


# We Know How You Feel



*By scanning your face, computers can decode your unspoken reaction to a movie, a political debate, even a video call with a friend. Credit Illustration by Bryan Christie*

Three years ago, archivists at A.T. & T. stumbled upon a rare fragment of computer history: a short film that Jim Henson produced for Ma Bell, in 1963. Henson had been hired to make the film for a conference that the company was convening to showcase its strengths in machine-to-machine communication. Told to devise a faux robot that believed it functioned better than a person, he came up with a cocky, boxy, jittery, bleeping Muppet on wheels. “This is computer H14,” it proclaims as

the film begins. “Data program readout: number fourteen ninety-two per cent H2SOSO.” (Robots of that era always seemed obligated to initiate speech with senseless jargon.) “Begin subject: Man and the Machine,” it continues. “The machine possesses supreme intelligence, a faultless memory, and a beautiful soul.” A blast of exhaust from one of its ports vaporizes a passing bird. “Correction,” it says. “The machine does not have a soul. It has no bothersome emotions. While mere mortals wallow in a sea of emotionalism, the machine is busy digesting vast oceans of information in a single all-encompassing gulp.” H14 then takes such a gulp, which proves overwhelming. Ticking and whirring, it begs for a human mechanic; seconds later, it explodes.

The film, titled “Robot,” captures the aspirations that computer scientists held half a century ago (to build boxes of flawless logic), as well as the social anxieties that people felt about those aspirations (that such machines, by design or by accident, posed a threat). Henson’s film offered something else, too: a critique—echoed on television and in novels but dismissed by computer engineers—that, no matter a system’s capacity for errorless calculation, it will remain inflexible and fundamentally unintelligent until the people who design it consider emotions less bothersome. H14, like all computers in the real world, was an imbecile.

Today, machines seem to get better every day at digesting vast gulps of information—and they remain as emotionally inert as ever. But since the nineteen-nineties a small number of researchers have been working to give computers the capacity to read our feelings and react, in ways that have come to seem startlingly human. Experts on the voice have trained computers to identify deep patterns in vocal pitch, rhythm, and intensity; their software can scan a conversation between a woman and a child and determine if the woman is a mother, whether she is looking the child in the eye, whether she is angry or frustrated or joyful. Other machines can measure sentiment by assessing the arrangement of our words, or by reading our gestures. Still others can do so from facial expressions.

Our faces are organs of emotional communication; by some estimates, we transmit more data with our expressions than with what we say, and a few pioneers dedicated to decoding this information have made tremendous progress. Perhaps the most successful is an Egyptian scientist living near Boston, Rana el Kaliouby. Her company, Affectiva, formed in 2009, has been ranked by the business press as one of the country's fastest-growing startups, and Kaliouby, thirty-six, has been called a "rock star." There is good money in emotionally responsive machines, it turns out. For Kaliouby, this is no surprise: soon, she is certain, they will be ubiquitous.


Affectiva is situated in an office park behind a strip mall on a two-lane road in Waltham, Massachusetts, part of a corridor that serves as Boston's answer to Silicon Valley. The headquarters have the trappings of a West Coast startup—pool table, beanbag chairs—but the sensibility is New England; many of the employees are from M.I.T. From a conference room, the Amtrak line to Boston is visible beyond a large parking lot.

When I visited in September, Kaliouby walked me past charts of facial expressions, some of them scientific diagrams, some borrowed from comics. Kaliouby has a Ph.D. in computer science, and, like many accomplished coders, she has no trouble with mathematical concepts like Bayesian probability and hidden Markov models. But she is also at ease among people: emotive, warm, even flirtatious. She is a practicing Muslim, and until two years ago she wore a head scarf, which had the effect of drawing the eye to her rounded, expressive features. Frank Moss, a former director of M.I.T.'s Media Lab, where she held a postdoctoral position, told me that she has a high "emotional intelligence." As a mother of two, she worries about technology's effects on her children.

Affectiva is the most visible among a host of competing boutique startups: Emotient, Realeyes, Sension. After Kaliouby and I sat down, she told me, "I think

that, ten years down the line, we won't remember what it was like when we couldn't just frown at our device, and our device would say, 'Oh, you didn't like that, did you?' ” She took out an iPad containing a version of Affdex, her company's signature software, which was simplified to track just four emotional “classifiers”: happy, confused, surprised, and disgusted. The software scans for a face; if there are multiple faces, it isolates each one. It then identifies the face's main regions—mouth, nose, eyes, eyebrows—and it ascribes points to each, rendering the features in simple geometries. When I looked at myself in the live feed on her iPad, my face was covered in green dots. “We call them deformable and non-deformable points,” she said. “Your lip corners will move all over the place—you can smile, you can smirk—so these points are not very helpful in stabilizing the face. Whereas these points, like this at the tip of your nose, don't go anywhere.” Serving as anchors, the non-deformable points help judge how far other points move.

Affdex also scans for the shifting texture of skin—the distribution of wrinkles around an eye, or the furrow of a brow—and combines that information with the deformable points to build detailed models of the face as it reacts. The algorithm identifies an emotional expression by comparing it with countless others that it has previously analyzed. “If you smile, for example, it recognizes that you are smiling in real time,” Kaliouby told me. I smiled, and a green bar at the bottom of the screen shot up, indicating the program's increasing confidence that it had identified the correct expression. “Try looking confused,” she said, and I did. The bar for confusion spiked. “There you go,” she said.



*“Have you been reading a lot of exhaustive biographies lately?”* [Buy the print »](#)

Like every company in this field, Affectiva relies on the work of Paul Ekman, a research psychologist who, beginning in the sixties, built a convincing body of evidence that there are at least six universal human emotions, expressed by everyone’s face identically, regardless of gender, age, or cultural upbringing. Ekman worked to decode these expressions, breaking them down into combinations of forty-six individual movements, called “action units.” From this work, he compiled the Facial Action Coding System, or *facs*—a five-hundred-page taxonomy of facial movements. It has been in use for decades by academics and professionals, from computer animators to police officers interested in the subtleties of deception.

Ekman has had critics, among them social scientists who argue that context plays a far greater role in reading emotions than his theory allows. But context-blind computers appear to support his conclusions. By scanning facial action units, computers can now outperform most people in distinguishing social smiles from those triggered by spontaneous joy, and in differentiating between faked pain and genuine pain. They can determine if a patient is depressed. Operating with unflagging attention, they can register expressions so fleeting that they are unknown even to the person making them. Marian Bartlett, a researcher at the

University of California, San Diego, and the lead scientist at Emotient, once ran footage of her family watching TV through her software. During a moment of slapstick violence, her daughter, for a single frame, exhibited ferocious anger, which faded into surprise, then laughter. Her daughter was unaware of the moment of displeasure—but the computer had noticed. Recently, in a peer-reviewed study, Bartlett’s colleagues demonstrated that computers scanning for “micro-expressions” could predict when people would turn down a financial offer: a flash of disgust indicated that the offer was considered unfair, and a flash of anger prefigured the rejection.

Kaliouby often emphasizes that this technology can read only facial expressions, not minds, but Affdex is marketed as a tool that can make reliable inferences about people’s emotions—a tap into the unconscious. The potential applications are vast. CBS uses the software at its Las Vegas laboratory, Television City, where it tests new shows. During the 2012 Presidential elections, Kaliouby’s team used Affdex to track more than two hundred people watching clips of the Obama-Romney debates, and concluded that the software was able to predict voting preference with seventy-three-per-cent accuracy. Affectiva is working with a Skype competitor, Oovoo, to integrate it into video calls. “People are doing more and more videoconferencing, but all this data is not captured in an analytic way,” she told me. Capturing analytics, it turns out, means using the software—say, during a business negotiation—to determine what the person on the other end of the call is not telling you. “The technology will say, ‘O.K., Mr. Whatever is showing signs of engagement—or he just smirked, and that means he was not persuaded.’ ”

Kaliouby created Affectiva with her mentor, Rosalind Picard, a professor at the M.I.T. Media Lab, whose early research laid the groundwork for the company. Picard, who has degrees in electrical engineering and in computer science, came to the Media Lab in 1990, to develop technology for image compression, but she soon reached a technical impasse. The models then in vogue worked independently of the

content: a landscape of the Grand Canyon and a Presidential portrait were compressed in the same way. Picard believed that the process could be improved if a computer recognized what it was looking at. But to do this it would need to be capable of vision, not merely sight; like the brain, it would need to distinguish objects, then determine which ones mattered.

One day, Picard picked up Richard Cytowic's "The Man Who Tasted Shapes," a book on synesthesia. Cytowic made the case that perception was partly processed in the brain's limbic system, an ancient part of neural anatomy that handles attention, memory, and emotion. Attention and memory seemed pertinent to the problems Picard sought to solve; emotion, she hoped, was extraneous. But as she delved into the neuroscience literature she became convinced that reasoning and emotion were inseparable: just as too much emotion could cause irrational thinking, so could too little. Brain injuries specific to emotional processing robbed people of their capacity to make decisions, see the bigger picture, exercise common sense—the very qualities that she wanted computers to have.

"I wanted to be taken seriously, and emotion was not a serious topic," Picard told me. Nonetheless, in 1995, she circulated an informal paper on her findings; laced with references to Leibniz and "Star Trek," Curie and Kubrick, it argued that something like emotional reasoning was necessary for true machine intelligence, and also that programmers should consider affect when writing software that interacts with people. At first, her ideas were met with perplexity. One scientist told her, "Why are you working on emotion? It's irrelevant!" Unmoved, Picard turned down hundreds of thousands of dollars in grants for research in image compression, and expanded her ideas into a book, titled "Affective Computing." Without realizing it, she had given a name to a new field of computer science.

Kaliouby was still in Cairo, an undergraduate at the American University. In 1998, she graduated at the top of her class, earning a merit scholarship to pursue a

master's. She aspired to teach computer science, but she knew that a tenured job would require doctorate work abroad. "My dad was like, 'Well, if you go, by the time you get back, you will be too old to get married.'" Uncertain, she applied for work at a local tech startup. "It was in a residential building," she said. "My dad drove me there, then wanted to come up, and I was like, 'Please, it will look awful,' so he waited in the car. I was wearing a skirt, and I looked very formal—it was my first interview—and I saw all these guys walking around in shorts, barefoot: typical software engineers. The guy who interviewed me said, 'We have run out of chairs,' and, pointing to my skirt, he said, 'We can either have this interview on the floor, or, if you are uncomfortable, we can reschedule.' I was like, 'O.K., I can sit on the floor.'"

A few days later, Kaliouby withdrew her application, and enrolled in the master's program. But she had made an impression; one of the company's founders, Wael Amin, had grown up an expat in Argentina, and sympathized with the social pressures that she faced. He tracked her down, and encouraged her to continue her education; they were married not long after. In graduate school, Kaliouby searched for focus. "The idea that computers can change the way we connect with one another—that was where I was being drawn," she recalled. One day, Amin passed along a review of Picard's book, and she ordered a copy. "It took four months to get to Egypt—it was held in customs for reasons that I don't understand," she said. "But eventually I read the book, and I was inspired." Without meeting Picard, she considered her a role model. "She was a female scientist, successful, and created this field that I found exciting." Kaliouby had settled on her direction: to create an algorithm that could read faces.





[Buy the print »](#)

The human face is a moving landscape of tremendous nuance and complexity. It is a marvel of computation that people so often effortlessly interpret expressions, regardless of the particularities of the face they are looking at, the setting, the light, or the angle. A programmer trying to teach a computer to do the same thing must contend with nearly infinite contingencies. The process requires machine learning, in which computers find patterns in large tranches of data, and then use those patterns to interpret new data.

From Cairo, Kaliouby contacted some of the early research teams for guidance and data. Ekman had begun working to automate faces, building systems designed to locate discrete action units. With nineties-era technology, this was painstaking work. Undergraduate students (or Ekman himself) would perform expressions in an exaggerated way, against a controlled background. Each frame of video took twenty-five seconds to digitize, and, in key frames, a person had to hand-label every facial movement. “There were so many challenges,” an early researcher told me; one version of his system struggled to track the deformable points. “It was always a little off, and as we processed more and more frames the errors started to accumulate.” Every ten seconds, he had to re-start.

Kaliouby hoped to create a system that was powerful enough to work in the real world. But when she began pursuing her Ph.D. at Cambridge University, in 2001, her adviser wasn't familiar with affective computing; nor were her peers. "There was a lot of curiosity, and also questioning: why would you ever want to do that?" she told me. During a presentation of her research goals, an audience member mentioned that the problem of training computers to read faces seemed to resemble difficulties that his autistic brother had. Kaliouby knew nothing about autism, so she began to look into it, searching for clues. At the time, Cambridge's Autism Research Centre was working on a huge project to create a catalogue of every human facial expression, which people on the autism spectrum could study to assist with social interactions. Rather than trying to break expressions into their constituent parts, as Ekman had, the center was interested in natural, easily understood portrayals; under the rubric of "thinking," it distinguished among brooding, choosing, fantasizing, judging, thoughtfulness. It hired six actors—of both genders, and a range of ages and ethnicities—to perform the emotions before a video camera. Twenty judges reviewed each clip, and near-consensus was required before an emotion was labelled. At the project's end, four hundred and twelve had been identified.

Kaliouby recognized at once that the catalogue presented an unprecedented opportunity: rich, validated data, ideal for a computer to learn from. By the time she completed her doctorate, she had built MindReader, a program that could track several complex emotions in relatively unstructured settings. As she considered its potential, she wondered if she could construct an "emotional hearing aid" for people with autism. The wearer would carry a small computer, an earpiece, and a camera, to scan people's expressions. In gentle tones, the computer would indicate appropriate behavior: keep talking, or shift topics.

While developing the idea, Kaliouby learned that Picard was planning to visit her lab. "That was the highlight of my summer," she recalled. "She was supposed to

spend ten minutes with every student. We ended up spending an hour.” Picard thought that Kaliouby’s system was the most robust anyone had created. The two women decided to collaborate on the emotional aid, and the National Science Foundation awarded them nearly a million dollars to build a prototype.


The Media Lab was devised as a refuge for tinkerers. Its founder had once commanded, “Forget technical papers and to a lesser extent theories. Let’s prove by doing.” Kaliouby embraced the ethos, and, though Picard was in a much more senior position, Frank Moss told me that the two women worked together in a “mind meld.” Just about everyone in the lab was playing with tiny, wearable cameras, and, Picard told me, “We talked a lot about ‘jacking in.’ ” During visits home to Egypt, Kaliouby would call to participate in meetings. Picard remembered a demonstration with a robot: “Rana was Skyping in, or something, through a laptop camera, and we left the camera on the floor while we walked over to see the demo. I felt bad, like leaving Rana’s body on the floor. So I thought, I need to put the camera on me. Then, when I walk around, Rana has the advantage of being on my body.”

While Kaliouby focussed on MindReader, Picard tested various devices—such as a computer mouse that could measure user frustration—that attempted to discern feelings by tracking physical responses. The most promising one, later called the Q, was strapped to the body, to record reactions like skin conductance. Picard wore one nearly continuously, and kept a diary to track the data against her experiences.

Kaliouby and Picard believed that their systems were complementary, and in 2007 they began testing at a facility for children with behavioral disabilities. Picard hoped that her biosensor would provide insight into the origins of tantrums and other outbursts; an autistic child might seem calm, even disengaged, but the Q would indicate that her skin conductance was twice normal. Kaliouby’s system helped navigate social situations.

“One day really stuck with me,” Kaliouby recalled. “There was this boy who was really avoiding eye contact. That is a problem that is very common with a lot of these kids—they are experiencing information overload. This boy—we were experimenting with something like an iPad, but that was before iPads—was wearing the camera, and getting the feedback, basically using the iPad to shield off face contact. He was seeing me through the screen.” The device reinforced when he was communicating well, and as they talked he gained confidence. “Then he actually started lowering the device, until he and I made eye contact. And it was this special moment. It was like, Wow, this technology can really help.”

As the team developed MindReader, Kaliouby uploaded the software onto a server, where corporate sponsors were invited to test whatever Media Lab products they found interesting. To her surprise, she said, it quickly became the most downloaded item. Pepsi was curious if it could use the software to gauge consumer preferences. Bank of America was interested in testing it in A.T.M.s. Toyota wanted to see if it could better understand driver behavior—and perhaps design a system to detect drowsiness. Inquiries flooded in—from Microsoft, H.P., Yamaha, Honda, Gibson, Hallmark, NASA, Nokia—and Kaliouby did what she could to accommodate each one. “They had lots of questions,” she recalled. “ ‘Exactly what do the data mean?’ ‘How can we adapt this to our particular environment?’ This V.P. at FOX basically said, ‘I want to test all our pilot shows with this.’ And I was like, ‘We don’t have the resources to do that. We’re a research lab.’ ”



[Buy the print »](#)

The requests began to overwhelm her autism research. Kaliouby built a spreadsheet, to keep track of which sponsors wanted what, and in November, 2008, she and Picard brought it to Frank Moss, the Media Lab’s director. “We said, ‘Here are all the things our sponsors need—we need to double our group size,’ ” Kaliouby told me. “And he was like, ‘No, the solution is not to add more researchers. The solution is to spin out.’ ” Kaliouby was reluctant to leave academia. “We really wanted to focus on the do-good applications of the technology,” she said. But Moss argued that the marketplace would make the technology more robust and flexible: a device that could work for FOX could also better assist the autistic. It was possible, he said, to build a company with a “dual bottom line”—one that not only did well but also changed people’s lives.

Kaliouby and Picard set out to create a “baby I.B.M.” for emotionally intelligent machines: a startup for myriad products based on affective computing. Government agencies started asking about the technology, but, Kaliouby told me, she turned them away. Some of the corporate interest alarmed them, too. Picard recalled, “We had people come and say, ‘Can you spy on our employees without them knowing?’ or ‘Can you tell me how my customers are feeling?’ and I was like, ‘Well, here is why

that is a bad idea.’ I can remember one wanted to put our stuff in these terminals and measure people, and we just went back to Affectiva and shook our heads. We told them, ‘We will not be a part of that—we have respect for the participant.’ But it’s tough when you are a little startup, and someone is willing to pay you, and you have to tell them to go away.”

MindReader had been trained with actors, rather than from real-life behavior, and the code had to be rebuilt entirely. In 2011, the company tested it with Super Bowl ads online, building up a database of authentic emotional responses; later, Kaliouby collaborated with Thales Teixeira, a professor at the Harvard Business School, on a more rigorous study, screening ads for two hundred and fifty respondents. Affectiva’s C.E.O., David Berman, a former sales executive, was steering the company away from assistive technology and toward market research, which helped attract millions of dollars in venture capital. “Our C.E.O. was absolutely not comfortable with the medical space,” Picard said. Tensions rose. After four\* years, Picard was pushed out, and her group was reassigned. Matthew Goodwin, an early researcher at Affectiva who now sits on its scientific board, told me, “We began with a powerful set of products that could assist people who have a very difficult time with perceiving affect and producing affect. Then they started to emphasize only the face, to focus on advertisements, and on predicting whether someone likes a product, and just went totally off the original mission.”

Kaliouby was upset by Picard’s departure, but the company’s new momentum was undeniable. In March, 2011, she and her team were invited to demonstrate MindReader to executives from Millward Brown, a global market-research company. Kaliouby was frank about the system’s limitations—the software still was having trouble distinguishing a smile from a grimace—but the executives were impressed. Ad testing often relies on large surveys, which deal in reasoned reflections, rather than in the spontaneous, even unconscious, sentiment that really interests marketers; new technology promised better results. A year earlier,

Millward Brown had formed a neuroscience unit, which attempted to bring EEG technology into the work, and it had hired experts in Ekman's system to study video of interviews. But these ideas had proved impossible to scale up. Now the executives proposed a test: if Affdex could successfully measure people's emotional responses to four ads that they had already studied, Millward Brown would become a client, and also an investor. "The stakes were so high," Kaliouby told me. "I remember, our C.E.O. said, 'This is all hands on deck.'"

One of the TV spots that Millward Brown had chosen was for Dove. Titled "Onslaught," it begins with an image of a young girl. The camera then shifts to her perspective as she is bombarded by a montage of video clips—a lifetime of female stereotypes compressed into thirty-two seconds—before the ad ends with the girl, all innocence, and the tagline "Talk to your daughter before the beauty industry does." The ad was critically acclaimed, but in surveys Millward Brown found that many people considered it emotionally difficult to sit through. Affdex scanned more than a hundred respondents watching the ad, and detected the same response. But it also found that at the moment of resolution this discomfort went away. "The software was telling us something we were potentially not seeing," Graham Page, a Millward Brown executive, told me. "People often can't articulate such detail in sixty seconds, and also, when it comes to negative content, they tend to be polite." Millward Brown's parent company, WPP, invested \$4.5 million in Affectiva. Soon Affdex was being used to test thousands of ads a year.

Kaliouby invited me to try the version of Affdex that Millward Brown uses, and one afternoon in her office she directed me to a MacBook loaded with the first fifteen minutes of Spike Jonze's movie "Her"—about a man who falls in love with an emotionally enabled computer operating system. After completing a short survey, I watched while the laptop's camera watched me. Fifteen other people did the same. Then I logged onto Affdex. Against a black background, the quantified sentiment—colorful jagged lines—appeared like plots on a lie detector. The software allowed me

to isolate smiles, disgust, surprise, concentration.

Kaliouby had seen “Her,” and she wondered if its mood was too muted to invoke responses that the software could measure. In the film’s opening, Theodore Twombly (Joaquin Phoenix), who works for a company called Beautiful Handwritten Letters, narrates a sentimental letter—from a woman to her husband, on their fiftieth wedding anniversary—which his computer then prints out in her handwriting. Leaving the office, he passes a receptionist, played by Chris Pratt, who says, “Who knew you could rhyme so many words with the name Penelope? It’s *badass*.” The goofball earnestness of Pratt’s delivery—the last word imparted as if congratulating a Navy SEAL on a successful mission—was amusing, and Affdex noticed: nearly all of us smiled. In fact, in many moments, everyone appeared to be reacting in synch. During a wordless transition, a pan through an empty apartment, our reactions dipped. I looked at video of myself. I had shifted in my seat.



“*To you it was fast.*” November 13, 2000 [Buy the print](#) »

We all became most expressive during a scene in which Twombly has phone sex with a woman identified as SexyKitten—the voice of the comedian Kristen Wiig. In a bizarre, funny moment, SexyKitten seizes control of the call, demanding that



Twombly strangle her with an imaginary dead cat, and, as he hesitantly complies, she screams with ecstasy. Affdex tracked us smiling—women, on the whole, more than men. Kaliouby noticed that the smiles came in three pulses. She suggested that this might indicate a micro-narrative worth exploring, and it turned out that there was a structure behind them: whenever Twombly, distraught, spoke of the dead cat, the smiles waned. The Affectiva team often strives to build a story from this kind of data, but the story remained unclear. Were the smiles waning out of empathy—because of Twombly’s distress? Or discomfort with the implied violence? Or was the scene simply less funny when Wiig stopped talking? In such cases, Kaliouby told me, market researchers would rely on old-fashioned human intelligence: interviews with respondents.

Spike Jonze spent months researching “Her,” and it’s not hard to find real-world intimations of the future he imagined. Recently, researchers at the University of Southern California built a prototype “virtual human” named Ellie, a digital therapist that integrates an algorithm similar to Affdex with others that track gestures and vocal tonalities. One of its designers sent me video of a woman talking with Ellie. “What were your symptoms?” Ellie asks, and the woman describes her trouble with weight gain, insomnia, and oversleeping. Ellie appears to listen, nodding. The woman explains that she often feels the need to cry. Her voice wavers, her eyes fill up, and Ellie sympathetically draws her brow into a frown, pauses, and says in a comforting tone, “I’m sorry.”

In October, Kaliouby took the Acela to New York to speak at a conference, called Strata + Hadoop World, at the Javits Center. More than five thousand specialists in Big Data had come from around the country—believers in the faith that transformative patterns exist in the zeros and ones that sustain modern life. The talks ranged from “Industrial Internet” to “How Goldman Sachs Is Using Knowledge to Create an Information Edge.” Some of the attendees wore badges for well-known corporations (Microsoft, Dell, G.E.); others were for companies I hadn’t

heard of (Polynumeral, Metanautix). While waiting to enter the main hall, I stood beside one of the few women there. Her badge simply said “U.S. Government.”

In the darkened hall, audience members opened laptops, and their screens glowed. In the greenroom, Kaliouby reviewed her notes and did breathing exercises. Onstage, she declared that it was the first time a scientist in her field had been invited to join “the Big Data conversation”: a throwaway line, but one with a remarkable implication—that even emotions could be quantified, aggregated, leveraged.

She said that her company had analyzed more than two million videos, of respondents in eighty countries. “This is data we have never had before,” she said. When Affectiva began, she had trained the software on just a few hundred expressions. But once she started working with Millward Brown hundreds of thousands of people on six continents began turning on Web cams to watch ads for testing, and all their emotional responses—natural reactions, in relatively uncontrolled settings—flowed back to Kaliouby’s team.

Affdex can now read the nuances of smiles better than most people can. As the company’s database of emotional reactions grows, the software is getting better at reading other expressions. Before the conference, Kaliouby had told me about a project to upgrade the detection of furrowed eyebrows. “A brow furrow is a very important indicator of confusion or concentration, and it can be a negative facial expression,” she said. “A lot of our customers want to know if their ad is offending people, or not really connecting. So we kicked off this experiment, using a whole bunch of parameters: should the computer consider the entire face, the eye region, just the brows? Should it look at two eyebrows together, or one and then the other?” By the time Kaliouby arrived in New York, Affdex had run the tests on eighty thousand brow furrows. Onstage, she presented the results: “Our accuracy jumped to over ninety per cent.”

From the Javits, Kaliouby caught a cab to the global headquarters of McCann Erickson, the ad agency. The company occupies eight floors in a midtown skyscraper, and her face brightened as she walked from the elevator and into a retro-modernist lobby with a lofty ceiling. Painted on a wall was McCann's credo: "Truth Well Told." Mike Medeiros, a vice-president of strategy, met Kaliouby and directed her to a conference room with a view across the city. He runs a group called Team America, which provided the U.S. military with the slogan "Army Strong." Tall with blond hair, he has a plainspoken manner that one could see putting a military officer at ease on Madison Avenue.

Within the firmament of consumer persuasion, companies like Millward Brown and McCann often act as adversaries: one seeks to evaluate what the other invents. Many advertising "creatives," Medeiros told Kaliouby, regard ad testing as antithetical to inspiration. "You know what they say about 'Seinfeld,' " he said. "By the metrics, it should have been killed in the first season." His group had just experimented with Affdex, to pitch an account worth millions of dollars; he seemed interested in the technology, but not convinced of its necessity. "I tend to go back to the methods that are lower tech, that are simpler—when you can sit people in a room with a script," he said. "More than anything, I am watching them to see: are they sitting back, leaning forward? Those cues tell you what is happening, and then you get them to talk about them. My own experience is that this is more important than what they say. Someone can say, 'Oh, I love that,' and at the same time they could not care less."

"That is where our technology can come in and quantify it," Kaliouby said.

Medeiros's superior, Steve Zaroff, stopped by, and Kaliouby gave him a demo. Sitting at her laptop with childlike enthusiasm, he mugged and twisted his face into brow furrows and lip curls. "A disgusted smile—I like it!" he said.

“We tested some really gross YouTube videos of people eating larvae,” Kaliouby said. “People smile, but they are also like, ‘Eww!’ ” The same sentiment, her software found, had animated a humorous gross-out ad that Doritos aired during the Super Bowl.

A few days later, Medeiros put me in touch with McCann’s Barcelona affiliate, which had used emotion-sensing technology in an unexpected way. In 2012, the Spanish government, facing a severe budget crisis, imposed strict austerity measures, including a thirteen-per-cent increase in the tax on theatre tickets. Teatreneu, a comedy club in Barcelona, lost a third of its nightly audience, so it approached the McCann affiliate for help. Instead of drafting an ad campaign, the agency recommended that the club outfit its seats with Affdex-like software, then open its doors for free, promising that visitors would be charged only .30 euro per laugh, with an eighty-laugh maximum. If anyone tried to cover up a laugh, or turn away, the system would charge the full fee: twenty-four euros. Revenue went up. Theatres in America, France, and South Korea contacted McCann, wanting to know more.



*“I still say it’s only a theory.”* May 23, 2005 [Buy the print »](#)

A Spanish economist aware of the Barcelona experiment had approached Affectiva, to set up a study based on the idea, with better technology. Already, Kaliouby saw in

the experiment the contours of the near-future—not merely P.R. gimmickry, or creative tax evasion, but the glimmer of an Emotion Economy. “Foreshadow where this technology is heading,” she once told me. Tech gurus have for some time been predicting the Internet of Things, the wiring together of all our devices to create “ambient intelligence”—an unseen fog of digital knowingness. (Imagine systems that adjust the temperature of your house, based on behavioral and physiological data that your car fed into the fog during your commute home.) Emotion would be a part of this. On our cab ride from the Javits, I had pointed out the screen on the seat back in front of us: intrusive and emotionally inert. Kaliouby saw it as a possibility, predicting that before long myriad devices will have an “emotion chip” that runs constantly in the background, the way geolocation works now in phones. “Every time you pick up your phone, it gets an emotion pulse, if you like, on how you’re feeling,” she said. “In our research, we found that people check their phones ten to twelve times an hour—and so that gives this many data points of the person’s experience.”

Foreshadowing the Emotion Economy, it turns out, takes little imagination. Already, most of us pay for services by a strangely intimate form of exchange, trading neurological activity for stuff. “Attention is the hard currency of cyberspace,” the authors Thomas Mandel and Gerard Van der Leun observed in their 1996 book, “The Rules of the Net”—an early recognition that the digital economy would largely be driven by new forms of targeted marketing. In 2008, Chris Anderson, the editor of *Wired*, wrote a triumphalist manifesto, “Free! Why \$0.00 Is the Future of Business,” proclaiming the dawn of a boundless giveaway world. But, of course, many of the services that we use without payment in dollars are paid for by other means: the attention we offer to providers who track our habits and preferences.

The free economy is, in fact, an economy of the bartered self. But attention can never be limitless. Kaliouby put me in touch with Thales Teixeira, the business

professor who collaborated with her, and we met at the Harvard Club in New York. “There are three major fungible resources that we as individuals have,” he said. “The first is money, the second is time, and the third is attention. Attention is the least explored.” Teixeira had recently tried to calculate the value of attention, and found that, like the dollar, its price fluctuated. Using Super Bowl ads as a rough indicator of the high end of the market, he determined that in 2010 the price of an American’s attention was six cents per minute. By 2014, the rate had increased by twenty per cent—more than double inflation. The jump had obvious implications: attention—at least, the kind worth selling—is becoming increasingly scarce, as people spend their free time distracted by a growing array of devices. And, just as the increasing scarcity of oil has led to more exotic methods of recovery, the scarcity of attention, combined with a growing economy built around its exchange, has prompted R. & D. in the mining of consumer cognition. “What people in the industry are saying is ‘I need to get people’s attention in a shorter period of time,’ so they are trying to focus on capturing the intensity of it,” Teixeira explained. “People who are emotional are much more engaged. And because emotions are ‘memory markers’ they remember more. So the idea now is shifting to: how do we get people who are feeling these emotions?”

Many companies are moving to take advantage of this shift. “We put together a patent application for a system that could dynamically price advertising depending on how people responded to it,” Kaliouby told me one afternoon. I found more than a hundred other patents for emotion-sensing technology, many of them tied to advertising. Represented: A.O.L., Hitachi, eBay, I.B.M., Yahoo!, and Motorola. Sony had filed several; its researchers anticipated games that build emotional maps of players, combining data from sensors and from social media to create “almost dangerous kinds of interactivity.” There were patents for emotion-sensing vending machines, and for A.T.M.s that would understand if users were “in a relaxed mood,” and receptive to advertising. Anheuser-Busch had designed a responsive beer

bottle, because sports fans at games “wishing to use their beverage containers to express emotion are limited to, for example, raising a bottle to express solidarity with a team.”

Not long ago, Verizon drafted plans for a media console packed with sensors, including a thermographic camera (to measure body temperature), an infrared laser (to gauge depth), and a multi-array microphone. By scanning a room, the system could determine the occupants’ age, gender, weight, height, skin color, hair length, facial features, mannerisms, what language they spoke, and whether they had an accent. It could identify pets, furniture, paintings, even a bag of chips. It could track “ambient actions”: eating, exercising, reading, sleeping, cuddling, cleaning, playing a musical instrument. It could probe other devices—to learn what a person might be browsing on the Web, or writing in an e-mail. It could scan for affect, tracking moments of laughter or argument. All this data would then shape the console’s choice of TV ads. A marital fight might prompt an ad for a counsellor. Signs of stress might prompt ads for aromatherapy candles. Upbeat humming might prompt ads “configured to target happy people.” The system could then broadcast the ads to every device in the room.

I had wondered if the Verizon system was an anomaly—perhaps dreamed up by overeager employees. But a number of its features are already available in Microsoft’s Xbox One system, which has a high-definition camera that can monitor players at thirty frames per second. Using a technology called Time of Flight, it can track the movement of individual photons, picking up minute alterations in a viewer’s skin color to measure blood flow, then calculate changes in heart rate. The software can monitor six people simultaneously, in visible or infrared light, charting their gaze and their basic emotional states, using technology similar to Affectiva’s. If people are moving, it can determine how much force their muscles are exerting. The system has tremendous potential for making digital games more immersive. But Microsoft has also been developing non-gaming applications, envisioning TV ads

targeted to your emotions, and programming priced according to how many people are in the room. Google, Comcast, and Intel have moved in a similar direction. Two years ago, Erik Huggers, a vice-president of Intel Media, expressed a common dissatisfaction: “Today, television doesn’t really know anything about you.”

In 2013, Representative Mike Capuano, of Massachusetts, drafted the We Are Watching You Act, to compel companies to indicate when sensing begins, and to give consumers the right to disable it. When Capuano publicly referred to Verizon, the company complained that it was not alone. An industry association, along with Samsung, quietly began to lobby. (Samsung declined to comment on its lobbying, but its goals are clear; as one of its researchers noted, “If we know the emotion of each user, we can provide more personalized services.”) Capuano couldn’t persuade colleagues to sign on to his bill for a hearing. “The most difficult part is getting people to realize that this is real,” he told me. “People were saying, ‘Come on. What are you, crazy, Capuano? What, do you have tinfoil wrapped around your head?’ And I was like, ‘Well, no. But if I did, it’s still real.’ ”



*“This is my stop. May I be released?”* April 27, 2009 [Buy the print »](#)

Kaliouby expects that notions of privacy will shift. She once wondered what it might



be like to scan everyone watching YouTube: “Somehow we figure out a way where people don’t mind or don’t care, so that every time you go on YouTube your camera turns on, and as you are watching it is collecting this data. It is like a cookie. Most people today have cookies running on their computers.” The business model for Affectiva’s main competitor, Emotient—which has Ekman on its board, and may well have a more powerful algorithm—makes similar assumptions. Emotient has already tested its technology on willing employees and consumers for a major fast-food company and a big-box store.

At McCann, Kaliouby told the admen, “In our space, you could very easily be perceived as Big Brother, as opposed to the gatekeeper of your own emotional data—and it is two very different positions. If we are not careful, we can very easily end up on the Big Brother side.”

Steve Zaroff nodded, and said, “It is interesting to think about wearable devices—those will become more open-ecosystem.” He pointed to a Nike FuelBand around his wrist, then he looked at Mike Medeiros, who was wearing a Fitbit. “Fitbit is better, because it reads your pulse,” he went on. “But there is a tremendous amount of data that the brand knows about you.”

He began to ask Medeiros if he would allow Fitbit to sell his data to marketers. Before he could finish, Medeiros, laughing, said, “Not yet.”

“But you can see it,” Zaroff said.

“You can see Apple going down that path, too,” Medeiros said. Apple had recently launched Health, a fitness app pre-installed on new iPhones. By gathering data from other apps and devices, or from medical providers, it can track weight, respiratory rate, sleep, even blood-oxygen saturation. This information could be used to build emotional profiles; researchers at Dartmouth demonstrated that smartphones can be configured to detect stress, loneliness, depression, and

productivity, and to predict G.P.A.s. In September, Apple also unveiled the Apple Watch, which can measure heart rate and physical activity, and link these data to your location. The new products complement another line of Apple's research: mood-targeted advertising. Medeiros said, "With the new watch, they are going to know everything about you."

By late fall, Kaliouby's Google calendar was a dense map of color blocks: every minute accounted for. The company was preparing to conduct research for Facebook, an experiment in placing ads in videos. There was a meeting with Samsung, which has licensed Affdex. A company in San Francisco wanted to give its digital nurse the ability to read faces. A Belfast entrepreneur wanted to know if the software would work in night clubs. Kaliouby forwarded me an e-mail that a member of her team wrote about a state initiative in Dubai, the Happiness Index, meant to measure social contentment: "Dubai is known to have one of the world's tightest CCTV networks, so the infrastructure to acquire video footage to be analyzed by Affdex already exists. I feel that's a promising opportunity."

Kaliouby doesn't see herself returning to autism work, but she has not relinquished the idea of a dual bottom line. "I do believe that if we have information about your emotional experiences we can help you be in a more positive mood and influence your wellness," she said. She had been reading about how to deal with difficult experiences. "The consistent advice was you have to take care of yourself, be in a good place, so that you can handle everything else," she said. "I think there is an opportunity to build a very, very simple app that pushes out funny content or inspiring content three times a day." Her tone brightened, as she began looking to wider possibilities. "It can capture the content's effect on you, and then you can gain these points—these happiness points, or mood points, or rewards—that can be turned into a virtual currency. We have been in conversations with a company in that space. It is an advertising-rewards company, and its business is based on positive moments. So if you set a goal to run three miles and you run three miles,

that's a moment. Or if you set the alarm for six o'clock and you actually do get up, that's a moment. And they monetize these moments. They sell them. Like Kleenex can send you a coupon—I don't know—when you get over a sad moment. Right now, this company is making assumptions about what those moments are. And we're like, 'Guess what? We can capture them.' ” ♦

\*An earlier version of this article misstated when Picard was pushed out.

An earlier version of this article misstated the name of the Apple Watch.