

Research on Privacy Boundaries in Machine-Vision Data Collection

Junfeng Wen^{1*}

¹ ZK Skynet Holdings Limited, Hong Kong, 999077, China

* wjf@zkskynet.com

<https://doi.org/10.70695/IAAI202603A1>

Abstract

After machine vision evolves from simple recording to recognition, tracking, and inference, image collection in public places is no longer merely a question of camera installation. It has become an issue concerning the intensity of personal-information processing and the boundaries of personality rights. Based on a review of recent English-language journal research and governance concerns, this paper concludes that public visibility is not equivalent to permission for continuous computation, and that deleting original images does not eliminate derivative risks. The paper proposes a judgment framework consisting of three dimensions: scenario sensitivity, identifiability, and purpose necessity. It also proposes strategies in five areas: impact assessment, minimization parameters, edge anonymization, permission auditing, and full-lifecycle accountability. The study argues that machine-vision collection must have a specific necessary purpose as a precondition and should use low-intrusion alternatives and edge-side reduction by default, so that a proportional balance can be achieved among security, efficiency, and personal dignity.

Keywords Machine Vision; Data Collection; Privacy Boundary; Personal Information Protection; Face Recognition; Edge Anonymization

1 Introduction

Machine vision is entering transportation hubs, campuses, hospitals, shopping malls, communities, and factories. After high-definition cameras, depth sensing, cross-camera tracking, and visual foundation models are combined, the system can not only see a person, but also calculate identity, dwell time, trajectory, posture, and group relationships. A recent study comparing surveillance in public spaces, hospitals, and schools found that privacy concerns and perceived trust vary by setting [1]. The privacy boundary therefore depends on both the processing purpose and the situation in which observation occurs.

This paper is organized as an introduction, a review of current research, a statement of the boundary problem, governance strategies, and a conclusion. Facial anonymization research emphasizes that re-identification risk depends on the wider processing context, including the actors and safeguards involved, rather than on the visual appearance of an image alone [2]. Accordingly, this study examines on-site perception, feature extraction, identity association, semantic inference, and subsequent circulation as distinct points of control.

2 Analysis of Current Research Status

Recent research has moved beyond treating notice and consent as sufficient safeguards. A lifecycle survey organizes privacy risks across face-data generation, model inference, and storage [3]. Experimental work on face de-identification shows that reducing identity disclosure while preserving image utility requires testing against recognition models beyond the one used during development [4]. Work on reversible anonymization likewise shows that authorized recovery creates a separate access-control requirement [5]. These findings support a graded boundary based on processing purpose, identifiability, and the possibility of later reuse.




Fig. 1. Machine-vision data lifecycle and privacy risk chain

Fig. 1 shows that risks are continuously amplified along the chain of perception, feature extraction, inference, circulation, and application. Governance should not stop at equipment installation. It should cover feature retention, cross-database association, algorithm updates, and result-based decisions to prevent derivative information beyond original images from continuing to identify individuals.

Governance must account for both organizational access and downstream reuse. A recent study of online photograph platforms proposes dynamic, person-specific access controls for faces appearing in shared images [6]. In public video systems, the same principle calls for documenting who can retrieve identifiable frames, what derived attributes can be exported, and which decisions may rely on those attributes.

Technical safeguards cannot be assessed only at capture. A recent review distinguishes appearance-based, identity-based, reversible, and recognition-system-oriented privacy protections and compares their evaluation requirements [7]. For machine-vision collection, this implies that raw frames, extracted features, linked trajectories, and model outputs each require a specific retention rule and an independently testable privacy control.

Table 1. Visual-data risk levels and boundary requirements

Risk level	Typical visual data	Boundary requirement
Low risk	Device status, anonymous counts, vehicle traffic	Edge-side statistics, aggregated results, and prohibition of backtracking to individuals
Medium risk	Clear images, trajectory IDs, behavior labels	Restricted field of view, shortest retention, conspicuous notice, and permission logs
High risk	Face templates, gait vectors, health or emotion inference	Sufficient necessity, alternative schemes, impact assessment, and independent audit
Prohibited zone	Private-space collection, irrelevant body parts, unjustified cross-scenario tracking	Prohibited in principle; generalized consent may not break the bottom line

Table 1 shows that machine-vision data risk becomes more stable and predictive as identification capability improves. Low-risk data can be processed in aggregated form. High-risk biometric features and sensitive inferences require proof of necessity, comparison with alternatives, and independent audit. Collection in private spaces should be directly excluded.

3 Problem Statement: Failure of Boundaries in Machine-Vision Collection

First, public visibility and computable use are conflated. The fact that a person is briefly seen by others on a street, in a shopping mall, or on a campus does not mean that the person's facial template, route of movement, and companion relationships may be recorded over the long term. The advantage of machine vision lies in continuous observation at low cost with repeatable characteristics, and in the ability to store information and quickly retrieve related data.

Second, the concept of collection often remains at the raw-pixel level. Many systems claim not to store video, but still retain derivative data such as vector templates, trajectory IDs, target boxes, logs, and model parameters. Such derivative data is more suitable for automatic matching than original images.

Third, consent is formalized in highly dependent environments. In public transportation, hospitals, campuses, and workplaces, individuals have no real ability to opt out, and broad notices at entrances cannot explain algorithmic functions, retention periods, and recipients. Fourth, purpose drift and fragmented responsibility coexist. Safety-inspection cameras may be upgraded for attendance control, and passenger-flow statistics may be converted into member profiling. Venue managers, equipment vendors, algorithm providers, and cloud-service providers jointly participate in processing, but contracts often avoid clarifying who decides the purpose, who stores the original images, and who undertakes deletion responsibility. Fifth, de-identification is treated as absolute safety. Background, clothing, time, location, and trajectory may continually combine into new identities, and visual models may also produce re-identification, membership inference, or model inversion.

4 Strategies: Graded Boundaries and Full-Lifecycle Governance

To overcome the above shortcomings, this paper proposes a model consisting of scenario sensitivity, identifiability, and purpose necessity. Scenario sensitivity distinguishes private spaces, semi-public high-sensitivity venues, and general public places, while incorporating the vulnerability of subjects such as minors, patients, and workers. Identifiability does not look only at whether a face is clear; it also examines whether the system can stably associate a person through gait, body posture, trajectory, or auxiliary databases. Purpose necessity requires processors to prove that the objective is specific and lawful, that collection has a direct functional link with the objective, and that no equally effective but less intrusive alternative exists. The higher these three dimensions are, the greater the proof and security obligations borne by the processor.

"Scenario + identifiability + necessity" determines the collection boundary




Fig. 2. Three-dimensional boundary determination for machine-vision data collection

Fig. 2 shows that the determination process transforms "whether filming occurs" into a comprehensive review of scenario, identifiability, and purpose necessity. The three indicators together define red, yellow, and green zones, avoiding the relaxation of high-risk recognition merely because a space is public or because broad notice has been given.

In terms of boundary conclusions, three thresholds can be set: red, yellow, and green. The red zone includes hidden collection in private spaces, unnecessary biometric recognition, emotion or health inference unrelated to the original purpose, and cross-scenario identity tracking without legal basis; these should all be prohibited. The yellow zone includes systems with legitimate purposes but high risks, such as campuses, hospitals, transportation hubs, and access control for important venues; these must meet requirements for impact assessment, conspicuous notice, alternative verification, short-term retention, and independent audit. The green zone includes low-identification tasks such as edge-side anonymous people counting, equipment-fault detection, and vehicle-flow analysis, but later functional upgrades must not be affected by relaxed initial review.

The governance strategy should first assess impact and register the purpose. The project list should record device location, field of view, algorithm model, input and output, recognition thresholds, recipients, retention period, and supplier permissions. Adding face or gait modules, increasing resolution, connecting external databases, changing recipients, or using samples for training should all be regarded as substantive changes and reassessed. Second, data minimization should be transformed into engineering parameters, including lens angle, focal length, resolution, frame rate, collection time period, region of interest, and event-triggering conditions. If only headcount statistics are needed, original images should not be uploaded. If only qualification verification is needed, cards, passwords, or human review should be prioritized.

Table 2. Full-lifecycle governance responsibilities and evidence

Stage	Main control measures	Verifiable evidence
Project approval	Purpose registration, comparison with low-intrusion alternatives, impact assessment	Assessment report, lens coverage map, algorithm list
Collection	Limitation of angle, resolution, frame rate, region, and time period	Device configuration, masking-area records, trigger logs
Processing	Edge anonymization, feature separation, encryption, and graded permissions	Model version, key records, access logs
Circulation	Recipient restrictions, interface approval, prohibition on supplier retraining	Contract clauses, call register, sharing list
Withdrawal	Scheduled deletion, backup cleanup, appeal review	Deletion certificate, review record, audit report

Table 2 shows that privacy boundaries must be transformed into verifiable evidence. Purpose registration, lens settings, permission logs, contract restrictions, and deletion certificates are connected with each other. This constrains both suppliers and internal personnel and can reduce the individual's evidentiary burden in disputes.

The technical architecture should be built on the principle of edge-side preprocessing, no raw data at the platform by default, feature generation on demand, and immediate deletion of results. The camera end or edge node can first perform regional cropping, real-time blurring, feature compression, and removal of irrelevant targets. The platform should accept only aggregated values or encrypted features. Identity mapping tables, business feature databases, and log systems should be stored separately, while keys and administrative privileges should be separated. De-identification schemes should also be tested from the attacker's perspective, including background association, cross-database matching, reverse recovery, and model inversion. One-time masking must not be treated as permanent anonymity.

Venue managers are responsible for the purposes and scenarios within their venues. Device and algorithm providers are responsible for device default settings, model security, and vulnerability repair. Cloud-service providers bear certain responsibilities for storage isolation and access auditing. Contracts should specify prohibited uses, data-storage locations, retention periods, approval of functional changes, incident reporting, exit migration, and destruction certificates. Individuals should be able to query collection rules, main algorithmic functions, and channels for exercising rights. When visual results affect access, disciplinary evaluation, insurance, or pricing, human review and reasons must also be provided.

Every substantive change triggers reassessment, logging, deletion, and appeal review




Fig. 3. Full-lifecycle accountability loop of a machine-vision system

Fig. 3 shows that the collection boundary is not a one-time approval. It is continuously updated as the purpose, algorithm, recipient, and retention status change. Every substantive change triggers reassessment, and logs, deletion, and appeals jointly form a continuous accountability chain.

5 Conclusion

The privacy boundary of machine-vision data collection is essentially the boundary of power from being seen to being identified, inferred, and decided upon. Visibility in public places is not equivalent to unlimited computation, deletion of original images is not equivalent to disappearance of derivative risks, and formal consent cannot replace necessity review. The three-dimensional model proposed in this paper places reasonable expectations in scenarios, technical identification capability, and purpose necessity within the same framework, and transforms abstract principles into project judgments through red, yellow, and green thresholds.

Future governance should further expand rules for gait, posture, emotion, and health inference, and create unified standards for re-identification testing, impact assessment, and independent audit. Only when privacy boundaries are embedded in device parameters, edge architectures, contractual responsibilities, and deletion certificates can machine vision protect the bottom line of human spiritual freedom and information autonomy while serving public security, production efficiency, and commercial innovation.

Acknowledgement

This work was supported without any funding.

Conflicts of Interest

The authors declare no conflicts of interest.

References

1. Choung, H., David, P., & Ling, T.-W. (2024). Acceptance of AI-powered facial recognition technology in surveillance scenarios: Role of trust, security, and privacy perceptions. *Technology in Society*, 79, 102721.
2. Hellmann, F., André, E., Benouis, M., Buchner, B., & Mertes, S. (2024). Anonymization of faces. *Datenschutz und Datensicherheit*, 48(6), 364-367.
3. Sun, Z., & Liu, Z. (2025). Ensuring privacy in face recognition: A survey on data generation, inference and storage. *Discover Applied Sciences*, 7, 441.
4. Hanawa, G., Ito, K., & Aoki, T. (2024). Face image de-identification based on feature embedding. *EURASIP Journal on Image and Video Processing*, 2024, 25.
5. Xu, S., Chang, C.-C., Nguyen, H. H., & Echizen, I. (2024). Reversible anonymization for privacy of facial biometrics via cyclic learning. *EURASIP Journal on Information Security*, 2024(1), 1-13.
6. Santoso, A., Huda, S., Koder, Y., & Nogami, Y. (2025). Facial privacy protection with dynamic multi-user access control for online photo platforms. *Future Internet*, 17(3), 124.
7. Wang, M., Li, S., Zhang, X., & Feng, G. (2025). Facial privacy in the digital era: A comprehensive survey on methods, evaluation, and future directions. *Computer Science Review*, 58, 100785.

Biographies

1. **Junfeng Wen** Master, Senior Big Data Engineer. Specializing in the research and application of artificial intelligence and big data, he has led and participated in multiple key provincial scientific research projects. His scientific and technological achievements in artificial intelligence have been appraised as reaching the internationally advanced level. He has published numerous papers focused on the commercialization of scientific and technological achievements.

機器視覺數據採集的隱私邊界研究

溫峻峯¹

¹中科天網控股有限公司，香港，中國，999077

摘要：機器視覺從簡單的記錄轉變為識別、跟蹤、推斷之後，公共場所的圖像採集就不只是一個攝像頭安裝的問題，而成了關乎到個人資訊處理強度和人格權益邊界的問題。本文在查閱近三年的英文期刊文獻的基礎上，得出結論認為公共可見並不同於可以被持續計算，刪除原圖也不意味著消除了衍生的風險。文章提出場景敏感度、可識別性、目的必要性三個維度的判定框架，從影響評估、最小化參數、邊緣匿名化、權限審計、全生命週期問責五個方面提出策略。研究認為機器視覺採集必須有具體的必要目的作為前置條件，以低侵擾代替、端側減量為預設方式，才能在安全、效率和個人尊嚴之間達到比例的平衡。

關鍵詞：機器視覺；數據採集；隱私邊界；人臉識別；邊緣匿名化

1. 溫峻峯，碩士，大數據高級工程師，專注人工智慧、大數據方向的研究和應用，有豐富的科技成果轉化工作經驗，牽頭和參與多個省級重點科研項目，其人工智慧的科技成果獲得國際先進的成果鑑定，有多篇關於科技成果轉化的論文，並轉化成產品取得不錯的市場反響。