Photo Fakery

Identifying falsified images can be straightforward if you know a few tricks.

By Robert Fiete

From oemagazine January 2005
31 January 2005, SPIE Newsroom

Like it or not, fake images are everywhere and have become a part of today's culture. Thanks to the popularity of digital cameras and the availability of desktop imaging software that allows users to easily manipulate images, fake images have become commonplace, especially on the Internet. We see many images that defy common sense and it is natural for us to question the authenticity of these images. Most of have seen images that are obvious fakes, such as the 80-foot grasshopper climbing the Empire State Building, but we naturally assume that these images are fake and know that they were created simply for our amusement. Unfortunately there are too many times when a fake image has been created but it is advertised as real, challenging us to decide for ourselves whether the image is real of fake.

A fake image can be defined as an image of an object or scene that wasn't captured as the image would imply. In general, fake images are created to generate a deception, but not all fake images are bad. The motivation may be simply for harmless entertainment, which accounts for most fake images generated today. Fake images can be generated for research and development purposes, e.g. to understand image quality issues with different camera designs. The fake images that concern us most are those that are created to perpetuate a lie. Some people will generate fake images for profit, such as a picture of an alien, a ghost, or an alien ghost of Elvis that they can then sell to a tabloid. Probably the most dangerous motive for generating fake images is to alter the public's perception of truth for political reasons. It would be nice a reliable method existed for determining if an image is real or fake, but unfortunately none exists. We can hope to catch most of the fake images, however, if we understand how fake image are made and what characteristics to look for.

Creating Fake Images

Although generating fake images historically originated with darkroom tricks, today almost every fake image is made using a computer. Even though it is getting more difficult to discern a real image from a fake image as image processing software improves, image analysis can still be used to detect traits that can expose many of them as fakes. To understand how fake images can be detected we must first understand how they can be made on computers. The two most common methods today for generating fake images are to "paint" a new image outright or to alter an existing image that has been captured by a camera.

A digital image is essentially a grid of numbers, where each number represents the brightness of each picture element, or pixel, in the image (see Figure 1). An 8-bit image can have $2^8=256$ gray-level values, with a value of 0 representing black and a value of 255 representing white. A color image is made by combining a red image, a green image, and a blue image. Adding together different gray-level values from the red, green, and blue image produces the various color values.

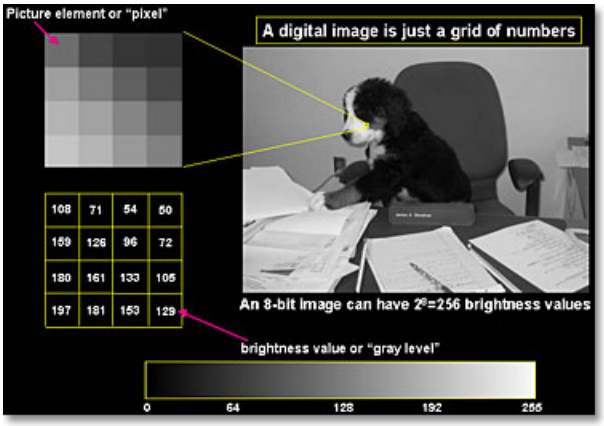

Figure 1: *A digital image is simply an array of numbers corresponding to brightness values.*

Since a digital image is simply a grid of numbers, it is conceivable for an artist to create a computer-generated image by "painting" a grid of numbers to represent any object or scene that could be captured with a digital camera. For a 24-bit color image composed of an 8-bit red, green, and blue image, there are almost 17 million possible colors for each pixel. A 4"x6" image at 300 dpi (dots per inch) will have over 2 million pixels, thus there are over 36 thousand billion numbers that can be considered to make the color digital image. Realistically all of the possible numbers do not need to be considered by the artist, but serious thought does need to be put into the values that will be used, especially when illumination and edges are considered. If the computer generated image is to look like a real photograph, then the image must be consistent with all of the laws of physics applicable to generating a real image.

Many of the classic painters, such as Leonardo Da Vinci, had an amazing talent to incorporate the proper shading, texture, tone, and color into their paintings that were consistent with the real world thus adding an amazing amount of realism to their work. However, their paintings do not look like modern photographs because they do not contain sufficient detail to match all of the physical properties associated with photographic imaging. (Actually, most artists probably would have been quite unhappy if their works of art looked like a modern photograph.)

In order to create a digital image that looks like a real photograph, the correct brightness values must be determined on a pixel-by-pixel basis to match the physical imaging properties, which could take months to years, depending on the image size, without the aid of computer software to perform the calculation. This problem was solved with the development of computer graphics software, designed to generate images of 3D objects with realistic illumination conditions. A rendering operation adds lighting, shading, colors, and texture to a mesh form of the object that is created by the artist. Ray tracing models produce the best quality by projecting many rays of light and modeling all of the physical qualities between the light and the objects in the scene, including reflection, refraction, transmission, scattering, absorption, and diffraction. The artist must simulate enough rays of light to cover every pixel in the image, which can be very time consuming if many rays of light are used. We have all seen the impressive results of computer animation in many feature films, creating dinosaurs or aliens that come to life on the screen. However, generating impressive detail in fake images using computer graphics, especially in a movie sequence, is still very difficult due to the complex calculations that need to be performed and the software is not accessible to the average PC user.

The most common method of generating a fake image, due to its simplicity, is to alter an existing image that was captured by a camera. The image can be altered by changing the *context* of the image, such as claiming that an actual image of a lampshade is actually an image of alien spaceship, or the image can be altered by changing the *content* of the image, such as superimposing an image of a cow onto an image of the moon (see Figure 2).

Figure 2: *Fake images generated by altering the context (left) or content (right) of an existing image.*

Figure 3: *The Surgeon's Photo, 1934, reportedly showing the Loch Ness Monster.*

Creating a fake image by altering the context of an image has historically been the preferred method for creating hoaxes because it requires no alterations and the image is an actual image captured by a camera; hence the image, and the film negative if it exists, will pass the scrutiny of scientific tests. A famous example of a faked image by altering the context is the "Surgeon's Photo" taken in 1934 by Robert Wilson who claimed it was a photograph of the Loch Ness Monster (see Figure 3). The image fooled many experts until an accomplice confessed in 1994 that the monster was nothing more than a toy submarine with the model of a serpent head attached.

Figure 4: *Many experts believed the Cottingley photographs of fairies were real.*

The Cottingley Hoax is another example of fake images created by altering the context (see Figure 4). In 1917 16-year-old Elsie Wright and her 10-year-old cousin Frances Griffiths took photographs of winged fairies near their home in England. Inspection of the images showed no alterations and Sir Arthur Conan Doyle, famous for creating the master sleuth Sherlock Holmes, deemed them authentic. Sixty years later the girls admitted that the fairies were paper cutouts held in place with hat pins.

Altering the content of an existing image most likely originated when early photographers were compelled to touch up the photographs of their paying customers to remove wrinkles and blemishes. Many people in the 19th century were accustomed to having flattering portraits painted of them and were not very tolerant at seeing the way they looked to the camera, which could not tell a lie. As dark room processes advanced, adding and removing people from images became a standard trick. When photographers were unable to get an entire family together for a family

portrait, they would set up the subjects such that the missing individuals could be added at a later time (see Figure 5). Altering images became routine for many political regimes in the 20th century, especially for propaganda. It was not uncommon for some governments to remove people from historic photographs when these people fell out of favor with the ruling party.

Figure 5: *My grandmother, second from the right in the back row, was added to this family portrait at a later time. (Note the unexplained shadow on her front.)*

Figure 6: *This fake image, created by altering the content of an image, took less than three minutes to make on a desktop PC.*

Today, altering the content of an image does not require dark room tricks but merely a PC with image editing software. Desktop software is readily available and easy to use, allowing anyone to quickly and creatively alter images. The easiest approach is to simply cut a section from one image and embed it into another image (see Figure 6). The desktop software allows the creator to modify the extracted image to the appropriate size and rotation. The software on the market today is so easy to use that that pre-school children have little difficulty creating impressive altered images.

Figure 7: *Our perception is the first line of defense at identifying fake images. The cat is obviously too big for this breed of cat and the man would need to lean backwards more to properly hold a cat of this weight.*

Figure 8: *Our perception can fail to detect a fake image if there is no cause for suspicion, such as when TV Guide used Ann-Margret's body for a picture of Oprah Winfrey.*

Figure 9: *A computer generated image from the movie Armageddon was circulated on the Internet as an actual image of the Columbia disaster taken from a satellite, even though it has a cartoon look to it.*

If an image is deemed suspicious, then we can first look for clues by visual inspection and then proceed with scientific inspection if necessary. The first line of defense for detecting a possible fake

image is our own perception. We have a keen ability to sense that something is wrong with an image and trusting our common sense works most of the time. If an image looks unbelievable, then it probably is unbelievable and is a fake (see Figure 7.). If an image looks real and similar images are easily obtained, then it probably is real since there would be no motive to warrant the time and effort to create the fake image. Unfortunately life isn't that simple. There are examples of fake images that we believe are real because they do not draw suspicion (see Figure 8) and there are examples of unbelievable images that are in fact real images. These real but unbelievable images are the ones that fascinate us but also make it harder to discount the images that we suspect are fake. Images that we believe to be real but are in fact fake are bothersome because they unfairly manipulate our sense of truth.

Using computer animation software to create a fake image works well in movies but generally does not fool the public when used to pass off a fake image as real. Our perception is very sensitive to subtle details in the composure and texture of objects in an image, especially when viewing images of people. Most computer-generated scenes, especially those involving people or animals, have a "cartoon look" about them when scrutinized. People generally look like mannequins and subtle details are missing. Images that circulated on the Internet claiming to be actual satellite images of the space shuttle Columbia exploding in space could easily be recognized as the work of computer animation when viewed closely (see Figure 9). The ability to generate realistic computer generated people is improving dramatically over time as software technology and mathematical models progress.

A fake image created by altering the context is the hardest to positively identify as fake since the image is real and will pass scientific tests on the validity of the image itself. Most fake UFO images cannot be immediately discounted as fake because they are indeed real photographs of objects that the viewer cannot properly identify, leaving the image subject to interpretation. The key to identifying a fake image when the context is altered is to identify aspects of the image that are inconsistent with the image description, i.e. catch the perpetrator in a lie. For example, the time and date claimed may be inconsistent with the sun's position or the known weather conditions for that date.

Figure 10: *1932 photograph from "Death in the Air: The War Diary and Photographs of a Flying Corps Pilot" was later discovered to be fake.*

Photographs published in 1932 reportedly showing scenes from WWI dogfights were amazing due to their sharpness and clarity (see Figure 10). But the amazing clarity was a clue that the images were probably fake because they appeared too sharp given the relatively long exposures required from cameras at the time and the amount of motion and vibration on the airplane. The images were not proven to be fakes until 1984 when the model airplanes used in the images were discovered.

When the image content has been altered, we focus on the aspect of the image that makes the image unbelievable. Images that have had their content altered will usually have physical

inconsistencies in the image that may be apparent under visual inspection. Unfortunately, these inconsistencies are not always apparent in the image and the image may not be proven to be fake until the original unaltered image is discovered (see Figure 11).

Figure 11: *Finding the original unaltered image is the best proof that image content has been altered.*

Understanding the image formation properties of a camera can help us to recognize fake images. The step-by-step physical process of forming an image is called the imaging chain and every image must adhere to the physics of an imaging chain, from the radiometric source to the final image product (see Figure 12). Image chain analysis methods are used to examine images for evidence that the laws of physics have been broken. Any inconsistencies found within the image can be an indication the image has been altered.

Figure 12: *A proper understanding of the image chain elements can expose inconsistencies in a fake image.*

The physical traits of the image that can be assessed include the illumination conditions, edge sharpness, resolution, tone, relative scale, and noise characteristics. Many of the computer animated scenes created for movies and electronic games do not adhere to the laws of physics, but this is usually intentional to save cost and to make the scenes more entertaining.

A common inconsistency found when the image content is altered is the mismatch of radiometric or illumination conditions between the altered part and the rest of the image. The altered part of the image may have shadowing that is not consistent; indicating that is was illuminated under different conditions from rest of the image (see Figure 13). This is commonly seen when an object captured at one time of the day is added to an image that was captured at a different time of the day. Also, the light illuminating the altered part may not be consistent with the diffuse or specular light illuminating the rest of the scene. This effect is commonly seen when an object captured with a photographic flash is added to an image that was acquired with outdoor or studio lighting. Color, contrast, and tone will also vary for different illumination conditions, thus creating a mismatch of these characteristics between different images.

Figure 13: *This image shows inconsistent illumination conditions. The shark is illuminated from the above left, whereas the rest of the image is illuminated from above and slightly behind.*

Figure 14: *Fake satellite image of the Northeast blackout in 2003*.

An image claiming to be a satellite image of the Northeast blackout in 2003 circulated on the Internet shortly after the blackout occurred (see Figure 14). The image was quickly identified as a fake because the blackout area is pure black compared to the other areas with no light sources. Other clues to this deception include the false satellite name, the unlikely lack of clouds anywhere over North and Central America, and the fact that the blackout was not total over the Northeast. The original image is a composite of many DMSP satellite images acquired between 1994 and 1995. Actual NOAA satellite images of the blackout area are shown in Figure 15

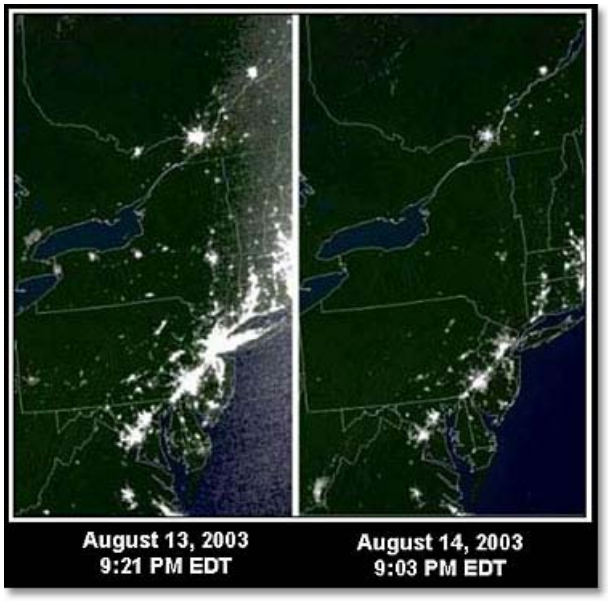

Figure 15: *Actual NOAA satellite images of the Northeast blackout.*

One must be very careful when analyzing the illumination characteristics of the scene. The shadows and illumination conditions can be misleading, especially if the three-dimensional aspects of the scene are not taken into account. The Apollo 11 moon landing images appear to contain "anomalies" that some people use to argue that the moon landing was staged in a studio. These "anomalies" include shadows on the lunar surface that are not parallel (see Figure 16) and objects that appear illuminated even though they are in the shadows, both suggesting that there were light sources other than the sun, as well as the lack of stars in the black sky, suggesting that a black back-drop was used on a studio set. Of course, all of these so-called anomalies are exactly what we expect to see in the images if we truly understand the imaging conditions on the lunar surface. The shadows are not parallel as seen in the images because the lunar surface is not flat and the objects are not necessarily parallel to one another in height, the shadows are illuminated from the light scattering off of the lunar surface, and the stars do not appear in the images because the camera exposure was set for the brightness of the lunar surface.

Figure 16: *The shadows appear to be inconsistent in this image until the terrain variations of the lunar surface are considered.*

Creators of fake images usually ignore the known physical properties of creating an image with a camera. The most significant camera effects are edge sharpness, influenced by the lens diffraction, focus, and motion blur; perspective geometry; and noise properties, usually from the detector and compression. Computer animated images are usually created without any camera effects since this will degrade the image quality and make the images less appealing to the audience. This, however, results in images that are physically impossible to capture with a camera in the real world.

When an object is added or deleted from an image, an edge is usually created that has a sharpness that is inconsistent with the rest of the image. Even an in-focus image will exhibit some blurring due to the diffraction of light from the camera aperture. The behavior of the blurring in the image is well understood and can be mathematically modeled if the camera design is known. Even if the camera design is not known, measurements within the image can produce a relatively accurate mathematical model of the camera that can provide reasonable predictions. Cutting an object from one image and inserting it into another image will create a sharp edge at the boundary of the inserted object that is sharper than physically possible. This sharpness is easily seen and creates an obvious sign that the image has been altered, so smudging tools in image processing software are usually used to reduce the visibility of these edges (see Figure 17). This smudging, however, will usually produce blurred edges around the object that are inconsistent with the rest of the image.

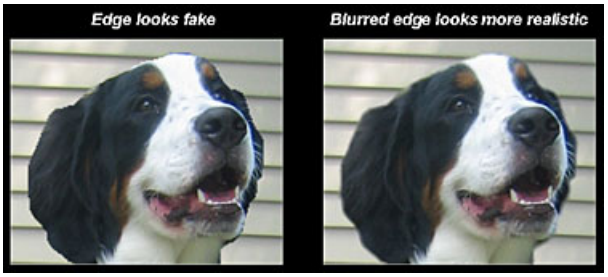

Figure 17: *The edges around the inserted object are noticeably sharper than the rest of the image and leave no doubt that the image has been altered. Blurring the edges reduces the visibility of the alteration, but the sharpness of the edges is still inconsistent with the rest of the image.*

All objects in an image must also contain the proper perspective and geometry. The perspective of three-dimensional objects in the two-dimensional image is dictated by the viewing geometry and the camera. If the geometry of an object is inconsistent with the other objects in the image, then it was probably added from another image. For example, lines that are parallel in the scene will converge to a point called the vanishing point in the image. If the parallel lines of an object do not converge to the same vanishing point as the rest of the image, then the object could not have been imaged with the same camera or perspective as the rest of the image (see Figure 18).

Figure 18: *Measuring the vanishing points reveals that a window has been added to this building.*

Most images will exhibit some amount of noise, primarily from the detector or from the image compression that was applied. The noise characteristics of an altered portion of an image can be inconsistent with the rest of the image. Magnifying digital images will generally exhibit graininess due to the detector noise and artifacts from the compression algorithm, depending on the level of compression. When images have been altered, the creator usually blurs the edges or other portions of the image to blend in the object, but this changes the noise characteristics, allowing the alteration to be detected (see Figure 19).

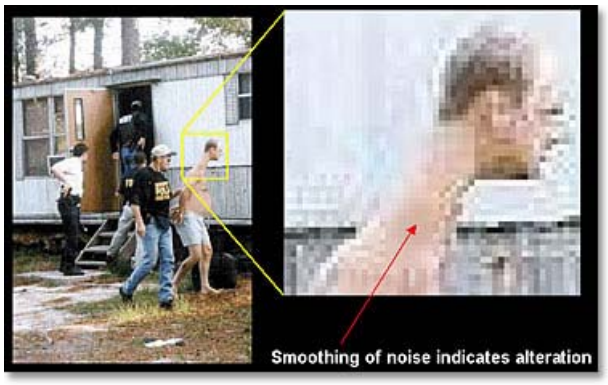

Figure 19: *Inconsistent noise properties may be apparent in altered images.*

Finally, an understanding of how image processing alters the image characteristics can lead to signs of alteration. For example, when the image contrast is enhanced, the resulting gray-level histogram of the image will usually display "holes" or gray-level values that contain are no longer present in the image. An object from one image that is inserted into a second image may exhibit a different histogram that will indicate that it was not originally part of the second image. However, if an image has been enhanced using an adaptive processing algorithm, then the image characteristics, such as the gray-level histogram or the edge sharpness, can change locally even though no other alteration have been made. Adaptive processing should not be used on real images if the integrity of the image is to be preserved. Unfortunately, if the image is processed after the alteration has been made, such as compressing the image, then the holes in the histogram may be filled in and the histogram will no longer look suspicious (see Figure 20).

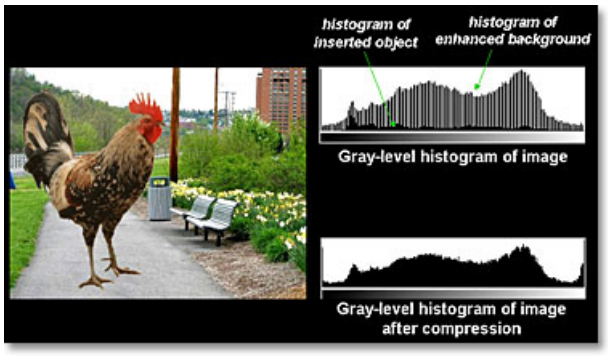

Figure 20: *The gray-level histogram may show signs that the image has been altered if it has not been processed after the alteration.*

The Difficulty of Detecting Fake Images

Most of the people generating fake images know little or nothing about the physics of the image chain, yet lots of fake images fool us because they seem to have properties that are consistent with real images. How is this possible? Images with altered context are actual images; hence image analysis will not show that the image itself is inconsistent with physics, only that the perpetrator is being untruthful. Images with altered content will usually show signs of alteration if the image is created quickly and carelessly. The anomalies created in an altered image can be reduced by having an understanding of the imaging chain properties and taking the time and effort to ensure that the entire image is consistent at the pixel level, but this is rarely performed due to the knowledge and time required.

The simplest method to reduce the detection of the anomalies in an altered image is to degrade the quality of the image of the alteration. The most common methods are blurring the edges, adding random noise, reducing the size of the image, or compressing the image, all of which will cover up

telltale signs of the manipulation. Many fake images have such poor quality that accurate measurements cannot be made to determine if inconsistencies exist. Admittedly, most creators of fake images do not reduce the quality with the intent of making image analysis more difficult, but instead reduce the quality by resizing and compressing the image simply to reduce the file size. However, reducing the image quality to hide the inconsistencies may reduce the impact that the creator of the altered image had hoped for. For example, inconsistent edge blurring can be reduced in altered images if the image is sub-sampled to a smaller size, but this could lead to unsatisfactory aliasing artifacts.

Image steganography offers a method for embedding hidden information into an image. Information pertaining to the unaltered image can be encoded and embedded into the image such that it is not visible. The information can also be encrypted, requiring a key to decode the embedded information so that unauthorized users cannot alter the information. The embedded information can withstand most alterations and processing such as scale, rotation, compression, and cropping. As an example, an edge map of the image can be created, encoded, and embedded into the image itself. If an image is suspected of being altered, then the embedded information can be extracted using the key and compared to the image. Any differences between the edge map of the current image and the edge map embedded in the image can prove if the image was altered (see Figure 21).

Figure 21: *An image has an encoded watermark that contains the edge information of original image, revealing an alteration that has been made to the original image.*

Although image analysis tools can help detect many fake images, currently there is no way to stop somebody from spending the time and resources to make a fake image that is not detectible. All one can do is hope that an inconsistency can be found, thus indicating that the image is fake. Methods currently being developed, such as image stenography and control coding in printers, can aid in the prevention and detection of altered images that are passed off as real images. Two great references for checking the authenticity of images being distributed on the Internet are *The Museum of Hoaxes* and *Urban Legends Reference Pages*. For further reading on fake images, a good reference is *Photo Fakery,* by Dino Brugioni.

Robert Fiete

Robert Fiete is chief technologist at the space systems division of ITT Industries, Rochester, NY.