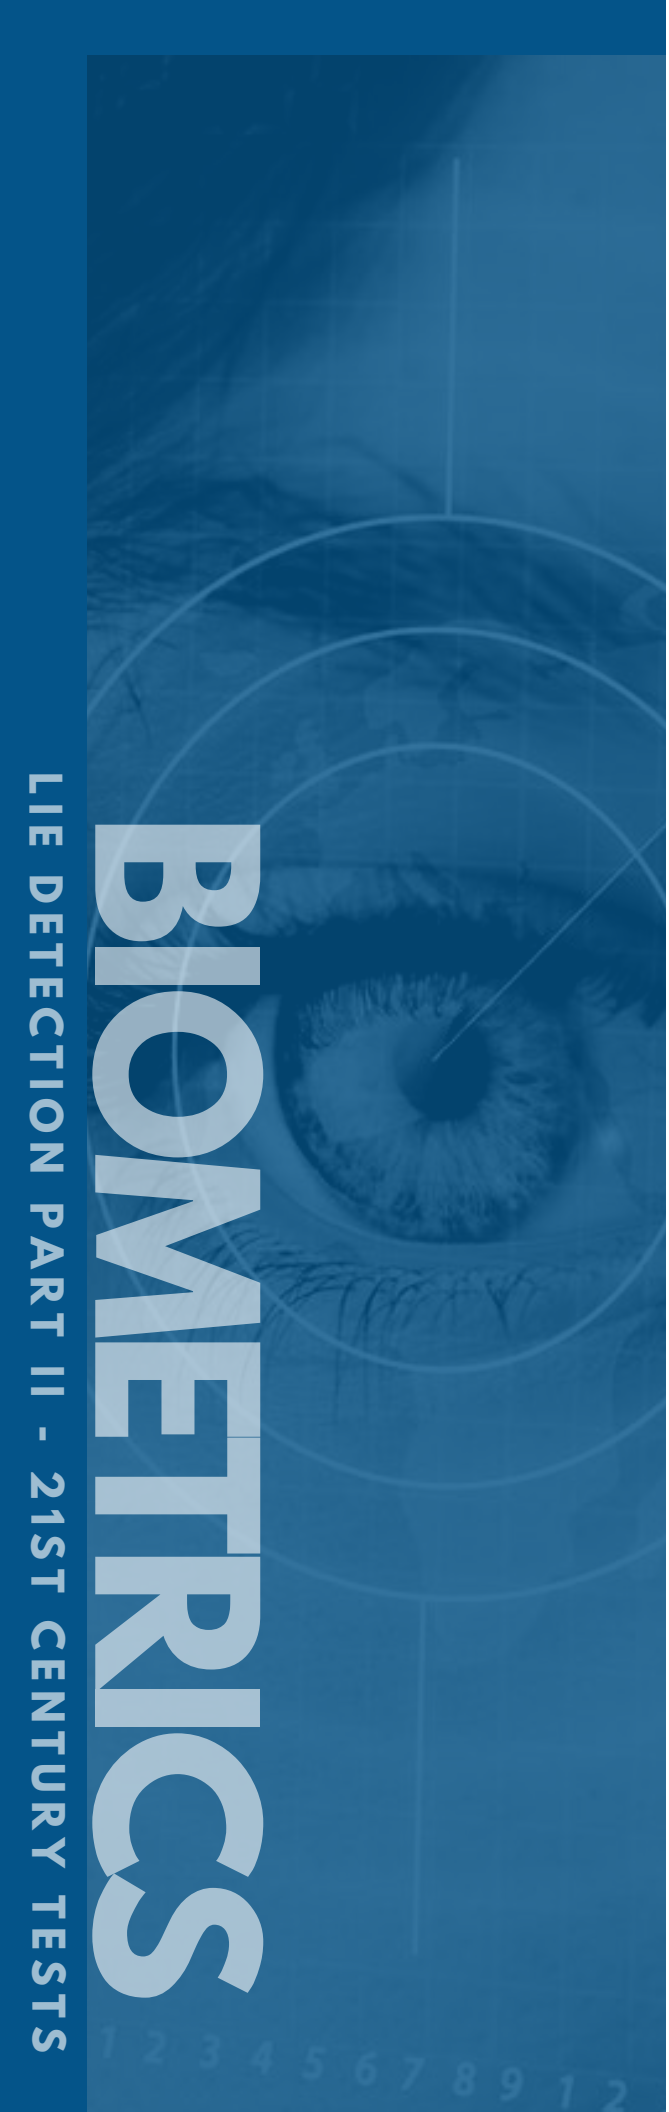

POLICING PROJECT FIVE-MINUTE PRIMERS:

21ST CENTURY LIE DETECTION

Written by Mira Joseph, Policing Project Extern and Allen Slater, Policing Project Student Fellow

The Policing Project's biometrics blog series explores basic information on one of the more complex—and rapidly changing—areas of policing: the use of biometric technologies. In each of these blogs we will explore four questions about a particular type of biometric technology:

- How are police using this technology?
- How does the technology work?
- How accurate is the technology?
- Should the public be concerned?

Previously we looked at the complicated history and usage of the polygraph, a machine that despite its generally chilly relationship with most courts is still utilized by law enforcement agencies in interrogating suspects or eliciting confessions.

While lie detectors have been around in some form for nearly all of human history, the 21st century has brought with it new variations that seem ripped from science fiction—from voice stress analysis software, to virtual policemen, to iris-scanning devices—many of which are driven by artificial intelligence (A.I.).

In this second part of our series on lie detection biometrics, we'll take a closer look at some of the latest iterations of duplicity detectors, with an emphasis on EyeDetect, AVATAR, and Silent Talker.

HOW ARE POLICE USING THESE NEW LIE DETECTION TECHNOLOGIES?

EyeDetect is a lie detection technology developed by the Utah-based company Converus that analyzes eye movements to assess whether a subject is being dishonest. At present, it has had only limited use by policing agencies. While U.S. law prohibits the use of lie-

detector tests for pre-employment screenings in the private sector, EyeDetect has been used in the public sector to vet policing recruits in multiple states, including Utah, Texas, Ohio, Louisiana, and Florida. Converus has pitched the test for use with parolees and probationers, and has lobbied the U.S. government to use the test in screening refugees, though it is unclear if the test is yet being used for these purposes. In at least one case, EyeDetect was used as part of a voluntary evaluation taken by a defendant in a child sexual abuse case.

Other new forms of lie detection technology, including AVATAR and Silent Talker, are aimed at automating and heightening border security. AVATAR, originally developed by researchers with the University of Arizona in collaboration with U.S. Customs and Border Protection and licensed by those researchers as Discern Science International, Inc., is essentially a highly-monitored interview with a digital border agent (the acronym stands for Automated Virtual Agent for Truth Assessments in Real-Time). Funding to develop the tech was provided by several public sector entities, including the U.S. Department of Defense, the National Science Foundation, the European Border and Coast Guard Agency (FRONTEX), and the U.S. Department of Homeland Security.

With AVATAR, travelers (and potentially asylum seekers) would be directed to a kiosk where the virtual border agent appears on a screen, asking interview questions while the system's camera and sensors monitor for signs of dishonesty. AVATAR has been deployed in field tests, including at Dennis DeConcini Port in Nogales, Arizona, and Henri Coandă International Airport in Bucharest.

The European Union has also funded a pilot program that used Silent Talker's technology for a similar virtual border agent screening system. The six-month border security pilot program, called iBorderCtrl, was coordinated by the Hungarian National Police and set up at four border crossing points in Hungary, Latvia, and Greece that connect to nations outside the EU. Unlike AVATAR's kiosk-based system, Silent Talker is accessed via home computer prior to the traveler's arrival at the border checkpoint. The virtual border agent commences an interview with the traveler while using the interviewee's webcam to record micro expressions, including facial movements, gaze, and posture. After the interview, the virtual border agent provides the interviewee a QR code to be shown to human border agents containing the Silent Talker's assessment of the interviewee's truthfulness.

During the pilot iBorderCtrl program, which ended in August 2019, the Silent Talker test was only deployed on a voluntary basis, but there are some signs these tests may become more widespread and mandatory in the future. Following completion of the iBorderCtrl program, the EU voted to create a centralized, searchable biometrics database called the Common Identity Repository. The database is expected to hold as many as 300 million records with biometric and biographic data, which includes names, fingerprints, photos, and home addresses, which would be incorporated into the iBorderCtrl system to strengthen its algorithm. (Such databases are already in use in the U.S., and though they are not linked to lie detection systems like iBorderCtrl, these databases have been linked to facial recognition algorithms in at least one pilot program.)

Beyond its border security applications, Silent Talker has also announced plans to lease its technology to the private sector, including to banks, law firms, and insurance companies to use in interviews and fraud screenings.

HOW DO THESE TECHNOLOGIES WORK?

As we explained in our previous discussion of polygraphs, there remains no universal physiological

indicator of deception. Humans have sought reliable methods for detecting lies for thousands of years, but from Lombrosso's Glove to Larson's polygraph to the A.I.-powered systems we're examining here, the core principle of lie detectors remains unchanged: these devices, no matter how technologically advanced, cannot detect objective truth. Instead, they can only measure physical responses that *may* be predictors of deception.

A.I.-powered systems detect patterns of movement typically associated with deception, but which may be too subtle or complex for a human eye to track (according to proponents). EyeDetect bases its test on changes in the interviewee's eyes, such as pupil dilation and reaction time. In comparison to a polygraph, which can take between one to three hours, EyeDetect takes fifteen to thirty minutes. The test is largely automated, with questions delivered via computer. Interviewees sit in front of an infrared camera that takes pictures of their eyes at 60 frames per second and records their response time to questions and error rates.

EyeDetect's test is based on the belief that liars will show more signs of cognitive load than truth tellers—and that EyeDetect's camera will be able to pick up these subtle signs as it captures the interviewee's eye dilation, movements, and response time. Unlike polygraphs, which require a human examiner to interpret the results, measurements collected during EyeDetect's test are analyzed by its proprietary algorithm (though as we'll discuss in a moment, the algorithm can be tweaked by human test administrators).

Like EyeDetect, AVATAR is also based on discerning deception based on micro-expressions and other alleged involuntary cues. As the system's virtual agent asks the traveler a series of questions about their journey, the system also scans their passport and fingerprints, and tracks their eye and body movements. AVATAR can pick up on as many as 50 different potential deception cues, including changes in pitch, gestures, posture, and eye movement. Its algorithm then interprets this data within 45 seconds and sends a verdict to a human patrol agent suggesting whether the traveler is safe to proceed or should be pulled for further questioning.

Silent Talker was one of the first A.I.-powered lie detection systems created, but in contrast to AVATAR, it relies entirely on micro-gestures in the face and head, utilizing only a camera as simple as a webcam. Representatives from Silent Talker have publicly confirmed that the company does not know how its algorithm arrives at its conclusions, though they do cite this as a reason for keeping a "human in the loop" when using the system.

HOW ACCURATE ARE THESE TECHNOLOGIES?

Converus, the company that developed EyeDetect, claims the test has an 86% accuracy rate. However, the only peer-reviewed studies of EyeDetect's accuracy were conducted by its in-house team, and the Policing Project was unable to find any independent academic studies of its accuracy by researchers without financial ties to the company.

Additionally, as WIRED reported, even among studies conducted by EyeDetect's own co-creator, the test's accuracy varied significantly among test groups. These variations may be linked to the test administrator's ability to tweak the algorithm—essentially, telling it to do a "softer" or "harder" assessment of the subject's performance depending on the administrator's preference. Researchers have also noted that EyeDetect uses control questions that can be easily identified by examinees. In other studies, these types of questions have caused innocent test subjects to create false positives due to anxiety.

Discern Science has similar accuracy rate claims for AVATAR of 80-85%, and Silent Talker claims a 80% accuracy rate. However, there are numerous

factors that affect the viability of these claims, including the size of the group being tested, and the racial and ethnic make-up of the sample group. For example, the researchers at Manchester Metropolitan University who developed iBorderCtrl used a test group of only 32 people in an academic paper that determined virtual border agents were "suitable" for detecting deception. The study group had twice as many men as women, only 10 participants of Asian or Arabic descent, and no Black or Hispanic participants. Facial recognition algorithms have been shown to perform poorly and struggle to recognize subjects of color when they are trained and tested on mainly white test participants.

Furthermore, many of these accuracy claims have not been independently replicated by scientists with no financial interest in the results. A significant complication with verifying the accuracy of these technologies is that their algorithms are proprietary, and companies are often not willing to share the inner workings of the systems during an independent audit.

SHOULD THE PUBLIC BE CONCERNED ABOUT THESE TECHNOLOGIES?

As mentioned in our post on polygraphs, traditional lie detectors are usually not considered reliable enough be used as evidence in most U.S. courts. However, this might not hold true for newer A.I.-powered models. In May 2018, a federal district court judge in New Mexico allowed the results of an EyeDetect test to be admitted as evidence on behalf of the defendant in a criminal trial. This presents concerns, as EyeDetect both relies on a similar underlying theory as the polygraph, and utilizes a proprietary algorithm that lacks independent peer-reviewed studies of its accuracy.

Due process concerns surrounding proprietary algorithms have already been raised in regard to algorithmic risk assessments used for sentencing, and similar concerns could be applied to A.I.-powered lie detectors. How can a defendant challenge the determination of a proprietary algorithm without adequate access to its underlying logic? Can a jury critically assess the validity of a test without knowing how the algorithm functions, and could that lack of understanding prejudice their view of the defendant?

The proprietary nature of the algorithms that govern these technologies presents additional bias concerns as well. AI algorithms develop complex models that are trained over huge data sets, making their conclusions nearly impossible to explain, and researchers have flagged many instances of inaccuracies and biases relating to characteristics such as race, age, and gender. Furthermore, there is no strong evidence that the technology works across different cultures, or on groups of people with atypical behavior whose non-verbal behavior diverges from normal expectations.

While Silent Talker's co-founder advocates keeping a "human in the loop" to verify the algorithm's findings, this is not a perfect solution. As we saw with polygraphs, human examiners may bring their own prejudices into the way they interpret lie detector results, or be swayed by more implicit biases. If these tests are also used in broader applications, it could present implications for inequality in matters such as public sector hiring, the investigation of insurance claims, or the monitoring of parolees.

Like polygraphs in the past, these new lie detection technologies may also pressure innocent subjects into making false confessions. There is not a lot of research surrounding this issue yet, but it raises ethical questions about how technology that claims such high accuracy rates and sophisticated procedures may affect the ability of law enforcement officers to misrepresent the test results during interrogations, or the ability of an attorney to effectively argue against them before a judge or jury.