Digital Doctoring: can we trust photographs?

Hany Farid
Dartmouth College

We may have the impression that photography can no longer be trusted. From the tabloid magazines to the fashion industry, main-stream media outlets, political campaigns, and the photo hoaxes that land in our email in-boxes, doctored photographs are appearing with a growing frequency and sophistication. The truth is, however, that photography lost its innocence many years ago. The nearly iconic portrait of the U.S. President Abraham Lincoln (circa 1860), for example, was a fake, and only the beginning of a long history of photographic trickery. I will briefly explore the history and more modern examples of photographic tampering and discuss recent technological advances that have the potential to return some trust to photographs.

Abraham Lincoln and Winged Fairies

In the early part of his career, Southern politician John Calhoun was a strong supporter of slavery. It is ironic, therefore, that the nearly iconic portrait of Abraham Lincoln is a composite of Calhoun's body and Lincoln's head [Figure 1]. It is said that this was done because there was no sufficiently "heroic-style" portrait of Lincoln available. While the creation of such an image required significant skill and time, it was by no means unique. In the early part of the 1900s Stalin famously had his political enemies air-brushed out of official photographs. Between 1917 and 1920 two young girls in Cottingley, Yorkshire created an international sensation when they released photographs purportedly showing tiny winged fairy creatures [Figure 1]. And it wasn't until 1984 that some of the most spectacular photographs of World War I aerial combat first published in 1933 were exposed as fakes. The Brown Lady of Raynham, perhaps one of the most famous "ghost images," was a sensation when published in 1936, but was later discovered to have been created by superimposing two pictures on top of each other. It is believed that a doctored photograph contributed to Senator Millard Tydings' electoral defeat in 1950: the photo of Tydings conversing with Earl Browder, a leader of the American Communist party, was meant to suggest that Tydings had Communist sympathies [Figure 1]. And the list goes on – history is riddled with photographic tampering.

Figure 1. A portrait of John Calhoun, from which the portrait of Abraham Lincoln was created; the Cottingley fairies and their creator; and Senator Millard Tydings (right) purportedly chatting with Communist party leader Earl Browder (left).

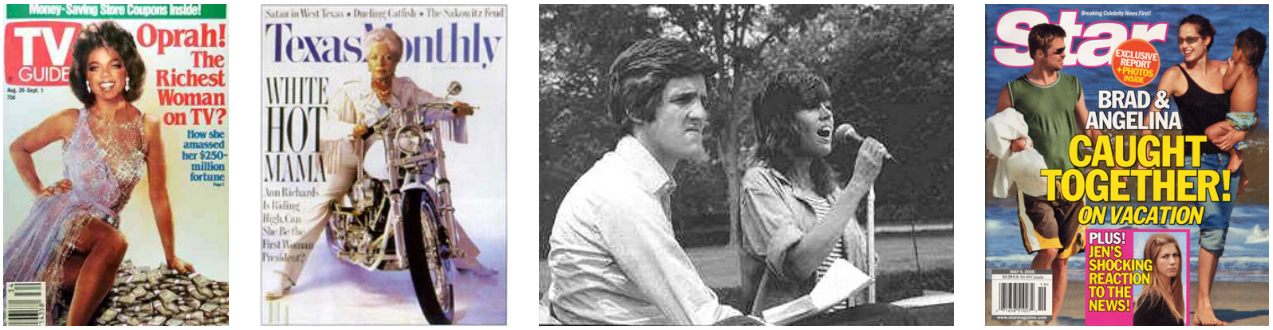

Figure 2. Oprah Winfrey's head and Ann-Margret's body; Governor Ann Richard's head and a model's body; and digital composites of Senator John Kerry with anti-war activist Jane Fonda and Brad Pitt with then rumored sweetheart Angelina Jolie.

Oprah Winfrey and Brad Pitt

With the advent of powerful computers and sophisticated software, the creation of photographic frauds has become increasingly easier. Interestingly the types of forgeries haven't changed much: attaching a person's head to another person's body, for example, remains a popular digital deception strategy. Among the best-known examples of this technique was the The August 1989 cover of *TV Guide*, which featured the head of popular daytime talk-show host Oprah Winfrey composited onto the body of actress Ann-Margret [Figure 2]. And in July of 1992, the cover of *TexasMonthly* showed Texas Governor Ann Richards astride a Harley-Davidson motorcycle: a picture created by splicing Richards' head onto the body of a model [Figure 2]; when asked if she objected to the image, Richards responded that since the model had such a nice body, she could hardly complain. The March 2005 cover of *Newsweek* featured a photograph of Martha Stewart with a headline that read "After Prison She's Thinner, Wealthier & Ready for Prime Time." The photograph, however, was a composite showing Stewart's head atop a (thin) model's body; its intent was apparently to illustrate what Stewart might look like when she was released from prison.

As with the Tydings fake, the use of compositing techniques to create the appearance of togetherness or relationship has also remained popular. In 1994, for example, *New York Newsday* published a composite of Olympic ice skaters Tanya Harding and Nancy Kerrigan in an improbable scene: practicing together at an ice rink shortly after Harding had an associate of her husband take Kerrigan out of competition with a blow to the leg. And in 2000, the University of Wisconsin at Madison – hoping to illustrate its diverse enrollment – doctored a brochure photograph by digitally inserting a black student in a crowd of white football fans (University officials said that they had spent the summer looking for pictures that would show the school's diversity – but had no luck). Reporters at the University's campus newspaper noticed lighting inconsistencies in the image and printed a story exposing the image as a fake. University officials apologized, calling the decision to use the image an "error in judgment". In the political arena, as Senator John Kerry was campaigning for the 2004 Democratic presidential nomination a doctored photo of Kerry sharing a stage with anti-war activist Jane Fonda was widely distributed [Figure 2]. Even after being revealed as a fake, the photograph did significant damage to Kerry's prospects by drawing attention to his controversial involvement in the anti-war movement following his service in Vietnam.

With the headline "Caught Together!", the April 2005 cover of *Star* magazine featured a photo that appeared to show actors Brad Pitt and Angelina Jolie – who were rumored to have started a romantic relationship – walking on the beach together [Figure 2]. The Star's readers were probably unaware that the picture was a composite of a photo of Pitt taken on a Caribbean island in 2005 and a picture of Jolie taken in Virginia a few years earlier.

Figure 3. The published (top) and original LA Times photographs showing a British soldier and Iraqi civilians.

Perhaps we have come accept and even expect a certain amount of photographic trickery when it comes to Hollywood and politics. When it comes to "hard news" like war-time reporting, however, the expectations have proven to be decidedly different. In March of 2003 a dramatic photograph of a British soldier in Basra, Iraq urging Iraqi civilians to seek cover was published on the front page of the *Los Angeles Times* [Figure 3]. The photograph was discovered to be a digital composite of two other images combined to "improve" the composition. In response, the outraged editors of the *Los Angeles Times* fired Brian Walski, a 20-year veteran news photographer. Similarly, in August of 2006, the Reuters news agency published a photograph showing the remnants of an Israeli bombing of a Lebanese town – an image that, in the week that followed, was revealed by hundreds of bloggers and nearly every major news organization to have been doctored with the addition of more smoke [Figure 4]. The general response was one of outrage and anger: the photographer, Adnan Hajj was accused of doctoring the image to exaggerate the impact of the Israeli shelling. An embarrassed Reuters retracted the photograph and removed from its archives nearly 1,000 photographs contributed by Hajj.

Figure 4. The published (left) and original (right) Reuters photograph showing the remnants of an Israeli bombing.

While historically they may have been the exception, doctored photographs today are increasingly impacting nearly every aspect of our society. While the technology to distort and manipulate digital media is developing at break-neck speeds, the technology to detect such alterations is lagging behind. To this end, I will describe some recent innovations for detecting digital tampering that have the potential to return some trust to photographs.

Exposing Digital Forgeries: lighting

A close examination of the *Star* cover of Pitt and Jolie reveals surprisingly obvious traces of tampering [Figure 5]. The setting and shadows suggest that this photograph was taken outdoors on a sunny day. There are several clues in this photograph as to the location of the sun. Jolie's shadow cast onto the sand, the shadow under her chin, her evenly illuminated face, and the lighting gradient around her right leg, all suggest that she is facing the sun. Given this position of the sun, we would expect the right side of Pitt's face to be illuminated. It is not. It is in shadow, which is impossible. It is clear that Pitt is facing the sun, which places the sun at a location at least 90 degrees away from the position of the sun illuminating Jolie. Were the lighting differences in this image more subtle, our manual analysis would most likely have been insufficient. We have, therefore, developed a computer program that automatically estimates the direction of an illuminating light source for each object or person in an image (Johnson and Farid 2005). By making some initial simplifying assumptions about the light and the surface being illuminated, we can mathematically express how much light a surface should receive as a function of its position relative to the light. A surface that is directly facing the light, for example, will be brighter than a surface that is turned away from the light. Once expressed in this form, standard techniques can be used to determine the direction to the light source for any object or person in an image. Any inconsistencies in lighting can then be used as evidence of tampering.

Figure 5. A composite of Brad Pitt and then rumored sweetheart Angelina Jolie (top); a composite of the American Idol host and judges (courtesy of Fox News and the Associated Press) (bottom).

A photograph of the host and judges for the popular television show *American Idol* was scheduled for publication when it caught the attention of a photo-editor [Figure 5]. Coming on the heels of several scandals that rocked major news organizations, the photo-editor was concerned that the image had been doctored. There was good reason to worry – the image was a composite of several photographs. A magnification of the host's and judge's eyes reveals inconsistencies in the shape of the specular highlight on the eyes suggesting that the people were originally photographed under different lighting conditions. We have shown that the location of a specular highlight on the eye can be used to determine the direction to the light source (Johnson and Farid 2007). Inconsistencies in the estimates from different eyes, as well as differences in the shape and color of the highlights, can therefore be used to reveal traces of digital tampering. In related work, Nishino and Nayar describe a technique for reconstructing, from the reflection on an eye, the image of the world surrounding a person and what they were looking at (Nishino and Nayar 2004).

Exposing Digital Forgeries: cloning

In order to create more smoke in his photograph, Hajj cloned (duplicated) parts of the existing smoke using a standard tool in Photoshop, a popular photo-editing software. In this case the duplication was fairly obvious because of the nearly identical repeating patterns in the smoke. When care is taken, however, it can be very difficult to visually detect this type of duplication. We have developed a computer program that can automatically detect image cloning (Popescu and Farid 2004) (a similar technique is described in (Fridrich, Soukal, and Lukáš 2003)). A digital image is first partitioned into small blocks. The blocks are then re-ordered so that they are placed a distance to each other that is proportional to the differences in their pixel colors. With identical and highly similar blocks neighboring each other in the re-ordered sequence, a region growing algorithm combines any significant number of neighboring blocks that are consistent with the cloning of an image region. Since it is statistically unlikely to find identical and spatially coherent regions in an image, their presence can then be used as evidence of tampering.

In a similar but more serious incident, Professor Hwang Woo-Suk and colleagues published what appeared to be ground-breaking advances in stem cell research (Hwang, et al. 2004). After its publication in *Science* in 2004, however, evidence began to emerge that the published results were manipulated and in places fabricated. After months of controversy, Hwang retracted the *Science* paper (Kennedy 2006) and resigned his position at Seoul National University. An independent panel investigating the accusations of fraud found, in part, that at least nine of the eleven customized stem cell colonies that Hwang had claimed to have made were fakes. Much of the evidence for those nine colonies, the panel said, involved doctored photographs of two other, authentic, colonies: the authors had digitally cloned their results. While the Hwang case garnered international coverage and outrage, it is by no means unique. In an increasingly competitive field, scientists are succumbing to the temptation to exaggerate or fabricate their results. Mike Rossner, the managing editor of the *Journal of Cell Biology* estimates that as many as 20% of accepted manuscripts to his journal contain at least one figure that has to be remade because of inappropriate image manipulation, and roughly 1% of figures are simply fraudulent (Pearson 2005).

Exposing Digital Forgeries: re-touching

While attending a meeting of the United Nations Security Council in September of 2005, U.S. President George W. Bush scribbled a note to Secretary of State Condoleezza Rice. The note read "I think I may need a bathroom break. Is this possible?" [Figure 6]. Because the original image was overexposed, a Reuters' processor selectively adjusted the contrast of the notepad prior to publication. This form of photo re-touching is quite common and can be used to alter a photograph in trivial or profound ways. We have developed a technique for detecting this form of tampering that exploits how a digital camera sensor records an image (Popescu and Farid 2005). Virtually all digital cameras record only a subset of all the pixels needed for a full-resolution color image. Instead, only a subset of the pixels are recorded by a color filter array (CFA) placed atop the digital sensor.

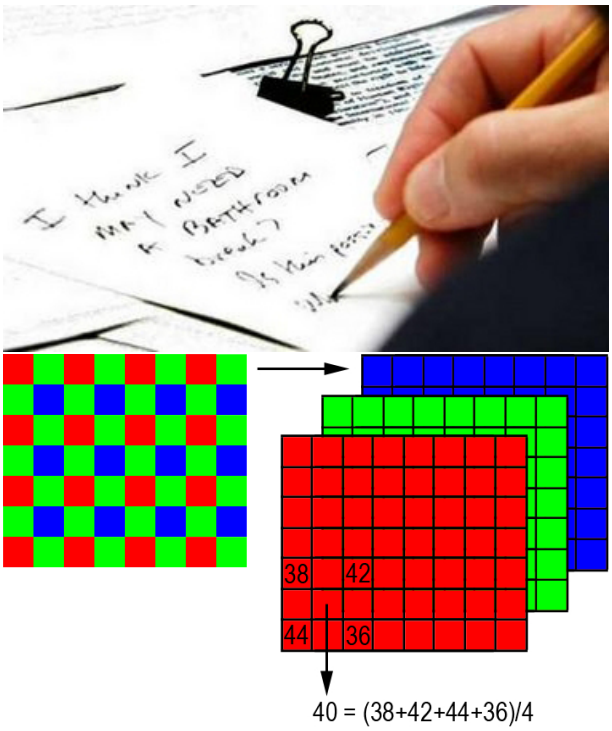

Figure 6. A note written by Bush was re-touched to improve readability, disrupting the color filter array correlations.

The most frequently used CFA, the Bayer array, employs three color filters: red, green, and blue [Figure 6]. Since only a single color sample is recorded at each pixel location, the other two color samples must be estimated from the neighboring samples in order to obtain a three-channel color image. The estimation of the missing color samples is referred to as CFA interpolation or demosaicking. In its simplest form, the missing pixels are filled in by spatially averaging the recorded values. Shown in Figure 6, for example, is the calculation of a red pixel from an average of its four recorded neighbors. Since the CFA is arranged in a periodic pattern, a periodic set of pixels will be precisely correlated to their neighbors according to the CFA interpolation algorithm. When an image is re-touched, it is likely that these correlations will be destroyed. As such, the presence or lack of these correlations can be used to authenticate an image, or expose it as a forgery.

Exposing Digital Forgeries: ballistics

Firearm ballistic experts routinely analyze bullets and bullet impacts to determine the type and caliber of a firearm. In some cases, unique grooves and scratches in the firearm barrel can be used to link a bullet to a specific weapon. In the field of camera ballistics, the goal is analogous: link an image to a specific camera, scanner, printer, etc.

Since the JPEG image format has emerged as a virtual standard, most devices and software encode images in this format. This compression scheme allows for some flexibility in how much compression is achieved. Manufacturers typically configure their devices differently to balance compression

and quality to their own needs and tastes. This choice is realized by adjusting the values in the JPEG quantization table, a set of 192 numbers. Smaller values yield less compression with higher quality, and larger values yield more compression with lower quality. Because camera manufacturers typically construct unique tables, the JPEG quantization table can be used to identify the source of an image (Farid 2006b). While this approach cannot distinguish between images taken with different cameras of the same make and model, it can be used to make a cruder distinction between different camera makes and models.

In related work, Lukáš and colleagues describe a more powerful technique to identify a specific camera based on the camera's pattern of noise (Lukas, Fridrich, and Goljan 2006; Lukáš, Fridrich, and Goljan 2006). This approach exploits the imperfections in a camera sensor that tend to be unique. Having measured these imperfections from a camera, they can be matched against the same imperfections extracted from an image of unknown origin.

Exposing Digital Forgeries: real vs. virtual

The child pornography charges filed against Police Chief David Harrison in 2006 shocked the small town of Wapakoneta, Ohio. At his trial, Harrison's lawyer argued that if the State could not prove that the seized images were real, then Harrison was within his rights in possessing the images. In 1996 the Child Pornography Prevention Act (CPPA) extended the existing federal criminal laws against child pornography to include certain types of "virtual porn." In 2002 the United States Supreme Court found that portions of the CPPA, being overly broad and restrictive, violated First Amendment rights. The Court ruled that "virtual" or "computer generated (CG)" images depicting a fictitious minor are constitutionally protected. In contrast, in the United Kingdom the possession or creation of such virtual images is illegal. The burden of proof in the Harrison case, and countless others, shifted to the State which had to prove that the images were real and not computer-generated.

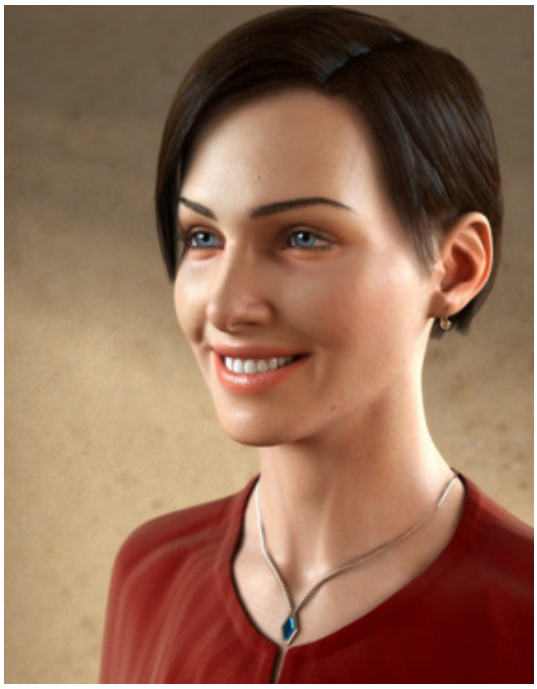

Figure 7. A computer generated (virtual) person (created by Mihai Anghelescu).

Given the sophistication of computer generated images, several state and federal rulings have further found that juries should not be asked to make the determination between real or virtual. And in 2006 at least one federal judge even questioned the ability of expert witnesses to make this determination. At the time, however, there were no data suggesting whether people could reliably make this distinction. To this end, my colleague Mary Bravo and I tested the ability of human observers to differentiate between CG and photographic images (Farid and Bravo 2007). We collected 180 high-quality CG images with human, man-made, or natural content, created over the previous six years. For each CG image, we found a photographic image that was matched as closely as possible in content. The 360 images were presented in random order to 10 observers from an introductory psychology subject pool. Observers were given unlimited time to classify each image. Observers correctly classified 83% of the photographic images and 82%

of the CG images, inspecting each image for an average of 2.4 seconds. Among the CG images, those depicting humans were classified with the highest accuracy rate: 93%. The observer with the longest inspection time (3.5 seconds/image) correctly classified 90% of all photographic images and 96% of all CG images. This observer correctly classified 95% of CG images depicting humans. It seems that, at least for now, even with the great advances in computer graphics technology, the human visual system is still very good at distinguishing between computer generated and photographic images.

As technology improves it is likely that it will become increasingly more difficult to distinguish the real from the virtual. To this end, we have developed a computer program that can distinguish between CG and photographic images (Lyu and Farid 2005). Because CG images are created using idealized lighting, surface geometries, optics, and sensors, they tend to exhibit statistical regularities different from photographic images. We have been able to quantify and measure these statistical differences, and use them to differentiate between photographic and CG images.

Photographs and Memories

Days before the 2004 U.S. presidential election, a voter was asked for whom he would vote. In reciting his reasons for why he would vote for George W. Bush, he mentioned that he could not get out of his mind the image of John Kerry and Jane Fonda at an anti-war rally. When reminded that the image was a fake, the voter responded "I know, but I can't get the image out of my head."

Several studies have shown that doctored photographs can implant and alter childhood and adult memories (Wade, Garry, Read, and Lindsay 2002; Garry and Wade 2005; Sacchi, Agnoli, and Loftus 2007). In a study by Wade and colleagues (Wade, Garry, Read, and Lindsay 2002), participants viewed doctored photographs of themselves and a family member taking a hot-air balloon ride, along with photographs of three real events from their childhood. After as few as three interviews 50% of participants reported remembering all or part of the hot-air balloon event. Similar results were reported in (Garry and Wade 2005), although the authors did find that images are not as powerful as narratives in stimulating false memories. Adult memories seem to be equally influenced by doctored images. In a study by Sacchi and colleagues (Sacchi, Agnoli, and Loftus 2007), participants were shown original and doctored photographs of memorable public events at which they were present (the 1989 Tiananmen Square protest in Beijing, and the 2003 protest in Rome against the Iraq war). The doctored images, showing either larger crowds or more violence, changed the way in which participants recalled the events. Images real or fake, have a very real and lasting impact.

Photographs and Trust

Schauer and Zeckhauser (this volume) explore the nature of paltering, which they define as "the widespread practice of fudging, twisting, shading, bending, stretching, slanting, exaggerating, distorting, whitewashing, and selective reporting." While some forms of photographic tampering certainly rise to the level of fraud – such as Hwang's fraudulent scientific paper – many forms of photographic doctoring might be classified as paltering – such as Hajj's addition of smoke to a war photograph. Both forms of deception, however, are equally damaging to our trust in photographs. Each form of deception creates uncertainty in a medium that is becoming increasingly more malleable, so that no matter how minor the palter, trust is eroded.

Perhaps this erosion of trust is inevitable in an increasingly digital age.

Hancock (this volume) explores how technology provides for new and sophisticated forms of deception in on-line personal interactions, such as on-line dating. In addition to under-reporting their weight, and over-reporting their height and income, some date-seekers have taken to posting digitally enhanced photographs. This erosion of trust will only increase with, for example, the next generation of cameras that automatically removes wrinkles (Panasonic) or 10 pounds (Hewlett-Packard) at the push of a button.

The mushroom cloud from the nuclear explosion over Nagasaki, a young girl fleeing from her village after being burned by napalm, prisoners being abused in Abu Ghraib prison in Iraq: these wartime photos have become ingrained in our collective memories and serve as powerful symbols of the horrors of war. Glenney (this volume) describes the potential for the use of doctored photographs in wartime to boost morale, demoralize or deceive the enemy, or justify military action. There is little doubt that these types of manipulations would raise serious ethical issues, and add to a continued degradation of trust in photography.

Digital technology is allowing us to manipulate, distort, and alter reality in ways that were simply impossible twenty years ago. And as the above examples illustrate, we are feeling the impact of this technology in nearly every corner of our lives. Tomorrow's technology will almost certainly allow us to manipulate digital media in ways that today seem unimaginable. As this technology continues to evolve it will become increasingly more important for us to understand its power, limits, and implications, and to possibly adopt a different relationship with digital media. It is my hope, nevertheless, that the science of digital forensics will keep pace with these advances, and in so doing return some trust to this wonderful, and at times puzzling, digital age.

Acknowledgments

Portions of this paper appeared in (Farid 2006a). This work was supported by a Guggenheim Fellowship, a gift from Adobe Systems, Inc., a gift from Microsoft, Inc., a grant from the United States Air Force (FA8750-06-C-0011), and by the Institute for Security Technology Studies at Dartmouth College under grant 2005-DD-BX-1091 from the Bureau of Justice Assistance and Award Number 2006-CS-001-000001 from the U.S. Department of Homeland Security. Points of view or opinions in this document are those of the author and do not represent the official position or policies of the U.S. Department of Justice, the U.S. Department of Homeland Security, or any other sponsor.

References

Farid, H. (2006a). Digital doctoring: How to tell the real from the fake. *Significance 3*(4), 162–166.

Farid, H. (2006b). Digital image ballistics from JPEG quantization. Technical Report TR2006-583, Department of Computer Science, Dartmouth College.

Farid, H. and M. Bravo (2007). Photorealistic rendering: How realistic is it? In *Vision Sciences*, Sarasota, FL.

Fridrich, J., D. Soukal, and J. Lukáš (2003). Detection of copy-move forgery in digital images. In *Proceedings of Digital Forensic Research Workshop*.

Garry, M. and K. Wade (2005). Actually, a picture is worth less than 45 words: Narratives produce more false memories than photographs. *Psychonomic Bulletin and Review 12*, 359–366.

Hwang, et al., W. (2004). Evidence of a pluripotent human embryonic stem cell line derived from a cloned blastocyst. *Science 303*(5664), 1669–1674.

Johnson, M. and H. Farid (2005). Exposing digital forgeries by detecting inconsistencies in lighting. In *ACM Multimedia and Security Workshop*, New York, NY.

Johnson, M. and H. Farid (2007). Exposing digital forgeries through specular highlights on the eye. In *Information Hiding*, Brittany, France.

Kennedy, D. (2006). Editorial retraction. *Science 211*(5759), 335.

Lukas, J., J. Fridrich, and M. Goljan (2006). Digital camera identification from sensor noise. *IEEE Transactions on Information Security and Forensics 1*(2), 205–214.

Lukáš, J., J. Fridrich, and M. Goljan (2006). Detecting digital image forgeries using sensor pattern noise. In *SPIE Electronic Imaging, Photonics West*.

Lyu, S. and H. Farid (2005). How realistic is photorealistic? *IEEE Transactions on Signal Processing 53*(2), 845–850.

Nishino, K. and S. Nayar (2004). The world in an eye. In *Computer Vision and Pattern Recognition*, pp. 444–451.

Pearson, H. (2005). Image manipulation: CSI: Cell biology. *Nature 434*, 952–953.

Popescu, A. and H. Farid (2004). Exposing digital forgeries by detecting duplicated image regions. Technical Report TR2004-515, Department of Computer Science, Dartmouth College.

Popescu, A. and H. Farid (2005). Exposing digital forgeries in color filter array interpolated images. *IEEE Transactions on Signal Processing 53*(10), 3948–3959.

Rossner, M. and K. Yamada (2004). What's in a picture? the temptation of image manipulation. *The Journal of Cell Biology 166*(1), 11–15.

Sacchi, D., F. Agnoli, and E. Loftus (2007). Doctored photos and memory for public events. *Applied Cognitive Psychology*, in press.

Wade, K., M. Garry, J. Read, and D. Lindsay (2002). A picture is worth a thousand lies. *Psychonomic Bulletin and Review 9*, 597–603.