# Can privacy concerns associated with Facial Recognition Technology be effectively addressed?


## Abstract

Facial Recognition Technology (FRT) has moved from pilot to infrastructure. It authenticates phone users; filters airport queues, unlocks doors, tags photographs, and—more controversially—locates people in public space. This paper evaluates whether the privacy concerns that trail this diffusion can be effectively addressed without sacrificing legitimate benefits. The analysis proceeds along three vectors of risk: (i) mass surveillance and function creep, (ii) data vulnerabilities and biometric irreversibility, and (iii) accuracy–bias disparities that amplify privacy and fairness harms. Drawing on error-rate studies, breach cases, and data-protection frameworks (e.g., GDPR-style regimes), it argues that meaningful protection is possible only when use is narrowly scoped, data flows are minimised and encrypted end-to-end, and accountability is enforceable through transparent metrics and independent audit. Privacy-enhancing technologies (PETs)—on-device matching, federated learning, differential privacy, and secure enclaves—materially reduce exposure but cannot substitute for purpose limitation or deletion by default. A structured risk→remedy map and a deployment checklist translate the discussion into practice. The conclusion sets out a realistic call-to-action, supplemented by a short reflexive note that situates everyday consumer use versus remote public-space identification.


## 1. Introduction and Scope

Research question**:** *Can privacy concerns associated with FRT be effectively addressed?*

**Why now?** Over the last decade, advances in face detection, feature extraction, and matching pipelines have coincided with cheap sensors, expansive cloud capacity, and widespread consumer normalization (phone unlock, payments, photo tagging). Public tolerance has grown through convenience, yet unease spikes when the same capability migrates to streets, classrooms, or workplaces. The ambivalence is the core normative problem: the same technical stack underwrites both benign authentication and intrusive population-level identification.

Definitions and taxonomy**:** For clarity, the paper distinguishes: (a) **1:1 verification** ("Are you who you say you are?"); (b) **1:N identification** ("Who is this person from a gallery?"); and (c) **analytics/inference** (attributes, emotion, behavior). Privacy stakes escalate from (a) to (b) and are highest for (c), which can drift into pseudoscientific claims with discriminatory outcomes. Liveness detection and anti-spoofing are security controls, not privacy safeguards; they reduce impostor risk but do not address over-collection or repurposing.

Scope**:** The analysis focuses on three domains where privacy harm is most plausible: (1) live or near-live identification in public spaces; (2) large-scale storage, sharing, and cross-border transfer of facial templates; and (3) automated decision-making in sensitive contexts (policing, immigration, welfare, employment, healthcare). Each domain is examined for risks, governance gaps, and workable mitigations.

Method and contribution**:** A structured literature and policy synthesis is paired with a risk→remedy mapping. The contribution is practical: it distils a deployment checklist and a visual risk–safeguard table that institutions can adopt without expanding programme scope. In keeping with feedback, the Introduction remains lean, the Overview avoids repetition, and later sections link examples directly to solutions.

### 1.1. Ethical and Legal Framing

The analysis is grounded in a right-respecting, risk-minimization approach that balances three claims often set in tension: **security and efficiency**, **individual autonomy and dignity**, and **social equity**. Privacy is treated as both an individual right and a precondition for civic participation—the ability to move, assemble, and express without pervasive tracking. The paper therefore employs two tests repeatedly: (i) **necessity and proportionality** (is FRT needed for this task, and is there a less intrusive tool that would suffice?), and (ii) **accountability and remedy** (if something goes wrong, who is answerable and what recourse exists?). This framing aligns with GDPR-style principles while remaining technology-agnostic.

## 2. Overview of Current Use

We live in a world saturated with cameras: billions of devices capture streets, shops, transit hubs, offices, and homes, with the densest deployments in global cities and transport corridors (Lu, 2022; Villaluz, 2021). Consumer adoption has normalised FRT for smartphone unlocking, payments, and photo management; enterprises use it for workforce access control, patient matching in hospitals, and fraud reduction in banking (Magocha et al., 2021; Maxiom Technology, 2024). Governments deploy FRT at borders, for event security, and—more controversially—for watch-list matching in public spaces. During the COVID-19 pandemic, several jurisdictions experimented with integrating biometric monitoring into public-health enforcement, including mask-aware identification and crowd analytics (Pollard, 2020).

These deployments generate genuine benefits: frictionless authentication reduces account-takeover risk; automated access control scales beyond what human guards can do; and post-event video search can accelerate criminal investigations. Yet the same capture that removes a password can, in other contexts, erode anonymity in public life. The privacy question is therefore not whether FRT can be useful, but whether governance and system design can confine the technology to legitimate, proportionate purposes while preventing function creep.

### 2.1. Benefits Commonly Cited

FRT's growth has been driven by tangible benefits that are not merely rhetorical. In consumer contexts, password fatigue and phishing have made knowledge-based authentication brittle; biometric verification offers a fast, low-friction alternative that is resistant to credential stuffing. In healthcare, patient-matching with facial verification can reduce duplicate records, help locate the right file in emergency admissions, and mitigate medication errors when combined with clinician checks (Maxiom Technology, 2024). In aviation, e-gates and biometric boarding have increased throughput and reduced queuing times while still meeting border-security mandates. Retailers report shrink reduction when high-theft incidents are correlated with identified repeat actors, though such claims require independent audit to check for over-reach and discriminatory spill-overs (Walmart, 2019). Importantly, these benefits generally stem from **1:1 verification** or post-event investigation rather than open-ended live scanning; the risk profile changes markedly when identification becomes ambient and continuous.

### 2.2. Sources of Public Concern

Public concern centres on four themes that recur across surveys and case studies: (1) opacity about when and where faces are processed; (2) retention and repurposing of images and templates; (3) third-party sharing and cross-border transfers; and (4) unequal error burdens. These are not Luddite fears. They reflect core data-protection principles—**lawfulness**, **necessity**, **proportionality**, **purpose limitation**, **data minimization**, **storage limitation**, and **integrity and confidentiality**—that pre-date modern FRT but map neatly onto it. Where institutions can show, with evidence, that deployments satisfy these principles, acceptance rises; where they cannot, trust corrodes quickly and litigation follows (Hill & Mac, 2021; Scarcella, 2024).

### 2.3. Short case vignettes

- Smartphone authentication**:** Users enroll their face once; subsequent 1:1 matches happen locally, inside a secure enclave (Cipriani, 2020). *Implication:* when designed this way, consumer FRT can be privacy-preserving because templates do not leave the device and there is no gallery.
- Border e-gates: Travellers are matched against a passport or visa record; retention is bound to immigration processing windows. *Implication:* risk arises primarily from secondary uses (law-enforcement mining or data sharing) rather than the matching event itself; transparency and deletion clocks are decisive.
- Retail loss prevention: Shops assemble galleries of prior incidents and run live checks on entry. *Implication:* proportionate use requires clear signage, an appeals process for misidentifications, tight access controls, and an outright ban on emotion or "shopper sentiment" inference.
- Municipal watch-lists: City police request live scanning across transport hubs. *Implication:* authorization and reporting should be external and specific (camera X, time Y, list Z), with non-hit deletion and after-action transparency to avoid becoming a de-facto people-tracking system.

These vignettes illustrate the spectrum from low-risk, privacy-preserving verification to high-risk, ambient identification. They also show how much hinges on purpose, retention, and independent oversight rather than purely on algorithmic prowess.

## 3. Analysis

### 3.1. Mass Surveillance and Function Creep

Problem framing**:** Live or near-live identification in public spaces compresses the cost of tracking to near zero. When outputs are fused with location histories or commercial profiles, authorities and firms can infer habits, associations, political participation, and health status. Even if single images are mundane, aggregates support powerful inference. Function creep—repurposing data beyond the original lawful basis—typically arises through incremental feature additions and vendor integrations rather than explicit policy shifts.

Exemplars linked to discrete risks and remedies:

1. **City-scale camera meshes used for live 1:N search.** Many global cities operate dense CCTV networks with capability upgrades that allow retrospective and, in some cases, live facial matching (Lu, 2022).
   *Risk:* covert tracking chills lawful assembly and speech; non-hit footage becomes a de-facto movement log.
   *Governance gap:* general "crime prevention" mandates can be stretched to warrantless live identification.
   *Remedies:* independent authorization for any live 1:N activation; event-bound targeting and short activation windows; deletion-by-default for non-hits within minutes; mandatory Data Protection Impact Assessments (DPIAs) and conspicuous signage with public reporting of activation counts.
2. **Integrated control systems in authoritarian contexts:** Highly integrated systems have reportedly linked payments, access control, and public shaming for minor offences (Mozur, 2018; Ng, 2020).
   *Risk:* biometric outputs feed coercive social sorting; redress is illusory.
   *Governance gap:* national-security exemptions swallow privacy rules.
   *Remedies:* bright-line bans on social-scoring and punitive emotion inference; human-rights impact assessments; export controls for high-risk deployments where due-process guarantees are absent.

3. **Commercial tagging, sentiment, and loss-prevention:** Retailers and platforms have trialed FRT for shoplifting deterrence, VIP recognition, and auto-tagging (Walmart, 2019; Hill & Mac, 2021; Scarcella, 2024).
   *Risk:* opaque cross-context profiling and consent fatigue; potential discriminatory flagging. *Governance gap:* bundled consent in terms of service; weak transparency obligations. *Remedies:* informed, unbundled opt-in; accessible logs of where and why a face was processed; third-party audits; simple consumer erasure and access rights; and a prohibition on emotion inference for commercial decision-making.
4. **Campus and workplace monitoring:** FRT has been introduced for attendance and access control.
   *Risk:* coercive "voluntariness" where opting out is costly; spill-over into productivity scoring and disciplinary decisions.
   *Governance gap:* employer interests often trump employee privacy; notice without genuine choice.
   *Remedies:* narrowly defined purposes; meaningful non-FRT alternatives; labour-law oversight; and a bar on emotion or behavioral inference for HR outcomes.

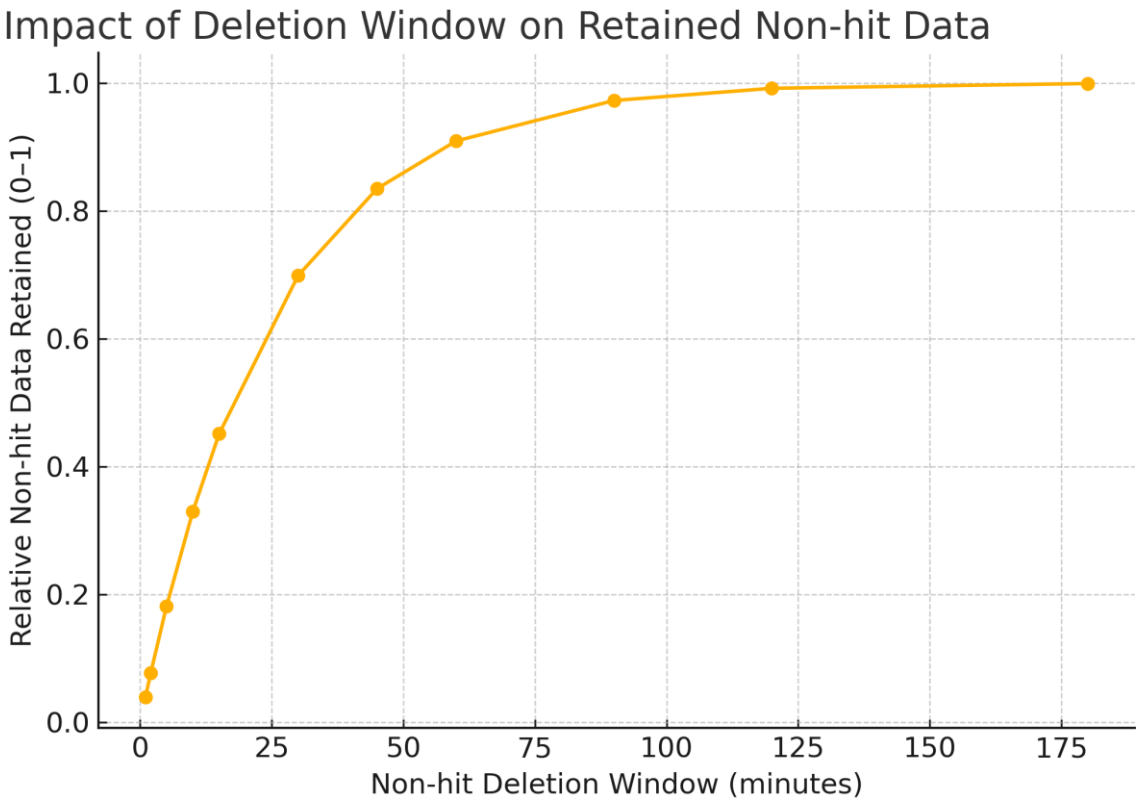


**Figure 1**: Effect of non-hit deletion windows on relative retained data volume (illustrative). Shorter windows materially reduce residual data.

Design and policy levers**:** Biometric data are typically treated as a "special category", requiring explicit consent or demonstrable public interest, strict retention limits, and security controls. Yet many frameworks lag on *operational specifics*: when live identification is permissible; what counts as "necessity and proportionality"; how non-hits must be deleted; and how to audit vendor claims. PETs—on-device matching, hashed templates, and secure enclaves—reduce raw-image exposure but cannot cure unlawful scope. Without independent authorization and deletion clocks, surveillance use becomes the default design outcome.

Community transparency mechanisms (practical tools)**:** To translate principles into practice, deploying bodies should maintain: (a) a **public deployment register** listing locations, purposes, activation windows, and the lawful basis; (b) **quarterly activation reports** that enumerate live 1:N

uses, hits, non-hits deleted, and outcomes; (c) **appeals and redress portals** that allow individuals to contest adverse actions and demand review; and (d) **stakeholder panels** including civil-society and technical experts to review DPIAs. These measures are inexpensive relative to system cost and materially increase accountability.
Interim conclusion: Privacy can be protected only when surveillance use-cases are narrowly scoped, time-bound, and independently reviewable, with rapid deletion of non-hits and transparency about activations. Otherwise, function creep is not an aberration but a system property.

### 3.2. Data Vulnerabilities and Biometric Irreversibility

Why biometric breaches are different**:** Passwords can be rotated; faces cannot. A compromised facial template enables persistent re-identification across services—especially if the same template or derived features are reused—creating durable identity risk. The ethical asymmetry is stark: collectors enjoy security and efficiency benefits, while individuals bear irreversible harm.

Normalization via consumer design**:** Modern smartphones use depth-sensing cameras to project structured light and build a high-dimensional template of the user's face for on-device matching. This design has been rightly praised for usability and security, but it also normalises everyday biometric capture and storage. The privacy question is not the geometry of dot-mapping per se; it is whether raw images or templates ever leave the device, whether retention is minimised, and whether third parties can use the template for unrelated purposes (Cipriani, 2020).

Where systems fail (without over-technical detours): The most salient failure points are organisational and architectural rather than algorithmic: centralised repositories aggregating templates from multiple services; permissive retention justified by vague "analytics" needs; weak key management; vendor chains with uneven controls; and cross-border transfers into weaker legal regimes. Documented corporate settlements and regulatory actions around large-scale face datasets illustrate how quickly misuse can scale and how slowly remediation follows (Hill & Mac, 2021; Scarcella, 2024).

Concrete breach scenarios and impacts: Consider three plausible events: (1) **Insider exfiltration** of a regional access-control gallery to a removable drive, later traded on criminal forums; (2) **Cloud misconfiguration** exposing a storage bucket containing raw enrolment images; and (3) **Cross-vendor API abuse** where a downstream analytics provider repurposes templates for model training. In each case, the individual cannot "change" their face; the realistic remedy is to revoke linked credentials, harden alternative factors, and pursue deletion and damages. These are costly, imperfect fixes—hence the emphasis on prevention and strict purpose limitation.

Security-by-design, expressed in privacy outcomes:

- **On-device matching by default.** Templates remain within secure enclaves; only a signed "match/no match" signal leaves the device. *Privacy outcome:* no central raw-image stores to breach.
- **Non-invertible, salted templates.** Even if exfiltrated, templates cannot reconstruct faces. *Outcome:* reduces stalking or persistent tracking risk.
- **Short retention with event-bound scopes.** Non-hits deleted within minutes; hits retained only for defined investigative or service windows. *Outcome:* limits function creep.
- **Federated learning and differential privacy for model improvement.** Aggregates patterns rather than raw samples; privacy noise prevents singling out (Reimsbach-Kounatze, 2023; Al-Tameemi et al., 2024). *Outcome:* reduces training-data exposure while sustaining performance.
- **Cautious use of homomorphic techniques.** Selective encrypted processing can permit matching or analytics without revealing templates, but performance and complexity trade-offs remain (Pasternack, 2022). *Outcome:* promising in narrow tasks; not yet a blanket solution.

Accountability and remedies: Regulators should require explicit, granular consent for consumer uses; prohibitions on repurposing without a new lawful basis; biometric-specific breach notification and

remediation (e.g., paid identity protection, rapid revocation of linked credentials); vendor audit rights; and adequacy or localisation requirements for cross-border transfers. In jurisdictions without comprehensive federal rules, a patchwork of state laws and settlements creates uneven protection; institutions should voluntarily harmonise to the stricter standard to reduce legal risk and to signal trustworthiness.

Interim conclusion: Technical hardening is necessary but not sufficient. The core ethical commitment is accountability for *irreversible* harm, operationalised through purpose limitation, deletion-by-default, audit, and user remedies.

### 3.3 Accuracy and Bias

Observed error patterns and why they matter for privacy: Even small average error rates translate into large absolute numbers of false matches in city-scale deployments. Empirical studies repeatedly show higher false-match rates for certain demographics, particularly women and people with darker skin tones (Buolamwini & Gebru, 2018; Crumpler, 2020). Privacy is implicated because erroneous identification triggers stop, searches, denials of service, or reputational harm; the burden of error is not evenly distributed.

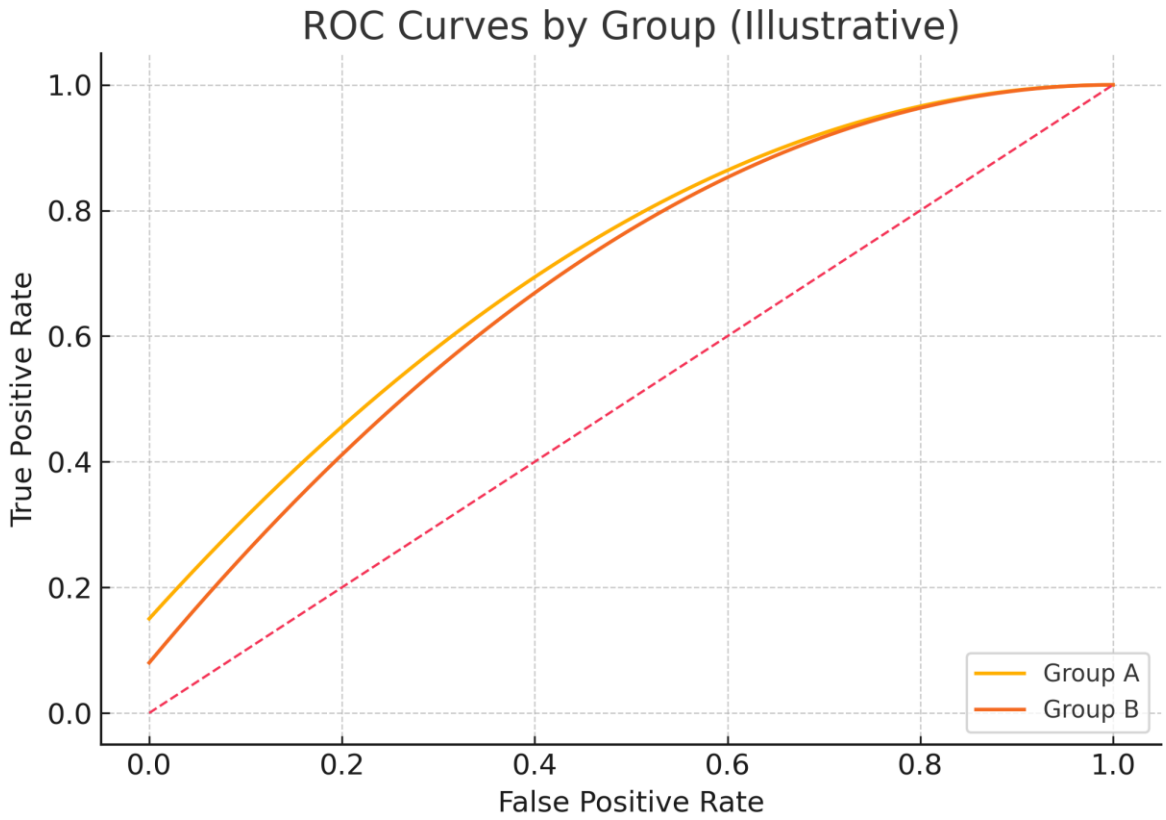


**Figure 2**: Receiver Operating Characteristic (ROC) curves by demographic group (illustrative). Deployers should publish ROC/DET curves disaggregated by group and context.

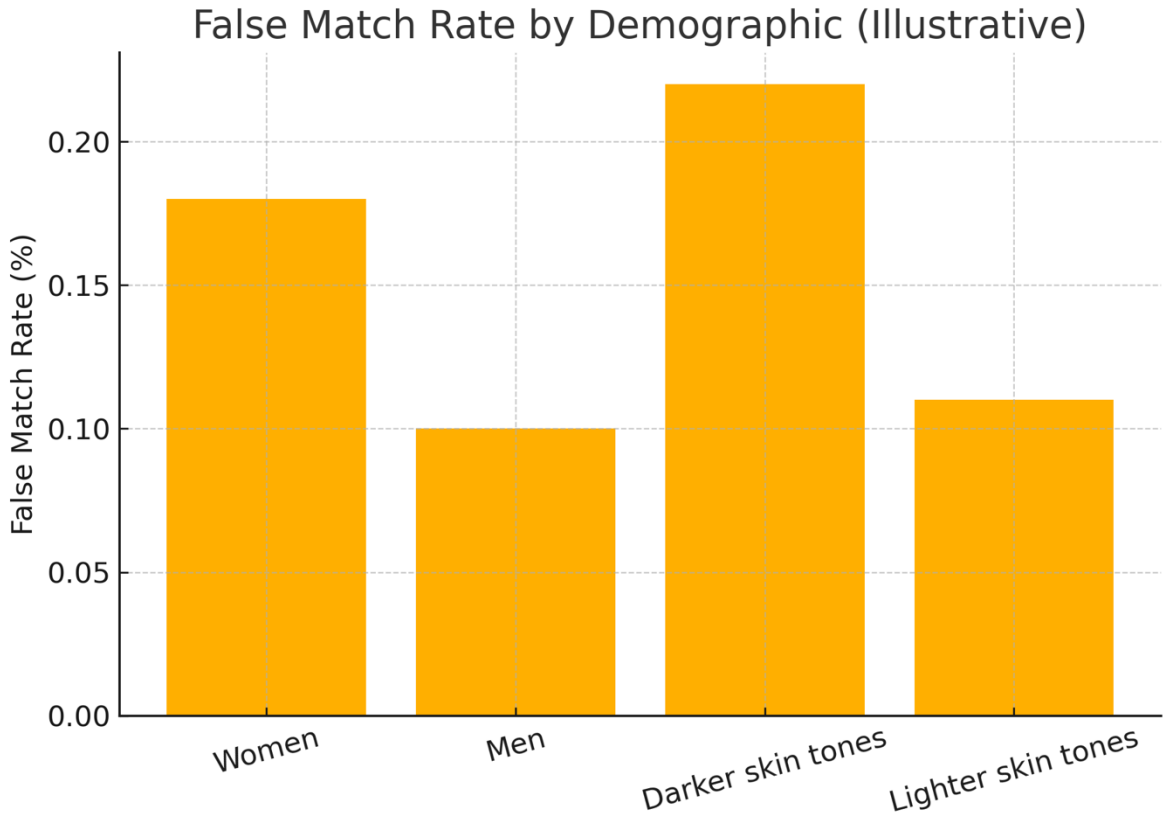


**Figure 3**: False-match rate by demographic (illustrative), reflecting the direction of disparities reported in the literature.

Why bias persists: Training corpora under-represent some groups; labels import historical prejudice; operational settings (thresholds, camera placement, lighting) differentially degrade performance; and domain drift (masks, angles, occlusions) shifts accuracy away from lab benchmarks. These are addressable design and governance choices.

Quantitative intuition for scale: Suppose a city runs live identification on 1 million faces per day with a false-match rate of 0.1% at the chosen threshold. That still yields **1,000 false matches daily**; if error rates are twice as high for a particular demographic group, that group shoulders a disproportionate share of intrusive checks. This arithmetic illustrates why high aggregate accuracy is not sufficient justification for high-stakes deployment without strong procedural safeguards.

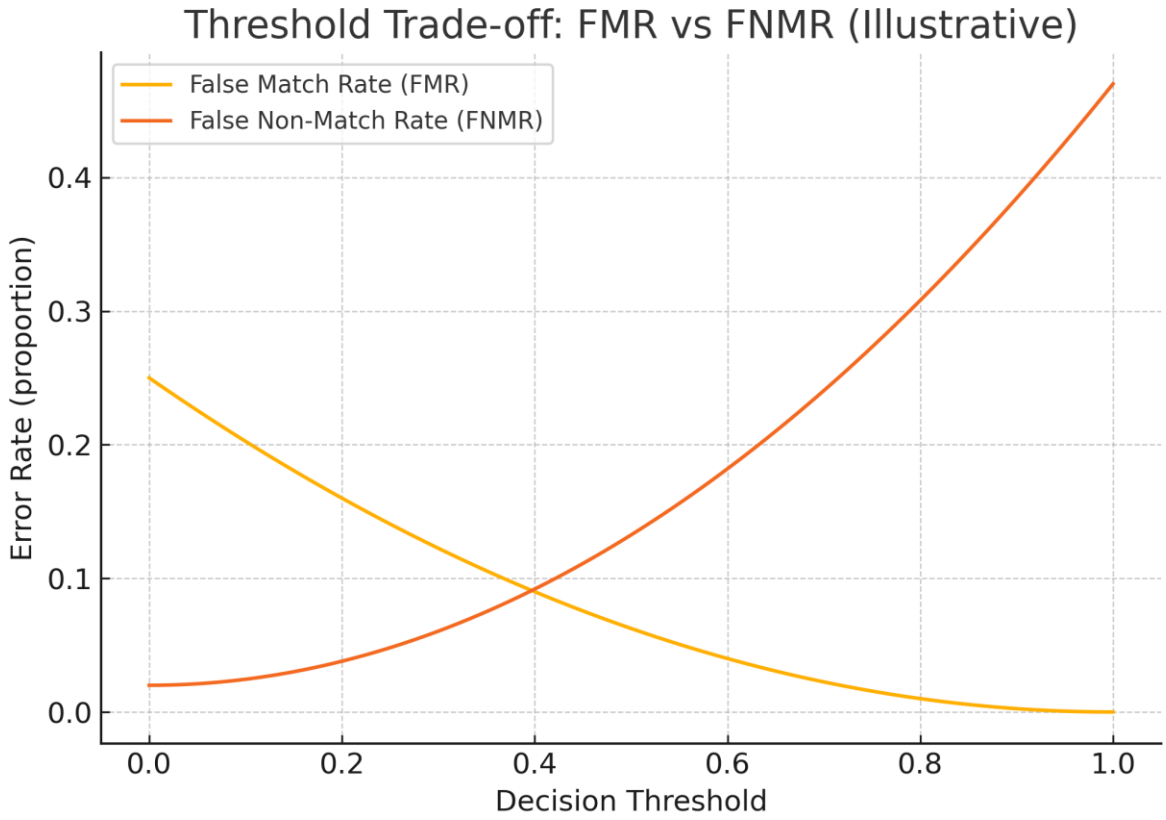


**Figure 4**: Illustrative decision-threshold trade-off between false-match rate (FMR) and false-non-match rate (FNMR). Chosen operating points should be justified and published.

Mitigations with evidence or strong plausibility:

1. **Representative, refreshed datasets** with demographic coverage analysis and documented collection ethics.
2. **Published, testable metrics disaggregated by group and context** (ROC curves, false-match/false-non-match rates at relevant thresholds), enabling external scrutiny and comparability.
3. **Conservative thresholds and "non-match by default"** for low-quality captures; secondary corroboration before any adverse action.
4. **Human-in-the-loop review** with recorded reasoning, especially for high-stakes decisions; no fully automated adverse outcomes.
5. **Independent, repeatable audits** tied to procurement and re-certification; failure to meet group-level performance floors pauses deployment.
6. **Transparent incident reporting** when misidentifications occur, with clear pathways for contestation and remedy.

Fairness–privacy tensions: Improving fairness by collecting more diverse training data can, paradoxically, increase privacy exposure by expanding the volume and sensitivity of the dataset. PETs help navigate the trade-off: federated learning reduces centralisation; differential privacy mitigates singling-out risk; and synthetic augmentation can improve coverage for under-represented conditions without collecting new faces. Nonetheless, PETs are not a panacea; they should be paired with strict data-retention limits and clear deletion guarantees.

Interim conclusion**:** Among the three vectors, accuracy and bias present the most immediate risk of wrongful, uneven harm. Privacy protection therefore requires fairness-first development, operational safeguards, and transparent reporting.

## 4. Conclusion and Call to Action

Judgement**:** FRT can be deployed in ways that respect privacy, but only when three pillars are simultaneously present: (1) narrowly scoped use with strict purpose limitation, independent authorisation for live 1:N searches, and deletion-by-default for non-hits; (2) encryption-first architectures that avoid central raw-image stores and keep templates on-device wherever feasible; and (3) enforceable accountability through group-disaggregated performance reporting, audit, user remedies, and public transparency about deployments. Where any pillar is absent, privacy harms are likely rather than hypothetical.

Practical call to action:

- **Developers:** design for on-device matching; publish disaggregated error metrics; disable emotion inference for consequential decisions; ship privacy-preserving defaults and clear data-deletion schedules; document dataset provenance and coverage; and integrate consent receipts into user flows.
- **Policymakers and regulators:** require independent authorisation and logging for live 1:N identification; mandate short retention and non-hit deletion; define biometric-specific breach remedies; set audit requirements and transparency registers for deployments; and prohibit social-scoring and punitive emotion inference.
- **Deploying institutions:** constrain use to necessary, proportional purposes; provide genuine alternatives where power imbalances exist (workplaces, schools); document human review for adverse actions; run annual DPIAs and publish summaries; test deletion timers and breach playbooks through regular drills.

Reflexive note (situating everyday use): In everyday contexts, unlocking a personal phone or authorising a payment via on-device matching can be compatible with privacy when participation is opt-in, templates remain local, and usage is narrowly scoped. By contrast, remote live identification in open public spaces should be exceptional, independently authorised, and time-bound; absent those conditions, a precautionary pause is warranted. This stance preserves convenience while avoiding irreversible harms from biometric misuse.

### 4.3 International policy snapshot

EU/UK (GDPR-style regimes): Biometric data are a special category: explicit consent or a narrow set of public-interest grounds is required; DPIAs are often mandatory; and cross-border transfers must meet adequacy or use standard clauses with appropriate safeguards. Supervisory authorities can levy significant fines for unlawful processing.

United States (patchwork): In the absence of a comprehensive federal privacy law, protection rests on state statutes and litigation. Biometric-specific laws and settlements have shaped platform behaviour (Hill & Mac, 2021; Scarcella, 2024). The result is uneven coverage and substantial compliance uncertainty for multi-state deployments.

Other jurisdictions**:** Some countries have enacted broad surveillance powers with limited transparency or redress, enabling integration of FRT with payments, public displays, and social-order enforcement (Mozur, 2018; Ng, 2020). Where due-process guarantees are weak, bright-line prohibitions on certain uses (e.g., social scoring, punitive emotion inference) are essential.

These snapshots underscore that technical safeguards travel better across borders than legal ones; hence the emphasis on on-device design, deletion clocks, and auditability as portable protections.

## 5. Reference list (Harvard)

Al-Tameemi, A.A., Mathew, S. and Bajaj, K. (2024) 'Privacy Preserving Technologies', *International Journal for Multidisciplinary Research (IJFMR)*, 6(6). Available at: https://www.ijfmr.com/papers/2024/6/30106.pdf (Accessed: 22 April 2025).

Buolamwini, J. and Gebru, T. (2018) 'Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification', *Proceedings of Machine Learning Research*, pp. 77–91. Available at: https://proceedings.mlr.press/v81/buolamwini18a.html (Accessed: 25 December 2024).

Cipriani, J. (2020) 'Learn how to use your iPhone or iPad's face unlock feature', *CNET*. Available at: https://www.cnet.com/tech/services-and-software/the-iphone-and-ipads-face-id-tech-is-pretty-darn-cool-heres-how-it-works-and-how-to-use-it/ (Accessed: 22 April 2025).

Crumpler, W. (2020) 'How Accurate Are Facial Recognition Systems—and Why Does It Matter?', *CSIS Strategic Technologies Blog*. Available at: https://www.csis.org/blogs/strategic-technologies-blog/how-accurate-are-facial-recognition-systems-and-why-does-it (Accessed: 22 April 2025).

Henning, C., దూలం, K., Peralta, D., Gilani, S.Z., Hadid, A. and Dugelay, J.-L. (2024) 'XAIface: a framework and toolkit for explainable face recognition'. Available at: https://xaiface.eurecom.fr/pdf/xaiface_CBMI_2024.pdf (Accessed: 22 April 2025).

Hill, K. and Mac, R. (2021) 'Facebook, Citing Societal Concerns, Plans to Shut Down Its Facial Recognition System', *The New York Times*, 2 November. Available at: https://www.nytimes.com/2021/11/02/technology/facebook-facial-recognition.html (Accessed: 8 August 2025).

Human Rights Research Initiative (2025) 'The Cost of Convenience: Biometric Data Collection and Privacy', *Student Human Rights Research*, 30 January. Available at: https://www.humanrightsresearch.org/post/the-cost-of-convenience-biometric-data-collection-and-privacy (Accessed: 23 April 2025).

Klosowski, T. (2020) 'Facial recognition is everywhere. Here's what we can do about it', *NYT Wirecutter*. Available at: https://www.nytimes.com/wirecutter/blog/how-facial-recognition-works/ (Accessed: 22 April 2025).

KPMG (2021) 'Facial recognition technology: privacy considerations in access control'. Available at: https://assets.kpmg.com/content/dam/kpmg/nl/pdf/2021/services/facial-recognition-privacy-considerations-in-access-control.pdf (Accessed: 22 April 2025).

Lively, T.K. (2021) 'Facial Recognition in the US: Privacy Concerns and Legal Developments', *ASIS Security Management*. Available at: https://www.asisonline.org/security-management-magazine/monthly-issues/security-technology/archive/2021/december/facial-recognition-in-the-us-privacy-concerns-and-legal-developments/ (Accessed: 22 April 2025).

Lu, M. (2022) 'Ranked: The World's Most Surveilled Cities', *Visual Capitalist*. Available at: https://www.visualcapitalist.com/ranked-the-worlds-most-surveilled-cities/ (Accessed: 22 April 2025).

Lyon, D. (2018) *The Culture of Surveillance: Watching as a Way of Life*. Cambridge: Polity Press. Laboratory accuracy figures are typically reported as a single number at a single threshold.

Magocha, T., Mutula, S. and Mabhiza, A. (2021) 'Facial Authentication as a Bank Security Measure in Zimbabwe', *Journal of Economics, Management and Policy*, 1(1), pp. 68–80. Available at: https://journal.ibs.ac.id/index.php/JEMP/article/download/522/477/1326 (Accessed: 22 April 2025).

Maxiom Technology (2024) ‘Facial Recognition Tech: 5 Key Benefits for Healthcare Security’. Available at: https://www.maxiomtech.com/facial-recognition-tech-health-security/ (Accessed: 22 April 2025).

McClellan, E. (2020) ‘Facial Recognition Technology: Balancing the Benefits and Concerns’, *Journal of Business and Technology Law*, 15(2), p. 363.

Mozur, P. (2018) ‘Inside China’s Dystopian Dreams: A.I., Shame and Lots of Cameras’, *The New York Times*, 8 July. Available at: https://www.nytimes.com/2018/07/08/business/china-surveillance-technology.html (Accessed: 22 April 2025).

Najibi, A. (2020) ‘Racial Discrimination in Face Recognition Technology’, *Science in the News* (Harvard). Available at: https://sites.harvard.edu/sitn/2020/10/24/racial-discrimination-in-face-recognition-technology/ (Accessed: 22 April 2025).

Ng, A. (2020) ‘How China uses facial recognition to control human behavior’, *CNET*. Available at: https://www.cnet.com/news/politics/in-china-facial-recognition-public-shaming-and-control-go-hand-in-hand/ (Accessed: 22 April 2025).

Pasternack, A. (2022) ‘Homomorphic encryption could revolutionise privacy—so what is it?’, *Fast Company*. Available at: https://www.fastcompany.com/90782408/what-is-homomorphic-encryption-and-why-is-it-a-privacy-holy-grail (Accessed: 22 April 2025).

Pollard, M. (2020) ‘Even mask-wearers can be ID’d, China facial recognition firm says’, *Reuters*, 9 March. Available at: https://www.reuters.com/article/technology/even-mask-wearers-can-be-idd-china-facial-recognition-firm-says-idUSKBN20W0WL/ (Accessed: 22 April 2025).

Reimsbach-Kounatze, C. (2023) ‘Emerging Privacy-Enhancing Technologies: Current Regulatory and Policy Approaches’.

Scarcella, M. (2024) ‘Meta to settle Texas lawsuit over Facebook facial recognition data’, *Reuters*, 5 June. Available at: https://www.reuters.com/legal/transactional/meta-settle-texas-lawsuit-over-facebook-facial-recognition-data-2024-05-31/ (Accessed: 22 April 2025).

Villaluz, K. (2021) ‘Number of Cameras Across the World Will Reach 45 Billion by 2022’, *Interesting Engineering*. Available at: https://interestingengineering.com/innovation/number-of-cameras-across-the-world-will-reach-45-trillion-by-2022 (Accessed: 22 April 2025).

Walmart uses AI cameras to spot thieves (2019) *BBC News*, 21 June. Available at: https://www.bbc.com/news/technology-48718198 (Accessed: 22 April 2025).

Hazlegreaves, S. (2021) ‘Why lack of encryption is putting public data at risk’, *ComputerWeekly*. Available at: https://www.computerweekly.com/news/252506382/Why-lack-of-encryption-is-putting-public-data-at-risk/101928/ (Accessed: 8 August 2025).