



*Permanently IR-modified Canon Rebel XT. Stitched panorama,
260Mpixel*

Digital Infrared Photography Part One

Marco Annaratone

Email: RLph3oto@res3onantlink.com

Claudio Ruscello

Email: IR3vision@li3bero.it

To email to one of the authors please remove first both digits 'three' from either email address above.

LIGHT AND PHOTOGRAPHY

Any radiation can be characterized by its *wavelength*, and light is of course no exception. What we consider as "light" is in fact *visible light*, i.e., that part of the radiation spectrum that the eye-brain system can indeed capture. We see various color hues, going from violet through blue, and then green, then yellow, then orange, and finally down to red and deep red. A wavelength in this area of the spectrum is typically measured in nanometers, where one nanometer is equal to one billionth of a meter. The wavelength of violet is about 400nm; going down to deep red the wavelength increases, from about 460nm (blue) to 540nm (green) and then yellow (600nm) and finally deep red (750nm). Beyond violet and below deep red there are radiations the human eye cannot see, i.e., *ultraviolet (UV)* and *infrared (IR)*. Although the human eye cannot see these radiations, the human *body* can certainly feel and react to them. Infrared is often perceived as heat and we know how harmful certain UV rays can be to the skin.

In what follows we will focus on infrared photography, a technique that can capture on film or on a digital sensor the IR radiation reflected by the scene we frame. IR radiation does start below a wavelength of 750nm¹ and continues down to 20,000nm or more. However, both film and sensor are seriously limited in their capability to record IR radiation. Digital sensors can go as far as 1300nm. Commercial IR films are unable to record radiations below about 900nm.

What we call "IR photography" then is the technique to capture IR radiation in the limited range between 750nm and 1300nm (in the case of digital sensors) and even less in the case of IR film.

IR photography is certainly not a new concept. On the contrary, it has been around for a long time: the 1935 "Leica Manual" by Morgan and Lester has a section dedicated to IR photography. IR photography was already carried out in the nineteenth century, but its appeal

¹ Note that all these numbers are approximations. It does not mean that, say, yellow stops at 600nm and that a wavelength of 601nm is 'another color.' The transition is not so abrupt! Some conventional threshold (albeit approximate) needs to be agreed upon to simplify the discussion.

to the general public grew significantly after 1931, when technology made it possible to shoot IR pictures using easy to handle infrared sensitive plates².

The reason why IR photography is undergoing a small Renaissance nowadays is that digital sensors are quite sensitive to IR radiation. This opens up, as we shall see below, interesting opportunities in utilizing again this old tool.

WHY INFRARED (IR) PHOTOGRAPHY?

IR photography opens up a new visual dimension for the photographer, a somewhat 'different' way of looking at the world around us. The spectrum of light is much wider than what the human eye can capture and, until CCD sensors became affordable to the general public, the only way to capture infrared radiation was to use special film, indeed sensitive to this part of the spectrum. Many are the applications of IR photography, from criminology to photomicrography and celestial photography. We shall focus here on landscape photography, but the reader interested in experimenting in other fields should be well aware that what is discussed here is the classical tip of the iceberg.

One of the fascinating features of IR photography is its ability to penetrate haze and light fog. As we have discussed above, infrared radiation has a longer wavelength than visible light and can 'go through haze' more easily. This is becoming, unfortunately, more and more important as the level of pollution in the air increases. Some even go as far as theorizing that 'moderate' monochrome IR photography³ may become the de facto standard in landscape black and white photography, as finding truly crisp and clear days is getting more and more difficult.



*Permanently IR-modified Canon Rebel XT. False color mode.
Stitched panorama, 490 Mpixel*

At any rate, if we use black and white (BN) IR film we obtain a black and white negative, while if we do the same (using methods we will see below) with a digital camera we will obtain an image in 'false colors' that can be then manipulated and modified as we wish (see above). When film was still the only game in town one could also buy 'false color' IR film. Nowadays, color IR film has all but disappeared.

In the past, IR pictures were often hand-painted to deliver "quasi natural colors". Here is an excerpt from "Leica Manual" by Morgan and Lester (1943): "A very interesting application of

² Handbook of Photography by Henney and Dudley, Whittlesey House, 1939.

³ By 'moderate' we mean a way to do IR photography without the strong "IR fingerprint." This can be accomplished in the digital domain by reducing or taming the "IR look" in post-production.


infra-red to landscape photography is to enlarge the photograph and tone the enlargement blue. If properly composed and toned the photograph will then show white clouds against a deep blue sky, white trees and grass, and various gray tones for buildings and pavements. The addition of oil coloring to the trees and grass and other parts of the picture will produce a surprisingly good imitation of a natural color photograph."

In order to really appreciate the beauty of IR photography, we need to limit ourselves to the part of the IR spectrum that the sensor (or film) can capture, i.e., we have to block visible light. This is accomplished by using an appropriate filter in front the lens. This filter will let the infrared portion of the spectrum go through while blocking at the same time the visible light part of the spectrum. This filter looks like any other filter used in black and white photography with one significant difference: it looks completely *black*, i.e., totally opaque, to the human eye!⁴

All the most important companies producing photographic filters offer at least one IR filter. The availability of IR filters may be in some geographies somewhat difficult though, especially if the diameter of the filter is not a common one (67mm or 72mm are relatively common diameters for filters, for instance). One may want to consider (also to contain costs) to buy one large-diameter IR filter (say, 77mm) and then use step-up rings to put it on lenses of different diameter. An IR filter that seems to be easier to find than others is the Hoya R72. This filter is a high-pass filter, i.e., it blocks all wavelength that are *below* a certain 'cut-off' wavelength . In the case of the Hoya R72 this wavelength is about 700nm (nanometers, a nanometer is one billionth of a meter). All the wavelengths below 700nm (i.e., all the visible light, ultraviolet light, and so on) will be blocked by the filter and will not go through it. Going back to the wavelength numbers presented in the introduction the reader will immediately notice that this filter lets also some visible light (deep red) go through. In fact, the filter does not look totally opaque to the eye but rather very dark red.

In the next page we show the same identical scene taken in three different ways, i.e., in color, in black and white visible light, and in black and white infrared radiation. These shots have been taken with a stock Nikon E5000 digital cameras and for the IR picture a Hoya filter has been used in front of the non-interchangeable lens.

It is quite easy to notice that, besides the obvious 'whitening' of the leaves (a classical effect of infrared photography due to the strong IR reflective power that the chlorophyll contained in plants has), the infrared shot suffers from almost no haze at all (see the mountain range on the extreme left of the picture) when compared to either the color picture or the visible light BN one. Note also the dramatic effect that capturing IR radiation has on the sky.



⁴ Referring to the wavelength numbers presented above we will block all the radiation below 750nm and let that above 750nm go through.



Color picture in visible light



Black and white picture in visible light



Black and white IR picture with Hoya R72 filter

ANALOG INFRARED PHOTOGRAPHY

The only way to do 'analog' IR photography is by using a film such as Kodak HIE Infrared or MACO IR820C (now Rollei 820c). These films tend to be quite sensitive to stray light, so it is better to load a 35mm film into the camera in the dark. This is instead a must and not just a recommendation with 120 film, where a changing bag is certainly a necessary accessory to carry around. Another thing to check is whether the camera has a *mechanical* frame counter; some cameras have this feature implemented via infrared light hitting an infrared sensor and this may in fact slightly fog the film around the area where the IR ray hits the film.



IR film: Maco 820c, now Rollei IR820c

One of the authors of this paper has been using the Maco 820c for years with a great deal of satisfaction, in spite of the constraints imposed by this film (that are quite typical of IR film and not at all unique to the Maco 820c, though). The need for long exposures is one of them: while the film is rated at around 200 ISO, the presence of the dark filter in front of the lens results in exposure times between 1/2s and 2s even in bright sunny days for reasonable f/stops (e.g., f/8)⁵. The tripod is therefore a must. Excellent results have been obtained from this film with abundant pre-wash to eliminate the anti-halation coating and by developing it in Kodak XTOL 1:2 for about 13 minutes at 20°C.

This film maintains its sensitivity, as the name suggests, up to around 820nm. Beyond 820nm its sensitivity drops rather abruptly. It is therefore a film sensitive to dark red and just the onset of infrared. For this reason the film should *not* be used in conjunction with a 'true' IR filter, but rather with 'very dark red' filters. Using it with IR filters results in a completely unexposed frame. While IR filters tend to block wavelengths below about 900nm, the Maco 820c is sensitive *up to* about 820nm. There is no intersection between the film sensitivity and the filter cut-off wavelength: where the film is sensitive the filter blocks the radiation from reaching the film; where the filter lets the radiation go through, the film is no longer sensitive. The result is therefore a completely unexposed piece of film!

An excellent match with the Maco 820c film is instead the 89B (092) dark red B&W filter. With this combination one can obtain interesting images like the ones in the next page.

⁵ In the past, film speed could be increased by hypersensitizing the film by mercury vapor treatment. This could deliver at most a 100% improvement in film speed, though. And we are now well aware of how terribly dangerous the exposure to mercury vapors can be.




Muir Woods, a redwood forest north of San Francisco. Maco 820c and Horseman SW612 with Rodenstock 65mm APO-Grandagon with B&W 89B filter. 1/2s exposure.



The Golden Gate from Lincoln Park. Maco820c and Fuji GSW690III, B&W 89B filter, 1s exposure. Note the excellent smoothness of the image with no trace of film grain in both pictures.

DIGITAL INFRARED PHOTOGRAPHY

The digital sensor in most cameras is quite sensitive to infrared radiation, and in fact manufacturers put right in front of it a special filter called "Hot Mirror" or "IR Cut-off Filter" to block radiation outside the visible range from reaching the sensor⁶. This in fact may create a loss of sharpness and also possibly inaccurate exposures. This filter blocks every radiation that falls below the deep red but does not eliminate it completely, and therefore we can use this at our advantage because in spite of the presence of the hot mirror there is still some IR radiation left that the sensor can capture. We will see below that this often comes with some serious disadvantages, however.




Characteristic response of a standard IR cut-off filter: nothing goes through it between about 700nm and 1100nm.



One of the filters used for the tests: the Hoya R72

⁶ We shall call this filter "hot mirror" from now on.



In order to eliminate visible light and let only IR radiation go through, one shall use an IR filter in front of the lens. This was true in analog photography and is true in digital photography as well. The Hoya R72 shown above blocks radiation below 720nm from going through and hence reaching the (sensor+Hot Mirror) component, as shown in the graph above.

Excellent filters are also the B&W (a bit more difficult to find in some geographies) and in particular the 87C (093). This filter lets only 1% of the radiation up to 800nm go through, while already at 900nm about 88% of the same goes through. It is a filter with a sharp transition therefore, somewhat more specialized than the Hoya filter just discussed. For this reason this B&W filter is to be avoided, as we said before, when using the Maco 820c film, and the B&W 89B (092) instead is the one to use.

TAKING THE PICTURE: FRAMING PROBLEMS

When we shoot IR with a TTL camera (be it analog or digital) the problem is that the viewfinder is blacked out by the IR filter in front of the lens, and therefore we can neither frame nor manually focus. The reason for having this filter on the lens is indeed that of blocking all visible light! There are two alternatives to circumvent this problem, neither one particularly effective nor elegant. The first one is to use an external viewfinder mounted on the flash hot shoe. The framing (apart from parallax errors) may be accurate enough but manual focusing is still not possible and we have to rely on the autofocus (we shall see in the second part of the article that this may or may not represent a wise choice). The second alternative is to mount the camera on the tripod, remove the filter, frame and focus, put the filter back on the lens and shoot. This requires a tripod but, considering that long exposures are the norm anyway, a tripod is required no matter what. It is a rather cumbersome and slow procedure, though.

For the above reasons serious IR photographers have always preferred twin-lens reflex (e.g., Rolleiflex) or rangefinder cameras (e.g., Leica in the 35mm format or Fuji or Mamiya in the medium format). Point-and-shoot cameras with a separate viewfinder can also be a solution, although the cheapest ones may not have a way to mount the IR filter on the lens.



Nikon FM with Hoya R72 filter



Separate viewfinder to be mounted on the flash hot shoe to allow the framing in a TTL reflex camera when an IR filter is present.



High-end digital point-and-shoot. The optical viewfinder (top right) allows us to frame the picture and shoot even though an IR filter has been mounted on the lens.

EXPOSURE

The light meters inside cameras are not designed to deliver accurate measurements outside the range of visible light. It is necessary therefore to carry out some experiments in controlled situations (using a grey card, for instance) to better understand their IR response curve in relation to that of the CCD sensor (or the film). The standard practice of bracketing becomes here a very useful tool, and when using digital cameras one does not even have the issue of wasting film. Having said this, after a while one develops a good feeling for the amount of infrared radiation present in an image and knows intuitively the amount of exposure compensation called for.

Moreover, a digital camera allows us to review the picture immediately after the shot and checking its histogram. This greatly simplifies the whole process of IR picture taking as far as exposure is concerned. We shall see that the other challenge in IR photography, i.e., that of focusing correctly, is a completely different and much more complicated story.

FOCUSING

The refractive indexes of a material (such as glass, for instance) are a function of the wavelength. One of the consequences of this is that rays of different wavelength follow different paths when going through a lens. In a simple lens design only one wavelength is correctly focused: all the others happen to focus correctly *behind* the plane where this one is in focus (longer wavelengths) or *in front* of it (shorter wavelengths). This situation is incompatible with the goal of forming an image of good quality on film (or on a sensor) so lens designers have come up more than a hundred years ago with lens designs that improve this situation. To translate this technical discussion into a more qualitative one that is closer to our photographic interests, let us talk about *colors* and not wavelengths now.

We have seen that visible light is made up of infinite shades of colors that go from violet down to deep red. To see what happens to visible light when going through a single lens we group for sake of simplicity all visible light into three colors, i.e., blue, red and green (see below).




When a ray of light goes through a lens it is subject to a phenomenon called *axial chromatic aberration* (see below). This distortion consists of bending a portion of the ray in different ways depending on its color (or, in more scientific terms, its wavelength). This is indeed a consequence of the refractive indexes being dependent on the wavelength, as mentioned before.⁷

To better understand how chromatic aberration varies depending on the portion of the visible spectrum going through a lens, the kind of material being traversed, and the angle of incidence, the following link is an excellent reference. It provides an interactive quick primer on the subject.

<http://micro.magnet.fsu.edu/primer/java/refraction/refractionangles/index.html>

If this distortion is not corrected we will have focusing problems because as the figure below shows it will be impossible to focus correctly *all the various colors*. If we assume that the green portion of visible light (i.e., the green ray in the figure below) is correctly focused the blue and red portions (i.e., the blue and red rays in the figure below) are not: the blue portion happens to be in focus *in front* of the film (or sensor) plane; the red portion ends up being in focus *behind* the film (or sensor) plane. In the case of the blue ray we have a *front focus* problem; in the case of the red ray we have a *back focus* problem.




Axial chromatic aberration: different portions of the visible light spectrum focus on different planes parallel to the cross section of the lens when going through the latter

When optical lenses for photographic applications are designed this problem is addressed because a lens-camera set-up shall make it possible to correctly focus all visible light: that is, the blue, green and red components shall end up in focus *exactly on the same plane*. In practice, economical and practical considerations conjure to pursue different goals. While a single lens will be able to focus correctly a single ray (i.e., wavelength) an *achromatic* lens like the one shown below is able to focus two rays. An *apochromatic* lens — whose simplest design consists of three lenses — will be able to have three rays in focus at the same time. Finally, those lenses called *superachromatic* can have multiple rays of light in focus at the same time.

⁷ The reality is more complicated than what we are presenting here and in the following sections, but for the practical needs of a photographer (as opposed to those of a lens designer) it suffices. Those interested in a rigorous and scientific treatment of all subjects related to lens design and optics are referred to “Applied Photographic Optics” by Sidney F. Ray, Focal Press, 2002.

Superachromatic lenses were theoretically introduced in the early '40s and became commercially available to the general public in the early '70s.



Simple correction of chromatic aberration via an achromatic doublet. Other two-lens designs can address this problem, the one shown here is just one possible approach.

We have limited ourselves to a discussion related to visible light. In fact, these considerations could be extended well beyond violet and well below deep red. That is, even apochromatic lenses can be significantly mistaken when focusing in the UV or IR portions of the spectrum. In practice, the vast majority of photographers have no interest in capturing the UV component or the IR component, film material responding to these wavelengths calls for special manufacturing, lens costs would increase for no reason, etc. Because of these economical and practical considerations, the corrections implemented in the design of a lens typically focus solely (pun intended) on visible light.

There are very few lenses available to the general public that have been designed to work outside the range of visible light. In the case of ultraviolet, two examples are the UV-Nikkor 105mm f/4.5 for 35mm cameras and the Zeiss UV-Sonnar 105mm f/4.3 for medium format, the latter focusing to around 215nm with no correction necessary. Both lenses utilize fluorite and quartz elements in their design and can go down to 700nm, i.e., they operate in visible light as well. As far as near infrared is concerned, two lenses that need no focus correction in this range are the Leitz Apo-Telyt R 280/4 for 35mm cameras and the Zeiss Sonnar 250/5.6 Superachromat for medium format cameras. The cost of some of these lenses can be as much as five times that of lenses of similar focal length and maximum aperture but designed solely for visible light.

Most manual focus lenses have some sign on the barrel to help re-focus for IR photography. This sign can be a red dot, a red line or rhomboid near the focus ring. It is therefore necessary to 'rotate' the barrel so that the distance that in visible light allowed us to obtain the best focus now corresponds to the red sign on the barrel. For sake of simplicity, let us assume that we are shooting a panorama, and therefore we focus at infinity. Well, to correct for IR focus the symbol ∞ shall now be made to coincide with the red dot on the barrel and NOT with the white central mark of the focusing ring. This is better explained by the pictures below, that show a Nikon AIS lens with a red dot indicating how the focus adjustment for IR photography shall be implemented.



The lens is focused at infinity and in visible light (black dot aligned with the ∞ symbol)



The lens is focused at infinity for IR photography (tiny red dot aligned with the ∞ symbol)

There is also a second chromatic aberration, called *lateral*, that happens outside the optical axis (hence the adjective 'lateral'). It results in a modification of the relative dimensions of an object being photographed. Discussing lateral and other aberrations is outside the scope of this article; it suffices to say that they too are directly proportional to the wavelength and therefore tend to affect IR photography more (where in fact the wavelength is greater), especially if the lens designer has not taken countermeasures to reduce their effect (in practice controlling lateral chromatic aberration for instance is quite complicated in visible light as well).

DIGITAL IR PHOTOGRAPHY WITH IR EXTERNAL FILTER

Quite simply, we proceed like we did in the analog case, i.e., we mount an IR (or very dark red) filter on the lens. In spite of the presence of the hot mirror in front of the sensor that is supposed to prevent IR radiation from reaching the sensor, some of it reaches the sensor anyway. The amount depends on how strong the filter is (i.e., how much attenuation the filter introduces in the stop-band portion, in technical terms) and varies from camera model to camera model and from manufacturer to manufacturer. IR photography is still possible therefore, in spite of the hot mirror. But it may come with some serious side effects.

First, the good news. Because the digital camera has not been modified at all, once the IR filter is taken off the lens the camera can shoot color in visible light , black and white, etc. The side effect we mentioned above is that the presence of the hot mirror results in exposure times that in some cameras with particularly strong hot mirrors can reach twenty or thirty seconds (in a sunny day). Canon cameras are known for having strong hot mirrors, for instance. While in analog IR photography 'long exposure times' meant a few seconds, here we may have a few tens of seconds. The problem goes beyond that of needing a strong and stable tripod and having even fewer subjects that can be shot with such long exposure times: the problem of these long exposure is called *digital noise, that results in objectionable grain in the picture*.

The picture below has been taken with a Canon 1DsMkII at f/8, 400 ISO, 4 seconds. It is the famous prison of Alcatraz in the San Francisco Bay Area.



Alcatraz. Canon 1DsMkII, 28-135mm zoom, F/8, 4 seconds with 400ISO.

The picture looks ok until one examines a 100% crop as the one shown here.



100% crop showing the noise level in the picture.

There is significant noise in the picture, and this example is also much more forgiving than others. The picture can be improved via software using (judiciously) programs like Noise Ninja or Neat Image. The latter has been used here, and has produced the result shown below.



Same crop after processing with Neat Image. Much less blur and some reduction in sharpness.

It should be clear at this point that a correct choice of the subject, the right lens and the right illumination and some serious post-processing can in fact deliver some good result in this case. But the long exposure times cannot be circumvented.

The presence of the hot mirror coupled with unwanted reflections inside the barrel has a second, somewhat surprising and truly annoying effect: *hot spots*. A hot spot is a circular (or shutter-shaped) portion of the image (in the center) that is brighter than the rest of the picture (see below). It is absolutely unacceptable and extremely difficult to remove with editing programs like Photoshop.



Hot-spot in a Nikon Coolpix E5000 with Hoya R72 IR filter (AF Auto t: 1/2s f:4,5 400ISO Zoom 21,4mm)



Hot-spot in a Canon 1DsMkII with IR B&W filter. Nikon 55mm/1.2 AIS lens. Similar behavior also occurs when using a Canon 16-35mm/2.8 USM L zoom and many other Canon 'L' lenses.

A number of experiments have provided us with some general idea of how hot spots behave and when (and when not) they may show up in a picture. In the case of the Nikon E5000 hot spots become more visible when the zoom is at 21mm (maximum extension) and/or when very small shutter apertures are used and/or when there is an unusually strong IR radiation.

In the case of the Canon 1DsMkII hot spots depend on the specific lens being used but also on where the lens is pointed and on the quantity of IR radiation present. It is quite difficult to prevent this phenomenon from happening. Reviewing pictures after the shot is advisable, although the circle in the middle of the picture can at times be quite difficult to detect by just looking at the small on-board LCD screen. There are a few simple rules that one can follow to improve the situation, however:

1. Fixed focal length lenses are better than zoom lenses (simpler optical path).
2. Slow lenses are to be preferred to fast lenses.
3. Lenses with the red mark on the barrel have been OK-ed by the manufacturer for use in the IR range and therefore should be *in theory* less susceptible to hot spots. We say 'in theory' because we have found examples of the contrary.

THE ALTERNATIVE: MODIFYING THE CAMERA FOR IR PHOTOGRAPHY

The one but critical modification consists of removing the hot mirror and letting IR radiation reach the sensor "full blast." The key question then is whether to substitute the hot mirror with something else or not. If not for reasons of protecting the sensor, something must be put in place of the hot mirror. The choice could be a clear (neutral) filter or one that blocks

all visible light and lets only IR radiation go through. The first option is typical in astrophotography, where photographs of a narrow portion of the spectrum are usually taken. The camera is left to capture a wide band of radiation but specific filters (H-alpha) are put in front of the lens to act as high-Q pass-band filters and then the overlapping of single images is performed. To have an idea of how the sky is transformed when represented via these wavelengths we recommend the reader to have a look at the paper "Near, Mid and Far Infrared" in

<http://www.ipac.caltech.edu/Outreach/Edu/Regions/irregions.html>


An excellent tutorial on IR celestial photography can be found at

http://coolcosmos.ipac.caltech.edu/cosmic_classroom/ir_tutorial/

Going back to our more down-to-earth activities, it is possible to replace the hot mirror with a glass that is optically transparent to both visible radiation as well as IR. The extraordinary plus in permanently removing the hot mirror is that hand-held IR photography becomes possible. Because of the sensitivity of the CCD sensor to infrared radiation, the exposure times drastically drop to the point that it is possible to shoot hand-held under normal illumination. Exposure times of 1/250s or shorter are possible. IR photography can now capture moving subjects, even sports events! Another advantage of this modification is that our camera will retain its capability of picture-taking in visible light.

The disadvantages are rather serious, though:

1. When shooting in visible light it is necessary to use a (clear) filter in front of the lens that blocks the infrared components of the spectrum, duplicating therefore what the hot mirror was doing in the camera before being modified. One has to buy the filter, the availability of which can be problematic. Step-up rings may have to be bought to adapt the filter to lenses with different front diameter. In other words, this is an extra cost. An alternative to purchasing the filter could be to use a color profiling of the camera so that the overall color balance can be brought back to its original settings.
2. When shooting in infrared an IR filter has to be used in front of the lens. The filter is the one we discussed already, i.e., a Hoya R72 or B&W 89C or 89B. As we now know this makes it impossible to focus and frame, as the view finder is completely blacked out by the IR filter. An external viewfinder becomes necessary, etc; in other words, *a major nuisance*.



Classical response of a filter to be used as a replacement of the hot-mirror that can be found from

What we have carried out therefore is the more complex and ambitious modification, i.e., to *specialize* the camera in such a way as to making it become a *high performance IR picture taking machine*. We have removed the hot mirror from within the camera and replaced it with an IR high-pass filter. This filter lets the radiation between 1000nm and 1300nm go through. Beyond 1300nm the sensitivity of the CCD sensor drops abruptly anyway, so there is no point in using a filter letting even deeper infrared radiation pass. While we understand that all these numbers can be intimidating to some, they are important because we will use them below to discuss the challenges in achieving accurate focusing. With this modification the camera is now capable of capturing *near infrared wavelengths*.



Hot mirror filter examples

The reader will have already realized the most significant advantage in such a solution: it is no longer necessary to use a filter in front of the lens! The filter is now right in front of the sensor and the photographer sees through the view finder a normal scene, and can therefore frame, focus, etc. No more step up rings, no more external viewfinder. Needless to say, this modification shares with the one above the ability of delivering hand-held IR picture taking to the photographer. The one disadvantage is that the camera is now unable to take pictures in visible light: as we said above it has become a *high performance IR picture taking machine*.



Canon 350D (Rebel XT) spectral sensitivity with hot mirror in place (Author: Chris Buil)



Canon 350D (Rebel XT) spectral sensitivity without hot mirror (Author: Chris Buil)



Carrying out the difference between the two spectral responses we can show the spectral response in a modified camera: in our camera -due to the passhigh filtering- only the infrared radiation is captured.

It is interesting to assess the new range of radiation that our IR-modified DSLR can now capture. The figures above show exactly this.

Moreover, experiments carried out by Terry Lovejoy have shown that when the hot mirror goes together with the anti-aliasing one, and therefore removing the former means also removing the latter, some intriguing side effects take place. First, the phenomenon of *moiré* becomes more evident (a minus), but if we process the image with a moderate *gaussian blur filter* to remove the moiré and then apply an *unsharp mask* to it the final picture is *sharper and snappier* than the image obtained by the unmodified camera. Incidentally, it appears that this is the path that the designers of the Leica M8 have undertaken.



The 'impossible' IR picture (see text). Permanently IR-modified Rebel XT, 50mm/1.4, 1/640s.

Permanently modifying the camera allows us to achieve a level of performance impossible before the advent of digital photography. The picture above, taken at 1/640 of a second and freezing the water dropping out of the fountain, would have been impossible before digital sensors became available.

Let us now have a look at what needs to be done to modify the camera. After that we shall return to the major issue of this kind of picture taking, i.e., managing the focus problem.

PERMANENTLY MODIFYING THE CAMERA: DETAILS

IMPORTANT:

Opening the camera voids the warranty of the manufacturer. If you want to carry out the modifications described below, please understand that you do it at your


own risk. The authors of this paper assume no responsibility whatsoever for the procedure explained below or its consequences. You have been warned.

What follows is a general description of the basic steps. The goal of this article is not that of describing in the most operational details all that needs to be done to carry out the modification described above, but only to document the work we have done to reach the final result. We have in fact carried out such modifications on different cameras and tested the results with lenses from different manufacturers.

The modification goes through the following steps:


1. READ AGAIN THE WARNING ABOVE!
2. Disassemble the camera
3. Remove springs, buttons, micro gears, etc
4. De-solder at the points where we need to intervene.
5. Detach all the cables
6. Remove the small printed circuit boards
7. Disassemble the CMOS sensor
8. Remove the hot mirror in front of the CMOS sensor. This is the clear glass that we clean when we notice dirt spots in our digital images.
9. Remove the tiny gaskets.
10. Replace the hot mirror with the IR filter (when the latter modification has to be carried out)
11. Correct the lens mount-to-sensor distance. Tuning/recalibration of the AF. This is to manage focusing issues.
12. Re-assembly going backward through the list.
13. Turn the camera on!

Please note that there are circuits inside a camera that are extremely sensitive to electrostatic discharge and can be destroyed by it. Also, in some cases it is required to remove some components and one ends up practicing electronic surgery in areas that are quite sensitive to thermal shock. Our advice is to always use a professional soldering station such as a 25W Weller Antistatic, for instance. An ESD wrist strap is strongly recommended, better safe than sorry. These operations are not particularly complex per se but require know-how and a steady hand.



Disassembling the camera

Disassembling the sensor calls for extra care because, for one thing, it must happen in an environment for all practical purposes *free of dust*. Once the hot mirror is removed the sensor is completely exposed to the environment until the IR filter is attached to it. Any particle that during this interval is deposited over the sensor will end up trapped between the sensor and the IR filter *and it will not be possible to remove it unless the sensor-IR filter is disassembled again* (and at the risk of having even more dust being deposited on the sensor!).



The sensor assembly of the Rebel XT with its original hot mirror

Every camera model has a design with its own unique quirks, Having said this, this kind of modification can be carried out on probably 90% of the digital cameras on the market. It is not possible when the hot mirror is completely mated with the sensor and removing the former would result in destroying the latter.




Comparison (in visible light!) between the IR filter (left) and

the Canon hot mirror (right). Note the significant difference in thickness of the two filters. In order to obtain a correct focusing these differences in thickness must be accounted for: the position of the new plane of focus after the modification shall be determined.



The modification has been completed. The sensor is now behind the IR filter (that looks black to the human eye).



The old hot mirror seen in the infrared range appears to be completely black.


THE MISTERY OF AUTOFOCUS

If the reader remembers the complexities encountered when discussing chromatic aberrations and the issue of focusing within and outside visible light, it is only natural to suspect that autofocus may present a number of not-so-hidden traps.

Indeed, the kingdom of uncertainty on our photographic planet is called "autofocus." Many of us have been often skeptical about the performance of autofocus (or its consistency) in visible light. Well, autofocus is not a fool-proof mechanism...on the contrary. When the photographer operates in the infrared portion of the spectrum the behavior of the autofocus gets even more difficult to comprehend (and master).

To interpret correctly the tables that will be presented and the problems that will be raised, we need to quickly introduce two basic concept of optics, i.e., *circle of confusion* and *depth of field*. The reader already familiar with these terms can skip what follows and go to the section RECALIBRATION.

The circle of confusion is defined as the smallest circle that the human eye can detect as a tiny dot and seen from a defined distance. In other words it is that percentage of out-of-focus that the human eye is not able to detect, and hence instead of a circle we see a tiny dot. Assuming an average healthy human eye, experiments have shown that the eye is able of distinguishing 5 line pairs per millimeter when looking at a negative of 20cm by 25cm seen from 25cm. This translates in ten lines per millimeters, that is 254 lines per inch, or 254lpi.



A "line pair" consists of a white line and a black line.



[5 line pairs = 10 lines in 1mm x 25,4 (inch conversion) = 254 lpi/dpi]

The CoC is equal to the inverse of the human eye's resolution and is 0.2 for a 20x25cm negative. From this value it is possible to derive the CoC for all other formats. In the case of digital formats one can start from the CoC of the 35mm ('Leica') format and introduce the crop factor of the sensor. The table below summarizes the results.

Manufacturer/format	Model	CoC (mm)
Canon	300D, 350D, 400D (1,6x)	0.019
Canon	1Ds MarkII, 1DS, 5D (1x)	0.030
Nikon	Coolpix E5000 (3,9x)	0.008
Nikon	D1H, D2H, D50, D70, D100, D200 (1,5x)	0.020
APS		0.025
24x36mm		0.030
6x6cm		0.045
4x5"		0.100
8x10"		0.200

Table with Coc's for the most common photographic formats and digital sensor sizes.

The depth of field is the area in front of us that will be perceived as "in focus" and it generally extends 1/3 in front of the focal plane and 2/3 behind it. The depth of this area depends on the focal length of the lens been used, on the distance from the focal plane, on the CoC and on the aperture (f-stop) being used. *It does not depend on the format of the film (or the size of the sensor)!*



A single dot is correctly focused on the plane (middle). Errors in focus create a circle on the plane of focus. The lens can front-focus (bottom) or back-focus (top). However, if the CoC is smaller than the one the human eye can detect the picture will still be perfectly acceptable.

The depth of field is also defined as the region where the size of the CoC is smaller than the ability of the human eye to disambiguate it. That is, all the circles smaller than that of the CoC will appear as in focus.

The hyperfocal distance is that distance where the closest object appears sharp and in focus when a lens is focused on *infinity*. When a lens is focused at its hyperfocal, everything that lies between half the hyperfocal distance and infinity will have an acceptable sharpness for photographic purposes. The table below summarizes in quantitative terms what just discussed^{8,9}.

	Standard equation	Simplified equation
Hyperfocal distance:	$H = \frac{f^2}{Nc} + f$	$H' = \frac{f^2}{Nc}$
Minimum distance with acceptable sharpness:	$D_n = \frac{s(H - f)}{H + s - 2f}$	$D'_n = \frac{sH'}{H' + s}$
Maximum distance with acceptable sharpness:	$D_f = \frac{s(H - f)}{H - s}$	$D'_f = \frac{sH'}{H' - s}$

⁸ Once the focal length is expressed in mm, all other parameters shall be expressed in mm for congruence reasons.

⁹ The simplified equations can be safely used when the focal length is much smaller than the focusing distance.

Where:	
H, H'	Hyperfocal distance, mm
f	Lens focal length, mm
s	Focusing distance, mm
D_n, D'_n	Minimum distance with acceptable sharpness
D_f, D'_f	Maximum distance with acceptable sharpness
N	f-stop
c	Circle of confusion, mm

RECALIBRATION

The step described as *tuning/recalibration of AF* is in fact the re-analysis of the optical design when the hot mirror is replaced by the IR filter or by the clear filter. All filters act like lenses and every material has its refraction index, as we have seen above. Refraction is responsible for bending a ray that goes through the material and the amount of bending is governed by the refraction index. Air has a refraction index slightly greater than 1 (the vacuum is a perfect 1). An optical glass like BK7 has a refraction index of 1.517.

The IR filter we have added is right in the optical path and plays a crucial role in making the image land exactly where the plane of the CCD detectors of the sensor is. A few tenths or hundredths of millimeters can make a huge difference.



Whenever the modification is carried out it is important to consider all these factors so that the plane of correct focus lands exactly (1) where the CCD sensors are and (2) for IR

wavelengths. Tiny adjustments are possible by tuning the AF in the mirror box, but in fact it is advisable to proceed in such a way that all the corrections balance themselves out rather than 'tuning' things like the AF.

In order to do that one has to compute the difference in the optical path before and after, using the following formula that is derived from the ones used to compute the hyperfocal distance:

$$\text{Adjustment} = (n - 1) / n * D\text{Thickness}$$

Where:

n	Refraction index (1.517 for BK7 glass)
DThickness	Difference in thickness between the original filter and the new one
Adjustment	Shift in the plane of correct focus

In the case of the modification for IR photography that we are discussing here the goal is to make sure that the visible light going to the viewfinder and the IR range of wavelengths reaching the sensor balance themselves in spite of the presence of axial chromatic aberration.

Unfortunately all lens designs are not made equal, i.e., all lens design (and this includes also the material being used and not just the element shapes and positions) behave differently in the way they respond to radiation going through them outside their intended design envelope. Even though high quality lenses have been designed to be apochromatic, in general this feature does not extend down to infrared and stays confined to visible light¹⁰. Once again, let us never forget that *we are working outside the design envelope of the optical system and that we have to cope with some degree of uncertainty*.

The focus correction when working in infrared that we have discussed above is not constant across focal lenses and optical designs but in fact varies from lens to lens. Most important, the amount of compensation depends on the *specific wavelength we want to have in focus*. For this reason Leitz has refused for many years to put an IR marker on the barrel of their lenses because it gives the scientifically incorrect notion to the photographer that "it compensates for focusing in IR," when in fact that compensation applies *exactly* to one specific wavelength. Because of all these considerations the operational advice when doing IR photography is therefore that of always closing down the shutter: the increased depth of field will compensate (hopefully) for an error in focusing. We shall see below, however, that this is not without unpleasant side effects.

THE PROBLEM OF DIFFRACTION

We have seen that working at small f-stop is a way to fight against the uncertainty in focusing within the IR range. Larger f-stops increase the depth of field, in fact. Unfortunately, large f-stops have a serious drawback when working in infrared. In general a lens should not be used at very large f-stops (e.g., f/32 or f/22 in lenses for the Leica format) because of diffraction problems that degrade the resolving power of a lens and hence the image quality. It is well known that diffraction is proportional to how much the shutter is closed down, that's why knowledgeable photographers shy away from extreme f-stops whenever possible. Fewer people instead know that diffraction is also proportional to the wavelength. This means that when shooting in the infrared range the onset of objectionable diffraction *happens for smaller f-stops* than when one works in visible light. There is no fast and hard rule to apply here; it suffices to say that we cannot address the focusing issue by simply cranking down the

¹⁰ We have seen that lenses such as the Leitz Apo-Telyt-R are corrected for near IR, but these are truly exceptions and certainly not the norm.

aperture: the image degradation because of diffraction would be more severe than that in visible light. Extreme caution in using large f-stops then.

CORRECTING FOR IR PHOTOGRAPHY

When we correct the focus to account for the fact that IR wavelengths are now of interest we move the lens' internal elements forward. That is, the distance between the elements and the sensor *increases*, i.e., as if the subject we are photographing were *closer* to us.

We have carefully measured the increased distance between the internal lens elements and the sensor with a resolution of 0.01mm. We have carried out this measurement on lenses with a simple optical design (50mm lenses). Groups of elements often move independently one from the other inside a lens when focusing and in this case the concept of "increased distance from the sensor" would have had no meaning.

Nikon's 50mm lenses have a single group of elements that move all together and this is the reason why we chose them for our measurements. The 55mm Micro-Nikkor is a special purpose lens and has been added to the table only as a reference.

We have carried out several measurements for each lens first focusing at 1m and infinity and then correcting the focus for IR wavelengths. Given the approximate nature of the IR correction we have averaged the results obtained to account for the unavoidable errors in rotating the focus ring of the lens.

Lens	Focal length	f-stop	IR Shift (mm)
			Distance increase from the focal plane
Nikon AiS 50mm/1.4 MF	50	1.4	0.15
Nikon AiS 50mm/1.8E MF (metal)	50	1.8	0.18
Nikon AiS 50mm/1.8 AF	50	1.8	0.16
Nikon AiS 55mm/2.8 MF Micro	55	2.8	0.15
Yashica 50mm f1.7	50	1.7	0.15

The table shows that the increase of the distance between lens elements and plane of focus is between 0.15 and 0.18mm. An average of 0.16mm of increased distance seems to be typical for this focal length and lens design. Interestingly, this corresponds in the case of the Nikkor AI 50mm/1.4 to moving the focus ring about 2.3mm, or better, to rotating it 1°18'.

EXPERIMENTS WITH CANON AND NIKON LENSES ON AN IR-MODIFIED CANON REBEL XT (350D)

In what follows we present the results of several experiments we carried out on a Canon Rebel XT (350D) that we have permanently modified for IR photography as described above. We have compared the results from this camera to those from a Canon 1DsMkII with an IR filter on the lens.

MODIFIED REBEL XT AND CANON 1DsMkII WITH IR FILTER

The filter we have used on the Canon 1DsMkII is the B&W 87C (093). The lens was a zoom Canon 16-35mm f/2.8 USM L. This lens has some issues as we shall see below, but serves our

purpose here perfectly. Because the sensor in the Rebel XT is smaller than that of the 1DsMkII (about 1.6x) we had to decide whether to keep the same angle of coverage in the pictures taken with both cameras or maintain constant the number of pixels. We opted for the latter and therefore we cropped the image of the 1DsMkII in such a way that it had the same resolution of the one from the Rebel XT.

Both images were recorded in RAW and converted in TIF using Canon's DPP without any change whatsoever. We have applied no sharpening. In general, Canon DSLRs generate pictures (with in-camera sharpening off) that are somewhat soft and require robust sharpening, but the goal here was not to evaluate the quality of the image in its minute details.

The image below has been taken with the Rebel XT:



And this is the one from the Canon 1DsMkII



The difference in colour is simply caused by the different way the three channels respond. The modified Rebel XT tends to overexpose the red channel with very fast lenses. In other words, the intuitive choice of picking the red channel as the 'best' one often proves to be wrong. This also means that it is not true either that the 'best' channel is the green one, as here: it all depends on the subject, how much it reflects IR radiation, the lens being used, etc.

This is the red channel, seriously washed out:



The green channel, quite balanced:



The blue channel, also pleasing the eye, but with more noise (impossible to see in these pictures) than the green channel:




We now show the three channels from the 1DsMkII. The red channel, with the classical IR look:



The green channel, darker and less 'IR-ish' than the red channel:



The blue channel, somewhat between the two, pictorially:



The first observation is that there is quite a lot of noise (impossible to see in the pictures above) in the image taken with the 1DsMkII. This was to be expected because the exposure time was 20 seconds, while that with the Rebel XT was 1/200s (!), same f-stop. A second observation is that it is easy to choose the 'best' channel (the green one) in the case of the Rebel XT, while this is not true in the case of the 1DsMkII. We always recommend to use the channel mixer in Photoshop to produce the most pleasing result, which is always a balance between IR-look and digital noise.

In the second part of this article we will consider the issue of IR focusing in much more detail by examining the behaviour of several lenses, classical manual Nikon AIS lenses as well as the latest Canon AF ones. We will discover that IR focusing is a rather complex issue and that the autofocus complicates matters even further. We will not provide a list of 'good' and 'bad' lenses because such a notion is superficial and completely meaningless. The goal will be instead to expose the reader to the more or less hidden traps so that when he or she decides to embark in this fascinating adventure of IR picture taking he or she will know more or less what to be aware of. This is the contribution that we hope the second part of this article will provide, i.e., not a quick set of fast and loose rules but that of surfacing *the issues* and discuss the ways to cope successfully with them.