

The Rise of Emotional Surveillance

Ellen Cushing

The good news, for me at least, is that the computer thinks I have a nice personality. According to an app called MorphCast, I was, in a recent meeting with my boss, generally “amused,” “determined,” and “interested,” though—sue me—occasionally “impatient.” MorphCast, you see, purports to glean insights into the depths and vagaries of human emotion using AI. It found that my affect was “positive” and “active,” as opposed to negative and/or passive. My attention was reasonably high. Also, the AI informed me that I wear glasses—revelatory!

The bad news is that software now purports to glean insights into the depths and vagaries of human emotion using AI, and it is coming to watch you. If it isn’t already: Morphcast, for example, has licensed its technology to a [mental-health app](#), a program that monitors schoolchildren’s [attention](#), and McDonald’s, which launched a promotional campaign in Portugal that scanned app users’ faces and [offered them personalized coupons](#) based on their (supposed) mood. It is one of many, many such companies doing similar work—the industry term is *emotion AI* or sometimes *affective computing*.

Some products analyze video of meetings or job interviews or focus groups; others listen to audio for pitch, tone, and word choice; still others can scan chat transcripts or emails and spit out a report about worker sentiment. Sometimes, the emotion AI is baked in as a feature in multiuse software, or sold as part of an expensive analytics package marketed to businesses. But it’s also available as a stand-alone product, and the barrier to entry is shin-high: I used MorphCast at no cost, taking advantage of a free trial, and with no special software. At no point was I compelled to ask my interlocutors if they consented to being analyzed in this way (though I did ask, because of my good personality).

Every successful technology needs to find a problem that people are willing to pay money to solve. In the case of emotion AI, that problem appears largely, so far, to be worker performance and productivity, especially in customer service and blue-collar labor. If you’ve ever been warned that your call “is being monitored for quality-assurance purposes,” chances are good that the person on the other end is being assessed by emotion AI: The insurance giant MetLife, like many other businesses, [uses](#) software to monitor call-center agents’ pitch and tone of voice. Trucking companies use eyeball trackers, high-sensitivity recording equipment, and brain-wave scanners to find signs of driver distress or fatigue. Burger King is [piloting](#) an AI chatbot embedded in employee headsets that will evaluate their interactions for friendliness. Her name is Patty.

[Read: AI’s next frontier: people skills](#)

In 2022, the writer Cory Doctorow [theorized](#) about what he called the “Shitty Technology Adoption Curve”: Extractive technologies, he wrote, come first to people in precarious circumstances—like, say, low-wage jobs—before they are refined and normalized and brought to people in greater positions of power. “Each disciplinary technology,” he [later wrote](#), “starts with people way down on the ladder, then ascends the ladder, rung by rung.”

Emotion AI’s next step is white-collar work. The Slack integration Aware advertises its ability to [continuously monitor](#) messages for “sentiment and toxicity”; Azure, Microsoft’s cloud-computing software, also [allows employers](#) to, theoretically, use AI to [batch-analyze](#) workers’ chat messages. MorphCast’s Zoom extension [tracks](#), in real time, meeting participants’ attention, excitement, and positivity. The emotion-AI company Imentiv [advises](#) clients on applying emotional analysis to the job-interview process, promising employers detailed analysis of candidates’ emotional engagement, intensity, and valence, as well as personality type. A number of HR companies are [turning toward](#) AI that applies sentiment analysis to employee surveys. Framery, which makes soundproof phone pods and [sells them](#) to companies such as Microsoft and L’Oreal, has [tested](#) outfitting its chairs with biosensors capable of measuring heart rate, breathing rate, and nervousness.

Last year, the European Union [banned](#) emotion AI in the workplace, except for when it’s used for medical or safety reasons. (The regulation prompted MorphCast, which was founded in Florence, to relocate to the Bay Area.) But still, according to [one estimate](#), the global emotion-AI market is expected to triple by 2030, to \$9 billion, as the technology becomes more sophisticated and more available. It is not that hard for me to imagine a near future in which workers in all industries are pushed to work not only harder and more, but more happily and more agreeably. This is the new era of employee surveillance: invisible, AI-supercharged, always on.

To have a job is, fundamentally, to trade some amount of freedom for some amount of money. “The idea that managers or corporations want to keep tabs on what their workers are up to is not a new concept,” Karen Levy, an associate professor of information sciences at Cornell, told me. Using new technologies to [track people’s emotions](#) without their consent is also not new—see Facebook in the 2010s. Nor is the lack of privacy protection for workers generally: Although regulations vary by state, U.S. federal law gives employers broad permission to monitor much of what an employee does on company time, property, and devices—to scan communication and record video and audio, even when employees are off duty.

For decades, workers were protected not by law but by reality: Their information may have been collectable, but analyzing such a huge amount of it was practically impossible. Not anymore. Over the past few years, a wave of companies has emerged to extract sophisticated and granular information about how employees spend their time, sometimes down to the minute, using tech such as location trackers, keystroke loggers, cameras, and microphones. (Employees have in turn figured out some work-arounds, such as [mouse jigglers](#) and keystroke simulators.) But the product is less the data than it is these companies’ ability to turn the data into narrative: “AI-powered systems can now analyze 100% of interactions rather than the typical 1-3% sample size of traditional approaches, ensuring nothing falls through the cracks,” the promotional copy on one call-center-monitoring firm’s website [reads](#).

[Read: When did the job market get so rude?](#)

And as the technological conditions for widespread employee surveillance have fallen into place, so have the cultural and economic conditions. The pandemic pushed more workers than ever before into remote work, out of sight of their bosses. Trust between employers and employees is [tanking](#). A recession has been promised for years, and while we wait, AI is upending the job market: The technologies currently surveilling workers such as call-center staff may soon replace them entirely, and in the meantime, corporations are laying off people by the tens of thousands and looking for other ways to replace them with machines. The availability of data, and tools with which to examine such information, has turned human resources, once a qualitative discipline, into “[people analytics](#).” After being bombarded for years with eerily targeted ads and news stories about data breaches, many Americans have settled into a state of [privacy nihilism](#), one in which we know that all of our data are being collected and exploited, even if we prefer not to think about it too much.

The companies selling digital surveillance advertise all manner of use cases: worker safety, mental health, organizational efficiency, burnout reduction in high-stakes fields such as medicine and transportation. (At First Horizon Bank, AI monitors call-center employees’ stress and [presents them with a montage of pictures](#) of their families when levels get too high.) In practice, these companies also seem to be selling an empirical assessment of worker productivity, down to the minute. A [2022 New York Times investigation](#) found that eight of the 10 largest private employers in the United States track individual workers’ productivity. In [one poll](#), 37 percent of employers said they had used stored recordings to fire a worker.

But the problem with many of these tools is that they’re not very good at doing the things they say they can. A keystroke tracker can’t necessarily know the difference between mindless typing and focused knowledge production; a breakdown of someone’s app usage doesn’t definitionally tell you much about the kind and quality of work they’re doing inside the app. At UnitedHealth Group, the [Times](#) found, a program used to monitor efficacy (and help set compensation) docked social workers for keyboard inactivity, even though they were offline for a good reason: They were in counseling sessions with patients. (UnitedHealth acknowledged to the *Times* that it monitored staff, but noted that multiple factors go into performance evaluations.)

If computers are flawed analysts of straightforward productivity, imagine, now, applying that same technology to something as complex as the constellation of emotions expressible by humans. Study after study shows that AI replicates the biases of the data it’s trained on. (In 2018, Lauren Rhue, then a professor of information systems and analytics at Wake Forest University, [studied](#) photographs of NBA players and emotion-recognition AI; she discovered that the tech found Black players to be angrier than their white teammates—even, in some cases, if they were smiling.) Many emotion-AI products base their rubrics on the clinical psychologist [Paul Ekman’s theory](#) of basic emotions, which holds that all people experience the same six core emotions: anger, disgust, fear, happiness, sadness, and surprise. That theory has been [widely challenged](#) as oversimplistic and methodologically flawed in the many decades since it was first published.

Body language is a metaphor that has become a cliché, but anyone who has spent much time at all around other people understands that everyone speaks in a different dialect. “Your movements,” the neuroscientist and psychologist Lisa Feldman Barrett told me, “whether it’s on your face or in your body or the tones that you emit, don’t have inherent emotional meaning. They have relational meaning.” They vary based on the context of the conversation, the physiognomy of the person making them, culture, room temperature, vibes.

[Read: The new age of performance anxiety](#)

Research suggests, Barrett said, that in the U.S., people scowl when angry about 35 percent of the time. This means a scowl is relatively likely to be an expression of anger. It also means that if you are looking only for a scowl, you miss about 65 percent of cases in which a person is angry. Half the time when people scowl, they aren't angry at all. "So imagine a situation where you're in a job interview," she said. "You're listening really carefully to the person, you're scowling as you're listening because you're paying really, really close attention, and an AI labels you as angry. You will not get that job."

A hospital call-center employee verbally expressing sadness when speaking with a patient about their condition could be read as conveying an inappropriate lack of warmth or cheer. A fast-food employee listening intently to someone's order could be perceived as upset. Although the MorphCast app liked me, I work in a newsroom in 2026—it's easy enough to imagine my little mood dial drifting into the "negative" quadrant for reasons having nothing to do with my personal pleasantness.

HireVue—a job-screening platform whose [clients include](#) Ikea, the pharmaceutical company Regeneron, and the Children's Hospital of Philadelphia—uses AI to interview and analyze job candidates and promotion-seeking employees. In a 2025 legal complaint, the ACLU [alleged](#) that HireVue's platform didn't provide adequate subtitles in a promotion interview for a deaf member of the accessibility team at Intuit, the financial-software company. The employee was denied her promotion; in the email that she got explaining the decision, she was advised to "practice active listening." (HireVue and Intuit [have disputed](#) these claims.)

Barrett has been studying the psychology of emotion for years. Toward the end of our conversation, I asked what she wished more people knew about emotion AI. First she asked if she was allowed to swear. "I have been talking about this for a fucking decade," she said. "There are—I mean, literally, at this point—hundreds and hundreds of studies involving thousands and thousands of people to show that when it comes to emotion, variation is the norm." The idea that emotions can be objectively measured or analyzed at all, in other words, is fantasy.

The companies packaging this technology—and the other companies buying it—do make some good points. Humans are biased, too, they say. In interviews, representatives of some companies told me about their algorithms' abilities to reveal patterns that impressions alone cannot. The tech will get better—this is the promise of AI: that it learns from its mistakes.

[Read: America isn't ready for what AI will do to jobs](#)

But if it gets better, then what? Most of the time, discussion of emotion AI and similar tools focuses on what can go wrong—the muddled signals, the imperfect analysis, the scowl of empathy, the junk science being leveraged to fire workers. The more I used MorphCast, the more I began to worry about the opposite: a world where the robot embedded in my inbox and my Zoom account *could* actually say something meaningful and true about my emotional state; a world where, in addition to my *job* job, I have the work of making the emotion robot think that I'm sufficiently cheerful; a world where my every unintentional facial expression has bearing on my ability to feed my family. I've always known that my workplace holds wide-ranging power over me, but I don't need it made quite so literal. "I mean, there's a reason there's a lot of sci-fi stories about this kind of thing," Levy, the Cornell information scientist, told me.

Levy wrote a book about the way affective computing and other forms of biometric surveillance have been deployed in the trucking industry—a field that, due to its mobile and distributed workforce, was long immune to surveillance. But in 2016, the federal government began mandating electronic logging, in an attempt to reduce overwork and ward off accidents. The constant surveillance added its own form of stress, however—without actually reducing crashes. Truckers, historically, have had a "really notable degree of pride," Levy said, and "had a lot of autonomy to kind of do the work in the way that they saw fit." That pride, she said, has been picked away at, as the computers have begun watching. "There really is, I think, a pretty strong dignitary concern to being watched in some fairly intimate ways, or pretty granular ways that have to do with people's bodies and their spaces." I am flattered the computer liked me, but I'd prefer it didn't know me at all.