SvedbergOpen
DISSEMINATION OF KNOWLEDGE

International Journal of Artificial Intelligence and Machine Learning

Publisher's Home Page: https://www.svedbergopen.com/

Research Paper Open Access

# AI-Based Biometric Surveillance in Public Spaces: Balancing Privacy Rights and National Security

Dr .Vaishali

Assistant Professor, School of Law, UPES. E-mail: vsingh3803@gmail.com

**Abstract**
Biometric surveillance is used in public spaces for identification purposes, as well as for public and state security, but raises concerns about legality, privacy, mass surveillance, and errors. This study constructs a Legal–Technical Proportionality Matrix to evaluate whether the use of artificial intelligence for biometric surveillance is legally-technically proportionate. A doctrinal-comparative analysis was conducted to examine the relevant privacy, data-protection, and AI governance frameworks. These were tested against technical performance assessments of facial-recognition technology that demonstrated varying degrees of accuracy and applicability. Four scenarios were evaluated using the matrix, resulting in compatibility scores of 90% for a time-bound survey of a missing person, 85% for a focused screening of a watch-list at an airport, 30% for a city-wide constant surveillance and 25% for identifying peaceful protestors. It was concluded that technical accuracy alone cannot overcome legal limitations, and that biometric surveillance should be implemented as a targeted, independently authorized, technically reliable, transparent, time-boxed, and remediable measure.

**Keywords:** Biometric Surveillance; Public Spaces; Privacy Rights; National Security; Facial Recognition; Artificial Intelligence; Data Protection; Proportionality

## 1. Introduction

Biometric scanning has become an essential component of modern-day security. Facial recognition, fingerprints, and iris scans, among others, can be utilized in safety, surveillance, policing and state security. Unlike traditional forms of surveillance, AI-enabled biometric scanning is able to transform images from the environment into information about a person's identity through detecting a face, extracting a biometric template and comparing it against one or more databases. This transforms surveillance into a scalable form of identification and, consequently, poses a more complex level of legal and technology-related risk.

The appeal of creating a biometric surveillance system is clear. A surveillance system of cameras, linked with machine-learning algorithms, can help a search for a missing person, confirm that a certain person of interest has not entered a secure area, create leads and enable security services to react to an imminent danger. However, the same system can be used to identify persons not suspected of any crime, track their whereabouts, associate them with specific locations or links and store their personal data for future references. The question that emerges is not whether such a system is beneficial to national security, but when and to which extent it is reasonable to allow state security to intrude on privacy.

A person cannot expect privacy in a public place. Modern technology allows biometrics to be captured without consent and often without awareness and additional data such as names or watch lists can be used to link what they were doing to their travel, purchases, and more. Where this is done continuously, is used for vague security reasons, or is using a database that contains many people who have no obvious connection to the issue at hand. The issue is further complicated with the technical aspect. Facial-recognition systems are probabilistic systems, which means that the similarity score depends upon an algorithm, operating threshold, image quality, camera position, lighting, database composition and gender. NIST has shown that there are demographic differentials for many face recognition algorithms and error patterns that vary materially across systems and operating conditions (Grother et al., 2019; NIST, 2026). A false positive can turn out to be a legal problem if a machine generated match contributes to questioning, detainment, a search, denial of access or any other coercive

measure. As a result, accuracy, bias, threshold selection and human verification are not only engineering issues, but can affect due process, equality and the proportionality of state action.

The legal domain has started addressing these risks. The GDPR provides for a specific protection of biometric personal data processed for the purpose of uniquely identifying an individual. The European Union Artificial Intelligence Act regulates more strictly real-time remote biometric identification in publicly accessible spaces for law-enforcement purposes and require narrowly defined purpose, appropriate safeguards and prior authorisation in permitted situations (European Union, 2016, 2024). The Supreme Court in Justice K.S. Puttaswamy (Retd.) v. Union of India acknowledged the right to privacy and elaborated a proportionality test for state interference. India's Digital Personal Data Protection Act, 2023 and the Digital Personal Data Protection Rules, 2025 now provide a general digital-personal-data framework, although a biometric-specific category similar to Article 9 of the GDPR is not present (Government of India, 2023, 2025).

The European Court of Human Rights also pointed out the strong connection between facial recognition, privacy, and freedom of speech. In the case of Glukhin v. Russia, the use of facial-recognition technology to identify one of the protesters led to the violation of Article 8 and Article 10 of the European Convention on Human Rights ("Judgment in Glukhin v. Russia", 2023). The cases presented above allow us to conclude that when considering biometric surveillance, a human rights lawyer has to go beyond the mere possibilities of the technology and evaluate whether its use in a particular case follows from a certain legal basis, pursues a legitimate and specific aim, and is therefore necessary, proportionate, and adequately accompanied by technical and organizational measures.

Existing literature has discussed algorithmic bias, facial-recognition auditing and regulation in public spaces (Buolamwini and Gebru, 2018; Raji et al., 2020; Solarova et al., 2023; Jasserand, 2023). However, legal and technical assessments are often separated. A legal test may assume a surveillance system is sufficiently accurate without specifying what technical validation is required, while an AI evaluation may report accuracy without asking whether the deployment itself is lawful or proportionate. This research addresses that gap by developing and applying an integrated Legal–Technical Proportionality Matrix (LTPM) that combines public-law requirements with AI/ML reliability and governance controls.

The present research is guided by the following four questions: (RQ1) Which legal conditions are generic across the major privacy, data-protection and biometric-surveillance authorities? (RQ2) Which AI/ML characteristics materially impact legality and proportionality of biometric surveillance? (RQ3) Can an integrated legal–technical scoring model distinguish narrowly targeted security uses from forms of mass or population surveillance? (RQ4) What governance controls should be incorporated into the design and operation of biometric surveillance systems?

Two propositions were tested during the scenarios. First, the one described as H1, which is the targeted and limited in time biometric surveillance applied to a clearly defined serious threat will have a higher legal-technical compatibility score than the continuous or non-targeted surveillance. Second, as described by H2, the high level of technical performance cannot offset deficiencies in other areas and the overall legal assessment.

## 2. Materials and Methods

This research utilizes an interdisciplinary legal-technical design, combining doctrinal legal analysis, comparative document coding, and a scenario-based AI governance assessment. The methodology was chosen due to the research question's nature, which regards both normatively relevant legal criteria and their technical preconditions in biometric identification. While the research does not aim to assess public opinion or the actual effectiveness of a particular police application, it focuses on examining conditional factors that determine whether a biometric surveillance application may be considered legitimate.

The corpus of law was selected on the grounds of authority, direct relevance to privacy or biometric identification and geographical diversity. GDPR, the EU Artificial Intelligence Act, India's Digital Personal Data Protection Act, 2023, the Digital Personal Data Protection Rules, 2025, the Supreme Court of India's Puttaswamy privacy judgment and Aadhaar judgment, the European Court of Human Rights decision in Glukhin v. Russia, the Council of Europe Guidelines on Facial Recognition and the United Nations report on the right to privacy in the digital age were included due to their relevance and authority. The corpus was supplemented with technical evidence from NIST face-recognition evaluations and peer reviewed papers on AI ethics and regulation. The legal and policy position was checked as of 5 September 2026.

Each legal or policy source was coded according to five dimensions that recurred across the sources: (1) lawful basis and purpose specification; (2) necessity and proportionality; (3) biometric or high risk data identification; (4) oversight, contestability, and remedy; and (5) AI/system governance. Each of the dimensions were coded

on a three-level scale: explicit (E), partial (P), and absent/not specific (A). The coding reflects varying levels of specificity of the requirements instead of a judgment of different jurisdictions. A particular source was coded explicit if the language directly stated the requirement, partially explicit if it could be inferred from more general privacy/data protection language, and absent/not specific if it did not mention the requirement or specify it in any form.

The second module yielded the Legal–Technical Proportionality Matrix (LTPM) containing 5 composite variables, including Legal Basis and Purpose (LB), Necessity (N), Proportionality (P), Data Governance (DG), and AI Reliability and Oversight (AO). Each criterion was evaluated using a 0–4 point scale, where 0 meant that the requirement was not satisfied or incompatible with the deployment scenario, whereas 4 indicated a clearly stated requirement or a reliable safeguard. In between, scores of 1–3 were given if requirements were partially met or a certain degree of protection was achieved.

The general calculation of the compatibility of the LTPM can be represented by the following formulae: LTPM = [(LB + N + P + DG + AO)/20] × 100. The results of the calculation can be evaluated based on a scale that includes 80-100 – strongly justified with continuing safeguards, 60-79 – conditionally justified with the need for additional control measures, 40-59 – high risk with the current level of justification being legally insufficient, and below 40 – legally disproportionate for the regular use of the technology. It must be stressed that the described compatibility score is not supposed to replace the evaluation by the judiciary. Instead, the scale was developed to make the analysis more transparent and ensure that the technical accuracy would not overshadow the legal evaluation.

Four analytically constructed public-space scenarios were used to test the matrix. The first scenario is a time limited search in a railway terminal for a specific missing child using a known reference image, limited field of view, and human confirmation with immediate deletion of irrelevant information. The second scenario involves a focused airport watch list for specific individuals associated with a credible and severe security threat with independent authorization, a defined temporal scope, and human intervention. The third scenario involves continuous facial recognition across an entire city for general public-safety purposes, which employs a broad watch list and extended retention of data. The fourth scenario involves facial recognition identification of participants in a peaceful public demonstration with no specific threat or severe crime envisioned. These scenarios were selected as they involve a progression of narrowly focused to population-level identification.

In order to minimize subjectivity in scoring, the assumptions of the scenarios and the scoring criteria were established before the matrix was developed. The analysis did not involve the use of any inferential statistical tests, since the scenarios are not a sample from an empirical population. The quantitative element is a decision model that has been developed on the basis of qualitative legal argument. The main way in which validity is assured is through triangulation: every score is related to at least one legal principle and one relevant technical or governance consideration.

## 3. Results and Discussion

### 3.1 Comparative Legal Coding

The comparative coding demonstrates a clear convergence in terms of legality, limitations of purpose, necessity, proportionality and safeguards, despite the use of different legal instruments by the involved parties. The EU framework is the most explicit in terms of the intersection between the personal data of a biometric nature and the use of AI. The GDPR grants special protection to biometric data subjected to unique identification, while the AI Act directly regulates the use of remote biometric identification and imposes strict conditions on its real-time use by law enforcement within publicly accessible spaces. This combination renders the justification of a deployment by the mere accuracy and utility of its technological components difficult. The coding comparison is summarised in Table 1.

The Indian constitutional law offers a significant rights-based test. Puttaswamy affirmed privacy as a fundamental right and required legality, legitimate state aim and proportionality of interference with procedural safeguards against abuse. The Aadhaar judgment shows that the Court can accept large-scale biometric processing when connected with a detailed statutory scheme and specific objectives but such decision must not be regarded as general approval of biometric surveillance in streets, transport hubs or demonstrations as it involves different questions of notice, involuntary capture, continuous observation and freedom of movement or association.

The DPDP Act and the 2025 Rules improve India's general data governance environment, including lawful processing, specific purpose, security obligations and institutional enforcement. Unlike the GDPR, however, the Act does not create a separate statute category for biometric data in its core definitions. For public space facial

recognition, this increases the importance of constitutional privacy analysis, sectoral legal authority and deployment-level safeguards rather than assuming that general data-protection compliance alone resolves the surveillance question.

The human-rights authorities add an additional point that surveillance could have an impact on the rights beyond privacy.

Glukhin provides an example of how the use of facial recognition to identify a peaceful protest participant could interfere with both private life and freedom of expression. Similarly, the UN and Council of Europe materials highlight the problem that indiscriminate or poorly regulated biometric surveillance could have. A chilling effect, discrimination and weak accountability are all potential consequences. The legal test must therefore take into account the context in which people are identified, not merely what the security agency has stated.

### 3.2 AI/ML Characteristics that Change the Legal Risk

The technical analysis discusses four characteristics that may significantly affect the risk of biometric surveillance to the law. Firstly, the ability to identify persons from a distance allows collecting personal data without touching and often without consent. Secondly, one-to-many identification differs fundamentally from the binary yes/no matching of one-to-one comparison. For one-to-many identification systems, the entire image gallery or watchlist is searched with the probe image, so the number and quality of candidates retrieved influence the false positive identification risk directly. Thirdly, thresholds are not only technical constants but often policy variables, so adjusting the decision threshold may be a way of regulating the risks by adjusting the number and quality of candidates. Fourthly, performance may vary substantially across different categories of users and images.

NIST's FRVT/FRTE work is important because it demonstrates that face-recognition accuracy depends on the algorithm and application, the data used, and that demographic differentials remain an evaluation concern(Grother et al., 2019; NIST, 2026). The early studies such as the Gender Shades also revealed substantial subgroup disparities in commercial facial analysis, although the gender classification is not directly comparable to the identity recognition (Buolamwini and Gebru, 2018). However, the legal problem is not that every modern algorithm is equal in its bias; rather, the public authority cannot consider the 'facial recognition' as a single uniform technology. The exact model, version, threshold, camera conditions, and operational population must be evaluated prior to its deployment.

A match does not suffice for identification in a criminal case and thus not for use as evidence without human review of image quality, context which supports the decision and the documentation of the decision process. Also it is up to the operators' training to recognize false positive and false negative results. In which case a match that goes on to instigate a coercive action the system's error rate becomes a part of the proportionality calculation in which we also factor in the increased expected harm from a bad decision.

A policy which states data use for a certain security purpose is of no value if the system does not have the technical means to prevent that data's use for other purposes.

### 3.3 Development of the Legal–Technical Proportionality Matrix

The LTPM which we put together using legal requirements and technical risk factors that came out of the source coding phase. Legal Basis and Purpose (LB) looks at what proper laws exist which explicitly allow for the biometric surveillance in question and also that the purpose of the biometric data collection is defined in such a way that it is amenable to oversight. Necessity (N) looks into if the biometric identification is in fact necessary to achieve the goal, or if there is a less intrusive alternative that will do the job. Proportionality (P) looks at the trade off between the security benefits which may come out of biometric identification and the privacy costs, the number of people affected, the level of the threat, and the issue of other rights. The scoring rubric is provided in Table 2 and the decision flow is depicted in Figure 1.

Data Governance (DG) assesses the use of small and relevant watchlists, collection of data on a minimalist basis, protection of templates, defining retention periods, preventing function creep whilst recording access. AI Reliability and Oversight (AO) assesses independent validation, testing against demographic breakdowns, operating threshold control, image quality monitoring, human confirmation, audit logs, explainability of the operational decision and a mechanism to challenge or correct an incorrect match. The five variables deliberately combine legal and technical requirements so that no single high score can dominate the result.A critical design choice is that the matrix does not include 'national security' as an automatic positive factor. Rather, it is the specificity and seriousness of any claimed purpose. This avoids a self-defeating result in which surveillance is considered proportionate by virtue of being a measure for national security. A generic assertion of national security, without specifying a threat, can therefore receive a low necessity/proportionality score even if the AI system is highly accurate.

### 3.4 Scenario-Based Results

Application of the matrix resulted in a clear distinction between targeted and indiscriminate surveillance. The "missing-person scenario" received a score of 18/20 (90%). The scenario was positively evaluated due to the presence of a clearly defined protective purpose, a specific person sought, a limited geographic area, short term and data deletion of unrelated data. The only significant risks associated with the scenario are image quality, incorrect matches and ensuring that the system does not retain the faces of ordinary passengers after the search is over. The scenario scores are displayed in Table 3 and visualized in Figure 2 below.

The targeted airport watchlist received a score of 17/20 (85%). The serious and credible security threat is a factor contributing to the necessity, and independent authorisation, a limited watchlist and human verification enhance proportionality and oversight. The score is slightly lower than that of the missing-person scenario because of the potentially more coercive consequences of a false match and the large number of travellers subject to surveillance at an airport. Technical validation and strict rules of action in the event of a match are therefore of particular importance.

Using continuous citywide surveillance scored 6 out of 20 (30%) — the weakness was not primarily model accuracy. The scenario has an insufficiently narrow purpose, affects a very large population, and is continuous and retains data for a long time. Even if the model had excellent benchmark performance, general use for population monitoring would be hard to justify on grounds of necessity and proportionality. This outcome supports H2.

Identification of peaceful protest participants scored 5/20 (25%), which was the lowest result. This scenario involves low necessity but substantial interference with privacy, expression, and association. According to Glukhin, there is a direct human-right issue involved which requires special caution. In the absence of a concrete threat or offence, identification of participants with facial recognition would create a high risk of chilling any legitimate activity. This score supports hypothesis 1 since the most targeted and time limited scenarios performed substantially better than the general surveillance scenarios.

The difference between the two high-scoring and two low-scoring scenarios does not arise from an assertion that one technology is 'good' while another is 'bad.' The same model could be used for all four scenarios. The results differ due to legality and proportionality of use, which depends upon the purpose, scope, and length of the surveillance, the criteria used to build the watchlists, data-retention practices, and the consequences of matches. The central finding of this research is that responsible biometric surveillance is a property of the whole socio-technical deployment, not just the machine-learning model.

### 3.5 Testing the Research Propositions

H1 is supported by the scenario analysis. The two targeted and time limited scenarios achieved compatibility scores above 80% while the two continuous or rights sensitive population-surveillance scenarios scored below 40%. It is consistent with the requirement in the EU AI Act that permitted real time remote biometric identification for law enforcement must be tied to specific objectives with temporal, geographic and personal limitations. It is also consistent with proportionality reasoning under Puttaswamy and the European human-rights framework.

H2 is also supported. Matrix structure prevents AI accuracy from curing a legal defect. A deployment that scored 4 on the AI Reliability and Oversight category but 0 or 1 in the legal categories of necessity or proportionality would not be "strongly justifiable." This is because technical capability answers the question of whether a system can identify an individual with a particular level of error, whereas law responds to the normative question about whether the state ought to do that in the first place.

The model also has a useful governance implication. Technical teams should not get a legal approval and go off on their own. The legal requirements need to be engineered into the architecture. If a court or the authorising agency allows surveillance to take place only within two hours of a specific terminal, the system should be engineered to shut down outside of that window. If the requirement is that non-matching data be deleted, it should be engineered to delete. If only a specific watchlist of names is legal for surveillance, the operators should not be allowed to quietly substitute a larger gallery. Proportionality is turned from a policy statement into engineering constraints.

### 3.6 Legal Analysis: Privacy, Consent and Public Space

One of the most challenging aspects is the initial assumption regarding the lack of privacy of the majority in public places. This statement is difficult to deny without the use of such methods. One can argue that being in a railway station involves the possibility of being observed by other people, but this may not necessarily mean that one is always subjected to biometric scans, comparisons with a watchlist, and the storage of one's data in

a convenient searchable database. This factor is essential, as well as the level of involvement of the processing of personal data of an individual.

Consent is also a weak foundation for most state-operated public-space surveillance. A pedestrian on a street, in a station or at a protest has little opportunity to negotiate or withhold consent to surveillance without sacrificing access to the public space. For this reason, public authorities have to rely on a clear statutory basis and on public-law safeguards rather than on artificial consent mechanisms: notice remains valuable for transparency where its absence does not undermine a legitimate operation, but should not be conflated with freely given consent.

The national-security justification requires similar discipline: National security is a legitimate state interest but the category includes many different risks of varying degrees of gravity. In that which is to be proportional we should see the concrete issue at hand, its probability and severity, the group which is the target and the requirement of biometric identification. A very broad aim like 'improving security' is not in itself a good enough reason for wide scale identification.

Indian constitutional interpretation at present is of import as privacy which we see through the Puttaswamy lens is a protected liberty which any surveillance that seeks to infringe on it must be proven as a necessary and proportionate to a legitimate aim. While the Aadhaar judgment does indeed open the door for biometrics in a well structured statutory framework it does not do away with the need to question the put forth goals and the design of a different surveillance which as we see India is to do with its own facial recognition regime.

### 3.7 Algorithmic Bias, False Matches and Equality

Algorithmic bias is present when we see error rates that are not distributed equally or issues which are not the same for all groups of people. Buolamwini and Gebru (2018) report that we see large intersectional differences in commercial gender classification systems. Also from NIST we see that what we get out of identity recognition software is that different demographics' error rates will vary according to the algorithm and use. What we should see is that there is put in place specific to each model and it's application correct legal action instead of the very general issue of bias or no bias in facial recognition software.

If what we get is a prompt for a trained professional to take it from there the damage is contained, but if the result of the algorithm is to do something like detention we see a much greater legal issue. The LTPM we put human element and post match processes into the technical score which in turn puts them into play as opposed to2 simply administrative issues.

Fairness assessments should also consider image acquisition. Poor lighting, camera angle, compression, motion blur and occlusion can affect different groups disproportionately in a particular environment. A model that has good performance on controlled benchmark images may not perform equally well on a crowded street video. Agencies should test the actual camera network they will deploy against representative operating conditions and the relevant demographic population prior to launch, and report or independently audit performance where security considerations permit.

### 3.8 Transparency, Explainability and Remedies

Explainability when it comes to biometric surveillance should be taken in an operational sense. Someone who has been matched against needs to be able to understand what system was used, what database it was matched against, what threshold or decision rule was used, whether a human counted the result and other factors. It should be contestable and subject to judicial review.

Auditability is a requirement equally important. The secure log should record the authorised purpose, time and location of deployment, watchlist version, model version, operator actions and disposition of each candidate match. The logs should be protected against alteration and reviewed by an independent authority. Without logs, it may be impossible to determine whether the surveillance has remained within its authorised limits or whether an incorrect identification resulted from a model error, operator error or misuse of the database.

An effective remedy must be possible after misuse or an incorrect match. Depending on the legal system, remedies may include correction or deletion of data, disclosure of the basis for a decision, administrative review, judicial challenge or compensation. A governance system which prevents abuse is preferable, but protection of rights is incomplete if individuals cannot challenge an adverse outcome after it has occurred.

### 3.9 Implications for AI/ML System Design

The research has direct implications for AI/ML engineering. The first is that the scope of the watch list should be considered a configurable legal control rather than an exclusively operational choice. Second, model thresholds should be approved for a specific use case and changes should be logged. Third, systems should support automatic expiry of authorisations and automatic deletion schedules. Fourth, user interfaces should

present candidate matches with uncertainty and should discourage operators from treating a similarity score as a definitive identity decision. The resulting control set is consolidated in Table 4.

Fifth, model documentation should include information about benchmark accuracy, false-match and false-non-match rates, and performance across different demographics; limitations and known shortcomings; and characteristics of the training or evaluation data; and the specific operational environments for which the model has been evaluated. Sixth, agencies should monitor for model drift and shifts in environmental conditions associated with the cameras. Seventh, privacy-enhancing design should aim to limit retention of raw imagery and separate out identity information from general video surveillance, where feasible.

These controls therefore illustrate the value of a legal–AI framework over a ban-versus-adoption binary. Certain narrowly defined applications can be made consistent with strong legal and technical safeguards, while population-level surveillance remains disproportionately coercive even with a highly accurate model. Regulation should therefore target the deployment context, purpose and architecture, as well as the underlying AI technology.

### 3.10 Limitations

The study has four limitations. Firstly, the LTPM is a normative analytical model, not a tool for proven predictive analytics. The reason for this is that the scores represent structured legal-technical reasoning and cannot be used as odds of lawfulness. Secondly, the four scenarios are constructed cases and do not include all the factual circumstances of real-world applications. Thirdly, legal rules vary geographically and are subject to change; thus, a score needs to be re-evaluated using the applicable local law. Fourthly, technical performance is a moving target; therefore, agencies should use fresh independent evaluations rather than relying on the results of benchmark findings to represent a current model's performance.

Further research might aim to confirm the matrix by multi-expert coding, inter-rater reliability testing, empirical case studies illustrating actual implementations and a sensitivity analysis of the weights. Researchers could also adapt the matrix for application to gait recognition and iris recognition and multimodal biometrics or compare different jurisdictions' approaches to retrospective and real-time identification.

## 4. Conclusion

Biometric surveillance in public spaces can serve public safety and national security, but the privacy of citizens should not be sacrificed for the sake of making the country safer. The key legal issue is whether a particular application has clear authority, a defined legitimate purpose, demonstrable necessity and a proportionate relationship between expected security benefit and interference with rights. The key technical issue is whether the AI system is sufficiently accurate, fair, secure and auditable for that specific context.

The research has produced an integrated Legal–Technical Proportionality Matrix necessary to make these requirements come true. The scenario outcome demonstrated the significant difference between the targeted and time-framed application of surveillance and its perpetual and rights-invasive use across populations. The two former examples of a single-watchman scenario for a missing person search and a serious-threat list had compatibility scores close to the maximum, while perpetual citywide monitoring and identification of peaceful protest watchers had scores below the threshold for a routine surveillance scenario. The study has validated both research propositions, demonstrating that technical precision alone cannot justify rights-damaging surveillance.

For all of AI/ML policy and practice biometric surveillance must be designed into purpose limitation, independent authorization, narrow watch lists, time and geography based access restrictions, model by model performance evaluation, demographic impact analysis, human in the loop approval systems, protected audit trails for all processing which is to be made transparent to affected parties, secure data storage practices, and remediation for affected parties. If these elements are not put in place this technology should not be deployed at all. A principled approach to policy does not require us to choose between privacy and security as base values; what is required is the implementation of security which at the same time is accountable to legal, technical and social parameters.

## References

1. Buolamwini, J., and Gebru, T. (2018). Gender shades: Intersectional accuracy disparities in commercial gender classification. Proceedings of Machine Learning Research, 81, 77–91. https://proceedings.mlr.press/v81/buolamwini18a.html

2. Council of Europe. (2021). Guidelines on facial recognition. Consultative Committee of the Convention for the Protection of Individuals with Regard to Automatic Processing of Personal Data (Convention 108).
3. European Court of Human Rights. (2023). Glukhin v. Russia, Application No. 11519/20, Judgment of 4 July 2023.
4. European Union. (2016). Regulation (EU) 2016/679 of the European Parliament and of the Council (General Data Protection Regulation). Official Journal of the European Union, L 119.
5. European Union. (2024). Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). Official Journal of the European Union.
6. Government of India. (2023). Digital Personal Data Protection Act, 2023 (Act No. 22 of 2023). Gazette of India.
7. Government of India, Ministry of Electronics and Information Technology. (2025). Digital Personal Data Protection Rules, 2025.
8. Grother, P., Ngan, M., and Hanaoka, K. (2019). Face Recognition Vendor Test (FRVT) Part 3: Demographic effects (NISTIR 8280). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.IR.8280
9. Jain, A. K., Ross, A., and Prabhakar, S. (2004). An introduction to biometric recognition. IEEE Transactions on Circuits and Systems for Video Technology, 14(1), 4–20. https://doi.org/10.1109/TCSVT.2003.818349
10. Jasserand, C. (2023). The future AI Act and facial recognition technologies in public spaces: Nice to have or strictly necessary? European Data Protection Law Review, 9(4), 430–443. https://doi.org/10.21552/edpl/2023/4/9
11. Justice K.S. Puttaswamy (Retd.) and Anr. v. Union of India and Ors., (2017) 10 SCC 1.
12. K.S. Puttaswamy (Aadhaar) v. Union of India, (2019) 1 SCC 1 (judgment delivered 26 September 2018).
13. National Institute of Standards and Technology. (2026). Face Recognition Technology Evaluation: Demographic effects in face recognition. U.S. Department of Commerce. Accessed September 5, 2026.
14. Office of the United Nations High Commissioner for Human Rights. (2022). The right to privacy in the digital age (A/HRC/51/17). United Nations Human Rights Council.
15. Raji, I. D., Gebru, T., Mitchell, M., Buolamwini, J., Lee, J., and Denton, E. (2020). Saving face: Investigating the ethical concerns of facial recognition auditing. Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society, 145–151. https://doi.org/10.1145/3375627.3375820
16. Solarova, S., Podroužek, J., Mesarčík, M., Gavornik, A., and Bielikova, M. (2023). Reconsidering the regulation of facial recognition in public spaces. AI and Ethics, 3, 625–635. https://doi.org/10.1007/s43681-022-00194-0

**Tables and Figures**

**Table 1. Comparative coding of selected legal and policy authorities**

| Authority | Legal basis / purpose | Necessity / proportionality | Biometric-specific protection | Oversight / remedy | AI / system governance |
|---|---|---|---|---|---|
| GDPR (EU, 2016) | E | P | E | E | P |
| EU AI Act (2024) | E | E | E | E | E |
| Puttaswamy privacy judgment (India, 2017) | E | E | P | E | P |
| Aadhaar judgment (India, 2018) | E | E | E | E | P |
| DPDP Act 2023 (India) | E | P | A | E | P |
| DPDP Rules 2025 (India) | P | P | A | E | P |
| Glukhin v. Russia (ECtHR, 2023) | E | E | E | E | P |

| Authority | Legal basis / purpose | Necessity / proportionality | Biometric-specific protection | Oversight / remedy | AI / system governance |
|---|---|---|---|---|---|
| Council of Europe Facial Recognition Guidelines (2021) | P | E | E | E | E |
| OHCHR Privacy in the Digital Age (2022) | P | E | E | E | P |

Note: E = explicit; P = partial or derived from general principles; A = absent/not specific. The coding is an analytical comparison and is not a ranking of jurisdictions.

**Table 2. LTPM scoring rubric**

| Score | Interpretation |
|---|---|
| 0 | Absent, unlawful, or fundamentally incompatible with the deployment |
| 1 | Weak basis or broad assertion with major unresolved risk |
| 2 | Partial compliance; significant safeguards or evidence still missing |
| 3 | Substantial compliance; limitations remain but are manageable |
| 4 | Clear, specific and strongly safeguarded condition |

Overall LTPM = [(LB + N + P + DG + AO) / 20] × 100.

**Table 3. Scenario assumptions and LTPM results**

| Scenario | LB | N | P | DG | AO | Total | Compatibility |
|---|---|---|---|---|---|---|---|
| S1: Missing-person search at railway terminal | 4 | 4 | 4 | 3 | 3 | 18/20 | 90% |
| S2: Targeted airport watchlist for serious threat | 4 | 4 | 3 | 3 | 3 | 17/20 | 85% |
| S3: Continuous citywide monitoring for general crime prevention | 1 | 1 | 1 | 1 | 2 | 6/20 | 30% |
| S4: Identification of peaceful protest participants | 1 | 1 | 0 | 1 | 2 | 5/20 | 25% |

LB = Legal Basis and Purpose; N = Necessity; P = Proportionality; DG = Data Governance; AO = AI Reliability and Oversight.

**Table 4. Recommended legal and AI/ML controls for public-space biometric surveillance**

| Control area | Minimum requirement |
|---|---|
| Legal authority | Specific statutory power; defined eligible purposes; documented authorisation |
| Scope | Named/defined targets; geographic boundary; start and end time; limited watchlist |
| Necessity | Written explanation of why less intrusive methods are insufficient |
| Proportionality | Assessment of persons affected, threat seriousness, rights impact and expected benefit |
| AI validation | Model/version-specific accuracy; false-match testing; demographic and environmental evaluation |
| Human oversight | No coercive action solely on automated match; trained verification; documented corroboration |
| Data governance | Encryption; access control; purpose limitation; retention schedule; non-match deletion |
| Transparency and audit | Public policy where possible; immutable logs; independent review; periodic reporting |
| Remedy | Correction/deletion process; challenge mechanism; investigation of incorrect matches |

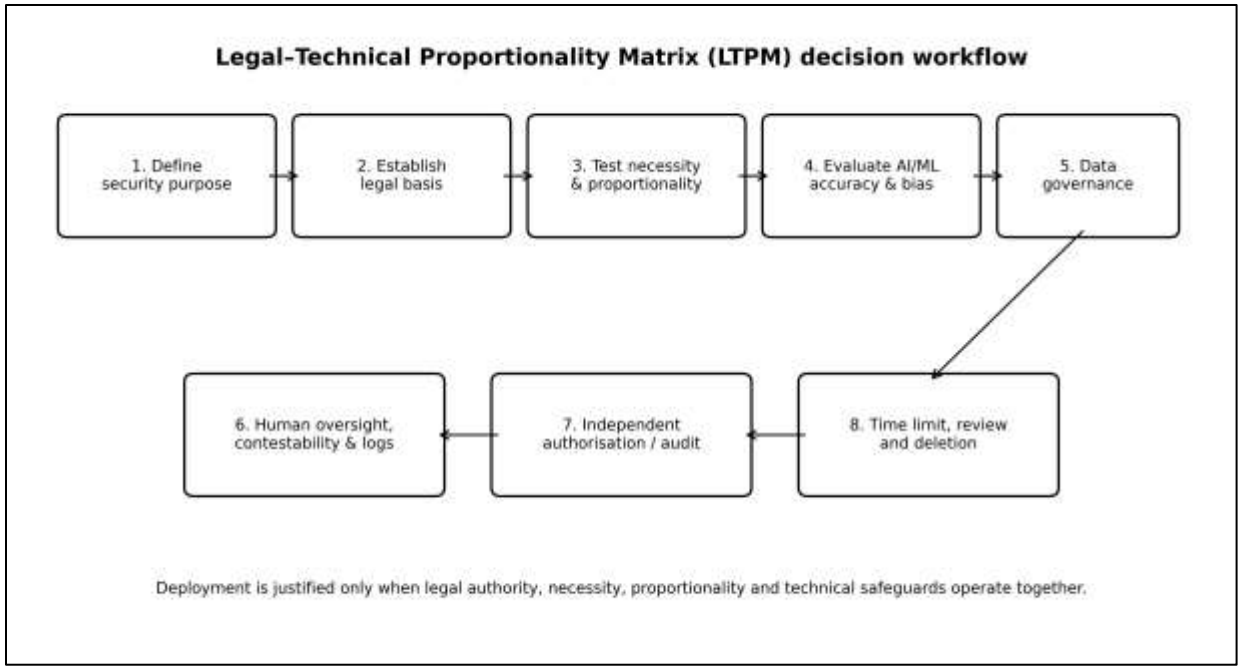


**Figure 1. Legal–Technical Proportionality Matrix (LTPM) decision workflow.**

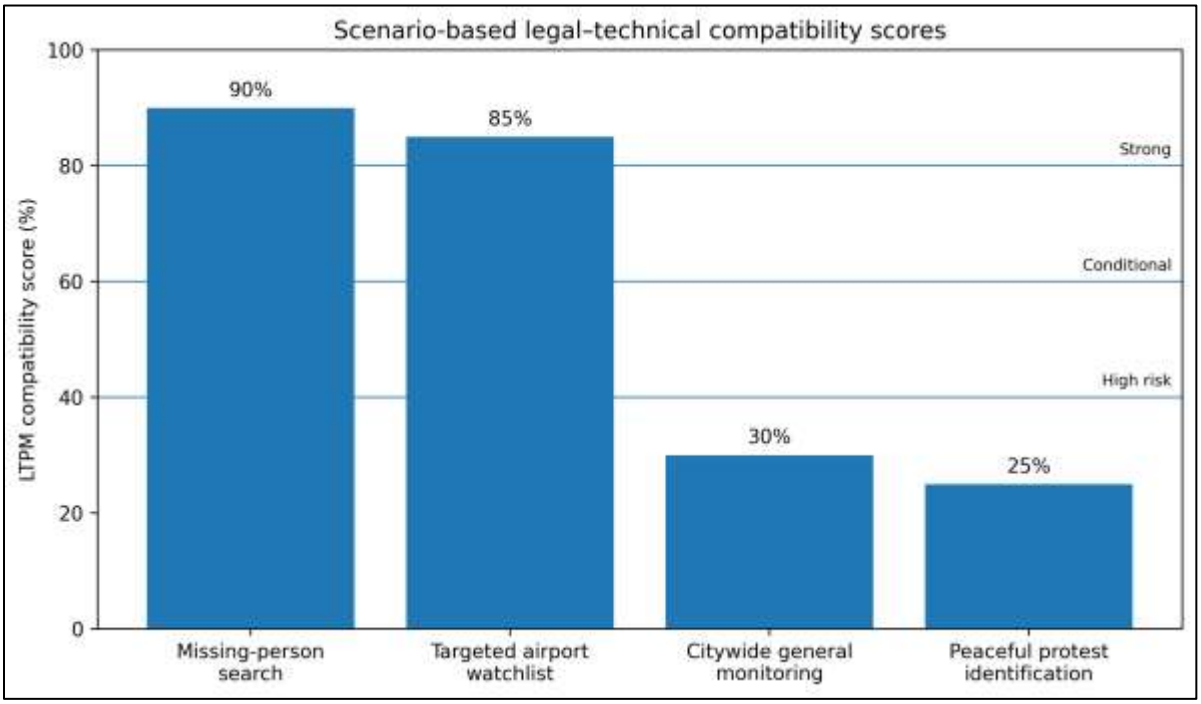


**Figure 2. Scenario-based legal–technical compatibility scores.**