

UNCATEGORIZED

Lie Detection

To a few human experts, our faces are open books. Now computer technology automates those abilities.

By Mark Williams Pontin

April 21, 2009

In the late 1960s, Paul Ekman—then a young psychology professor at the University of California, San Francisco, School of Medicine and just commencing his life’s work—filled a San Francisco Victorian with a library of films showing 40 psychiatric patients’ faces as they were interviewed. Ekman, now a leading figure in his profession, wanted to know whether he could isolate facial expressions to help diagnose mental disorders. A woman named Mary, who had attempted suicide three times before, smiled and spoke cheerily on her tape. As it happened, she was angling for a weekend pass—so that she could go home and kill herself.



Lie detection is tricky: Is Richard Nixon, here interviewed by David Frost in 1977, on the level? Research finds that “microexpressions” reveal our basic emotions, whether we like it or not.

“Mary was how I first discovered microexpressions,” Ekman told me when I caught up with him on the set of *Lie to Me*, the Fox television drama inspired by his decades of research into how facial expressions, gestures, and other nonverbal behaviors reveal our emotions and—most pertinently—our deceptions. “Some young psychiatrists I was teaching asked whether I could help identify when a suicidal patient was telling the truth or lying about improving,” he said. “Some of their patients had left the hospital and killed themselves within an hour. Mary, however, had confessed before she left that she’d been lying during a [previous] interview I’d filmed. Looking at the film, I couldn’t see any evidence. So I went through it frame by frame for a week, and these microexpressions showed up—two instances, each a 25th of a second, out of 12 minutes.”

In Mary's case, her features had fleetingly exhibited despair when the interviewing doctor asked about her plans. Ekman learned that the human subjects he studied betrayed their emotional state through microexpressions, however much they tried to suppress them. He identified 46 facial-muscle movements that, across cultures, signal such basic emotions as fear, distrust, and distress.

"What I didn't know at the beginning," Ekman told me, "was you could train people to recognize these microexpressions in real time." He developed the Facial Action Coding System, or FACS, in the 1970s as an exhaustive taxonomy of all facial expressions, including these telltale muscle behaviors. Since then, trained FACS users have generally demonstrated better than a 75 percent success rate in reading faces. *Lie to Me*—which stars the estimable Tim Roth as Dr. Cal Lightman, the character based on Ekman—is very average entertainment in the genre of Fox's great success *House*, where a maverick expert solves cases that establishment types cannot. In reality, however, a lot of FACS users are establishment types—cops, FBI agents, members of the U.S. Secret Service.

RESOURCES:

- Computer Expression Recognition Toolbox (CERT)
- *Lie to Me*, FOX

It requires no innate gift to apply Ekman's research in practice. "You could go online now www.mettonline.com and learn the microexpression recognition, which is one part, in an hour," Ekman says. With practice, most of us could decode these fleeting expressions in real time. "Initially, everybody believes they'll never do it," he says. "By the end, they're asking, 'Are you slowing these things down?' We're not, but your eyes have learned to see them."

Other studies bear out Ekman's claims. In research conducted in 2006, neuroscientist Tamara Russell of the University of London's King's College showed that an hour of microexpression training enabled people with schizophrenia to identify facial expressions as accurately as healthy people.

Some, however, are much better than others at reading microexpressions. Ekman's University of San Francisco colleague Maureen O'Sullivan has tested 20,000-odd folks over two decades and identified 50 individuals among that number who consistently demonstrate over 80 percent accuracy in detecting when others are lying, with a very few approaching perfect accuracy. Clearly, some specific, optimal set of capabilities underlies these rare individuals' success.

Since trained FACS experts generally replay footage for three hours in order to analyze just a single minute of a subject's facial twitches and blinks on video, it made sense to ask whether a computer system could automate the process of microexpression analysis and match O'Sullivan's human "wizards." Ekman first considered the challenge in the late 1980s. On a sabbatical in London, he visited Brunel College, where an engineer who had developed one of the first parallel-processing computers was training an artificial neural network to recognize terrorists. The engineer's problem was that subjects' varied facial expressions made it difficult for his system to recognize their identities, while Ekman's difficulty tended to be the reverse: he needed to disregard his subjects' individual physiognomies to recognize the emotions revealed by their expressions. So the two men worked together. "Within three days, we taught the machine to recognize three different emotions on different people," he says. "Back in the U.S., I wrote up a grant proposal for the NIH, who turned it down, claiming parallel-processing computers didn't exist." Ekman

expressed his frustration to a friend who was a Nobel Prize-winning physicist; the friend contacted Terry Sejnowski, the cross-disciplinary eminence of computational neurobiology at the Salk Institute, whose lab had the necessary computers. Ekman and Sejnowski teamed up and got the grant.

Mark Frank, a former postdoctoral student of Ekman's and now a professor at the University at Buffalo, in New York, has had the greatest success automating FACS. Frank, working out of Buffalo's Center for Unified Biometrics and Sensors, has worked with a group of computer scientists at the University of California, San Diego—mostly former students of Sejnowski's—to turn FACS into a technology called the Computer Expression Recognition Toolbox (CERT). I asked him how the project was going.

"We've done it," Frank told me. "We have a system that operates in real time. In terms of machine learning, we had to give the machines good audiovisual material with real emotions and expressions. Then it's just a matter of training, testing, training, testing." CERT works about as well as a human expert, he says, but it's a little faster.

A technology that accurately detects people's true emotions possesses tremendous political, social, and commercial potential. What if political commentators had applied it to footage of last year's U.S. presidential debates, for instance, to reveal if McCain or Obama was lying? Or if lawyers used it to analyze video depositions presented during court trials to determine whether a witness had lied, a finding that could be cited as evidence? Indeed, since the technology mines ordinary video, it might be commodified as a cheap Web service so everybody could use it: people might record job interviews, business negotiations, prenuptial-agreement signings, wedding ceremonies, or any other kind of civil transaction, with an eye toward reviewing them to ascertain the good faith of those involved. "You wonder what you do when the cat comes out of the bag," Frank says. "And can you get it back in?"

The argument for admitting such evidence in court seems straightforward. To be admissible, a technology must satisfy one of two legal standards; the Daubert test (from the 1993 U.S. Supreme Court case *Daubert v. Merrell Dow Pharmaceuticals*) is the one used in most jurisdictions. "Daubert requires that scientific testimony must qualify as reliable 'scientific knowledge,'" says Edward Imwinkelried, a law professor at the University of California, Davis, who is an expert on the admissibility of scientific evidence. "The Supreme Court defines 'scientific knowledge' as knowledge validated by a specific methodology, which it described in classic terms as, firstly, the formulation of an hypothesis and, secondly, the subsequent controlled experimentation or systematic field observation to verify or falsify the hypothesis." Given FACS's three decades of acceptance and CERT's record of accuracy, automated facial-expression analysis might well meet those criteria.

Making this argument, however, would require the support of expert witnesses like Frank or Ekman, and that's not forthcoming. Frank, for instance, supports CERT's use by the U.S. government for purposes of national security— it may happen by 2011, he guesses—but he doesn't want to see the technology spread much further: "Though we get a call every two weeks from people wanting to make the big bucks by marketing this as lie detection, I'm proud that nobody involved in the science has thus far gone beyond what it supports."

What the science confirms is that both FACS and CERTS can reveal much of any human subject's real emotions, but those results must be construed intelligently—especially in the context of detecting deception. Otherwise, Ekman summed up, users risk what he calls "Othello's error":

“Othello read Desdemona’s fear accurately. But he didn’t recognize the fear of being disbelieved is just like the fear of being caught. Yes, our faces reveal what emotions we’re experiencing, if you can read the signs. What our faces don’t necessarily reveal is what triggered that emotion.” If you don’t know that, interpretation can go far astray. “Rule out all the possible explanations before you conclude that what you’re seeing is a sign of lying about a criminal act,” Ekman warns. “Because very often, it’s not.”

Mark Williams is a contributing editor to *Technology Review*.