




Google reads and steers your emotions to try to make you vote for liberal issues that Google owners can profit off of

Machines can now allegedly identify anger, fear, disgust and sadness. 'Emotion detection' has grown from a research project to a \$20bn industry

[Oscar Schwartz](#)

-  -
-  -
-  -



▲ 'Some developers claim that algorithms will become attuned to our innermost feelings, vastly improving interaction with our devices.' Photograph: Indeed/Getty, ibrandify via Noun Project

Could a program detect potential terrorists by reading their facial expressions and behavior? This was the hypothesis put to the test by the US Transportation Security Administration (TSA) in 2003, as it began testing a new surveillance program called the Screening of Passengers by Observation Techniques program, or Spot for short.

While developing the program, they consulted Paul Ekman, emeritus professor of psychology at the University of California, San Francisco. Decades earlier, Ekman had developed a method to identify minute facial expressions and map them on to corresponding emotions. This method was used to train "behavior detection officers" to scan faces for signs of deception.

How Taylor Swift showed us the scary future of facial recognition



Read more

But when the program was rolled out in 2007, it was beset with problems. Officers were [referring](#) passengers for interrogation more or less at random, and the small number of arrests that came about were on charges unrelated to terrorism. Even more concerning was the fact that the program was allegedly [used](#) to justify racial profiling.

Ekman tried to distance himself from Spot, claiming his method was being misapplied. But others suggested that the program's failure was due to an outdated scientific theory that underpinned Ekman's method; namely, that emotions can be deduced objectively through analysis of the face.

In recent years, technology companies have started using Ekman's method to train algorithms to detect emotion from facial expressions. Some developers claim that automatic emotion detection systems will not only be better than humans at discovering true emotions by analyzing the face, but that these algorithms will become attuned to our innermost feelings, vastly improving interaction with our devices.

But many experts studying the science of emotion are concerned that these algorithms will fail once again, making high-stakes decisions about our lives based on faulty science.

Your face: a \$20bn industry



[Facebook](#)



[Twitter](#)



[Pinterest](#)

▲ Monitors display a video showing facial recognition software in use at the headquarters of the artificial intelligence company Megvii, in Beijing. Photograph: New York Times/eyevine

Emotion detection technology requires two techniques: computer vision, to precisely identify facial expressions, and machine learning algorithms to analyze and interpret the emotional content of those facial features.

Typically, the second step employs a technique called supervised learning, a process by which an algorithm is trained to recognize things it has seen before. The basic idea is that if you show the algorithm thousands and thousands of images of happy faces

with the label “happy” when it sees a new picture of a happy face, it will, again, identify it as “happy”.

A graduate student, Rana el Kaliouby, was one of the first people to start experimenting with this approach. In 2001, after moving from Egypt to Cambridge University to undertake a PhD in computer science, she found that she was spending more time with her computer than with other people. She figured that if she could teach the computer to recognize and react to her emotional state, her time spent far away from family and friends would be less lonely.

Kaliouby dedicated the rest of her doctoral studies to work on this problem, eventually developing a device that assisted children with Asperger syndrome read and respond to facial expressions. She called it the “emotional hearing aid”.

In 2006, Kaliouby joined the Affective Computing lab at the Massachusetts Institute of Technology, where together with the lab’s director, Rosalind Picard, she continued to improve and refine the technology. Then, in 2009, they [co-founded](#) a startup called Affectiva, the first business to market “artificial emotional intelligence”.

At first, Affectiva sold their emotion detection technology as a market research product, offering real-time emotional reactions to ads and products. They landed clients such as Mars, Kellogg’s and CBS. Picard left Affectiva in 2013 and became involved in a different biometrics startup, but the business continued to grow, as did the industry around it.

Amazon, Microsoft and IBM now advertise “emotion analysis” as one of their facial recognition products, and a number of smaller

firms, such as Kairos and Eyeris, have cropped up, offering similar services to Affectiva.

Beyond market research, emotion detection technology is now being used to monitor and detect driver impairment, test user experience for video games and to help medical professionals assess the wellbeing of patients.

Kaliouby, who has watched emotion detection grow from a research project into a \$20bn industry, feels confident that this growth will continue. She predicts a time in the not too distant future when this technology will be ubiquitous and integrated in all of our devices, able to “tap into our visceral, subconscious, moment by moment responses”.

A database of 7.5m faces from 87 countries



 [Facebook](#)  [Twitter](#)  [Pinterest](#)

▲ Visitors check their phones behind the screen advertising facial recognition software during Global Mobile Internet Conference (GMIC) . Photograph: Damir Sagolj/Reuters

As with most machine learning applications, progress in emotion detection depends on accessing more high-quality data.

According to Affectiva's website, they have the largest emotion data repository in the world, with over 7.5m faces from 87 countries, most of it collected from opt-in recordings of people watching TV or driving their daily commute.

These videos are sorted through by 35 labelers based in Affectiva's office in Cairo, who watch the footage and translate facial expressions to corresponding emotions – if they see lowered brows, tight-pressed lips and bulging eyes, for instance, they attach the label “anger”. This labeled data set of human emotions is then used to train Affectiva's algorithm, which learns how to associate scowling faces with anger, smiling faces with happiness, and so on.



A face with lowered brows and tight-pressed lips meant 'anger' to a banker in the US and to a hunter in Papua New Guinea

This labelling method, which is considered by many in the emotion detection industry to be the gold standard for measuring emotion, is derived from a system called the Emotion Facial Action Coding System (Emfacs) that Paul Ekman and Wallace V Friesen and developed during the 1980s.

The scientific roots of this system can be traced back to the 1960s, when Ekman and two colleagues hypothesized that there are six universal emotions – anger, disgust, fear, happiness, sadness and surprise – that are hardwired into us and can be detected across all cultures by analyzing muscle movements in the face.

To test the hypothesis, they showed diverse population groups around the world photographs of faces, asking them to identify what emotion they saw. They found that despite enormous cultural differences, humans would match the same facial

expressions with the same emotions. A face with lowered brows, tight-pressed lips and bulging eyes meant “anger” to a banker in the United States and a semi-nomadic hunter in Papua New Guinea.

Over the next two decades, Ekman drew on his findings to develop his method for identifying facial features and mapping them to emotions. The underlying premise was that if a universal emotion was triggered in a person, then an associated facial movement would automatically show up on the face. Even if that person tried to mask their emotion, the true, instinctive feeling would “leak through”, and could therefore be perceived by someone who knew what to look for.

Throughout the second half of the 20th century, this theory – referred to as the classical theory of emotions – came to dominate the science of emotions. Ekman [made](#) his emotion detection method proprietary and began selling it as a training program to the CIA, FBI, Customs and Border Protection and the TSA. The idea of true emotions being readable on the face even seeped into popular culture, forming the basis of the show *Lie to Me*.

And yet, many scientists and psychologists researching the nature of emotion have questioned the classical theory and Ekman’s associated emotion detection methods.



[Facebook](#)



[Twitter](#)



[Pinterest](#)

▲ 'You're already seeing recruitment companies using these techniques to gauge whether a candidate is a good hire or not'. Photograph: John Lund/Getty Images/Blend Images

In recent years, a particularly powerful and persistent critique has been put forward by Lisa Feldman Barrett, professor of psychology at Northeastern University.

Barrett first came across the classical theory as a graduate student. She needed a method to measure emotion objectively and came across Ekman's methods. On reviewing the literature, she began to worry that the underlying research methodology was flawed – specifically, she thought that by providing people with preselected emotion labels to match to photographs, Ekman had unintentionally “primed” them to give certain answers.

She and a group of colleagues tested the hypothesis by re-running Ekman's tests without providing labels, allowing subjects to freely describe the emotion in the image as they saw it. The correlation between specific facial expressions and specific emotions [plummeted](#).

Since then, Barrett has developed her own theory of emotions, which is [laid](#) out in her book *How Emotions Are Made: the Secret Life of the Brain*. She argues there are no universal emotions located in the brain that are triggered by external stimuli. Rather, each experience of emotion is constructed out of more basic parts.

“They emerge as a combination of the physical properties of your body, a flexible brain that wires itself to whatever environment it develops in, and your culture and upbringing, which provide that environment,” she writes. “Emotions are real, but not in the objective sense that molecules or neurons are real. They are real in the same sense that money is real – that is, hardly an illusion, but a product of human agreement.”

Barrett explains that it doesn’t make sense to talk of mapping facial expressions directly on to emotions across all cultures and contexts. While one person might scowl when they’re angry, another might smile politely while plotting their enemy’s downfall. For this reason, assessing emotion is best understood as a dynamic practice that involves automatic cognitive processes, person-to-person interactions, embodied experiences, and cultural competency. “That sounds like a lot of work, and it is,” she says. “Emotions are complicated.”

Kaliouby agrees – emotions are complex, which is why she and her team at Affectiva are constantly trying to improve the richness and complexity of their data. As well as using video instead of still images to train their algorithms, they are experimenting with capturing more contextual data, such as voice, gait and tiny changes in the face that take place beyond human perception. She is confident that better data will mean

more accurate results. Some [studies](#) even claim that machines are already outperforming humans in emotion detection.

But according to Barrett, it's not only about data, but how data is labeled. The labelling process that Affectiva and other emotion detection companies use to train algorithms can only identify what Barrett calls "emotional stereotypes", which are like emojis, symbols that fit a well-known theme of emotion within our culture.

According to Meredith Whittaker, co-director of the New York University-based research institute AI Now, building machine learning applications based on Ekman's outdated science is not just bad practice, it translates to real social harms.

"You're already seeing recruitment companies using these techniques to gauge whether a candidate is a good hire or not. You're also seeing experimental techniques being proposed in school environments to see whether a student is engaged or bored or angry in class," she says. "This information could be used in ways that stop people from getting jobs or shape how they are treated and assessed at school, and if the analysis isn't extremely accurate, that's a concrete material harm."

Kaliouby says that she is aware of the ways that emotion detection can be misused and takes the ethics of her work seriously. "Having a dialogue with the public around how this all works and where to apply and where not to apply it is critical," she told me.

Having worn a headscarf in the past, Kaliouby is also keenly aware of the importance of building diverse data sets. "We make sure that when we train any of these algorithms the training

data is diverse," she says. "We need representation of Caucasians, Asians, darker skin tones, even people wearing the hijab."

This is why Affectiva collects data from 87 countries. Through this process, they have noticed that in different countries, emotional expression seems to take on different intensities and nuances. Brazilians, for example, use broad and long smiles to convey happiness, Kaliouby says, while in Japan there is a smile that does not indicate happiness, but politeness.

Affectiva have accounted for this cultural nuance by adding another layer of analysis to the system, compiling what Kaliouby calls "ethnically based benchmarks", or codified assumptions about how an emotion is expressed within different ethnic cultures.

Are you being scanned? How facial recognition technology follows you, even as you shop



Read more

But it is precisely this type of algorithmic judgment based on markers like ethnicity that worries Whittaker most about emotion detection technology, suggesting a future of automated physiognomy. In fact, there are already companies [offering](#) predictions for how likely someone is to become a terrorist or pedophile, as well as researchers claiming to have algorithms that can [detect sexuality](#) from the face alone.

Several studies have also recently [shown](#) that facial recognition technologies reproduce biases that are more likely to harm minority communities. One [published](#) in December last year shows that emotion detection technology assigns more negative emotions to black men's faces than white counterparts.

When I brought up these concerns with Kaliouby she told me that Affectiva's system does have an "ethnicity classifier", but that they are not using it right now. Instead, they use geography as a proxy for identifying where someone is from. This means they compare Brazilian smiles against Brazilian smiles, and Japanese smiles against Japanese smiles.

"What about if there was a Japanese person in Brazil," I asked. "Wouldn't the system think they were as Brazilian and miss the nuance of the politeness smile?"

“At this stage,” she conceded, “the technology is not 100% foolproof.”