Makeit

Spotting faked photos

Manipulating photos is easy – spotting the fakery much less so

Like it or not, photographic manipulation is now easier and the results better than ever before. Usually, the alterations are fairly benign – removing red-eye or fixing the light balance, for example – but sometimes the change is designed to affect our emotions and our understanding. Given that we're bombarded with images throughout the day, how can we teach ourselves to recognise the fakery when we come across it?

Faked images have been with us since the start of the photographic era. Who can forget the five photos of the 'Cottingley Fairies' (**www.bit.ly/hfBWaX**) from 1917 that duped the author of Sherlock Holmes, Sir Arthur Conan Doyle? With today's cynical eye, the photos are amateurish – laughable, even – but they did fool a lot of people for some time. Sceptic James Randi came up with the proof in 1978: the fairies were traced from drawings in Princess Mary's Gift Book (a publication from 1917) and mounted on cardboard.

Back in the pre-digital days, modifying photographs by removing existing artifacts or adding elements that weren't there was hard work. There were therefore more photos where the fakery was due to a staged subject (Robert Wilson's infamous 'Surgeon's Photograph' of the Loch Ness monster (**www.bit.ly/fy7dQV**), or Robert Doisneau's famous 'Le Baiser de l'Hôtel de Ville' (**www.bit.ly/hCKyPR**) rather than a deliberately altered original (like Yevgeny Khaldei's 'Raising a Flag over the Reichstag' (**www.bit.ly/ghFaQQ**), where the smoke was enhanced and a (presumably looted) watch on a soldier's wrist was edited out of the picture).

Digital manipulation

However, as soon as programs like Adobe Photoshop arrived on the scene and digital cameras became affordable enough for everyday use, the propensity for digital manipulation of photographs took off. Alongside this explosion of digital photography, we saw the creation of a specialised field called digital forensics, which attempts to uncover instances of deliberate falsification in digital images. It's important to remember that photos are often altered for simple aesthetic reasons. When taking a photo of a group with a flash, it's pretty certain that some of the people will exhibit what's known as 'red-eye', where the light from the flash is reflected off the retina and seen in the image as a red spot on the pupil. This is simple to fix (change the colour from red to black), and many photo-editing programs include this functionality as a matter of course. Many modern cameras are also able to manipulate images to remove red-eye, alongside their capability for face-detection.

Falsifying skill

There's also the ability to modify the white balance of the photo, or to adjust the brightness and contrast to create a more pleasing image. Again, it's not really falsifying the photograph (more falsifying the photographer's skill), and even in the pre-digital era, photographers would do the same kind of processing in the dark room.

Another popular visual fix for photos is to crop the image to remove any jarring or distracting elements along the edges. That way the viewer's attention is concentrated on what the photographer wants to say with the picture rather than extraneous detail. This is hardly a problem unless, for example, the picture is of a house for sale and the cropping removes the derelict car in the neighbour's drive.

The clone brush

Of course, all these changes are fairly benign in the grand scheme of things, and I'm sure we've all been guilty of them to a greater or lesser extent. How about this next one: using the clone brush in your favourite photo editor to remove that intrusive telegraph pole and wires from an otherwise perfect landscape? Yes, guilty. If the cloning has been done carefully, this kind of change is hard to detect visually (unless you know the original view, for example, or can find similar photos taken from the same vantage point online), and can be very hard to detect without some kind of software. We'll take a look a little later on to see how this kind of change could be spotted.

Professor Hany Farid, of the Department of Computer Science in Dartmouth College in New Hampshire, has been researching how we can visually spot changes in photographs. Although we as humans are extremely good at face recognition, quickly understanding a ▸

▲ **Figure 1: A JPG photo saved directly from a RAW photo from a Canon Rebel XTi.**

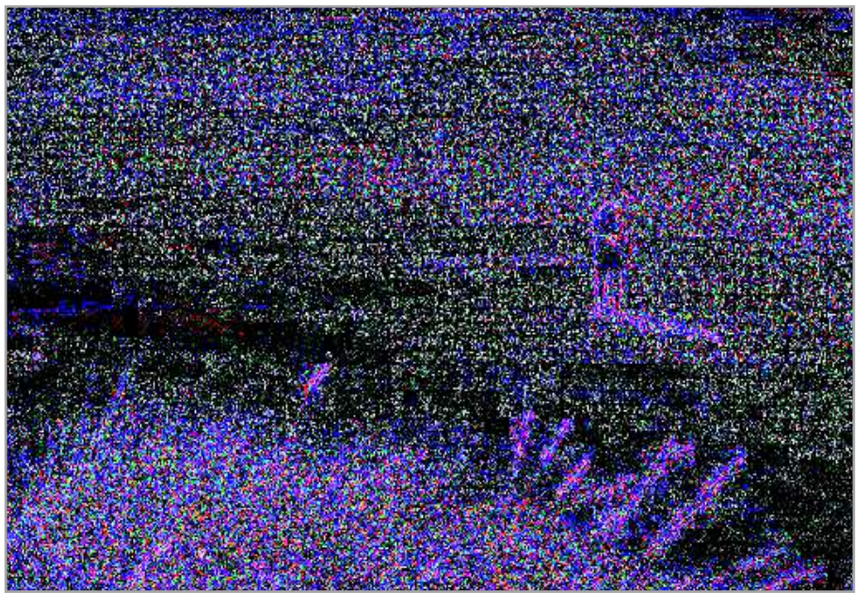

▲ **Figure 2: The Image Error Analysis image from the photo in Figure 1. Note the even, random noise.**

Using highlights

Another technique for spotting doctored photos is to look at highlights. Are the highlights on surfaces all pointing' in the same direction as the assumed light source? By this, I mean that the perpendicular line from the highlighted surface should be pointing to the light source. If these perpendiculars are pointing in different directions, something's not right. Note however that this analysis is very difficult to do: sometimes it is not obvious at all where the light is coming from (a cloudy day, for example), and trying to ascertain the perpendiculars can be hard.

If the photo is of a group of people you can use the highlights or hot spots on their eyes: if you can zoom in, these highlights should all be in the same direction. Note that staged photos may have more than one light source, in which case everyone's eyes should show the same number of highlights each from the same direction. ■

► scene, and determining direction and speed of motion, it turns out we're worse at spotting problems with lighting and shadows, and distortions caused by perspective.

Lighting and perspective

This is somewhat unfortunate, since the first step in judging the authenticity of a photograph is to check shadows, lighting and perspective. Since our visual acuity is easily fooled, we have to pay attention. Take, for example, the rather disturbing 'Accidental Tourist' photo that made the email rounds just after 9/11 (**www.bit.ly/i3J8ng**). It purportedly shows a man standing on the observation deck just before a plane hit one of the World Trade Center towers, but if you can avoid the initial sense of horror this creepy photo gives you and look at it more carefully, you'll start to notice incongruities in the various shadows. For example, both the plane and the man are 'nose-on' in perspective. The man has some very sharp shadows on his face – dark enough that it's hard to determine any details about his right side – yet the plane looks as if the sun is directly overhead: there are virtually no shadows on the upper surface. So, our initial revulsion means that we don't pay attention to the clues that indicate fakery in the image.

Professor Farid points out that there is a simple technique for detecting incongruities with regard to shadows. Because light travels in a straight line, you can easily draw a straight line between a point on an object being lit and the same point on its shadow. Extend the line in the direction of the light source. All such extended lines in an image should meet at some point, or in the case of the sun, should all be parallel to each other. The problem is finding some obvious point of the object and the same point on its shadow.

Follow the light

Another technique that can help you identify manipulated photos is to look at highlights. Are the highlights on surfaces all 'pointing' in the same direction as the assumed light source? A line perpendicular to any highlight's surface should point towards the light source. If you notice that these perpendiculars are pointing in different directions, something's not right. Of great use here are the highlights in people's eyes: if you zoom in, these highlights should all be in the same direction.

Spotting the fakes: Stourhead

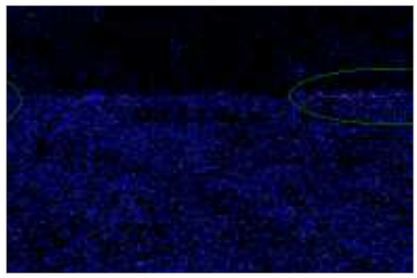

1 Inspect the photo: This is lovely image showing a view of the changing autumn leaves at Stourhead in Wiltshire. My first thought is to see whether the reflections in the lake show something that's not in the rest of the photo. Has the editor cloned out some people, or removed a car from the photo, perhaps? Note that even though it's autumn, there are no leaves visible on the grass in the foreground, but there are some elsewhere.

2 Zoom in, if possible: If you have the image available in a large enough resolution, you can zoom in on areas of interest identified during the visual inspection. Here, for example, I can see that there's a strange-looking blue reflection in the lake by the ruin that's not matched by anything at the ruin. Also, coincidentally in the same area, there are some obvious masked-out splotches that might have been ducks on the water.

3 Use the JPG image error analysis tool: Make sure the photo is online, then go to **http://errorlevelanalysis.com**. I cropped part of the foreground to the right (the tool only accepts small photos) and fed it in to give the difference image above. The randomness isn't even – the grassy area has darker splotches (possible evidence of a cloning brush). Note the hot spots along the dividing line between water and land: something's been changed here too.

When you're looking for altered photos, also keep an eye out for obvious repetitions in the image. One of the most famous examples of this (mainly because it got past the photo desk at Reuters with no problem before being called out by the world at large) is the photo by Adnan Hajj of smoke over Beirut after an Israeli bombing raid (**www.bit.ly/e2cITN**). The smoke from two separate plumes shows distinct patterns of duplication, probably through use of a cloning tool; it's so badly done, it looks obviously fake. Professor Farid is trying to extend this type of visual check into a software tool that tries to spot cloned cells (that is, blocks of copied pixels) in a photo.

Understanding formats

The images we see online are usually in JPG format. This format is extremely popular, mainly because of its small file sizes and near-universality. Many point-and-shoot cameras only record their photos in JPG format, but larger more expensive DSLRs have an option to save images in either RAW or JPG format (or both at the same time).

The most important point to understand about JPGs is that they're compressed to create the smallest possible file sizes. Not only that, but the compression used is a lossy format; in other words, the image we see in a JPG file is not exactly the same as the original photo, because some information has been discarded. If you like, information is averaged out in a strict mathematical way to make the image data more compressible.

What this means is that, although a JPG file is a very close representation of the original image, it's not quite the same. If a JPG is then saved again, the errors caused by the compression format are multiplied and become even more visible. A third-generation JPG is even worse. This is why the recommendation when you're editing a photo is to only work on and change the original image (and use 'Save as', of course). Never work on a second or third-generation JPG – the smearing will become obvious quite quickly.

Identifying edited JPGs

How does this behaviour help when you're trying to identify edited photos? Imagine we have a JPG file, and we re-save it using our photo editing software to give two versions: the first generation and the second. We make no other changes – it's the same image at the same size. The two images will look very similar to the naked eye. Now we 'subtract' the second image from the first.

The second compression will produce a set of data that's very close, but not exactly the same as the original. This means, in effect, that the subtraction of the pixel values will produce numbers close to, but not exactly zero. Since zero is black, we get a darkish image with random 'noise', and that randomness will be spread evenly across the difference image. Areas that encapsulate edges will be coloured differently, but still will show even, random noise. Figure 1 shows an original image (it was converted to JPG format from RAW) and Figure 2 the shows difference image created by comparing it with a second-generation copy. As you can see, it's just random noise, evenly spread throughout the image.

What researcher Dr Neal Krawetz discovered is that, if the image is altered, the altered areas become much brighter in the difference image – there will be visible 'hot spots'. This happens because the alteration leaves artifacts in the pixel data, which are magnified by the compression process.

Consider the workflow: you have a JPG, you make changes to it and save it (it is therefore a second-generation image). Krawetz's process then creates a third-generation image in order to calculate the difference. The changes show up more.

At least that's the theory. In my own experiments (for example, cloning out the person in Figure 1), it's still hard to discern changes, Projects Editor Alex Cox asked me to try out this theory on some 'after' images he provided (see 'Spotting the fakes', below).

All in all, forensic image analysis is in its infancy, with much work needed to better what sharp-eyed viewers can do. And it's an arms race: with Adobe Photoshop's automated tools for touching up photos, it's becoming harder and harder to spot the fakes. **PCP**

Julian M Bucknall has worked for companies ranging from TurboPower to Microsoft and is now CTO for Developer Express.
feedback@pcplus.co.uk

Lossy compression

JPG compression is lossy, which means that although a JPG file is a faithful rendition of an original photo, it is not the same as the original. Firstly, the colour spaces are different (images on your screen are in the RGB space, JPG uses the YCbCr space) and the conversion is not exact. The colour channels in the converted image can be downsampled from 256 levels to 128 (in essence, we see less detail in the colour channels than the Y luminance channel). The image is then parcelled up into 8 x 8 pixel squares, and a mathematical conversion similar to a Fourier series is applied to each block to approximate the pixel data.

This entire conversion is carried out to reduce the amount of information in the image, in order to make it more compressible. The particular process flow is designed so that, when the data is decompressed at the end, the resulting image looks the same as the original to the naked eye, to a high degree of fidelity. ■

Spotting the fakes: Pulteney Bridge

1 Inspect the photo: This is a shot of Pulteney Bridge in Bath. This time the reflections in the water aren't as valuable: the water is too turbulent. The most obvious edit here would be to remove some people, but that doesn't seem to be the case. Perhaps the changes (if any) are more subtle. Since it's a well-known landmark, I pulled up some images from Flickr to cross-check – has the editor removed or added anything architectural?

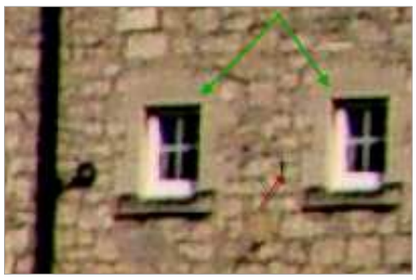

2 Zoom in: Having compared this image with another shot taken from the same angle, it seems that the building on the right has gained an extra window. Notice that when zoomed in, the added window (on the right) looks exactly like the real window (on the left). Notice another thing – the editor has copied just a bit too much of the wall: there is a speck of black to the left of the cloned window that's been duplicated from the drainpipe.

3 Zoom in again: My find on Flickr shows me there's been another, more subtle change. Just above the bridge roof on the right, in the distance, there should be a church tower. It's missing from the doctored photo, and at high magnification you can just see the edit. It must be noted, of course, that I shouldn't trust the Flickr photo either – it might have been edited too – but should find another image taken by someone else to double-check. ■