

What Caricatures Can Teach Us About Facial Recognition

By Ben Austen [✉](#) July 15, 2011 | 9:26 pm | [Wired August 2011](#)



Wired asked four top caricaturists to sketch the writer of this story. The results are shown here and throughout the story. To read about how writer Ben Austen reacted to the images, see the end of the story.

Photo: Joshua Anderson; caricature: Court Jones


Our brains are incredibly agile machines, and it's hard to think of anything they do more efficiently than recognize faces. Just hours after birth, the eyes of newborns are drawn to facelike patterns. An adult brain knows it's seeing a face within 100 milliseconds, and it takes just over a second to realize that two different pictures of a face, even if they're lit or rotated

in very different ways, belong to the same person. Neuroscientists now believe that there may be a specific region of the brain, on the fusiform gyrus of the temporal lobe, dedicated to facial recognition.

Perhaps the most vivid illustration of our gift for recognition is the magic of caricature—the fact that the sparest cartoon of a familiar face, even a single line dashed off in two seconds, can be identified by our brains in an instant. It's often said that a good caricature looks more like a person than the person himself. As it happens, this notion, counterintuitive though it may sound, is actually supported by research. In the field of vision science, there's even a term for this seeming paradox—the caricature effect—a phrase that hints at how our brains *misperceive* faces as much as perceive them.

Human faces are all built pretty much the same: two eyes above a nose that's above a mouth, the features varying from person to person generally by mere millimeters. So what our brains look for, according to vision scientists, are the outlying features—those characteristics that deviate most from the ideal face we carry around in our heads, the running average of every visage we've ever seen. We code each new face we encounter not in absolute terms but in the several ways it differs markedly from the mean. In other words, to beat what vision scientists call the homogeneity problem, we accentuate what's most important for recognition and largely ignore what isn't. Our perception fixates on the upturned nose, rendering it more porcine, the sunken eyes or the fleshy cheeks, making them loom larger. To better identify and remember people, we turn them into caricatures.

Ten years ago, the science of facial recognition—until then a somewhat esoteric backwater of artificial-intelligence research—suddenly became a matter of national security. The hazy closed-circuit images of Mohamed Atta, taped breezing through an airport checkpoint in Portland, Maine, enraged Americans and galvanized policymakers to fund research into automated recognition systems. We all imagined that within a few years, as soon as surveillance cameras had been equipped with the appropriate software, each face in a crowd would stand out like a thumbprint, its unique features and configuration offering a biometric key that could be immediately checked against any database of suspects.



Pawan Sinha, director of the Sinha Laboratory for Vision Research at MIT, thinks caricature is the key to better computer vision. For his Hirschfeld Project, to be started this year, Sinha's lab will analyze hundreds of caricatures by dozens of different artists, in order to isolate the facial proportions that are most important for recognition. This chart shows some of the myriad measurements that might prove crucial, like distance from pupil to pupil, distance from bottom lip to chin, or the area of the forehead.

But now a decade has passed, and face-recognition systems still perform miserably in real-world conditions. It's true that in our digital photo libraries, and now on Facebook, pictures of the same person can be automatically tagged and collated with some accuracy. Indeed, in a recent test of face-recognition software sponsored by the National Institute of Standards and Technology, the best algorithms could identify faces more accurately than humans do—at least in controlled settings, in which the subjects look directly at a high-resolution camera, with no big smiles or other displays of feature-altering emotion. To crack the problem of real-time recognition, however, computers would have to recognize faces as they actually appear on video: at varying distances, in bad lighting, and in an ever-changing array of expressions and perspectives. Human eyes can easily compensate for these conditions, but our algorithms remain flummoxed.

Given current technology, the prospects for picking out future Mohamed Attas in a crowd are hardly brighter than they were on 9/11. In 2007, recognition programs tested by the German federal police couldn't identify eight of 10 suspects. Just this February, a couple that accidentally swapped passports at the airport in Manchester, England, sailed through electronic gates that were supposed to match their faces to file photos.

All this leads science to a funny question. What if, to secure our airports and national landmarks, we need to learn more about caricature? After all, it's the skill of the caricaturist—the uncanny ability to quickly distill faces down to their most salient features—that our computers most desperately need to acquire. Better cameras and faster computers won't be enough. To pick terrorists out of a crowd, our bots might need to go to art school—or at least spend some time at the local amusement park.

In the 19th century, law enforcement knew that exaggerated art could catch crooks. When New York's Boss Tweed, on the lam in Spain, was finally arrested in 1876, he was identified not with the aid of a police sketch but with a Thomas Nast caricature from *Harper's Weekly*. Today, though, most police departments use automated facial-likeness generators, which tend to create a bland, average face rather than a recognizable portrait of the guilty party. Paul Wright, the president of Identi-Kit, one of the most commonly used composite systems in the US, concedes that the main value of his product is in ruling out a large fraction of the population. "Half the people might say a composite sketch looks like Rodney Dangerfield, another half like Bill Clinton. But it's not useless. It doesn't look like Jack Nicholson."

Visit the annual convention of the International Society of Caricature Artists and you'll find people who describe their face-depiction skills in far less modest terms. Take Stephen Silver, who began his career 20 years ago as a caricaturist at Sea World and is now a character designer for TV animation studios. "If they used caricatures for police composites today," Silver says, "people would be like, 'What is this, a joke?' But the cops would catch the guy. If I drew a caricature, the guy would be shit out of luck."



Daniel Almariei's caricature of author Ben Austen
Photo: Joshua Anderson; caricature: Daniel Almariei

Silver is one of 188 artists from 13 different countries who attended the most recent ISCA gathering, in Las Vegas. Over five days, and sometimes late nights, these artists draw one another's faces over and over again, often in orgiastic clusters, the artist-subject pairings shifting repeatedly and assuming every conceivable angle. The caricatures produced are eventually displayed and voted on by the attendees, with the first-place winner awarded a Golden Nosey trophy. Silver won the prize in 2000, and it's easy to see why. As he scans a room, he can size up faces and get the drop on each at a glance.

"I don't care how many wrinkles there are around the eye or if there's stubble," he says. "Those features aren't going to help me. You know who a person is from basic shapes." He spies a red-haired woman across the room, takes aim at her head. "Do you see how its meat is all on the outside?" he asks. "With the features crammed into the center?" Next his sights shift to an African-American woman drawing busily at a foldout table. Her head is actually tiny, Silver points out, but the span from her bottom lip to the base of her neck is immense.

This sort of instant insight is precisely what computers have trouble generating. "The miraculous thing about caricature artists is that they're able to zero in on the most distinctive aspect of somebody," says Erik Learned-Miller of the Computer Vision Laboratory at the University of Massachusetts, Amherst. "We still don't know how to do that in computer vision. People are working very hard to write programs that find just that combination of two or three things that give a person away."

At the University of Central Lancashire in England, Charlie Frowd, a senior lecturer in psychology, has used insights from caricature to develop a better police-composite generator. His system, called EvoFIT, produces animated caricatures, with each successive frame showing facial features that are more exaggerated than the last. Frowd's research supports the idea that we all store memories as caricatures, but with our own personal degree of amplification. So as an animated composite depicts faces at varying stages of caricature, viewers respond to the stage that is most recognizable to them. In tests, Frowd's technique has increased identification rates from as low as 3 percent to upwards of 30 percent.

“A lot of people think that caricature is about picking out someone's worst feature and exaggerating it as far as you can. That's wrong. Caricature is finding the truth.”

To achieve similar results in computer face recognition, scientists would need to model the caricaturist's genius even more closely—a feat that might seem impossible if you listen to some of the artists describe their nearly mystical acquisition of skills. Jason Seiler, the 2008 Golden Nosey winner, recounts how he trained his mind for years, beginning in middle school, until he gained what he regards as nothing less than a second sight. “You know at the end of *The Matrix* when Keanu Reeves sees the code falling everywhere, and all of a sudden he knows he's the One?” Seiler says with utter earnestness. “It's a lot like that.” For Roger Hurtado, an ISCA master from Chicago, the transformation was similar. “Suddenly everyone turned into a caricature,” he says. “I couldn't turn it off. You become incredibly sensitive to little details about people's faces that others wouldn't catch.” He adds: “It makes it hard to date.”

But when you talk to these artists about their process, you realize that the psychologists have gotten the basics down pretty well. When Court Jones, the 2005 Golden Nosey winner, describes how he teaches the craft to younger artists, he lays out exactly the algorithm that vision scientists believe humans use to identify faces. Students, he says, should imagine a generic face and then notice how the subject deviates from it: “That's what you can judge all other faces off of.”

Also, just as a vision scientist would predict, symmetrical faces—those close to our internal average—are especially difficult to caricature. People at the convention mention struggles with Katy Perry and Brad Pitt; the animator Bill Plympton, a guest speaker at the convention, tells me that Michael Caine has long been a *bête noire*. The same principle explains why the person at the convention with maybe the least symmetrical of faces appears by week's end in no fewer than 33 works of art on the ballroom walls. Kerim Yildiz, a 3-D designer from Montreal, possesses not merely a big-nose face, one caricaturist informs me, but a “big-nose-thick-eyebrow-glasses-ponytail-crazy-facial-hair” face. As a caricaturist himself, Yildiz understands his appeal. “With me it's all on the surface,” he says. “It's cool. It's my thing.”

“A lot of people think that caricature is about picking out someone's worst feature and exaggerating it as far as you can,” Seiler says. “That's wrong. Caricature is basically finding the truth. And then you push the truth.”

The caricaturists in Vegas who can bring a face to life this way are an eclectic bunch. They come from all over, make their living wherever they can find it. Angie Jordan, who started caricaturing while in the Army when someone asked her to “draw the captain but make it

funny,” skips one day of the convention to work a four-hour corporate gig in Atlanta. Between sketches for the conference, Roger Hurtado busies himself with big-head drawings of the 1985 Chicago Bears—they of “The Super Bowl Shuffle”—for a pizzeria promotion back home. (“People think Ditka’s eyes are close together, but they’re not,” Hurtado says. “It’s the wideness of his head. He does have a tiny chin close to his mouth.”) There’s a large contingent from a Tokyo caricature school that seems to toil sweatshop hours. Another Japanese artist, who registers for the conference under the name Toramaru, dresses each day in a full-body fleece tiger suit, from footie tiger toes to a tiger head that flops down over his forehead. In the caricatures of him drawn by his peers, he is depicted (not inaccurately) as a man-child of beatific innocence.



Jason Seiler's caricature of author Ben Austen
Photo: Joshua Anderson; caricature: Jason Seiler

From the various amusement parks around San Diego comes a group of artists who call themselves the Beastheads, a name that reflects their ethos of (as one of the members, Sea World cartoonist Andy Urzua, puts it) “just jacking faces, being the most extreme, *beasting* it.” And in fact, the theme park caricaturists tend to be among the most daring in their treatments. Brian Oakes, a Beasthead who did time behind the counter of a Taco Bell in Buffalo before taking his first amusement park job, draws Toramaru mounting Hello Kitty from behind, his tiger mitts on her cartoon rump, the lifeless costume head juxtaposed with the straining, sweating face beneath it.

On one afternoon, the artists assemble in the hotel ballroom to participate in a likeness competition, in which everyone will get five minutes to draw the same set of photographed

faces. In the moments before the contest begins, there is a flurry of confusion over rules: the acceptable paper size, the number and variety of drawing implements. An artist at his first ISCA convention uneasily raises his hand to ask if the point is to draw more realistically than everyone else. Robert Bauer, the organization's outgoing president, tries to set the newbie straight. "It's a likeness competition—for caricature artists. That's the whole point."

One of the first photographs is of a young woman with an enormous smile and a small tiara. The artists get to work forming the diamond shape of her head and the slits of her eyes, filling at least a third of the page with her ballooned cheeks or neat rows of towering teeth. From the front of the room, Bauer shouts, "No cheating. No looking on someone else's paper." One participant adds a halo of stars around the drawn face. In Brian Oakes' caricature, the woman has a small head and a giant breast tumbling out of her shirt. An artist named Daniel Almariei, originally from Romania and now living in West Nyack, New York, soon attracts a crowd of onlookers—both for his picture, which is shaping up as masterful, and for the sheer entertainment value of his performance. He abruptly flips the page sideways to shade or to get a curved line just so; he holds the drawing above him, as if to the light, studying it with a squint and a single raised brow, his dramatic frown forming an upside-down U. Almariei says of his art at one point, "I work with a lot of expression," which seems a description as much of his sketching as of his sketches.

The next photo is of the hotel's chef, and the best likenesses all capture the longness of his face, its rectangular shape, and the way his features are crowded down toward the bottom of it. An artist known as Big Al curses what he's committed to paper. "If I take longer than three minutes, it's awful. I'm used to capturing personality quickly." People again hover behind Almariei as he works. His chef conveys a certain attitude—betraying a slight smirk, with a hyperbolic twist of the cheeks that seems exactly right—and Almariei will, in due course, be announced as the likeness competition winner.

With so many renditions of the same face, I also see examples of what vision scientists call the other-race effect, which addresses why peoples throughout history have believed that those of a different race or ethnicity all look alike. The theory goes that we form our internalized average face by gazing upon people who resemble ourselves. Because members of another race can differ greatly—and all in the same way—from our prototype, we end up disregarding important information for distinguishing among those "others." In several of the coarser caricatures of the chef, an African-American, he is given stereotypically thick lips and a broad nose that don't match his own.

The final face to flash on the screen is of a well-coiffed elderly white woman. This, Robert Bauer announces, is his own mother, who has been undergoing radiation therapy for the past several months; he plans to present her with the bulk of these drawings over the Christmas holiday. Later, when the completed caricatures of her are arranged on the floor for viewing, Bauer sidles over and proclaims one of them a perfect likeness. "I see that one and I see my mom," he says. Bauer, who owns a caricature business that books artists for events across the country, explains how capturing a likeness has less to do with the depiction of individual features than with their placement in relationship to one another. "It's how the human brain recognizes a face: When the ratios between the features are correct, you see that face instantly."

While he speaks, I notice two other drawings in the arrangement—both, it turns out, by San

Diego Beastheads—that Bauer is pretending not to notice. One depicts his mother as a lizardly geriatric, standing bowlegged, birthing from the folds of her skirt an adult’s head that resembles almost perfectly the man in his forties standing at my side, with his deeply receding hairline, trim goatee, and russet-potato nose. The other drawing shows only a distended plot of earth, a tombstone atop it, rip carved into its front. “It’s all about proportions,” Bauer is saying.



Glenn Ferguson's caricature of author Ben Austen
Photo: Joshua Anderson; caricature: Glenn Ferguson

Did you hear the one about the vision scientist who used only caricaturists as his test subjects? He exaggerated his findings! Pawan Sinha, director of MIT’s Sinha Laboratory for Vision Research, and one of the nation’s most innovative computer-vision researchers, knows that caricatures are meant to be humorous, grotesque, and outlandish—he dabbles as a caricaturist himself, drawing occasionally for university publications. But Sinha also contends that these simple, exaggerated drawings can be objectively and systematically studied and that such work will lead to breakthroughs in our understanding of both human and machine-based vision. His lab at MIT is preparing to computationally analyze hundreds of caricatures this year, from dozens of different artists, with the hope of tapping their intuitive knowledge of what is and isn’t crucial for recognition. He has named this endeavor the Hirschfeld Project, after the famous *New York Times* caricaturist Al Hirschfeld.

Quite simply, the Hirschfeld Project would reverse-engineer the caricaturist’s art. By analyzing sketches, Sinha hopes to pinpoint the recurring exaggerations in the caricatures that most

strongly correlate to observable deviations in the original faces. The results, he believes, will ultimately produce a rank-ordered list of the 20 or so facial attributes that are most important for recognition: “It’s a recipe for how to encode the face,” he says. In preliminary tests, the lab has already isolated what seem to be important ingredients—for example, the ratio of the height of the forehead to the distance between the top of the nose and the mouth.

On a given face, four of 20 such Hirschfeld attributes, as Sinha plans to call them, will be several standard deviations greater than the mean; on another face, a different handful of attributes might exceed the norm. But in all cases, it’s the exaggerations that hold the key. As matters stand today, an automated recognition system must compare its target faces against the millions of continually altering faces it encounters. But so far the software doesn’t know what to look for amid this onslaught of variables. Armed with the Hirschfeld attributes, Sinha hopes that computers can be trained to focus on the features most salient for recognition, tuning out the torrent of noise. “Then,” Sinha says, “the sky is the limit—all the applications one can think of with a face-recognition system open up.”

On a visit to MIT a few weeks after the convention, I bring along some footage of the Las Vegas gathering to show Sinha and a few of his students. They are particularly interested in the speed competition, an even more intense and raucous affair than the likeness contest. In half a dozen heats, groups of artists scramble to see who can draw the greatest number of recognizable caricatures in a five-minute span. The artists’ hands flutter like hummingbirds over their pages, and quick-sketched faces pile up on the floor. A diminutive Japanese woman in some sort of vintage Pinky Tuscadero biker-chick outfit draws Freddy Krueger-style, with special ink extensions on all five fingers of one hand. Toramaru, among the fastest eight to reach the final round, tilts his head back at the three-minute mark and emits a tigerly roar. Steve Dorris, the event’s 2008 winner, who has a Teenage Mutant Ninja Turtles T-shirt stretched over his protruding belly, breathes heavily from the exertion. Another competitor, Beejay Hawn, increases her speed by drawing according to a strict regimen, starting with a single, counterclockwise line to represent one side of the face, then depicting eyes, nose, and mouth, in that order, unless one of the features is so large that it allows her to skip another altogether and save a valuable second. The eventual winner, to be announced at the convention’s closing banquet, is Yonie Woo, the Usain Bolt of caricaturists, who whips off 21 drawings in five minutes, or one every 14 seconds.

Sinha is keen to witness that magical moment when the likeness is achieved, when the combination of just a few minimally represented features suddenly captures what’s unique in a visage. It happens so fast that we watch the footage again, in slo-mo. A contestant places a big smile inside a pointed chin, and the resemblance is already apparent. Another sketches a helmet of hair, connects an L-shaped nose to a left eye, and—*bam!*—there’s the seated subject. A third depicts nothing more than a jawline, a top lip, and a droopy eye—and suddenly it’s unmistakable. Watching all this instantaneous mastery of the face, Sinha gets an idea: What if the lab analyzed the caricatures produced at the convention?

Better yet, counter his students, why don’t they study the brains of the caricaturists? With an electroencephalogram, they could see exactly which features are required in a caricature to elicit the occurrence of an N170, a key neural response that people exhibit when looking at a face. An fMRI, meanwhile, could watch while a caricaturist differentiated between familiar and unfamiliar faces, detecting whether the artist’s brain displayed any uncommon neural activity.

Yuri Ostrovsky, a postdoctoral fellow, mentions recent research on “super-recognizers” that suggests they may have a larger area of the brain dedicated to identifying faces. “It could be that caricature artists have those larger areas, too,” he says. “We should look into this, Sinha. We just need to bring our scanner to the next convention.”

Author’s reaction:

When Ben Austen saw the giant yellow teeth on an artist’s rendering of his face for his story on caricatures and facial recognition, his immediate reaction was, “I’m gonna go brush my teeth. Better use the whitening toothpaste.” Austen then showed the drawings to his family. Gazing at a portrait of Ben with a misshapen face, his wife, Danielle, blurted, “That doesn’t look like a human ... At least not one I want to be married to.” Ben and Danielle eventually got (mostly) comfortable with the exaggerated versions of his actually pretty-hip-looking self, but daughter Lusia, 6, never did. “This one makes me feel good that you’re my father and I might look like you,” she said. “These other ones make me feel bad.” Us, too, Lusia. Us, too.

Ben Austen (b_austen@gmail.com) is a contributing editor at Harper’s Magazine. His work will be included in The Best American Travel Writing 2011.