Welcome to the United States Border.

PLEASE ANSWER EACH QUESTION TRUTHFULLY.

Are you a citizen of the United States of America?

YES

Results
- Facial expression nominal
- Voice response nominal
- Eye movement nominal

PROCEED

Have you ever been arrested?

NO

Results
- Facial expression nominal
- Delayed verbal response
- Pupils dilated

CAUTION

Have you been employed in the past five years?

YES

Results
- Eyebrows elevated
- Tonal rise at end of verbal response
- Eye movement nominal

CAUTION

Have you ever used illegal drugs?

NO

Results
- Eyebrows elevated
- Delayed verbal response
- Tonal rise at end of response
- Pupils dilated

ALERT

SUSPICIOUS BEHAVIOR DETECTED.

Wait for an agent.

Immigration authorities are testing a sensor-studded, computer-controlled lie detector that does something remarkable: actually detect lies.

BY ADAM HIGGINBOTHAM
PHOTOGRAPHS BY TIM FLACH

A

Alan Bersin, commissioner of Customs and Border Protection, arrives at the gloomy US border post in Nogales, Arizona, early one winter morning wearing an expression of mildly pained concentration.

He got up before dawn and now looks as if he'd rather be anywhere else. In the immigration lanes downstairs, a procession of pickups and SUVs nudge dejectedly in from Mexico, taillights blinking through a relentless drizzle. Bersin arrived late, and he seems in no mood to assess the state of the art in automated psychophysiological evaluation technology. Yet there it is, pushed up against the wall of a cramped back office at the DeConcini Port of Entry: a gray metal box about the size and shape of an ATM, with two softly glowing video monitors, one on top of the other.

Bersin, a self-assured bureaucrat and a Rhodes Scholar who studied at Oxford with Bill Clinton, approaches the device. The lower monitor displays an icon of an oversize red button; the upper screen shows the head and shoulders of a smoothly rendered, computer-generated young man blinking and occasionally suffering a slight electronic shudder. He appears to be in his twenties and has an improbably luxuriant head of blue-black hair combed back in a sumptuous pompadour. This is the Embodied Avatar, the personification of the latest software developed to help secure the nation's frontiers by delivering what its creator calls "a noninvasive credibility assessment"—sifting dishonest travelers from honest ones. Which is to say, this late-model Max Headroom is a lie detector.

Bersin taps the red button to start the test, and in an agreeable Midwestern voice, the avatar asks Bersin a series of questions.

"Are you a citizen of the United States of America?"

"Yes," Bersin says.

"Have you visited any foreign countries in the past five years?"

"Yes."

"Do you live at the address you listed on your application?"

"Yes."

When the interview is over, Bersin turns to the other people in the room—his entourage, a delegation from the Canadian border agency, and the engineers who are anxiously overseeing this most critical test yet of their invention.

One technician explains to Bersin how the kiosk has instantly analyzed his responses, displayed on a rubber-jacketed iPad and broken down into categories of risk: green, yellow, and red. Bersin's mask of barely suppressed boredom does not crack.

But then the technician points out that one of his answers is flagged in red: The machine is suspicious about his address. Bersin acknowledges that, yes, what he usually describes as his home is not actually where he lives, and that he was thinking about something else when he was answering—it's just that he has a work residence in Washington, DC, but of course his family home remains back in San Diego and—

Bersin's counterpart from Canada, a former intelligence officer, interrupts, cracking an interrogator's indulgent smile: "Do you have a lawyer?"

Afterward, Jay Nunamaker, the sardonic computer engineer overseeing the Embodied Avatar project, allows himself a low chuckle. "I don't think it could have gone better," he says. Within a few hours, the young man with the improbable hair is interviewing members of the public. The first field tests of the US government's state-of-the-art computer-controlled lie-detection device have begun.

COURTESY OF JOYCE P. CHAN/THE UNIVERSITY OF ARIZONA

SINCE SEPTEMBER 11, 2001, federal agencies have spent millions of dollars on research designed to detect deceptive behavior in travelers passing through US airports and border crossings in the hope of catching terrorists. Security personnel have been trained—and technology has been devised—to identify, as an air transport trade association representative once put it, "bad people and not just bad objects." Yet for all this investment and the decades of research that preceded it, researchers continue to struggle with a profound scientific question: How can you tell if someone is lying?

That problem is so complex that no one, including the engineers and psychologists developing machines to do it, can be certain if any technology will work. "It fits with our notion of justice, somehow, that liars can't really get away with it," says Maria Hartwig, a social psychologist at John Jay College of Criminal Justice who cowrote a recent report on deceit detection at airports and border crossings. The problem is, as Hartwig explains it, that all the science says people are really good at lying, and it's incredibly hard to tell when we're doing it.

In fact, most of us lie constantly—ranging from outright cons to minor fibs told to make life run more smoothly. "Some of the best research I've seen says we lie as much as 10 times every 24 hours," says Phil Houston, a soft-spoken former CIA interrogator

ADAM HIGGINBOTHAM (adam@adamhigginbotham.com) wrote about the Chernobyl Exclusion Zone in issue 19.05.

The Interrogation Bot

Just three sensors tell the Embodied Avatar kiosk everything it needs to know about whether someone is telling the truth. An infrared camera records eye movement and pupil dilation at up to 250 frames per second—the stress of lying tends to cause the pupils to dilate. A high-definition video camera captures fidgets such as shrugging, nodding, and scratching, which tend to increase during a deceptive statement. And a microphone collects vocal data, because lies often come with minute changes in pitch. Future versions of the machine might go even further—a weight-sensing platform could measure leg and foot shifts or toe scrunches, and a 3-D camera could track the movements of a person's entire body. *—Sara Breselor*

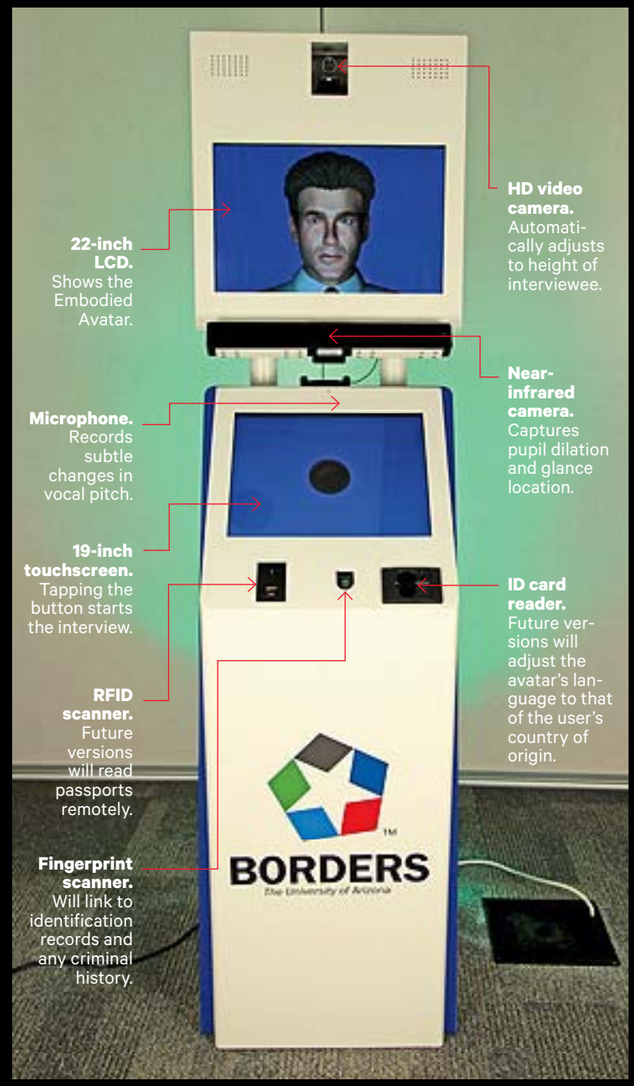

who is now CEO of QVerity, a company selling lie-detecting techniques in the business world. "There's some research on college students that says it may be double and triple that. We lie a ton." And yet, statistically, people can tell whether someone is telling the truth only around 54 percent of the time, barely better than a coin toss.

For thousands of years, attempts to detect deceit have relied on the notion that liars' bodies betray them. But even after a century of scientific research, this fundamental assumption has never been definitively proven. "We know very little about deception from either a psychological or physiological view at the basic level," says Charles Honts, a former Department of Defense polygrapher and now a Boise State University psychologist specializing in the study of deception. "If you look at the lie-detection literature, there's nothing that ties it together, because there's no basic theory there. It's all over the place."

Despite their fixation on the problem of deceit, government agencies aren't interested in funding anything so abstract as basic research. "They want to buy hardware," Honts says. But without an understanding of the mechanics of lying, it seems that any attempt to build a lie-detecting device is doomed to fail. "It's like trying to build an atomic bomb without knowing the theory of the atom," Honts says.

Take the polygraph. It functions today on the same principles as when it was conceived in 1921: providing a continuous recording of vital signs, including blood pressure, heart rate, and perspiration. But the validity of the polygraph approach has been questioned almost since its inception. It records the signs of arousal, and while these may be indications that a subject is lying—dissembling can be stressful—they might also be signs of anger, fear, even sexual excitement. "It's not deception, per se," says Judee Burgoon, Nunamaker's research partner at the University of Arizona. "But that little caveat gets lost in the shuffle."

The US Army founded a polygraph school in 1951, and the government later introduced the machine as an employee-screening tool. Indeed, according to some experts, the polygraph can detect deception more than 90 percent of the time—albeit under very strictly defined criteria. "If you've got a single issue, and the person knows whether or not they've shot John Doe," Honts says, "the polygraph is pretty good." Experienced polygraph examiners like Phil Houston, legendary within the CIA for his successful interrogations, are careful to point out that the device relies on the skill of the examiner to produce accurate results—the right kind of questions, the experience to know when to press harder and when the mere presence of the device can intimidate a suspect into telling the truth. Without that, a polygraph machine is no more of a lie-detector than a rubber truncheon or a pair of pliers.

As a result, although some state courts allow them, polygraph examinations have rarely been admitted as evidence in federal court; they've been dogged by high false-positive rates, and notorious spies, including CIA mole Aldrich Ames, have beaten the tests. In 2003 the National Academy of Sciences reported that the evidence of polygraph accuracy was "scanty and scientifically weak" and that, while the device might be used effectively in criminal investigations, as a screening tool it was practically useless. By then, other devices and techniques that had been touted as reliable lie detectors—voice stress analysis, pupillometry, brain scanning—had also either been dismissed as junk science or not fully tested.

But spooks and cops remain desperate for technology that could boost their rate of success even a couple of points above chance. That's why, in 2006, project managers from the Army's polygraph school—by then renamed the Defense Academy for Credibility Assessment—approached Nunamaker and Burgoon. The government wanted them to build a new machine, a device that could sniff out liars without touching them and that wouldn't need a trained human examiner: a polygraph for the 21st century.

A

A FORMER COLLEGE WRESTLER from Pittsburgh, Nunamaker is a leathery 75-year-old who trained as a mechanical engineer, and his methodical approach to problem-solving has carried him through four decades of developing software. He became interested in deception in the '90s, while building teleconferencing and collaboration software for corporate-scale behemoths including IBM and the US Army and Air Force. His clients suspected that many of their employees' contributions were often deliberately misleading, warped by self-interest and interdepartmental rivalry. Nunamaker discovered that he could pick out liars by looking for a statistical prevalence of evasive language and "hedging words," and he became fascinated with the ways deceitful employees betray themselves.

Burgoon, 64, a brisk and polished psychologist with cropped silver hair, had already done a decade of military-funded deception-detection work when she began collaborating with Nunamaker 12 years ago—both were at the University of Arizona working, it turned out, on similar projects. Burgoon specialized in examining deceit as part of interpersonal communication, a laborious, time-consuming, and—compared with the twitching needles and brain scanning at the other end of the field—unglamorous area of research. Nunamaker suggested they collaborate; her psychology background complemented his engineering expertise.

Instead of simply measuring signs of physiological arousal, Burgoon analyzed liars' body movement—expressions and gestures—and linguistic cues. Like Nunamaker, she had established that liars tend to hedge, equivocate, or fail to deny things directly. One recent study, using publicly available recordings of 911 calls from Florida and Ohio, found that using vague language about things like location and details of the crime often correlated with the caller being the perpetrator. Other research has shown that dishonest stories tend to be better structured than honest ones, although in the end the true story may seem more coherent. True narratives feature richer sensory detail, more direct speech, and more spontaneous corrections. "Deceivers are not going to say, 'Well, I can't remember that, I forgot that,'" Burgoon says. "They'll make something up."

PHOTO RETOUCHING BY JEFF LYSGAARD

For years, Burgoon collected linguistic data the hard way, transcribing interviews and marking them up by hand. Analyzing body movement was even more painstaking. Trained coders watched video of experimental subjects for hundreds of hours and logged each cue they saw—one blink, a slight smile—manually, using a pen and paper. One research project involved 300 videotaped interviews; the coding took three years. Analyzing the audio was even harder. Burgoon worked with specialists who attempted to hear the changes in vocal pitch in individual phonemes, the units of sound that make up words. Expanding to whole conversations proved almost impossible. "They were used to dealing with a phoneme, not with an entire utterance, much less an entire interview," Burgoon says.

When they first began working together in 2000, Nunamaker found it hard to believe that nobody had tried using machines to simplify this data collection. "It drove me crazy," he says. So Burgoon and Nunamaker started tracking body movement with software that superimposed computer-generated blobs over the video of interviewees; now computer vision tools can find a human being in an image and track more than 80 different landmarks on the face alone. Using transcripts of interviews with indicted Air Force personnel—whose lies were pretty well documented—they worked on artificial-intelligence tools to analyze language, counting hedging words and tracking pronoun use. (This evolved into a suite of software Burgoon dubbed the Agent99 Analyzer—she liked the idea of naming it after the female sidekick from Get Smart, famously more competent than her male boss.)

Eventually, Nunamaker and Burgoon came to believe that no single technology could solve the problems of lie detection. "There is no silver bullet answer—which is what everybody wants," Nunamaker says. "It's going to be this basket of cues and figuring out whether you've got the right cues in the basket." But they also knew that computers could detect the signs they'd identified. The researchers decided to combine as many sensors as possible in a single lie-detecting toolbox. By monitoring potentially hundreds of different psychophysiological, linguistic, and verbal cues, their hypothetical machine would spot tells in even the most polished liar. "A human can only control three or four at a time—so cues leak out no matter how hard you try," Nunamaker says.

One of the first government agencies interested in Nunamaker and Burgoon's work was the Department of Homeland Security. DHS paid for early data collection at the border station at Nogales, a project in which the researchers filmed travelers during screening interviews and then compared their linguistic and physical cues to the way customs officers rated them after screening. But the Science and Technology Directorate at DHS believed that even a working lie detector wouldn't be good enough to fight terrorism. They didn't just want to know when someone was lying—they wanted to look for signs that the person intended to do bad things, or "malintent." So before Nunamaker and Burgoon finished their fieldwork in Nogales, they say, DHS asked them to abandon it and instead study the relationship between emotions, physical cues, and malintent—specifically incorporating the microexpression theories of Paul Ekman.

Ekman is a divisive figure. Now 78, his work on lie detection has made him a rock star among behavioral psychologists, with a best-selling book, a profitable consulting business, and a network TV drama—Lie to Me—inspired by his research. In 1969 he theorized that facial muscles that expressed seven human emotions also created "microexpressions" that could reveal concealment, despite the fact that these microexpressions last just 0.04 second. Ekman claims that with his training, it's possible to spot microexpressions and successfully detect deception 70 percent of the time, increasing this to almost 100 percent if other body movements are taken into account.

In 2006 Ekman and his team spent 30 days training TSA officers to read microexpressions as part of a program called SPOT—Screening Passengers by Observational Techniques. These officers deployed to 161 airports across the

US. According to the TSA, from January 2006 through March 2012, SPOT officers referred more than 331,280 travelers for secondary screening, but the merest fraction were arrested—just 2,270. In that time, at least 16 people involved in terrorism cases passed unchallenged through airport checkpoints manned by SPOT personnel—some of them more than once. Nobody outside the TSA knows what SPOT officers are looking for, since the details remain classified. The agency admits to being uncertain that it has yet detained a single terrorist. Charles Honts, who was trained by Ekman, says that all his attempts to replicate Ekman's experiments have failed, and in 2009 the researchers studying airports and border crossings found no evidence that microexpressions reliably betray concealed emotion or can be used to detect deceit. The next year the Government Accountability Office reported that the TSA's scheme had never been scientifically tested. (Ekman disputes these criticisms. Of the 9/11 hijackers, for example, he says, "If the behavior that people reported them showing did occur, it would certainly have been picked up by SPOT.")

Nunamaker and Burgoon didn't want to abandon their fieldwork, and they didn't want to focus on microexpressions. "There's not a lot of science to back up Ekman's claims," Nunamaker says. "Applying them to deception detection is a reach." The project manager pulled their funding—because, Nunamaker says, he wouldn't switch the focus of his work. DHS moved ahead with Ekman's research. The new behavioral forecasting program—Future Attribute Screening Technology—is so secret that even Burgoon has no idea what it does.

WE LIE 10 TIMES A DAY.

But we can tell someone is lying just half the time. Using physiological and vocal data, computers can do better.

AFTER FALLING OUT WITH DHS, Nunamaker and Burgoon pressed on. They won new funding from the Pentagon and other agencies. Customs and Border Protection, for example, wanted to help overburdened customs officers screen immigration lines at borders, so the two decided to combine their lie-detecting toolbox with an idea other deception researchers were already playing with: a computer-generated interrogator.

An avatar interrogator has many advantages over its human counterparts. It's consistent, tireless, and susceptible to neither persuasion nor bribery. Douglas Derrick, a researcher at the University of Nebraska who studies human-computer interaction and has worked on the Embodied Avatar since 2007, even suspects that people fear the power they feel it embodies. "They view it as the personification of the system," Derrick says. "They believe they're talking to the computer." One early version was a menacing shaven-headed character nicknamed Scary Guy. On the other hand, Nunamaker says, Las Vegas casinos, which fund their own deception research to catch cheats, have had more success giving casino-goers screen-based directions and advice with avatars that resemble cartoons. Derrick even tried using a camera and morphing software so that an avatar would increasingly resemble the person in front of it, reflecting research that suggests you're more likely to trust someone who looks like you. The one thing they all had in common was skirting the edges of the uncanny valley, where characters look just human enough to be disturbing. "I think we're close with this one," Derrick says of the thick-haired young man used in Nogales. "It's realistic, but we're not in the valley."

The team dug into commercially available lie-detection technology, but most proved unusable. A thermal-imaging camera was enormous and required a cooling fan so noisy that it drowned out the other equipment. The laser Doppler vibrometer, which could monitor blood pressure from 10 feet away, could be circumvented by anyone wearing a turtleneck or even a beard. And the lie/truth analyzer, built into vocal dynamics software provided by the Israeli company Nemesysco, was hopeless under experimental conditions.

Despite those failings, Nunamaker and Burgoon thought that some of the gadgets had potential. Arizona grad student Aaron Elkins found that the Nemesysco software really was finding a correlation between vocal stress and deception, so he wrote his own algorithms to do the same thing—to measure cues like hesitation, changes in tempo and intonation, and spoken errors. It worked; Elkins' approach can identify deceitful speech 75 percent of the time in an experimental setting, and speech dynamics now provide key data points used by the technology being tested in Nogales.

Now, using just three sensors, they can collect as many as 50 different psychophysiological and vocal deception cues. A microphone gathers vocal information. An HD video camera captures body movement—for example, the sudden freezing of a liar attempting to control physical tells. And an infrared camera monitors pupil dilation and gaze pattern. Some of the team's most successful experiments have shown that eye flicker correlates to deception: Examining images of falsified documents, for example, subjects often cannot help looking repeatedly at the details they've doctored.

Separately, these streams of data can provide a good picture of when test subjects might be lying—in the lab, information from the eyes alone correctly flagged liars 60 percent of the time. But when the avatar kiosk combines the data from eye and voice analysis, its accuracy spikes. In an experiment in Poland last year using 37 EU border guards, some of whom were asked to present false documents, the kiosk identified every one of the liars. Taking into account two false positives, the machine scored 94 percent. Human agents asked to perform the same task failed to stop a single impostor.

Yet the success of such experiments has depended on the context of the interrogation. In October 2011, as Nunamaker's team began readying a version of the device for the field tests in Nogales, he admitted that he still wasn't certain how it would perform in the outside world: "We really don't know, until we test it at the border with real people who don't have a vested interest in the system working."

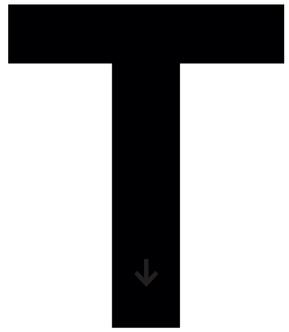

TWO MONTHS LATER IN NOGALES, a uniformed customs officer introduces the avatar kiosk to the public for the first time. In line are a mother and daughter from Tucson, a well-dressed couple from over the border in Mexico, and a portly retiree in a baseball cap there to renew his trusted-traveler card. Each is here as part of a fast-track border crossing program and is first screened by a cheerful immigration specialist trained in interviewing and behavioral analysis techniques. Then they meet the young man with the luxuriant hair. The avatar is set up to deliver the standard final set of questions asked of anyone trying to join a trusted-traveler program. Giving their yes-or-no answers to a five-minute robotic catechism, they seem curious or bemused or visibly anxious. One girl, eager to meet the machine after she heard there was a lie detector in the building, behaves as if she's trying her luck on a carnival midway, giddy and excited. On his way out, the old man remarks, "For guys, you might want to make it look like Salma Hayek."

The Nogales field test, intended to reveal the kind of limitations only everyday use can show, has led to further revisions of the kiosk. It's now bilingual, speaking both English and Spanish, and new lab versions have a camera that can collect eye data regardless of the height of the person it's interrogating. Eventually, if the machine flags a traveler as potentially deceptive, that person will be questioned further by a human customs officer. If the traveler triggers no alert, the machine will tell them they're free to go.

When the avatar catches him out, even commissioner Bersin—soon afterward promoted to assistant secretary of international affairs and chief diplomatic officer at DHS—seems to see its potential. He tells a group of customs officers at the DeConcini crossing station that he hopes the kiosk will soon check more and more people coming across the border. "We start off in this more controlled setting, but eventually the payoff is getting it into the lanes," he says.

Customs and Border Protection initially expressed interest in installing five kiosks in each of nine different customs stations, where they would conduct preliminary screening for the Nexus and Sentri programs. Budget issues have now postponed those plans, but last year the research team spent a month showing the machine to several DHS agencies in Washington, DC, including Immigration and Customs Enforcement, TSA, and the Secret Service. Nunamaker, Burgoon, and their colleagues now have funding to research countermeasures and identify the ways people might successfully beat the machine. Meanwhile, the TSA's FAST program, built around Paul Ekman's ideas, has been beset by controversy and technical difficulties. TSA officials now say the agency has no plan to deploy it.

Back in December, the last interviewee of the day at the DeConcini crossing station in Nogales is a stocky Mexican engineer wearing an Otis Elevator ID card around his neck. "So this is the future, huh?" he says at the end of his five-minute interrogation, his face unreadable. "Nice."

It is not yet 5 pm, but it's been a long day of cross-examinations. The customs officer pulls her uniform jacket on over her gun and equipment belt and heads into the rain for home. In the corner of the office, the tech from the university clicks on a wireless mouse a few times. The screen on the Embodied Avatar kiosk flickers, and the device goes to sleep where it stands.