

## An Eye-Scanning Lie Detector Is Forging a Dystopian Future

Mark Harris

Wired, Black Channel, 12/04/2018

Sitting in front of a Converus EyeDetect station, it's impossible not to think of *Blade Runner*. In the 1982 sci-fi classic, Harrison Ford's rumpled detective identifies artificial humans using a steam-punk Voight-Kampff device that watches their eyes while they answer surreal questions. EyeDetect's questions are less philosophical, and the penalty for failure is less fatal (Ford's character would whip out a gun and shoot). But the [basic idea](#) is the same: By capturing imperceptible changes in a participant's eyes—measuring things like pupil dilation and reaction time—the device aims to sort deceptive humanoids from genuine ones.

It claims to be, in short, a next-generation lie detector. [Polygraph tests](#) are a \$2 billion industry in the US and, despite their inaccuracy, are widely used to screen candidates for government jobs. Released in 2014 by Converus, a Mark Cuban-funded startup, EyeDetect is pitched by its makers as a faster, cheaper, and more accurate alternative to the notoriously unreliable polygraph. By many measures, EyeDetect appears to be the future of lie detection—and it's already being used by local and federal agencies to screen job applicants. Which is why I traveled to a testing center, just north of Seattle, to see exactly how it works.

Jon Walters makes an unlikely Blade Runner. Smartly dressed and clean cut, the former police chief runs Public Safety Testing, a company that conducts preemployment tests for police forces, fire departments, and paramedics in Washington State and beyond. Screening new hires used to involve lengthy, expensive polygraph tests, which typically require certified examiners to facilitate them. Increasingly, however, Walters tells me, law enforcement agencies are opting for EyeDetect.

Unlike a polygraph, EyeDetect is fast and largely automatic. This bypasses one of the pitfalls of polygraphs: human examiners, who can carry their biases when they interpret tests. According to Walters, biases don't really "come into play" with EyeDetect, and the test takes a brisk 30 minutes as opposed to the polygraph's 2- to 4-hour-long slog. Moreover, EyeDetect is a comfortable experience for the test subject. "When I was wired up for the polygraph, it was kind of intimidating," Walters told me. "Here you just sit and look into the machine."

*"When I was wired up for the polygraph, it was kind of intimidating," Walters told me. "Here you just sit and look into the machine."*

I settle in for a demonstration: a swift 15-minute demo where the test will guess a number I'm thinking of. An infrared camera observes my eye, capturing images 60 times a second while I answer questions on a Microsoft Surface tablet. That data is

fed to Converus' servers, where an algorithm, tuned and altered using machine learning, calculates whether or not I'm being truthful.

The widely accepted assumption underlying all of this is that deception is cognitively more demanding than telling the truth. Converus believes that emotional arousal manifests itself in telltale eye motions and behaviors when a person lies.

Converus [claims](#) that EyeDetect is "the most accurate lie detector available," boasting 86 percent accuracy. By comparison, many academics consider polygraph tests to be 65 to 75 percent accurate. The company already claims close to 500 customers in 40 countries, largely using the EyeDetect for job screening. In the US, this includes the federal government as well as 21 state and local law enforcement agencies, according to Converus. The Department of State [recently paid](#) Converus \$25,000 to use EyeDetect when vetting local hires at the US Embassy in Guatemala, WIRED's reporting revealed. Converus says its technology has also been used in an internal investigation at the US Embassy in Paraguay.

In documents [obtained through public records](#) requests, Converus says that the Defense Intelligence Agency and the US Customs and Border Protection are also trialing the technology. Converus says that individual locations of Best Western, FedEx, Four Points by Sheraton, McDonald's, and IHOP chains have used the tech in Guatemala and Panama within the last three years. (A 1988 [federal law prohibits](#) most private companies from using any kind of lie detector on staff or recruits in America.) WIRED reached out to all five companies, but none were able to confirm that they had used EyeDetect.

However, a close reading of records of EyeDetect's use, obtained through public records requests, suggest that a reliable, useful, and equitable lie detector is still the stuff of science fiction. WIRED found that like polygraphs, EyeDetect's results may introduce human bias and manipulation into its results. "Converus calls EyeDetect a next-generation lie detector, but it's essentially just the same old polygraph," says Vera Wilde, a transparency activist and independent researcher who has been [studying polygraphs](#) for many years. "It's astounding to me that there are paying customers deploying this technology and actually screening people with it," adds William Iacono, professor of psychology, psychiatry, neuroscience, and law at the University of Minnesota.

But the fact that EyeDetect is cheaper and faster than a polygraph might make Converus' new lie detector a tantalizing option for hiring offices across the country—a technology that could move into widespread use just as quietly as it leapt into existence.

Taking an EyeDetect test is as painless as Jon Walters promised. He asks me to pick a number between 1 and 10 and write it on a scrap of paper before I sit down in front of the EyeDetect camera. Walters instructs me to lie about my chosen number, to allow the system to detect my falsehood. If I beat it, Walters promises to give me \$50. (Journalistic ethics mean I'd pass any winnings along to a charity.)

A series of questions flash across a screen, asking about the number I picked in straightforward and then roundabout ways. I click true or false to each question. The EyeDetect camera feels no more intrusive than a normal webcam, and I do my best to keep my face and expression neutral, whether I'm lying or telling the truth.

Almost immediately after the test is over, the screen flashes a prediction based on my eye motions and responses. EyeDetect thinks that I chose the number 3. I had, in fact, picked the number 1. But when I reach for Walters' crisp \$50 note, he stops me. It turns out that Walters' interpretation of "a number between 1 and 10" includes only the digits 2 through 9. I had fooled the machine, but only by not playing by its rules. On my next attempt, the system correctly detects my hidden number.

Having my mind read is unsettling, and makes me feel vulnerable. It's like I've been tricked by a magician—but that doesn't mean I'd trust an illusionist to vet my local police chief.

Converus derives its 86 percent accuracy rate from a number of lab and field studies. But an upcoming academic book chapter written by the company's chief scientist and cocreator of EyeDetect, John Kircher, shows that from study to study the accuracy rates can vary quite a bit, even dipping as low as 50 percent for guilty subjects in one experiment.

The only peer-reviewed academic studies of Converus' technology have been carried out by the company's own scientists or students in their labs. These present largely positive results. "This is a huge problem," [says](#) John Allen, a professor of psychology at the University of Arizona. "If the only evidence in a medical trial came from a researcher with a financial interest in the product, no one would dare to think it has proven efficacy."

Even so, some in-house experiments reveal potential flaws with the device. In a study from 2013, the National Security Agency used an early version of EyeDetect to identify NSA employees who had taken a cellphone into a secure area, a minor security violation. The test accurately identified just 50 percent of those guilty of the mistake (the same as you would expect from chance) and just over 80 percent of those innocent.

*The test accurately identified just 50 percent of those guilty of the mistake—the same as you would expect from chance.*

In his book chapter, Kircher writes that the NSA's study, which promised an hour off work to those who passed, did not produce what Kircher calls a meaningful incentive. "In order for these tests to work, there needs to be jeopardy and proper protocols must be followed," Converus president and CEO Todd Mickelsen told Wired.

It's difficult to instill a feeling of peril in a study subject, a condition that presumably makes testing difficult. In 2016, a Converus marketing manager wrote to an

investigator at the Kent police department, in the suburbs of Seattle: “Please note, when an EyeDetect test is taken as a demo ... the results are often varied from what we see when examinees taking the test under real test circumstances where there are consequences.”

Jeopardy is a slippery concept, Wilde says: “There are many things, such as anxiety about results for both liars and truth-tellers, which could conceivably influence the physiological responses at issue.”

For the past four and a half years, Converus has been researching countermeasures that subjects might use to beat EyeDetect, such as squinting, using eye drops, or failing to respond. Based on that research and the belief that rapid-fire questioning allows little opportunity for deception, Converus says that its system have been tuned to “virtually eliminate” these countermeasures’ effectiveness.

After reading two of EyeDetect’s academic papers, Allen told WIRED: “My kindest take is that there is some promise, and that perhaps with future independent research this test might provide one measure among many for formulating a hypothesis about deceptive behavior. But even that would not be definitive evidence.”

Even assuming Converus’ most optimistic accuracy rating, an EyeDetect screening would turn out a large number of false positives when used to evaluate a large group of people for a rare crime, like terrorism. Kircher himself advises against relying solely on EyeDetect, or any single screening technology for detecting such offenses. “Even if a test is 90 percent accurate, about 10 percent of the tested population would fail it, and the vast majority of those individuals who fail the test would be innocent of the crimes,” he writes. (Converus says that EyeDetect’s false positive rate of 10 percent is the lowest of any credibility assessment technology on the market today, including polygraph.)

The company decided not to publish results of their first field experiment in Colombia, a study that appeared to show EyeDetect working erratically. “Although the data were limited, the [test] appeared to work well when we tested well-educated people who had applied to work for an airline, but the [test] was ineffective when we tested less well-educated applicants for security companies,” Kircher writes. Kircher speculated that the aspiring security guards might have had reading problems that meant they could not understand the test, and Converus says it now accounts for reading ability during testing. But without published data, other researchers aren’t able to evaluate what exactly caused the system’s spotty performance.

Correspondence with law enforcement investigators, supplied as part of WIRED’s public records requests, reveals EyeDetect has given surprising results in real life, too. In January 2017, Alan McCarty, a sergeant at the Columbus, Georgia, Police Department, wrote to Converus’ vice president of marketing and operations, Russ Warner, about an applicant who had admitted to using marijuana within the previous two years but still passed the EyeDetect test, which normally asks about illicit drug

use. (In his response at the time, Warner suggested that perhaps the applicant had problems with his left eye, which could have affected the results.)

Over at the Salt Lake City Police Department, Converus' first law enforcement customer, a sergeant told Warner about a similar case, where an applicant admitted to a disqualifying action but still aced the EyeDetect test with a score of 78. (50 is a pass.) Warner detailed a way this could happen: "We set the scoring algorithm to be less sensitive for [this person]. If we had used a standard algorithm, that person would have scored less than 49 (deceptive)."

*"If you're going to administer tests to existing sworn officers, we should create a new test, with a softer algorithm. This is what we've done in other agencies."*

Emails show that Converus has encouraged police departments to set an easier test for personnel transferring from other law enforcement agencies. "If you're going to administer tests to existing sworn officers, we should create a new test, with a softer algorithm. This is what we've done in other agencies," Warner told Columbus' McCarty early in 2017. Mickelsen says that modifying the base rate of guilt for some examinees "improves accuracy" and is "a standard practice" in polygraphy.

McCarty did not seem to be convinced: "We don't differentiate in the [polygraph] between [law enforcement officers and civilians]. [The applicant, a deputy sheriff] was asked about committing serious crimes, drugs use, theft and violating her oath as a law enforcement officer. Not really following the logic on this one," he wrote.

Not only can departments choose between administering a hard or soft test, another email exchange appears to show Converus changing test results when asked to do so. In January 2017, Alan McCarty had a candidate who passed an EyeDetect test, scoring 61. "I called him deceptive on the questions concerning drugs, theft and affiliation with gangs, terrorist organizations or subversive groups," McCarty wrote to Warner. "This is a 23-year-old kid who grew up in Atlanta that could have very well had some affiliation with gangs. Give me your thoughts."

After looking over the data, Warner responded. "His pupil data doesn't reveal deception. However, his linear eye movement does indicate some deception," he wrote. "The algorithm we are using right now to score the tests assumes a base rate of guilt of 20-25% ... If we modify the algorithm to consider a higher rate of test failure for the group in general, I believe [the applicant] would have scored less than 50 (fail)."

(McCarty later wrote to WIRED, "The fact that the candidate was from Atlanta played no bias nor did their socioeconomic status or race. The comment about Atlanta was only meant because gangs are more prevalent there than here in Columbus so the opportunity to be exposed could be greater.")

*“This is a 23-year-old kid who grew up in Atlanta that could have very well had some affiliation with gangs. Give me your thoughts.”*

Converus is proud of the fact that its system is designed, according to Mickelsen, “to accommodate varying historical levels of test failure by its applicant pool.” That is to say, subjects can be judged on how people of similar backgrounds have fared on credibility tests in the past. Law enforcement officers may get an easier ride, while those from the wrong part of town may face an uphill battle. Converus sees no problem with this kind of institutional bias. Mickelsen told WIRED, “Sensitivity can be adjusted for specific groups. This gives all examinees a fairer chance of being classified correctly. Most organizations can make good estimates of base rates by considering the number of previously failed background checks, interview data, confessions, evidence, etc.”

John Allen worries that it is a dangerous practice. “[You] would need to have a very good database on which to estimate rates of guilt,” he says. “Otherwise, leaving this up to the individual examiner will create a situation of high variability across examiners and the very real possibility of bias.”

Civil liberties groups are also wary of EyeDetect. “The criticism of technologies like lie detectors is that they allow bias to sneak in,” says Jay Stanley of the ACLU’s Speech, Privacy, and Technology Project. “But in this case it sounds like bias isn’t sneaking in—it’s being welcomed with open arms and invited to stay for dinner.”

While the polygraph may be shockingly [unreliable century-old](#) technology, at least critics can reinterpret and discuss test results out in the open. EyeDetect is a closed system using a proprietary algorithm, whose results can effectively be altered at the operator’s discretion. Its low price and automated operation also allow it to scale up in a way that time-consuming and labor-intensive polygraph tests never could.

Converus told WIRED that a Middle Eastern country has purchased EyeDetect and is planning to use it to check whether people entering the country are associated with terrorist activity. In an email to the Salt Lake City Police Department last year, obtained through WIRED’s public records requests, a Converus [executive](#) wrote that the company had “been identified as the solution for ‘extreme vetting’ by the new [Trump] administration.” (Though there were discussions with the Trump administration about using EyeDetect for vetting, Converus says the administration never committed to using EyeDetect.)

And while polygraphs remain banned from most US courts, EyeDetect appears poised to enter the legal system. In May, a district court in New Mexico became the first court to admit an EyeDetect test, in the trial of a former high school coach [accused](#) of raping a 14-year-old girl. The defendant passed the test, and the jury failed to agree on a verdict. Hearings on the admissibility of EyeDetect are due in at least four other states, the company tells me.

Unlike the polygraph, which is typically a one-off purchase, Converus earns money from every single test that each of its \$3,500 [EyeDetect stations](#) runs. According to emails, in the fall of 2017 Converus was charging law enforcement agencies between \$60 and \$80 per test. If EyeDetect could replace even a fraction of the estimated 2.5 million polygraph tests conducted annually in the US, Converus would have a reliable revenue stream for years to come. Whether it will prove as reliable for those who take the test remains a more troubling question.