*Journal of Computers and Applications, 2(2):54–65, July-December: 2026.*
*DOI: https://doi.org/10.58613/jca226*

**Journal of Computers and Applications**
*ISSN (Online): 3139-2024*

*Journal homepage: https://www.gpub.org/journal-details.php?journal-id=44*

***Review Article***

# The Privacy Paradox of Smart Environments: Reconciling Continuous Sensing with Privacy-Preserving Digital Intelligence

**Emmanuel Toochukwu Ogbodo [1*], Damilare Timothy Ogunjobi [2], Abayomi Robert Onawole [3] and Mariam Toyin Balogun [4]**

[1] *Department of Media, School of Culture and Communication, Swansea University, Swansea, Wales, UK.*
[2] *Department of Data Science and Analytics, BI Norwegian Business School, Oslo, Norway.*
[3] *Department of Electrical and Electronics Engineering, Federal University of Technology, Akure, Nigeria.*
[4] *Department of Political Science, Clemson University, South Carolina, USA.*
[*] *Corresponding author: tooemma4all@gmail.com*

**Article Info**

***Keywords:*** *Smart environments, Continuous sensing, Privacy-preserving artificial intelligence, Ambient intelligence, Internet of Things, Privacy–utility trade-off, Federated learning, Digital governance.*

***Received:*** *28.09.2026;*
***Accepted:*** *20.10.2026;*
***Published:*** *31.10.2026*

*© 2026 by the author's. The terms and conditions of the Creative Commons Attribution (CC BY) license apply to this open access article.*

**Abstract**

Smart environments increasingly rely on continuous sensing and artificial intelligence to interpret human activity, anticipate needs, and automate decisions across homes, healthcare facilities, workplaces, and public infrastructure. Although these capabilities can improve safety, responsiveness, and efficiency, persistent data collection also enables behavioral profiling, sensitive inference, secondary data use, and surveillance of individuals who may have limited control over monitoring. This narrative review critically examines the privacy–intelligence tension by integrating technical, empirical, behavioral, and governance perspectives. It considers risks arising across the data lifecycle and evaluates the capabilities and limitations of data minimization, edge computing, federated learning, differential privacy, and cryptographic safeguards. The analysis highlights three recurring challenges: privacy guarantees that address isolated threats rather than cumulative inference, inconsistent measures of privacy and system utility, and insufficient attention to institutional power and the circumstances of affected individuals. It argues that privacy protection and intelligent functionality are not inherently incompatible, but their reconciliation depends on proportional sensing, clearly defined purposes, realistic threat models, and accountable deployment. To integrate these requirements, the review proposes a Context-Adaptive Privacy–Intelligence Assurance Framework linking necessity assessment, purpose-limited sensing, protected inference, context-sensitive action, and continuous oversight. The framework offers an agenda for future empirical evaluation rather than a validated solution, emphasizing longitudinal testing, bystander protection, and meaningful user control across diverse settings and changing technological and institutional conditions.

## 1. Introduction

Smart environments are reshaping everyday spaces. Homes adjust lighting and temperature in response to occupancy; hospitals use connected devices to monitor patients and detect deterioration; workplaces deploy sensors to manage buildings and assess activity; and cities draw on distributed data to regulate traffic, energy use, and public services [1]. What connects these applications is their ability to respond without

**How To Cite:** E.T. Ogbodo, D.T. Ogunjobi, A.R. Onawole, and M.T. Balogun. The Privacy Paradox of Smart Environments: Reconciling Continuous Sensing with Privacy-Preserving Digital Intelligence. *Journal of Computers and Applications*, 2(2):54–65, 2026. https://doi.org/10.58613/jca226

constant human input. Their appeal lies in convenience, efficiency, safety, and more timely decisions, especially where needs change rapidly [2]

These capabilities depend increasingly on continuous sensing and machine-learning systems that convert environmental signals into predictions and actions. Cameras, microphones, wearable devices, motion detectors, smart meters, and wireless signals provide different views of human activity. When combined over time, such data can reveal patterns that are difficult to observe from isolated measurements [3]. Yet the same analytical capacity that allows a system to recognize a fall, anticipate an occupant's needs, or optimize resource use can disclose routines, relationships, health-related characteristics, and periods of absence [4]. Privacy concerns therefore extend beyond the unauthorized disclosure of personal records to information that algorithms can derive from apparently mundane observations.

This creates a fundamental tension in the design of smart environments. More data may improve an application's responsiveness, but greater collection is not always necessary or proportionate to the benefit achieved. The term "privacy paradox" is commonly used to describe the gap between expressed privacy concerns and actual behavior; in continuously sensing environments, a related but distinct problem concerns the dependence of intelligent services on data that individuals may be unable to meaningfully avoid sharing [5]. Visitors, patients, employees, and household members may be monitored without comparable opportunities to understand or contest the process [6]. Treating consent as a one-time decision overlooks these unequal circumstances and the changing uses of accumulated information.

Existing safeguards offer important, but incomplete, responses. Encryption, access controls, edge processing, federated learning, and differential privacy address different sources of exposure, yet none alone determines whether sensing is justified, whether sensitive inferences should be made, or who bears responsibility for harmful uses [7]. Technical evaluations also do not always capture lived experience, institutional power, or the effects of long-term deployment [8]. Consequently, privacy protection risks being assessed as a property of an algorithm rather than of the wider system in which it operates.

This review critically examines how continuous sensing generates privacy risks, what current privacy-preserving approaches can and cannot achieve, and why their effectiveness varies across settings. By integrating technical evidence with behavioral and governance perspectives, it proposes a context-adaptive privacy–intelligence assurance framework that links the necessity of sensing to protected inference, meaningful human control, and ongoing accountability.

## 2. Review Approach and Analytical Scope

This article adopts a critical narrative review approach to examine the relationship between continuous sensing, privacy protection, and intelligent decision-making in smart environments. Rather than attempting an exhaustive assessment of published research, the review draws on relevant empirical investigations, technical studies, conceptual contributions, and policy literature to develop an integrated understanding of the subject. Emphasis is placed on evidence that clarifies how sensing technologies create privacy risks, how protective mechanisms operate in practice, and the conditions under which privacy and intelligent functionality can be reconciled.

The analysis is organized around five interconnected areas: the development of continuous sensing and ambient intelligence; privacy risks arising from data collection and algorithmic inference; the capabilities and limitations of privacy-enhancing technologies; empirical evidence concerning privacy, utility, and user trust; and the governance challenges associated with intelligent environments. Evidence from residential, healthcare, workplace, and urban settings is considered to identify both shared challenges and context-specific differences.

Rather than describing individual studies in isolation, the review compares their underlying assumptions, methodological approaches, reported findings, and limitations. Particular attention is given to differences between experimental evaluations and real-world applications, the adequacy of privacy and performance measures, and inconsistencies in reported outcomes. These comparisons inform the identification of unresolved questions and the development of a context-adaptive privacy–intelligence assurance framework. As a narrative synthesis without systematic literature retrieval, the review does not claim comprehensive coverage or formal evidence grading. Its purpose is to provide a theoretically informed, critically grounded interpretation of existing knowledge and identify priorities for future investigation.

## 3. Continuous Sensing and the Privacy Paradox

### 3.1. Evolution of Continuous Sensing and Ambient Intelligence

Conventional digital systems largely depend on deliberate human interaction, requiring users to enter commands, provide information, or initiate specific actions. Smart environments represent a shift toward systems that continuously observe their surroundings, interpret changing conditions, and respond with limited direct human involvement [9]. Through interconnected cameras, motion detectors, wearable devices, microphones, and wireless sensors, these environments generate information that machine-learning models can use to recognize activities, anticipate needs, and support automated decisions [10]. Intelligence consequently emerges not from individual devices alone but from their ability to combine observations and respond to contextual changes.

The integration of multiple sensing technologies has expanded these capabilities. For example, combining movement patterns with physiological measurements and environmental information may improve the detection of health-related changes or support more responsive assisted-living services [11]. However, multimodal sensing also creates new privacy vulnerabilities. Information that appears harmless when collected independently may reveal sensitive details about personal routines, health conditions, relationships, or behavior when analyzed collectively [12]. Continuous observation further allows such inferences to develop over time, making the distinction between necessary environmental awareness and excessive surveillance increasingly difficult to maintain.

### 3.2. Understanding the Privacy Paradox

The privacy paradox traditionally describes the apparent inconsistency between individuals' expressed concerns about personal information and their willingness to use technologies that collect it [13]. However, this relationship is more complex than a simple contradiction between attitudes and behavior. Individuals may value privacy while accepting continuous monitoring because it offers convenience, safety, accessibility, or personalized services [14]. Such decisions are influenced by trust in technology providers, awareness of data practices, perceived benefits, and the availability of alternatives. In workplaces, healthcare facilities, and shared households, participation may be

difficult to avoid, meaning that apparent acceptance does not necessarily indicate informed or voluntary agreement.

A related but distinct tension exists between intelligent functionality and privacy protection. Extensive sensing may improve contextual understanding and predictive accuracy while increasing opportunities for sensitive inference and secondary data use. Nevertheless, greater data collection does not invariably produce better intelligence. Carefully selected sensors, localized processing, and purpose-specific analytical models may preserve useful functionality while reducing unnecessary exposure. The central challenge is therefore not simply whether individuals are willing to exchange privacy for convenience, but whether intelligent systems can deliver meaningful benefits without collecting or inferring more personal information than their intended purposes require.

### 3.3. Why Conventional Privacy Protection Is Insufficient?

Conventional privacy safeguards were largely developed for systems in which data collection occurs through identifiable interactions and access can be restricted to authorized users. Their limitations become more apparent in continuously sensing environments, where information is collected repeatedly and its future uses may be difficult to anticipate. Consent mechanisms often assume that individuals understand the implications of data collection, yet complex processing practices and changing analytical capabilities make meaningful consent difficult to sustain [15]. Anonymization offers protection against direct identification but may be undermined when longitudinal observations are combined with other datasets or used to infer distinctive behavioral patterns [16]. Similarly, access controls restrict who can retrieve information without necessarily limiting the sensitive conclusions that authorized systems can generate. These limitations suggest that privacy cannot be adequately protected through isolated safeguards alone. Effective protection requires continuous assessment of data necessity, inferential risks, and accountability throughout the information lifecycle. Figure 1 illustrates how continuous sensing and data integration can simultaneously enhance intelligent functionality and increase privacy exposure, highlighting the need to balance technological benefits with proportionate data collection and privacy protection.

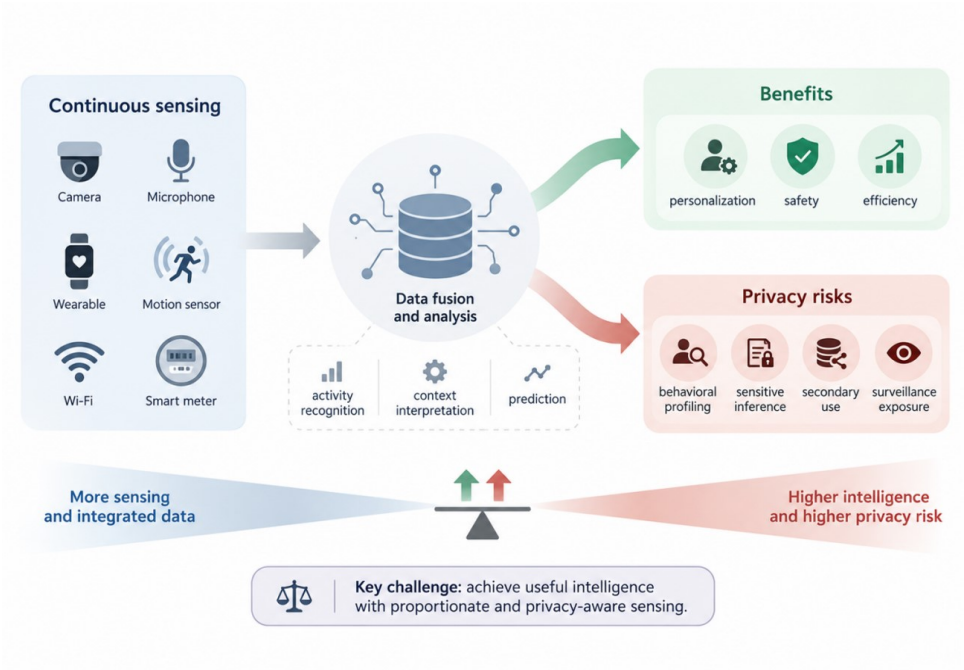


**Figure 1:** The Continuous Sensing–Intelligence–Privacy Paradox

This figure illustrates the central tension in smart environments between the functional benefits of continuous sensing and the privacy risks it creates. Diverse sensing modalities, including cameras, microphones, wearables, motion sensors, Wi-Fi, and smart meters, generate data that are integrated through fusion and analysis to support activity recognition, context interpretation, and prediction. These capabilities enable important benefits such as personalization, safety, and efficiency, but they also increase the potential for behavioral profiling, sensitive inference, secondary data use, and surveillance exposure. The figure highlights the core challenge addressed in this review: how to achieve useful digital intelligence while ensuring that sensing remains proportionate, context-sensitive, and privacy-aware.

## 4. Privacy Risks Across the Smart-Environment Data Lifecycle

### 4.1. Data Collection, Behavioral Profiling, and Inferential Privacy

Smart environments collect information that is often more revealing than its original purpose suggests. Cameras and microphones can capture identifiable faces and voices, while wearable devices and location sensors produce records of movement, physiological activity, and proximity to other people. Even systems designed to avoid direct identifiers may disclose sensitive characteristics [17]. Changes in household electricity consumption, patterns of wireless-signal disturbance, or repeated room occupancy can indicate sleeping routines, working hours, or periods when someone is away [4]. The privacy implications therefore depend not only on the type of sensor but also on the analytical questions applied to its outputs.

Artificial intelligence intensifies this concern by turning observations into classifications and predictions. Activity-recognition models

may distinguish everyday movements, but similar techniques can also be used to estimate health status, infer relationships, or monitor habits that individuals have not chosen to disclose [3]. These inferences need not be correct to cause harm: inaccurate labels may influence service provision, employment judgments, or perceptions of personal risk. Meanwhile, accurate inferences can reveal highly private information even when the underlying measurements are not conventionally categorized as sensitive [18].

Combining data introduces an additional problem. Records that provide little identifying information individually may become distinctive when linked across devices, places, and time. Re-identification is particularly difficult to prevent when movement traces or behavioral routines remain sufficiently unique [19]. Removing names or replacing device identifiers does not necessarily eliminate these linkages [20]. Privacy assessments focused only on whether raw data contain explicit personal identifiers may therefore overlook the more consequential question of what a system can learn, infer, and act upon. Effective risk assessment must consider plausible inference pathways, who can exploit them, and the consequences of both correct and erroneous predictions.

### 4.2. Persistent Surveillance, Secondary Data Use, and Third-Party Exposure

The persistence of sensor data changes privacy exposure from a single collection event into an accumulating condition. Longitudinal records allow systems to recognize changes and improve predictions, but they also preserve detailed accounts of everyday life. When previously separate sources are linked, the resulting profiles may disclose patterns that were invisible within any one dataset. Information originally collected for safety, comfort, or maintenance can subsequently be repurposed for productivity monitoring, targeted advertising, insurance decisions, or other uses outside the expectations of affected individuals [21]. This function creep may occur without a new sensor being installed or an obvious change in how the environment operates [22].

Exposure also extends beyond the person who purchases or formally uses a device. Visitors may enter monitored rooms without noticing active sensors; household members may disagree about acceptable recording; patients may depend on monitoring for care; and employees may face practical penalties for declining workplace surveillance [6]. These situations complicate appeals to notice and consent, especially when monitoring is built into essential services or shared spaces [23]. The ethical concern is not simply whether information is protected from outsiders, but whether people subject to sensing have a meaningful opportunity to question its necessity, restrict its use, or seek redress.

Protecting privacy across this lifecycle requires more than secure storage and controlled access. Collection should be limited to defined purposes, processing should be evaluated for unintended inferences, and sharing arrangements should prohibit incompatible secondary uses. Retention periods, deletion practices, auditability, and accountability for model outputs require comparable scrutiny. Deleting a raw record does not automatically erase information incorporated into derived profiles or trained models, so disposal policies must address downstream artefacts as well. The appropriate standard is not the elimination of every risk, which may be unrealistic, but defensible limits on what is observed, retained, inferred, and acted upon.

## 5. Privacy-Preserving Digital Intelligence: Capabilities and Limitations

### 5.1. Data Minimization, Edge Computing, and Local Intelligence

Data minimization offers a fundamental approach to privacy protection by restricting information collection to what is necessary for a defined purpose [24]. In smart environments, this may involve reducing sensor resolution, limiting monitoring periods, extracting only relevant features, or replacing identifiable recordings with less revealing measurements [25]. Such measures can reduce exposure before information enters storage or analytical systems. However, their effectiveness depends on whether the remaining data are genuinely sufficient for the intended task without enabling unintended inferences.

Edge computing extends this principle by processing information near its point of collection rather than routinely transmitting raw data to centralized servers [26]. For example, a home-monitoring system may identify unusual movement patterns locally and transmit only an alert when intervention is required. This architecture can reduce network dependence, communication delays, and exposure during data transmission [27]. Nevertheless, local processing does not automatically guarantee privacy. Compromised devices, insecure software, retained sensor records, and revealing model outputs may still expose personal information.

Practical implementation also involves trade-offs. Resource-constrained devices may struggle to support computationally demanding models, while simplified algorithms or reduced sensing resolution can compromise accuracy in complex environments. Local processing may also limit opportunities to combine information across devices when broader contextual interpretation is necessary. Importantly, even minimal data can disclose sensitive characteristics when analyzed repeatedly or combined with other observations. Consequently, data minimization and edge intelligence are valuable preventive measures, but their effectiveness must be evaluated against specific inference risks, operational requirements, and the consequences of reduced analytical performance.

### 5.2. Federated Learning and Decentralized AI

Federated learning offers an alternative to conventional centralized machine learning by allowing multiple devices or institutions to train a shared model without transferring their raw datasets to a common repository [28]. In smart environments, this approach can support collaborative intelligence across households, healthcare facilities, and connected devices while reducing the need to concentrate sensitive information [7]. Its principal advantage is that data remain under local control during training, although this architectural separation should not be mistaken for a guarantee of privacy.

Experimental research has demonstrated that information about training data can sometimes be reconstructed or inferred from exchanged model updates [29]. Gradient inversion attacks may reveal characteristics of local training examples, while membership inference attacks can indicate whether particular records contributed to model development [30]. These vulnerabilities depend on factors such as model architecture, adversarial access, training procedures, and the safeguards employed. Consequently, the privacy benefits of federated learning vary considerably across implementations.

Practical deployment introduces further difficulties. Devices differ in processing capacity, connectivity, and available training data, creating computational and communication burdens. Differences in local data distributions may also reduce model performance or produce

uneven accuracy across participating environments [31]. Repeated model updates can increase communication costs, particularly in resource-constrained sensor networks. Although techniques such as secure aggregation can reduce certain disclosure risks, they do not address every vulnerability [32]. Federated learning therefore represents a useful architectural strategy rather than a complete privacy solution, requiring complementary safeguards and evaluation under realistic operating conditions.

### 5.3. Differential Privacy and Cryptographic Safeguards

Differential privacy and cryptographic techniques provide complementary approaches to protecting sensitive information during intelligent data processing. Differential privacy introduces carefully calibrated randomness into analytical outputs or learning procedures, limiting how much the participation of an individual record can influence disclosed results. Unlike conventional anonymization, it offers a formal privacy guarantee under specified assumptions [33]. However, the strength of this protection depends on the privacy budget, sensitivity of the analysis, and cumulative effects of repeated queries. Stronger privacy guarantees may reduce statistical accuracy or predictive performance, particularly when datasets are small or contain underrepresented populations [34]

**Table 1:** Comparative Evaluation of Privacy-Preserving Technologies in Smart Environments

| **Technology** | **Privacy mechanism** | **Application** | **Strengths** | **Limitations** | **Residual privacy risks** |
|---|---|---|---|---|---|
| Data minimization | Restricts data collection to information necessary for specific purposes | Smart homes, healthcare , monitoring occupancy detection | Reduces unnecessary collection, storage, and exposure | May restrict analytical detail and predictive accuracy | Sensitive inferences remain possible from limited data [24] |
| Edge computing and local intelligence | Processes sensor data locally rather than routinely transmitting raw information | Wearables, assisted living, intelligent buildings | Reduces raw-data transmission, network dependence, and latency | Limited computational resources, maintenance challenges, device heterogeneity | Device compromise, local data leakage, revealing outputs [26] |
| Federated learning | Trains shared models using decentralized datasets and exchanged model updates | Distributed healthcare, smart homes, IoT networks | Reduces centralized data collection; supports collaborative learning | Communication overhead, heterogeneous data, uneven model performance | Gradient inversion, membership inference, model-update leakage [7] |
| Differential privacy | Introduces calibrated randomness to limit disclosure about individual records | Behavioral analytics, statistical reporting, collaborative AI | Offers mathematically defined privacy guarantees under specified assumptions | Privacy-utility trade-offs, privacy-budget management, accuracy concerns | Information leakage permitted within the chosen privacy guarantee [33] |
| Secure aggregation | Combines participants' contributions without revealing individual updates to the aggregator | Federated learning, distributed sensor networks | Protects individual contributions during aggregation | Protocol complexity, communication costs, participant dropout | Leakage through aggregated outputs or compromised endpoints [32] |
| Homomorphic encryption | Enables supported computations directly on encrypted data | Encrypted cloud analytics, remote health monitoring | Protects data during outsourced computation | High computational costs, implementation complexity, scalability challenges | Output inference, endpoint compromise, unprotected metadata [35] |
| Secure multiparty computation | Enables joint computation without directly revealing each participant's private inputs | Cross-institutional analytics, collaborative healthcare | Supports confidential computation across multiple parties | Communication demands, protocol complexity, deployment costs | Output leakage, collusion beyond assumed threat models, implementation vulnerabilities [36] |

Note: IoT = Internet of Things. The comparative strengths and limitations are context-dependent. Privacy guarantees vary according to implementation, adversarial assumptions, system architecture, and operational conditions.

Cryptographic approaches address different vulnerabilities. Secure aggregation allows participating devices to contribute model updates

without exposing individual contributions to the coordinating server. Homomorphic encryption enables selected computations on encrypted information, while secure multiparty computation permits several parties to jointly evaluate functions without directly revealing their private inputs [35]. These techniques can reduce exposure during collaborative processing but introduce practical challenges involving computational overhead, communication demands, implementation complexity, and scalability [36]. Furthermore, protecting intermediate information does not necessarily prevent sensitive conclusions from being inferred through legitimate model outputs.

The broader evidence suggests that privacy-preserving intelligence cannot be achieved simply by selecting a technically sophisticated safeguard. Different mechanisms protect against different threats, and their effectiveness depends on system architecture, adversarial capabilities, deployment conditions, and acceptable performance costs. Combining data minimization, local processing, decentralized learning, and formal privacy protections may provide stronger safeguards, but integration also introduces operational complexity. Privacy protection must therefore be evaluated across the entire sensing and analytical process rather than inferred from the presence of individual technologies. The principal privacy-preserving approaches differ in the information they protect, the threats they address, and the practical constraints they introduce. Table 1 compares these technologies in terms of their operating mechanisms, applications, advantages, limitations, and residual privacy risks, highlighting why no single approach provides comprehensive protection in continuously sensing environments.

## 6. Reconciling Privacy, Utility, and Trust: Critical Evidence Synthesis

### 6.1. Empirical Evidence on Privacy–Utility Trade-offs

Empirical research increasingly challenges the assumption that stronger privacy protection necessarily requires substantial sacrifices in intelligent-system performance. Experimental evaluations of federated learning for human activity recognition demonstrate that decentralized training can support accurate classification without transferring raw sensor records to a central repository [31]. However, performance varies with model architecture, aggregation strategy, and computational resources. Similarly, local processing and selective data collection can reduce information exposure while preserving useful functionality for defined tasks [37]. These findings suggest that privacy and utility are not inherently opposing objectives; their relationship depends on system design, the information required for prediction, and how performance is measured.

Nevertheless, preserving predictive accuracy does not establish that a system adequately protects privacy. Some technical evaluations treat the absence of centralized raw-data storage as evidence of privacy preservation without directly examining information leakage through model updates or outputs. Others report performance under controlled conditions that may not reflect the variability of occupied homes, healthcare facilities, or workplaces. Consequently, favorable experimental results should be interpreted as evidence of technical feasibility rather than comprehensive privacy assurance.

Research involving smart-environment users reveals another important dimension. Studies examining household monitoring, device preferences, and bystander experiences indicate that willingness to share information depends on its purpose, the intended recipient, perceived benefits, and opportunities for control [38]. Participants may accept limited activity monitoring for safety while expressing stronger reservations about audio recording or unrestricted secondary use. Preferences can also differ between device owners and visitors, particularly when those exposed to monitoring have little influence over device settings [39].

Taken together, these findings suggest that trust cannot be inferred solely from technology adoption or reported satisfaction. It depends partly on whether individuals understand what information is collected, consider its use legitimate, and possess meaningful control over subsequent processing. The most promising approaches therefore combine technical safeguards with clearly defined purposes and context-sensitive participation mechanisms.

### 6.2. Explaining Contradictions and Methodological Weaknesses

Differences in reported privacy–utility outcomes partly reflect variations in what researchers define and measure as privacy [40]. Technical studies commonly evaluate classification accuracy, processing time, communication costs, or formal privacy guarantees, whereas behavioral investigations assess perceived risk, trust, willingness to disclose information, and acceptance of monitoring [5]. These outcomes address different questions and cannot be treated as interchangeable. A system may achieve strong predictive performance under a specified privacy mechanism while remaining unacceptable to users because of its surveillance capabilities or limited transparency.

Methodological differences further complicate comparisons. Laboratory experiments frequently rely on benchmark datasets, simulated participants, or controlled sensing conditions that simplify the complexity of everyday environments. Such designs allow reproducible comparisons but may underestimate challenges associated with irregular device connectivity, changing household routines, multiple occupants, and diverse user characteristics. Small or homogeneous samples also restrict generalizability, particularly when privacy expectations differ across cultural, institutional, and socioeconomic contexts. Conversely, surveys and hypothetical scenarios provide insight into stated preferences but cannot reliably establish how individuals will respond to continuous monitoring over extended periods.

The evaluation of privacy-preserving technologies presents additional weaknesses. Studies may compare predictive accuracy while applying different threat models, privacy budgets, or assumptions about adversarial access [41]. Some examine resistance to specific attacks without establishing protection against other plausible forms of inference [29]. Short observation periods may also overlook cumulative privacy exposure, changes in user behavior, and the consequences of linking information across devices. Accordingly, apparently contradictory findings may reflect differences in evaluation conditions rather than genuine disagreement about a technology's effectiveness.

These limitations expose a broader gap between technical feasibility and demonstrated protection in operational settings. Future evaluations must examine privacy leakage, predictive utility, computational demands, and user experience together, preferably through longitudinal deployments involving realistic threats and diverse participants. Standardized reporting of privacy assumptions and performance measures would improve comparisons without obscuring contextual differences. Most importantly, claims of privacy preservation should be supported by evidence about what information remains inferable, who can access it, and how affected individuals experience the system. Such evidence is essential for determining whether privacy-enhancing approaches deliver meaningful protection beyond controlled experimental environments. The empirical literature provides important evidence that privacy protection and intelligent-system performance do not always require an absolute trade-off. However, the outcomes of privacy-enhancing interventions vary according to sensing modality, analytical

methods, deployment conditions, and the measures used to assess privacy and utility. Table 2 synthesizes selected empirical investigations, highlighting their principal findings and methodological limitations.

**Table 2:** Empirical Evidence on Privacy–Utility Trade-offs in Smart Environments

| Study | Setting | Methodology | Privacy intervention | Principal finding | Evidence limitation |
|---|---|---|---|---|---|
| Study Malekzadeh et al. [25] | Setting Wearable sensing and activity recognition | Methodology Experimental evaluation using sensor-based activity and gesture recognition datasets | Privacy intervention Autoencoder-based sensor-data transformations | Principal finding Reduced sensitive inferences and re-identification risks while maintaining activity-recognition performance with relatively small accuracy losses | Evidence limitation Benchmark evaluations do not establish long-term effectiveness in deployed environments [25] |
| Ek et al. [31] | Smartphone-based human activity recognition | Comparative evaluation of federated-learning algorithms | Decentralized model training | Demonstrated the feasibility of federated approaches for activity recognition, while identifying performance and implementation trade-offs | Results depend on evaluated algorithms, data distributions, and experimental conditions [31] |
| Alshehri et al. [39] | Smart homes and bystander privacy | Interactive negotiation study involving 460 participants | Negotiation of device data-collection, storage, and sharing preferences | Owners and bystanders demonstrated different willingness to compromise, revealing tensions between device utility and privacy preferences | Simulated negotiations with digital agents may not reflect actual household behavior [39] |
| Shi et al. [17] | Wi-Fi sensing and behavioral inference | Deep-learning experiments using two public Wi-Fi sensing datasets | Privacy-leakage assessment; no protective intervention tested | Demonstrated that Wi-Fi sensing signals could reveal private attributes, including gender, height, and weight | Dataset-specific results do not establish equivalent inference accuracy across real-world settings [17] |
| Stirapongsasuti et al. [37] | Smart-home activity monitoring | Experimental comparison using two smart-home datasets | Rényi differential privacy and feature-merging anonymization | Reduced person-identification accuracy while retaining relatively high activity-recognition accuracy | Evaluation relied on benchmark datasets and selected privacy and utility measures [37] |
| Cardell-Oliver and Carter-Turner [41] | Smart water-meter monitoring | Experimental evaluation using real-world and synthetic water-consumption datasets | Activity-aware data protection and temporal-resolution adjustments | Privacy-utility outcomes varied substantially with activity type, sampling resolution, and population size | Findings are specific to water-meter data and the evaluated disclosure scenarios [41] |
| Frik et al. [38] | Smart homes, shared housing, and short-term accommodation | Contextual vignette survey involving 761 US participants | Evaluation of privacy preferences across different information-sharing scenarios | Data recipients, sharing purposes, and information types significantly influenced privacy judgments | Hypothetical scenarios and US-based sampling limit conclusions about actual behavior and wider applicability [38] |

**Note:** The table includes both evaluations of technical privacy-preserving mechanisms and empirical investigations of privacy risks or user preferences. These study types provide complementary evidence but should not be interpreted as measuring the same outcomes. The findings reflect the specific datasets, participants, methods, and assumptions used in each investigation.

## 7. Contextual Differences and Unequal Privacy Consequences

Privacy risks in smart environments cannot be assessed independently of where sensing occurs, who controls it, and what individuals stand to lose by refusing participation. In private homes, connected devices may support convenience, security, and independent living, yet decisions about installation and data sharing are rarely made by everyone exposed to them [15]. A resident who values voice-controlled assistance may have different expectations from a visiting relative or domestic worker [38]. Device-level permission settings offer limited protection when the person granting permission is not the only person being monitored.

Healthcare facilities introduce different responsibilities. Continuous monitoring can help clinicians detect deterioration and respond promptly [11], but patients may have little practical opportunity to decline surveillance without affecting their care. Safeguards must therefore address clinical necessity, access to sensitive inferences, retention, and accountability for secondary use, rather than relying principally on nominal consent. In assisted-living settings, movement and activity sensing may promote safety and independence while also exposing intimate routines [42]. An approach that maximizes detection sensitivity may be unacceptable if it results in constant observation or removes residents' ability to determine when privacy matters most.

Workplaces and smart cities further complicate the meaning of choice. Employers may justify occupancy, location, or productivity monitoring as necessary for operational efficiency, although employees can face unequal bargaining power and uncertainty about how data influence managerial decisions [21]. In urban settings, cameras and networked sensors can collect information about people who neither sought a service nor received meaningful notice [1]. Public benefits do not eliminate obligations to limit collection, restrict repurposing, and scrutinize discriminatory or disproportionate impacts.

Privacy consequences also differ across populations, without following simple demographic categories. Some older adults may prioritize emergency detection, whereas others may reject monitoring that compromises independence. Individuals with disabilities may benefit from adaptive sensing while facing additional risks if systems misinterpret their behavior. Children often have limited influence over decisions made by caregivers or institutions, and bystanders may be captured without any direct relationship with the system operator. These concerns intersect with income, digital literacy, institutional dependence, and access to effective remedies. They should not be taken to imply that members of any group share a uniform preference or level of vulnerability.

Universal privacy settings cannot adequately respond to these competing interests. Effective design requires a context-specific assessment of necessity, authority, exposure, and likely harm, together with options for proportional sensing, understandable controls, independent oversight, and contesting consequential inferences. The relevant question is not simply whether a technology can protect data, but whether its deployment protects the people whose lives it renders observable.

## 8. Toward Context-Adaptive Privacy-Preserving Digital Intelligence

### 8.1. Conceptual Foundations for Integrating Privacy and Intelligence

The preceding analysis demonstrates that privacy protection in smart environments remains fragmented across technological, behavioral, and institutional dimensions. Existing approaches primarily address particular vulnerabilities, such as unauthorized data access, centralized storage, or information leakage during computation. Although these safeguards can reduce specific risks, they rarely establish whether continuous monitoring is necessary, whether algorithmic inferences are appropriate, or whether affected individuals possess meaningful control over their information. This fragmentation limits the effectiveness of privacy protection when intelligent systems operate across interconnected devices, institutions, and social contexts.

A more comprehensive approach must recognize that privacy and intelligent functionality are shaped by decisions throughout the system lifecycle. The necessity of collecting particular information should be assessed before sensing begins, while analytical procedures should be evaluated for both intended functionality and unintended disclosure. Privacy protection must also account for contextual differences. Monitoring practices considered proportionate in emergency healthcare may be inappropriate in domestic or workplace settings, even when similar technologies are involved. Consequently, technical performance alone cannot determine whether an intelligent system is responsibly designed.

These considerations support an integrated model in which data necessity, inferential risk, contextual sensitivity, individual agency, and institutional accountability operate as interconnected requirements. Rather than treating privacy as a constraint introduced after system development, this approach incorporates protection into decisions about what information is collected, how intelligence is generated, and which actions are permitted. It also recognizes the need for ongoing reassessment as technologies, purposes, and user circumstances change. To translate these principles into a coherent approach, Figure 2 presents the proposed Context-Adaptive Privacy–Intelligence Assurance Framework. The framework integrates five interconnected stages of privacy protection and intelligent-system operation, supported by continuous feedback, cross-cutting governance principles, and measurable evaluation criteria.

The framework illustrates a five-stage approach to integrating privacy protection into the design and operation of intelligent smart environments. The stages encompass context and necessity assessment, purpose-limited sensing and data minimization, privacy-enhanced processing and inference, context-sensitive access and intelligent action, and audit, accountability, and adaptive safeguards. Data necessity, individual agency, contextual sensitivity, and institutional accountability operate as cross-cutting principles. A continuous feedback loop supports reassessment and adaptation, while evaluation criteria address privacy leakage, system utility, fairness, computational feasibility, contextual appropriateness, and user control. The framework is conceptual and requires empirical validation across different smart-environment settings.

## 9. Governance Implications and Future Research Priorities

### 9.1. Implications for Responsible Smart-Environment Governance

Privacy-preserving technologies cannot substitute for effective governance of how intelligent systems are designed, deployed, and used. Privacy-by-design provides a starting point by requiring organizations to consider data necessity, proportionality, and foreseeable harms

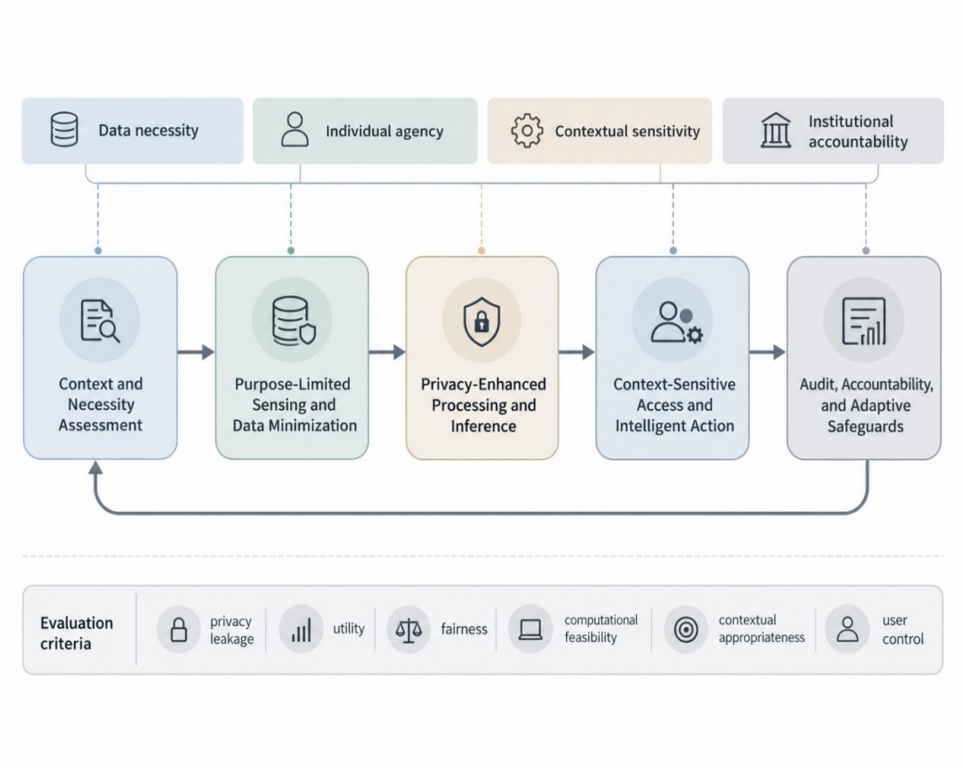


**Figure 2:** Context-Adaptive Privacy–Intelligence Assurance Framework

before monitoring begins [8]. Data protection impact assessments can further identify risks associated with continuous sensing, automated inference, and secondary information use [43]. However, these processes offer limited protection when treated as administrative requirements rather than mechanisms for influencing system design and operational decisions.

Responsible governance requires transparent explanations of sensing purposes, clearly assigned institutional responsibilities, enforceable restrictions on data use, and accessible procedures for challenging consequential decisions. Where applicable, individuals should be able to exercise rights concerning access, correction, objection, and deletion. These protections are particularly important in healthcare and workplaces, where dependence on services or unequal bargaining power may undermine meaningful consent. Independent oversight and periodic reassessment are therefore necessary to ensure that technological safeguards remain effective as institutional practices and analytical capabilities evolve.

### 9.2. Priority Directions for Future Research

The most pressing research challenge is determining whether privacy-preserving approaches remain effective beyond controlled experimental conditions. Longitudinal studies in occupied homes, healthcare facilities, and workplaces are needed to examine changing privacy risks, system performance, and user experiences over time. Such investigations should evaluate not only predictive accuracy but also information leakage, computational demands, and unintended consequences.

A second priority concerns the absence of sufficiently comparable evaluation methods. Standardized privacy–utility benchmarks should specify adversarial assumptions, data characteristics, privacy parameters, and performance outcomes while remaining sensitive to contextual differences. Research must also investigate how multimodal sensing enables sensitive inferences that cannot be anticipated from individual data sources.

Finally, greater attention is needed to bystander privacy, contextual consent, and unequal access to protective technologies. Participatory studies involving diverse users and non-users could clarify how privacy preferences differ under conditions of limited choice. Field experiments and independently audited deployments should test whether adaptive safeguards meaningfully reduce exposure without compromising essential functionality. Together, these approaches would provide the empirical foundation needed to assess and refine the proposed framework.

## 10. Conclusion

Continuous sensing has expanded the capacity of smart environments to anticipate needs, support timely decisions, and automate services. Yet the same data streams can expose personal routines, enable sensitive inferences, and extend surveillance beyond those who knowingly participate. This review demonstrates that privacy and intelligent functionality are not necessarily incompatible; their relationship depends on system design, contextual requirements, threat models, and governance arrangements. Data minimization, edge computing, federated learning, differential privacy, and cryptographic safeguards can reduce specific vulnerabilities, but none independently resolves concerns about proportionality, secondary data use, or institutional accountability.

The proposed context-adaptive privacy–intelligence assurance framework integrates these considerations across sensing, processing, access, and continuous oversight. Its practical value requires empirical validation through longitudinal deployments, comparable privacy–utility assessments, and evaluations involving diverse affected populations. Ultimately, progress in smart environments should be judged not only

by the sophistication of their sensing and predictive capabilities but also by their ability to protect personal privacy, preserve individual autonomy, and ensure accountability for how digital intelligence is developed and used.

### Article Information

**Ethics Approval:** Not applicable.

**Consent to Participate:** This study is a narrative review of previously published literature and does not involve human participants, animal subjects, or the collection of primary data.

**Consent for Publication:** Not applicable.

**Data Availability:** No new datasets were generated or analyzed during this review. The discussion is based on previously published literature.

**Conflict of Interest:** The authors declare that they have no competing interests.

**Funding:** No external funding was received for this study.

**Authors' Contributions:** All authors contributed to the conceptualization, development, critical review, and revision of the manuscript. All authors read and approved the final version and agreed to be accountable for its content.

**Acknowledgments:** Not applicable.

**AI Use Disclosure:** Generative artificial intelligence was used to assist with manuscript drafting and language refinement. The authors are responsible for reviewing the scientific content, verifying the accuracy of references and supporting evidence, and approving the final manuscript.

## References

[1] R. Kitchin. The ethics of smart cities and urban science. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 374(2083), 2016. URL https://doi.org/10.1098/rsta.2016.0115. Article 20160115.

[2] D. J. Cook, J. C. Augusto, and V. R. Jakkula. Ambient intelligence: Technologies, applications, and opportunities. *Pervasive and Mobile Computing*, 5(4):277–298, 2009. URL https://doi.org/10.1016/j.pmcj.2009.04.001.

[3] J. L. Kröger, P. Raschke, and T. R. Bhuiyan. Privacy implications of accelerometer data: A review of possible inferences. In *Proceedings of the 3rd International Conference on Cryptography, Security and Privacy*, pages 81–87. ACM, 2019. URL https://doi.org/10.1145/3309074.3309076.

[4] A. Molina-Markham, P. Shenoy, K. Fu, E. Cecchet, and D. Irwin. Private memoirs of a smart meter. In *Proceedings of the 2nd ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Building*, pages 61–66. ACM, 2010. URL https://doi.org/10.1145/1878431.1878446.

[5] S. Barth and M. D. T. de Jong. The privacy paradox – Investigating discrepancies between expressed privacy concerns and actual online behavior – A systematic literature review. *Telematics and Informatics*, 34(7):1038–1058, 2017. URL https://doi.org/10.1016/j.tele.2017.04.013.

[6] Y. Yao, J. R. Basdeo, O. R. Mcdonough, and Y. Wang. Privacy perceptions and designs of bystanders in smart homes. In *Proceedings of the ACM on Human-Computer Interaction*, volume 3, pages 1 – 24, 2019. URL https://doi.org/10.1145/3359161.

[7] T. Li, A. K. Sahu, A. Talwalkar, and V. Smith. Federated learning: Challenges, methods, and future directions. *IEEE Signal Processing Magazine*, 37(3):50–60, 2020. URL https://doi.org/10.1109/MSP.2020.2975749.

[8] S. Spiekermann. The challenges of privacy by design. *Communications of the ACM*, 55(7):38–40, 2012. URL https://doi.org/10.1145/2209249.2209263.

[9] L. Atzori, A. Iera, and G. Morabito. The Internet of Things: A survey. *Computer Networks*, 54(15):2787–2805, 2010. URL https://doi.org/10.1016/j.comnet.2010.05.010.

[10] L. Chen, J. Hoey, C. D. Nugent, D. J. Cook, and . Z. Yu. Sensor-based activity recognition. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 42(6):790–808, 2012. URL https://doi.org/10.1109/TSMCC.2012.2198883.

[11] P. Rashidi and A. Mihailidis. A survey on ambient-assisted living tools for older adults. *IEEE Journal of Biomedical and Health Informatics*, 17(3):579–590, 2013. URL https://doi.org/10.1109/JBHI.2012.2234129.

[12] J. H. Ziegeldorf, O. Garcia Morchon, and K. Wehrle. Privacy in the Internet of Things: Threats and challenges. *Security and Communication Networks*, 7(12):2728–2742, 2014. URL https://doi.org/10.1002/sec.795.

[13] S. Kokolakis. Privacy attitudes and privacy behaviour: A review of current research on the privacy paradox phenomenon. *Computers & Security*, 64:122–134, 2017. URL https://doi.org/10.1016/j.cose.2015.07.002.

[14] J. Lau, B. Zimmerman, and F. Schaub. Alexa, are you listening? Privacy perceptions, concerns and privacy-seeking behaviors with smart speakers. In *Proceedings of the ACM on Human-Computer Interaction*, volume 2, pages 1 – 31, 2018. URL https://doi.org/10.1145/3274371. Article 102.

[15] S. Zheng, N. Apthorpe, M. Chetty, and N. Feamster. User perceptions of smart home IoT privacy. In *Proceedings of the ACM on Human-Computer Interaction*, volume 2, pages 1 – 20, 2018. URL https://doi.org/10.1145/3274469. Article 200.

[16] L. Rocher, J. M. Hendrickx, and Y.-A. de Montjoye. Estimating the success of re-identifications in incomplete datasets using generative models. *Nature Communications*, 10, 2019. URL https://doi.org/10.1038/s41467-019-10933-3. Article 3069.

[17] Y. Shi, X. Zhang, L. Fu, and H. Zhang. An investigation of the private-attribute leakage in WiFi sensing. *High-Confidence Computing*, 4(4), 2024. URL https://doi.org/10.1016/j.hcc.2024.100209. Article 100209.

[18] M. Kosinski, D. Stillwell, and T. Graepel. Private traits and attributes are predictable from digital records of human behavior. *Proceedings of the National Academy of Sciences*, 110(15):5802–5805, 2013. URL https://doi.org/10.1073/pnas.1218772110.

[19] Y.-A. de Montjoye, C. A. Hidalgo, M. Verleysen, and V. D. Blondel. Unique in the crowd: The privacy bounds of human mobility. *Scientific Reports*, 3, 2013. URL https://doi.org/10.1038/srep01376. Article 1376.

[20] A. Narayanan and V. Shmatikov. Robust de-anonymization of large sparse datasets. In *2008 IEEE Symposium on Security and Privacy*, pages 111–125. IEEE, 2008. URL https://doi.org/10.1109/SP.2008.33.

[21] K. Ball. Workplace surveillance: An overview. *Labor History*, 51(1):87–106, 2010. URL https://doi.org/10.1080/00236561003654776.

[22] B.-J. Koops. The concept of function creep. *Law, Innovation and Technology*, 13(1):29–56, 2021. URL https://doi.org/10.1080/17579961.2021.1898299.

[23] C. Geeng and F. Roesner. Who's in control? Interactions in multi-user smart homes. In *In Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, pages 1–13. ACM, 2019. URL https://doi.org/10.1145/3290605.3300498.

[24] A. J. Biega, P. Potash, H. Daumé, III, F. Diaz, and M. Finck. Operationalizing the legal principle of data minimization for personalization. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 399–408. ACM, 2020. URL https://doi.org/10.1145/3397271.3401034.

[25] M. Malekzadeh, R. G. Clegg, A. Cavallaro, and H. Haddadi. Privacy and utility preserving sensor-data transformations. *Pervasive and Mobile Computing*, 63, 2020. URL https://doi.org/10.1016/j.pmcj.2020.101132. Article 101132.

[26] W. Shi, J. Cao, Q. Zhang, Y. Li, and L. Xu. Edge computing: Vision and challenges. *IEEE Internet of Things Journal*, 3(5):637–646, 2016. URL https://doi.org/10.1109/JIOT.2016.2579198.

[27] M. Satyanarayanan. The emergence of edge computing. *Computer*, 50(1):30–39, 2017. URL https://doi.org/10.1109/MC.2017.9.

[28] Q. Yang, Y. Liu, T. Chen, and Y. Tong. Federated machine learning: Concept and applications. *ACM Transactions on Intelligent Systems and Technology*, 10(2):1 – 19, 2019. URL https://doi.org/10.1145/3298981. Article 12.

[29] M. Nasr, R. Shokri, and A. Houmansadr. Comprehensive privacy analysis of deep learning: Passive and active white-box inference attacks against centralized and federated learning. In *2019 IEEE Symposium on Security and Privacy*, pages 739–753. IEEE, 2019. URL https://doi.org/10.1109/SP.2019.00065.

[30] R. Shokri, M. Stronati, C. Song, and V. Shmatikov. Membership inference attacks against machine learning models. In *2017 IEEE Symposium on Security and Privacy*, pages 3–18. IEEE, 2017. URL https://doi.org/10.1109/SP.2017.41.

[31] S. Ek, F. Portet, P. Lalanda, and G. Vega. Evaluation and comparison of federated learning algorithms for human activity recognition on smartphones. *Pervasive and Mobile Computing*, 87, 2022. URL https://doi.org/10.1016/j.pmcj.2022.101714.

[32] K. Bonawitz, V. Ivanov, B. Kreuter, A. Marcedone, H. B. McMahan, S. Patel, D. Ramage, A. Segal, and K. Seth. Practical secure aggregation for privacy-preserving machine learning. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, pages 1175–1191. ACM, 2017. URL https://doi.org/10.1145/3133956.3133982.

[33] C. Dwork and A. Roth. The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3-4):211–407, 2014. URL https://doi.org/10.1561/0400000042.

[34] M. Abadi, A. Chu, I. Goodfellow, H. B. McMahan, I. Mironov, K. Talwar, and L. Zhang. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, pages 308–318. ACM, 2016. URL https://doi.org/10.1145/2976749.2978318.

[35] A. Acar, H. Aksu, A. S. Uluagac, and M. Conti. A survey on homomorphic encryption schemes: Theory and implementation. *ACM Computing Surveys*, 51(4):1 – 35, 2018. URL https://doi.org/10.1145/3214303.

[36] Y. Lindell. Secure multiparty computation. *Communications of the ACM*, 64(1):86–96, 2021. URL https://doi.org/10.1145/3387108.

[37] S. Stirapongsasuti, F. J. Tiausas, Y. Nakamura, and K. Yasumoto. Preserving data utility in differentially private smart home data. *IEEE Access*, 12:56571–56581, 2024. URL https://doi.org/10.1109/ACCESS.2024.3390039.

[38] A. Frik, X. Zhan, N. Abdi, and J. Bernd. Who cares? Contextual privacy judgments from owner and bystander perspectives in different smart home situations. *Proceedings on Privacy Enhancing Technologies*, 2025(3):106–134, 2025. URL https://doi.org/10.56553/popets-2025-0091.

[39] A. Alshehri, E. Pahk, J. Spielman, J. T. Parker, B. Gilbert, and C. Yue. Exploring the negotiation behaviors of owners and bystanders over data practices of smart home devices. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, pages 1–27. ACM, 2023. URL https://doi.org/10.1145/3544548.3581360.

[40] I. Wagner and D. Eckhoff. Technical privacy metrics: A systematic survey. *ACM Computing Surveys*, 51(3):1 – 38, 2018. URL https://doi.org/10.1145/3168389. Article 57.

[41] R. Cardell-Oliver and H. Carter-Turner. Activity-aware privacy protection for smart water meters. In *Proceedings of the 8th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation*, pages 31–40. ACM, 2021. URL https://doi.org/10.1145/3486611.3486650.

[42] P. Pirzada, A. Wilde, G. H. Doherty, and D. Harris-Birtill. Ethics and acceptance of smart homes for older adults. *Informatics for Health and Social Care*, 47(1):10–37, 2022. URL https://doi.org/10.1080/17538157.2021.1923500.

[43] D. Wright. The state of the art in privacy impact assessment. *Computer Law & Security Review*, 28(1):54–61, 2012. URL https://doi.org/10.1016/j.clsr.2011.11.007.