by HANY FARID

SEEING IS NOT BELIEVING

Doctoring digital photos is easy. Detecting it can be hard

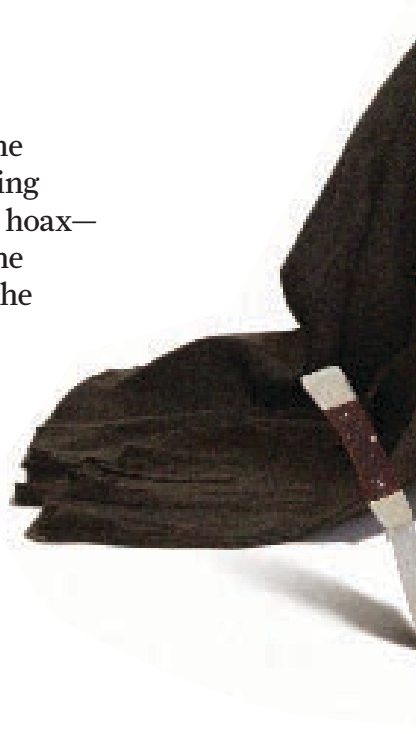

JUST DAYS AFTER SARAH PALIN'S SELECTION last August as the Republican vice presidential candidate, a photo of a bikini-clad, gun-toting Palin blitzed across the Internet. Almost as quickly, it was revealed as a hoax—a crude bit of Photoshop manipulation created by splicing an image of the Alaska governor's head onto someone else's body. From start to finish, the doctoring probably took no more than 15 minutes.

Altering digital imagery is now ubiquitous. People have come to expect it in the fashion and entertainment world, where airbrushing blemishes and wrinkles away is routine. And anyone surfing the Web is routinely subjected to crude photographic mashups like the Palin hoax, whose creators clearly aren't interested in realism but in whatever titillation or outrage they can generate.

But other photo manipulations demonstrate just how difficult it has become to tell altered images from the real thing. For example, in 2005 Hwang Woo-Suk, a South Korean professor, published a paper in one

ILLUSTRATION: MICHAEL ELINS;
BILL GATES: PETER MACDIARMID/GETTY IMAGES

ROCKET DUSTUP: A July 2008 photo [left] shows four Iranian missiles streaking skyward. But only three of those rockets actually left the ground; a fourth was digitally added. The altered image was first posted on the Web site of Sepah News, the media arm of Iran's Revolutionary Guard, and then published by media outlets around the world. Careful observers pointed out that portions of the faked rocket's exhaust plume and dust cloud had obviously been duplicated from its neighbors'. Sepah News soon replaced the faux photo with the original [right] without explanation.

of the most prestigious scientific journals, *Science*, claiming groundbreaking advances in stem-cell research. But at least 9 of the 11 uniquely tailored lines of stem cells that Hwang claimed to have made were fakes. Much of the evidence for those 9 lines of stem cells involved doctored photographs.

Apparently, Hwang's fabrication was not an isolated occurrence. Mike Rossner, then the managing editor of *The Journal of Cell Biology*, estimated that 20 percent of the manuscripts his journal accepted contained at least one image that had been inappropriately manipulated. Since then, a number of scholarly journals have implemented new fraud-detection procedures, such as software that makes it easier to compare images within or between documents. The incidence of image fraud in scholarly publishing has not declined, though; indeed, it seems to be on the rise.

A more recent example of photo tampering came to light in July 2008. Sepah News, the media arm of Iran's Revolutionary Guard, celebrated the country's military prowess by releasing a photo showing the simultaneous launch of four missiles. But one of those missiles had, in fact, failed to launch. The truth emerged after Sepah circulated the original photo showing three missiles in flight—but not before the faked image appeared on the front pages of the *Chicago Tribune,* the *Financial Times*, and the *Los Angeles Times*. If the world could be fooled by such a photo, then what's to prevent any country or militant group from using doctored images to intimidate?

To be sure, photographic alterations have existed about as long as photography itself. But before the digital age, such deceptions required mastery of complex and time-consuming darkroom techniques. Today anyone with a modicum of computer skills can call on powerful and inexpensive software to alter digital images. And as sophisticated forgeries appear with alarming frequency, people's belief in what they see has been eroded.

Over the past few years, the field of digital-image forensics has emerged to combat this growing problem and return some level of trust in photographs. By using computer methods to look at the underlying patterns of pixels that make up a digital image, specialists can detect the often-subtle signatures of manipulated images that are invisible to the naked eye.

NEARLY EVERY DIGITAL FORGERY starts out as a photo taken by a digital camera. The camera's image sensor acts as the film. It consists of a two-dimensional array of photoelectric elements that become electrically charged when exposed to light, which is why this type of light sensor is called a charge-coupled device, or CCD. The amount of charge is proportional to the light's intensity, so the electrical pattern captured by the CCD faithfully represents the light pattern striking the sensor.

Although exquisitely sensitive to intensity, the CCD elements can't detect the light's wavelength—that is, its color. So a device called a color-filter array is overlaid on the CCD, enabling each element to record a limited range of wavelengths corresponding to red, green, or blue.

After taking a picture, the camera transfers the pattern of electrical charges to the camera's memory, where it is represented as an array of pixels. A 6-megapixel camera, for example, has a CCD sensor with 6 million elements and takes digital images of up to 6 million pixels each. The charge or light intensity is translated into a number, 0 being the minimum and 255 the maximum. In a full-resolution color image, each pixel is assigned three such numbers, one for the intensity of red, one for green, and one for blue. But as noted above, the color-filter array initially assigns each CCD element just one color. So the camera fills in the missing color values by interpolating across neighboring pixels. These three values can yield more than 16 million colors.

Digital images can be stored in a number of formats. The most basic is raw format, in which the pixel values are stored exactly as they're recorded by the CCD, with no interpolation. This format is efficient,

COURTESY HANY FARID

as only one number is stored per pixel, but it requires any subsequent photo-editing software to perform the interpolation. The remaining image formats fall into one of two categories: nonlossy and lossy. Nonlossy formats, such as TIFF, PNG, and BMP, compress an image file by representing redundant or repetitive data using a kind of digital shorthand; when the file is subsequently expanded, the redundant data can be retrieved, so there is no loss of information. The lossy formats all compress their files by permanently removing data. The GIF format, for instance, limits the number of colors in a compressed image from millions to typically a few hundred. The JPEG format, perhaps the most popular lossy format, compresses by removing some color and image details.

THE FIRST RULE in any forensic analysis must surely be "preserve the evidence." So you might think that lossy image compression, which deletes information, would be a forensic analyst's worst enemy. In fact, it's a great aid: The unique properties of lossy compression can be exploited to track manipulations.

Take the ubiquitous JPEG format. It uses a compression algorithm that transforms the underlying pixel values into a map of low, middle, and high frequencies, where the low frequencies correspond to areas where the color changes very little (a blue sky and white clouds, for instance) and the high frequencies correspond to rapidly changing colors (as in a flamboyant Hawaiian shirt). Human eyes are less sensitive to the minute details of the high frequencies, so in JPEG files these areas are compressed more than the lower frequencies.

The JPEG image format also specifies how compression and memory consumption are balanced. This balance is represented as a matrix, called a quantization table, with each value specifying how much each of 64 distinct frequencies in the image, in each of three channels specifying brightness and color, has been compressed.

If you've shopped for a digital camera lately, you know that different models of cameras often produce very different results, even if they have the same overall pixel count. That's because camera manufacturers balance image compression and quality in subtly different ways, which creates differences in the quantization table. That means the photos taken by a given model of camera will have a signature of sorts embedded within each JPEG file it produces. The quantization tables used by Photoshop and other photo-manipulation software are also distinct, so you can tell whether one of those programs has been used to alter the image. This allows for a crude form of digital-image "ballistics." In many cases, you can figure out what type of camera the photographer used to take the shot.

Another related technique that my group at Dartmouth College, in Hanover, N.H., is now studying makes use of the thumbnail image that every digital camera automatically creates along with each full-resolution photo. The thumbnail, which has a resolution of about 160 by 120 pixels, is the tiny image you see when you preview a photo you've just taken. To create the image, the camera takes the full-resolution image, filters it, selectively removes pixels, filters it again, and then adjusts the brightness and contrast. Our research shows that this image processing relies on algorithms that appear to vary between different camera models. In our experiments, we've been able to estimate the parameters used to create a given thumbnail. The next step will be to build a database of thumbnail parameters from a large array of camera makes and models, which can then be used to authenticate the source of an image. In addition, the thumbnail is itself saved as a JPEG, using a different quantization table, and this information can be used to further refine the camera's signature.

FONDA SPEAKS TO VIETNAM VETERANS at Anti-War Rally" reads the headline, and the accompanying photograph, purportedly from 1970, shows a young Jane Fonda sharing a stage with a fresh-faced John

HOW HE DID IT

Michael Elins specializes in illustrations that blend digital imagery with his own photography. His image of a down-at-the-heels Bill Gates, an homage to a *Saturday Evening Post* cover from 1924, started with a stock photograph of the Microsoft billionaire. Elins then took photos of a male model dressed as a hobo (note the knife held by an assistant) and of Zippy the dog. He blended the three photos together using professional editing software, digitally adding the flames and smoke to create the makeshift campfire. The scenario may be improbable, but the image is deceptively realistic.

IDOL CURIOSITY: The cast of the TV show "American Idol" posed for this shot at different times. You can see it in their eyes—literally. The white specks, or specular highlights, indicate the direction of the lighting. Randy Jackson and Paula Abdul [seated] were likely photographed together, while Simon Cowell and Ryan Seacrest [standing] each posed separately. The yellow arrows at right indicate the lighting source for Cowell [left], Seacrest [middle], and Jackson and Abdul [right].

AMERICAN IDOL: ASSOCIATED PRESS; COURTESY HANY FARID

Kerry. The image, used to discredit Kerry during his unsuccessful U.S. presidential campaign in 2004, was a fake, composited from two unrelated photos taken at different places in different years.

To create such a composite, it is often necessary to resize, rotate, or stretch portions of an image. Let's say you're creating a composite image by grafting one person's head onto another person's body. It's unlikely that the relative sizes of the two images match exactly, so you'll have to enlarge or shrink one of them. In the process, you'll alter the underlying pattern of pixels in a distinct and detectable way.

Consider a small 3- by 3-pixel patch. Each of those pixels has a number corresponding to its brightness. Now let's say you want to enlarge, or up-sample, that patch by a factor of two. Enlarging an image requires adding extra pixels; in this case, an extra row of pixels would be added after each original row, and you'd end up with a 3- by 6-pixel patch. The computer software automatically assigns each new pixel's brightness by averaging the values of its immediate neighbors. As a result, the new pixels are perfectly correlated to their neighbors. Such correlations are unlikely to occur naturally, so a forensics expert detecting their presence knows the image has been manipulated.

But what if the resized image has been enlarged to a lesser degree or reduced in size? In those cases, the periodic correlations among the pixels are trickier to spot, but they do exist. My group has developed a computer

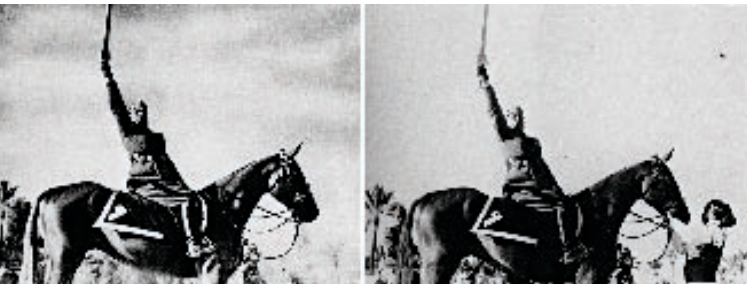

LEADING MAN: To create a more heroic portrait of himself, Italian dictator Benito Mussolini ordered the horse handler removed from the original 1942 photograph.

program that can detect such patterns by iteratively looking for pixels that are correlated to their neighbors. If detected, the correlations are then used to determine which portion of the image has been resized.

IN APRIL 2005, months before the romance between actors Brad Pitt and Angelina Jolie had been confirmed, *Star* magazine featured a cover photograph of the two über-celebrities strolling down the beach. The photo was a fake. The telltale sign was that the lighting on Pitt's and Jolie's faces was inconsistent with a single light source—in this case, the sun. To judge by the photo, you might conclude there had been two suns shining that day.

The amount of light that strikes a surface depends on the 3-D orientation of the surface relative to the 3-D position of the light source. But if you use photo-editing software to alter the image, you're dealing with a 2-D image, so it can be difficult to match the lighting conditions exactly. Studies show that our eyes are often insensitive to such lighting inconsistencies. But where human eyes fail, computers excel.

My group has developed a technique that can estimate the direction of the light source in an image by looking at a given object's 2-D surface contour, such as a person's jawline and chin. There, the surface orientation to the light source is always perpendicular to the contour.

By measuring the brightness and orientation at several points along the contour, we can estimate which direction the light is coming from. Then we can compare the lighting direction for that object to those of other objects in the photo. Any inconsistency in the lighting direction is evidence of tampering.

THE ASSOCIATED PRESS PLANNED to run a photo of the cast of the television series "American Idol." A photo editor at the wire service had doubts about the photo's authenticity, however, and contacted my lab for a second opinion.

When my colleagues and I examined the image closely, we immediately noticed that the small white specks of reflected light in each person's eyes, known as specular highlights, were inconsistent. To us, that was an obvious sign that the cast members had posed at different times and that the individual photos had been melded together.

The eyes are a beautiful tool for digital forensics, because they act as tiny windows into the world in which the photo was taken. By looking at the shape, color, and location of the specular highlights, we can learn quite a bit about the lighting that was used to take the photograph.

The location of the bright spot on the eye, for example, can indicate where the light source was positioned; multiple spots indicate more than one light source. The precise position of the specular highlights depends on both the curve of the eyeball and the relative orientations of the eye, the camera, and the light. The curve of the eye, it turns out, is remarkably similar from person to person, and there are very accurate 3-D models of eyeball shape. To calculate the relative orientation between a person's eyes and the camera, we can compare the shapes of the circular boundaries between the iris and the white part of the eye, known as the limbus. For example, if the person is directly facing the camera, the limbus in each eye will appear to be perfectly circular. As the orientation of the eyes changes relative to the camera, the limbus becomes more elliptical.

We can then use the shape and orientation of the limbus to estimate the direction to the light. Any inconsistencies in the lighting are evidence of tampering.

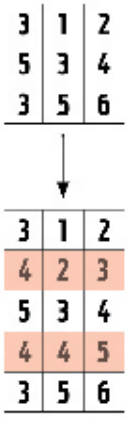

BIGGER: To enlarge an image, editing software adds extra pixels [pink], assigning each new pixel a value that is perfectly correlated to its neighbors'—a telltale sign of image tampering.

EVEN AS EXPERTS CONTINUE to develop techniques for exposing photographic frauds, new techniques for creating better and harder-to-detect fakes are also evolving. As in the battle against spam and computer viruses, it seems inevitable that the arms race between the forger and the forensic analyst will continue to escalate, with no clear victor. Improved image forensics will never be able to eradicate or prevent digital tampering, but these techniques can make it more time-consuming and difficult for forgers to ply their trade. Tomorrow's technology will almost certainly enable digital manipulations that today seem unimaginable, and the science of digital forensics will have to work hard to keep pace. It is my hope that these new techniques, along with a greater awareness of the technological possibilities and sensible updates in policy and law, will help the media, the courts, and the public contend with the exciting but often baffling events of our digital age. o

TO PROBE FURTHER The author's home page at Dartmouth College (http://www.cs.dartmouth.edu/farid) has more information about digital photo tampering as well as more examples.

RUBBED OUT: Cuban dictator Fidel Castro and Carlos Franqui fell out over the Soviet intervention in Czechoslovakia. Franqui went into exile, and Cuban authorities had his image expunged from photographs.

COVER UP: O.J. Simpson's 1994 mug shot following his arrest for murder was digitally darkened on *Time*'s cover, but not on *Newsweek*'s. *Time* apologized and issued a second cover but said the intent wasn't racist.