effect estimate; the manuscript uses them only for the shared data-reuse mechanism.

3.4. Source-Quality-Weighted Evidence Episodes

The strongest administrative evidence arises from official audits, statutory studies, and court records. Los Angeles's inspector general evaluated approximately 2000 cameras and more than 210.5 million reads during a two-month review window. The report documented recoveries and arrests alongside erroneous stops associated with stale source information and significant contract, security, privacy, access-control, and oversight limitations [39]. Virginia's 2026 statutory review found retention practices exceeding the state's 21-day rule among a subset of responding agencies [51]. Court decisions in Schmidt, Robinson, and Bartholomew demonstrate that system configuration and statutory procedure alter legal outcomes [27,30,31].

Investigative reporting and civil-society records identify activated cross-jurisdictional pathways. Illinois reviewed abortion- and immigration-related searches [52]. Flock paused federal pilots following concern about access scope and safeguards [53]. Reporting on federal travel-pattern monitoring demonstrates pressure toward mobility inference rather than known-plate matching [54]. California's attorney general sued El Cajon over out-of-state sharing [55]. These materials have different evidentiary status from an audit or judicial finding, so Table 4 states what each source supports and what it does not.

Table 4. Evidence episodes, source quality, supported inference, and prohibited overclaim.

Episode	Source Quality	Supported Inference	Overclaim Excluded
Los Angeles audit, 2026	Tier A government audit	Benefits and control failures coexisted at large scale; termination remained available.	Does not isolate vendor-specific causal crime effects.
Virginia review, 2026	Tier A legislative study	Statutory retention text did not assure uniform operational compliance.	Does not prove intentional violation by every listed agency.
Schmidt/Robinson/Bartholomew	Tier A judicial records	Legal outcomes depend on coverage, retention, query facts, and statutory duties.	Does not establish one nationwide constitutional rule for every ALPR configuration.
Illinois/federal pilots/El Cajon	Tier A/D mixed official and investigative records	Cross-jurisdictional purpose drift and oversight lag are activated pathways.	Does not establish prevalence across all queries or platforms.
Protest-search log analysis	Tier D civil-society analysis with records	Protected-activity search is an observed use path.	Does not establish that every query was unlawful or caused participation decline.
ALPR effectiveness studies	Tier B peer-reviewed empirical research	Selected patrol/investigative benefits are plausible; evidence is mixed and context-dependent.	Does not justify population-scale collection or every platform feature.

3.5. STRIDE-Informed Threat Model and Availability

The security target includes event records, images, plate translations, vehicle attributes, hotlists, user identities, permissions, query logs, exports, case links, retention settings, encryption keys, backups, and deletion evidence. Trust boundaries lie between edge device and network, network and cloud, vendor tenant and customer tenant, local agency and partner agency, platform and external hotlist, user interface and exported file, and operational system and independent auditor.

Threat actors include external attackers, malicious or careless insiders, compromised partner accounts, unauthorized private users, vendor administrators exceeding role, and legitimate users pursuing impermissible purposes. STRIDE supplies a security taxonomy [38]; privacy and constitutional analysis adds purpose limitation, association exposure, and lawful authority. Table 5 links threats to controls.

Table 5. STRIDE-informed ALPR threat model, including availability.

STRIDE Class	Representative Event	System Consequence	Required Controls
Spoofing	Plate clone; forged user identity; compromised partner credential	False match, wrongful stop, unauthorized search	MFA, device identity, plate/vehicle/context verification, credential recertification
Tampering	Altered hotlist, query log, retention setting, export, or model output	Evidence corruption; concealed misuse; extended retention	Signed updates, immutable logs, separation of duties, configuration attestation
Repudiation	User denies query or administrator denies policy change	No attribution or sanction	Named accounts, tamper-evident audit, reason/case fields, independent custody
Information disclosure	Breach, overbroad sharing, sensitive-location query, exported dataset	Movement and association exposure	Encryption, least privilege, query bounding, masking, export controls, notice/remedy
Denial of service	Camera/network/cloud/hotlist/audit outage	Missed alert, stale reliance, unsafe workaround, incomplete accountability	Continuity plan, health monitoring, cached-data expiry, fail-safe procedures, recovery audit
Elevation of privilege	Stale account, vendor admin misuse, permissive integration	Broader query/share power than authorized	Role limits, just-in-time privilege, periodic recertification, dual approval, rapid revocation

Availability requires explicit treatment. If cameras, networks, cloud search, hotlists, or audit services fail, officers may miss time-sensitive alerts, rely on stale local information, bypass normal verification, or lose evidence needed for accountability. A rights-preserving design treats ALPR as a lead system rather than a sole basis for coercive action. It maintains tested continuity procedures, records offline operations, expires cached hotlists, queues deletion safely, restores audit completeness, and requires incident review after recovery [45]. Safe degradation is intended to reduce both operational harm and emergency-based expansion of unchecked discretion.

3.6. Seven Falsifiable Propositions


Table 6 operationalizes the propositions. Each states a relationship for later testing, identifies evidence that would count against it, and avoids presenting conceptual derivation as empirical verification.

Table 6. Falsifiable propositions, indicators, methods, and disconfirming conditions.

Proposition	Relationship	Unit	Observable Indicators	Suitable Method	Disconfirming Condition
P1 Infrastructure complementarity	Coverage $\times$ retention $\times$ query richness $\times$ interoperability	Agency-platform-month	Camera/location maps; retention; query functions; partner reach; sensitive exposure	Multilevel audit-log or longitudinal comparative study	Interaction terms are null/negative; camera count alone explains capability and exposure.
P2 Persistence and secondary use	Search cost/retention $\rightarrow$ secondary-purpose query share	Query or agency-month	Retention window; historical-query volume; purpose codes; case linkage	Audit-log time series; policy discontinuity	Longer retention/lower search cost does not increase non-primary or weakly linked queries.
P3 Authority diffusion	Origin–use–oversight distance $\rightarrow$ policy nonconformity	Interagency request	Jurisdictions crossed; rule differences; approval; denial; violation	Network analysis; matched request study	Greater distance predicts equal or stronger compliance after controls.
P4 Feedback asymmetry	Output-only dashboards $\rightarrow$ expansion/weak correction	Jurisdiction-year	Benefits reported; harms/exposure reported; camera growth; pauses; sanctions	Longitudinal policy comparison	Balanced indicators do not alter expansion, correction, or renewal decisions.
P5 Potential democratic chill	Perceived durable ALPR exposure $\rightarrow$ reduced protected participation, moderated by vulnerability	Person/event/place	Coverage awareness; attendance; route choice; speech willingness; trust	Survey experiment; panel; natural experiment; qualitative study	Awareness has no negative effect, or increases participation consistently.
P6 Institutional oversight	Decision rights/capacity $\rightarrow$ fewer uncorrected violations than nominal human review	Agency-year	Audit independence; refusal protection; remedies; sanctions; time to correction	Comparative case study; audit study	Nominal review performs as well as empowered multi-level oversight.
P7 Bounded hybrid intelligence	Technical bounds + human authority + independent control $\rightarrow$ benefits with lower rights burden	Deployment-year	Validated assists; uninvolved reads; false stops; violations; trust; termination feasibility	Matched multi-site or stepped-wedge evaluation	Bounded architecture reduces serious-case performance without reducing rights burden.

3.7. Rights-Preserving Hybrid-Intelligence Architecture

Figure 3 places automated detection inside five nested control layers. Democratic authorization defines necessity, alternatives, purpose, placement, and sunset. Technical boundaries enforce minimization, short default retention, sensitive-location rules, and narrow sharing. Protected human judgment verifies source, context, proportionality, and lawful basis before coercive action. Independent institutions audit, hear complaints, correct records, impose sanctions, and report publicly. Continuity and termination controls preserve safe operation during failure and credible exit when governance proves inadequate.



Automated detection remains subordinate to institutional authority over purpose, coercion, review, and shutdown.

Figure 3. Rights-preserving hybrid-intelligence architecture. The layers represent linked control authorities rather than a sequential software pipeline.

Meaningful human oversight is therefore distributed. Residents and legislatures define acceptable purposes. Engineers encode retention and permission boundaries. Officers verify alerts and retain protected refusal authority. Supervisors investigate anomalies. Judges authorize sensitive historical searches. Auditors inspect logs and configuration. Councils or other lawful authorities pause integrations, decline renewal, and require deletion. The architecture does not romanticize discretion; it assigns bounded decision rights and evidence duties at each layer.

Bounded rationality explains why formal discretion without information, time, and institutional protection remains weak [56]. Value-sensitive design supplies a method for translating privacy, security, association, due process, and equal treatment into technical and organizational requirements [57]. The NIST Privacy Framework and AI Risk Management Framework add life-cycle governance, mapping, measurement, and risk-management functions [43,44]. RPHI adds constitutional authorization, public contestability, correction, and termination authority.

The design narrows operational variety and strengthens governance response. Sharing defaults, deletion automation, purpose blocks, short retention, and validated hotlists

reduce reachable states. Independent audit, judicial review, notice, appeal, sanctions, tested recovery, and exit enlarge the response set. Neither strategy is sufficient alone. Unlimited architecture overwhelms oversight; strong technical restrictions without external authority remain vulnerable to quiet reconfiguration. Table 7 specifies the automated role, protected human authority, independent or public control, and evidence generated at each life-cycle stage.

Table 7. Rights-preserving controls across the ALPR life cycle.

Stage	Automated Role	Protected Human Authority	Independent/Public Control	Evidence Generated
Authorization	Propose need, coverage, alternatives	Define purpose, risk tolerance, sunset	Impact assessment; public hearing; legislative approval	Decision record; map; alternatives analysis
Capture/classification	Read plate/attributes; calculate confidence	Approve placement; verify uncertainty	Sensitive-zone rules; accuracy and bias testing	Placement map; confidence; rejected reads
Storage/deletion	Encrypt, index, enforce retention	Authorize case-specific preservation	Key governance; deletion and backup audit	Retention state; order; deletion receipt
Query/sharing	Authenticate, match, log, apply rules	State factual basis; approve elevated query	Warrant/independent review; partner controls	User; case; purpose; scope; recipient
Enforcement	Display alert context and source status	Confirm plate, vehicle, hotlist, lawful basis	Supervisory review; exclusion of unverified alert	Verification; action; outcome; correction
Availability/recovery	Monitor health; queue protected operations	Apply safe fallback; authorize restoration	Continuity test; independent incident review	Outage; fallback; restored logs; lessons
Review/renewal	Aggregate benefits, burdens, anomalies	Recommend redesign, pause, or end	Audit; complaint; public report; termination	Outcomes; exposure; violations; trust; deletion proof

4. Discussion

4.1. From Product Critique to Infrastructure Governance

The revised analysis moves beyond a Flock-centered title without diluting the empirical case. Flock's record reveals how inexpensive distributed cameras, cloud search, and cross-agency sharing produce infrastructure. Axon/Fusus demonstrates convergence between ALPR, livestreaming, vehicle attributes, and real-time operational views [4,5]. Motorola Vigilant demonstrates the significance of commercial-scale detections, association analysis, and predictive vehicle-location features [6]. Federal access to commercial databases further weakens the assumption that local procurement defines the whole system [46].

A vendor ban alone would therefore invite substitution. Governance should follow data and capability across contracts, platforms, exports, partner agencies, and successors. Minimum rules should attach to collection, retention, search, sharing, inference, enforcement, audit, and deletion regardless of brand or ownership model.

4.2. Interpreting the RCD Without False Precision

For applied use, the RCD yields a four-domain profile—capture; persistence and inference; access, sharing, and integration; and action and error propagation—showing where operational demand exceeds effective response. It is not an index and does not rank jurisdictions on a universal scale. The profile preserves non-compensability: excellent

cybersecurity cannot offset absent lawful purpose limitation, and public reporting cannot offset the absence of correction or sanction.

Existing state frameworks illustrate elements of higher governance response. Virginia's statutory review reveals the distance between formal retention rules and implementation [51]. Minnesota law addresses purpose, role-based access, audit trails, public logs, sharing, destruction, warrants for tracking, and suspension after substantial noncompliance [42]. National legislative tracking documents substantial interstate variation [58]. These sources inform control anchors but do not validate a universal numerical index.

The three documented configurations show descriptive discrimination. Los Angeles combined broad operational capability with weak fine-grained control but retained exit authority [39,40]. Norfolk presents a narrower and legally contested configuration in which incomplete coverage influenced judicial analysis [27,30]. Austin demonstrates termination after public controversy [41]. The bounded benchmark clarifies design ambition without pretending that all safeguards presently coexist in one site.

The sensitivity analysis is intentionally modest. Reclassifying every uncertain judgment toward lower operational demand or stronger governance does not erase the central Los Angeles access-and-integration mismatch. Norfolk remains uncertain, which is itself a result: public records do not support confident claims about every integration or control. Future multi-coder field research should test reliability and predictive validity across a larger jurisdiction sample.

4.3. Free Expression and Public Confidence as Performance Criteria

Police technology evaluation traditionally privileges crime outputs, cost, and technical accuracy. Recent ALPR scholarship argues for evaluation against both public safety and public confidence [10]. A democratic systems account adds protected participation. If credible evidence shows that a deployment assisting selected cases also deters lawful association or contact with sensitive institutions, that outcome is not merely an externality; it is impairment of a core public value.

Measurement should include total uninvolved exposure, query-to-action ratio, false and stale hotlist events, sensitive-purpose searches, external requests, denied requests, user suspensions, complaints, corrections, geographic placement, public trust, and independently measured participation effects. Direct measurement remains difficult because avoidance leaves little administrative trace. Panel surveys, survey experiments, event studies, and qualitative interviews should be designed with privacy protections so research does not reproduce the surveillance under study.

4.4. Security, Availability, and Civil Liberties Are Coupled

Confidentiality and integrity failures receive deserved attention, but availability also carries rights implications. An outage may encourage officers to act on stale caches, bypass verification, or shift to undocumented personal channels. Loss of audit availability may conceal otherwise lawful activity and prevent review. Recovery that restores operational data but not logs produces an accountability gap.

Continuity design should specify maximum tolerable outage, cached-data expiration, manual verification, incident command, restoration priority, audit reconciliation, queued deletion, and post-incident reporting [45]. No person should face coercive action solely because an unavailable system prevents confirmation. Conversely, service interruption should not become a pretext for indefinite emergency exceptions.

4.5. Procurement, Contracts, and Organizational Drift

Procurement is a constitutional and systems-design control point because options narrow after deployment. Contracts should, at minimum, define ownership and custody of

records, derived information, logs, and configuration; prohibit sale, advertising, and unrelated model training; enforce short retention and case-specific preservation; disable sharing by default; require named users and periodic recertification; provide immutable logs to independent auditors; govern keys, subcontractors, exports, backups, and breaches; require deletion proof; preserve defense and public-record obligations; and establish financial, disciplinary, pause, continuity, and termination remedies.

Implementation can be staged. Before procurement or renewal, a public authority should approve purpose, placement, sensitive-location exclusions, retention, sharing defaults, audit and export rights, continuity, exit, and deletion certification. During initial deployment, administrators should configure named users, multifactor authentication, roles and purpose fields, hotlist-freshness rules, independent alert verification, and outage-safe procedures. Quarterly, an independent reviewer should sample queries and configuration against the approved state. Annually, the authorizing body should publish benefit-and-burden indicators, hear complaints, and decide whether to renew, narrow, pause, or terminate the system.

Responsibility should follow authority. Vendors enforce configured technical boundaries and provide exportable logs and deletion evidence; agencies manage access, verification, and discipline; independent reviewers test use and configuration; courts or other designated authorities approve sensitive historical searches; and elected or statutorily designated bodies hold renewal and termination power. Smaller jurisdictions can use state model clauses, pooled independent audit, and shared technical expertise, but limited oversight capacity should narrow deployment scope rather than weaken minimum controls.

Organizational drift often appears mundane: a retention exception becomes routine, a purpose field becomes generic, an account remains active after transfer, or a product update changes an integration. Normalization of deviance explains why repeated departures become ordinary when immediate catastrophe does not follow [59]. Quarterly configuration review and annual democratic reauthorization should compare actual system state with the publicly approved state.

Austin and Los Angeles illustrate exit as a real control action [40,41]. Cancellation does not prove that every jurisdiction should abandon every ALPR use. It proves that consultation is meaningful only when an institution holds authority to change or end the system and require verified deletion.

4.6. Proportionate Public-Safety Use

Benefits and burdens should be compared at the same level. A recovered vehicle should be evaluated alongside offense seriousness, alternative methods, total reads, retention, external access, false alerts, enforcement outcomes, and less intrusive configurations. For a serious and time-sensitive case, real-time matching against a current validated hotlist may have strong justification. The burden falls when non-hit records expire rapidly, historical search requires legal process, sharing is narrow, and officers independently verify every alert.

Justification weakens as retention lengthens, descriptive or association search expands, commercial datasets enter the query, low-level offenses trigger access, or placement exposes protected activity. Evidence on deterrence and investigative effects remains mixed [7–10]. A narrower deployment may preserve most validated outcomes if benefits concentrate at selected locations or offense types. Independent evaluations should use matched sites, phased implementation, exposure denominators, and counterfactual designs.

4.7. Research Agenda

The seven propositions define the next empirical program. Multi-site audit-log research should measure coverage, retention, query richness, partner reach, purpose, case linkage, denial, and action. Network analysis should test whether jurisdictional distance predicts policy nonconformity. Longitudinal comparison should test whether balanced dashboards alter growth, discipline, pause, or renewal. Community studies should measure awareness, trust, participation, route choice, and avoidance without collecting new movement histories.

RCD reliability requires trained independent coders, expert review of anchors, transparent disagreement resolution, and preregistered sensitivity rules. Predictive validity would require evidence that unmatched domains precede documented violations, erroneous stops, public-trust loss, or termination. The present article supplies the construct and coding protocol; it does not report validation that has not occurred.

Legal research should track how courts define the relevant mosaic after *Chatrle* and how private or commercial repositories affect state-action analysis. Procurement research should compare contract language with deployed configuration and post-termination deletion. Security research should test not only breach resistance but safe degradation, audit continuity, and restoration after cloud or network failure.

5. Conclusions

The title's "systemic erosion" denotes a conditional systems-level risk, not a measured universal behavioral effect. That risk arises when governance treats cameras as isolated tools even though the relevant infrastructure extends through classification, networks, cloud archives, commercial repositories, historical and attribute search, real-time crime centers, permissions, police action, audit, continuity, and democratic review. Its power emerges from complementarity among coverage, persistence, inference, and interoperability.

The evidence does not support the categorical claim that ALPR lacks public-safety value. It also does not support an assumption that selected recoveries justify population-scale collection or broad retrospective access. Audit findings, litigation, statutory compliance reviews, cross-state controversies, protest-related searches, and termination decisions establish activated governance pathways while leaving prevalence and causal magnitude open.

The RCD identifies unmatched control domains without reducing heterogeneous rights failures to a score. Rights-preserving hybrid intelligence offers a constructive architecture: narrow automation; protected human authority over context and coercion; judicial review for sensitive search; independent audit; short retention; restrictions that travel with data; contestability, correction, sanctions, continuity, and credible exit.

The governing principle is straightforward: no surveillance capability should exceed public institutions' demonstrated capacity to authorize, bound, inspect, contest, correct, recover, and end it. Applying that principle across Flock, Axon/Fusus, Motorola Vigilant, commercial repositories, and future ALPR platforms would preserve carefully justified uses while resisting the conversion of local observations into an unaccountable infrastructure of movement and association.

Supplementary Materials: The following supporting information can be downloaded at: <https://www.mdpi.com/article/10.3390/systems14081019/s1>, Supplementary Material S1, Documentary review protocol, search log, source-quality rules, source inventory and evidentiary-role matrix, central claim-to-source crosswalk, domain-level RCD sensitivity display, and case-selection and comparability record.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable. The study used public documents and did not involve human participants.

Informed Consent Statement: Not applicable.

Data Availability Statement: The documentary audit trail is provided in Supplementary Material S1. No personal ALPR records or new personally identifiable dataset were created.

Acknowledgments: The author acknowledges the public officials, researchers, journalists, and civil-society organizations whose documented work supports independent evaluation of surveillance systems.

Conflicts of Interest: The author declares no conflict of interest.

Abbreviations

ALPR	automated license plate readers
AWS	Amazon Web Services
D_{RCD}	set of Rights Control Deficit domains
MFA	multifactor authentication
RCD	Rights Control Deficit
RPHI	rights-preserving hybrid intelligence
STRIDE	spoofing, tampering, repudiation, information disclosure, denial of service, and elevation of privilege

References

1. U.S. Department of Homeland Security, Science and Technology Directorate. *Automated License Plate Readers Market Survey Report*; National Urban Security Technology Laboratory: Washington, DC, USA, 2025. Available online: https://www.dhs.gov/sites/default/files/2025-06/25_0606_st_lprmsr.pdf (accessed on 3 August 2026).
2. Flock Safety. License Plate Readers. Available online: <https://www.flocksafety.com/products/license-plate-readers> (accessed on 3 August 2026).
3. Flock Safety. Flock Safety License Plate Reader Usage and Privacy Policy. Available online: <https://www.flocksafety.com/legal/lpr-policy> (accessed on 3 August 2026).
4. Axon Enterprise, Inc. Axon Outpost: Fixed Infrastructure Camera with ALPR and Vehicle Attribute Recognition. Available online: <https://www.axon.com/products/axon-outpost> (accessed on 3 August 2026).
5. Axon Enterprise, Inc. Fusus ALPR Experience. Available online: <https://www.axon.com/help/fusus/software/fusus/alpr/alpr-experience.htm> (accessed on 3 August 2026).
6. Motorola Solutions. Vigilant VehicleManager License Plate Recognition Software. Available online: https://www.motorolasolutions.com/en_us/video-security-access-control/license-plate-recognition-camera-systems/vehiclemanager-lpr-analytics-software.html (accessed on 3 August 2026).
7. Lum, C.; Hibdon, J.; Cave, B.; Koper, C.S.; Merola, L. License plate reader police patrols in crime hot spots: An experimental evaluation in two adjacent jurisdictions. *J. Exp. Criminol.* **2011**, *7*, 321–345. [CrossRef]
8. Koper, C.S.; Taylor, B.G.; Woods, D.J. A randomized test of initial and residual deterrence from directed patrols and use of license plate readers at crime hot spots. *J. Exp. Criminol.* **2013**, *9*, 213–244. [CrossRef]
9. Koper, C.S.; Lum, C. The impacts of large-scale license plate reader deployment on criminal investigations. *Police Q.* **2019**, *22*, 305–329. [CrossRef]
10. Lum, C.; Koper, C.S.; Lee, H.; Nagin, D.S.; Sherman, L.W. Measuring the cost-effectiveness of new technologies in policing: The case of automatic license plate readers. *Camb. J. Evid.-Based Polic.* **2025**, *9*, 4. [CrossRef]
11. Willis, J.J.; Koper, C.S.; Lum, C. The adaptation of license-plate readers for investigative purposes: Police technology and innovation re-invention. *Justice Q.* **2018**, *35*, 614–638. [CrossRef]
12. Lum, C.; Koper, C.S.; Willis, J.J.; Happeny, S.; Vovak, H.; Nichols, J. The rapid diffusion of license plate readers in U.S. law enforcement agencies. *Polic. Int. J.* **2019**, *42*, 376–393. [CrossRef]
13. Baxter, G.; Sommerville, I. Socio-technical systems: From design methods to systems engineering. *Interact. Comput.* **2011**, *23*, 4–17. [CrossRef]
14. Alibašić, H. *Hybrid Intelligence for Effective Digital Governance: AI in Administration*; Springer Nature: Cham, Switzerland, 2026. Available online: <https://link.springer.com/book/9783032280855> (accessed on 3 August 2026).

15. Nissenbaum, H. *Privacy in Context: Technology, Policy, and the Integrity of Social Life*; Stanford University Press: Stanford, CA, USA, 2010.
16. Lyon, D. (Ed.) *Surveillance as Social Sorting: Privacy, Risk, and Digital Discrimination*; Routledge: London, UK, 2003.
17. Brayne, S. *Predict and Surveil: Data, Discretion, and the Future of Policing*; Oxford University Press: New York, NY, USA, 2021.
18. Ferguson, A.G. *The Rise of Big Data Policing: Surveillance, Race, and the Future of Law Enforcement*; New York University Press: New York, NY, USA, 2017.
19. Stoycheff, E. Under surveillance: Examining Facebook's spiral of silence effects in the wake of NSA internet monitoring. *J. Mass Commun. Q.* **2016**, *93*, 296–311. [\[CrossRef\]](#)
20. Penney, J.W. Chilling effects: Online surveillance and Wikipedia use. *Berkeley Technol. Law J.* **2016**, *31*, 117–182. Available online: <https://scholarship.law.berkeley.edu/btlj/vol31/iss1/5/> (accessed on 3 August 2026).
21. Büchi, M.; Festic, N.; Latzer, M. The chilling effects of digital dataveillance: A theoretical model and an empirical research agenda. *Big Data Soc.* **2022**, *9*. [\[CrossRef\]](#)
22. Murray, D.; Fussey, P.; Hove, K.; Wakabi, W.; Kimumwe, P.; Saki, O.; Stevens, A. The chilling effects of surveillance and human rights: Insights from qualitative research in Uganda and Zimbabwe. *J. Hum. Rights Pract.* **2024**, *16*, 397–412. [\[CrossRef\]](#)
23. Merola, L.M.; Lum, C.; Murphy, R.P. The impact of license plate recognition technology on trust in law enforcement: A survey-experiment. *J. Exp. Criminol.* **2019**, *15*, 55–66. [\[CrossRef\]](#)
24. United States v. Jones, 565 U.S. 400 (2012). Available online: <https://supreme.justia.com/cases/federal/us/565/400/> (accessed on 3 August 2026).
25. Carpenter v. United States, 585 U.S. 296 (2018). Available online: <https://supreme.justia.com/cases/federal/us/585/16-402/> (accessed on 3 August 2026).
26. Chatrue v. United States, 609 U.S. ____ (2026), No. 25-112, Slip Opinion (29 June 2026). Available online: https://www.supremecourt.gov/opinions/25pdf/25-112_0am4.pdf (accessed on 3 August 2026).
27. Schmidt v. City of Norfolk, No. 2:24-cv-00621-MSD-RJK, Document 191 (E.D. Va. 27 January 2026). Available online: <https://law.justia.com/cases/federal/district-courts/virginia/vaedce/2%3A2024cv00621/562101/191/> (accessed on 3 August 2026).
28. American Civil Liberties Union; American Civil Liberties Union of Virginia; Electronic Frontier Foundation. Brief of Amici Curiae in Support of Plaintiffs-Appellants, Schmidt v. City of Norfolk, No. 26-1227 (4th Cir. 20 April 2026). Available online: <https://assets.aclu.org/live/uploads/2026/04/Schmidt-v.-Norfolk-file-stamped-amicus.pdf> (accessed on 3 August 2026).
29. Leaders of a Beautiful Struggle v. Baltimore Police Department, 2 F.4th 330 (4th Cir. 2021) (en banc). Available online: <https://law.justia.com/cases/federal/appellate-courts/ca4/20-1495/20-1495-2021-06-24.html> (accessed on 3 August 2026).
30. Robinson v. Commonwealth, Record No. 1912-24-1 (Va. Ct. App. 2026). Available online: <https://law.justia.com/cases/virginia/court-of-appeals-published/2026/1912-24-1.html> (accessed on 3 August 2026).
31. Bartholomew v. Parking Concepts, Inc., 118 Cal. App. 5th 438 (Cal. Ct. App. 2026). Available online: <https://law.justia.com/cases/california/court-of-appeal/2026/a171546m.html> (accessed on 8 August 2026).
32. Checkland, P. *Systems Thinking, Systems Practice*; John Wiley & Sons: Chichester, UK, 1981.
33. Meadows, D.H. *Thinking in Systems: A Primer*; Chelsea Green Publishing: White River Junction, VT, USA, 2008.
34. Ashby, W.R. *An Introduction to Cybernetics*; Chapman & Hall: London, UK, 1956.
35. Leveson, N.G. *Engineering a Safer World: Systems Thinking Applied to Safety*; MIT Press: Cambridge, MA, USA, 2011.
36. Ulrich, W. *Critical Heuristics of Social Planning: A New Approach to Practical Philosophy*; Haupt: Bern, Switzerland, 1983.
37. Bowen, G.A. Document analysis as a qualitative research method. *Qual. Res. J.* **2009**, *9*, 27–40. [\[CrossRef\]](#)
38. Shostack, A. *Threat Modeling: Designing for Security*; John Wiley & Sons: Indianapolis, IN, USA, 2014.
39. Office of the Inspector General, Los Angeles Board of Police Commissioners. Review of Los Angeles Police Department Automatic License Plate Reader System; BPC #26-184; Los Angeles, CA, USA, 10 July 2026. Available online: https://www.lapdpolicecom.lacity.org/071426/BPC_26-184x.pdf (accessed on 3 August 2026).
40. NBC Los Angeles. LAPD Watchdog Urges Reconsideration of Automatic License Plate Readers. 13 July 2026. Available online: <https://www.nbclosangeles.com/investigations/lapd-watchdog-urges-reconsideration-of-automatic-license-plate-readers/3916867/> (accessed on 3 August 2026).
41. Mir, R. Victory! Austin Organizers Cancel City's Flock ALPR Contract. Electronic Frontier Foundation, 6 June 2025. Available online: <https://www.eff.org/deeplinks/2025/06/victory-austin-organizers-cancel-citys-flock-alpr-contract> (accessed on 3 August 2026).
42. Minnesota Legislature. Minnesota Statutes § 13.824: Automated License Plate Readers. Available online: <https://www.revisor.mn.gov/statutes/cite/13.824> (accessed on 3 August 2026).
43. National Institute of Standards and Technology. *NIST Privacy Framework: A Tool for Improving Privacy through Enterprise Risk Management*; Version 1.0; NIST: Gaithersburg, MD, USA, 2020. [\[CrossRef\]](#)
44. Tabassi, E. *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*; NIST AI 100-1; National Institute of Standards and Technology: Gaithersburg, MD, USA, 2023. [\[CrossRef\]](#)

45. Swanson, M.; Bowen, P.; Phillips, A.W.; Gallup, D.; Lynes, D. *Contingency Planning Guide for Federal Information Systems*; NIST SP 800-34 Rev. 1; National Institute of Standards and Technology: Gaithersburg, MD, USA, 2010. [CrossRef]
46. U.S. Department of Homeland Security, Immigration and Customs Enforcement. DHS/ICE/PIA-039(b): Acquisition and Use of License Plate Reader Data from a Commercial Service; Washington, DC, USA, 21 May 2021. Available online: https://www.dhs.gov/sites/default/files/publications/privacy-pia30b-ice-acquisitionanduseoflprdatafromacommercialservice-june2021_0.pdf (accessed on 3 August 2026).
47. NAACP v. Alabama ex rel. Patterson, 357 U.S. 449 (1958). Available online: <https://supreme.justia.com/cases/federal/us/357/449/> (accessed on 3 August 2026).
48. Maass, D.; Alajaji, R. How Cops Are Using Flock Safety's ALPR Network to Surveil Protesters and Activists. Electronic Frontier Foundation, 20 November 2025. Available online: <https://www.eff.org/deeplinks/2025/11/how-cops-are-using-flock-safety-s-alpr-network-surveil-protesters-and-activists> (accessed on 3 August 2026).
49. Richardson, R.; Schultz, J.M.; Crawford, K. Dirty data, bad predictions: How civil rights violations impact police data, predictive policing systems, and justice. *N.Y.U. Law Rev. Online* **2019**, *94*, 192–233. Available online: <https://www.nyulawreview.org/online-features/dirty-data-bad-predictions-how-civil-rights-violations-impact-police-data-predictive-policing-systems-and-justice/> (accessed on 3 August 2026).
50. Selbst, A.D.; Boyd, D.; Friedler, S.A.; Venkatasubramanian, S.; Vertesi, J. Fairness and abstraction in sociotechnical systems. In Proceedings of the Conference on Fairness, Accountability, and Transparency, Atlanta, GA, USA, 29–31 January 2019; pp. 59–68. [CrossRef]
51. Virginia State Crime Commission. 2025 Annual Report—Law Enforcement Use of ALPR Update; Richmond, VA, USA, January 2026. Available online: <https://vsc.virginia.gov/Annual%20Reports/2025%20VSCC%20Annual%20Report%20-%20Law%20Enforcement%20Use%20of%20ALPR%20Update.pdf> (accessed on 8 August 2026).
52. Associated Press. Illinois Officials Investigate License-Plate Data Shared with Police Seeking Woman Who Had Abortion. 12 June 2025. Available online: <https://apnews.com/article/abortion-access-immigration-license-plate-readers-surveillance-13fac7c045df3c5e5145f6d4e4c4db28> (accessed on 3 August 2026).
53. Associated Press. License Plate Camera Company Halts Cooperation with Federal Agencies Amid Investigation Concerns. 25 August 2025. Available online: <https://apnews.com/article/immigration-abortion-license-plates-cameras-cc5f29df94a29ee2c6c2feb2151c8f5e> (accessed on 3 August 2026).
54. Tau, B.; Burke, G. Border Patrol is Monitoring U.S. Drivers and Detaining Those with 'Suspicious' Travel Patterns. Associated Press, 20 November 2025. Available online: <https://apnews.com/article/immigration-border-patrol-surveillance-drivers-ice-trump-9f5d05469ce8c629d6fecf32d32098cd> (accessed on 3 August 2026).
55. California Department of Justice, Office of the Attorney General. Attorney General Bonta Sues El Cajon for Illegally Sharing License Plate Data with Out-of-State Law Enforcement. 3 October 2025. Available online: <https://www.oag.ca.gov/news/press-releases/attorney-general-bonta-sues-el-cajon-illegally-sharing-license-plate-data-out> (accessed on 3 August 2026).
56. Simon, H.A. *Administrative Behavior: A Study of Decision-Making Processes in Administrative Organizations*, 4th ed.; Free Press: New York, NY, USA, 1997.
57. Friedman, B.; Hendry, D.G. *Value Sensitive Design: Shaping Technology with Moral Imagination*; MIT Press: Cambridge, MA, USA, 2019.
58. National Conference of State Legislatures. Automated License Plate Readers: State Statutes. Updated 3 February 2022. Available online: <https://www.ncsl.org/technology-and-communication/automated-license-plate-readers-state-statutes> (accessed on 3 August 2026).
59. Vaughan, D. *The Challenger Launch Decision: Risky Technology, Culture, and Deviance at NASA*; University of Chicago Press: Chicago, IL, USA, 1996.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.