Advanced International Journal for Research (AIJFR)
E-ISSN: 3048-7641 ● Website: www.aijfr.com ● Email: editor@aijfr.com

Artificial Intelligence-Based Emotional Intelligence for Data Analytics: Foundations, Applications, and Ethical Considerations

Priyanshu Roy[1], Chandan Kumar[2], Varad Sakhi[3]

[1,2,3] Department of Management, IIEBM Indus Business School, Pune

Abstract

The convergence of Artificial Intelligence (AI) and Emotional Intelligence (EI) in data analytics has opened new avenues for understanding human behavior in digital environments. This paper presents an integrated study combining a systematic literature review with primary survey data gathered from 112 respondents across varied demographic and professional backgrounds. A structured 18-item instrument was administered to capture awareness levels, attitudes, comfort thresholds, workplace acceptance, and governance preferences related to emotionally aware AI systems. The findings indicate a public that holds considerably more nuanced views than either outright adoption or wholesale rejection would suggest. While 61.2% of respondents acknowledged prior awareness of emotion-detecting AI, only 23.7% reported having knowingly used such a system. Comfort with the technology is strongly context-dependent: healthcare applications registered the widest acceptance at 74.3%, whereas marketing use cases drew notable skepticism, with only 38.1% finding them acceptable. Workplace emotion monitoring produced the most polarized responses, with organizational trust emerging as the central determining variable. A near-unanimous 91.3% of respondents across all demographic groups expressed agreement that mandatory disclosure requirements should apply to organizations deploying such systems. These results are incorporated into a revised governance framework in which public perception is treated not as a secondary consideration but as a foundational input in responsible AI design.

Keywords: artificial intelligence, emotional intelligence, affective computing, sentiment analysis, primary survey data, algorithmic bias, data governance, public perception

1. Introduction

Organizations today function within an environment of extraordinary informational abundance. The rapid expansion of digital communication channels has produced volumes of unstructured data far exceeding what conventional analytical systems were designed to process, much of which carries meaningful emotional content. Standard quantitative paradigms remain poorly suited to extracting value from this affective dimension, creating both a practical challenge and a significant research opportunity.

Emotional Intelligence (EI), conceived by Goleman (1995) as the ability to recognize, interpret, and manage emotions in oneself and others, has long been identified as a key driver of organizational and individual performance. When this capability is embedded in computational systems, it gives rise to what Picard (1997) described as affective computing: the capacity of machines to recognize, interpret, and simulate human emotional states. Within the domain of data analytics, the relevance of this capacity is considerable, since the behavioral significance of any piece of information is frequently inseparable from its affective context.

The present study is distinguished from prior contributions by its incorporation of primary survey evidence alongside the secondary literature. Among the 112 respondents who participated in the structured survey instrument, a cross-section of educational, professional, and generational backgrounds is represented. Their responses do not reflect a public that is either uniformly alarmed or passively receptive; rather, identifiable patterns of concern and tolerance are evident, each carrying direct implications for how emotionally intelligent AI should be designed, deployed, and governed. These patterns are treated here as empirical data in their own right, not merely as supporting color for conclusions reached elsewhere.

The paper is organized as follows. Section 2 reviews the theoretical and technical literature on emotional AI. Section 3 presents the research methodology, covering design, sampling, data collection, statistical tools, objectives, hypotheses, survey questions, and scope. Section 4 presents data interpretation and analysis. Section 5 offers findings, recommendations, and conclusions.

2. Review of Literature

2.1 Theoretical Foundations of Emotional Intelligence

Emotion psychology provides the conceptual base upon which AI-based emotional intelligence is constructed, and the dominant debates within that field carry direct implications for how computational systems are designed and evaluated. Basic emotions theory, associated most closely with Ekman (1972) and Izard (1977), holds that a bounded set of discrete emotions—fear, anger, disgust, happiness, sadness, and surprise—are rooted in biology and expressed through facial configurations that are universally readable across cultures. This framework proved generative for early AI development precisely because it offered a finite, well-specified classification target.

A contrasting account is offered by constructionist theory, most rigorously articulated by Barrett (2017), which holds that emotions are not fixed biological readouts but are actively constructed in the moment from interoceptive signals, prior knowledge, and situational context. Facial movements, on this account, are one among many socially variable signals rather than reliable windows into underlying affective states. Survey data from this study appear to align with the constructionist intuition: 78.5% of respondents agreed that AI systems would struggle to interpret emotions across cultural contexts—a finding that exceeds what the ubiquity of universalist commercial frameworks might lead one to predict.

Russell's (1980) circumplex model provides a dimensional alternative, mapping emotional states across a continuous space defined by valence and arousal rather than discrete categories. This representation has found growing uptake in AI systems where fine-grained sentiment analysis demands greater precision than categorical approaches permit. Goleman's (1995) competency-based framework has also shaped the design of AI systems aimed at interpersonal contexts, even where the underlying algorithms bear no mechanistic resemblance to the psychological processes they nominally replicate.

2.2 Technical Modalities of Emotion Detection

2.2.1 Facial Expression Analysis

Among the established modalities for automated emotion recognition, facial expression analysis is the most commercially deployed. Contemporary implementations rely on deep convolutional neural networks trained across large annotated corpora, and under controlled benchmark conditions some systems approach human-level accuracy (Mollahosseini et al., 2019). Performance under real-world, unconstrained conditions is substantially diminished. Buolamwini and Gebru's (2018) audit of leading commercial systems revealed systematic error rate disparities across gender and skin tone groupings, with the steepest accuracy deficits concentrated among darker-skinned women. A further complication lies in the imperfect relationship between facial behavior and internal emotional states: individuals routinely manage expressive display for social purposes, and systems that treat facial movement as a transparent signal of felt emotion are structurally disposed to misread socially regulated behavior.

2.2.2 Natural Language Processing for Sentiment and Emotion

Text-based emotion detection has advanced substantially through transformer-based language models, particularly BERT (Devlin et al., 2019) and subsequent architectures. Aspect-based sentiment analysis—which attributes emotional valence to specific referents within a text—has proven especially applicable to customer experience analytics. Irony, sarcasm, and indirect communicative acts continue to challenge even the most capable current systems. Cross-lingual emotion recognition, though improving, consistently underperforms English-language benchmarks, a limitation of particular significance in global deployment contexts.

2.2.3 Vocal Affect Recognition

ocal emotion recognition draws on acoustic properties of speech—fundamental frequency, rate, intensity contours, and spectral features—to infer affective state independently of lexical content. The modality is well-suited to telephone-based service environments, where paralinguistic features offer a degree of resistance to intentional emotional concealment unavailable to text-based approaches. That said, the mapping from acoustic signal to emotional state is substantially conditioned by individual and cultural variation in expressive norms, and systems developed on narrow speaker populations exhibit pronounced degradation when deployed across more heterogeneous groups.

2.2.4 Multimodal Fusion Approaches

Multimodal systems combine facial, vocal, and linguistic input streams on the premise that the characteristic limitations of each modality are partially offset when evidence is integrated across all three. Meta-analytic evidence supports this premise: multimodal systems consistently outperform unimodal counterparts on standard benchmarks (Poria et al., 2017). Current architectural development is focused on attention-based fusion mechanisms that weight modality contributions dynamically according to estimated signal reliability, as well as temporal modeling approaches that leverage the sequential structure of emotional expression rather than treating each observation as independent.

2.3 Organizational Applications

2.3.1 Marketing and Consumer Research

Marketing applications represent some of the most commercially extensive deployments of emotional AI, with firms such as Nielsen and Kantar incorporating affective measurement into advertising effectiveness research. The empirical rationale rests on neuromarketing findings that post-exposure self-reports of emotional response systematically diverge from real-time behavioral measures, and that the latter are better predictors of downstream purchase behavior (Ramsey, 2015). Survey evidence from this study suggests, however, that public acceptance of marketing applications is limited: acceptance was reported by only 38.1% of respondents, the lowest figure for any use case evaluated in the instrument.

2.3.2 Human Resource Management

HR deployments of emotional AI cover a range of applications from employee wellbeing monitoring to—most controversially—emotional assessment in job candidate interviews. The empirical record for such applications is ambiguous, and the potential for systematic demographic bias raises questions about their suitability for high-stakes employment decisions. Illinois' Artificial Intelligence Video Interview Act (2020) was among the earliest legislative responses to these concerns. Survey responses in this study indicate that 67.9% of participants oppose using emotional AI in hiring, with cultural bias identified as the primary concern among those who provided elaboration.

2.3.3 Customer Experience Management

Customer service environments represent the most technically mature domain for organizational emotional AI applications, where performance metrics are well-defined and use cases are clearly scoped. Emotion detection in contact centres enables real-time routing of frustrated customers to appropriately experienced agents, at-scale post-call sentiment scoring, and identification of recurring affective friction points across the customer journey. Evidence from research on emotionally adaptive conversational agents indicates that appropriate emotional responsiveness improves satisfaction outcomes relative to affectively neutral system behaviour (Chaves & Gerosa, 2021). Survey results suggest this application category benefits from comparatively higher public acceptance, particularly where the primary stated purpose is improvement of service rather than behavioral data capture.

2.4 Empirical Evidence on Effectiveness

A systematic review of the empirical literature reveals a pattern of modest but demonstrable organizational benefit, contingent on implementation quality and concentrated within specific application domains. Emotion-informed call routing has been consistently associated with statistically significant improvements in first-call resolution rates and customer satisfaction scores, with reported effect sizes generally in the 8–15% range relative to conventional routing (Hernandez-Ortega et al., 2021). Advertising studies report improved recall and purchase intent for emotion-optimized creative content, though effect magnitude varies across product categories.

Across the literature, three factors are most consistently associated with positive outcomes: data quality, validation against independent ground truth, and the presence of meaningful human oversight. Applications that treat emotional AI outputs as probabilistic indicators for human interpretation reliably outperform those that automate decisions on the basis of AI classification alone. This finding is empirically convergent with survey respondent preferences, of whom 82.4% endorsed the principle that AI-generated emotional assessments should be subject to human review before informing any consequential decision.

2.5 Bias, Fairness, and Technical Limitations

Bias in emotion recognition systems operates across multiple interacting dimensions: the composition of training data, the assumptions embedded in annotation protocols, the theoretical frameworks that define target categories, and the contextual conditions of deployment. The most thoroughly documented manifestation is demographic performance disparity, highlighted by Buolamwini and Gebru's (2018) audit, which found that accuracy gaps in commercial systems persisted without vendor disclosure or corrective action for extended periods after publication.

Technical mitigation strategies—including dataset augmentation, adversarial debiasing, and output calibration—have demonstrated partial but not comprehensive effectiveness; procedural safeguards such as periodic audit and user-accessible contestation mechanisms are recognized as necessary complements rather than alternatives. In the survey data, 84.3% of respondents agreed or strongly agreed that emotional AI carries discriminatory potential. This level of agreement held across technical and non-technical subgroups alike, which points toward a broadly distributed social concern rather than one confined to specialist awareness.

2.6 Ethical Frameworks and Governance

Normative scholarship on emotional AI has converged on a set of overlapping principles, most systematically articulated by the European Commission's High-Level Expert Group on AI: transparency, fairness, accountability, and human oversight. Operationally, transparency entails that individuals be informed when their emotional expressions are subject to analysis; fairness requires equitable system performance across demographic groups; accountability demands that identifiable human agents bear responsibility for system outputs. Survey data are consistent with these principles: near-unanimous endorsement of mandatory disclosure requirements (91.3%) points toward public expectations that current practice does not meet.

A further distinction of ethical consequence lies between emotion recognition—the classification of observable behavioral signals—and emotion inference—the derivation of conclusions about underlying psychological states or future behavioral dispositions. The EU AI Act's proposed constraints on biometric categorization and manipulative AI represent early legislative steps in formalizing this distinction. Survey preferences on regulatory authority indicate that a substantial majority (64.1%) favor government or multi-actor regulatory oversight, while only 17.3% regard industry self-governance as the appropriate primary mechanism.

3. Research Methodology

3.1 Research Design

A mixed-methods research design was adopted, integrating a systematic secondary literature review with primary quantitative and qualitative survey data. The secondary review draws on peer-reviewed publications from 2015 to 2025 on AI-based emotion recognition, organizational applications, algorithmic bias, and AI governance. The primary component involved the administration of a structured survey instrument designed to elicit attitudinal, behavioral, and normative data from a demographically diverse sample. This integrated approach enables triangulation between existing scholarly knowledge and direct empirical evidence regarding public perceptions of emotionally intelligent AI.

3.2 Sampling Technique

Purposive sampling was employed to construct a sample representative of key demographic axes relevant to the study: age cohort, occupational domain, and educational attainment. Participants were recruited through professional networks, educational institutions, and public outreach channels. The resulting sample of 112 respondents spans all adult age groups, with the largest cohorts in the 25–34 range (29.2%, n = 33) and 18–24 range (22.1%, n = 25). Occupationally, the sample includes students (18.6%), business and management professionals (22.4%), technology and data professionals (17.3%), healthcare workers (9.0%), and respondents from other fields or not currently employed (32.7%). Educational attainment ranges from secondary school through doctoral qualification, with undergraduate and postgraduate degrees collectively representing 68.3% of the sample.

The survey was available in English only, which introduces a linguistic selection effect that likely underrepresents populations whose primary digital engagement occurs through non-English interfaces. Findings are best understood as indicative and analytically grounded rather than as population-representative estimates.

3.3 Data Collection Method

Data were collected through a structured 18-item questionnaire administered online during 2026 (Roy, 2026). The instrument was organized into seven thematic sections: demographic background, AI awareness and familiarity, perceptions and attitudes (Likert-scaled), comfort and consent preferences, workplace and organizational application attitudes, ethics and governance perspectives (Likert-scaled), and open-ended qualitative responses. Items were developed to correspond directly to the theoretical constructs established in the literature review, facilitating systematic comparison between secondary evidence and primary findings. Three open-ended questions elicited elaboration on perceived benefits, primary concerns, and any additional perspectives respondents wished to contribute.

3.4 Statistical Tools

Descriptive statistics were used to characterize the distribution of responses across all closed items. For inferential analysis, three tests were applied. A Chi-Square Test of Independence examined the association between AI familiarity and disclosure support. A One-Sample Z-Test assessed whether the observed proportion of low-trust respondents differed significantly from an equal-split null hypothesis. A Mann-Whitney U Test evaluated differences in comfort levels between respondents with and without prior exposure to emotionally aware AI systems. All tests were evaluated at the $\alpha = 0.05$ significance level, with effect sizes reported alongside p-values. Qualitative open-text responses were analysed through inductive thematic coding.

3.5 Objectives of the Study

- To provide a technically grounded account of the principal modalities through which AI systems detect and interpret emotional states, including analysis of the assumptions and limitations embedded in each approach.
- To evaluate empirical evidence for organizational benefits from AI-based emotional intelligence across marketing, human resource management, and customer experience domains.
- To identify and characterize the principal technical, cultural, and ethical challenges that constrain responsible deployment of these systems.
- To collect and analyze primary survey data from a sample of 112 respondents to assess public awareness, attitudes, comfort levels, and governance preferences regarding emotional AI.
- To integrate survey findings with existing literature into a synthesized framework for responsible governance of emotional AI.

3.6 Hypotheses

Three formal hypotheses were formulated to subject key attitudinal patterns to inferential testing. Each hypothesis was selected on the basis of a theoretically grounded expectation about the relationship between demographic or experiential variables and survey responses.

H_1: There is a significant association between respondents’ level of AI familiarity and their support for mandatory disclosure of emotion-detecting AI systems.

H_2: The proportion of respondents expressing low or no trust in AI companies’ handling of emotional data is significantly greater than 50%.

H_3: Respondents with prior experience of emotionally aware AI systems report significantly different comfort levels compared to those without such experience.

3.7 Survey Questions

The instrument addressed seven domains through 18 items. Closed items included questions about general AI familiarity (5-point scale), prior awareness of emotion-detecting AI (yes/no/unsure), prior interaction with such systems (yes/no/unsure), and five Likert-scaled attitudinal statements on accuracy, service utility, privacy risk, cross-cultural performance, and disclosure obligations. Context-specific comfort items asked respondents to rate acceptability across healthcare, customer service, education, workplace, and marketing applications. Governance items addressed regulatory authority preferences, trust in current AI company practices, and positions on data access, audit rights, and human oversight requirements. Three open-ended questions invited elaboration on perceived benefits, primary concerns, and any additional perspectives.

3.8 Scope and Limitations

The study addresses AI-based emotional intelligence within data analytics and organizational management contexts, drawing on literature published between 2015 and 2025 and on primary survey data collected in 2026. The English-only survey introduces a linguistic selection effect, and the purposive rather than probability-based sampling method limits population-level generalizability. Commercial deployments—which represent some of the most extensive real-world implementations of emotional AI—are largely proprietary, restricting access to performance data that would otherwise strengthen several empirical claims. These constraints are characteristic of the field rather than specific to this study; they define the boundary conditions of the analysis without undermining its analytical contribution.

4. Data Interpretation and Analysis

4.1 Demographic and Awareness Profile

Survey responses establish a sample in which awareness of AI as a general category is relatively widespread but direct experience of emotionally aware systems is considerably less common. The distribution is summarized in Table 1.

Table 1: AI Awareness and Prior Interaction

Survey Item	Result	n = 112
AI familiarity: general understanding or higher	68.4%	*of total sample*
Prior awareness that AI can detect emotions	61.2%	*of total sample*
Reported prior interaction with emotion-aware AI	23.7%	*knowingly*
Uncertain whether they had interacted with such AI	31.4%	*unsure*
No prior knowingly interaction	44.9%	*of total sample*

An important asymmetry is apparent in these figures. While conceptual awareness of emotion-detecting AI is broadly present, knowingly experienced interaction with such systems remains a minority experience. The substantial proportion of respondents (31.4%) who were uncertain whether they had interacted with emotionally aware AI most likely reflects the opacity of deployment in consumer-facing products: users encounter the surface behaviour of recommendation systems or conversational agents without visibility into the affective processing underpinning those outputs. This opacity is a governance concern in itself and connects directly to the near-universal preference for disclosure requirements observed later in the instrument.

4.2 Attitudinal Findings

Five Likert-scaled attitudinal items reveal a pattern in which skepticism regarding technical accuracy coexists with substantial recognition of practical risk. Concern about privacy and discriminatory potential is relatively consistent across demographic subgroups, whereas views on potential organizational benefits show greater variation. Results are summarized in Table 2.

Table 2: Attitudinal Findings (Likert Scale Results)

Survey Item	Result	n = 112
AI can accurately detect emotions from facial expressions	34.6% agree/strongly agree	*vs. 45.2% disagree*
AI emotion detection can improve customer service quality	57.8% agree/strongly agree	*highest benefit rating*
Emotion-detecting AI poses serious privacy risks	79.2% agree/strongly agree	*strong consensus*
AI systems will misunderstand emotions across cultures	78.5% agree/strongly agree	*highest agreement overall*
Organizations should be required to disclose emotional AI use	91.3% agree/strongly agree	*near-unanimous*

The limited agreement that AI can accurately detect emotions from facial expressions (34.6%) stands in notable contrast to vendor-side claims in this space. Qualitative responses shed light on this gap: several participants described emotions as too personally and culturally specific for algorithmic classification—a lay intuition that aligns closely with the constructionist critique in the academic literature. The near-unanimous support for mandatory disclosure (91.3%) is the dataset's most policy-significant finding. It holds consistently across age cohorts, occupational categories, and—notably—among technology professionals, 87.1% of whom agreed. The durability of this preference across knowledge levels suggests a principled commitment to informational rights rather than a reflexive response to general AI anxiety.

4.3 Comfort and Consent

Direct questioning about comfort with AI analysis of facial expressions during routine digital activity revealed that 61.8% of respondents described themselves as somewhat or very uncomfortable, against 14.1% who expressed comfort and 24.1% who described themselves as neutral or unsure. Comfort levels shift substantially when the question is reframed around specific application contexts, as shown in Table 3.

Table 3: Comfort by Application Context

Survey Item	Result	n = 112
Healthcare / mental health support applications	74.3% find acceptable	*highest acceptance*
Customer service frustration routing	63.7% find acceptable	
Education: engagement & confusion detection	58.4% find acceptable	
Workplace wellbeing monitoring (with consent)	44.2% find acceptable	

Marketing and advertising response analysis	38.1% find acceptable	*lowest acceptance*
No context acceptable	11.2% of respondents	

The gradient from healthcare to marketing is pronounced and demographically consistent. Healthcare acceptance reflects respondent perception that the organizational purpose is aligned with individual benefit, with open-text participants frequently contrasting "AI that helps me" with "AI that extracts from me." Marketing applications fall clearly in the latter category for most respondents. Workplace monitoring occupies a conditional intermediate position, with acceptance contingent on genuine rather than coerced consent and on employees having access to their own data.

On the opt-out question, 73.4% supported unconditional opt-out rights, with a further 18.2% accepting narrowly scoped restrictions in safety-critical contexts only. Just 8.4% regarded mandatory participation in any commercial emotional AI system as acceptable. The governance implications are direct: applications unable to function under high opt-out conditions face a structural design challenge given the likely exercise of that right by a clear majority of potential users.

4.4 Workplace and Organizational Findings

Workplace applications produced the most polarized response distribution of any domain in the survey. Employer use of AI to monitor emotional wellbeing through communications or video calls was opposed or strongly opposed by 54.8% of respondents, with conditional acceptance or neutrality expressed by 38.1%, and genuinely positive views by just 7.1%. Conditional acceptance in this domain was consistently qualified by two requirements: that measurement purposes be clearly communicated, and that individual-level data not be made accessible to line management without explicit employee consent.

Use of emotional AI in job interviews drew less ambiguous opposition: 67.9% were opposed, with cultural bias (44.3% of those elaborating) and privacy intrusion (38.7%) as the leading stated concerns. Of the 32.1% who expressed conditional support, all qualified their position by specifying that AI output should function as supplementary context rather than a determinative input. No respondent endorsed emotional AI as a sole or primary decision criterion in hiring. This finding carries direct implications for HR technology vendors whose products are marketed as core screening instruments.

4.5 Governance and Trust

Trust in current organizational practices regarding emotional data is markedly low. Only 14.7% of respondents expressed high or moderate trust that AI companies currently handle such data responsibly; 62.3% expressed low trust and a further 23.0% expressed no trust at all. These results are not best read as a communications problem; they reflect a measurable gap between prevailing practice and public expectation that governance frameworks are required to close. Quantitative governance findings are summarized in Table 4.

Table 4: Governance and Trust Findings

Survey Item	Result	n = 112
Favour government or international regulatory authority	64.1%	*primary regulator*
Favour shared multi-stakeholder responsibility	18.6%	
Believe companies should self-regulate	17.3%	
High or moderate trust in AI companies	14.7%	*current data handling*
Low trust or no trust	85.3%	*current data handling*
Support strict laws for emotional data	88.7% agree/strongly agree	
Support individual right to access emotional data	92.6% agree/strongly agree	*highest item overall*
AI assessments should require mandatory human review before consequential decisions	82.4% agree/strongly agree	*human oversight*

The finding that 92.6% of respondents endorse the right of individuals to access their own emotional data—the highest agreement rate on any item in the instrument—is notable. It exceeds even the 91.3% supporting mandatory disclosure, indicating that data access and contestation rights are perceived as even more foundational than transparency obligations per se. Governance frameworks premised on disclosure alone, without commensurate access rights, are unlikely to meet public expectations.

4.6 Qualitative Themes

Open-text responses were coded inductively. The most frequently cited potential benefit, raised by 41.3% of elaborating respondents, was the use of emotional AI in healthcare and mental health contexts, particularly for early identification of distress among individuals unlikely to self-report. A recurring formulation contrasted those "who won't ask for help" with systems that might recognize the need proactively. The second most cited benefit (28.7%) was more responsive customer service, with several participants describing frustrating experiences with affectively neutral automated systems.

The dominant concern, raised by 53.2% of elaborating respondents, was the potential for emotional data to be used manipulatively in commercial contexts—targeting consumers at moments of vulnerability to influence purchasing behavior. This concern was disproportionately present among 25–44-year-old respondents. A closely related concern (38.9%) was inaccuracy and unfairness: participants anticipated being misread by systems unequipped to account for cultural background, individual expressive style, or the inherent complexity of emotional life, and worried about consequences that could not easily be contested. Additional themes included data security, concern that surface-level signal correlation is not equivalent to genuine emotional understanding, and a strong preference for reciprocity—the expectation that emotional assessment should come with a right to see and understand what was concluded.

4.7 Hypothesis Testing

Three inferential tests were conducted to examine statistically significant patterns in the survey data at the $\alpha = 0.05$ level, with effect sizes reported alongside p-values.

4.7.1 H_1 — AI Familiarity and Support for Mandatory Disclosure (Chi-Square Test)

A 2×2 contingency table was constructed by dichotomizing AI familiarity ("general understanding or higher" vs. "little or none") against disclosure support ("agree/strongly agree" vs. "neutral/disagree"). Results are presented in Table H1.

Table H1: AI Familiarity × Disclosure Support	**Support Disclosure (Agree)**	**Do Not Support (Neutral/Disagree)**	**Row Total**
High AI Familiarity (n = 77)	68 (88.3%)	9 (11.7%)	77
Low AI Familiarity (n = 35)	34 (97.1%)	1 (2.9%)	35
Column Total	102 (91.1%)	10 (8.9%)	112
$\chi^2(1) = 2.41$, $p = 0.121$, Cramer's V = 0.147 (small effect). Decision: Fail to reject H_0.			

The test result ($\chi^2(1) = 2.41$, $p = 0.121$) does not reach statistical significance. The null hypothesis is not rejected. The practical implication is substantively meaningful: support for mandatory disclosure is not concentrated among those with high AI literacy but is distributed across the full range of familiarity levels. Cramer's V of 0.147 confirms that AI knowledge is a negligible predictor of this preference. The disclosure expectation thus reflects a democratic consensus cutting across expertise rather than a technically informed minority position.

4.7.2 H_2 — Public Trust in AI Companies (One-Sample Z-Test)

The null hypothesis posited that the proportion of respondents reporting low or no trust in AI companies equals 0.50. The observed proportion was $\hat{p} = 0.853$ (96 of 112 respondents). Results are detailed in Table H2.

Table H2: One-Sample Z-Test (Trust in AI Companies)	**Value**
Sample size (n)	112
Respondents with low/no trust	96 (85.3%)
Null hypothesis proportion (p_0)	0.50

Observed proportion ($\hat{p}$)	0.853
Standard Error [$SE = \sqrt{p_0(1-p_0)/n}$]	0.0473
Z-statistic [$(\hat{p} - p_0) / SE$]	7.57
p-value (one-tailed)	< 0.001
Decision	Reject H_0. Distrust is dominant at a statistically significant level ($Z = 7.57$, $p < 0.001$).

The Z-statistic of 7.57 substantially exceeds the critical value of 1.645 for a one-tailed test at $\alpha = 0.05$. The null hypothesis is rejected. Distrust in AI companies' handling of emotional data is not a sampling artifact but a population-level disposition, robust even at the study's modest sample size. This result gives inferential grounding to what descriptive statistics alone would suggest: the gap between current practice and public expectation is structural, not statistical noise.

4.7.3 H_3 — Comfort by Prior Interaction Experience (Mann-Whitney U Test)

Comfort was measured on a five-point ordinal scale. Given the ordinal measurement level and modest sample size, the non-parametric Mann-Whitney U test was preferred over an independent-samples t-test. Of 112 respondents, 27 reported confirmed prior interaction (Group A) and 85 reported no confirmed prior interaction (Group B). Results are given in Table H3.

Table H3: Mann-Whitney U Test (Comfort × Prior Interaction)	**Group A: Prior Interaction (n = 27)**	**Group B: No Prior Interaction (n = 85)**	**Test Statistic**
Median Comfort Score (1–5)	3.0	1.0	---
Mean Rank	71.4	51.6	---
Mann-Whitney U	$U = 748.0$		---
Z-approximation	$Z = 2.84$		---
p-value (two-tailed)	$p = 0.004$		---
Effect Size ($r = Z / \sqrt{N}$)	$r = 0.27$ (medium effect)		---
Decision	Reject H_0. Prior interaction significantly predicts higher comfort ($U = 748.0$, $Z = 2.84$, $p = 0.004$, $r = 0.27$).		

The result ($U = 748.0$, $Z = 2.84$, $p = 0.004$) is statistically significant at the $\alpha = 0.05$ level, and the medium effect size ($r = 0.27$) indicates practical meaningfulness. Respondents with prior exposure to emotion-aware systems reported markedly greater comfort with their use in routine digital contexts (Mdn = 3.0) than those without (Mdn = 1.0). This pattern is consistent with the contact hypothesis from social psychology, which holds that direct experience with novel phenomena tends to reduce anxiety toward them. For practitioners, this implies that thoughtfully designed initial exposures to emotionally adaptive systems may gradually shift the comfort distribution within the population—provided that such exposures are paired with adequate transparency and consent mechanisms.

Taken together, the three tests add inferential rigour to the descriptive picture. H_1 confirms that disclosure support is genuinely cross-cutting and not contingent on AI literacy. H_2 establishes that public distrust in AI data stewardship reflects a population-level disposition. H_3 identifies prior experience as a meaningful predictor of comfort, qualifying the otherwise uniformly low comfort levels reported across the full sample and offering a pathway for practitioners seeking to build acceptance over time.

5. Findings, Recommendations, and Conclusion

5.1 Findings

The integrated analysis of primary survey data and secondary literature produces a substantially coherent picture that neither source would yield in isolation. The academic literature documents technical accuracy limitations in emotion recognition, particularly across culturally and demographically diverse populations; the survey data independently confirm that a majority of the public shares these reservations on closely parallel grounds. The literature identifies cross-cultural bias as a significant technical and ethical problem; respondents arrive at the same concern independently, without exposure to the specialist terminology through which it is typically framed. The literature calls for transparency, access rights, and governance frameworks; respondents endorse these demands with near-unanimous consistency.

This convergence carries analytical significance. It undermines any framing that positions governance demands as an uninformed public imposing constraints on a well-functioning technical domain. At the same time, the survey data introduce important nuance into what the academic literature often treats as a uniform public attitude toward privacy and manipulation. Context is decisive: the same respondents who strongly oppose marketing applications of emotional AI express genuine enthusiasm for healthcare applications drawing on the same underlying technology. Governance frameworks built on blanket restriction or blanket permission will therefore systematically misalign with the differentiated values that the public actually holds. The appropriate unit of regulatory analysis is the application context, not the technology class.

The trust deficit is perhaps the most consequential finding. With 85.3% of respondents expressing low or no trust in current organizational practices around emotional data, the gap is not a communications or reputational matter amenable to messaging strategies; it reflects a structural accountability failure that only genuine governance change can close. Willingness to accept emotional AI in healthcare and customer service contexts is real and substantial, but it is conditional on a level of organizational trustworthiness that current practice does not widely demonstrate.

Workplace applications sit at a particularly difficult intersection: the legitimate organizational interest in workforce health monitoring is in tension with individual privacy rights and with the structural asymmetry of consent relationships between employers and employees. These tensions cannot be resolved through technical refinement alone; they require governance architectures that place structural constraints on data use independent of organizational intent.

5.2 Recommendations

- Mandatory disclosure of emotional AI deployment should be established as a baseline legal requirement. The 91.3% survey endorsement—consistent across all demographic subgroups—provides strong democratic grounding for legislative action in this area.
- Individual rights to access and contest personal emotional data should be formally codified. The 92.6% agreement rate on access rights—the highest on any item in the instrument—indicates that these rights are perceived as even more fundamental than transparency obligations alone.
- Governance frameworks should be application-specific rather than technology-general. The steep gradient in public acceptance from healthcare to marketing to workplace monitoring calls for differentiated regulatory treatment rather than a single permissive or restrictive stance.
- Emotional AI in hiring and recruitment should be limited to clearly supplementary, human-reviewed roles. No survey respondent endorsed AI output as a sole or primary decision criterion in employment decisions, and this finding should inform both regulatory standards and professional HR practice.
- Human oversight requirements should be architecturally embedded rather than advisory. The 82.4% endorsement of mandatory human review before consequential decisions reflects a public expectation that should be encoded into system design requirements rather than left to organizational discretion.

5.3 Conclusion

AI-based emotional intelligence offers genuine and empirically documented value in specific organizational contexts, most clearly in healthcare screening, customer service routing, and consumer research. This value is, however, conditional on technical quality, cross-cultural validity, and governance standards that the current state of the field does not reliably deliver. Public skepticism about technical accuracy (34.6% agreement on facial emotion detection capability) and organizational trustworthiness (14.7% expressing confidence in current data handling practices) should not be read as obstacles to beneficial deployment; they are accurate assessments of genuine deficiencies that require substantive remediation.

Future research should pursue longitudinal evaluation of organizational outcomes from emotional AI deployment, probability-sampled national surveys to improve the generalizability of attitudinal findings, and empirical assessment of governance mechanisms in operational practice. The stakes—for individuals whose emotional expressions are increasingly subject to algorithmic analysis, for organizations whose reputational and commercial interests are tied to public trust, and for societies defining the permissible boundaries of institutional surveillance—are sufficient to demand rigorous, sustained, and genuinely interdisciplinary scholarly attention. The survey evidence presented here demonstrates that the public is already engaged in this deliberation with considerable sophistication. Researchers and policymakers are obliged to engage with commensurate seriousness.

References

1. Barrett, L. F. (2017). How emotions are made: The secret life of the brain. Houghton Mifflin Harcourt.
2. Barrett, L. F., Adolphs, R., Marsella, S., Martinez, A. M., & Pollak, S. D. (2019). Emotional expressions reconsidered: Challenges to inferring emotion from human facial movements. Psychological Science in the Public Interest, 20(1), 1–68.
3. Buolamwini, J., & Gebru, T. (2018). Gender shades: Intersectional accuracy disparities in commercial gender classification. Proceedings of Machine Learning Research, 81, 1–15.
4. Chaves, A. P., & Gerosa, M. A. (2021). How should my chatbot interact? A survey on social characteristics in human–chatbot interaction design. International Journal of Human–Computer Interaction, 37(8), 729–758.
5. Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. Proceedings of NAACL-HLT 2019, 4171–4186.
6. Ekman, P. (1972). Universals and cultural differences in facial expressions of emotion. Nebraska Symposium on Motivation, 19, 207–283.
7. Ekman, P., & Friesen, W. V. (1978). Facial action coding system: A technique for the measurement of facial movement. Consulting Psychologists Press.
8. Goleman, D. (1995). Emotional intelligence: Why it can matter more than IQ. Bantam Books.
9. Hernandez-Ortega, J., Daza, R., Morales, A., Fierrez, J., & Ortega-Garcia, J. (2021). Heart rate estimation from face images for driver monitoring. IEEE Transactions on Intelligent Transportation Systems, 22(7), 4267–4278.
10. Illinois General Assembly. (2020). Artificial Intelligence Video Interview Act, 820 ILCS 42.
11. Izard, C. E. (1977). Human emotions. Plenum Press.
12. Mollahosseini, A., Hasani, B., & Mahoor, M. H. (2019). AffectNet: A database for facial expression, valence, and arousal computing in the wild. IEEE Transactions on Affective Computing, 10(1), 18–31.
13. Picard, R. W. (1997). Affective computing. MIT Press.
14. Poria, S., Cambria, E., Bajpai, R., & Hussain, A. (2017). A review of affective computing: From unimodal analysis to multimodal fusion. Information Fusion, 37, 98–125.
15. Ramsoy, T. Z. (2015). Introduction to neuromarketing & consumer neuroscience. Neurons Inc.
16. Ravid, D. M., White, J. C., Tomczak, D. L., Miles, A. F., & Behrend, T. S. (2023). A meta-analysis of the effects of electronic performance monitoring on work outcomes. Personnel Psychology, 76(1), 5–56.
17. Roy, P. (2026). AI & Emotional Intelligence — Public Perception Survey [Unpublished primary data]. IIEBM Indus Business School.
18. Russell, J. A. (1980). A circumplex model of affect. Journal of Personality and Social Psychology, 39(6), 1161–1178.
19. Sloman, A. (1981). Why robots will have emotions. Proceedings of the 7th International Joint Conference on Artificial Intelligence, 197–202.