

Privacy-Aware Computer Vision for Ambient Assisted Living: Balancing Older Adults' Autonomy, Safety, and Trust

Maurice C. Dunn

School of Computing, Clemson University, Clemson, SC, USA.

mauricedunn@clemson.edu

Abstract

Ambient assisted living systems increasingly rely on computer vision to support older adults in maintaining independence while enabling timely detection of falls, health deterioration, and unsafe conditions. However, camera-based monitoring introduces substantial privacy risks that can undermine autonomy, dignity, and trust. This paper examines privacy-aware computer vision for ambient assisted living as a systems problem spanning architecture, data governance, edge intelligence, robustness, fairness, and policy. It argues that meaningful privacy protection cannot be achieved solely through algorithmic modifications or encryption, but requires a layered sociotechnical design in which raw visual data are interpreted close to the source, sensitive information is abstracted before transmission, and residents retain control over how monitoring outputs are used. The discussion considers structural trade-offs between utility and privacy, the role of edge processing in reducing data leakage, the importance of dynamic consent infrastructures, and the challenges of ensuring equitable performance across diverse older populations. The paper also addresses deployment sustainability, regulatory alignment, and the need for continuous auditing of vision systems in domestic settings. By integrating technical and governance perspectives, the paper provides a system-level framework for balancing safety, autonomy, and trust in real-world ambient assisted living deployments.

Keywords

ambient assisted living, computer vision, privacy, older adults, autonomy, trust, edge computing, data governance.

1. Introduction

The global growth of ageing populations has intensified interest in ambient assisted living (AAL) systems that enable older adults to remain in their homes while receiving timely assistance and health monitoring. These systems integrate sensing, networking, data analytics, and service delivery to detect falls, monitor daily activities, and identify behavioral changes that may indicate emerging health problems. Existing AAL frameworks vary widely in their sensing modalities, architectural models, and quality attributes, but many prioritize functional performance while treating privacy as a secondary constraint [1]. At the same time, the field has matured through systematic surveys of ambient-assisted living tools, revealing that successful deployments depend not only on recognition accuracy but also on user acceptance, reliability, and compatibility with domestic life [2]. Vision-based activity recognition has emerged as a particularly powerful approach because it can capture detailed physical context, support unobtrusive fall detection, and track complex motion patterns [3]. Yet the presence of

cameras in private living spaces raises deep concerns about surveillance, embarrassment, data misuse, and loss of personal control.

The central problem in privacy-aware computer vision for AAL is not merely technical; it is a structural tension between the demand for rich situational information and the need to preserve older adults' autonomy and dignity. A system that continuously transmits raw video to a remote caregiver or cloud service may maximize safety but severely erode trust. Conversely, a system that avoids cameras entirely may fail to detect dangerous events or provide sufficient contextual understanding. This paper therefore approaches privacy as an architectural and governance challenge. It examines how vision pipelines can be organized to transform raw imagery into abstract, task-relevant representations at the edge, how consent mechanisms can remain meaningful over time, and how system-level accountability can be maintained. The analysis focuses on structural trade-offs rather than on isolated algorithms, and it positions fairness, robustness, and sustainability as essential dimensions of privacy-aware design.

2. Sociotechnical Tensions in Ambient Assisted Living

Video-based AAL technologies offer considerable potential for automated lifelogging, health monitoring, and emergency detection, but their deployment intersects with personal identity, bodily exposure, and the meaning of home as a private space [4]. Older adults often hold nuanced attitudes toward smart home technologies. They may recognize safety benefits while simultaneously expressing discomfort about being watched, concern about burdening family members, and uncertainty about who can access monitoring data [5]. These tensions indicate that privacy cannot be reduced to compliance with a set of data protection rules. Instead, privacy is relational and contextual; it depends on what information is collected, how it is interpreted, who receives it, and whether the older adult can influence these processes. Privacy by design frameworks have long emphasized that privacy should be embedded proactively into systems rather than added as a retrospective patch, but translating those principles into concrete vision-based AAL architectures remains difficult [6].

The lived experience of camera monitoring can be particularly intrusive because visual data reveal not only discrete events but also continuous bodily appearance, movement, clothing, visitors, and domestic routines. Even when a system is intended to detect only falls, it may capture intimate details that are not relevant to the safety goal. This overcollection creates a structural mismatch between the stated purpose and the actual data exposure. Older adults may accept camera-based systems more readily when they perceive clear benefits, when monitoring is limited to specific risk scenarios, and when they can control or observe the resulting data flows. However, these conditions are rarely stable in long-term deployments, because care needs evolve, caregivers change, and technical systems require remote maintenance. A sociotechnical perspective therefore suggests that privacy-aware vision systems must be designed for continuous negotiation among residents, informal caregivers, clinical providers, and technical operators rather than for one-time informed consent.

3. Privacy-Preserving Vision System Architecture

A privacy-aware vision architecture for AAL should be organized around the principle of data minimization at every layer. The General Data Protection Regulation establishes that personal data processing must be adequate, relevant, and limited to what is necessary for specified purposes, and this principle has direct architectural implications for camera-based monitoring [7]. Rather than transmitting raw video to a central server, a privacy-preserving system can

embed visual analysis in a local device that extracts only task-relevant features such as motion trajectories, posture states, or event labels. Differential privacy offers a mathematically grounded approach for bounding the information leaked by such features by adding calibrated noise to query outputs or model updates [8]. Although originally developed for statistical databases, differential privacy has been extended to deep learning, where it can be applied during model training to reduce the risk that learned representations memorize sensitive visual details [9]. However, differential privacy introduces a trade-off between privacy budget and model accuracy, and its guarantees are weaker when the same visual scene is processed repeatedly over long periods.

System architecture must also address the transformation of raw imagery into abstract representations that are sufficient for safety tasks but less revealing of personal identity. Common strategies include extracting skeletal poses, computing room-level occupancy maps, detecting motion regions, and generating event summaries rather than storing or transmitting full frames. A review of machine learning and privacy highlights that privacy-preserving techniques range from data anonymization and perturbation to federated learning and homomorphic encryption, each with different trust assumptions and computational costs [10]. In the AAL context, the physical device boundary is particularly important because domestic networks may be vulnerable to tampering or unauthorized access. Security mechanisms for medical cyber-physical systems emphasize authentication, integrity protection, and secure boot processes that must be integrated into edge devices used for visual monitoring [11]. The system should assume that any raw frame escaping the local device increases the risk surface, so architectural choices should favor local processing even at the cost of higher hardware requirements.

Several lessons can be drawn from activity recognition research that uses unobtrusive and anonymous sensing. For example, binary sensors combined with deep convolutional networks can recognize daily activities without requiring continuous video, demonstrating that meaningful behavioral inference can be achieved with coarse, non-visual or low-resolution inputs [12]. Similarly, behavior models built from multimodal observations can represent human routines at a sufficiently abstract level to support anomaly detection and safety alerts [13]. Privacy-aware vision architectures can adopt a similar abstraction philosophy by converting video into event histories, activity graphs, and risk scores. These representations preserve safety-relevant information while reducing the identifiability of residents and the likelihood that incidental private moments are exposed. The challenge is to ensure that the abstraction layer does not remove context needed for accurate fall detection, medication adherence monitoring, or emergency response coordination. For this reason, architecture design should be guided by explicit use cases and acceptable risk thresholds rather than by a generic desire to minimize all data.

4. Edge Processing and Distributed Intelligence

The shift from centralized cloud analytics to edge processing is one of the most consequential structural decisions in privacy-aware computer vision for AAL. Edge devices can perform feature extraction, event detection, and model inference locally, allowing only aggregated alerts or anonymized behavioral summaries to leave the home. This architectural pattern reduces the amount of personal data in transit and limits the exposure associated with cloud storage. Wearable and environmental sensing systems have already demonstrated the value of local data fusion for rehabilitation and activity monitoring, where continuous raw data streams would be impractical and sensitive [14]. Older adults' acceptance of technology is

also shaped by perceptions of control, trust, and perceived usefulness, which can improve when monitoring appears to operate within the home rather than in a distant infrastructure [15]. Yet edge processing introduces constraints related to computational capacity, energy consumption, environmental heat, and the difficulty of updating models deployed in diverse homes.

A balanced system may use a hybrid architecture in which privacy-sensitive tasks are executed on the edge device while computationally intensive but less sensitive tasks are selectively offloaded. For instance, a local model might detect that a person has fallen and immediately communicate a high-level alert, while a more detailed diagnostic analysis is performed later on anonymized features if needed. This separation of real-time safety functions from retrospective analytics is structurally important because it aligns latency, privacy, and reliability requirements. The edge layer must be capable of functioning during network outages or cloud service failures, which are realistic conditions in domestic environments. It must also support secure over-the-air updates, remote configuration, and diagnostic logging without exposing raw video. Distributed intelligence therefore requires not only model compression and hardware acceleration but also a lifecycle management framework that treats each home device as a long-lived, safety-critical component.

5. Data Governance, Consent, and Infrastructure

Privacy-aware vision systems need governance mechanisms that extend beyond the moment of installation. Older adults often hold deeper attitudes toward camera privacy than simple digital privacy preferences, with concerns shaped by social expectations, bodily exposure, and the symbolic meaning of cameras in the home [16]. Meaningful consent cannot be obtained once and then assumed indefinitely, because cognitive capacity, living arrangements, and care relationships may change. Dynamic consent models allow residents to pause monitoring, restrict the types of events that may be shared, or specify which caregivers receive alerts. The infrastructure must therefore include identity management, access control, audit logs, and transparent data flow records. These mechanisms should be designed to be usable by older adults and their caregivers without requiring deep technical expertise. Consent interfaces that rely on dense legal language or hidden settings are unlikely to produce genuine autonomy.

The governance layer must also define retention periods, deletion procedures, and conditions under which data may be shared with healthcare providers or emergency services. Visual abstractions such as skeleton trajectories or activity labels may be less sensitive than raw video, but they can still reveal behavioral patterns and vulnerabilities. Infrastructure decisions therefore depend on whether the system is provided by a health service, a municipality, a commercial vendor, or a family member. Each organizational context brings different accountability structures, funding models, and regulatory obligations. A sustainable deployment must include provisions for device maintenance, software updates, technical support, and independent privacy auditing. Without such infrastructure, privacy protections embedded in algorithms may degrade over time as devices become obsolete, models drift, or caregivers work around inconvenient controls.

6. Fairness, Robustness, and Algorithmic Accountability

Visual analytics models used in AAL are subject to the same robustness and fairness challenges as other machine learning systems. Deep learning architectures have accelerated progress in visual recognition, but their performance can vary across environments, camera placements, lighting conditions, and user populations [17]. Older adults exhibit diverse

movement patterns due to mobility aids, chronic conditions, and cultural differences in daily routines. A fall detection model trained on a narrow demographic may perform poorly for wheelchair users or people with unsteady gait, producing both safety risks and unfair visibility. Adversarial manipulation is another concern, because visual systems can be misled by small perturbations that are imperceptible to human observers [18]. While adversarial attacks may seem unlikely in domestic AAL, the increasing connectivity of home devices means that compromised or misconfigured equipment can produce erroneous alerts or leak behavioral information. Security and privacy challenges in cloud-based IoT further emphasize the need for end-to-end protection, device attestation, and data integrity checks [19].

Algorithmic accountability requires monitoring model outputs across different households and subpopulations, especially because fairness metrics for safety systems are not straightforward. A naive equal accuracy target may be inappropriate when risk distributions differ across groups; instead, fairness should be assessed in relation to the clinical and social consequences of missed events and false alarms. False alarms can undermine trust and lead older adults to disable monitoring, while missed events may delay emergency response. Accountability therefore depends on audit trails, explainable alerts, and processes for reviewing system failures. The system should record which abstract features triggered an alert and how the decision was presented to caregivers. Such records support post-incident analysis and foster institutional learning. Robustness and fairness are not purely technical properties; they are consequences of how data are collected, how models are validated, and how deployment decisions are made.

7. Policy, Standards, and Long-Term Sustainability

Older adults' willingness to adopt ambient intelligence technologies is strongly influenced by perceived usefulness, privacy risks, and the availability of social support [20]. Privacy concerns and acceptance of assistive technologies are highly contextual, varying with the type of monitoring, the recipient of data, and the level of user control [21]. Policy instruments such as data protection law, medical device regulation, and accessibility standards shape the conditions under which vision-based AAL systems can be deployed. However, standards for camera-based monitoring in domestic settings remain fragmented. Fall detection systems, for example, are often evaluated against laboratory datasets and short-term studies, but real-world deployment requires sustained performance across seasons, room layouts, and household compositions [22]. Regulatory frameworks should encourage transparent reporting of privacy properties, accuracy metrics, failure modes, and maintenance requirements.

Long-term sustainability is another structural concern. Privacy-aware architecture may require more expensive edge hardware, ongoing software updates, and specialized technical support. Funding models must account for these costs, or privacy features will be compromised under resource constraints. Interoperability standards can reduce vendor lock-in and enable local data control, but they also require coordination among device manufacturers, healthcare providers, and housing organizations. A sustainable ecosystem should include procurement rules that prioritize privacy-by-design and open auditability. It should also recognize that older adults are not a homogeneous group; individual preferences and capacities must be respected through configurable monitoring modes. Without such policy and operational support, technical privacy mechanisms may remain isolated prototypes rather than durable elements of care infrastructure.

8. Future Research Directions

Future research should investigate how privacy-aware vision systems can adapt to changing household conditions and evolving care needs over years rather than weeks. The Internet of Things has raised broader questions about social and ethical behavior in interconnected environments, including the responsibilities of device manufacturers and data processors [23]. These questions are especially urgent when monitoring is embedded in private homes and used to make inferences about vulnerable individuals. Privacy scholarship has long recognized that privacy is not a single unified value but a family of concerns involving information collection, processing, dissemination, and intrusion [24]. Applying this multidimensional perspective to camera-based AAL can help designers distinguish between different types of harm and design more targeted protections. For example, a resident may accept the collection of motion features for fall detection while strongly objecting to the storage of identifiable video or the automatic notification of family members without consent.

Promising directions include federated learning across households, which can improve model generalization without centralizing raw visual data, and privacy-preserving video representations that retain action semantics while obscuring identity. Another important direction is the development of certified vision pipelines whose privacy properties can be independently verified, perhaps through formal abstractions or standardized test suites. Co-design methods involving older adults, caregivers, disability advocates, and health professionals can help align system goals with lived experience. Longitudinal studies should evaluate not only recognition accuracy but also trust, perceived autonomy, and the stability of consent preferences. As camera-based AAL moves from experimental prototypes to routine care infrastructure, research must grapple with the organizational, regulatory, and economic conditions that determine whether privacy protections are maintained over time.

9. Conclusion

Privacy-aware computer vision for ambient assisted living requires a systemic integration of local data abstraction, edge intelligence, dynamic governance, algorithmic accountability, and sustainable operational practices. The tension between surveillance and care cannot be resolved by any single technique, because privacy risks arise from the interaction of camera placement, data flow, model outputs, caregiver access, and the changing circumstances of older adults. Architectures that process raw video close to the source and emit only task-relevant abstractions can reduce the exposure of intimate domestic life, but they also introduce reliability and fairness challenges. Effective governance must preserve meaningful consent, allow configuration and revocation, and ensure that safety benefits do not depend on continuous raw video collection. Robustness and equity require ongoing evaluation across diverse populations, while policy and standards must support long-term maintenance and independent audit. By treating privacy as a structural property of the entire sociotechnical system, researchers and practitioners can design ambient assisted living technologies that support older adults' autonomy, safety, and trust in mutually reinforcing ways.

References

1. Memon, M., Wagner, S. R., Pedersen, C. F., Beevi, F. H. A., & Hansen, F. O. (2014). Ambient assisted living healthcare frameworks, platforms, standards, and quality attributes. *Sensors*, 14(3), 4312-4341.
2. Rashidi, P., & Mihailidis, A. (2013). A survey on ambient-assisted living tools for older adults. *IEEE Journal of Biomedical and Health Informatics*, 17(3), 579-590.

3. Chen, L., Hoey, J., Nugent, C. D., Cook, D. J., & Yu, Z. (2012). Sensor-based activity recognition. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 42(6), 790-808.
4. Climent-Pérez, P., Spinsante, S., Mihailidis, A., & Florez-Revuelta, F. (2020). A review on video-based active and assisted living technologies for automated lifelogging. *Expert Systems with Applications*, 139, 112847.
5. Demiris, G., Rantz, M. J., Aud, M. A., Marek, K. D., Tyrer, H. W., Skubic, M., & Hussam, A. (2004). Older adults' attitudes towards and perceptions of smart home technologies: A pilot study. *Medical Informatics and the Internet in Medicine*, 29(2), 87-94.
6. Cavoukian, A. (2011). Privacy by design: The 7 foundational principles. Information and Privacy Commissioner of Ontario.
7. European Union. (2016). Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data (General Data Protection Regulation). *Official Journal of the European Union*, L119, 1-88.
8. Dwork, C. (2008). Differential privacy: A survey of results. In *International Conference on Theory and Applications of Models of Computation* (pp. 1-19). Springer.
9. Abadi, M., Chu, A., Goodfellow, I., McMahan, H. B., Mironov, I., Talwar, K., & Zhang, L. (2016). Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security* (pp. 308-318). ACM.
10. Liu, B., Ding, M., Shaham, S., Rahayu, W., Farokhi, F., & Lin, Z. (2021). When machine learning meets privacy: A survey and outlook. *ACM Computing Surveys*, 54(2), 31.
11. Kocabas, O., Soyata, T., & Aktas, M. K. (2016). Emerging security mechanisms for medical cyber physical systems. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 13(3), 401-416.
12. Gochoo, M., Tan, T. H., Liu, S. H., Jean, F. R., Alnajjar, F. S., & Huang, S. C. (2019). Unobtrusive activity recognition of elderly people living alone using anonymous binary sensors and DCNN. *IEEE Journal of Biomedical and Health Informatics*, 23(2), 693-702.
13. Brdiczka, O., Langet, M., Maisonnasse, J., & Crowley, J. L. (2009). Detecting human behavior models from multimodal observation in a smart home. *IEEE Transactions on Automation Science and Engineering*, 6(4), 588-597.
14. Patel, S., Park, H., Bonato, P., Chan, L., & Rodgers, M. (2012). A review of wearable sensors and systems with application in rehabilitation. *Journal of NeuroEngineering and Rehabilitation*, 9, 21.
15. Arning, K., & Ziefle, M. (2009). Different perspectives on technology acceptance: The role of technology type and age. In *Human-Computer Interaction. Ambient, Ubiquitous and Intelligent Interaction* (pp. 20-29). Springer.
16. Zhang, W., Yin, J., Chan, K. I., Sun, T., Jin, T., Jeung, J., & Gong, J. (2024). Beyond digital privacy: Uncovering deeper attitudes toward privacy in cameras among older adults. *International Journal of Human-Computer Studies*, 192, 103345.
17. LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436-444.

18. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., & Bengio, Y. (2014). Generative adversarial nets. In *Advances in Neural Information Processing Systems* (pp. 2672-2680). Curran Associates.
19. Zhou, J., Cao, Z., Dong, X., & Vasilakos, A. V. (2017). Security and privacy for cloud-based IoT: Challenges, countermeasures, and future directions. *IEEE Communications Magazine*, 55(1), 26-33.
20. van Hoof, J., Kort, H. S. M., Rutten, P. G. S., & Duijnste, M. S. H. (2011). Ageing-in-place with the use of ambient intelligence technology: Perspectives of older users. *International Journal of Medical Informatics*, 80(5), 310-331.
21. Schomakers, E. M., & Ziefle, M. (2019). Privacy concerns and acceptance of assistive technologies for older adults. In *Proceedings of the 5th International Conference on Information and Communication Technologies for Ageing Well and e-Health* (pp. 276-283). ScitePress.
22. Mubashir, M., Shao, L., & Seed, L. (2013). A survey on fall detection: Principles and approaches. *Neurocomputing*, 100, 144-152.
23. Berman, F., & Cerf, V. G. (2013). Social and ethical behavior in the Internet of Things. *Communications of the ACM*, 56(2), 32-34.
24. Solove, D. J. (2006). A taxonomy of privacy. *University of Pennsylvania Law Review*, 154(3), 477-560.