MIT Technology Review

SIGN IN

SUBSCRIBE

Better Face-Recognition Software

Computers outperform humans at recognizing faces in recent tests.

By Mark Williams Pontin

May 30, 2007

For scientists and engineers involved with face-recognition technology, the recently released results of the Face Recognition Grand Challenge–more fully, the Face Recognition Vendor Test (FRVT) 2006 and the Iris Challenge Evaluation (ICE) 2006–have been a quiet triumph. Sponsored by the National Institute of Standards and Technology (NIST), the match up of face-recognition algorithms showed that machine recognition of human individuals has improved tenfold since 2002 and a hundredfold since 1995. Indeed, the best face-recognition algorithms now perform more accurately than most humans can manage. Overall, facial-recognition technology is advancing rapidly.

Face facts: The top 3-D image only shows the information associated with the shape of a man’s face. The lower image shows the texture as well as the shape.

Jonathon Phillips, program manager for the NIST tests and lead author of the agency’s report, says that the intended goal of the Face Recognition Grand Challenge was always an order-of-magnitude improvement in recognition performance over the results from 2002. Phillips believes that the necessary decrease in error rate to achieve that goal was due in large measure to the development of high-resolution still-images and 3-D face-recognition algorithms. “For the FRVT 2006 and the ICE 2006, sets of high-resolution face images, 3-D face scans, and iris images were collected of the same people," Phillips says. "The FRVT 2006 for the first time measured the performance of six 3-D algorithms on a set of 3-D face scans. The ICE 2006 measured the performance of ten algorithms on a set of iris images. 3-D face recognition has come into its own in the last few years because 3-D sensors for face recognition have become available only recently. What 3-D face recognition contributes is that it directly captures information about the shapes of faces."

Among other advantages, 3-D facial recognition identifies individuals by exploiting distinctive features of a human face's surface–for instance, the curves of the eye sockets, nose, and chin, which are where tissue and bone are most apparent and which don't change over time. Furthermore, Phillips says, "changes in illumination have adversely affected face-recognition performance from still images. But the shape of a face isn't affected by changes in illumination." Hence, 3-D face recognition might even be used in near-dark conditions.

According to Ralph Gross, a researcher at the Carnegie Mellon Robotics Institute, in Pittsburgh, 3-D facial recognition can also recognize subjects at different view angles up to 90 degrees–in other words, faces in profile. "Face recognition has been getting pretty good at full frontal faces and 20 degrees off, but as soon as you go towards profile, there've been problems." Gross says that the explanation for face-recognition software's difficulties with profiles may be no more complicated than the fact that no one was focusing on the problem. The main applications of face recognition have been in contexts like ID cards and face scanners, for which the aim has been recognition of the full frontal faces of cooperative subjects under controlled lighting.

High-resolution still images have been another factor in the improvement of face-recognition technology, in part because highly detailed skin-texture analysis has also become possible. With such analysis, any patch of skin–called a skin print–can be captured as an image, then broken up into smaller blocks that algorithms turn into mathematical, measurable spaces in which lines, pores, and the actual skin texture are recorded. "It can identify differences between identical twins, which isn't yet possible using facial-recognition software alone," Gross explains. "By combining facial recognition with surface-texture analysis, accurate identification can increase by 20 to 25 percent."

What about the FRVT report's claim that some face-recognition algorithms equal or exceed humans' recognition capabilities? Phillips explains: "Humans are very good at recognizing faces of familiar people. However, they aren't so good at recognizing unfamiliar people." Since many proposed face-recognition systems would complement or replace humans, the FRVT's comparative tests of the face-recognition capabilities of humans and software–the first such testing–were important for measuring the potential effectiveness of applications. Phillips says that at low false accept rates (a false accept rate is the measure of the likelihood that a biometric security system will incorrectly accept an access attempt by an unauthorized individual), six out of seven automatic face-recognition algorithms were comparable to or better than human recognition. These were algorithms from Neven Vision, Viisage, Cognitec, Identix, Samsung Advanced Institute for Technology, and Tsinghua University. Unfortunately, Phillips adds, "because the majority of FRVT 2006 participants haven't disclosed the details of their methods, it's not possible yet to assess what's distinctive about these algorithms."

How does the commercial payoff for face recognition look? Quite promising, because dozens of companies aim to cash in on face recognition's potential as a biometric for credentialing and verification purposes. For the FRVT, venerable corporations like Toshiba and Samsung competed alongside companies like Neven Vision–just acquired by Google–andViisage and Identix (which have just merged into L1 Identity Solutions), as well as alongside researchers from universities as diverse as Beijing, Cambridge, and Carnegie Mellon. What applications does a company like Google foresee for the technology developed by its recent acquisition, Neven Vision? According to a Google PR person, "We believe it offers promising integration possibilities with Google's services, such as Picasa and Picasa Web Albums, particularly in terms of helping users organize and search their own photos."

At Carnegie Mellon, Ralph Gross says that among other efforts, he and his colleagues have been "involved with local DMVs in order to scan images for driver's licenses. I've gotten reports from the state level to say that, using face-recognition technology, they caught quite a number of people who applied for licenses in either different states or in the same state under a different name because their previous license got suspended." It's a growing trend. States using such technology include Massachusetts, Illinois, West Virginia, Wisconsin, Colorado, North and Southern Carolina, Oklahoma, North Dakota, Arkansas, and Mississippi. Nevertheless, Gross stresses, applying face-recognition technology to ID photos is a long way from having the capability that would let law enforcement search a city's webcam networks for specific individuals. "With driver's license photos, you have a controlled background, an operator telling you exactly how to position your face; the images are collected under comparable conditions. It's much more restricted than the random-face-in-the-crowd problem, where you're sticking a camera on a building."

Still, Gross says, "you can already see the path building." Until recently, the video-surveillance industry still mostly relied on analog cameras, requiring cable to be set up for long distances to connect those cameras to monitoring equipment. Now, "the industry is switching to IP-based cameras, with which you can pretty easily tap into already existing Ethernet networks," Gross says. "So you have wireless cameras and cameras using POE [Power over Ethernet technology allows IP telephones, wireless LAN Access Points, and other appliances to receive power as well as data over existing LAN cabling] where you don't need a separate power plug. You can buy commercial solutions that are essentially a TiVo for these cameras, with motion sensors built in so they only record when there's motion happening. With digital storage, you can keep the data indefinitely and enhance it in ways that you can't with analog images. So all these things are coming together."

In principle, therefore, as face-recognition software continues its rapid advance, it will likely be possible to search for specific faces across a network of webcams. Accordingly, Gross's recent work at Carnegie Mellon, in conjunction with colleagues at the Data Privacy Lab there, has been the development of algorithms to *protect* individuals' privacy while under video surveillance. The usual methods that thwart human recognition of an individual's features on video–for example, those pixelated fields sometimes covering faces and body parts on reality-TV shows–already won't fool much face-recognition software. Completely blacking out each face in a video clip *would* do the job, but this would be of limited use if law-enforcement agencies wanted to follow up evidence of suspicious behavior once they had a court warrant. The function of the privacy-preserving algorithms that Gross is helping to create, he explains, is to automatically take the average values of individuals' faces and, from those, synthesize new facial images, then superimpose those new images over the originals. "It may seem like the opposite technology," Gross says, "but actually, it's just the other side of face recognition."

by Mark Williams Pontin