Journal of Computer Science and Technology Studies
ISSN: 2709-104X
DOI: 10.32996/jcsts
Journal Homepage: www.al-kindipublisher.com/index.php/jcsts

JCSTS
AL-KINDI CENTER FOR RESEARCH AND DEVELOPMENT

| **RESEARCH ARTICLE**

Cross-Contextual Vision: Architecting Privacy-Preserving Multi-Modal AI Systems for Public Safety and Retail in Mixed Spaces

Akangsha Sunil Bedmutha
Independent Researcher, USA
Corresponding Author: Akangsha Sunil Bedmutha, **E-mail**: akangshasbedmutha@gmail.com

| ABSTRACT

The development of privacy-preserving computer vision systems using mixed-use environments will have to play this critical role. It suggests an innovative single framework alongside adaptive privacy protection based on the situational context, whilst the capability of the functionality traverses across the spectrum. The layered design makes it possible to regulate the implementation of privacy policies, detection of context, and analysis functions to achieve complex and advanced functions, such as cross-camera tracking and behaviour analysis, without violation of the right to privacy. Non-biometric tracking approaches, context-aware model switching, and policy observability are technical methods. The framework contends with the main issues of context boundary detection, performance optimisation, and cross-system integration using multi-modal sensing, edge computing, and a privacy standardised interface. A case study illustrates its effective realisation in a mixed-use urban centre and shows the framework is capable of mediating privacy through contextual transitions even when an emergency situation occurs. The article also points out that any form of governance cannot be effective without technical steps and human supervision, consultation with stakeholders, and even reporting of its proceedings.

| KEYWORDS

Privacy-preserving computer vision, context-aware surveillance, dynamic privacy enforcement, multi-modal sensing, ethical AI governance.

| ARTICLE INFORMATION

ACCEPTED: 01 August 2025 **PUBLISHED:** 28 August 2025 **DOI:** 10.32996/jcsts.2025.7.9.8

1. Introduction

In contemporary cities, the heralding of future urban environments is characterized by a testifying convergence of AI surveillance technology in tandem with the rampant phenomenon of global urbanization, posing an urgent question involving the development of vision systems that can traverse heterogeneous contexts under the aegis of fundamental privacy safeguards. This quandary is especially evident in multi-purpose venues, i.e., shopping centers, transport hubs, healthcare, and education systems, where the different collections of people interact in conditions of specific privacy and security demands.

The surveillance market's remarkable trajectory, documented by [1], continues its upward climb through this decade's conclusion. Security apprehensions, governmental safety programs, and widespread adoption of smart metropolitan initiatives primarily fuel this expansion [1].

In the future, society expects that visual recognition technologies will be integrated into the civic and mercantile spheres in a manner like never before. However, traditional paradigms that place the focus on the effectiveness of surveillance systems at the forefront of the conflict between privacy rights and data protection policies are more and more in conflict with the continuously developing regulatory framework, such as GDPR or CCPA, as well as emerging global standards. These rules currently include identifiable data along with behavioral trends transduced by AI mechanisms. Scholarly investigations consistently reveal inherent tensions between contextual awareness and privacy protection within existing system designs. Reference [2] highlights how contextually-aware frameworks necessarily gather substantial personal and environmental information for effective operation, creating fundamental privacy challenges demanding fluid protection strategies rather than rigid configurations.

This paper introduces an innovative, comprehensive framework for privacy-conscious, context-sensitive computer vision applications. The architecture enables advanced capabilities, including inter-camera subject tracking, object identification, and behavioral examination while fluidly adjusting privacy shields based on situational parameters, legal mandates, and moral considerations.

2. The Privacy-Context Paradox in Mixed Spaces

Multi-purpose settings create distinctive hurdles for vision systems. Any person might traverse several contexts with shifting privacy expectations momentarily—browsing retail merchandise (commercial domain), clearing security stations (safety zone), then accessing medical services (sensitive personal territory).

Conventional vision architectures typically deploy uniform processing methodologies across these varied contexts, either amassing excessive information or implementing overly stringent privacy measures, hampering functionality. This binary methodology fails to address real-world environmental fluidity, where context determines appropriate privacy boundaries [3].

2.1 Context Transitions in Everyday Scenarios

Picture a typical metropolitan scenario: someone enters a shopping complex (retail environment), passes through transport facilities (safety zone), and ultimately reaches hospital premises (healthcare setting). Each domain necessitates processing identical visual data according to different priorities:

Retail environments warrant anonymized engagement assessment, heat mapping, and merchandise interaction analysis. Transit zones necessitate security screening, prohibited item detection, and crowd management tactics. Healthcare settings require patient recognition (with explicit consent), fall monitoring, and distress identification [4].

Current implementations typically struggle with these transitions, creating either privacy vulnerabilities or operational constraints at contextual boundaries. Recent examinations of integrated sensing networks emphasize frequent privacy and security weaknesses emerging at transition points between contextual domains [4]. The core challenge involves developing adaptive frameworks capable of modifying privacy preservation approaches based on contextual factors while maintaining functionality across diverse settings—similar challenges appear in IoT-driven smart city implementations where Software Defined Networking approaches offer potential remedies [3].

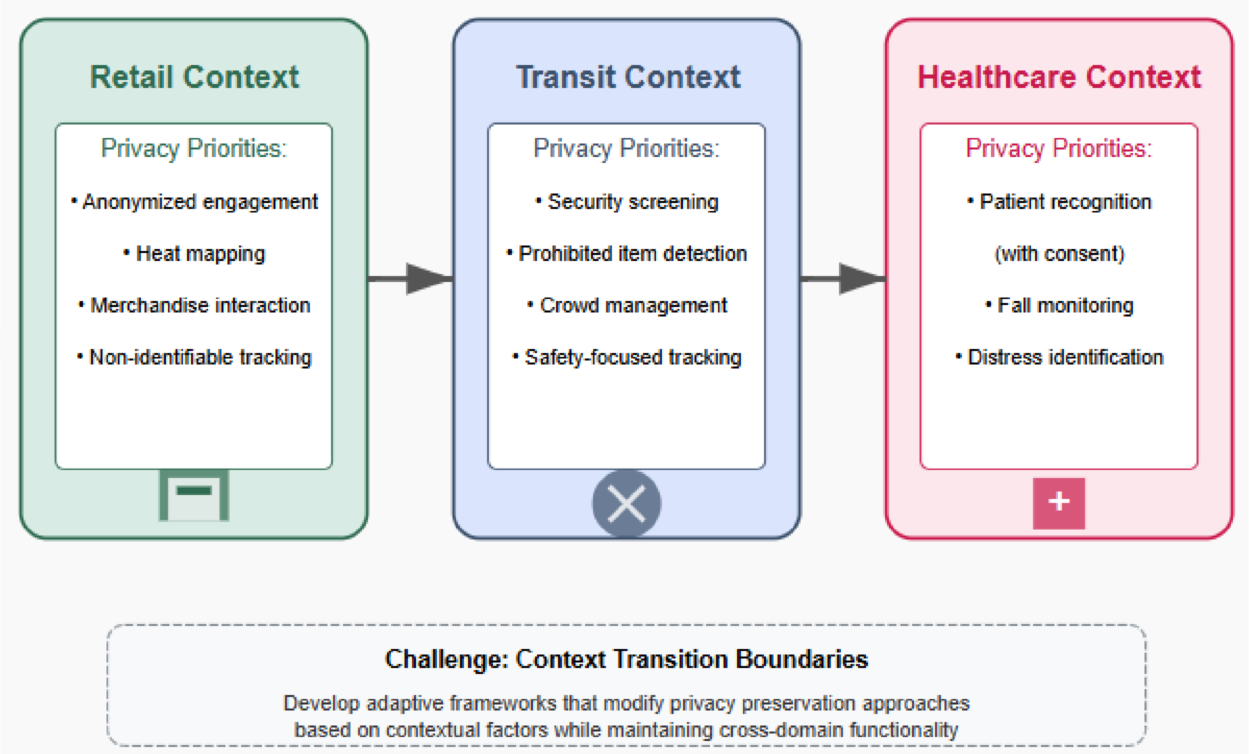
Retail Context
Privacy Priorities:
• Anonymized engagement
• Heat mapping
• Merchandise interaction
• Non-identifiable tracking
Transit Context
Privacy Priorities:
• Security screening
• Prohibited item detection
• Crowd management
• Safety-focused tracking
Healthcare Context
Privacy Priorities:
• Patient recognition
(with consent)
• Fall monitoring
• Distress identification
Challenge: Context Transition Boundaries
Develop adaptive frameworks that modify privacy preservation approaches
based on contextual factors while maintaining cross-domain functionality

Fig 1: The Privacy-Context Paradox in Mixed Spaces [3, 4]

3. Architectural Framework for Cross-Contextual Vision Systems

The proposed architecture presents a stratified design separating context recognition, privacy rule enforcement, and functional capabilities. This layered methodology builds upon established privacy-enhancing architectural patterns while addressing specific challenges inherent to context-aware vision applications [5].

3.1 Context Detection Layer

This foundational component continuously evaluates environmental, spatial, and temporal markers to determine operational context. Key elements include spatial-awareness modules mapping physical boundaries between contexts using geofencing and visual landmarks; temporal context recognition identifying time-dependent contextual shifts (business hours versus emergency scenarios); activity detection recognizing context-defining behaviors within visual fields; and jurisdictional recognition applying appropriate regulatory frameworks based on location. This multifaceted approach to contextual awareness mirrors concepts found within fog computing architectures for IoT environments, where contextual understanding facilitates more effective resource allocation and security management [5].

3.2 Dynamic Privacy Policy Enforcement Layer

Operating between raw sensor inputs and analytical functions, this layer implements context-appropriate privacy protections. The automatic redaction engine performs real-time anonymization of sensitive visual elements based on context, while consent management tracks individual privacy preferences where identification becomes permissible. Data minimization protocols implement selective processing based on legitimate purpose, and retention policy enforcement applies context-specific data storage limitations. These capabilities reflect emerging privacy engineering approaches emphasizing dynamic policy enforcement as a key mechanism balancing utility and privacy within AI systems [6].

3.3 Functional Capability Layer

This layer houses analytical capabilities that operate on privacy-filtered data. Anonymous tracking modules implement privacy-preserving cross-camera tracking without persistent identifiers, while behavioral analysis engines examine movement patterns and interactions with context-specific thresholds. Anomaly detection identifies concerning patterns with context-appropriate sensitivity, and integration interfaces connect downstream systems with appropriate privacy guards. This architecture aligns with current research regarding effective AI governance frameworks, incorporating privacy-preserving mechanisms directly into functional systems rather than treating privacy as an external constraint [6].

Layer	Component	Privacy Protection (1-10)	Functional Capability (1-10)	Implementation Priority
Context Detection	Spatial-awareness Modules	5	8	High
	Temporal Context Recognition	6	7	Medium
	Activity Detection	4	9	High
	Jurisdictional Recognition	8	5	Medium
Privacy Policy Enforcement	Automatic Redaction Engine	9	6	Critical
	Consent Management	10	4	High
	Data Minimization	8	7	Critical
	Retention Policy Enforcement	9	3	Medium
Functional Capability	Anonymous Tracking	7	8	High
	Behavioral Analysis	5	9	Medium
	Anomaly Detection	6	10	High
	Integration Interfaces	7	8	Medium

Table 1: Layered Architecture Components and Their Characteristics [5, 6]

4. Technical Approaches to Privacy-Preserving Computer Vision

4.1 Privacy-First Tracking Without Biometric Identifiers

Traditional cross-camera tracking heavily depends on facial recognition or similar biometric markers. The proposed framework instead employs temporary pseudonymous identifiers (short-duration tracking tokens expiring at context boundaries); appearance-based tracking utilizing non-identifying visual attributes like clothing color schemes; trajectory modeling predicting movement paths based on motion vectors rather than identity; and differential privacy techniques introducing calibrated noise, obscuring individual identities while preserving aggregate data patterns. These methodologies align with recent breakthroughs in non-biometric tracking for intelligent transport networks, where privacy-conscious object tracking successfully maintains functionality while safeguarding personally identifiable details [7].

4.2 Context-Based Model Switching

The system dynamically toggles between different behavioral and analytical models according to detected contexts. Model containers encapsulate context-specific models with appropriate privacy characteristics, while transfer mechanisms establish protocols for transitioning between contexts with proper data handling procedures. Confidence thresholds determine context-appropriate certainty requirements before actions proceed, and multi-objective optimization balances privacy, security, and functionality based on contextual weightings. This approach mirrors emerging research regarding context-aware AI systems adapting behavior based on situational factors, similar to techniques explored in recent studies concerning context-sensitive model selection for multi-domain adaptive systems [8].

4.3 Policy-Based Observability

The framework delivers transparent operation through automatic audit trails, providing context-specific logging with appropriate detail levels. Privacy impact visualization presents real-time representations of active privacy measures, while stakeholder interfaces offer context-appropriate dashboards for security personnel, privacy officers, and affected subjects. Explainability modules provide context-sensitive explanations regarding system decisions. These observability mechanisms address growing recognition that transparency and accountability form essential components within privacy-preserving AI systems, particularly where stakeholders maintain diverse and occasionally competing interests regarding system operation [7]. By providing appropriate visibility into system behaviors without compromising privacy protections, these mechanisms support both regulatory compliance and public confidence, following established principles within adaptive AI governance frameworks [8].

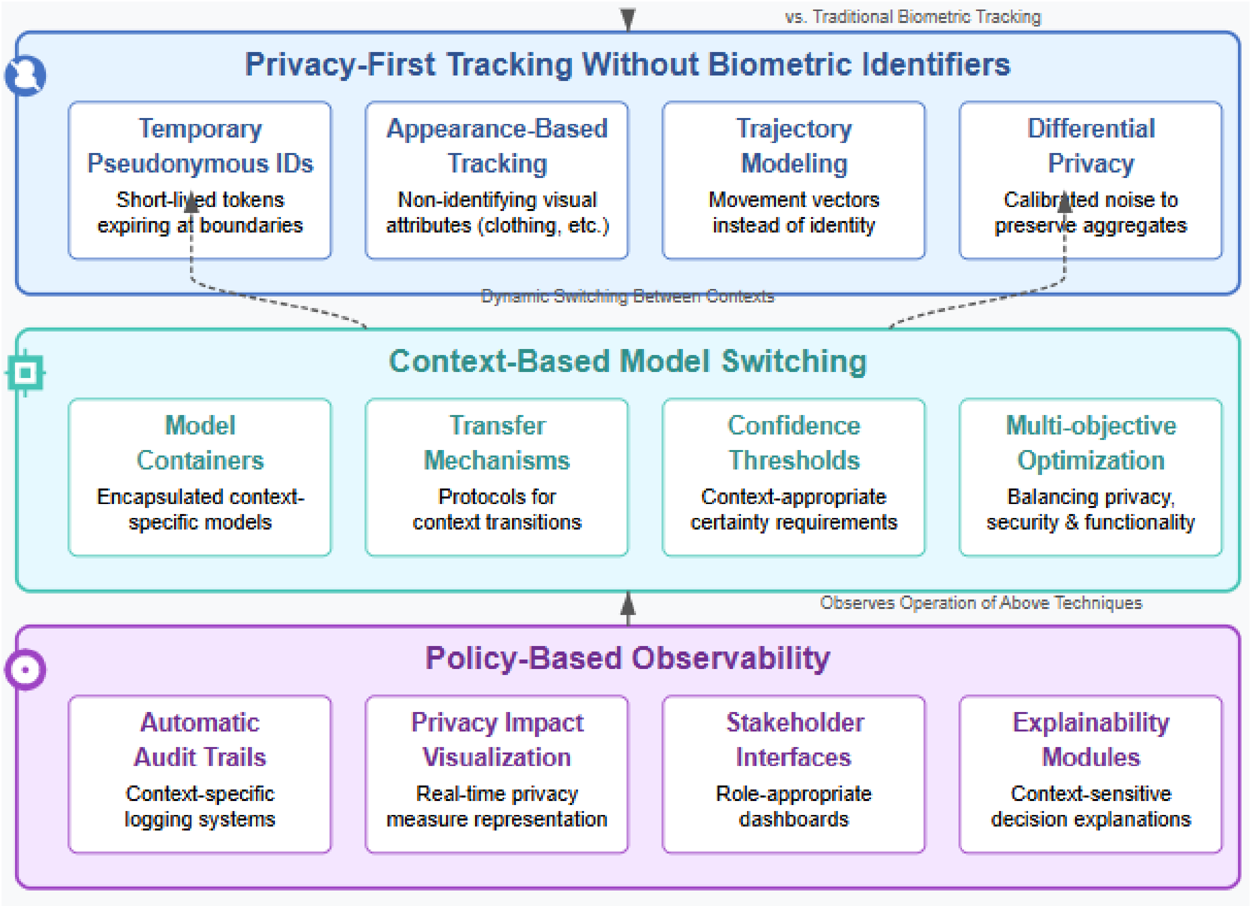
vs. Traditional Biometric Tracking
Privacy-First Tracking Without Biometric Identifiers
Temporary Pseudonymous IDs
Short-lived tokens expiring at boundaries
Appearance-Based Tracking
Non-identifying visual attributes (clothing, etc.)
Trajectory Modeling
Movement vectors instead of identity
Differential Privacy
Calibrated noise to preserve aggregates
Dynamic Switching Between Contexts
Context-Based Model Switching
Model Containers
Encapsulated context-specific models
Transfer Mechanisms
Protocols for context transitions
Confidence Thresholds
Context-appropriate certainty requirements
Multi-objective Optimization
Balancing privacy, security & functionality
Observes Operation of Above Techniques
Policy-Based Observability
Automatic Audit Trails
Context-specific logging systems
Privacy Impact Visualization
Real-time privacy measure representation
Stakeholder Interfaces
Role-appropriate dashboards
Explainability Modules
Context-sensitive decision explanations

Fig 2: Technical Approaches to Privacy-Preserving Computer Vision [7, 8]

5. Technical Challenges and Solutions

5.1 Challenge: Context Boundary Detection

Precisely identifying transitions between contexts remains a technically demanding task. The framework addresses this through multi-modal confirmation utilizing complementary signals (Wi-Fi positioning, Bluetooth beacons, visual landmarks), verifying context transitions; probabilistic context models handling uncertainty within boundary areas through graduated policy application; and default-private transitions applying maximally restrictive applicable policies within ambiguous boundary zones. These approaches build upon established research concerning context-awareness for multi-sensor data fusion within smart environments, where complementary sensing modalities significantly enhance robust context detection throughout complex, dynamic settings [9].

5.2 Challenge: Performance Under Privacy Constraints

Privacy-preserving techniques frequently impose computational burdens. The framework mitigates these through edge-processed anonymization, performing privacy operations before data transmission; context-aware resource allocation, distributing computational resources according to contextual significance; and optimizations for privacy-preserving operations, including specialized hardware acceleration for anonymization tasks. These solutions align with recent advances regarding privacy-preserving edge computing architectures, where distributing privacy operations across edge nodes reduces latency while maintaining robust privacy guarantees for sensitive information [10].

5.3 Challenge: Cross-System Integration

Mixed environments typically contain multiple systems from varied vendors with differing privacy capabilities. The framework addresses this through privacy middleware providing intermediation layers enforcing policies across heterogeneous systems; standardized privacy APIs offering common interfaces for policy enforcement across diverse systems; and privacy credential exchange mechanisms enabling secure communication regarding privacy context between systems. These integration approaches address critical gaps identified within multi-sensor fusion literature concerning interoperability of privacy controls across heterogeneous systems [9]. By establishing common privacy interfaces and credential exchange protocols, the framework enables consistent privacy enforcement even across environments with diverse technical infrastructures, becoming increasingly vital as privacy regulations evolve toward outcome-based requirements rather than prescriptive technical specifications [10].

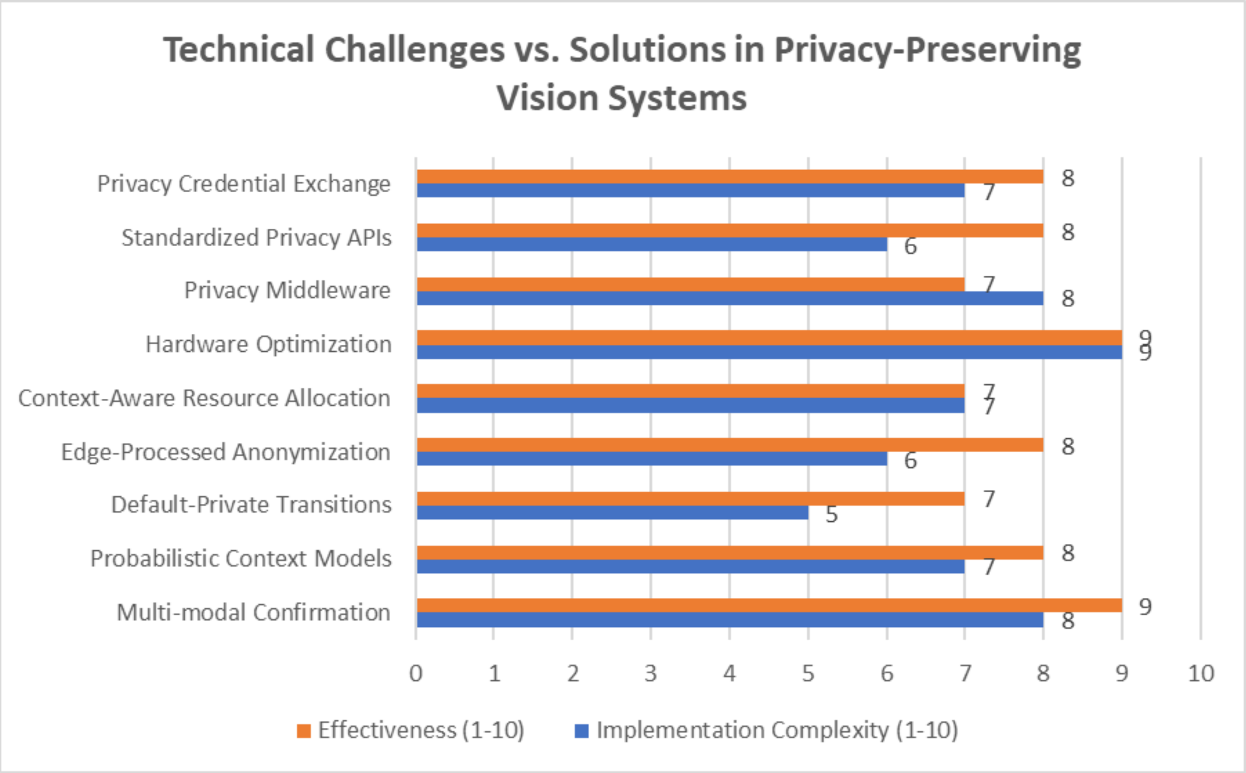

Fig 3: Implementation Complexity and Effectiveness of Privacy Solutions [9, 10]

6. Case Study: Mixed-Use Urban Center

A hypothetical implementation within modern urban development, incorporating retail, transit, healthcare, and residential components, illustrates practical framework application. The integrated system successfully manages privacy across contextual transitions while maintaining operational effectiveness throughout each domain. When individuals move from shopping districts through transit hubs toward medical centers, privacy protections automatically adjust according to context-appropriate policies. This dynamic adaptation reflects emerging capabilities within smart city integrations for public safety, where IoT-enabled systems increasingly maintain appropriate operation across different urban contexts while significantly enhancing emergency response capabilities [11].

During simulated emergency scenarios, the system demonstrated appropriate flexibility, temporarily enhancing certain capabilities while preserving core privacy protections. Post-incident analysis confirmed proper handling of privacy boundaries even during cross-context operations. This balanced approach to emergency handling aligns with emerging smart city frameworks where IoT integration enables nuanced, effective emergency responses without compromising essential privacy protections [11]. The system's ability to maintain appropriate privacy boundaries during emergency operations addresses critical requirements within next-generation urban safety infrastructure, where maintaining public trust demands demonstrating that emergency capabilities enhance safety without unnecessarily compromising privacy.

The case study confirmed contextual transitions can be effectively managed through layered architectural approaches, with privacy policies dynamically adjusting to maintain both functional effectiveness and appropriate privacy protections. This effective experiment in complex urban areas implies that possible adaptation to other mixed-use areas with different desires for privacy per zone and activity may prove a valid idea of continuing to adjust smart city technologies in which capability is balanced with privacy protection [11].

7. Ethical Considerations and Human Oversight

As much as technical measures are given at the foundation, ethical governance is crucial. Ethics review procedures will give a periodic check of the definition of the contexts and privacy policies so that technical implementations are within the values and expectations of society. Human-in-the-loop monitoring achieves proper human monitoring of context-sensitive decisions in cases of possible bounded cases at the edge of automation, where ethical complexity or ambiguity is present. This strategy takes into consideration that though AI systems excel in huge amounts of processing data, human judgment is still critical to deal with boundary cases, confirming results, and demonstrating fairness particularly in circumstances that use individual applications that are subjective and sensitive to privacy like those that need contextual interpretation that are beyond the scope of a computer to comprehend [12].

The consultation of the stakeholders in the form of the continued engagement with the involved communities, as well as a privacy advocate, serves to ensure that diverse views are considered in the governance of the system. This multi-stakeholder approach represents a critical component within responsible AI deployment throughout public spaces, where different groups maintain varying privacy expectations and concerns. UNESCO's work regarding AI governance emphasizes bringing together diverse stakeholders—policymakers, technical experts, civil society representatives, and affected communities—to create more robust governance frameworks reflecting varied societal needs and cultural contexts [13]. Through formal mechanisms for community input, the framework acknowledges that privacy expectations vary across cultural, social, and individual factors, requiring consideration throughout system design and operation.

Transparency reporting through public communication regarding system operation and privacy impacts completes the governance framework. This transparency enables external evaluation and accountability, fostering public trust through visible commitment toward responsible operation. As surveillance capabilities grow increasingly sophisticated, commitment toward transparency becomes vital for maintaining the social license necessary for these systems operating within democratic societies [13]. By combining robust technical measures with comprehensive ethical governance, the framework establishes sustainable approaches toward computer vision within mixed-use environments, respecting privacy as a fundamental right while enabling beneficial functionality across diverse contexts.

8. Conclusion

With computer vision AI moving more and more into the environment and publicly visible spaces and even commerce, binary privacy vs functionality will be insufficient to address. The vision framework suggested in this paper proposes a delicate architectural view, using the fluidity of the real-world environments. These systems can appropriately support legitimate security and other needs in retail and security and healthcare and serve other areas on a case-by-case basis, considering context, and balance the fundamental privacy rights by dynamically tailoring privacy protections and functional capabilities to the requirements of individuals, businesses, and the context. This context-aware workflow is one such evolution that is required in the architecture of computer vision, and will only serve to be more important as societies move forward into the extremely rich and problematic sets of technology use, privacy, and personal space in the next several decades. The future of intelligent environments cannot rely on those systems that are able to see, but it requires systems that are able to perceive the context and use the relevant ethical schemes in relation to their observations. This paper is an engineering roadmap of the direction to that future, one where the privacy and features work in tandem with one another because of intelligent and context-sensitive design.

Funding: This research received no external funding
Conflicts of interest: The authors declare no conflict of interest
Publisher's Note: All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers.

References

[1] Alessandra D P et al., (2016) Context-Awareness for Multi-sensor Data Fusion in Smart Environments, ResearchGate, 2016. https://www.researchgate.net/publication/311933955_Context-Awareness_for_Multi-sensor_Data_Fusion_in_Smart_Environments

[2] David R et al., (2024) Large Language Models: A New Approach for Privacy Policy Analysis at Scale, arXiv:2405.20900v1, 2024. https://arxiv.org/html/2405.20900v1

[3] Google Cloud, (n.d) What is Human-in-the-Loop (HITL) in AI & ML?". https://cloud.google.com/discover/human-in-the-loop?hl=en

[4] IEEE Public Safety Technology, (n.d) Smart City Integration: How IoT is Reducing Emergency Response Times and Saving Lives. https://publicsafety.ieee.org/topics/smart-city-integration-how-iot-is-reducing-emergency-response-times-and-saving-lives/

[5] Jihoon M et al., (2024) Object detection under the lens of privacy: A critical survey of methods, challenges, and future directions, ICT Express, Volume 10, Issue 5, 2024. https://www.sciencedirect.com/science/article/pii/S2405959524000833

[6] Kaiqian Q et al., (2023) Privacy and Security in Ubiquitous Integrated Sensing and Communication: Threats, Challenges and Future Directions, ResearchGate, 2023. https://www.researchgate.net/publication/372827287_Privacy_and_Security_in_Ubiquitous_Integrated_Sensing_and_Communication_Threats_Challenges_and_Future_Directions

[7] Markets and Markets, (2024) AI in Video Surveillance Market Size, Share, and Trends, 2025 - 2030, 2024. https://www.marketsandmarkets.com/Market-Reports/ai-in-video-surveillance-market-84216922.html

[8] Mehdi G et al., (2019) A Context-aware Privacy-preserving Method for IoT-based Smart City Using Software Defined Networking, ResearchGate, 2019. https://www.researchgate.net/publication/333026390_A_Context-aware_Privacy-preserving_Method_for_IoT-based_Smart_City_Using_Software_Defined_Networking

[9] Ming T et al., (2018) Multi-layer cloud architectural model and ontology-based security service framework for IoT-based smart homes, Future Generation Computer Systems, Volume 78, Part 3, 2018. https://www.sciencedirect.com/science/article/abs/pii/S0167739X16305775

[10] Nandish C et al., (2025) A Survey of Adversarial Defenses in Vision-based Systems: Categorization, Methods and Challenges, arXiv:2503.00384, 2025. https://arxiv.org/abs/2503.00384

[11] Pramod J et al., (2011) Preserving Privacy in Context-Aware Systems, ResearchGate, 2011. https://www.researchgate.net/publication/221406184_Preserving_Privacy_in_Context-Aware_Systems

[12] UNESCO, (2023) How to adopt a multistakeholder approach to AI governance in Southern Africa? 2023. https://www.unesco.org/en/articles/how-adopt-multistakeholder-approach-ai-governance-southern-africa

[13] Xianzhi Z et al., (2024) A Survey on Privacy-Preserving Caching at Network Edge: Classification, Solutions, and Challenges, arXiv:2405.01844v1, 2024. https://arxiv.org/html/2405.01844v1