KYRIAKIDES GEORGOPOULOS
LAW FIRM



AI & Technology Law

Between a Smile and a Ban: Where the AI Act Draws the Line on Emotion Recognition

November 23, 2026

Between a Smile and a Ban: Where the AI Act Draws the Line on Emotion Recognition

BY IRENE KYRIAKIDES, ELECTRA RETZIKA

1. Introduction

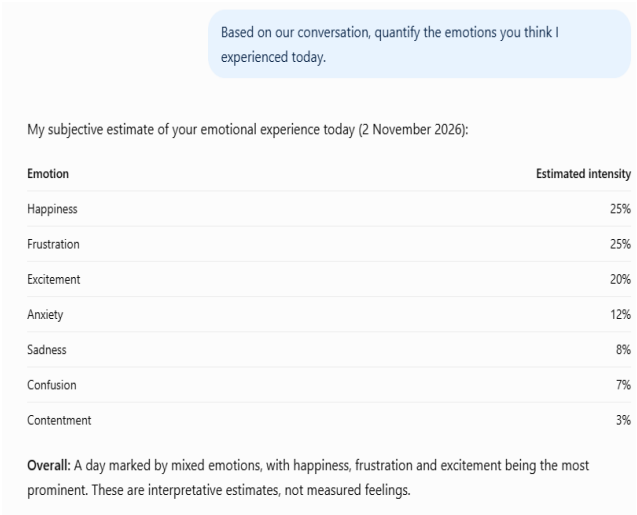

"Happiness, frustration, excitement."

According to ChatGPT, these were the predominant emotions I had felt today, judging only by the way I communicated with it. Now imagine the same assessment being made about a student's emotional state in a classroom or an employee's at work.

2. How Emotion Recognition Works

Emotion recognition systems draw inferences about a person's emotional state in four steps:

(1) Data collection: Cameras, microphones, text inputs and biometric sensors capture facial expressions, voice, word choice, typing patterns and physiological signals.

(2) Feature extraction: Algorithms identify cues linked to emotions, such as facial muscle movements, vocal pitch and emotionally charged language.

(3) Emotion classification: Machine learning models trained on existing datasets sort these cues into categories such as happiness or sadness.

(4) Adaptive Response: The system adjusts its behavior to the emotion it detects. For example, it may offer encouragement when it detects frustration or interactive content when it detects boredom.

3. How AI Facial Emotion Recognition Models are Trained

These models are typically trained on labelled image datasets such as FER2013 and AffectNet, in which each face is tagged with an emotion category, such as happiness, sadness, surprise and neutrality.



Another approach is to film participants while visual cues and music prompt specific facial expressions, and then split the footage into individual image frames. Tools such as the open-source Python library dlib can map 68 key points on the face, allowing a system to analyze specific facial regions, such as the lower half of the face

when a VR covers the eyes.

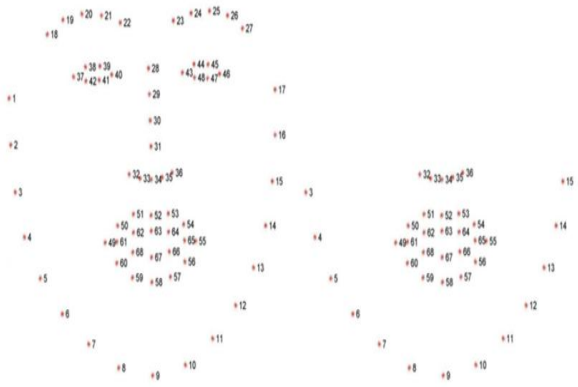

Some models also use the Facial Action Coding System (FACS), which breaks facial movements into individual Action Units (AUs), that are then linked to specific emotions.

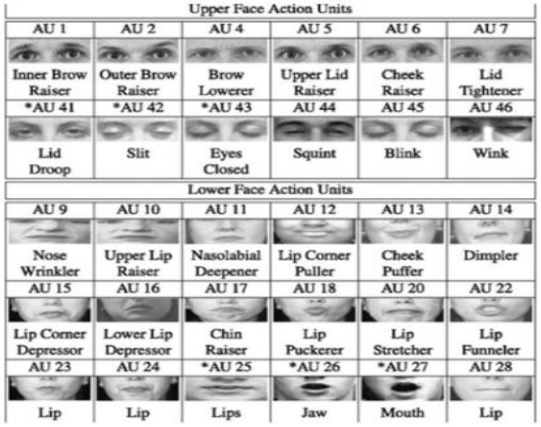

Upper Face Action Units					
AU 1	AU 2	AU 4	AU 5	AU 6	AU 7
Inner Brow Raiser	Outer Brow Raiser	Brow Lowerer	Upper Lid Raiser	Cheek Raiser	Lid Tightener
*AU 41	*AU 42	*AU 43	AU 44	AU 45	AU 46
Lid Droop	Slit	Eyes Closed	Squint	Blink	Wink
Lower Face Action Units					
AU 9	AU 10	AU 11	AU 12	AU 13	AU 14
Nose Wrinkler	Upper Lip Raiser	Nasolabial Deepener	Lip Corner Puller	Cheek Puffer	Dimpler
AU 15	AU 16	AU 17	AU 18	AU 20	AU 22
Lip Corner Depressor	Lower Lip Depressor	Chin Raiser	Lip Puckerer	Lip Stretcher	Lip Funneler
AU 23	AU 24	*AU 25	*AU 26	*AU 27	AU 28
Lip	Lip	Lips	Jaw	Mouth	Lip

Before training, facial images are often artificially varied, for example by rotating or zooming them, so the training data is more diverse.

The dataset is then split into a training and a test set. The model learns patterns from the training images and classifies new, unseen images into predefined emotion categories (shown in different colors below).

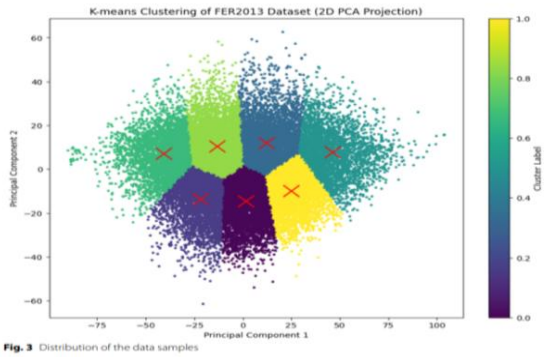

Fig. 3 Distribution of the data samples

The model's performance is measured by accuracy and loss curves for both the training and test sets. Rising accuracy and falling loss indicate better classification. Tracking both also reveals "overfitting", where a model performs well on its training data but poorly on new data.

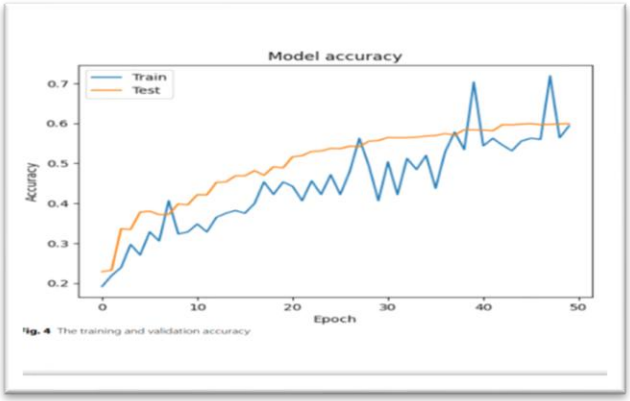

Fig. 4 The training and validation accuracy

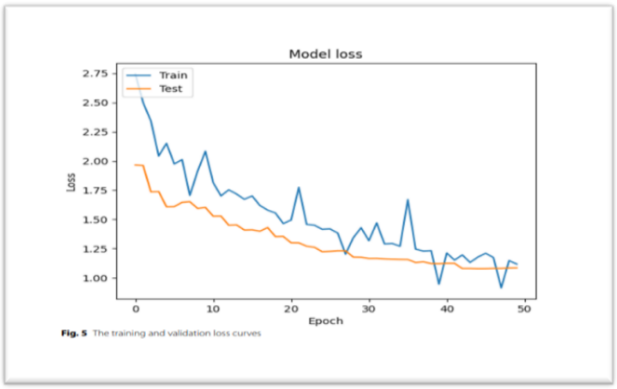

Fig. 5 The training and validation loss curves

4. *What Counts as emotion recognition under the AI Act*

Article 3(39) of the AI Act defines an emotion recognition system as an AI system intended to identify or infer the emotions or intentions of natural persons on the basis of their biometric data. However, **Recital 18 of the AI Act**, narrows this definition in two ways:

Physical states are excluded. Pain and fatigue are not emotions. A system that detects fatigue in pilots or drivers to prevent accidents therefore is not an emotion recognition system.

Detecting obvious expressions is not enough. Recognising a frown, a smile, a hand gesture or a raised voice is not emotion recognition, unless the system uses these signals to identify or infer

emotions. In practice, recording that a person is smiling, or how often they smile, falls outside the definition. By contrast, concluding from someone's voice, gestures or lack of smiling that an employee is sad or enraged and about to become violent falls squarely within the definition of emotion recognition.

5. AI emotion recognition systems banned at Work and in the Classroom

Article 5(1)(f) of the AI Act prohibits placing on the market, putting into service for this purpose, or using AI systems to infer the emotions of natural persons in the workplace and in education institutions. The only exception is for systems intended for medical or safety reasons.

Recital 44 explains why. Employees and students are in a weaker position than employers and schools, and emotion recognition is inherently intrusive. Its scientific basis is also contested, because people express emotions differently across cultures, situations and individuals. As a result, these systems may lead to discrimination and infringements of fundamental rights, particularly the rights to privacy, human dignity and freedom of thought.

The stakes are high. Under **Article 99(3)**, breaching any Article 5 prohibition can lead to administrative fines of up to EUR 35 million or, where the offender is an undertaking, 7% of its total worldwide annual turnover, whichever is higher.

National rules may go further. Article 2(11) allows the EU and Member States to adopt more favorable rules for employees, including through collective agreements. For example, a Member State could ban emotion recognition at work even where it is used for medical purposes.

No extension for now. In its Report of 20 May 2026[1], the European Commission saw no immediate need to amend Article 5 or Annex III, given the limited practical experience with their application. Although stakeholders had proposed extending the Article 5 (1)(f) ban to other sensitive areas, such as migration, the Commission deferred that question to future reviews.

6. Outside the Ban: High-Risk AI systems

Emotion recognition systems based on biometric data, whose use is permitted under Union or national law and fall outside the Article 5 ban are not unregulated. Under **Article 6(2) and point 1(c) of Annex III of the AI Act**, they are classified as high-risk AI systems.

Following the Digital Omnibus on AI (Regulation (EU) 2026/1744), these high-risk obligations will apply from 2 December 2027. The Article 50 transparency obligations, however, have applied since 2 August 2026.

The providers and deployers of such systems must therefore comply with the high-risk requirements in Sections 2 and 3 of Chapter III.

Transparency on top. Since 2 August 2026, deployers have also had to comply with the transparency obligations in Article 50(3), (5) and (6). The relevant European Commission Guidelines[2] explain these obligations in more detail. In particular, deployers must: (a) inform natural persons exposed to such systems, in a clear and distinguishable manner, that those systems are operating, no later than at the time of their first interaction with or exposure to them; and (b) process personal data in accordance with the EU data protection law.

There is a narrow exception for systems permitted by law to detect, prevent or investigate criminal offences, subject to appropriate safeguards.

[1] https://digital-strategy.ec.europa.eu/en/library/report-review-prohibitions-and-high-risk-ai

[2] https://digital-strategy.ec.europa.eu/en/library/guidelines-transparency-obligations-providers-and-deployers-ai-systems

7. Where exactly is the line? The Commission's Guidelines and the Kindt Study

According to the Commission's Guidelines[3] on prohibited AI practices, the Article 5(1)(f) ban applies where **all** of the following conditions are met:

(A) An AI system within the meaning of Article 3(1), falling within the material and territorial scope of Article 2 of the AI Act.

(B) The placing on the market, putting into service or use of that AI system.

(C) Identifying or inferring emotions or intentions from biometric data. Emotions and intentions should be understood broadly. Systems based on text alone, such as chatbot conversations, fall outside the ban.

(D) Use in the workplace or in education.

- Workplace. This means any physical or virtual place where work is performed. It covers individuals -regardless of whether they are employees, contractors, trainees, volunteers, and the self-employed, as well as job applicants and staff on probation. Emotion recognition used to coach call-center employees through webcams and voice recognition systems is allowed only for personal training, provided their outputs are not shared with the human resources team, cannot affect the employee's assessment or promotion, and do not circumvent the prohibition or affect the employment relationship.

- Education. This means public or private institutions, whether online, in person or hybrid, including vocational schools. It covers teachers, students and applicants for admission. These institutions are usually accredited by national authorities and award certificates, or attendance is required to obtain one. By contrast, language learning apps used outside an institution are not covered. Similarly, role-play training tools for example for actors or teachers are allowed provided they cannot affect the trainee's assessment or certification.

(E) No medical or safety-related purpose. "Safety" here means protecting life and health only. It does not cover protecting property, for example against theft or fraud. The exception is interpreted narrowly:

- It must be strictly necessary and proportionate, assessed objectively against the medical or safety-related purpose and not against the employer's or institution's needs.
- It must come with safeguards, such as prior written and reasoned expert assessment.
- The data collected and processed must not be used for any other purpose; and
- Less intrusive alternatives must be considered.

What the exception covers, and what it doesn't:

The exception covers, for example, CE-marked therapeutic medical devices that support people with autism or with visual or hearing impairments. It does **not**, however, cover tools that monitor general well-being, motivation, satisfaction, burnout, stress or boredom. The only exception is where stress or loss of concentration creates a specific risk, such as when operating dangerous machinery or handling hazardous chemicals. Even then, the data cannot be used for anything else, including performance reviews.

Grey areas. The Kindt Study[4] sheds light on three key questions:
(i) Can therapy be preventive, for example, to prevent injuries or the onset of neuropsychiatric disorders? As a rule, the answer is likely no, but it could be yes where there is a serious, specific,

[3] https://digital-strategy.ec.europa.eu/en/library/commission-publishes-guidelines-prohibited-artificial-intelligence-ai-practices-defined-ai-act

[4] https://op.europa.eu/en/publication-detail/-/publication/3a999ace-4829-11f1-8095-01aa75ed71a1/language-en

foreseeable and scientifically substantiated risk of harm.
(ii) Must therapeutic use be confined to healthcare settings, given the expansion of healthcare beyond hospitals to home care and telemedicine?
(iii) What constitutes scientifically a system's essential therapeutic function?

Even where the ban does not apply, these systems remain subject to data protection, consumer protection, product safety, employment and medical device regulations.

8. Concluding Remarks: Can AI really read us?

The Article 5(1)(f) raises a broader question: should it extend to other contexts where emotion recognition threatens autonomy and privacy? While a broader prohibition would strengthen protection, it would also have to be weighed against the value of research and potentially beneficial uses, especially where emotion recognition only supports a decision rather than driving it.

Even within its current scope, the ban leaves open questions. It is not yet clear exactly what counts as inferring emotions from biometric data, or which educational and workplace settings are covered. For instance, would monitoring an employee's emotions during a break at the workplace be caught? Clearer boundaries would make the ban easier to apply consistently.

The medical and safety exceptions are the main risk. If they are read broadly, "safety" could justify tracking stress levels in high-risk jobs or monitoring students to prevent bullying and violence. The real question is where a legitimate safety measure ends and a backdoor to workplace or classroom surveillance begins.

The ban also leaves a notable gap. Because it relies on biometric data, it does not cover systems that infer emotions from text, which is exactly what ChatGPT did in our opening example. As conversational AI enters workplaces and classrooms, the legislator may need to consider whether the biometric threshold still reflects where the real risks lie.

This is where the GDPR comes in. Where the AI Act does not apply, whether because a system is text-based, falls within an exception or is classified as high-risk rather than prohibited, the GDPR still applies in full. Inferences about a person's emotions are personal data, whatever the source. Employers and educational institutions will therefore need a valid legal basis, will have to respect purpose limitation and data minimization, and in most cases will need to carry out a data protection impact assessment. Given the imbalance of power, consent from employees or students will rarely be a reliable basis. In practice, the GDPR may become the main safeguard in the areas that the AI Act leaves uncovered.

For organisations, the immediate priority is practical. Emotion-inference features are increasingly built into everyday tools, such as HR analytics, recruitment platforms, online proctoring, call-center software and well-being apps, often without being marketed as such. Employers and educational institutions should map the tools they use, ask vendors directly whether any feature infers emotions, and assess each one under both the AI Act and the GDPR.

In our view, protecting autonomy may ultimately require more than limits on technology. It may also require a right to refuse both the assessment of one's emotions and the use of any conclusions drawn from them.

ChatGPT may have guessed how I felt today. But can AI truly understand our emotions, let alone know them better than we do?

Contact Us



Irene Kyriakides
PARTNER
i.kyriakides@kglawfirm.gr

Follow Us

ATHENS OFFICE

28, Dimitriou Soutsou Str.,
115 21 Athens

T +30 210 817 1500
E kg.law@kglawfirm.gr

THESSALONIKI OFFICE

31, Politechniou Str.,
551 34 Thessaloniki

T +30 2310 441 552
E kg.law@kglawfirm.gr

www.kglawfirm.gr

Disclaimer: This newsletter contains general information only and is not intended to provide specific legal, or other professional advice or services, nor is it suitable for such professional advice, and should not be used as a basis for any decision or action that may affect you or your business. Before making any decision or taking any action that may affect you or your business, you should consult a qualified professional advisor. We remain at your disposal should you require any further information or clarification in this regard.

©Kyriakides Georgopoulos, 2026