DIGITAL IMAGE FORENSICS

Modern software has made manipulation of photographs easier to carry out and harder to uncover than ever before, but the technology also enables new methods of detecting doctored images

By Hany Farid

KEY CONCEPTS

- Fraudulent photographs produced with powerful, commercial software appear constantly, spurring a new field of digital image forensics.
- Many fakes can be exposed because of inconsistent lighting, including the specks of light reflected from people's eyeballs.
- Algorithms can spot when an image has a "cloned" area or does not have the mathematical properties of a raw digital photograph.

—The Editors

History is riddled with the remnants of photographic tampering. Stalin, Mao, Hitler, Mussolini, Castro and Brezhnev each had photographs manipulated—from creating more heroic-looking poses to erasing enemies or bottles of beer. In Stalin's day, such phony images required long hours of cumbersome work in a darkroom, but today anyone with a computer can readily produce fakes that can be very hard to detect.

Barely a month goes by without some newly uncovered fraudulent image making it into the news. In February, for instance, an award-winning photograph depicting a herd of endangered Tibetan antelope apparently undisturbed by a new high-speed train racing nearby was uncovered to be a fake. The photograph had appeared in hundreds of newspapers in China after the controversial train line was opened with much patriotic fanfare in mid-2006. A few people had noticed oddities immediately, such as how some of the antelope were pregnant, but there were no young, as should have been the case at the time of year the train began running. Doubts finally became public when the picture was featured in the Beijing subway this year and other flaws came to light, such as a join line where two images had been stitched together. The photographer, Liu Weiqing, and his newspaper editor resigned; Chinese government news agencies apologized for distributing the image and promised to delete all of Liu's photographs from their databases.

In that case, as with many of the most publicized instances of fraudulent images, the fakery was detected by alert people studying a copy of

THIS IMAGE HAS BEEN MODIFIED in several places. The digital forensic techniques described on the following pages could be used to detect where changes were made. The answers are given on the final page.

the image and seeing flaws of one kind or another. But there are many other cases when examining an image with the naked eye is not enough to demonstrate the presence of tampering, so more technical, computer-based methods—digital image forensics—must be brought to bear.

I am often asked to authenticate images for media outlets, law-enforcement agencies, the courts and private citizens. Each image to be analyzed brings unique challenges and requires different approaches. For example, I used a technique for detecting inconsistencies in lighting on an image that was thought to be a composite of two people. When presented with an image of a fish submitted to an online fishing competition, I looked for pixel artifacts that arise from resizing. Inconsistencies in an image related to its JPEG compression, a standard digital format, revealed tampering in a screen shot offered as evidence in a dispute over software rights.

As these examples show, because of the variety of images and forms of tampering, the forensic analysis of images benefits from having a wide choice of tools. Over the past five years my students, colleagues and I, along with a small but growing number of other researchers, have developed an assortment of ways to detect tampering in digital images. Our approach in creating each tool starts with understanding what statistical or geometric properties of an image are disturbed by a particular kind of tampering. Then we develop a mathematical algorithm to uncover those irregularities. The boxes on the coming pages describe five such forensic techniques.

The validity of an image can determine wheth-

[THE AUTHOR]

Hany Farid has worked with federal law-enforcement agencies and many other clients on uncovering doctored images. Farid is David T. McLaughlin Distinguished Professor of Computer Science and Associate Chair of Computer Science at Dartmouth College and is also affiliated with the Institute for Security Technology Studies at Dartmouth. He thanks the students and colleagues with whom he has developed digital forensic methods, in particular Micah K. Johnson, Eric Kee, Siwei Lyu, Alin Popescu, Weihong Wang and Jeffrey Woodward.

CHRISTOPHE ENA *AP Photo (bicyclists)*; LILI YANG *Redlink/Corbis (woman's head)*; SHEARER IMAGES/CORBIS *(fire hydrant)*

[LIGHTING]

IN A DIFFERENT LIGHT

Composite images made of pieces from different photographs can display subtle differences in the lighting conditions under which each person or object was originally photographed. Such discrepancies will often go unnoticed by the naked eye.

For an image such as the one at the right, my group can estimate the direction of the light source for each person or object (*arrows*). Our method relies on the simple fact that the amount of light striking a surface depends on the relative orientation of the surface to the light source. A sphere, for example, is lit the most on the side facing the light and the least on the opposite side, with gradations of shading across its surface according to the angle between the surface and the direction to the light at each point.

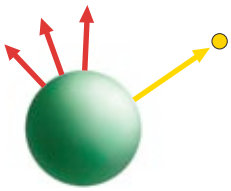

To infer the light-source direction, you must know the local orientation of the surface. At most places on an object in an image, it is difficult to determine the orientation. The one exception is along a surface contour, where the orientation is perpendicular to the contour (*red arrows above*). By measuring the brightness and orientation along several points on a contour, our algorithm estimates the light-source direction.

For the image above, the light-source direction for the police does not match that for the ducks (*arrows*). We would have to analyze other items to be sure it was the ducks that were added. —***H.F.***

HUGHES LÉGLISE-BATAILLE (*riot*); CHARRO BADGER (*inTheSunStudio* (*ducks*); LISA APFELBACHER (*illustration*)

[SHAPES]

EYES AND POSITION

Because eyes have very consistent shapes, they can be useful for assessing whether a photograph has been altered.

A person's irises are circular in reality but will appear increasingly elliptical as the eyes turn to the side or up or down (*a*). One can approximate how eyes will look in a photograph by tracing rays of light running from them to a point called the camera center (*b*). The picture forms where the rays cross the image plane (*blue*). The principal point of the camera—the intersection of the image plane and the ray along which the camera is pointed—will be near the photograph's center.

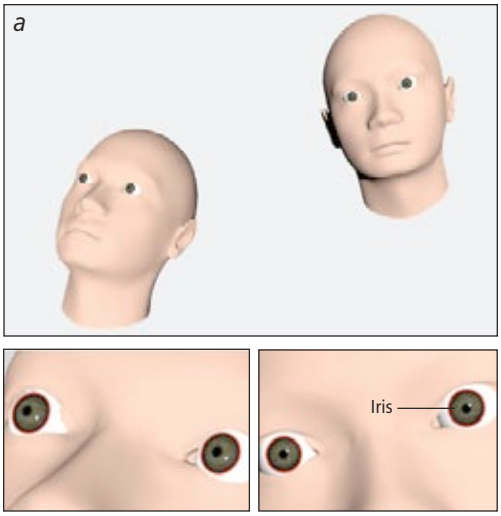

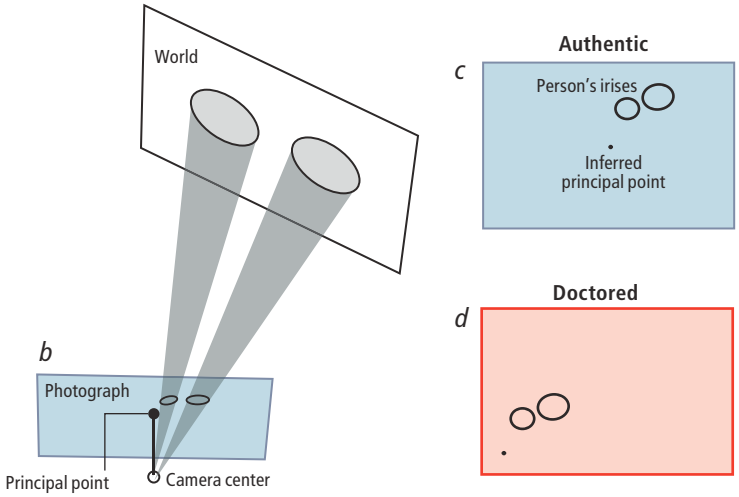

My group uses the shape of a person's two irises in the photograph to infer how his or her eyes are oriented relative to the camera and thus where the camera's principal point is located (*c*). A principal point far from the center or people having inconsistent principal points is evidence of tampering (*d*). The algorithm also works with other objects if their shapes are known, as with two wheels on a car.

The technique is limited, however, because the analysis relies on accurately measuring the slightly different shapes of a person's two irises. My collaborators and I have found we can reliably estimate large camera differences, such as when a person is moved from one side of the image to the middle. It is harder to tell if the person was moved much less than that. —***H.F.***

COURTESY OF HANY FARID (*a*); LISA APFELBACHER (*b–d*)

er or not someone goes to prison and whether a claimed scientific discovery is a revolutionary advance or a craven deception that will leave a dark stain on the entire field. Fake images can sway elections, as is thought to have happened with the electoral defeat of Senator Millard E. Tydings in 1950, after a doctored picture was released showing him talking with Earl Browder, the leader of the American Communist Party. Political ads in recent years have seen a startling number of doctored photographs, such as a faux newspaper clipping distributed on the Internet in early 2004 that purported to show John Kerry on stage with Jane Fonda at a 1970s Vietnam War protest. More than ever before, it is important to know when seeing can be believing.

Everywhere You Look

The issue of faked images crops up in a wide variety of contexts. Liu was far from the first news photographer to lose his job and have his work stricken from databases because of digital fakery. Lebanese freelancer Adnan Hajj produced striking photographs from Middle Eastern conflicts for the Reuters news agency for a decade, but in August 2006 Reuters released a picture of his that had obviously been doctored. It showed Beirut after being bombed by Israel, and some of the voluminous clouds of smoke were clearly added copies.

Brian Walski was fired by the *Los Angeles Times* in 2003 after a photograph of his from Iraq that had appeared on the newspaper's front page was revealed to be a composite of elements from two separate photographs combined for greater dramatic effect. A sharp-eyed staffer at another newspaper noticed duplicated people in the image while studying it to see if it showed friends who lived in Iraq. Doctored covers from newsmagazines *Time* (an altered mug shot of O. J. Simpson in 1994) and *Newsweek* (Martha Stewart's head on a slimmer woman's body in 2005) have similarly generated controversy and condemnation.

Scandals involving images have also rocked the scientific community. The infamous stem cell research paper published in the journal *Science* in 2005 by Woo Suk Hwang of Seoul National University and his colleagues reported on 11 stem cell colonies that the team claimed to have made. An independent inquiry into the case concluded that nine of those were fakes, involving doctored images of two authentic colonies. Mike Rossner estimates that when he was the managing editor of the *Journal of Cell Biol-*

[SPECULAR HIGHLIGHTS]

TELLTALE TWINKLES

Surrounding lights reflect in eyes to form small white dots called specular highlights. The shape, color and location of these highlights tell us quite a bit about the lighting.

In 2006 a photo editor contacted me about a picture of *American Idol* stars that was scheduled for publication in his magazine (*above*). The specular highlights were quite different (*insets*).

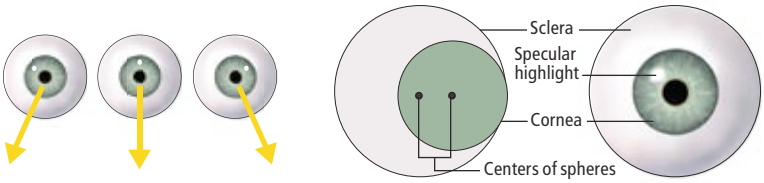

The highlight position indicates where the light source is located (*above left*). As the direction to the light source (*yellow arrow*) moves from left to right, so do the specular highlights.

The highlights in the *American Idol* picture are so inconsistent that visual inspection is enough to infer the photograph has been doctored. Many cases, however, require a mathematical analysis. To determine light position precisely requires taking into account the shape of the eye and the relative orientation between the eye, camera and light. The orientation matters because eyes are not perfect spheres: the clear covering of the iris, or cornea, protrudes, which we model in software as a sphere whose center is offset from the center of the whites of the eye, or sclera (*above right*).

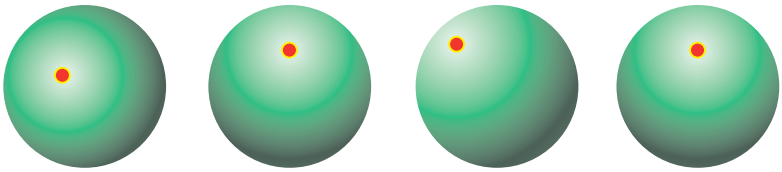

Our algorithm calculates the orientation of a person's eyes from the shape of the irises in the image. With this information and the position of the specular highlights, the program estimates the direction to the light. The image of the *American Idol* cast (*above; directions depicted by red dots on green spheres*) was very likely composed from at least three photographs. —***H.F.***

FOX NEWS (*American Idol*); LISA APFELBACHER (*eyes*); MELISSA THOMAS (*specular highlights*)

[DUPLICATION]

SEND IN THE CLONES

Cloning—the copying and pasting of a region of an image—is a very common and powerful form of manipulation.

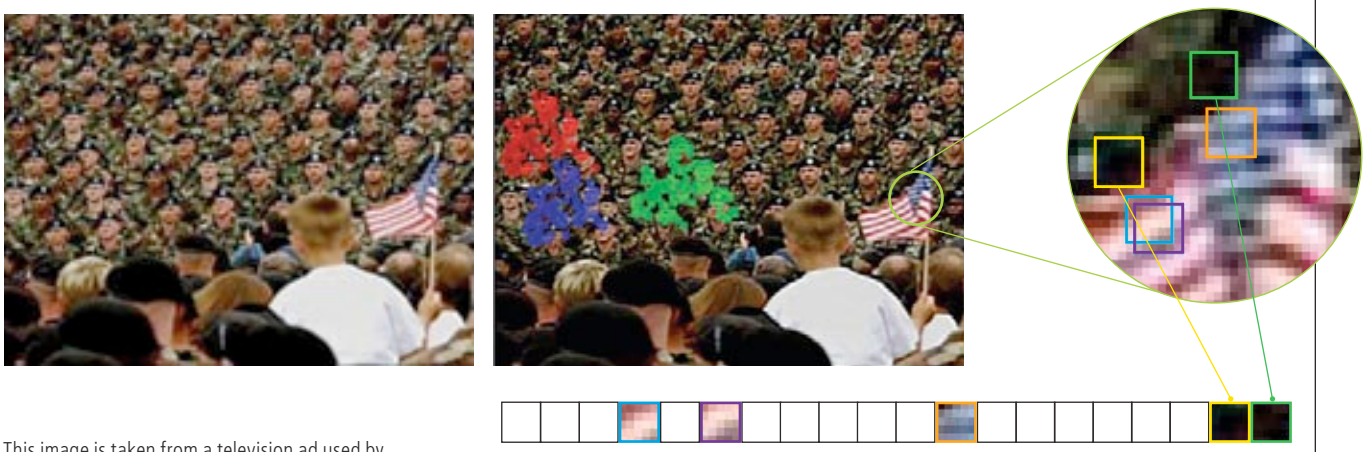

COURTESY OF HANY FARID

This image is taken from a television ad used by George W. Bush's reelection campaign late in 2004. Finding cloned regions by a brute-force computer search, pixel by pixel, of all possible duplicated regions is impractical because they could be of any shape and located anywhere in the image. The number of comparisons to be made is astronomical, and innumerable tiny regions will be identical just by chance ("false positives"). My group has developed a more efficient technique that works with small blocks of pixels, typically about a six-by-six-pixel square (*inset*).

For every six-by-six block of pixels in the image, the algorithm computes a quantity that characterizes the colors of the 36 pixels in the block. It then uses that quantity to order all the blocks in a sequence that has identical and very similar blocks close together. Finally, the program looks for the identical blocks and tries to "grow" larger identical regions from them block by block. By dealing in blocks, the algorithm greatly reduces the number of false positives that must be examined and discarded.

When the algorithm is applied to the image from the political ad, it detects three identical regions (*red, blue* and *green*). —**H.F.**

[RETOUCHING]

CAMERA FINGERPRINTS

Digital retouching rarely leaves behind a visual trace. Because retouching can take many forms, I wanted to develop an algorithm that would detect any modification of an image. The technique my group came up with depends on a feature of how virtually all digital cameras work.

A camera's digital sensors are laid out in a rectangular grid of pixels, but each pixel detects the intensity of light only in a band of wavelengths near one color, thanks to a color filter array (CFA) that sits on top of the digital sensor grid. The CFA used most often, the Bayer array, has red, green and blue filters arranged as shown at the right.

Each pixel in the raw data thus has only one color channel of the three required to specify a pixel of a standard digital image. The missing data are filled in—either by a processor in the camera itself or by software that interprets raw data from the camera—by interpolating from the nearby pixels, a procedure called demosaicing. The simplest approach is to take the average of neighboring values, but more sophisticated algorithms are also used to achieve better results. Whatever demosaicing algorithm is applied, the pixels in the final digital image will be correlated with their neighbors. If an image does not have the proper pixel correlations for the camera allegedly used to take the picture, the image has been retouched in some fashion.

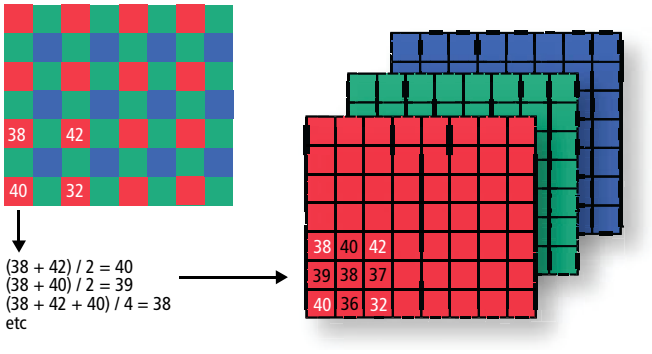

RICK WILKING *Reuters* (*note*); LISA APFELBACHER (*grids*)

My group's algorithm looks for these periodic correlations in a digital image and can detect deviations from them. If the correlations are absent in a small region, most likely some spot changes have been made there. The correlations may be completely absent if image-wide changes were made, such as resizing or heavy JPEG compression. This technique can detect changes such as those made by Reuters to an image it released from a meeting of the United Nations Security Council in 2005 (*left*): the contrast of the notepad was adjusted to improve its readability.

A drawback of the technique is that it can be applied usefully only to an allegedly original digital image; a scan of a printout, for instance, would have new correlations imposed courtesy of the scanner. —**H.F.**

ogy, as many as a fifth of the accepted manuscripts contained a figure that had to be remade because of inappropriate image manipulation.

The authenticity of images can have myriad legal implications, including cases involving alleged child pornography. In 2002 the U.S. Supreme Court ruled that computer-generated images depicting a fictitious minor are constitutionally protected, overturning parts of a 1996 law that had extended federal laws against child pornography to include such images. In a trial in Wapakoneta, Ohio, in 2006, the defense argued that if the state could not prove that images seized from the defendant's computer were real, then he was within his rights in possessing the images. I testified on behalf of the prosecutor in that case, educating the jurors about the power and limits of modern-day image-processing technology and introducing results from an analysis of the images using techniques to discriminate computer-generated images from real photographs. The defense's argument that the images were not real was unsuccessful.

Yet several state and federal rulings have found that because computer-generated images are so sophisticated, juries should not be asked to determine which ones are real or virtual. At least one federal judge questioned the ability of even expert witnesses to make this determination. How then are we to ever trust digital photography when it is introduced as evidence in a court of law?

Arms Race

The methods of spotting fake images discussed in the boxes have the potential to restore some level of trust in photographs. But there is little doubt that as we continue to develop software to expose photographic frauds, forgers will work on finding ways to fool each algorithm and will have at their disposal ever more sophisticated image manipulation software produced for legitimate purposes. And although some of the forensic tools may be not so tough to fool—for instance, it would be easy to write a program to restore the proper pixel correlations expected in a raw image—others will be much harder to circumvent and will be well beyond the average user. The techniques described in the first three boxes exploit complex and subtle lighting and geometric properties of the image formation process that are challenging to correct using standard photo-editing software.

As with the spam/antispam and virus/antivirus game, not to mention criminal activity in general, an arms race between the perpetrator and the forensic analyst is inevitable. The field of image forensics will, however, continue to make it harder and more time-consuming (but never impossible) to create a forgery that cannot be detected.

Although the field of digital image forensics is still relatively young, scientific publishers, news outlets and the courts have begun to embrace the use of forensics to authenticate digital media. I expect that as the field progresses over the next five to 10 years, the application of image forensics will become as routine as the application of physical forensic analysis. It is my hope that this new technology, along with sensible policies and laws, will help us deal with the challenges of this exciting—yet sometimes baffling—digital age. ■

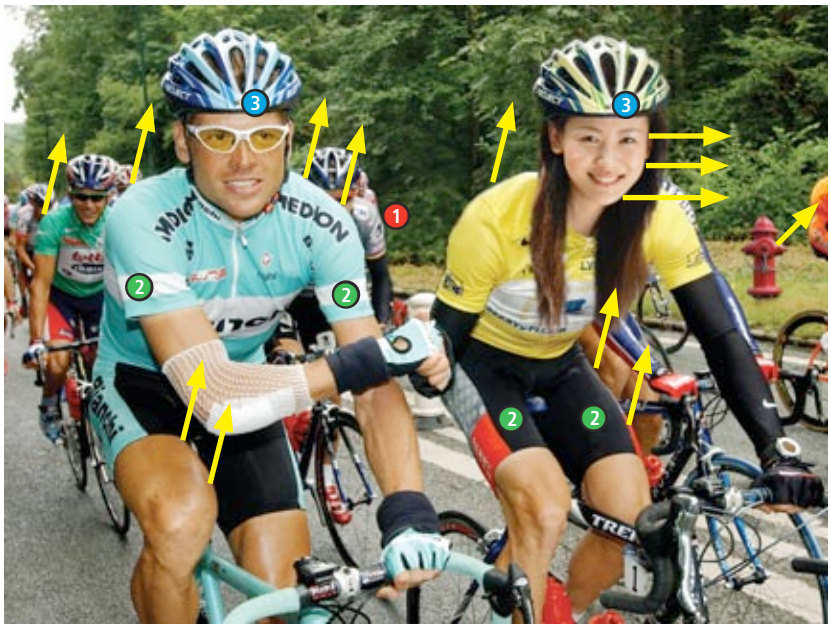

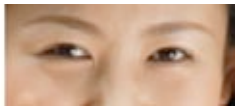

OPENER ANSWER: Inconsistent specular highlights (*bottom*) indicate the two leading cyclists were not photographed together. The light-source direction (*arrows*) for the girl's face conflicts with that of "her" body and the other cyclists. The added fire hydrant has yet another light-source direction. Cloned shrubs, grass and the curbside (1) cover cyclists in the background. Spoiled pixel correlations might reveal areas where retouching removed logos (2) and that the girl's helmet is doctored (3); it is copied from the man's but also has been recolored. The original photograph can be seen at www.SciAm.com/jun2008

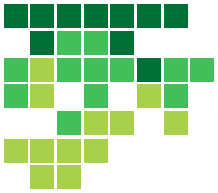

MORE TO EXPLORE

Exposing Digital Forgeries in Color Filter Array Interpolated Images. Alin C. Popescu and Hany Farid in *IEEE Transactions on Signal Processing,* Vol. 53, No. 10, pages 3948–3959; October 2005. Available at **www.cs.dartmouth.edu/farid/publications/sp05a.html**

Detecting Photographic Composites of People. Micah K. Johnson and Hany Farid. Presented at the 6th International Workshop on Digital Watermarking, Guangzhou, China, 2007. Available at **www.cs.dartmouth.edu/farid/publications/iwdw07.html**

Lighting and Optical Tools for Image Forensics. Micah K. Johnson. Ph.D. dissertation, Dartmouth College, September 21, 2007. Available at **www.cs.dartmouth.edu/farid/publications/mkjthesis07.html**

Hany Farid's Web site: **www.cs.dartmouth.edu/farid**