American Journal of Advanced Multidisciplinary Innovation and Research (AJAMIR)

E-ISSN: XXXX-XXXX ● Website: www.ajamir.com ● Email: editor@ajamir.com

Ethical Governance of Emotion-Recognition Technologies in Public Spaces

Dr. Elena Fischer

Associate Professor, University of Amsterdam, Netherlands

Abstract

Emotion-recognition technologies are increasingly positioned as tools capable of interpreting human affect from facial movements, voice characteristics, posture, physiological signals, and other behavioral indicators. Their proposed deployment in transportation facilities, shopping environments, stadiums, civic buildings, health-service entrances, and other publicly accessible spaces creates governance challenges that are qualitatively different from those associated with voluntary consumer applications. Individuals moving through public environments may have limited practical ability to avoid sensing infrastructure, may not understand that emotional states are being inferred, and may be unable to challenge inaccurate classifications.

At the same time, scientific debate continues over the reliability of inferring internal emotional states directly from observable facial behavior, especially when cultural, situational, linguistic, and interpersonal context is insufficiently represented. Barrett and colleagues have emphasized that facial movements cannot consistently be treated as universal diagnostic indicators of particular emotions, strengthening the case for caution when such inferences influence consequential decisions.

This study develops an ethical-governance framework for emotion-recognition technologies operating in public spaces. The research adopts a conceptual-methodological design supported by structured policy analysis and a transparent simulation-based risk assessment. Six deployment settings are evaluated using six governance dimensions: privacy intrusion, meaningful consent limitations, inference-validity uncertainty, demographic and contextual bias, function-creep potential, and accountability weakness.

The analysis indicates that high-density and unavoidable environments, particularly transportation hubs, may generate especially elevated ethical risk because surveillance scale, low practical opt-out capacity, and possible security-related consequences interact simultaneously. The paper proposes a governance architecture based on necessity, proportionality, purpose limitation, data minimization, independent validation, public transparency, human review, contestability, and continuous auditing. The proposed model treats emotion recognition not merely as a technical classification problem but as a form of probabilistic behavioral inference requiring heightened institutional responsibility.

Keywords: emotion recognition, affective computing, public spaces, artificial intelligence ethics, biometric governance, algorithmic accountability, privacy, surveillance, human rights, responsible AI

1. Introduction

Emotion-recognition technology occupies a distinctive position within contemporary artificial intelligence because its objective extends beyond recognizing externally observable objects or identities and attempts to infer aspects of an individual's internal affective condition. Contemporary systems may analyze facial movements, vocal characteristics, text, gestures, posture, gaze, or combinations of behavioral signals to estimate categories such as happiness, anger, fear, stress, engagement, or frustration. NIST distinguishes facial analysis, including attempts to estimate attributes such as emotion, from facial recognition used for verification or identification. This distinction is important because emotion inference can generate consequential judgments even when a person's identity is unknown. In public environments, therefore, the ethical problem is not limited to whether a camera recognizes who a person is; it also concerns whether institutions should computationally interpret what that person is presumed to feel.

Deployment in publicly accessible environments intensifies these concerns because consent is difficult to operationalize meaningfully. A person entering a railway station, government office, hospital entrance, stadium, shopping center, or pedestrian area may theoretically receive signage explaining automated sensing, yet practical avoidance may be impossible. The individual may need transportation, public services, medical care, or access to essential infrastructure. Formal notification cannot automatically be equated with voluntary consent under such conditions. Mohammad's ethics framework for automatic emotion recognition similarly emphasizes that emotion technologies raise questions extending from privacy and social-group effects to the more fundamental issue of why emotion should be automated in a particular context. Public-space deployment must therefore be evaluated not merely through conventional software performance indicators but through principles of necessity, power asymmetry, proportionality, individual autonomy, and institutional legitimacy.

A second concern involves the epistemic foundations of emotion inference. Emotion-recognition systems frequently depend on assumptions that particular facial configurations or observable behaviors provide dependable indications of specific emotional states. However, major reviews of the scientific literature have challenged simple one-to-one mappings between facial movements and discrete emotions. Barrett and colleagues concluded that facial configurations vary substantially according to context, culture, situation, and individual characteristics, meaning that observable facial movement should not automatically be treated as a reliable diagnostic signature of internal emotion. This limitation becomes especially important when public authorities or private operators use probabilistic emotional classifications to trigger security interventions, customer targeting, behavioral monitoring, access restrictions, or other actions.

Regulatory developments demonstrate growing concern about emotionally inferential AI. The European Union's AI Act prohibits emotion-recognition systems in workplace and educational contexts except in limited medical or safety circumstances and treats certain other biometric, emotion-recognition, law-enforcement, migration, and justice applications as high-risk where they are not otherwise prohibited. It also restricts real-time remote biometric identification in publicly accessible spaces for law-enforcement purposes, subject to narrow exceptions. Although emotion recognition in every public-space context is not categorically prohibited under the EU framework, the risk-based architecture illustrates a broader regulatory principle: AI applications affecting fundamental rights require governance that reflects their actual social consequences rather than simply their technical functionality.

Against this background, the present study develops an ethical-governance framework specifically for emotion-recognition technologies in public spaces. Rather than presenting hypothetical scores as empirical findings, the paper uses an explicitly labeled simulation to demonstrate how governance risks may be systematically compared across deployment environments. The principal contribution is a context-sensitive model in which deployment permissibility depends upon the interaction of scientific validity, legitimate purpose, necessity, proportionality, privacy, bias, accountability, data governance, human oversight, and meaningful opportunities for contestation.

2. Objectives of the Study

2.1 Assessment of Ethical and Human-Rights Risks

The first objective is to identify the principal ethical risks created when emotion-recognition technologies operate in publicly accessible environments. Particular attention is given to privacy intrusion, surveillance effects, restricted consent, behavioral manipulation, discriminatory error, contextual misinterpretation, chilling effects, and the possibility that inferred emotional states may influence institutional decisions without adequate evidentiary justification.

2.2 Development of a Context-Sensitive Governance Framework

The second objective is to formulate a governance model capable of distinguishing lower-risk experimental or assistive applications from deployments that create significant threats to autonomy, dignity, equality, privacy, freedom of movement, freedom of expression, and procedural fairness.

The framework is intended to convert broad ethical principles into operational requirements concerning purpose limitation, necessity testing, data minimization, independent validation, impact assessment, retention controls, human oversight, auditability, and remedies.

2.3 Simulation-Based Evaluation of Public-Space Applications

The third objective is to demonstrate how a structured risk-index methodology can be applied across representative public-space scenarios without falsely representing synthetic observations as real field evidence.

Six environments are considered:

- transportation hubs;
- stadiums and event venues;
- shopping malls;
- hospital entrance areas;
- civic service centers; and
- public parks.

The simulation is illustrative and methodological. It demonstrates how institutions could organize ethical assessments before deployment rather than claiming measurement of actual deployed systems.

3. Review of Literature

3.1 Scientific Validity and the Problem of Inferring Emotion

Modern emotion-recognition research emerged partly from affective computing, which sought to create computational systems capable of recognizing and responding to affective information. The field has contributed important research in human-computer interaction, accessibility, health applications, communication, and adaptive systems. Nevertheless, the transition from controlled research environments to consequential public surveillance raises fundamental questions about construct validity.

Barrett, Adolphs, Marsella, Martinez, and Pollak critically examined the proposition that particular facial configurations reliably reveal specific emotional states. Their review demonstrates substantial variation in how emotions are expressed and perceived, challenging the idea that facial movements can consistently function as universal fingerprints of internal affect. Context is therefore not a peripheral variable but part of the interpretive process.

This finding has direct governance implications. An algorithm may correctly identify facial muscle configurations while still incorrectly interpreting the individual's emotional state. A person narrowing the eyes could be responding to sunlight, concentration, fatigue, pain, cultural display rules, interpersonal circumstances, or numerous other conditions. Technical accuracy measured against a labeled training dataset does not automatically establish psychological validity in an uncontrolled public environment.

Mohammad consequently argues that ethical evaluation of automatic emotion recognition must address assumptions embedded in data construction, task formulation, evaluation, privacy, social-group effects, and deployment purposes. His ethics sheet presents fifty considerations designed to encourage developers to question whether automation is justified before optimizing the system itself.

Recent research similarly continues to treat identity and contextual differences as central to public acceptance and ethical evaluation of emotion AI rather than peripheral matters. These developments support the proposition that public-space emotion recognition should not be governed as a conventional camera-analytics application.

3.2 Privacy, Consent, Bias, and Behavioral Power

Emotion-recognition systems create privacy challenges because they transform observable behavior into inferred psychological information. An individual's face may be visible in a public environment, but visibility does not necessarily imply reasonable acceptance of computational emotional profiling.

This distinction can be described as the difference between observation and inference. Ordinary observation permits another person to notice behavior within the limits of human attention and memory. Automated systems can continuously capture thousands of individuals, transform their behavior into structured data, aggregate repeated observations, and connect inferred states with time, location, transaction, or identity records. The scale and persistence of this process alter the nature of the privacy relationship.

The ethical literature has consequently emphasized privacy, consent, and algorithmic bias as recurring concerns in emotion-recognition deployment. Technical privacy mechanisms may reduce certain risks, but privacy-preserving computation alone cannot answer whether the underlying emotional inference is scientifically valid, socially legitimate, or necessary. Research reviewing technical approaches to emotion-AI privacy likewise indicates that technological mitigation must be considered within a broader ethical framework rather than treated as a complete solution.

Bias can enter at several stages. Training datasets may underrepresent demographic, cultural, linguistic, disability, neurodiversity, age, or situational variation. Labels may encode annotator assumptions. Sensor quality may differ with lighting, camera position, facial coverings, assistive devices, or environmental noise. Finally, decision rules may attach unequal consequences to identical error rates.

NIST notes more generally that artificial-intelligence systems can amplify harmful biases and reproduce disadvantages affecting individuals and organizations. In the context of emotional inference, this problem is particularly serious because the underlying construct itself can be uncertain.

3.3 Emerging Governance and Regulatory Approaches

International governance frameworks increasingly emphasize human rights, transparency, fairness, privacy, accountability, and human oversight in the development and use of artificial intelligence. UNESCO's Recommendation on the Ethics of Artificial Intelligence provides a global normative framework addressing the ethical implications of AI and emphasizing human dignity and human rights.

NIST's AI Risk Management Framework similarly provides a voluntary, rights-preserving and use-case-agnostic structure intended to help organizations manage AI risks throughout design, development, deployment, and use. Its broader emphasis on governance, mapping, measurement, and management provides a useful foundation for operationalizing emotion-AI oversight.

European governance has moved further toward binding restrictions. The European Commission identifies emotion recognition in workplaces and educational institutions among prohibited AI practices, except for specified medical or safety uses. It also places non-prohibited biometric identification, biometric categorization, and emotion-recognition systems within high-risk regulatory categories in certain settings.

Related Council of Europe guidance on facial recognition emphasizes lawfulness, supervisory oversight, data quality, reliability, accountability, data security, ethical governance, and protection of data-subject rights. Although facial identification and emotion inference are technically distinct processes, these principles remain relevant because both can involve automated analysis of bodily characteristics in environments where individuals may possess limited bargaining power.

The literature therefore demonstrates an emerging consensus around stronger safeguards, yet an implementation gap remains. Broad principles such as fairness and transparency require context-specific decision mechanisms before they can reliably govern emotional inference in public space. The present research responds by constructing such a mechanism.

4. Research Methodology

4.1 Research Design

The study adopts a conceptual-methodological research design supported by simulation-based scenario analysis. It does not claim that the reported scores represent observations collected from operating surveillance systems.

The research process consisted of four stages.

American Journal of Advanced Multidisciplinary Innovation and Research (AJAMIR)

E-ISSN: XXXX-XXXX ● Website: www.ajamir.com ● Email: editor@ajamir.com

First, recurring ethical concerns were identified from peer-reviewed emotion-recognition literature and recognized AI-governance frameworks. These concerns were consolidated into measurable governance dimensions.

Second, six public-space deployment scenarios were selected to represent variation in crowd density, ability to avoid monitoring, institutional power, security implications, and sensitivity of the environment.

Third, each scenario was assigned an illustrative score from 0 to 100 across the defined governance dimensions.

Fourth, a weighted composite ethical risk score was calculated to demonstrate how a pre-deployment governance assessment could function in practice.

The methodology is intentionally transparent so that future empirical studies may replace simulated scores with audited field measurements, public surveys, regulatory assessments, technical validation data, or independent impact-assessment evidence.

4.2 Ethical Risk Dimensions and Weighting

Six dimensions were incorporated into the model.

Privacy intrusion measures the degree to which behavioral observation could expose or generate sensitive psychological information.

Consent limitation assesses whether individuals can meaningfully refuse or avoid monitoring without losing access to an important place or service.

Inference-validity risk considers the possibility that observable behavior is incorrectly translated into an emotional classification.

Bias and unequal-impact risk examines whether errors or interventions could disproportionately affect particular social groups.

Function-creep risk evaluates the likelihood that information originally collected for one purpose could later be reused for advertising, security profiling, employee assessment, law enforcement, insurance, or other unrelated objectives.

Accountability weakness assesses uncertainty concerning responsibility, explanations, auditability, contestation, and remedies when automated inference causes harm.

Table 1: Ethical-Governance Risk Assessment Framework

Dimension	Weight	Principal Governance Question	Illustrative Indicators
Privacy Intrusion	20%	Does the system create sensitive behavioral or psychological inference beyond ordinary observation?	Continuous monitoring, linkage with identity, persistent profiles
Consent Limitation	15%	Can individuals realistically avoid the system without losing	Opt-out feasibility, alternative routes,

American Journal of Advanced Multidisciplinary Innovation and Research (AJAMIR)

E-ISSN: XXXX-XXXX ● Website: www.ajamir.com ● Email: editor@ajamir.com

Dimension	Weight	Principal Governance Question	Illustrative Indicators
		essential access?	service dependency
Inference-Validity Risk	20%	Is the inferred emotional state scientifically and contextually defensible?	External validation, contextual variables, uncertainty reporting
Bias and Unequal Impact	20%	Are errors or consequences distributed unfairly across groups?	Subgroup performance, accessibility, cultural variation
Function-Creep Potential	15%	Could data or inferences migrate to unrelated purposes?	Secondary use, retention, data sharing, cross-system linkage
Accountability Weakness	10%	Can affected individuals understand, challenge, and obtain remedy?	Audit logs, human review, complaint channels, named controller
Total	**100%**	**Overall ethical permissibility**	**Composite pre-deployment assessment**

Source: Proposed by the author on the basis of the present conceptual analysis and established responsible-AI governance principles. The weighting is illustrative rather than empirically validated.

P = privacy intrusion score C = consent limitation score V = inference-validity risk B = bias and unequal-impact risk F = function-creep potential A = accountability weakness
All component variables are normalized to a 0–100 scale.

4.3 Governance Interpretation and Decision Rules

For methodological demonstration, the composite scores are interpreted using four governance bands. Scores below 40 indicate comparatively limited ethical risk, although standard privacy and accountability controls remain necessary.

Scores from 40 to 59 indicate moderate risk requiring documented safeguards and periodic evaluation.

Scores from 60 to 74 indicate high risk requiring formal impact assessment, independent validation, stringent data minimization, human oversight, and demonstrable necessity.

Scores of 75 or higher indicate very high risk. In such cases, deployment should ordinarily be rejected or suspended unless a compelling legitimate purpose, strong legal basis, independent evidence of effectiveness, and exceptional safeguards can be established.

This structure follows the broader logic of risk-based AI governance, while recognizing that the NIST AI RMF itself is voluntary and use-case agnostic rather than a legal authorization system.

5. Analysis

5.1 Comparative Ethical Risk Across Public-Space Contexts

The simulation suggests that identical emotion-recognition technology can generate substantially different governance concerns depending on the environment in which it is deployed.

Table 2: Simulated Ethical Risk Assessment by Deployment Context

Public-Space Context	Composite Ethical Risk Index	Risk Category	Principal Governance Concern	Suggested Governance Position
Transit Hub	82	Very High	Limited avoidance, mass surveillance scale, security consequences	Presumptive non-deployment unless exceptionally justified
Stadium/Event Venue	76	Very High	Crowd profiling and behavioral intervention	Require strict necessity and independent authorization
Civic Service Center	74	High	Power imbalance and access to public services	Avoid decision-making based on emotional inference
Shopping Mall	71	High	Behavioral targeting and commercial manipulation	Prohibit covert profiling; require genuine opt-in for nonessential uses
Public Park	68	High	Ambient surveillance and chilling effects	Strong presumption against continuous emotional profiling
Hospital Entry Area	66	High	Sensitive context and misinterpretation of distress	Restrict to narrowly defined assistive uses with human review

Note: All values are simulated for methodological demonstration. They are not results of field deployment or measurements of particular commercial systems.

The transportation-hub scenario records the highest simulated index at 82. This result is driven by the convergence of unavoidable infrastructure access, persistent camera coverage, high population throughput, potential integration with security systems, and significant consequences if emotional classifications are interpreted as indicators of suspiciousness or dangerousness.

The stadium scenario also falls into the very-high-risk category. Although event attendance may be more discretionary than using essential transportation, mass monitoring at entrances or inside venues could classify large groups without meaningful individualized consent.

Civic service centers create a different risk profile. Individuals may need access to licenses, welfare assistance, documentation, taxation, migration services, or other public functions. Emotional classification in these environments could influence staff perceptions of cooperativeness, stress, aggression, truthfulness, or credibility. Such applications would be difficult to justify where inference uncertainty could affect substantive rights or access to services.

Shopping environments present strong risks of commercial manipulation. Emotion inference combined with advertising infrastructure could allow organizations to adjust promotions or persuasive content according to presumed vulnerability, frustration, excitement, or anxiety.

Hospital entrance areas receive the lowest score in the simulation, although the value of 66 still represents high risk. Carefully designed assistive applications intended to identify visible distress may appear socially beneficial, yet systems should not substitute inferred emotional states for clinical evaluation.

5.2 Simulated Ethical Risk Graph

Graph 1: Simulated Ethical Risk of Emotion-Recognition Deployment in Public Spaces

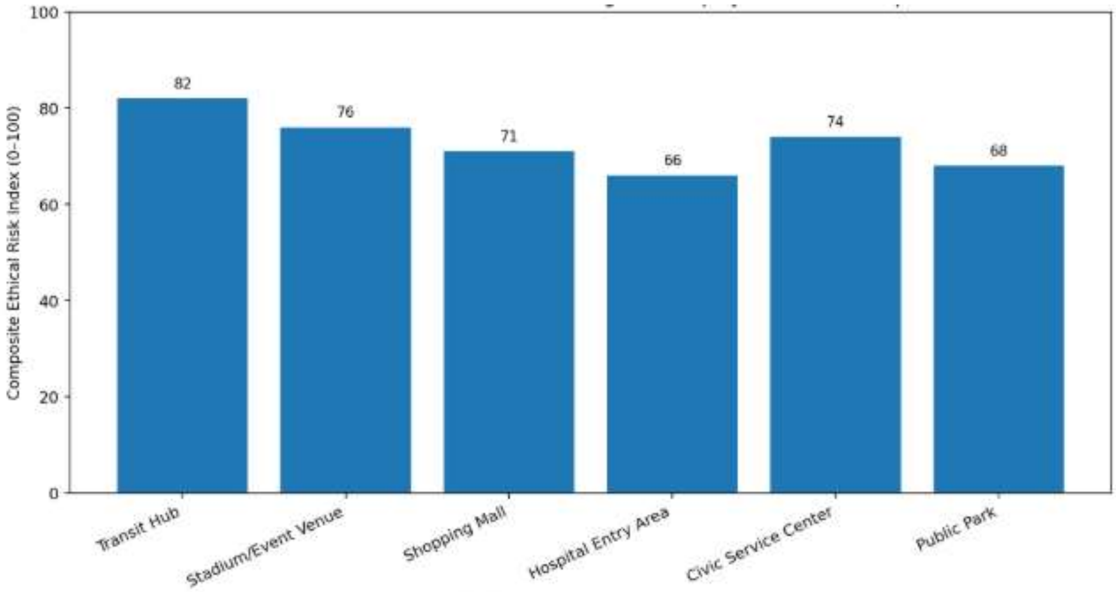

Interpretation: The graph demonstrates that public-space governance cannot rely on a single technological classification. Transit hubs and event venues exhibit the highest simulated risks because emotion inference is combined with scale, low practical avoidance, and potentially consequential institutional responses. Civic service environments also produce substantial risk because citizens can face a strong power imbalance when interacting with public institutions. Even the lowest simulated value remains above the proposed high-risk threshold, suggesting that routine public-space emotional profiling should not be normalized merely because sensing technology is technically available.

5.3 Proposed Governance Architecture

The analysis supports a governance structure containing eight sequential safeguards.

1. **Legitimate-Purpose Test:** The deploying organization must identify a specific public interest or clearly legitimate purpose. General objectives such as improving safety or understanding customers are too broad to justify emotional surveillance.
2. **Necessity Test:** The organization must demonstrate why the intended objective cannot reasonably be achieved through less intrusive alternatives.
3. **Scientific-Validity Test:** Independent evidence must establish that the targeted emotional construct can be inferred with sufficient validity in the intended population and environmental conditions. Accuracy against a benchmark dataset alone should not satisfy this requirement.
4. **Proportionality Assessment:** Expected benefits must be weighed against privacy loss, autonomy restrictions, discrimination risk, chilling effects, and possible downstream consequences.
5. **Data-Minimization Architecture:** Raw facial or behavioral recordings should not be retained merely because storage is inexpensive. Processing should minimize identifiability, granularity, retention, and linkage.
6. **Human Decision Safeguard:** An emotional classification should not independently trigger detention, denial of services, punishment, removal, or another significant adverse action.

6. Discussion

6.1 Emotion Recognition as Inference Rather Than Measurement

A central governance error occurs when emotion recognition is described as though an algorithm directly measures an objective internal quantity. Cameras measure reflected light. Microphones record acoustic signals. Algorithms detect statistical features and generate probabilistic classifications. The transition from observable signal to psychological interpretation involves assumptions.

This distinction matters because consequential governance requires more than predictive correlation. If an algorithm detects a facial configuration associated with anger in a training dataset, this does not establish that every person exhibiting the configuration is angry, that anger implies dangerousness, or that intervention is justified.

The scientific critique developed by Barrett and colleagues therefore has normative significance beyond laboratory psychology. Weak construct validity can become a fairness, accountability, and due-process problem when institutional authority attaches consequences to the inferred label.

Public authorities should consequently require emotion-recognition vendors to disclose uncertainty, validation boundaries, known failure conditions, population limitations, and the distinction between observed behavior and inferred affect.

6.2 Meaningful Consent and the Special Character of Public Space

Public-space surveillance challenges conventional notice-and-consent models because individuals frequently cannot negotiate the terms under which infrastructure operates.

A sign stating that emotion analytics are in use does not necessarily make participation voluntary. A commuter needing a railway service, a citizen visiting a government department, or a patient entering a hospital may have no realistic alternative.

Governance should therefore distinguish three categories of public environments:

- Essential-access environments include transportation, government services, health facilities, and other spaces individuals may need to enter. Emotion recognition should face the highest necessity threshold in such contexts.
- Conditional-access environments include ticketed venues and commercial premises. Organizations may possess greater operational discretion, but consumers still require meaningful alternatives and protection from covert behavioral profiling.
- Open civic environments include streets, squares, parks, and public gathering areas. Continuous emotional inference in these locations can alter the social meaning of public presence by turning ordinary participation into a source of behavioral intelligence.
- Council of Europe guidance on biometric technologies emphasizes lawfulness, privacy, data protection, accountability, data security, and protection of fundamental rights. These principles reinforce the conclusion that mere physical visibility in public should not be interpreted as unlimited permission for psychological profiling.

6.3 From Ethical Principles to Institutional Accountability

Many AI policies articulate principles such as fairness, transparency, privacy, and human oversight. The practical problem is translating them into institutional obligations.

NIST's AI RMF is valuable precisely because it frames risk management as an organizational process extending across governance, system context, measurement, and management rather than as a single technical test. For emotion-recognition systems, this approach should be strengthened through domain-specific requirements.

First, every deployment should identify a named accountable authority rather than distributing responsibility ambiguously among a public agency, system integrator, cloud provider, camera vendor, and model developer.

Second, independent auditing must examine both technical performance and the legitimacy of the use case. A highly accurate system may remain ethically unacceptable if its purpose is disproportionate.

Third, procurement agreements should restrict secondary use. Emotional data collected for crowd-management research, for example, should not quietly migrate into advertising, policing, insurance, employee assessment, or commercial profiling.

Fourth, incident reporting should include false classifications and inappropriate human reliance on model outputs, not merely cybersecurity breaches.

Fifth, affected individuals require accessible mechanisms for complaints, explanation, correction, deletion where applicable, and review of consequential decisions.

The European Commission's current implementation guidance demonstrates that AI governance increasingly distinguishes prohibited practices from high-risk and lower-risk applications. Public-space emotion-recognition governance should follow the same underlying logic: a technology should not receive blanket authorization merely because some applications may be beneficial.

7. Conclusion

Emotion-recognition technologies illustrate a broader transformation in artificial intelligence from systems that classify external objects toward systems that attempt to interpret human intention, affect, disposition, and behavior. When these systems enter publicly accessible environments, the ethical stakes become especially significant because individuals may be observed without meaningful negotiation, emotional classifications may be scientifically uncertain, and institutional actors may possess considerably greater power than those being analyzed.

The present study argues that ethical governance must begin before deployment rather than after harm occurs. Scientific validity, legitimate purpose, necessity, proportionality, privacy, meaningful consent, bias, data minimization, function creep, human oversight, transparency, contestability, and institutional accountability should form an integrated authorization framework.

The simulation demonstrates the practical value of context-sensitive risk assessment. Transportation hubs and large event venues produce very high illustrative risk scores because mass-scale monitoring interacts with limited avoidance and potential security consequences. Civic service centers, shopping environments, public parks, and healthcare entrances generate somewhat different risk configurations, yet all remain ethically significant.

The primary contribution of the paper is therefore a shift from a narrow question—“Can an algorithm recognize emotion?”—toward a more appropriate governance question: “Under what circumstances, if any, should an institution be permitted to infer emotional states from people occupying shared public environments?”

Responsible governance requires evidence that the emotional construct is valid, that the application is genuinely necessary, that less intrusive alternatives have been considered, that risks are proportionate to the anticipated benefit, and that affected individuals retain enforceable rights.

Where these conditions cannot be demonstrated, non-deployment should be regarded not as technological resistance but as a legitimate outcome of responsible AI governance.

8. References

1. Brey, P. A. E. (2004). Ethical aspects of face recognition systems in public places. *Journal of Information, Communication and Ethics in Society, 2*(2), 97–109.
2. Sedenberg, E., & Chuang, J. (2017). Smile for the camera: Privacy and policy implications of emotion AI. *Proceedings of the 2017 IEEE International Symposium on Technology and Society*, 1–8.
3. Crawford, K., & Paglen, T. (2019). Excavating AI: The politics of images in machine learning training sets. *AI Now Institute*.
4. Buolamwini, J., & Gebru, T. (2018). Gender shades: Intersectional accuracy disparities in commercial gender classification. *Proceedings of Machine Learning Research, 81*, 77–91.
5. Raji, I. D., & Buolamwini, J. (2019). Actionable auditing: Investigating the impact of publicly naming biased performance results of commercial AI products. *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, 429–435.
6. Leslie, D. (2020). *Understanding Bias in Facial Recognition Technologies*. The Alan Turing Institute.
7. McStay, A. (2018). *Emotional AI: The Rise of Empathic Media*. SAGE Publications.
8. McStay, A. (2020). Emotional AI, soft biometrics and the surveillance of emotional life. *Philosophy & Technology, 33*, 257–276.
9. Lepri, B., Oliver, N., Letouzé, E., Pentland, A., & Vinck, P. (2018). Fair, transparent, and accountable algorithmic decision-making processes. *Philosophy & Technology, 31*, 611–627.
10. Stark, L., & Hoey, J. (2021). The ethics of emotion in artificial intelligence systems. *Proceedings of the ACM Conference on Fairness, Accountability, and Transparency*, 782–793.
11. Mohammad, S. M. (2021). Ethics sheet for automatic emotion recognition and sentiment analysis. *Computational Linguistics, 48*(2), 239–278.
12. Calvo, R. A., D'Mello, S., Gratch, J., & Kappas, A. (Eds.). (2015). *The Oxford Handbook of Affective Computing*. Oxford University Press.
13. Poria, S., Cambria, E., Bajpai, R., & Hussain, A. (2017). A review of affective computing research based on function-component representation framework. *IEEE Transactions on Affective Computing, 8*(2), 144–168.
14. Picard, R. W. (1997). *Affective Computing*. MIT Press.
15. Andrejevic, M., & Selwyn, N. (2020). Facial recognition technology in schools: Critical questions and concerns. *Learning, Media and Technology, 45*(2), 115–128.
16. Introna, L. D., & Nissenbaum, H. (2009). Facial recognition technology: A survey of policy and implementation issues. *Social Research, 76*(2), 385–418.
17. Garvie, C., Bedoya, A. M., & Frankle, J. (2016). *The Perpetual Line-Up: Unregulated Police Face Recognition in America*. Georgetown Law Center on Privacy & Technology.
18. European Union Agency for Fundamental Rights. (2019). *Facial Recognition Technology: Fundamental Rights Considerations in the Context of Law Enforcement*. European Union Agency for Fundamental Rights.
19. Biometrics and Forensics Ethics Group. (2019). *Police Use of Live Facial Recognition Technology: Ethical Issues*. UK Government.
20. Wendehorst, C., & Duller, Y. (2021). *Biometric Recognition and Behavioural Detection: Assessing the Ethical Aspects of Biometric Recognition and Behavioural Detection Techniques with a Focus on Their Current and Future Use in Public Spaces*. European Parliament