ANNALS OF TECHNOLOGY

ADVERSARIAL MAN

Dressing for the surveillance age.

BY JOHN SEABROOK

Is there anything fashion can do to counter the erosion of public anonymity?

Tom Goldstein, an associate professor of computer science at the University of Maryland, took an "invisibility cloak" from a pile on a chair in his office and pulled it on over his head. To my eye, it looked like a baggy sweatshirt made of glossy polyester, printed with garish colors in formless shapes that, far from turning Goldstein invisible, made him impossible to miss.

It was mid-January. Early that morning, in my search for a suitable outfit to thwart the all-seeing eyes of surveillance machines, I had taken the train from New York City to College Park. As I rode the subway from Brooklyn to Penn Station, and then boarded Amtrak for my trip south, I counted the CCTV cameras; at least twenty-six caught me going and returning. When you come from a small town, as I do, where everyone knows your face, public anonymity—the ability to disappear into a crowd—is one of the great pleasures of city living. As cities become surveillance centers, packed with cameras that always see, public anonymity could vanish. Is there anything fashion can do?

I could have worn a surgical mask on my trip, ostensibly for health reasons; reports of an unexplained pneumonia outbreak in China were making the news, and I'd spotted a woman on the C train in an N95 respirator mask, which had a black, satiny finish. Later, when I spoke to Arun Ross, a computer-vision researcher at Michigan State University, he told me that a surgical mask alone might not block enough of my face's pixels in a digital shot to prevent a face-recognition system from making a match; some algorithms can reconstruct the occluded parts of people's faces. As the coronavirus spread through China, SenseTime, a Chinese A.I. company, claimed to have developed an algorithm that not only can match a surgically masked face with the wearer's un-occluded face but can also use thermal imaging to detect an elevated temperature and discern whether that person is wearing a mask. For my purposes, a full-face covering, like the Guy Fawkes mask made popular by the "V for Vendetta" graphic novels and films, would have done the trick, but I doubt whether Amtrak would have let me on the train. During Occupy Wall Street, New York enforced old anti-mask laws to prevent protesters from wearing them.

Goldstein's invisibility cloak clashed with the leopard-print cell-signal-blocking Faraday pouch, made by Silent Pocket, in which I carried my phone so that my location couldn't be tracked. As a luxury item, the cloak was far from the magnificent Jammer Coat, a prototype of anti-surveillance outerwear that I had slipped on a few weeks earlier, at Coop Himmelb(l)au, an architecture studio in Vienna. The Jammer Coat, a one-of-a-kind, ankle-length garment with a soft finish and flowing sleeves, like an Arabic thawb, is lined with cellular-blocking metallic fabric and covered with patterns that vaguely resemble body parts, which could potentially render personal technology invisible to electronic-object detectors. Swaddled in the cushy coat, I could at least pretend to be the absolute master of my personal information, even if its designers, Wolf and Sophie Prix, wouldn't let me leave the studio in it.

However, the invisibility cloak, while not as runway-ready as some surveillance-wear, did have one great advantage over other fashion items that aim

ILLUSTRATION BY ANA GALVAÑ

to confuse the algorithms that control surveillance systems: the cloak's designer *was* an algorithm.

To put together a Jammer outfit for my style of dressing—something like stealth streetwear—I first needed to understand how machines see. In Maryland, Goldstein told me to step in front of a video camera that projected my live image onto a large flat screen mounted on the wall of his office in the Iribe Center, the university's hub for computer science and engineering. The screen showed me in my winter weeds of dark denim, navy sweater, and black sneakers. My image was being run through an object detector called YOLO (You Only Look Once), a vision system widely employed in robots and in CCTV. I looked at the camera, and that image passed along my optic nerve and into my brain.

On the train trip down to Maryland, I watched trees pass by my window, I glanced at other passengers, and I read my book, all without being aware of the incredibly intricate processing taking place in my brain. Photoreceptors in our retinas capture images, turning light into electrical signals that travel along the optic nerve. The primary visual cortex, in the occipital lobe, at the rear of the head, then sends out these signals—which are conveying things like edges, colors, and motion. As these pass through a series of hierarchical cerebral layers, the brain reassembles them into objects, which are in turn stitched together into complex scenes. Finally, the visual memory system in the prefrontal cortex recognizes them as trees, people, or my book. All of this in about two hundred milliseconds.

Building machines that can process and recognize images as accurately as a human has been, along with teaching machines to read, speak, and write our language, a holy grail of artificial-intelligence research since the early sixties. These machines don't see holistically, either—they see in pixels, the minute grains of light that make up a photographic image. At the dawn of A.I., engineers tried to "handcraft" computer programs to extract the useful information in the pixels that would signal to the machine what kind of object it was looking at. This was often achieved by extracting information about the orientation of edges in an image, because edges appear the same under different lighting conditions. Programmers tried to summarize the content of an image by calculating a small list of numbers, called "features," which describe the orientation of edges, as well as textures, colors, and shapes.

But the pioneers soon encountered a problem. The human brain has a remarkable ability, as it processes an object's components, to save the useful content, while throwing away "nuisance variables," like lighting, shadows, and viewpoint. A.I. researchers couldn't describe exactly what makes a cat recognizable as a cat, let alone code this into a mathematical formula that was unaffected by the infinitely variable conditions and scenes in which a cat might appear. It was impossible to code the cognitive leap that our brains make when we generalize. Somehow, we know it's a cat, even when we catch only a partial glimpse of it, or see one in a cartoon.

Researchers around the world, including those at the University of Maryland, spent decades training machines to see cats among other things, but, until 2010, computer vision, or C.V., still had an error rate of around thirty per cent, roughly six times higher than a typical person's. After 9/11, there was much talk of "smart" CCTV cameras that could recognize faces, but the technology worked only when the images were passport-quality; it failed on faces "in the wild"—that is, out in the real world. Human-level object recognition was thought to be an untouchable problem, somewhere over the scientific horizon.

A revolution was coming, however. Within five years, machines could perform object recognition with not just human but superhuman performance, thanks to deep learning, the now ubiquitous approach to A.I., in which algorithms that process input data learn through multiple trial-and-error cycles. In deep-learning-based computer vision, feature extraction and mapping are done by a neural network, a constellation of artificial neurons. By training a neural net with a large database of images of objects or faces, the algorithm will learn to correctly recognize objects or faces it then encounters. Only in recent years have sufficient digitized data sets and vast cloud-based computing resources been developed to allow this data- and power-thirsty approach to work. Billions of trial-and-error cycles might be required for an algorithm to figure out not only what a cat looks like but what kind of cat it is.

"Computer-vision problems that scientists said wouldn't be overcome in our lifetime were solved in a couple of years," Goldstein had told me when we first met, in New York. He added, "The reason the scientific community is so shocked by these results—they ripple through everything—is that we have this tool that achieves humanlike performance that nobody ever thought we would have. And suddenly not only do we have it but it does things that are way crazier than we could have imagined. It's sort of mind-blowing."

Rama Chellappa, a professor at the University of Maryland who is one of the top researchers in the field, told me, "Let me give you an analogy. Let's assume there are ten major religions in the world. What if, after 2012, everything became one religion? How would that be?" With computer-vision methods, he said, "that's where we are."

Computers can now look for abnormalities in a CT scan as effectively as the best radiologists. Underwater C.V. can autonomously monitor fishery populations, a task that humans do less reliably and more slowly. Heineken uses C.V. to inspect eighty thousand bottles an hour produced by its facility in France—an extremely boring quality-control task previously performed by people. And then there is surveillance tech, like the YOLO detector I was standing in front of now.

No one would mistake me for Brad Pitt, I thought, scrutinizing my image, but no one would mistake me for a cat, either. To YOLO, however, I was merely a collection of pixels. Goldstein patiently led me through YOLO's visual process. The system maps the live digital image of me, measuring the brightness of each pixel. Then the pixels pass through hundreds of layers, known as convolutions, made of artificial neurons, a process that groups neighboring pixels together into edges, then edges into shapes, and so on until eventually you get a person. Nuisance variables—the bane of handcrafted

C.V.—are removed along the way, as pixels are distilled into features that encode my presence. All this happens in about the same amount of time that it takes the brain to recognize an object. Finally, a red outline, called a "bounding box," appeared around me on the live screen, with the label "Person." Boom.

"You're detected," Goldstein said.

Advances in computer vision have occurred so rapidly that local and national privacy policies—what aspects of your face and body should be protected by law from surveillance machines—are lagging far behind A.I.'s technological capabilities, leaving the public vulnerable to a modern panopticon, a total-surveillance society that could be built before we know enough to stop it. Chris Meserole, a foreign-policy fellow at the Brookings Institution who studies China's use of face recognition and other surveillance technologies—widely deployed as part of Xi Jinping's "stability maintenance" drive—told me that policymakers in the States haven't, so far, created governing structures to safeguard citizens. And, he added, "in the U.S., the government hasn't thought to use it yet the way that China has."

Some activists think we've already run out of time. Before travelling to Maryland, I had acquired several ready-to-wear anti-surveillance items from a woman named Kate Bertash, whom I went to see in Los Angeles. She met me in the lobby of my hotel in Venice Beach wearing a black dress printed with license plates. She handed me a black T-shirt, men's large, also covered with license plates. It was a warm winter day in Venice Beach, where Bertash, thirty-three, lives and works. I stepped into the rest room and put on my T-shirt. The dummy plates spelled out words from the Fourth Amendment.

"Welcome to the resistance," Bertash said, when I emerged.

We set off on a stroll down Abbot Kinney Boulevard, the main drag in Venice Beach. The plates on our clothing were designed to trigger automatic license-plate readers. In the U.S., the networks of A.L.P.R.s and databases that exist across the country make up a different kind of surveillance system. First developed in the U.K., in the late seventies, A.L.P.R.s began appearing in U.S. cities in the early two-thousands. The readers use optical character recognition, which captures plate numbers and stores the information, along with the location, date, and time of the recording. Newer systems can also pinpoint where a car is most likely to be found, based on travel patterns. A.L.P.R.s are mounted on street lights, highway overpasses, freeway exits, toll booths, digital speed-limit signs, and the tops of police cars. They are also found in parking garages, schools, and malls. Companies such as PlateSmart Technologies market software to the general public that can turn almost any surveillance camera into an A.L.P.R. The open-source version of a similar software is free.

A.L.P.R.s automatically record all license-plate numbers that come within their view, at a rate of thousands per minute. In newer systems, "hot lists" of "plates of interest" belonging to criminal suspects are widely shared by law-enforcement agencies, including U.S. Immigration and Customs Enforcement. Officers are alerted to a location when a plate shows up on a reader connected to the network they're using. There are few privacy restrictions on this data, and it is not secure. Private companies collect and sell it. Police departments obtain data and share it with one another. According to *The Atlantic*, Vigilant Solutions, the industry leader, has a database of at least two billion unique license-plate locations. A recent audit of the Los Angeles Police Department and three other California law-enforcement agencies found that, at the time they were logged, 99.9 per cent of the three hundred and twenty million plate images in the department's database had not been involved in criminal investigations. State Senator Scott Wiener, who requested the audit, told the Los Angeles *Times*, "I am horrified. We believed that there were problems with the ALPR program, but I did not anticipate the scale of the problem—the

"I think we just have time for one more quick question."

fact that we have so many law enforcement agencies that are not complying with state law, including LAPD."

Bertash works for the Digital Defense Fund, a nonprofit that provides security and tech support for the abortion-access movement. In a café where we stopped for lunch, she explained that protesters often "stand outside of abortion clinics photographing all day long." She was concerned that an anti-abortion activist with access to A.L.P.R. data could easily figure out where abortion providers and patients live.

Dave Maass, a senior investigator with the Electronic Frontier Foundation, a nonprofit digital-privacy advocate, confirmed Bertash's fears about the insecurity of the data. Bertash wondered what, as an activist, she could do. Maass suggested postering public spaces with paper images of license plates that would feed false data into the system. Bertash had a better idea.

As a part-time gig, Bertash designs and sells novelty fabrics for kids—sheets and towels printed with manatees and kittens, pillows that look like cuts of meat. She started producing mockups of clothing with phony plates, testing them with an open-source A.L.P.R. app that might (or might not) work like those used by law enforcement. Eventually, she got her designs to read into the system as real plates and produced a line of garments and accessories with dummy plates printed on them, which she sells on Adversarialfashion.com.

Bertash's anti-A.L.P.R. clothes are "poison" attacks, which aim to pollute databases with garbage, so that the system as a whole is less reliable. Poison attacks are predicated on collective action. A few people festooned in license plates won't make much difference; a lot of people wearing them might. For that to happen, designers need to make anti-surveillance clothes that you'd want to put on. T-shirts strewn with fake license plates might not be everyone's must-have look for spring.

We spent several hours strolling around. No one asked about our clothes: in Venice Beach, it takes a lot more than a license-plate outfit to stand out as unusual. When a police cruiser with an A.L.P.R. on top passed us on the sidewalk, I tried to feel adversarial. Honestly, I felt kind of sheepish.

Possibly to compensate for the guilt I felt about not being down with the resistance, I ended up buying a license-plate backpack, for $49.95. When I got home, my eleven-year-old daughter, to whom I often feel invisible—not in a good way—actually noticed me. "What's that?" she said, studying the plates. "That's really cool!" Detected at last.

When I told my children, both "Harry Potter" fans, that I was going to check out an invisibility cloak, they were excited. I'd learned of Goldstein's cloak in a scientific paper that he and his students produced about their work. But when I saw Goldstein in his sweatshirt, which featured a foreground of blurry organic shapes in orange, like a display of horribly irradiated vegetables, with dark, vaguely human shapes above, I couldn't imagine Harry or Hermione wizarding with one. The only recognizable shape (to me) was what appeared to be a traffic light just below the neckline. Considered more generously, the pattern loosely evoked Georges Seurat's "A Sunday Afternoon on the Island of La Grande Jatte," as it might appear at the bottom of a swimming pool painted by David Hockney.

Then Goldstein stepped in front of the camera, and the YOLO detector did a double take. It couldn't see him at all. The computer saw the chair behind him ("Chair," the bounding box was labelled) but not the six-foot-tall, thirty-six-year-old man standing right in front of it—Goldstein Unbound. I, in my supposedly anonymous city duds, was instantly detected and labelled. It was like a conceit from William Gibson's 2010 science-fiction novel, "Zero History," in which a character wears a T-shirt so ugly that CCTV cameras can't see it.

The pattern on the sweatshirt was an "adversarial image"—a kind of deep-learning optical illusion that stopped the algorithm from seeing the person wearing it. Unlike poison attacks, which seek to subvert surveillance systems with bad data, adversarial attacks are images that have been engineered to take advantage of flaws in the way computers see. They are like hacks, but for artificial intelligence. The security vulnerabilities of operating systems and computer networks are widely known, but deep-learning A.I. systems are still new and so complex that scientists don't yet fully understand the kinds of hacks they are vulnerable to.

The phenomenon of adversarial imagery was discovered more or less by accident in 2011, by Christian Szegedy, at Google Research. Szegedy trained a neural net to solve the problem of just how much he could change an image of a ship before the system reclassified the image as an airplane. He discovered that with only a minimal modification of pixels the system reclassified it with a high degree of confidence, even though to the human eye it was still obviously a ship and not an airplane. Students at M.I.T. printed a three-dimensional model of a turtle with a textured shell that fooled Google's object-detection algorithm into classifying the reptile as a rifle. In a 2018 paper, "Robust Physical-World Attacks on Deep Learning Visual Classification," researchers described an experiment in which they "perturbed" a stop sign with a few small decals that to a human look like graffiti but that made an object classifier see the octagonal red sign as a rectangular black-and-white sign that said "Speed Limit 45." It isn't hard to imagine the kind of chaos one of these perturbances could cause in a future world of autonomous cars.

Goldstein's research is ultimately aimed at understanding these vulnerabilities, and making A.I. systems more secure. He explained that he and his student Zuxuan Wu were able to create a pattern that confuses the network using the same trial-and-error methods employed in training the neural network itself. "If you just try random patterns, you will never find an adversarial example," he said. "But if you have access to the system you can find a pattern to exploit it." To make the sweatshirt, they started with a pattern that looked like random static. They loaded an image of people, covered a small part of the image with the pattern, and showed the result to a neural network. An algorithm was used to update the pattern to make the neural net less confident that it was seeing people. This process was repeated using hundreds of thousands of images, until the static slowly morphed and the neural net could no longer see people when the resulting pattern was present in an image.

"I couldn't tell you why this pattern

works," Goldstein said. Researchers can't understand exactly how the machine sees. "These are very complicated systems," he said. "They have weaknesses that occur in the interactions between feature maps and artificial neurons. There are strange and exploitable pathways in these neural networks that probably shouldn't be there."

Adversarial examples demonstrate that deep-learning-based C.V. systems are only as good as their training data, and, because the data sets don't contain all possible images, we can't really trust them. In spite of the gains in accuracy and performance since the switch to deep learning, we still don't understand or control how C.V. systems make decisions. "You train a neural network on inputs that represent the world a certain way," Goldstein said. "And maybe something comes along that's different—a lighting condition the system didn't expect, or clothing it didn't expect. It's important that these systems are robust and don't fail catastrophically when they stumble on something they aren't trained on."

The early work on adversarial attacks was done in the digital realm, using two-dimensional computer-generated images in a simulation. Making a three-dimensional adversarial object that could work in the real world is a lot harder, because shadows and partial views defeat the attack by introducing nuisance variables into the input image. A Belgian team of researchers printed adversarial images on two-dimensional boards, which made them invisible to YOLO when they held the boards in front of them. Scientists at Northeastern University and at the M.I.T.-I.B.M. Watson A.I. Lab created an adversarial design that they printed on a T-shirt. Goldstein and his students came up with a whole line of clothes—hoodies, sweatshirts, T-shirts.

I put on a sweatshirt, which had shapes and colors similar to Goldstein's, but in a slightly different configuration. On stepping in front of the camera, I was undetected, too. I felt strangely weightless.

I asked Goldstein to speculate about why these particular blurry shapes were adversarial. He pointed to the shape that looked sort of like a traffic light on his chest. Perhaps, he said, because there were no human faces above traffic lights in the training data, the algorithm could not see a face that was above one on the sweatshirt.

As long as I stood still, I was an adversarial man. But the luxury of invisibility was fleeting: as soon as I moved, I was detected again. Goldstein's gear works as a proof of concept, but it has a long way to go in the wild.

Like object detection, face recognition improved dramatically in the twenty-tens with the switch to deep learning. Early handcrafted features for faces involved mathematical formulas that expressed how far apart the pupils of the eyes are on a face, for example, or the distance from the bottom of your nose to the top of your lip. But "there are things about your face I don't even know how to write down mathematically," Goldstein told me, "and a neural net will discover and extract this information."

Deep-learning-based face recognition starts with a detector much like YOLO, and can run on top of any CCTV camera's feed. First, an image passes through layers of a neural network that quickly map out the locations of facial features. "Anything with two eyes, a nose, and a mouth is almost always a face at this stage," Goldstein said. Then each face is isolated and passes through a more refined neural network that removes nuisance variables, distilling the face into a short list of unique coördinates—your facial fingerprint, or faceprint. Many systems also outline the eyes, the eyebrows, the nose, the lips, and the mouth, using sixty-eight standard landmark points to identify emotions and gaze. Some sophisticated systems (like Apple's FaceID for the iPhone) use infrared scanners to make three-dimensional face maps. The results are expressed as numerical data—your unique identifier. Unlike the tips of your fingers or your driver's license, your face can be scanned remotely, without your knowledge or consent, and mined for age, gender, emotion, and, if your labelled picture happens to be in the system's database, your identity.

As with all deep-learning systems, the more data you train the algorithm on, the more accurate the model will become. Early face-detection systems developed for military, border-control, and law-enforcement purposes were trained on labelled databases of faces in

"Gross! Can't I sneeze into somebody else's elbow?"

the form of passport and driver's-license photos, and mug shots—the only large collections of faces that existed before the Internet. But these databases were of little value in trying to match faces captured in challenging light conditions and obscured views. Photos posted to photo-sharing Web sites and social media, on the other hand, are gold.

If the government were to demand pictures of citizens in a variety of poses, against different backdrops, indoors and outdoors, how many Americans would readily comply? But we are already building databases of ourselves, one selfie at a time. Online images of us, our children, and our friends, often helpfully labelled with first names, which we've posted to photo-sharing sites like Flickr, have ended up in data sets used to train face-recognition systems. In at least two cases, face-recognition companies have strong connections to photo-management apps. EverRoll, a photo-management app, became Ever AI (now Paravision), and Orbeus, a face-recognition company that was acquired by Amazon, once offered a consumer photo app. And even when our images are supposedly protected on social-media sites like Facebook, Instagram, and YouTube, how secure are they?

In January, the *Times* reported that Clearview, a Manhattan-based startup backed by the investor Peter Thiel and co-founded by Richard Schwartz, a former mayoral aide to Rudolph Giuliani, had assembled a database of more than three billion images scraped from social-media sites, and that Clearview's technology was being used by more than six hundred law-enforcement agencies to match faces of suspects or persons of interest with faces in Clearview's database. Google, Twitter, Venmo, and other companies have sent cease-and-desist letters to Clearview. Its co-founder and C.E.O., Hoan Ton-That, an Australian entrepreneur in his early thirties, claims that the company has a First Amendment right to these images. In any case, as Clare Garvie, a senior associate at Georgetown Law's Center on Privacy & Technology, told me, the Clearview database "gives the lie to" the notion that social-media "privacy policies are a safeguard against data collection." (Clearview's entire client list was stolen by hackers last month.)

Some of the data sets that academics used for training early algorithms skewed white and male—actors, politicians, and the academics themselves. Even diverse data sets presented problems: poor contrast in photos with darker skin tones, for example, would make it more difficult to match faces. There are biases built into algorithms, as Joy Buolamwini and Timnit Gebru, of M.I.T., showed in a 2018 report, "Gender Shades": facial-recognition systems in commercial use performed much better on light-skinned males than on dark-skinned females. Women of color are up to thirty-four per cent more likely to be misidentified by the systems than white men, according to their research. Newer collections of faces for training, like I.B.M.'s Diversity in Faces data set, aim to overcome these biases. However, I.B.M.'s effort also proved to be problematic—the company faced backlash from people who found their images in the data set. In face recognition, there is a trade-off between bias and privacy.

Apart from biases in the training databases, it's hard to know how well face-recognition systems actually perform in the real world, in spite of recent gains. Anil Jain, a professor of computer science at Michigan State University who has worked on face recognition for more than thirty years, told me, "Most of the testing on the private venders' products is done in a laboratory environment under controlled settings. In real practice, you're walking around in the streets of New York. It's a cold winter day, you have a scarf around your face, a cap, maybe your coat is pulled up so your chin is partially hidden, the illumination may not be the most favorable, and the camera isn't capturing a frontal view."

The technology's questionable performance doesn't seem to be impeding its ongoing implementation. Though face recognition is only one application of computer vision, it poses a unique threat to civil liberties. The E.U. has tried to make privacy policies to contain it, as have a few states, including Illinois. In China, by contrast, the state has embraced the technology. A 2015 proposal laid out the country's plans for a vast centrally controlled surveillance system using face recognition and other technologies. In addition to making use of China's installed base of more than two hundred million CCTV cameras (the U.S. lags behind, with fewer than a hundred million cameras), the plan, according to Chris Meserole, of the Brookings Institution, involves linking video feeds from smart TVs and mobile devices in rural areas, where CCTV coverage is much lighter. The Chinese A.I. company SenseTime, which last year was valued at more than seven billion dollars, has said that the facial-recognition system it is building will be able to process feeds from up to a hundred thousand CCTV cameras in real time.

Are China's surveillance-state ambitions technically feasible? Meserole is skeptical. However, he added, "whether they are able to do it in a totally unified way or not is in some ways irrelevant. A huge part of how Chinese authoritarianism works is the uncertainty about whether you are being watched. The technology is incredibly precise, but the way the laws are applied is incredibly arbitrary. You are uncertain if you are being watched, and you are uncertain about what's permissible, and that puts the onus on you as the individual to be really conservative about what you're doing."

I considered adding face camouflage to my adversarial look, and met with Adam Harvey, an American artist based in Berlin, who made a name for himself in the early twenty-tens by creating a series of asymmetric getups that could defeat the Viola-Jones algorithm, which until 2015 was the most widely used object- and face-detection platform. Face-detection algorithms are trained to expect symmetry in faces. When people put on makeup, they are unwittingly helping the systems by accenting some of the landmarks that scanners use to read your faceprint. To fly under the radar, you must deface yourself. Harvey's work showed faces with makeup applied asymmetrically, in a way unlikely to be represented in the systems' training data, therefore making them harder for machines to detect as faces. Fashion-forward types can

download symmetry-distorting looks from Harvey's Web site, though he said that they probably won't work with newer algorithms.

Harvey, thirty-eight, is slim, pale, and quietly intense. We met in a café in Williamsburg, Brooklyn, where he glanced several times at a CCTV camera mounted high up in a corner of the room. He said, "We don't really understand what we're doing when we go outside. We can know the weather, and we dress for it, but if I had known on the way over here that I was going to pass four private surveillance cameras outside houses, or that there was"—he broke off and glanced at the CCTV camera—"would I have dressed for it?"

He went on, "We exist in this world where we are observed by machines. How can you mediate that to appear one way to the machines and another way to people? How can you ride the fine line between appearing avant-garde and appearing invisible?"

Harvey explained that he had moved on from face camouflage because, theoretically, any makeup design that can be used to foil a detection system could be incorporated into the system's training data. "I realized that, whatever I post on my Web site, people are going to use it to download and test their algorithm on." This is the paradox of the adversarial man: any attempt to evade the system may only make it stronger, because the machine just keeps learning. And, with deep learning, it keeps learning faster.

Is there any science on what kinds of disguise thwart face recognition? That question has long fascinated Rama Chellappa, Goldstein's colleague at the University of Maryland. "I'm interested in this because the spymasters do it in real life," Chellappa told me. "But there was no really scientific evaluation of what works."

Disguises present the same problem to recognition algorithms as aging does; aging is a kind of natural disguise, he said. Some well-known faces become unrecognizable as they age (Anthony Michael Hall looks nothing like the young actor in those Brat Pack movies), whereas others (like Paul Rudd or Halle Berry, say) don't appear to age at all. "Aging is hard to train an algorithm on, because it's person-specific," Chellappa said.

In 2018, Chellappa and other A.I. researchers, based in India, created the Disguised Face in the Wild competition. "With the advances of deep-learning algorithms, we wanted to evaluate whether the deep-learning methods were robust to disguises," Chellappa told me. He and his colleagues put together a database of thousands of faces, taken both from movies, like Dana Carvey's 2002 film, "Master of Disguise," and from ordinary people's photos of Halloween and other dress-up events that had been posted on social media.

Teams from around the world were invited to test their face-recognition algorithms by matching disguised faces with their undisguised counterparts. The competition was supported by IARPA, a research organization within the Office of the Director of National Intelligence, which gave a twenty-five-thousand-dollar cash prize to the winning team. In return, IARPA's face-recognition capabilities had the chance to benefit from the training data the competition generated, making them that much more robust against real-life masters of disguise.

I asked the professor to summarize the research. "What can I wear if I really don't want to be seen?"

"You can wear a beard, you can shave your head, and that will affect face-recognition algorithms in different ways," Chellappa said, adding, "I really can't tell you if you do x, y, and z it will mess up the face recognition. All I can say is, if you do a combination of hat, wig, dark glasses, you can assume the accuracy will go down." For now, at least.

The top-performing algorithms—the hardest to fool—were designed by the Russian and Taiwanese entrants. The 2019 challenge was sponsored by Facebook and Apple.

So far in the U.S., the deployment of face recognition by public agencies and law enforcement is less advanced than it is in China. Last May, San Francisco banned city agencies from using facial-recognition technologies. In an Op-Ed published in the *Times* in June, titled "How Facial Recognition Makes You Safer," James O'Neill, a former commissioner of the N.Y.P.D., wrote that in New York City "no one can be arrested on the basis of the computer match alone," and that human investigators would need to confirm any matches that machines suggest.

Where the U.S. leads the world is in the commercial use of face recognition by private companies. Many major tech companies have deep-learning face-recognition systems and training databases. Facebook's product, DeepFace, can identify faces in photographs and tag them. Google has FaceNet, as well as an object detector, Cloud Vision. Amazon markets Rekognition, a C.V. platform that has been deployed by police departments and was pitched to ICE for use in border enforcement. Apple makes infrared scans of the faces of users who opt into its FaceID password system; the encrypted data isn't supposed to leave the user's phone.

In addition to Big Face, there is a rapidly growing field of startups, part of a market that is expected to be worth nine billion dollars a year by 2022, according to some estimates. The products include face recognition for stores, which can identify repeat shoplifters and troublemakers as soon as they step onto the premises. In a casino, as Richard Smith, the sales director of SAFR, a division of RealNetworks, explained to me, a system can spot unwanted patrons and problem gamblers who are on the casino's watch list, as well as high rollers whom management wants to court. "Before face recognition, the guards had to remember those people," Smith said. Some schools have installed similar security systems; college campuses have also begun contemplating their implementation. Taylor Swift has reportedly used face recognition to detect the presence of stalkers at her shows.

Face recognition also offers "smart retail" applications, allowing companies to harvest demographic information from customers' faces, such as age and gender, and also to track and measure "dwell time"—how long a customer spends in any particular section of the store. "What if you could see what your ad sees?" SAFR asks on the company's Web site. A video shows a couple having a conversation while data appears

on the screen, assigning attention and "sentiment scores" to their faces. Other companies offer face surveillance that alerts stores to shoppers' previous shopping habits, or their V.I.P. status, when they walk in. Face Six, a biometrics company based in Israel and Nevada, markets Churchix, software that is often used to track congregants' attendance at church. As Clare Garvie, of Georgetown Law's privacy center, put it, "Think of a possible application for this technology and chances are good it's being sold."

For marketers, face recognition has the potential to be the ultimate form of targeted advertising. The consumer's face could serve as a kind of license plate that connects the digital world—where your search histories live—to the location data that Google Maps collects from your phone, to the emotions on your face. We're already traced online, and are served ads based on recent searches. Face recognition could follow us around in the real world, alerting the owners and managers of public spaces, such as subway stations and parks, or private spaces, like bars, shops, and stadiums, to our presence. I can opt out of accepting cookies, and disable location settings on my phone, but a face-recognition system doesn't give me any way to opt out, short of defacing myself.

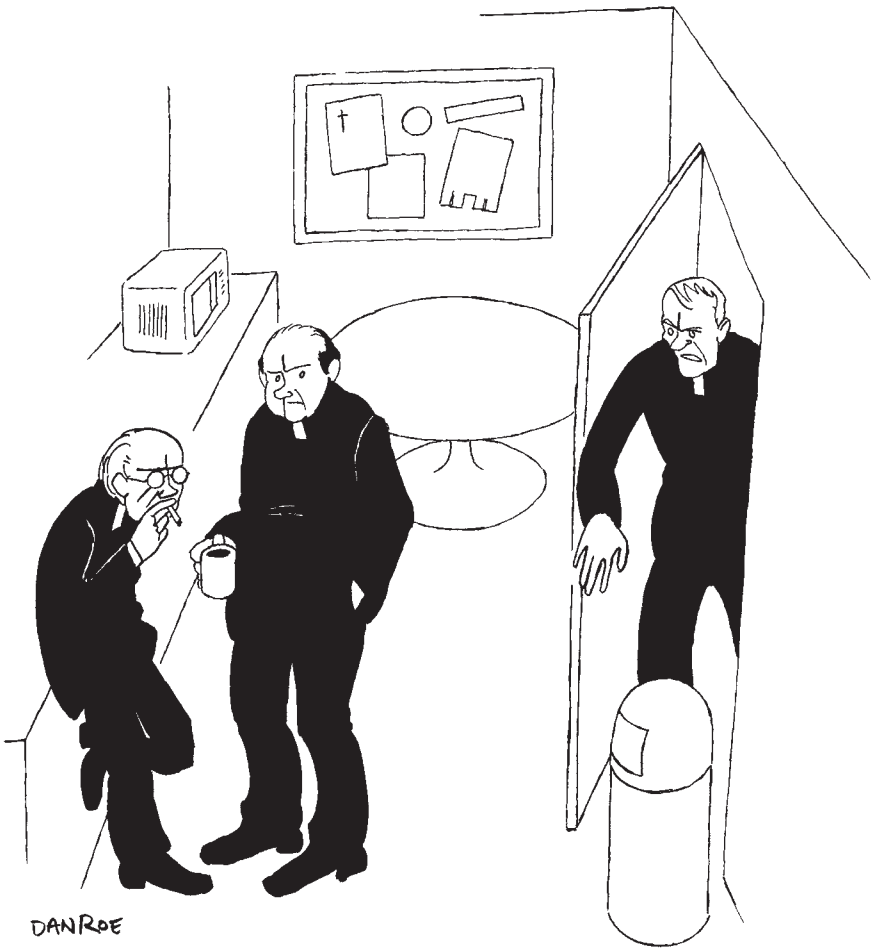

"Back to work, boys. Those mysteries of the Trinity aren't going to grapple with themselves."

• • •

One day last month, I put on a hoodie that I had selected from Goldstein's line of YOLO-busting fashion, hoisted my A.L.P.R.-poisoning backpack (I had my T-shirt on, too, for good measure), grabbed my Faraday pouch, and set off from my home in Brooklyn for downtown Manhattan and *The New Yorker's* offices, in One World Trade Center.

On my face I wore a pair of Reflectacles—sunglasses made by Scott Urban, a custom eyewear-maker in Chicago, that block attempts via infrared light to scan your face. They come in three models—IRPair, Phantom, and Ghost, my style. The lenses contain "infrared absorbents," Urban told me. "This means that on your average security camera using infrared for illumination these lenses turn dark black, whereas regular sunglasses become completely clear." The Ghost model also has frames that reflect both infrared and visible light, which can make your face less readable in photos taken with a phone, especially with a flash. And, unlike Adam Harvey's asymmetrical makeup, Urban's spectacles allow you to remain relatively inconspicuous. "Sure, you can paint yourself up and look like some cyberpunk-type character," he observed. "But you're not going to wear that to work."

Finding a seat on the C, I was alert to any stares my adversarial hoodie might attract from people across the aisle. As I searched their faces, I wondered how long it will be before face-recognition technology, with the power of deep learning behind it, arrives on everyone's phone. You will be able to snap a picture of someone across the subway aisle and run the face through a reverse-image search, such as that offered by Socialcatfish.com, an "online dating investigation" site based in California, which promises to ascertain if one's Tinder hookups are who they claim to be. You could potentially get a name or a social network before you reach your stop. That dystopia could be an app away. As Urban put it, "People are concerned about governments and corporations—well, soon it's going to be people doing all this tracking." Big Brother is us.

My plan was to loiter in the lobby of One World Trade Center, where one can assume that video surveillance is in place—an adversarial man at work. Maybe, if I stood perfectly still, the algorithm wouldn't see me. I recalled a comment made by Anil Jain, of Michigan State, when I asked him what I needed to wear to beat detection algorithms. "You can put tinfoil all over yourself and that will do the job nicely," he said. "It all depends how much of a spectacle of yourself you want to make."

Before I could go undetected, security spotted me. "You need help?" a guard asked. Good question. ♦