AMBIENT PERSONAL DATA AND THE RIGHT TO BYSTANDER PRIVACY: AI, SMART GLASSES, SOCIAL MEDIA, AND THE LEGAL ARCHITECTURE OF UNCHOSEN SURVEILLANCE

Lotus Khanna, Christ University

Shashank Soni, Christ University

ABSTRACT

This manuscript argues that the forthcoming evolution of privacy law pertains not to the interaction between a data subject and the entity to whom she voluntarily reveals information, but rather to the pervasive phenomenon of individuals being transformed into data by external devices, platforms, or models. Smart eyewear, mobile cameras, house doorbells, auto sensors, workplace analytics, social media integration, and AI training frameworks have rendered personal data ubiquitous: passively created, relationally interconnected, machine-readable, and perpetually reused. Current privacy regimes provide limited instruments for addressing this issue, such as data minimization, restrictions on special categories, biometric regulations, wiretap legislation, and constitutional principles acknowledging the mosaic-like intrusiveness of aggregated data. However, these instruments are inadequately developed for bystanders, who do not fit the roles of traditional consumers or suspects, users or workers, contractual parties or typical litigants. The book introduces an innovative idea of the right to bystander privacy: a right to acceptable non-participation in data collection, identification, inference, and dissemination. It then advocates for an Ambient Capture Duty, underpinned by Bystander Impact Assessments, device-specific capture minimization, anti-identification defaults, context-sensitive no-capture zones, auditable provenance, and collective remedies. The primary assertion is that privacy law ought to see bystander exposure as a data externality generated by socio-technical systems, rather than as a deficiency in person permission.

Keywords: ambient personal data | bystander privacy | AI | smart glasses | biometric surveillance | consent | social media | GDPR | EU AI Act | DPDP Act | CCPA | data protection reform.

I. Introduction: From Chosen Disclosure to Unchosen Datafication

Privacy legislation is structured around a common scenario: an individual supplies information to an organization, which articulates a goal, and the subject is invited to accept, decline, or subsequently contest the processing. That scene remains significant. However, it no longer characterises the typical privacy infringement of interconnected existence. The prevalent scenario is incidental: an individual traverses behind a tourist equipped with camera glasses, occupies a doorstep while a smart doorbell captures footage, emerges in the backdrop of a social media post, converses near an AI assistant, sits adjacent to a commuter whose phone transmits location and Bluetooth data, or is deduced from the contact lists, photographs, and messages of others. The individual has not shown herself; she has been revealed by her environment.

The transition is not solely technological. It alters the legal framework of privacy. Conventional doctrine frequently enquires if an individual choose to disclose information, whether the expectation of privacy was justifiable, whether a firm provided sufficient warning, or whether a platform user granted consent. Ambient data environments address each of those enquiries. They transform privacy from a bilateral exchange into a geographical and relational state. They also render privacy violations probabilistic: the harm may not manifest at the moment of data acquisition, but subsequently, when a face is identified, a voice is transcribed, a scene is catalogued, a location history is compiled, or a training dataset is accessed by a model.

This paper employs the term ambient personal data to denote information regarding identifiable or reasonably linkable individuals that is generated passively by devices and infrastructures integrated into commonplace surroundings, frequently without direct engagement between the data subject and the collector. Ambient data is not merely background information; it is data generated as a result of the background transforming into a sensing surface. A pavement, café, school corridor, train platform, apartment hallway, livestream, and comment thread can all serve as input layers for identification and inference.

The bystander is the constitutional and statutory orphan of this environment. She is present enough to be recorded but absent from the notice interface; visible enough to be identified but invisible to the contractual relationship; socially connected enough to be inferred but legally remote from the platform's consent flow. The bystander is therefore not an edge

case. She is the modal data subject of ambient computing.

The article's central argument is that bystander privacy should be recognized as a distinct legal interest: the right of reasonable non-participation in the capture, identification, inference, retention, and propagation of one's personhood by ambient data systems. The right is not an absolute veto over observation in public; democratic societies require journalism, documentation, accessibility aids, artistic expression, public safety, and interpersonal memory. Rather, it is a right against disproportionate conversion of ordinary co-presence into durable, searchable, and action-guiding data. This distinction matters because the privacy loss created by smart glasses or AI training is not identical to being seen. It is being made persistently computable.

The manuscript proceeds in nine parts. Part II defines ambient personal data and distinguishes it from transactional personal data. Part III develops the right to bystander privacy. Part IV explains why consent collapses in ambient environments. Part V examines the technical roles of AI, smart glasses, social media, and digital surveillance. Part VI evaluates legal frameworks, including the GDPR, the EU AI Act, the CCPA/CPRA, Illinois's biometric law, India's DPDP framework, wiretap law, tort law, and constitutional privacy doctrine. Part VII identifies structural gaps. Part VIII proposes a reform architecture centered on an Ambient Capture Duty. Part IX concludes by arguing that the question is no longer whether privacy survives in public, but whether law can prevent public life from becoming a compulsory training set.

II. Ambient Personal Data: Definition, Properties, and Stakes

Ambient personal data has four defining properties. First, it is passive. The data subject often does nothing that looks like disclosure. Her face, gait, voice, posture, clothing, companions, route, or proximity to objects is sensed because a device near her is sensing by default. Second, it is relational. It is frequently collected by a device owned by someone else or by a platform to which someone else uploaded content. Third, it is inferential. Its value often lies not in the raw image, audio, or coordinate, but in the predictions made from it. Fourth, it is durable and mobile. It can migrate from the moment of capture into search indexes, biometric templates, model weights, advertising profiles, law enforcement tools, workplace dossiers, or social graphs.

These properties distinguish ambient personal data from the data categories around which most privacy instruments were designed. A purchase record, health form, bank application, or app registration creates an administrable relationship: there is a controller, a declared purpose, a user interface, an account, a notice, and a mechanism for rights. Ambient data often has none of those features. It arises from encounter, not enrollment.

The social stakes are significant as ambient data diminishes the practical distinction between being observable and being identifiable. Historically, a one may encounter another in public without possessing knowledge of her name, address, social media presence, employment background, or political affiliations. The old friction of identification served a privacy function. AI weakens that friction by making face recognition, speech transcription, object recognition, geolocation, re-identification, and cross-database linkage cheaper and more scalable.[1] The important legal change is that obscurity, once supplied by practical difficulty, now requires affirmative design or regulation.

Ambient data also destabilizes the category of 'publicly available' information A face observable on a sidewalk is not synonymous with a face registered in a searchable biometric database; a conversation perceivable to a nearby individual is not correlated with a transcript archived in a cloud service; a photograph shared by a friend is not synonymous with the same image incorporated into a foundational model. Public accessibility does not settle downstream purpose, scale, or persistence. The manuscript calls this the publicity fallacy: the assumption that exposure to some humans in a context authorizes processing by any machine in any future context.

A useful way to understand ambient data is by analogy to environmental externalities. A factory may create value for its customers while imposing pollution on neighbors who never contracted with it. Ambient data systems create convenience, safety, entertainment, and AI capabilities for users and firms while imposing capture risk on nearby people. The cost is not distributed through the transaction that produces the benefit. Bystanders pay in identifiability, behavioral chilling, reputational exposure, and loss of practical anonymity.

The environmental analogy should not be exaggerated. Data is non-rivalrous and can

1 See Alessandro Acquisti, Curtis Taylor & Liad Wagman, The Economics of Privacy, 54 J. Econ. Literature 442, 448-52 (2016) (describing how technological change alters the costs of collecting and using personal information).

be replicated without exhaustion. Yet that very feature makes the externality more severe: once an ambient trace exists, it can be recombined indefinitely. The injury is not merely that data has been collected, but that a person has been made available for secondary uses that she could neither predict nor contest.

Ambient personal data therefore requires a shift from data possession to data event governance. Law should ask not only who holds the data, but what event caused a person to become data, who had the capacity to prevent unnecessary capture, what inferences are enabled, and what future audiences may access the trace. That event-oriented approach makes bystanders visible as rights holders.

III. The Right to Bystander Privacy

The right to bystander privacy is the interest of a person who is incidentally subjected to sensing, identification, inference, retention, or disclosure because of the actions of another person, organization, or device. The bystander may be physically near the sensor, socially connected to the primary user, or merely represented in a dataset derived from others. The right is not reducible to secrecy. Bystanders are often visible. The injury lies in transformation: ordinary visibility is transformed into persistent datafication.

Classical privacy theory provides partial foundations. Warren and Brandeis framed privacy as a response to technologies of image reproduction and mass circulation, not simply as secrecy.[2] Nissenbaum's theory of contextual integrity explains why information flows can violate privacy even when the information is not secret: appropriateness depends on actors, attributes, transmission principles, and context.[3] Solove's taxonomy shows that privacy harm includes aggregation, identification, insecurity, secondary use, exclusion, disclosure, and intrusion.[4] Bystander privacy draws from all three traditions but adds a distinct emphasis: the injured person did not initiate the information flow.

The bystander right has five components. The first is capture minimization: ambient systems should avoid collecting identifiable bystander data when the primary function can be achieved with less data. The second is anti-identification: recording should not automatically

2 Samuel D. Warren & Louis D. Brandeis, The Right to Privacy, 4 Harv. L. Rev. 193, 195-96 (1890).
3 Helen Nissenbaum, Privacy in Context: Technology, Policy, and the Integrity of Social Life 127-29 (2010).
4 Daniel J. Solove, A Taxonomy of Privacy, 154 U. Pa. L. Rev. 477, 488-91 (2006).

imply facial, voice, gait, or social identity matching. The third is context limitation: data captured in one social setting should not be repurposed into another without a strong justification. The fourth is contestability: bystanders should have meaningful routes to challenge retention, indexing, and high-impact uses. The fifth is collective accountability: where individual notice is impossible, the law should impose duties ex ante through impact assessments, technical standards, and enforceable defaults.

The right also requires a vocabulary for bystander status. A proximate bystander is physically captured by a device, as when camera glasses record a stranger. A derivative bystander is inferred from another person's data, as when a contact list, tag, or group photo reveals relationships A latent bystander exists inside training data or surveillance records but remains unrecognized until a subsequent inquiry. An institutional bystander is recorded by systems implemented in public or semi-public environments, including retail establishments, educational institutions, airports, transportation systems, and workplaces. These categories show why privacy cannot be managed solely at the account level.

Bystander privacy should be understood as a right of reasonable non-participation. The phrase matters. It avoids the impossibility of a total right not to be seen, while insisting that people should not be silently conscripted into biometric matching, behavioral analysis, AI training, or perpetual indexing. It also prevents a false choice between privacy and social life. People cannot participate in society if every participation becomes a raw material for unknown models and watchlists.

Recognising bystander privacy would elucidate the relationship between individual autonomy and democratic liberty. Privacy transcends consumer preference; it facilitates experimentation, disagreement, association, closeness, and the acceptance of error.. Julie Cohen has argued that privacy shelters the processes by which people develop selves capable of democratic participation.[5] Bystander privacy extends that insight from the home and device to the shared spaces where social life unfolds.

IV. Consent in the Age of AI: Why Notice-and-Choice Fails for Bystanders

Consent remains an important legal basis for many data practices. The GDPR describes consent as a voluntarily provided, explicit, informed, and unequivocal expression of

5 Julie E. Cohen, What Privacy Is For, 126 Harv. L. Rev. 1904, 1906-08 (2013).

preferences. The DPDP Act of India mandates that consent must be free, specific, informed, unconditional, and clear, and it acknowledges the right to withdraw consent when it serves as the basis for processing..[6] These definitions assume an identifiable data principal and an administrable request. Ambient data collection violates that assumption.

Bystander consent fails for at least six reasons. First, the bystander may not perceive capture. A tiny LED on a pair of glasses, a doorbell camera, a microphone array, or a vehicle sensor may not communicate enough to trigger a real choice. Second, the bystander may know capture is occurring but have no practical exit. Eschewing streets, cafes, schools, transit, or jobs does not constitute a substantive opt-out. The onlooker may object to one downstream application while permitting another; she may approve a tourist shot but oppose facial recognition or AI training. Fourth, the capture may occur before notice can be supplied. Fifth, the collector and user may be different actors. Sixth, the data may later be repurposed into uses no one could explain at the moment of collection.

The result is not simply imperfect consent but consent displacement. Consent is obtained, if at all, from the wearer, homeowner, platform user, employer, or store operator, while the costs fall on the bystander. The system treats one person's agreement as a license over another person's data. This is consent laundering: the transformation of a primary user's assent into apparent legitimacy for processing non-users.

AI intensifies the problem because AI systems convert ambient traces into predictions and actions. A raw image may reveal little to a human observer, but the same image may support face matching, emotion classification, age estimation, identity clustering, object recognition, location inference, and social graph analysis. The relevant consent would need to cover not only capture but model development, deployment, inference, retention, auditing, and sharing. Such consent cannot be obtained through ordinary signage or platform terms.

The challenge has compelled regulators to use non-consent frameworks, including legitimate interests, public tasks, or statutory exceptions. The European Data Protection Board's 2024 ruling on AI models underscores that the consideration of legitimate interest must take into account necessity, balancing, fairness, transparency, and the implications of unlawful training data..[7] This is useful but insufficient. Legitimate interest assessments tend to be

6 Digital Personal Data Protection Act, No. 22 of 2023, sec. 6, India Code (2023).
7 European Data Protection Board, Opinion 28/2024 on Certain Data Protection Aspects Related to the

controller-centered. They ask what the controller needs and what safeguards it proposes. Bystander privacy requires a different baseline: if a person did not initiate the relationship, the burden should shift to the collector to justify why identifiable capture was necessary at all.

The future of consent should therefore be narrower and stronger. Consent should govern genuine relationships in which the data subject can understand and reject the processing without disproportionate penalty. Ambient data should be governed primarily by duties, prohibitions, minimization, and impact assessments. Put differently, the failure of consent is not a reason to remove protection; it is a reason to move protection upstream.

V. Technology as a Legal Fact: AI, Smart Glasses, Social Media, and Digital Surveillance

AI is not only a downstream user of ambient data; it changes what ambient data is. A photograph once functioned as a record. In an AI environment, it is also an input for recognition, classification, synthesis, and model training. Audio is no longer only sound; it is searchable text, speaker identity, sentiment signal, and behavioral marker. Location is not only where someone was; it is routine, association, workplace, health inference, worship practice, or political participation. The same raw trace becomes legally different because its machine-readable affordances change.

Smart glasses concentrate the bystander problem because they relocate high-resolution sensing from fixed infrastructure to the face. A phone camera is visible as a held object. Glasses are worn continuously, socially normalized, and aimed by attention itself. A bystander cannot easily tell whether the wearer is looking, recording, livestreaming, translating, asking an AI to identify objects, or merely wearing ordinary eyewear. Meta's Ray-Ban Meta product materials describe a capture LED intended to signal photo and video capture, but the existence of an indicator does not settle whether notice is salient, durable, or meaningful in ordinary social settings.[8]

The I-XRAY demonstration by Harvard students illustrates the legal significance of composability. The glasses themselves did not need to ship with a face-recognition feature to become part of a real-time identification pipeline. Video capture, livestreaming, face search services, data broker information, and mobile notifications could be combined to identify

Processing of Personal Data in the Context of AI Models 19-24 (Dec. 17, 2024).
8 Meta, Ray-Ban Meta AI Glasses, Meta Store (last visited June 4, 2026).

strangers and surface personal details.[9] This reveals an important regulatory point: privacy risk resides in ecosystems, not only devices. A product may appear compliant when evaluated alone but become invasive when connected to public databases and AI services.

Social media adds a second layer. Platforms mediate the upload, tagging, compression, ranking, and reuse of images and text created by users Bystanders are represented in images they did not upload, tagged by acquaintances, mentioned in comments, and recorded in live streams. When such material is subsequently utilized for AI training, the onlooker may not even be a user of the platform. The Irish Data Protection Commission openly acknowledged that Meta had notified them of intentions to train a big language model utilizing public pornographic content from Facebook and Instagram throughout the EU/EEA, which raised concerns that postponed the implementation..[10] The controversy shows that 'public content' is often relational content: one user's public post may contain another person's face, name, home, child, or association.

Digital surveillance infrastructures are broader still. Retail facial recognition, school monitoring, workplace analytics, home security networks, vehicle cameras, and smart-city systems create archives of people who may never become customers, suspects, employees, residents, or platform users. The FTC's Rite Aid action, which resulted in a five-year prohibition on using facial recognition technology for surveillance purposes, demonstrates how private-sector surveillance can produce humiliation, misidentification, and discriminatory risk when deployed without reasonable safeguards.[11]

These technologies share a pattern: they compress time, space, and identity. They make present observations available to future actors; they make local events accessible globally; and they convert ambiguous human appearance into actionable identity claims. Legal analysis must treat those transformations as part of the privacy injury, not merely as later uses.

VI. Existing Legal Frameworks and Their Bystander Gaps

The GDPR is the most sophisticated general data protection regime for ambient data

9 Victoria Song, College Students Used Meta's Smart Glasses to Dox People in Real Time, The Verge (Oct. 2, 2024); John Koetsier, Meta's Ray-Ban Smart Glasses Used To Instantly Dox Strangers In Public Thanks To AI And Facial Recognition, Forbes (Oct. 3, 2024).
10 Data Protection Commission, DPC Statement on Meta AI (May 21, 2025).
11 Rite Aid Banned from Using AI Facial Recognition After FTC Says Retailer Deployed Technology Without Reasonable Safeguards, Fed. Trade Comm'n (Dec. 19, 2023).

because it defines personal data broadly, regulates processing rather than only disclosure, requires lawful bases, restricts special categories, imposes data minimization, recognizes data protection by design, and requires data protection impact assessments for high-risk processing.[12] Its strengths matter. A bystander's face, voice, or location can be personal data when identifiable. Biometric data used for unique identification may receive special protection. Controllers must explain purposes and limit retention. Supervisory authorities can sanction unlawful scraping and facial recognition practices.

Yet GDPR implementation often struggles with ambient capture. The first problem is notice. Articles 13 and 14 assume that controllers can provide information to data subjects, but ambient systems may encounter thousands of transient people. The second problem is lawful basis elasticity. Legitimate interests can become a fallback for collection that cannot practically obtain consent. The third problem is household and journalistic contexts. A private person's camera glasses or social media upload may fall partly outside ordinary controller obligations, even though the resulting data may later enter commercial pipelines. The fourth problem is remedy. Bystanders often cannot know which controller captured them, much less exercise access, erasure, or objection rights.

The EU AI Act adds an important system-level layer. ***It regulates AI systems based on risk, forbids designated practices, handles biometric classification and emotion detection in specific contexts, governs high-risk systems, and enforces transparency requirements for certain AI interactions and synthetic content.***.[13] But it is not a comprehensive bystander privacy code. It regulates AI systems, not every sensor that supplies AI systems. It can require transparency for AI outputs while leaving unresolved the earlier stage of bystander capture. It may prohibit or constrain some biometric uses while allowing many consumer or low-risk uses that nevertheless normalize unchosen datafication.

The CCPA, as amended by the CPRA, gives California consumers rights to know, delete, correct, opt out of sale or sharing, and limit use of sensitive personal information.[14] It also treats precise geolocation, biometric information used for identification, and certain sensitive categories as sensitive personal information. The framework is valuable against

12 GDPR, supra note 11, arts. 4(1), 5, 6, 9, 25, 35.
13 Regulation (EU) 2024/1689 of the European Parliament and of the Council arts. 5, 6, 50, 2024 O.J. (L 2024/1689) 1 [hereinafter EU AI Act].
14 Cal. Civ. Code secs. 1798.100-.199.100 (West 2024).

businesses that process bystander data at scale. Its limitations are equally important: it is consumer-facing, business-threshold dependent, and not designed for real-time public-space capture. A bystander may have a formal right after the fact but no practical way to identify the business, prove capture, or stop the initial collection.

Illinois's Biometric Information Privacy Act is one of the strongest targeted statutes because it requires informed written consent before collecting or obtaining biometric identifiers or information, mandates retention policies, and provides a private right of action.[15] BIPA's strength is also its narrowness. It reaches biometric identifiers and biometric information, not all ambient data. It does not govern ordinary image capture unless the capture is processed into covered biometrics. It therefore addresses one important endpoint of ambient data but not the broader ecology of passive sensing.

India's DPDP Act and DPDP Rules create a modern data protection framework for digital personal data. The Government of India stated that the DPDP Rules, 2025 were notified on November 14, 2025 and gave effect to the 2023 Act after nationwide consultation.[16] The DPDP Act's consent provisions and rights framework are relevant to ambient data when identifiable digital personal data is processed. Nevertheless, the bystander gap is widely acknowledged: identification and consent pose difficulties in public surveillance, exemptions may be considerable, and the existing framework lacks a device-level or spatial architecture for unintentional non-users.

Legislation regarding wiretapping and eavesdropping concerns audio recording, especially in jurisdictions requiring consent from all participants in the conversation. They are significant because smart eyewear and AI helpers frequently utilise microphones in addition to cameras. Wiretap legislation is disjointed, concentrating on communications instead of including all types of personal data, and frequently neglects to consider visual capture, biometric inference, or artificial intelligence training. Privacy tort law is structurally weak for bystanders. Intrusion upon seclusion typically requires an intrusion into a private place or affairs that would be highly offensive. Public disclosure of private facts requires publicity of private information. False light and appropriation have their own narrow elements. A bystander recorded in public and never individually publicized may suffer real datafication without

15 740 Ill. Comp. Stat. 14/1-99 (2024).
16 Press Information Bureau, Government of India, DPDP Rules, 2025 Notified: A Citizen-Centric Framework for Privacy Protection and Responsible Data Use (Nov. 17, 2025).

satisfying tort doctrine. The law's public/private spatial distinction is poorly adapted to technologies that make public observations searchable, linkable, and permanent.

Constitutional privacy doctrine offers important analogies, especially in search-and-seizure law. Carpenter recognized that cell-site location records can reveal the privacies of life even though held by third-party carriers.[17] That reasoning matters because ambient data creates similar mosaic effects. A single sidewalk sighting may reveal little; a network of doorbells, cameras, vehicle sensors, and data brokers can reveal everything. ***India's Supreme Court in KS Puttaswamy case grounded privacy in dignity, autonomy, and informational self-determination, which provides a constitutional vocabulary for challenging pervasive digital surveillance***.[18] ***The UK Court of Appeal in Bridges likewise found unlawful aspects of live automated facial recognition deployment by police, emphasizing inadequate legal framework and safeguards***.[19]

The comparison lesson is evident: current legislation has elements of bystander protection, although these elements are dispersed throughout data protection, AI regulation, consumer privacy, biometric law, wiretap law, tort, and constitutional doctrine. None alone supplies a general right of reasonable non-participation in ambient data systems.

VII. Structural Failures: Why Current Law Underprotects Bystanders

The first structural failure is the subject gap. Privacy law often recognizes the user, consumer, employee, patient, subscriber, or suspect. The bystander is not easily one of these Although she theoretically possesses rights as a data subject, the practical means to access those rights is absent. A right that cannot be identified, asserted, or upheld is not a practical right. The second failure is the moment gap. Law often intervenes after collection, through access, deletion, correction, objection, or damages. Bystander privacy is most effectively protected before or at capture. Once a face has been extracted into a template, once audio has been transcribed, once a scene has been uploaded, the bystander has already lost practical control. Post-hoc rights remain important but cannot substitute for capture minimization.

The third failure is the purpose-drift gap. Ambient data collected for one apparent

17 Carpenter v. United States, 138 S. Ct. 2206, 2217-20 (2018).
18 Justice K.S. Puttaswamy (Retd.) v. Union of India, (2017) 10 S.C.C. 1 (India).
19 R (Bridges) v. Chief Constable of S. Wales Police [2020] EWCA Civ 1058 (Eng.).

purpose can become useful for another. A doorbell camera installed for package security becomes neighborhood surveillance; a store camera becomes facial analytics; a social media post becomes AI training material; smart glasses designed for memory aids become real-time identification tools. Purpose limitation doctrines exist, but they struggle when datasets and models are reused across complex supply chains.

The fourth failure is the inference gap. Many statutes regulate identifiable records but not the production of sensitive inferences from seemingly ordinary data. Yet ambient data is valuable because it produces inferences: who met whom, what mood someone displayed, whether a person is pregnant, whether a worker is distracted, whether a protester attended a march, whether a child has a routine route to school. The legal object should be not only raw data, but inferential capability.

The fifth failure is the collective-action gap. A single bystander rarely has the information, resources, or incentive to sue over incidental capture. The harm is diffuse and cumulative. Meanwhile, the collector's benefits are concentrated. This is the classic structure of an externality. Without collective enforcement, impact assessment, public registers, and regulator capacity, ambient privacy harms will be systematically under-remedied.

The sixth failure is the design gap. Law often commands transparency but not architecture. A notice sign, LED, or privacy policy may inform people that sensing occurs without altering what is collected. In ambient systems, design choices are law-like: whether a device has a hardware indicator, whether it can be disabled, whether face recognition is off by default, whether bystander faces are blurred on device, whether audio is locally processed, whether uploads contain provenance metadata, whether retention is short, and whether third-party integrations are blocked. Privacy law should treat these choices as regulatory sites.

The seventh failure is the scale gap. Observation by humans has natural limits. Observation by machines does not. The legal system has often treated public exposure as a waiver because human observation was ephemeral and practically limited. AI systems eliminate the conditions that made that waiver tolerable. They turn exposure into extraction, extraction into inference, and inference into action.

VIII. Reform Architecture: The Ambient Capture Duty

This manuscript proposes an Ambient Capture Duty: a legal duty imposed on entities that design, deploy, or commercialize devices, platforms, or systems capable of collecting identifiable data about non-users in shared environments. The duty would not prohibit all recording. It would require organizations to justify, minimize, signal, protect, and account for bystander capture. The duty should attach to commercial and institutional actors, not ordinary memory or casual photography, except where ordinary devices are configured to enable systematic identification, commercial reuse, or surveillance.

The first element is capture necessity. A system should collect identifiable bystander data only when necessary for a legitimate and proportionate function. If a smart device can perform a task with on-device processing, ephemeral buffers, low-resolution sensing, object detection without identity, or automatic bystander blurring, it should do so. Capture minimization must be assessed at the sensory layer, not merely at the database layer.

The second element is anti-identification by default. Devices and platforms should not perform or facilitate face, voice, gait, or other biometric identification of bystanders unless a strict legal basis applies. This default would address the gap revealed by composability: even if a device maker does not itself identify bystanders, it should not design easy pathways for third-party identification without safeguards.

The third element is visible statefulness. Ambient sensors should communicate not merely their existence but their state: inactive, locally processing, recording, livestreaming, identifying, or uploading. Current indicators often compress these states into a binary light. For bystanders, the legally relevant question is not only whether a camera exists, but what the system is doing. Statefulness should be standardized, multimodal where appropriate, non-disableable for high-risk capture, and paired with enforceable penalties for circumvention.

The fourth element is contextual no-capture and limited-capture zones. Educational institutions, medical facilities, religious establishments, residential interiors, shelters, judicial facilities, restrooms, changing areas, polling sites, labour organising venues, and sensitive public services necessitate enhanced standards. The legislation already acknowledges context in secrecy, wiretap regulations, and data protection. Ambient privacy must implement context via device policies, local signage, reliable geofencing, institutional regulations, and penalties

for systematic infractions.

The fifth element is Bystander Impact Assessment. Existing data protection impact assessments are useful but should be expanded for ambient systems. A Bystander Impact Assessment would ask: Who may be incidentally captured? Can the system operate without identity? What sensitive contexts are likely? What downstream inferences are enabled? What retention periods apply? What third-party integrations exist? What opt-out or objection mechanisms are realistic? How will vulnerable groups be affected? What collective remedies are available?[20] The assessment should be required before deployment of smart glasses with recording functions, large-scale biometric analytics, social media AI-training pipelines, and surveillance systems in semi-public environments.

The sixth element is provenance and propagation control. Ambient data must include information that specifies capture context, consent status, bystander treatment, AI-training eligibility, retention deadline, and identification constraints. Provenance will not resolve all issues; nonetheless, without it, downstream processors are unable to differentiate between a consensual selfie and a group image featuring non-consenting spectators. Legal duties should follow the data across supply chains.

The seventh element is collective enforcement. Regulators should be able to investigate ambient capture without waiting for individualized complaints. Civil society organizations should have representative action authority. Statutory damages or administrative penalties should be available for prohibited biometric identification, circumvention of indicators, and unlawful retention of bystander data. A bystander right without collective enforcement will underperform because the persons most affected are least likely to know.

The eighth element is public procurement discipline Governments ought to refrain from acquiring or using ambient surveillance systems unless vendors can prove capture minimization, anti-identification defaults, auditability, bias testing, deletion methods, and public reporting. Public procurement can establish market standards more rapidly than individual enforcement actions.

The ninth element is interoperability with expression and documentation rights. A

20 Cf. GDPR, supra note 11, art. 35 (requiring data protection impact assessments for processing likely to result in high risk).

bystander right must not become a tool to suppress journalism, evidence gathering, police accountability, artistic practice, or accessibility technologies for people with disabilities. The better line is not capture versus no capture, but proportional capture versus identity extraction. Law should protect recording for accountability while restricting face matching, automated profiling, commercial reuse, and retention unrelated to the original expressive or evidentiary purpose.

IX. Application to Smart Glasses

Smart glasses should be regulated as ambient capture devices when they include cameras, microphones, live AI assistance, or cloud upload. The baseline duty should require hardware recording indicators, state-specific signals, anti-tamper design, short default retention, on-device processing where feasible, and no biometric identification of bystanders by default. Manufacturers should publish bystander-facing privacy specifications, not only user-facing privacy policies.

A smart-glasses regime should also distinguish ordinary recording from enhanced capture. Ordinary recording is the storage of photos, video, or audio. Enhanced capture includes livestreaming, real-time AI analysis, translation of nearby speech, face or voice matching, scene summarization, object recognition tied to location, and automatic upload. Enhanced capture should trigger stronger duties because it makes the bystander immediately computable. A person may tolerate being in a vacation photo but object to being described, translated, indexed, and cross-referenced by a cloud model.

The legislation ought to mandate social-context protections. In semi-private environments like clinics, classes, salons, businesses, and residences, venue operators ought to implement no-enhanced-capture regulations. Manufacturers ought to endorse such regulations via gadget modes and venue beacons, while obstructing clandestine overrides. The objective is not to prohibit smart glasses but to avert their acceptance as unobtrusive monitoring devices.

Ultimately, legislation regarding smart glasses must acknowledge the accessibility conundrum. Wearable AI can aid those with blindness or impaired vision by elucidating scenes, reading text, or recognizing things.. A bystander privacy regime should not disable such uses. Instead, it should prefer local processing, non-retention, no identity matching unless specifically authorized by the identified person or necessary for accessibility in a narrow

setting, and clear indicators that distinguish assistive perception from recording and uploading. Accessibility and bystander privacy are not enemies; both depend on trustworthy design.

X. Application to Social Media and AI Training

Social media is the main laundering channel for bystander data. A person may never join a platform yet appear in thousands of user uploads. Platforms often treat uploader consent, account settings, or public-post status as the relevant authorization. That approach fails where the content contains other identifiable people. The uploader controls publication of the post, but not all downstream uses of every person represented in it.

Platforms should therefore separate publication consent from model-use eligibility A public post may be accessible on a site without qualifying for biometric indexing, AI training, advertising inference, or external data licensing. Bystanders in photographs and videos should activate specific protections: automatic exclusion from facial recognition, prohibition of biometric template generation without a definitive legal justification, and procedures for removal or masking upon objection from a non-uploader.

AI training raises a harder question because removal after training is technically and legally complicated. A model may not store a discrete copy of each training example, but it may retain patterns, memorized fragments, or representational traces. The legal response should not wait for perfect machine unlearning. Instead, it should impose dataset governance before training: source documentation, exclusion of high-risk ambient data, honor of opt-outs where feasible, minimization of facial and voice data, privacy-preserving training methods, and audit trails showing that public content was not treated as unrestricted raw material.

The right to be forgotten offers a useful analogy but also a warning. Google Spain recognized that search indexing can alter the practical accessibility of lawful information.[21] AI training alters information differently: it may abstract, memorize, infer, or synthesize. Bystander privacy therefore needs both delisting-like remedies for outputs and dataset-level restraints on input. Otherwise, the law will repeatedly chase downstream outputs after the social cost has already been absorbed.

21 Google Spain SL v. Agencia Española de Protección de Datos (AEPD), Case C-131/12, ECLI:EU:C:2014:317 (May 13, 2014).

XI. Application to Retail, Home, and Urban Surveillance

Retail surveillance illustrates why consumer privacy is not enough. A store's customer may have a loyalty account or receive a privacy notice, but a passerby, child, delivery worker, protester, or companion may be captured without becoming a customer The implementation of facial recognition for theft deterrence may result in misidentification, perpetuate bias, and engender enduring stigma. The Rite Aid incident is not merely an isolated enforcement issue; it serves as a cautionary tale that security justifications might legitimise private watchlists in public areas.

Home surveillance networks produce a different asymmetry. A homeowner may purchase a doorbell camera for legitimate security reasons, but the camera also records neighbors, pedestrians, delivery workers, domestic workers, children, and guests. The home's boundary becomes a sensing zone that extends into shared space. Law should respond through retention limits, neighborhood aggregation limits, restrictions on law enforcement access without legal process, visible indicators, and obligations on manufacturers to prevent mass sharing by default.

Urban surveillance poses the most significant threat to democracy. Smart-city sensors, transportation cameras, license plate scanners, drones, and public-private camera networks can transform urban movement into a searchable database. Even when each camera is legal, the cumulative effect can suppress dissent, religion, medical visits, personal relationships, and the movement of minorities. Carpenter’s insight about the whole being more revealing than the parts should guide urban surveillance governance.[22]

The reform principle is proportional legibility. Some public-space sensing is necessary for safety, accessibility, traffic management, and accountability. But legibility should be proportional to the public function. Counting crowd density is not the same as identifying every face. Detecting a vehicle hazard is not the same as retaining driver identities. Managing a transit platform is not the same as building a citywide movement dossier.

XII. Toward a Legal Test for Bystander Privacy

Courts and regulators need administrable criteria. This manuscript proposes a six-factor

22 Carpenter, 138 S. Ct. at 2217.

test for bystander privacy risk. First, identifiability: does the system capture or enable identification of a person, directly or through linkage? Second, initiation: did the person meaningfully initiate the data relationship? Third, context: did the capture transpire in a sensitive, semi-private, mandatory, or susceptible environment? Fourth, transformation: does the system convert human-observable knowledge into searchable, persistent, inferred, or automated decision-making inputs? Fifth, propagation: is it possible for the data to transfer among actors, objectives, databases, or models? Sixth, power asymmetry: can the bystander realistically avoid, contest, or obtain remedy?

The test deliberately avoids asking whether the person was in public as a threshold bar. Publicness remains relevant to context and expectations, but it should not defeat protection where technology alters scale and persistence. A person may be visible in public without consenting to biometric enrollment, AI training, or indefinite retention.

The test would also help distinguish legitimate from illegitimate uses. A journalist recording a public official at a public event scores differently from a retailer enrolling every visitor into a facial recognition database. A blind user's local scene description scores differently from a commercial platform uploading bystander faces for model improvement. A one-time family photograph scores differently from a social app that extracts face embeddings from every upload. The examination emphasizes transformation and power, rather than moral hysteria around surveillance cameras.

Regulators may integrate the test into guidelines, DPIA templates, AI compliance evaluations, consumer protection regulations, and procurement criteria. Courts may utilize it to elucidate reasonable expectations, inequity, legitimate interests, or proportionality. Legislatures could codify it as a risk-classification trigger for ambient capture systems.

XIII. Objections and Responses

The first objection is that bystander privacy is impractical because public life necessarily involves observation. The response is that the proposed right does not prohibit ordinary observation. It targets the conversion of observation into durable, searchable, inferential, and transferable data. Law already distinguishes seeing from recording, recording from publishing, and publishing from biometric identification. Ambient systems require extending those distinctions.

The second objection is that notice and user choice are sufficient. The response is that bystanders are not users. They do not receive the notice, cannot negotiate the terms, and often cannot exit the space. User choice can legitimize processing of the user's data, but it cannot fully authorize extraction from non-users.

The third objection is that stronger bystander rights would burden innovation. The response is that privacy-preserving architecture is innovation. On-device processing, data minimization, privacy indicators, differential privacy, face blurring, short retention, and provenance systems create trustworthy markets. The alternative is a race to the bottom in which consumer convenience depends on socializing the privacy costs onto everyone nearby.

The fourth objection is that bystander privacy could suppress speech or accountability recording. The response is that the reform architecture expressly distinguishes documentation from identity extraction and commercial reuse. The recording of police misbehavior, prejudice, or journalistic activities should continue to be safeguarded. Automated identification, persistent indexing, and unrelated AI training of bystanders require stronger justification.

The fifth objection is that technical measures may be circumvented. That is accurate, however not conclusive. Seatbelts, locks, encryption, and speed limits can also be bypassed. Law works by setting defaults, creating incentives, enabling enforcement, and raising the cost of harmful conduct. Ambient privacy is no different.

XIV. Conclusion: Privacy for the People Who Never Clicked Accept

The privacy crisis of the AI era is not only that companies know more about their users. It is that non-users are increasingly made available to systems they never chose. Ambient personal data transforms co-presence into computability. Smart glasses make attention into a sensor. Social media makes friendship into a data pipeline. AI makes raw traces into prediction. Digital surveillance makes movement into institutional memory.

The law's inherited tools remain valuable but incomplete. Consent is meaningful where relationships are genuine; data protection principles matter where controllers can be identified; biometric statutes matter where identity templates are created; constitutional doctrines matter where the state aggregates the privacies of life. But bystanders require a new legal center of gravity: the right of reasonable non-participation in unchosen datafication.

Recognizing bystander privacy would not end public observation, photography, accessibility, journalism, or security. It would instead restore a boundary that technology has quietly removed: the boundary between being present and being processed. The future of privacy will be decided not only at login screens and cookie banners, but in the spaces between people. A legal system that protects only the person who clicked accept will fail the person standing beside her.