Can Computers Be Racist?

Kelly Gates

September 11, 2014

Kelly Gates is Associate Professor in Communication and Science Studies at the University of California San Diego.

Can computers be racist? I want to explore this question as it relates to facial recognition technology by focusing on the "accuracy" of this technology. Simply raising the question—"Can computers be racist?"—implies that my answer is "Yes." And putting scare quotes around the word "accuracy," as I do, implies that I am going to argue that facial recognition technology is not accurate. But rather than providing a simple yes-or-no answer to this question, I want to provide a more complete answer that complicates how we understand the accuracy of computers and the work they do. In doing so, I want to challenge assumptions about the imminent perfectibility of facial recognition technology—the idea that in the very near future the technology will work perfectly and will be able identify human beings better than humans do.

We can put the question of accuracy in the context of peace and conflict studies. Donald MacKenzie's 1993 book, *Inventing Accuracy*, addresses this question in relation to the history of nuclear missile guidance systems.[1] He argues that it was very important for developers of intercontinental nuclear missile systems to convince a range of actors—including policy makers, funders, and the public—that it was possible to accurately project a nuclear warhead across great distance so that the bomb would hit a specific target like a city on another continent. In other words, the construction of the accuracy of the intercontinental guidance systems was as important as the construction of the guidance systems and the nuclear missiles themselves. In fact, the viability of the nuclear program depended very centrally on this invention of the accuracy of nuclear missile guidance.

I argue that a similar case can be made about facial recognition technology. As with nuclear missile guidance, it is critically important to the success of facial recognition technology that its proponents construct its accuracy in addition to developing the technology itself. Certain constituencies, including funders and policy makers as well as potential users in the government and private sectors, must be convinced that this technology is accurate—or accurate enough to be useful—and that research and development is achieving greater and greater levels of accuracy in matching faces and identifying people.

I will say more about what it means to invent the accuracy of facial recognition technology, but it is the anniversary of 9/11, so I think it is appropriate to first reflect on those events. I want to do this by

referring to an image that circulated in the press in the immediate aftermath of the events of September 11, 2001. (The image is reproduced in the opening chapter of my book, *Our Biometric Future*,[2] and is also available from the *Washington Post* at http://www.washingtonpost.com/wp-srv/photo/postphotos/asection/2005-02-13/11.htm.) There was a great deal of interest in and attention to facial recognition technology in the aftermath of 9/11, and I want to use this image to revisit that moment of intensified interest. (Not everyone is familiar with this image; in fact, I find a diminishing number of people who recognize it—an indicator that the events and their aftermath are receding into the past.) The image, which was taken from a surveillance video feed, appears to depict two of the 9/11 hijackers passing through security in the Portland, Maine, airport on the morning of September 11, 2001. It is a banal type of surveillance image that we are now quite familiar with seeing—images taken from a surveillance video that circulate as press photos. It is also a chilling image because of what it depicts.

Importantly, this particular surveillance camera image was often accompanied by a claim about facial recognition technology. This claim held that if the technology had been in place, it could have identified the men in the image as wanted terrorist suspects and alerted airport security, conceivably preventing the attacks from happening. Take, for instance, the following quote from the December 2001 issue of the *MIT Technology Review*:

> Of all the dramatic images to emerge in the hours and days following the September 11 attacks, one of the most haunting was a frame from a surveillance-camera video capturing the face of suspected hijacker Mohamed Atta as he passed through an airport metal detector in Portland, ME. Even more chilling to many security experts is the fact that, had the right technology been in place, an image like that might have helped avert the attacks. According to experts, face recognition technology that's already commercially available could have instantly checked the image against photos of suspected terrorists on file with the FBI and other authorities. If a match had been made, the system could have sounded the alarm before the suspect boarded his flight.[3]

HOW DOES FACIAL RECOGNITION TECHNOLOGY WORK?

This was a big claim, and one of the goals of my book is to unpack this claim—to ask the question of what would have had to happen to make this claim a reality. The answer is much more complicated than it seems on the surface. To understand why it is a complicated question, it is important to understand something about how facial recognition systems work. First of all, there is an important distinction between verification and identification—the verification of known identities, on the one hand, versus the identification of unknown individuals. Verifying a known identity using facial recognition technology—or any kind of identification technology, like fingerprinting—means making a one-to-one comparison between the person whose identity is being verified and other known information about that specific person. Verification involves a one-to-one match of two images of a particular individual (including typically one live image and one stored image) to determine if those images depict the same person.

It took a long time for computer scientists to come up with a way of doing that reliably, and it can still fail at times, especially when they are taken under different lighting and other conditions. But it is not as difficult as identifying a person when there is no information to go on other than an image or a fingerprint. In other words, verifying a known or professed identity is not nearly as challenging for computers as identifying unknown individuals—a process usually called identification to distinguish it from verification. Identification involves a one-to-many comparison of images in search of a match (or more precisely, the likelihood of a match), which is a much more challenging problem. For example, when an unknown individual passes through the field of view of a surveillance camera, and her face is grabbed from a video feed and processed by a facial recognition system, the problem is one of identification—identifying an unknown individual. And trying to match that unknown face against a database to figure out who that person is, without having a known identity to work with, is a much more challenging problem. When the 9/11 attacks happened, computer algorithms were not able to identify faces with much accuracy, and identification remains a challenging problem today (especially when the aim is to identify not one person, but many thousands or even hundreds of thousands of unknown people).

Another related problem is the effort to program computers to recognize facial expressions. Automated facial recognition and automated facial expression analysis are in some ways distinct technological efforts and in some ways connected. Where facial recognition involves programing computers to identify faces and distinguish them from one another, facial expression analysis involves getting computers to parse movements that are happening on the surface of the face and make meaningful connections between those movements and emotions. With facial recognition technology, the face is used as a kind of fingerprint or index of identity. Facial expression technology, on the other hand, involves using the motion of the face to try to see inside the person, to make sense of what the person is thinking or feeling. In fact, one of its potential applications, the most attractive one for state security and law enforcement purposes, is for deception detection.

SOME PROBLEMS WITH ACCURACY

It became obvious after 9/11 that facial recognition technology did not work that well. It certainly would not have been ready to be plugged into surveillance systems at airports so that it could identify terrorist suspects from video feeds. One place where this became obvious was in Tampa, Florida, where, in the summer just before 9/11, the Tampa Police began a project to integrate facial recognition technology with a video surveillance system in an entertainment and shopping district called Ybor City. For those trying to implement this project—to create a so-called “Smart CCTV” (closed-circuit television) system that would automatically identify wanted suspects from video—it was clear right away that this was going to be a much more difficult undertaking than they initially thought. The smart CCTV

system ostensibly relied on a watchlist database of people with outstanding arrest warrants in the area. The idea was to compare facial images of unknown individuals taken from video feeds generated by a CCTV system installed in Ybor City with images stored in the watchlist database.

I went to Tampa to interview some of the people involved with implementing and operating the Ybor City smart CCTV system, and they very graciously took me into the control room to give me a demonstration. In the control room that day, I witnessed that not only was the system not turned on, but as soon as they booted it up, the system did not work well. It seemed to start registering false alarms. The American Civil Liberties Union (ACLU) later confirmed that the Tampa Police were not really using the system in earnest, largely because it did not work well—a conclusion the ACLU came to after analyzing documents obtained through Florida's open-records law.[4] In fact, the effort to implement a facial recognition system in Ybor City ended up being a complete failure. The Tampa Police abandoned the project entirely after two years.

This case underscores that what you can achieve in a laboratory can mean very little once you attempt the same thing in a real-world setting. This is partly a problem of scale or scalability—a small-scale system that works fairly well within limited constraints often does not work nearly as well (if at all) when it is scaled to a larger size, especially when increase in scale involves orders of magnitude more data. What you can achieve on a small scale with a small dataset and relatively closed systems has almost nothing to do with what you can achieve on a large scale with large datasets and much larger systems or infrastructures. For example, if you were to develop a facial recognition system to recognize the one hundred or so people in this room, enrolling everyone here into a database, this limited facial recognition system may be able to recognize everybody here consistently, with a good accuracy rate. However, in order to scale that to 10,000 people, 100,000 people, or a million people, the system would need to be completely reinvented. It is not simply a matter of making the system bigger, adding more images or more cameras or more wires. It becomes a completely different problem, or set of problems, with orders of magnitude more complexity.

Scalability is one of the issues that complicates the question of accuracy. The real world is a lot messier than laboratories—that is why we have experimental research—to try to control for the complexity of the world. In the case of Ybor City, this was a very small-scale operation of only thirty-six cameras. It was, in a sense, a real-world laboratory where the developers and the police were conducting an experiment. And it still failed, because there was a lot that could not be controlled for, like lighting conditions, movement of people within and beyond the area, distance of people from the camera, and even how police officers sitting in front of the monitors would use and react to the system. Of course, even though it was a failure, there were no doubt many lessons learned from the experiment that could be taken forward in future efforts to create functioning facial recognition systems.

The point I want to emphasize here is that fully formed, functional, ready-to-be-deployed, automated facial recognition technology was not a reality when the 9/11 attacks occurred. And while that became painfully obvious fairly quickly, we continued to hear persistent claims about the technology's imminent perfection. An implicit acknowledgement of the technology's limitations was always combined with a projection of the technology's future perfectibility: it may not work perfectly now, but it very soon will, with more work and more processing power, with better algorithms and more experience, and with more real-world testing. And this claim about the inevitability of the technology, and about its imminent perfectibility, is itself an important part of the effort to push the development of facial recognition technology forward by gaining increased investment in research and development.

What we have witnessed since 9/11 is the repetition of the claim that in the near future, computers will be able to recognize faces and facial expressions, if not perfectly, then near perfectly—and certainly better than humans are able to do so. That facial recognition technology will soon work better than humans at identifying people is another claim that one frequently hears. But what exactly does it mean to get computers to recognize facial expressions "better than humans do?" I will return to this question in a moment.

IDENTIFYING BIN LADEN'S BODY

There have been some more recent developments around facial recognition technology since 9/11, and I am going to mention a few of them to bring us to the present moment, with the caveat that I am just scratching the surface. One moment when facial recognition technology appeared again in the news was after the killing of Osama bin Laden by U.S. forces in 2011. After bin Laden was killed, a very important question arose—not explicitly, but beneath the surface—about how the commandos could be certain that they had the right man, that they had actually killed bin Laden and not someone who looked like him. His dead body was gone, disappeared, apparently buried at sea, yet confirming the identity of the person who had been killed was a priority. It was a priority for the national security state in particular, so that U.S. authorities could make the claim to have killed bin Laden without any skepticism about whether in fact they had the right person.

That U.S. government officials felt the need to prove that U.S. forces had in fact killed bin Laden was apparent in reporting on the verification of the identity of bin Laden's body.[5] In fact, the media reported that the Central Intelligence Agency (CIA) used multiple methods to verify that the dead body was bin Laden. One of these methods was old-fashioned human facial recognition: human beings using their own perceptual abilities. It was reported that the CIA brought in a woman who had apparently been one of bin Laden's wives at one time, and she identified him. Then the actual commandos who found him and killed him also said it was him—they recognized the man they killed as Osama bin Laden. These

were actually two different kinds of human-based identification, because the U.S. forces who identified the dead body would have only ever seen images of him before that. They were not likely ever in bin Laden's actual presence before then, so they were comparing images they had seen of his face to the dead body. The woman, on the other hand, was someone who had obviously been in bin Laden's immediate physical presence before, since she was married to him. Thus she was performing a different kind of identity verification based on her own, firsthand experiences of what he looked like.

A second technique used to identify bin Laden, as reported by the media, was DNA typing. U.S. government officials had DNA taken from the dead man's body compared to DNA taken from some of bin Laden's known relatives, which ostensibly established a match. And the third technique was facial recognition technology, which was used to compare the face of the dead body to some known images of bin Laden. Why would they use so many techniques to identify bin Laden's body? And why use facial recognition technology?

One answer is that government officials had to be very sure, and very convincing, about actually having killed bin Laden—to put the nail in bin Laden's coffin, so to speak. But why use facial recognition technology in particular? For me, it seemed that bin Laden's killing presented an opportunity for facial recognition technology to finally fulfill the promise made after 9/11 that it would have identified the terrorists had it been in place. Now, it was actually being used to identify the uber-terrorist—Osama bin Laden himself. By identifying bin Laden's dead body, facial recognition technology fulfilled the promise made of it in the immediate aftermath of 9/11—to identify the enemy of the state, in this case, the ultimate 9/11 terrorist.

IS THE TECHNOLOGY GETTING MORE ACCURATE?

This year has seen a lot of press attention given to facial recognition technology. One headline came over the summer, from *The Guardian*: "Suspected Child Abuse Fugitive Caught by Facial Recognition after Fourteen Years."[6] Another one, also from *The Guardian*, had nothing to do with identifying suspects: "CIA Facial Software Uncovers the Artist Francis Bacon—In Drag."[7] The story included an accompanying image titled "*Unknown Woman*" taken by a photographer named John Deacon in the 1930s. Apparently it is Francis Bacon, although no one could explain why the image appears to show cleavage. In any case, articles like these provide important publicity for the technology. They allegedly show that the technology works and that developers are making improvements that are enabling the technology to have value in the real world for addressing real problems—like identifying fugitives and famous men posing as women.

Another moment when facial recognition made the news was earlier this year, when Google publicly announced that it would not be incorporating the facial recognition app NameTag into the

Google Glass platform. Apparently, it was controversial enough that people wearing Google Glass could take video of those around them without their knowledge. Google wisely recognized it would just add to the controversy if Glass wearers could also identify people around them with a facial recognition app; people were already creeped out enough by the glasses themselves. There was a general sentiment that people were not keen on that idea because it went too far, especially in the context of Edward Snowden's National Security Agency (NSA) revelations.

Snowden's revelations provided further occasion to report on developments in facial recognition technology. In May 2014, a *New York Times* headline proclaimed: "N.S.A. Collecting Millions of Faces from Web Images."[8] According to reporters James Risen and Laura Poitras, the NSA is now intercepting millions of images per day from emails and other communications, including about 55,000 facial recognition-quality images per day, according to documents obtained from Edward Snowden. Along with this coverage, we again find repeated references to the improvements in the accuracy of the technology. The article quotes Alessandro Acquisti, a researcher from Carnegie Mellon University: "There are still technical limitations on it," he says, "but the computational power keeps growing, the databases keep growing, and the algorithms keep improving."[9]

Interestingly, in this case it is the scale of the data that fuels the technology's improvement. With the rise of what we now call Big Data, it seems that scale has become the solution to creating better, more accurate systems. In other words, now the way to overcome the difficulties associated with scaling up facial recognition systems is precisely by working with and designing algorithms that process millions of faces. Facebook in particular is seen as having a big scale advantage. In March 2014, the company publicized the results of research on its facial recognition system, called DeepFace. The Facebook researchers claimed that DeepFace achieved "human-level performance in facial recognition."[10] They trained DeepFace to work on "the largest facial data set to date, an identity-labeled data set of four million facial images belonging to more than 4,000 identities."[11] Here Facebook is claiming to use the enormous scale of their facial image database (containing hundreds of millions of images uploaded by users) as the basis from which to improve the technology by orders of magnitude.

An article in *MIT Technology Review* also reported on the Facebook DeepFace findings: "Facebook Creates Software that Matches Faces Almost as Well as You Do."[12] This is a common approach to constructing the accuracy of facial recognition technology—comparing computer-matching algorithms to human beings' capacity for identifying faces. According to this article, when asked to identify two photos showing the same person, a human will get it right 97.53% of the time. (Any human? Every human?) The new software from Facebook scores a close second—97.25% accuracy—on the same challenge, regardless of variations of lighting, or whether the person in the picture is directly facing the camera. *Forbes* also covered this story: "Facebook's DeepFace Software Can Match Faces with 97.25%

Accuracy."[13] And again from the tech website *ExtremeTech*: "Facebook's Facial Recognition Software is Now as Accurate as a Human Brain."[14] The claim from a specific, limited study gets repeated again and again, becoming a statement of fact about the software's comparability to the human brain.

What exactly does it mean to say that facial recognition technology is as accurate as a human brain? On the one hand, you could say it is simple. How often can a program accurately identify a human being, and how does that compare to how accurately human beings identify other human beings? Can a program accurately identify people 100% of the time, or 97.25%, or some other percentage of the time? How does that compare to the percentages achieved when human beings try to identify people's faces? But we need to ask where these percentages come from, how are they generated, and what they actually mean. Can we really go from the percentages reached in specific studies to the claim that facial recognition software is now as accurate as the human brain?

What I want to put a fine point on is that accuracy rates are always generated from experiments. When someone makes a claim about accuracy, that person is typically referring to the accuracy rate that was generated in an experiment or a set of experiments. Experiments involve specific kinds of software applications being used with specific parameters around them and limitations as to what they measure. With the DeepFace program that Facebook claims is 97.25% accurate, the claim is based on an experiment or set of experiments that Facebook used to generate this accuracy rate, and in fact, it is not facial identification. Facebook researchers are not talking about facial *identification* in this experiment; they are talking about *verification*. DeepFace performs what researchers call facial verification by making a one-to-one comparison and recognizing that two images show the same face. In other words, this is not the more challenging form of facial recognition, where one-to-many comparison are made, and names are put to unknown faces. So you can see the way limited research findings start to circulate as accuracy claims that go beyond what the experiments are actually showing.

You could say this a simple problem of validity. In other words, does an experiment measure what it claims to measure? And certainly, the researchers themselves are careful to avoid these errors, as one finds if one reads the technical literature. The validity problem is not in the experiment itself, but in how the experimental findings get taken up and circulated as bigger claims about the increasing accuracy of the technology more broadly. Experimental findings about accuracy can never be accurately extended to facial recognition technology in general, to the entirety of facial recognition systems and all of their applications.

The point is that there are so many different types of systems, different projects, different approaches being taken, and different experiments involved in the effort to develop automated facial recognition that one could never make the claim that in every context, the technology has a certain

accuracy rate. What you achieve in any particular experiment or set of experiments does not extend to what you will achieve in the real world.

Of course, there are also claims that the technology is working with greater accuracy in the real world. A few months after Facebook reported the accuracy rates that it was achieving with DeepFace, the Chinese University of Hong Kong claimed to beat Facebook with its program, called DeepID. According to an article in the magazine *Fast Company*, "a few months after Facebook's breakthrough, the Chinese University of Hong Kong claims to have smashed Facebook's record by building a recognition system that achieves a massive 99.15% accuracy rate based on a truly innovative deep learning model."[15] What makes the Hong Kong labs research different is that it used images taken "in the wild" (their term) rather than in a laboratory setting under controlled conditions.

In this case, along with the usual claim of greater accuracy, we find an implicit acknowledgement of the limitations of laboratory research. Here the Chinese researchers are claiming even greater accuracy rates and suggesting that their research findings apply beyond the computer science laboratory. However, despite this claim, their findings are still based on experiments conducted in a university research setting. Again, you could say this is simply a measure of validity: does the research measure what it purports to measure? But in fact, it is not simply a matter of validity, strictly speaking. Using findings based on a limited study to make a statement about the reality of the world is a classic error, sometimes called the fallacy of misplaced concreteness—assuming that a measure or representation of something can stand in for or become the thing that it represents.

A similar thing could be said about automated facial expression analysis. Much like with facial recognition technology, there is a near compulsion to make the claim that experiments with automated facial expression analysis are achieving greater and greater accuracy at identifying facial expressions, and in fact doing so more accurately than humans. But what does it mean for technology to "see" expressions better than humans do? Expression *is* human, and to suggest that computers can see human expressions better than humans do is a strange claim. How can you say that any measure of emotion accurately represents the emotion itself? This is a validity problem, and again, it tends to be dealt with very carefully in the psychology literature on facial expression. The people studying facial recognition in psychology are constantly trying to grapple with the question of what facial expressions mean, and what the connection is between the expressions on a face and what is going on inside the person.

Researchers have tried to resolve this problem by taking other kinds of measures—other physiological measures in particular—and correlating them. So, for example, having an elevated heart rate when making a certain facial expression may be correlated to an emotion. But such correlations always require interpretive judgments about what emotions are indicated by physiological measures. It is also the case that different disciplines understand and define emotion differently. The physiological view

of emotion—that emotion is essentially an embodied, physiological process—seems highly compatible with the computational analysis of emotion. If emotions are something that can be measured physiologically, it would seem that they are also something that can be modeled in a computational system.

However, another school of thought conceptualizes emotion as something relational, intersubjective, and emergent. According to this view, my emotional state right now is not mine alone. Instead, it is happening in relation to the "vibe" I perceive from other people in the room, while at the same time, they are having emotional responses to what I am saying and doing. There is a live, intersubjective process of emotional activity that is based on emergent relationships taking shape in this room, right now. To try to measure this complicated process using physiological indicators, and then to claim to be doing so more accurately than human beings do it, is committing the error of misplaced concreteness: assuming a measure of something is the thing it represents. Again, people who actually do this research tend to be much more careful about avoiding this error, but then their research gets taken up and circulated in deceptive ways, and even used for applications like deception detection (an application that demands precision and accuracy where no such thing is possible).

DO DATA HAVE POLITICS?

Finally, I want to point to another concept that is useful for thinking about problems with the so-called accuracy of facial recognition technology. The idea that facial recognition and emotional recognition can be modeled in computational systems—that computers can be made to do these things with great accuracy, and much better than humans do them—is an example of what José van Dijck calls dataism–the belief that Big Data offers a perfect, accurate representation of the world, and that data science is the ultimate source of truthful information about everything.[16] Dataism also includes the assumption that there is a direct and self-evident relationship between data and people—that the digital traces we leave behind on servers and hard drives reveal our essential traits and qualities, including our spiritual beliefs, our intellect, and even our predisposition to diseases. Dataism is the belief that we can discover everything we need to know in the data that flows over the Internet, and in the data contained in the world's databases. These vast volumes of data represent a perfect, mirror-reflection of reality. Van Dijck also suggests that dataism represents a resurgence of a flawed faith in the objectivity of quantification and a misguided view of metadata as the raw material of social life.

A particularly good example that shows the problems with this data-centric view of the world is that there appears to be a dearth of data on police killings in the United States. There is no easy way of tracking the number of police killings over time, or comparing them across precincts, or analyzing the demographics of people who get killed by the police. You cannot do this kind of statistical analysis

because no dataset exists. Creating the dataset is not an easy thing to do, and it is not a job for one person. It has to be an institutional endeavor; it has to be systematic, and there has to be a policy mandate. We can not just assume that the data we need exist somewhere. Sometimes the data do not exist; in many cases, the data *purposefully* do not exist.

Another term for the fallacy of dataism is computationalism, or the idea that the entire world operates computationally, that everything can be explained or modeled in computation. Efforts to model something computationally are almost always accompanied by the belief that computational models can perform that function with precision and with complete technical neutrality. The computationalist view holds that vision can be explained and modeled via computational models—and more specifically, the way we see each other's faces can be modeled this way. While I am not suggesting that it is impossible to develop computational techniques that simulate visual processes, I am arguing those techniques do not and cannot offer a perfect model of vision.

Let me just wrap up by returning to the question I started with: can computers be racist? I have not addressed this question directly. Some of you are probably thinking, "Of course computers can be racist!" Others of you are probably thinking, "How could computers possibly be racist?" And maybe some of you are asking, "What do *you* think? Just tell us already!" Let me explain the connection between this question and the question of accuracy.

I want to suggest that these accuracy and technical neutrality claims are essentially a post-racial way of dealing with race, by not addressing it directly or explicitly. In other words, implicit in efforts to establish the accuracy of facial recognition technology is the claim that it cannot be racist—it cannot have any social or cultural bias built into it. Despite the inescapable connection between the face and race—the face as a key component of racial identification—it is actually quite rare to hear anyone in the biometrics industry talk about race explicitly. I do have one example, from 2001, in the immediate post-9/11 moment when the technology was everywhere in the press and in policy discussions. The example comes from Joseph Atick, the one-time CEO of a company called the Visionix Corporation. Speaking about his company's product at the time, he said that FaceIt "performs matches on the face based on analytical measurements that are independent of race, ethnic origin, or religion. It is free of the human prejudices of profiling."[17]

It should be clear by now that I think we should be skeptical of these sorts of claims. I think it is important to hold out the possibility that something like facial recognition technology could be used in certain ways that might in fact mitigate racism or structural inequalities; we should not rule that out. However, it is also important to understand that even if a perfectly accurate facial recognition system were possible—which it is not—it still could be used in ways that reproduce structural inequalities. All

kinds of technologies reproduce structural inequalities in intentional and unintentional ways, whether or not they are designed to do so.

In short, asking whether computers can be racist is another way of asking the philosopher Langdon Winner's famous question, "Do artifacts have politics?"[18] This question is a foundational one for the field of science and technology studies, where a great deal of work aims to identify and unpack the politics of artifacts. A central question that animates this field is the question of whether technologies can be neutral, or whether they inevitably embody certain social values, cultural assumptions, and political priorities. For his part, Langdon Winner argues that yes, artifacts do have politics, and some artifacts have more politics—and more troubling politics—than others. I would argue that the effort to program computers to see the face falls in this latter category.

NOTES

1. Donald Mackenzie, *Inventing Accuracy: A Historical Sociology of Nuclear Missile Guidance* (Cambridge, MA: MIT Press, 1993).
2. Kelly Gates, *Our Biometric Future: Facial Recognition Technology and the Culture of Surveillance* (New York: NYU Press, 2011).
3. Alexandra Stikeman, "Recognizing the Enemy," *MIT Technology Review*, December 1, 2001, http://www.technologyreview.com/featuredstory/401300/recognizing-the-enemy/.
4. "Drawing a Blank: Tampa Police Records Reveal Poor Performance of Face-Recognition Technology," American Civil Liberties Union, January 3, 2002, https://www.aclu.org/technology-and-liberty/drawing-blank-tampa-police-records-reveal-poor-performance-face-recognition-t.
5. Madison Park and Sabriya Rice, "How Did U.S. Confirm the Body Was Bin Laden's?" CNN.com, May 3, 2011, http://www.cnn.com/2011/HEALTH/05/02/bin.laden.body.id/.
6. Samuel Gibbs, "Suspected child abuse fugitive caught by facial recognition after 14 years," *Guardian*, August 13, 2014, http://www.theguardian.com/technology/2014/aug/13/suspected-child-abuse-fugitive-caught-facial-recognition.
7. Gordon Comstock, "CIA Facial Software Uncovers the Artist Francis Bacon – In Drag," *Guardian*, June 16, 2104, http://www.theguardian.com/artanddesign/2014/jun/16/cia-software-unveils-francis-bacon-in-drag.
8. James Risen and Laura Poitras, "N.S.A. Collecting Millions of Faces from Web Images," *New York Times*, May 31, 2014, http://www.nytimes.com/2014/06/01/us/nsa-collecting-millions-of-faces-from-web-images.html.
9. Risen and Poitras, "N.S.A. Collecting Millions of Faces from Web Images."
10. Yaniv Taigman, Ming Yang, Marc Aurelio Ranzato, and Lior Wolf, "DeepFace: Closing the Gap to Human-Level Performance in Face Verification," June 24, 2014, https://research.facebook.com/publications/480567225376225/deepface-closing-the-gap-to-human-level-performance-in-face-verification/
11. Taigman et al., "DeepFace: Closing the Gap to Human-Level Performance in Face Verification."
12. Tom Simonite, "Facebook Creates Software that Matches Faces Almost as Well as You Do," *MIT Technology Review*, March 17, 2014, http://www.technologyreview.com/news/525586/facebook-creates-software-that-matches-faces-almost-as-well-as-you-do/.

13. Amit Chowdhry, "Facebook's DeepFace Software Can Match Faces with 97.25% Accuracy," *Forbes,* March 18, 2014, http://www.forbes.com/sites/amitchowdhry/2014/03/18/facebooks-deepface-software-can-match-faces-with-97-25-accuracy/.
14. Sebastian Anthony, "Facebook's Facial Recognition Software is Now as Accurate as a Human Brain, But What Now?" ExtremeTech, March 19, 2014, http://www.extremetech.com/extreme/178777-facebooks-facial-recognition-software-is-now-as-accurate-as-the-human-brain-but-what-now.
15. Luke Dormehl, "How Machines Learned to Recognize Our Faces So Well—And What's Next," *Fast Company*, June 25, 2014, http://www.fastcolabs.com/3032386/how-machines-learned-to-recognize-our-faces-so-well-and-whats-next.
16. José van Dijck, "Datafication, Dataism, and Dataveillance: Big Data between Scientific Paradigm and Ideology." *Surveillance and Society*, 12 (2014): 197-208.
17. *Biometric Identifiers and the Modern Face of Terror: New Technologies in the Global War on Terrorism: Hearing before the Subcommittee on Technology, Terrorism, and Government Information of the Committee on the Judiciary, United States Senate*, 107th Cong. (2001) (statement of Joseph J. Atick, Chairman and Chief Executive Officer, Visionics Corp., Jersey City, NJ). http://www.gpo.gov/fdsys/pkg/CHRG-107shrg81678/html/CHRG-107shrg81678.htm.
18. Langdon Winner, "Do Artifacts Have Politics?" *Daedalus* 109 (1980): 121-136.