ASPEN POLICY ACADEMY
aspen institute

Operational Guidance

Proposed New York City Computer Vision Deployment Framework

By Julia Lane and Blair Palmer

November 2025

This guidance was developed by Aspen Policy Academy fellows while participating in the Science and Technology Policy Fellowship. The Fellows were asked to recommend updates to the New York City Office of Technology and Innovation's Artificial Intelligence Action Plan. This document outlines a framework for implementing computer vision technologies across city agencies while protecting privacy and vulnerable populations. The full project, including a policy brief explaining the fellows' core recommendations, is available here. Please note that all authors' opinions published here are their own. This publication does not reflect the views of the Aspen Policy Academy or the Aspen Institute.

Executive Summary

The City of New York's Office of Technology and Innovation (OTI) and other city agencies are increasingly exploring computer vision technologies for transportation, infrastructure monitoring, and public safety applications. However, agencies lack structured guidance for responsible deployment that balances innovation with privacy protection and community trust.

We recommend that OTI incorporate a 3-tier deployment framework into New York City's Artificial Intelligence (AI) Action Plan to provide agencies with clear guidance on implementing computer vision technology. This framework categorizes applications by risk level and data sensitivity, enabling the deployment of beneficial low-risk applications while establishing appropriate safeguards for community-impacting deployments.

The framework builds on New York City's existing privacy policies and AI governance structure, providing practical guidance that agencies can implement within current authority and resource constraints.

Research Methodology

This framework was developed through an analysis of computer vision policies in multiple cities, consultation with academic researchers specializing in computer vision technologies and algorithmic governance, and a review of New York City's existing AI Action Plan, Citywide Privacy Protection Policies and Protocols, and technology oversight frameworks. The recommendations are designed to complement the city's current privacy and cybersecurity policies while providing specific guidance for computer vision applications.

Framework Philosophy

Our approach builds on New York City's AI Action Plan and existing city policies related to privacy and cybersecurity, recognizing that computer vision applications exist on a spectrum of risk, feasibility, and community impact. By categorizing deployments into 3 distinct tiers based on data sensitivity and privacy implications, we can:

- **Enable Innovation:** Streamline beneficial applications with minimal privacy concerns;
- **Ensure Accountability:** Apply appropriate oversight for community-impacting technologies;
- **Build Trust:** Require comprehensive engagement for applications affecting civil liberties; and
- **Show Leadership:** Position New York City as a national model for responsible computer vision governance.

This framework is designed to complement the city's 4-level information confidentiality classification system (public, sensitive, private, confidential) outlined in the Agency Privacy Officer Toolkit. The 3-tier computer vision framework provides technology-specific guidance that aligns with existing data classifications, ensuring consistency with current privacy laws and cybersecurity requirements while offering practical deployment guidance for agencies.

3-Tier Deployment Framework

Tier	Characteristics	Governance Requirements	Example Applications
TIER 1: FOUNDATIONAL APPLICATIONS *Streamlined deployment for infrastructure and operational efficiency*	• No identifying information collected at any point[1] • Privacy-by-design systems with edge processing only[2] • Clear public benefit with no privacy impact • Limited access to sensitive information	• Standard privacy assessment • Internal OTI policy compliance review • Public notification via agency websites	• **Traffic flow optimization:** Anonymous vehicle counting for traffic reports and signal timing • **Infrastructure monitoring:** Automated assessment of bridge integrity, road conditions, or building exteriors using privacy-by-design sensors • **Public space analytics:** Aggregated occupancy patterns for park and facility resource allocation • **Waste management:** Monitoring for excessive trash buildup or illegal dumping using privacy-by-design cameras with automatic anonymization

3-Tier Deployment Framework (continued)

Tier	Characteristics	Governance Requirements	Example Applications
TIER 2: OPERATIONAL APPLICATIONS *Enhanced oversight for community-facing deployments*	• Information collection with engineered anonymization before storage or transmission[3] • Moderate privacy implications with established safeguards • Direct impact on public services or safety • Potential for minor algorithmic bias requiring monitoring[4]	• Enhanced review involving OTI privacy experts • Community notification through multiple channels • Annual performance and bias audits[5] • Regular stakeholder consultations • Clear data retention and deletion policies	• **Smart traffic management:** Real-time pedestrian and vehicle detection for signal optimization with automatic face blurring • **Parking enforcement:** License plate reading with strict data retention limits and purpose limitations • **Emergency response:** Crowd density monitoring for evacuation planning with edge processing • **Infrastructure alerts:** Automated detection of maintenance needs using existing CCTV with privacy safeguards[6]

3-Tier Deployment Framework (continued)

Tier	Characteristics	Governance Requirements	Example Applications
TIER 3: SENSITIVE APPLICATIONS *Comprehensive engagement for civil liberties considerations*	• Collection and retention of identifying information • Cloud processing acceptable with comprehensive safeguards • Significant privacy or civil liberties implications • Potential for substantial algorithmic bias or discrimination[7] • Use in law enforcement or sensitive public spaces • Impact on vulnerable populations	• Rigorous multidisciplinary oversight committee review • Community engagement and public hearings • Independent bias testing and algorithmic audits[5] • Ongoing monitoring with community representation • Annual public impact assessments • Strict data sharing limitations and retention policies	• **Facial recognition systems:** Any biometric identification in secure facilities • **Individual tracking systems:** Gait analysis or behavioral pattern monitoring that can identify specific individuals across locations • **Predictive analytics:** Systems that analyze behavior to predict law enforcement needs or influence resource deployment • **Surveillance integration:** Systems that combine computer vision with other identification technologies for comprehensive monitoring

Technical Implementation Definitions

[1] **No identifying information:** Systems are designed so that faces, license plates, and other personal identifiers are never captured by sensors or are automatically obscured at the hardware level before any processing occurs.

[2] **Edge processing only:** All data analysis occurs on the device itself; only aggregate statistics or metadata (e.g., "15 vehicles passed") leave the device. Video streams are never transmitted or stored off-device.

[3] **Engineered anonymization:** Systems may initially capture identifying information but automatically blur, redact, or aggregate these data before storage or transmission. This requires technical verification that the original identifying data cannot be reconstructed.

[4] **Minor versus substantial bias:** Minor bias affects service quality or efficiency but does not result in discriminatory enforcement or denial of services. Substantial bias results in measurable disparate impacts on protected groups or affects civil liberties.

[5] **Audit requirements:** Audits must be conducted by qualified independent evaluators with expertise in algorithmic assessment. Funding responsibility and audit frequency are to be determined through the OTI procurement guidelines development process.

[6] **Existing CCTV considerations:** When using existing camera infrastructure, agencies must address resolution limitations and camera positioning constraints and ensure that privacy safeguards are compatible with legacy systems.

[7] **Substantial algorithmic bias:** Bias that results in measurably different outcomes for protected groups, affects access to city services, or influences law enforcement actions. This requires immediate remediation and may result in system suspension.

Vulnerable Population Protections

Computer vision deployments must incorporate enhanced protections for communities that face heightened risks from surveillance technologies.

Core Principles

- **Data Sovereignty:** Prohibit sharing with immigration enforcement agencies in accordance with sanctuary city policies and implement technical safeguards to prevent inadvertent access. Balance city data transparency goals with confidentiality requirements through appropriate data governance frameworks.
- **Community Control:** Require enhanced OTI privacy review and community notification for Tier 3 deployments in areas with significant vulnerable populations. Oversight committees should include affected community representation rather than relying solely on community boards.
- **Bias Prevention:** Mandatory algorithmic testing with ideally qualified independent evaluators; clear remediation protocols when discrimination is detected.
- **Privacy by Design:** Automatic deletion of personal identifiers where technically feasible, with edge processing strongly preferred to minimize data transmission and reduce potential for misuse.
- **Enhanced Community Engagement** in areas with vulnerable populations, including pre-deployment impact assessments and public meetings with interpretation services. Translation and interpretation costs should be incorporated into project budgets, with accessibility accommodations including American Sign Language interpretation as required by city policy.
- **Multilingual Notification** and consultation processes in the top languages spoken in each deployment area, with clear timelines and multiple notification channels to ensure broad community awareness.
- **Independent Monitoring with Community Advocates** who have defined oversight roles, including access to aggregate performance data and authority to file formal complaints through established city grievance processes. Advocates cannot access individual case data but can review system performance metrics and bias testing results.
- **Precise Opt-out Mechanisms** where technically feasible (primarily for Tier 1 and some Tier 2 applications), plus individual data access and correction rights in accordance with existing privacy policies where personal data are collected.

- **Regular Equity Impact Assessments** with mandatory system modifications when measurable disparate outcomes are identified. When bias is detected, agencies must suspend affected system components pending remediation, following established civil rights compliance procedures and timelines for resolution.

Agency Implementation Road Map

Implementing computer vision systems in a city environment demands deliberate planning, cross-agency coordination, and a strong foundation of public trust. The computer vision framework outlined here is intended as a first step to help OTI and New York agencies responsibly explore and implement these technologies with clarity and accountability. We provide 7 implementation recommendations:

1. **Start with Tier 1 projects when possible:** Determine whether privacy-by-design systems, which do not collect identifying information, can address the core problem before moving on to complex or sensitive applications.
2. **Engage early with OTI:** Early consultation can help refine the use cases, avoid potential issues, and align with citywide AI guidance and best practices.
3. **Focus on proven solutions:** Be transparent about the technology's true capabilities and focus on established technologies that address specific agency needs and serve genuine public interests, rather than pursuing speculative or experimental applications.
4. **Take into account technology life-cycle costs:** When initiating a computer vision project, consider the full costs of deployment and ongoing maintenance. Given the rapid pace of technological advancement, deployed hardware such as cameras can become difficult to replace, while software and AI algorithms are easily upgraded, creating compatibility challenges. High deployment costs should be evaluated in relation to the expected lifespan of the technology. Consider whether existing technology infrastructure or data sources can be leveraged to achieve project goals
5. **Engage independent evaluators:** The computer vision technology landscape is evolving rapidly, making it difficult for agencies to stay current with emerging techniques and vendors. Engaging academic experts or other independent evaluators can provide valuable insight and support informed decision-making when assessing new technologies or potential vendors.
6. **Consider security:** What are the security implications of this project? Will new risks emerge, and if so, how will they be managed?

7. **Keep in mind privacy and transparency:**
    a. Whenever possible, edge processing rather than remote data centers should be used to minimize privacy risks by ensuring that video data are processed locally and remains on the device.
    b. In cases in which cloud processing is necessary, agencies should implement safeguards such as blurring personally identifiable information – including faces, license plates, and other sensitive features – before any data are transmitted or stored off-device.
    c. Computer vision pilot projects should be accompanied by robust public engagement efforts. This engagement includes providing clear signage in areas where technology is deployed (where appropriate) and hosting community information sessions to explain the project's purpose, scope, and safeguards.
    d. Finally, all deployments should be aligned with and integrated into the city's broader AI risk assessment framework.

Recommendations and Next Steps

This computer vision deployment framework will help New York City agencies navigate the complex landscape of visual analytics technologies while maintaining public trust and protecting civil liberties. We recommend that agencies consider adopting this tiered structure to evaluate computer vision projects based on data sensitivity and privacy implications, enabling beneficial innovation while ensuring appropriate safeguards.

Key Recommendations

- **Integrate framework guidance** into existing AI governance processes and agency technology planning;
- **Prioritize Tier 1 applications** to build operational experience and demonstrate clear public benefit with minimal privacy impact;
- **Engage early with OTI** for consultation on computer vision projects to ensure alignment with citywide policies; and
- **Establish regular review processes** to update this guidance as technologies evolve and community feedback emerges.

The framework's effectiveness will depend on collaborative implementation between OTI, city agencies, and community stakeholders. We recommend regular assessment of these guidelines to reflect new technological capabilities, emerging best practices, and evolving community needs.

As agencies consider computer vision deployments, we encourage starting with lower-risk applications that demonstrate clear public value while building institutional knowledge and community trust. This approach positions the city to thoughtfully expand into more complex applications when appropriate safeguards and oversight mechanisms are fully established.

We recommend that agencies contact the Office of Technology and Innovation early in their planning process for consultation on specific computer vision projects to ensure alignment with this framework and broader citywide technology policies.

If you'd like to learn more, see the full project, including a policy brief explaining the fellows' core recommendations, at aspenpolicyacademy.org/project/nyc-computer-vision-deployment-2025.

About the Aspen Policy Academy

The Aspen Institute's Policy Academy helps community leaders and experts across the political spectrum elevate their voices, influence key decisions, and strengthen democracy from the ground up. Our innovative training programs and resources equip people across sectors — from tech to the environment, science to civic engagement — with the skills to shape critical policy efforts. Learn more at aspenpolicyacademy.org.