Notes on Photography for Beginners

by

Kyriakos Kalorkoti

Web site: kkpictures.com
Email: kk@kkpictures.com

September 2006
(Revised May 2016)

Contents

1 Introduction

The idea for these notes arose out of assisting at a workshop for beginners given by me and Joe Cornish. They are based on various questions asked there and some other relevant points. They consist of a sequence of fairly brief sections most of which are devoted to a single topic.

In writing the notes my preference has been to provide explanations (at least to some level) of the various concepts involved. In some cases diagrams or intuitive arguments are sufficient. Inevitably there are cases where mathematics is necessary, e.g., in understanding why the series of f-numbers looks the way it does. If you find that an explanation is hard going, or just of no interest to you, then you can safely skip it. Even some of the best photographers would be hard pushed to explain why some of the numbers look the way they do; it really doesn't matter the point is that they know how to use them.

I must explain that I do almost all my photography with large format transparency film. I have not done any serious photography with digital cameras but practically all the technical details are the same. The main difference is the brightness range that can be covered by transparency film as opposed to digital sensors (or negative film for that matter). This will be discussed in the appropriate section.

NOTE: All cross references are hyper links, so just clicking on one will take you to the appropriate place; hyper links appear coloured.

2 Understanding exposure basics

Exposure is simply the total amount of light that you allow to fall on the film or sensor. Too little light means a very dark picture and too much means a very light one. At its extreme this can mean a totally black or totally white result. There is no such thing as a correct exposure (in the sense of a precise single one) because for creative reasons you might want to render one scene slightly dark or slightly lighter. There are of course limits, the critical thing is to understand how to use exposure to get the effect you want.

Automatic exposure makes the decisions for you and for many situations the result is pretty good. This is much more the case now because manufactures use matrix (or evaluative) metering almost as standard: this measures the light in several parts of the scene and then compares the result to a collection of patterns derived from real scenes. From this the camera can make a good guess at a good exposure. However no matter how well this works there are always situations where it fails and in any case making an informed decision about exposure is one of the key controls you have over how your photograph looks.

Unfortunately a myth has grown up about exposure that makes people dread it as being incredibly difficult hence the lure of automation. (To be fair automation is very useful in situations where you must react very quickly, e.g., young children aren't going to freeze in mid play while you adjust things precisely!) When I was starting I bought, and read several times, a book on exposure that left me none the wiser. The problem was that it confused the *theory* of exposure (which is indeed complicated) with the *practice* which I hope to show you is in fact refreshingly simple. To do that we need to understand some basic ideas, so bear with me.

To introduce the next three topics let's use an analogy; instead of trying to capture a total amount of light we want to fill a container with 1 pint of water using a hose. To do this we can use a hose with a very narrow bore in which case we have to pour water for a fairly long time. Alternatively we can use a hose with a much wider bore in which case we need less time. Of course the actual amount of time also depends on water pressure. In any case here are the three things on which exposure depends:

1. How strongly lit our scene is.

2. How much light can pass through the lens in a fixed period of time.

3. How long we allow the light to pass through the lens and fall onto the film or sensor.

2.1 Exposure meters

These measure the light for you. Most cameras have a built in meter. Strictly speaking these meters measure the amount of light *reflected* by your subject rather then the amount of light falling on it. Think of a white washed cottage with a black door; the wall and door receive the same illumination but of course the wall will appear very bright under most conditions and the door, well, black. If you point your camera just at the white wall it will see a lot of light coming back at it and assume that it should give reduced exposure. If you point it just at the black door it will see much less light and assume that it should use more exposure. How can these be correct? After all there is only one scene and it is illuminated in the same way.

The reason for the discrepancy is that exposure meters for reflected light are calibrated to assume that the scene is reflecting 18% of the light falling on the scene (because this is what happens on the average in mixed situations). Any colour that reflects this amount of light is classed as *mid-tone*. When meters calculate an exposure they produce one that would result in a mid tone version of whatever you are looking at. Put briefly when you point your meter at a uniformly lit object what it is saying to you is:

> I don't know what you are looking at but if you want it rendered as a mid tone version then use this exposure; now adjust to taste.

This is the key to understanding exposure and taking control over how your photographs look. We will discuss the phrase "adjust to taste" in more detail in §3.1.

Try the following experiment. Take a matt black card or piece of paper (around A4 size is good) and put it next to a clear white one, taking care that they are equally illuminated;

the best thing is to put them near a window on a fairly bright overcast day. Measure the exposure your camera gives for the black card and for the white one (take care that you do not cast a shadow by using a telephoto lens, e.g., zoom in). In each case your viewfinder should just show nothing but the relevant card. You will see that the black card is given *much* more exposure than the white one. In fact the black is given as much extra exposure as it needs to become mid tone and the exposure for the white one is *reduced* by an equal amount so that it also becomes mid tone.

You can buy from any good photographic shop a Kodak (or other make of) gray card[1]. The colour is not really relevant the point is that this is manufactured to reflect 18% of the light. Now put this with your black and white cards (or paper) and measure the light. I did this experiment at the time of writing and here are my measurements

Black EV6,   Gray EV8,   White EV10.

EV stands for *exposure value*, an increase of 1 means doubling the exposure, and of course a decrease of 1 means halving it. Your camera is more likely to show you time and aperture (see below) but as long as you keep one of these constant (by selecting either aperture or time priority) your experiment will be just as easy to carry out. To conclude, if you base your exposure on the reading for the gray card then the black will indeed be black (receiving a quarter of the amount of light used for the gray card), the gray will be mid tone and the white pure white (receiving four times as much light as for the gray card). If you use a digital camera (or even 35mm film) take the photographs and compare. First take a picture of each card separately, based on the exposure recommended for it, and then all three based on the exposure for the gray card. Your results might not be exactly as described here (e.g., due to variations in meter response or actual colour of paper used) but they should be fairly similar. We will continue this discussion in §3.1

2.2 f-stops

Imagine a steadily lit scene with your camera set on a tripod (for consistency) and lens fixed to your desired focal length. One way of controlling the amount of light that passes through the lens in a given time period is by varying the aperture available for it to go through (recall the water analogy in §2). In the old days photographers had various discs, or stops, with different size circular holes cut in them; the smaller the hole the less light that went through. For a given lens, if we halve the area of the hole then we halve the amount of light. In modern lenses the separate discs have been replaced by iris diaphragms with blades that close to a desired setting giving an appropriate size hole, but the principle is exactly the same.

[1]A good gray card is manufactured to a precise tolerance and will therefore cost at least a few pounds. It is a useful item to have but not essential so as an alternative just use a mid-grey coloured card or cloth such as a fleece. Your figures might not show the exact relationships described here but they should be close enough to make the point.

This explains the use of "stop" but why "*f*–stops"? This refers to the ratio of the focal length of the lens to the size of the stop (diameter of the hole). What this buys us is the fact that the same f-stop used on any two lenses receiving the same amount of light indicates exactly the same luminosity given to the film or sensor. So in arriving at an exposure we only need to know the f-stop being used in addition to the light intensity (as well as the sensitivity of our film or sensor, this is discussed in §2.5). The rest of this section, apart from the final paragraph, is devoted to an explanation of this and how the series of f-numbers arises. While technicalities are kept to a minimum it is the most technical section of the notes and can safely be omitted.

To understand how f-stops work, imagine two lenses one of focal length 50mm and the other of focal length 100mm both focused at infinity. The light from the 100mm lens has to travel twice as far as that from the 50mm lens in order to be focused. So the same amount of light is spread over a larger area (of course our film or sensor only sees a fixed portion); in fact the area is four times bigger for the 100mm lens compared to the 50mm lens. As a result our film or sensor receives one fourth of the light from the 100mm lens than from the 50mm lens. Now suppose we put a stop in the 50mm lens that reduces its light by eight times. What size of stop should we put in the 100mm lens so as to make the image have the same intensity? Well the 50mm lens now has 1/8 of what it had while the 100mm has 1/4 of what the 50mm had. So we need to reduce the light for the 100mm to 1/2 of what it currently is. The stop we used for the 50mm lens had an area that reduced light to 1/8 so to get a stop that reduces the light by only 1/2 we need one with a circular hole that is 4 times the area. Now if the original stop had a diameter of d mm then to get one with 4 times the area we need one with a diameter $2d$ (the area of a circle is proportional to the square of the diameter[2]). Now note the following ratios for focal length divided by diameter of relevant stop:

- for 50mm we have $50/d$;
- for 100mm we have $100/2d$.

But the second ratio is of course the same as $50/d$. This reasoning applies generally and we have found that as long as the ratio of focal length to diameter (or aperture) of stop is the same for two lenses then they will give equal illumination to our film or sensor irrespective of focal length[3]. This simplifies things for us as photographers immensely; as noted above,

[2] You have probably have met at some time the formula πr^2 (i.e., $\pi \times r \times r$) for the area of a circle, where r is the radius and π is a constant which can be approximated by 22/7. Of course if the diameter is d then $r = d/2$ so an equally good formula is $\pi(d/2)^2$ which is the same as $\pi d^2/4$. So if we double a given diameter from d to $2d$ then the area goes from $A = \pi d^2/4$ to $\pi(2d)^2/4 = (\pi \times 2 \times d \times 2 \times d)/4 = \pi \times d \times d = \pi \times d^2$, which is $4 \times A$. Note that if we just want to double the area then we need to multiply the diameter with a number s such that $(s \times d)^2 = 2 \times d^2$, i.e., $s^2 \times d^2 = 2 \times d^2$. This just means that we need to have $s^2 = 2$ and this gives us $s = 1.414\ldots$, which we can approximate by 1.4 for our purposes.

[3] Strictly speaking this applies so long as our lenses are focused at infinity, or for practical purposes on something not too close. For close up photography this reasoning has to be adjusted, there is a simple way

in our exposure calculations we don't need to know focal length and size of stop (just as well) all we need is the ratio. So we define

$$\text{f-number} = \frac{\text{Focal length}}{\text{Diameter of effective apperture}},$$

(we use “effective” to account for compound lens designs where things are a little more complicated but the ideas are essentially as described above). So given an f-number and a focal length the diameter of our stop is obtained by dividing the focal length by the f-number (this is why the larger the f-number is the smaller the aperture).

So far so good but where do those strange numbers come from? We can start at 1, so effective aperture is the same as the focal length. To continue, we want a series of f-numbers that reduces the exposure by a half each time (modern cameras and meters do it in 1/3 of a stop but let's keep things simple). So we need to divide the diameter by a number that halves the area of the stop each time. This means we need to multiply the f-number by the square root of 2 each time, i.e., by approximately 1.4. So if we start at $f/1$ we then go to $f/1.4$. After that we of course go to $f/2$ (the square root of 2 is equal to 2 when squared!). We then go to $f/2.8$ and so on. In this way we get the familiar sequence:

$$1,\ 1.4,\ 2,\ 2.8,\ 4,\ 5.6,\ 8,\ 11,\ 16,\ 22,\ 32, \ldots$$

Any numbers close to a whole number are rounded for simplicity which accounts for $f/11$ since $8 \times 1.4 = 11.2$ and similarly for $f/22$. If you want to find the sequence in steps of 1/3 of a stop then the multiplier is $\sqrt[3]{\sqrt{2}} = \sqrt[6]{2}$ which is little larger then 1.2.

The phrase *stopping down* is now pretty clear, it refers to the choice of stop that we make in order to control exposure.

2.3 Time

We have seen that one way to control the amount of light getting to the film or sensor is by stopping down. Of course the other control we have is the length of time during which we allow light to register on the film or sensor. At the start of photography several minutes or even hours were needed so a leather cap (or even a top hat) could be used to mark the interval. Nowadays when we might want to expose for as short a time as $1/8000''$ we have mechanical or electronic shutters. In fact for landscape photography we never use very short times but the shutters are useful anyway. Just as we built a series of f-stops to reduce exposure by a half each time, shutter times have been standardized to do the same so we have:

$$1'',\ 1/2'',\ 1/4'', \ldots$$

to do it. This is only necessary though if you are using a separate hand held light meter, through the lens metering takes care of things for you.

but in order to avoid fractions these are usually written just with the denominator as:

$$1,\ 2,\ 4,\ 8,\ 15,\ 30,\ 60, 125,\ 500,\ 1000, \ldots$$

As you can see each number is double the previous one or near enough to make no practical difference (this is done just to make them easier to remember, though Computer Scientists might well question this!). Obviously if light (or water) is flowing in at an even rate then allowing it to flow for twice the time gives us twice the amount.

The convention of not using fractions means that we need to have some way of distinguishing 2 meaning 1/2 from 2 indicating 2 seconds. Most cameras show a plain figure 2 for the first speed and $2''$ for the second (sometimes different colours are used, e.g., on the Pentax spot meter).

2.4 Shutter speed and aperture reciprocity

We have now seen that closing down the lens aperture by one stop halves the exposure. On the other hand opening the shutter for twice as long doubles the exposure. So the shutter speed and aperture reciprocity is now clear. Thus the following combinations all lead to the same exposure:

Apeture	1.4	2	2.8	4	5.6	8
Shutter	30	15	8	4	2	1

It is now easy to see how to draw up a table for other combinations: for a given starting aperture and shutter speed we close down the aperture by one stop (remember this means increasing the f-number) and double the shutter speed (remember that only the denominators of fractions is shown so the number is halved, e.g., $2 \times 1/8 = 1/4$.).

This reciprocity is correct without exceptions in the sense of actual exposure given. However the film or sensor might not behave as expected for exposures lasting longer than several seconds. For example if I am using Velvia 50 then I must adjust the time for anything over $4''$ (for this figure I use $5''$ instead). This effect is known as *reciprocity failure*. It is caused by a change of sensitivity with exposure time; most films are better behaved in this regard than Velvia 50. Digital sensors do not exhibit this phenomenon but you can end up with a noisy image instead (consult your handbook).

2.5 ISO (or ASA)

So far in our discussion of exposure we have mostly left out one obvious factor: the sensitivity of the recording medium. This sensitivity is indicated by the ISO number (ISO simply abbreviates International Standards Organization). This number is sometimes referred to as the ASA number (American Standards Association) though strictly speaking this has been superceeded; in any case there is no difference. The number is also called EI (Exposure Index) by some authors. In the past another way of indicating the sensitivity was by

ISO 50						
Apeture	1.4	2	2.8	4	5.6	8
Shutter	30	15	8	4	2	1
ISO 100						
Apeture	1.4	2	2.8	4	5.6	8
Shutter	60	30	15	8	4	2
ISO 200						
Apeture	1.4	2	2.8	4	5.6	8
Shutter	125	60	30	15	8	4

Figure 1: Equivalent exposures for varying ISO numbers.

the DIN number (Deutsche Industrie Norm); this was a different scale. For ISO a doubling of the number indicates doubling of the sensitivity (i.e., half the exposure is needed for the same result)[4]. For film the manufacturers recommend an ISO speed though photographers often choose to alter this somewhat to suit their taste (for colour transparency only small changes are made but for black and white negative film the changes can be quite large). Of course for digital photography the camera only has one sensor but the concept holds all the same, when you increase the ISO for your digital camera you are asking it to derive information from less light (i.e., smaller electrical charge per pixel). No matter whether we use film or digital, the higher the ISO rating the more grain or noise the image has. Thus for maximum quality a low ISO rating is preferred (there are also creative reasons, low ISO makes slow shutter speeds are easier to attain).

So now we have a third actor in our shutter speed and aperture reciprocity. This is still pretty straight forward though. If we add a stop we can either double the shutter speed or the ISO (but *not* both!) to get the same exposure. So if we assume that the table in §2.4 is for ISO 50 then the tables in Figure 1 all give equivalent recorded results (in terms of luminosity).

2.6 Stops

If you talk to a photographer it isn't long before phrases like *increase exposure by a stop* or *under expose by a stop* crop up. When used in this sense all that is meant is to give *twice* as much exposure or *half* as much respectively. So put simply each additional stop doubles the exposure and naturally each subtracted one halves it. Nowadays this gets refined to half or even one-third of a stop but the idea is the same. This change can, in principle, be achieved by varying any of our three controls (aperture, shutter speed, ISO rating) though in practice what is intended is the use of either aperture or shutter speed. This shorthand

[4]For the DIN scale 3 was added to indicate twice the sensitivity (see §3.4)

Adjustment	Result
$+2\frac{1}{2}$ stops	textureless white
$+2$ stops	extremely light
$+1\frac{1}{2}$ stops	very light
$+1$ stop	light
$+\frac{1}{2}$ stop	slightly light
Meter reading	mid tone
$-\frac{1}{2}$ stop	slightly dark
-1 stop	dark
$-1\frac{1}{2}$ stops	very dark
-2 stops	extremely dark
$-2\frac{1}{2}$ stops	textureless black

Figure 2: The effect of under- and over-exposure.

keeps things simple really.

3 Creative control

We now have the basics of how to create an acceptably exposed image but we still need to understand how to make an intelligent choice of the controls available to us. Sometimes automation is a perfectly intelligent choice *provided* we understand the situations for which this is the case. For landscape work it is often a hindrance and once exposure is thoroughly understood it really does seem rather unnecessary.

3.1 Choosing your exposure

Let's return to our experiment with the black, white and gray cards discussed in §2.1. How should we decide on our exposure to produce a beautiful picture of these three cards? Well one starting point is to use the exposure based on the gray card. However our type of photography is not documentary and we are free to vary our interpretation of a scene for dramatic effect. So we might want to make the gray card darker. We could use an exposure that is half the suggested one in which case our gray card will be dark grey but not black. The black card will be a richer black and the white one will still be fairly white but not so bright. This is the time to discuss in more detail the phrase "adjust to taste" used in the quote of §2. Figure 2 is a rough guide of the effect of using more or less exposure than suggested by the meter.

The verbal description of this tonal scale gives a reasonable idea but it would be best for you to make a scale. Just use a uniformly coloured card in uniformly lit conditions that are not changing (e.g., overcast around mid day) and make the exposures, you could just do them in 1 stop increments to save time (and film cost if you are not using digital) though half stops are a good idea. Your scale will look something like the following:

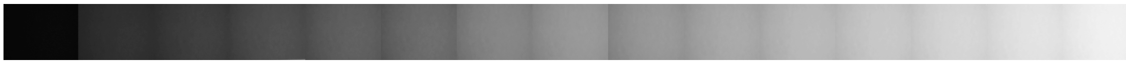

The difference is that this was made with a digital camera that allowed only 1/3 of a stop increments. Also the quality of reproduction will not give you a very accurate scale (in particular the right hand end should be lighter) but you should see a gradual lightening from left to right with mid tone around the middle.

Of course all the trouble over cards is not very exciting but replace the cards with parts of a scene and you get the point. One reason for underexposing in the way described might be because there is some clutter in a dark part of the scene that is best rendered invisible (and we are happy that the black will still work in the composition). Alternatively we might just want to create a more mysterious mood. As a bonus under exposed colours are richer.

Here, briefly, is how I arrive at my exposures. I use a Pentax spot meter to measure the light because this enables me to read it precisely on any part of the scene (it is also refreshingly free of gimmicks, it just gives me a reading, no fancy modes etc.[5]) This meter gives exposure in EV numbers (see §2.1) and a scale on the meter can then be used to convert them to aperture and shutter speed for any given ISO rating. I scan the scene taking note of the variation in EV numbers and in particular the highest and lowest. If the highest and lowest are within 4 stops of each other then things are fairly easy (because transparency film can record detail over 4 stops roughly speaking). Let's assume this is the case and for the sake of definiteness the lowest number was EV9 while the highest was EV13. This means that if I base my exposure on EV11 all the scene will record. My next step is to double check the critical parts of my composition and decide how dark or light they will appear in relation to this choice. For example it might well be that there is a large part of a mountain that reads at EV12. This means that in the final photograph based on EV11 this large part will be fairly bright and attract the eye. Assuming that this is undesirable, one thing I could do is to base my exposure on EV12. This will make the previously bright part just mid tone but of course those parts of the scene at EV9 will now be textureless black so I must be happy that this is acceptable, and with luck even desirable. The parts of the scene at EV12 will now be reduced in brightness but will have the same relation to all those parts of the scene that do actually record, relative brightness is the important thing. Of course a lot depends on the situation at hand but under the conditions described an exposure based on EV12 or perhaps EV11$\frac{2}{3}$ is likely to be my choice. If I go with EV12 then using my favourite film Velvia 50 the scale on my meter shows that at $f/32$ I should use a shutter speed of $1/2''$, quiet fast for me. There is

[5]I'm not anti technology, the simplicity of the meter means there there is very little to go, or be set, wrong which is critical when working under pressure.

another way to control the brightness problem if going to EV12 is undesirable, this will be discussed in §3.4.

This discussion raises the question of what to do if the brightness range (i.e., contrast) is more than 4 stops. We will discuss this in §3.4. Digital sensors (and negative film) have a wider range, depending on the type. Unfortunately I cannot report on any serious experiments with digital but it costs you nothing to try. All the discussions apply, just alter the figures as necessary.

So do you have to buy a spot meter? Unfortunately good ones are very expensive. Luckily most modern cameras have a spot meter setting. In most cases this is a rather optimistic description of what the setting does[6] but it will do a reasonable job. The main downside is that having chosen your composition you have to move the camera to determine exposure and then recompose. Of course if you use a tripod all you need to do is clip the camera back on it; use of a tripod is highly recommended for this and other reasons (see §4.2). The main drawback is in fast changing light but with practice you can develop a usable routine.

Finally, in this section, we now have the solution to a problem that beginners often come across. You have in front of you a superb misty setting so you go ahead and photograph it using automatic mode. When you look at the picture the bright mist is at best a murky grey, not good. Of course the reason is that as most of the scene is filled with mist the exposure is determined by the light reflected by (or transmitted through) it. The light meter has done its job and set the camera to produce a mid tone version of the mist. But mist is anything but mid tone, it will be at least one stop brighter and probably more. So at the time of exposing an extra $1\frac{1}{2}$ stops would probably be a good choice. If you forget to do this you can still rescue the photograph in Photoshop by using Curves under Adjustments in the Image menu (start by increasing the brightness of the mid tones) but obviously getting it right in the first place is preferable (the overall quality will be better as well as your pride in a job well done).

3.2 Choosing your aperture

So far we have discussed aperture purely as a means of controlling exposure. However it has an equally, if not more, important function in controlling how much of the scene is in focus (or, more accurately, acceptably so). In modelling the behaviour of lenses we think of the lens as being used to project mathematical points in 3 dimensional space onto a plane; this is an idealization of reality but a very useful one[7]. We say that a point in space

[6] A true spot meter reads light over an angle of 1°; camera ones use a much bigger angle. To some extend you can overcome this by using a long focal length.

[7] In reality light consists of wave packets and cannot be focused down to a mathematical point. In any case a physical plane will consist of atoms and these are not points. Of course this makes no real difference to us as photographers since the scales involved are considerably smaller than anything we can perceive directly. However you might come across articles making absurd claims such as "there is no such thing as depth of field" based on an over literal application of the idealized mathematical model, ignore them.

is in focus if it is reproduced as a point on the projection plane (i.e., our film or sensor). Points out of focus are reproduced as discs, called *circles of confusion*. Human vision can only resolve to a certain level, so if the discs are small enough they are perceived as being in focus. How small do these discs have to be? Well that depends on what enlargement we intend to make of our image and of course the viewing distance. There are standard ways of making these calculations but we will not go into details, they are not difficult but a little involved[8]. Clearly the bigger our original the bigger the circles of confusion can be since we need less magnification to achieve any given size of print.

The points that are in focus all lie on a plane of the scene. Assuming that the image plane (i.e., film or sensor) is parallel to the lens plane then the plane of focused points is also parallel to them. Points in front of the plane of focus are focused behind the film or sensor and points behind it are focused in front. Of course points very close to the plane of focus will have very small circles of confusion and so will be acceptably focused. Thus the plane of focus is, for practical purposes, more of a slab, i.e., it has some thickness. The thickness is called *depth of field*. To understand how stopping down increases the depth of field consider Figure 3 where the aperture is wide open and the diameter of the circle of confusion is shown as a thick line on the film or sensor. The circle of confusion is created by a cone of light whose apex is in front of the film or sensor; for simplicity we are examining a point in the scene along the axis of the lens but the principle is the same no matter what point we use. Now let's stop the lens down, see Figure 4. As you can see the diaphragm is stopping the outer parts of the cone and allowing in only a part of it. This results in a smaller circle of confusion. The narrower the hole in the diaphragm, the smaller is the circle of confusion[9].

Let's now look at how far depth of field can extend in front of and behind the focused plane. First of all let's recall a simple fact: as points go further and further away (towards infinity) they are focused on a screen that gets closer and closer to the focal point of the lens (on the other side of the lens of course). Conversely as points in the scene get closer to the focal point so the image plane on the other side goes further and further away (towards infinity); there is a symmetry here. Now let's fix the position of the image plane (film or sensor). If we take a point on the focused plane and move it towards the lens it becomes increasingly defocused (larger and larger circle of confusion). On the other hand if we take the same point and move it further away (towards infinity) it again becomes increasingly defocused. Any given size of circle of confusion can be achieved by these two experiments. However for movement towards the lens we just have a finite distance; from

[8] A fairly reliable way of calculating the size of an acceptable circle of confusion is to divide the length of the diagonal of the film by 1730. For 35mm film this about 43.27/1730 which is around 0.025mm.

[9] There is a limit to this process due the the phenomenon of *diffraction*. Whenever waves hit an edge they are bent and spread, setting up interference. This is normally negligible but becomes increasingly significant as waves are passed though sufficiently small openings. This is the reason that pinhole camera images exhibit a certain softness that cannot be improved by making the pinhole smaller; that just makes them worse.

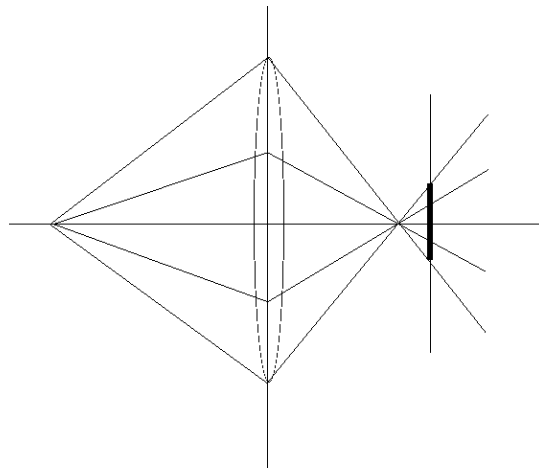

Figure 3: A point behind the plane of focus; lens aperture full.

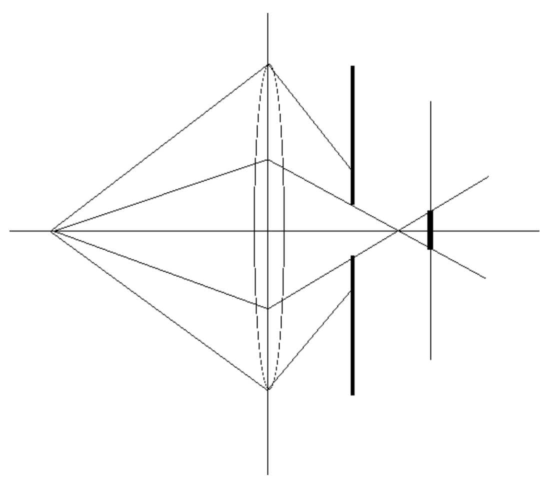

Figure 4: A point behind the plane of focus; lens aperture stopped down.

the focused plane to the focal point. For movement away from the lens we have an infinite distance. Thus moving towards the lens defocuses points at a faster rate than moving away[10]. This argument suggests that depth of field extends much further behind the plane of focus than the front and it can be verified by more precise geometry or algebra. For sufficiently distant scenes around 2/3 of the depth of field is behind the plane of focus and the other 1/3 is in front. This rule of thumb becomes more and more incorrect as we focus on closer and closer subjects. Of course the depth of field depends on the aperture chosen, becoming larger and larger as we stop down. Older lenses used to have a scale on the front showing the depth of field. Each f-number was shown either side of the focusing mark so after focusing the distances against the two chosen f-numbers gave the closest and furthest planes in acceptable focus. Progress has seen to it that this is now quite rare (you will find it on professional prime lenses).

For most of the time in landscape photography we want maximum depth of field. This suggests using a small aperture, which is fine. However for any choice of aperture there is a simple way to maximize depth of field. Imagine setting your focus very close, this will mean that parts of the scene that are very far away are likely to be out of focus while depth of field means that there are parts of the scene in front of and behind the chosen plane of focus that are acceptably focused. Now focus further away; as this is done an increasing amount of the scene behind the plane of focus is within the depth of field. At a certain distance this depth of field will extend to infinity, this is called the *hyperfocal distance*. It can be shown by straightforward mathematics that depth of field extends in front of the plane of focus by half the hyperfocal distance.

Going into further detail, the hyperfocal distance depends on the focal length of the lens, the f-number and the diameter of the largest acceptable circle of confusion; the formula is

$$\text{Hyperfocal distance} = \frac{(\text{Focal length})^2}{\text{f-number} \times (\text{diameter of circle of confusion})}.$$

There is no point in trying to use this formula in the field but it is worthwhile understanding its implications. Firstly the diameter of the circle of confusion is fixed for a given format so that leaves two variables: focal length and f-number. Note that each doubling of focal length results in multiplying the hyperfocal distance by 4 (it goes much further away from us). On the other hand, each doubling of the f-number results in halving the hyperfocal distance (note that doubling the f-number means an increase of 2 stops). So, assuming the 35mm format, with a 28mm wide angle lens we can focus from very close (just over 1 meter) to infinity even at f/11.

A frequently repeated piece of advice for landscape photography is to use a small aperture (large f-number remember) and set the lens to the hyperfocal distance for maximum

[10]For mathematicians we should note that the correctness of this observation relies, amongst other things, on continuity.

sharpness[11]. Nothing wrong with this *unless* your image does not include anything particularly close. If that is the case then the sharpest part of your image is nothing but air which doesn't really register. So the advice is fine if (as is often the case) you have foreground very close to you and a beautiful scene going off into the distance. Otherwise you are better off focussing about one third of the way into your scene. Even this must be modified if there are parts that must be absolutely sharp (e.g., a craggy silhouette).

All well and good but how do you find the hyperfocal distance? If you are using an old fashioned lens with depth of field indicated on the barrel this is easy. Rotate the focus ring so that the infinity symbol (∞) is against your chosen f-number on that side of the barrel; the second occurrence of this f-number will now show you the closest distance in focus[12]. If your lens does not have the markings you can find tables of hyperfocal distances on the web for most formats. As observed above, if you do not have anything very close to your lens you can focus about one third of the way in. For close things though accurate setting is very useful (think of a large boulder acting as the main foreground of your composition fairly close to your wide angle lens). Trying to guess in such cases is not easy, for example in the 35mm format with a 24mm lens at f/22 the hyperfocal distance is just under 1 meter which means you get sharpness from around 50cm to infinity. Of course if the boulder really is very close to you then focusing on it (with a small aperture) will be about right, but I like to be certain of the outcome whenever that is feasible. I should add that large format photographers rarely concern themselves with these things because the ability to tilt the lens provides a much better way to get sharpness all the way.

Before we leave this section we should address one common confusion. Look at Figure 3 and move the film or sensor towards the lens an equal distance to the other side of the point of focus. What is the size of the circle of confusion? Exactly the same of course (the twin cone is symmetrical about the apex). So, if the circle represents acceptable focus, what this experiment shows is that we have as much distance in front of as behind the true position of focus to get the film or sensor misplaced and still get acceptable sharpness (not that we would do this on purpose, it just shows the available margin of error). This distance is called *depth of focus* and should not be confused with depth of field (not that this stops even very eminent photographers from doing so).

[11] I am starting to question the received wisdom that everything should be sharp which presumably has its origins in the American tradition, e.g., the f/64 group. It is the safest way to proceed for those starting out but many great photographers such as Hill and Adamson or Sutcliffe, all of whom pre-dated the American school, did not feel so hide bound. It must be said though that it is considerably harder to make good photographs with this approach maybe because of the prevailing party line. In particular Sutcliffe often used misty conditions with the foreground sharply focused and the background somewhat out of focus; of course this appears perfectly natural.

[12] In some cases these scales are rather optimistic so many photographers err on the side of caution and set the lens to one number higher than the scale; experience will show if this is necessary for your lenses.

3.3 Choosing your shutter speed

For hand held photography we need a shutter speed that is fast enough to avoid camera shake (unless that is the intention). However in landscape photography fast shutter speeds are rarely used. One practical reason for this is that for maximum quality we use the lowest ISO possible. Furthermore we use small apertures for good depth of field. To compound all this we usually work when the light is not particularly strong, e.g., early morning or late evening. Put these factors together and shutter speeds running into seconds are not unusual.

In fact even if a fast shutter speed is available we might well choose a slow one for expressiveness. If photographing water, a shutter speed from $1/2''$ and slower is very effective. If we have trees or grass in a breeze we might well choose shutter speed of several seconds so that the energy of the scene is captured; contrasting with the immobile parts. Think about the resulting overall shapes and flows that you want to have on your photograph, frozen movement is unlikely to be the best choice. This is not an exact science but a little practice goes a long way.

3.4 ND graduate filters

We will now pick up once again the concerns of §3.1 on controlling contrast. Suppose your chosen composition has a glen, a perfectly placed meandering stream and ideally placed mountains. The sun is low from the side giving great modelling and even the clouds in the sky are superb. You meter for the land and everything of importance is within 4 stops, as is quite likely. But then you meter for the sky and find that it is much too bright thanks to those magnificent clouds. What can you do? Well one thing you can do is follow 19th century practice (really) and make two exposures: one for the land, one for the sky and combine them later in your print (called "printing in" back in the old days). All right, nowadays you might go high tech and use photoshop but otherwise this is nothing new. An alternative is to decrease the brightness of the sky using a *neutral density graduate* filter (say *ND grad* if you don't want to sound ignorant). These indispensable filters start with a dark part that gradually gets lighter till around the middle when they become clear (either very quickly for *hard grads* or more slowly for *soft grads*.) Figure 5 illustrates a hard and a soft grad next to each other[13].

Filters are made to various densities so that we can take off illumination from 1 stop to as much as 6 stops or even more (by combining one or more filters) in steps of half a stop (from 1 stop upwards). To use the filters you need a proprietary holder and a step ring that screws into the front filter thread of your lens (regrettably most compact cameras do not allow this now even when the lens does not retract). I have sets of both hard and soft filters but I use hard step ones for most of the time. The idea is that hard step ones are used for situations where the transition zone (usually the horizon) is fairly straight while

[13] These illustrations are used with permission from `http://www.robertwhite.co.uk`.

Figure 5: A Lee hard grad filter on the left and the equivalent soft grad on the right.

soft ones are used for other situations. However I find that in the conditions I photograph I need the filter to stay dense quite close to the transition zone. As a result sometimes a mountain top is made fairly dark, but in fact this is just right because when we look the brightness of the sky causes our pupils to contract so we perceive the top as dark. Soft step filters do of course have their uses, e.g., with gradual reflections on wet sand at a beech with a bright sky. Sometimes I have combined a hard step filter with a soft one.

Suppose you have decided to take off two stops of light from the sky. How do you place the filter? If you have no way of viewing the actual image (e.g., through the viewfinder) then you can estimate what proportion of your picture is given to the sky and put the filter transition zone correspondingly down. In all likelihood you will need to tilt the filter as well if there are mountains (high on one side and low on the other). By far the best method is to look at the image while you move the filter up and down to place it. This task is made much easier if you stop the lens down (i.e., cause the diaphragm to close down thus dimming the image). Watch the zone move up or down, it is almost impossible to tell where it is if you just place the filter and then look (just as well). As you would expect precise placement of a hard step filter is easier than for a soft step one.

Returning to the discussion in §3.1 on controlling the over bright part of a mountain we now have the alternative solution. We can cover that part with a grad that takes off 1 stop. This is likely to reduce the brightness of other parts of the scene but with a bit of luck this is acceptable.

Finally in this section let's look at how the strength (light reduction) of ND grads is indicated. A grad that reduces light by one stop is described as .3 and one that reduces it by one and a half stops as .45, thus one that reduces light by three stops is .9. Where do these numbers come form? They give us the *density* of the strongest part of the grad. A grad that cuts out one stop is twice as dense as a clear one. Traditionally density is measured using logarithms (to base 10) and $\log_{10}(2)$ is approximately .3 hence the designation[14]. By the way this explains why in the DIN scale 3 was added for each doubling of sensitivity (the scale was multiplied by 10 to avoid decimals).

3.5 Composition

A poorly composed landscape photograph is never satisfying (the situation is of course different in documentary or personal family albums). But what constitutes good composition? This is a question to which there is no formulaic answer and indeed what is considered good practice changes with time. For example in the late 19th century painters were influenced by Japanese art and employed compositions that would traditionally have been rejected in western art. This section can only scratch the surface but to leave it out might give the impression that good landscape photography resides in technical ability. Such a notion makes no more sense than believing that the artistic quality of recorded music depends on the quality the HI-FI used!

Despite the intractability of the topic there are some simple steps you can take when composing a photograph that are very effective. First of all think deeply about the aim of your photograph and make sure that your composition conveys this aim. Exclude anything that confuses the issue, e.g., highlights that draw the eye away from the aim and to nothing of any importance. When I examine a composition (on the ground glass in my case) I ask myself why anything of any significance is included and if there is no good explanation then the composition needs rethinking. Avoid the trap of including something just because it looks good, in fact this might be the very reason for excluding it because it detracts attention from the main point.

I have seen many published photographs that fall apart just because of an overlooked distraction. If you look at an example that makes you feel uneasy for some reason try covering the appropriate area with your thumb and see the difference. The point here is that your eyes are well used to your thumb and will not be drawn to it unnecessarily, unlike the unwisely included pool, rock or whatever.

Make a practice of looking past the immediate "wow" factor (amazing tree, rock, mountain etc.) and consider the underlying shapes. Is the mass of the rock you are using in the foreground truly balancing and echoing the distant mountain or is it actually unbalancing the composition? Are the overall shapes organized to lead the eye into the frame or is the eye given contradictory signals? On the subject of lead in lines, a common mistake is to establish strong lines that lead the eye to nothing of any importance. An excellent way

[14]Strictly speaking the density factor is $\log_{10}(1/2)$ but this is just $-\log_{10}(2)$ and to avoid tedium the negation is dropped.

of separating composition from content is to look at it upside down (a common practice amongst artists). Large format photographers have an advantage here since the image is upside down (and laterally reversed) on the ground glass but of course you can examine prints upside down and you'll soon learn what to look for.

Another pitfall is to include as much as possible, e.g., putting in all of a mountain range or peak. This *can* be good if it creates an overall satisfying set of shapes but quite often it leaves nothing for the viewer's imagination. Establishing a strong theme and then leaving the rest to suggestion and imagination is a very powerful technique used in all the arts, including music. Try including only 3/4 of a high mountain peak with the slope leading the eye into the frame at, e.g., a stream meandering into the scene or a waterfall. You'd need to balance the other side of the photograph, perhaps with a lower peak taking up the movement and essentially resolving it. After this, recompose to include all of the high peak and watch the composition fall apart.

Look out for large dead areas, usually caused by the pitfall of trying to include everything, e.g., a low broad waterfall. Bear in mind that when we look at the actual scene our eyes are continually scanning it with only a relatively small area being seen in detail. This way we keep our attention on the action and essentially ignore the dead areas. When we look at a photograph though we scan the whole thing and the dead areas are only too apparent. Learn to look photographically, your photograph will not be viewed in the same way as the actual scene.

It is a good idea to avoid visual complexity. Busy scenes can be made to work but they are considerably more difficult to handle and should only be attempted for good reason. Even in a fairly sparse scene care has to be taken. Look at the photograph in Figure 6 of Carnethy Hill (from near the top of Turnhouse Hill) in the Pentlands just outside Edinburgh. This is the second version made just over a year after the first. The reason is that for my first version I concentrated almost completely on getting the foreground and main silhouetted tree on the left of the frame just right. As a result I overlooked the fact that the more distant tree was not placed at the precise intersection point of the two slopes but a fraction to the left. This was unacceptable to me for two reasons: in the first place the misplacement created an extra intersection (more complexity) for no good reason. Secondly the eye is naturally led to the intersection point but with the incompetent version it then had to shift left a bit to take up the distant tree, just terrible. By the time I got the processed transparency back the snow had melted hence the wait (as it happens I got a better sky the second time). In any case this was a lesson well learnt and I now examine all intersection points minutely.

On a similar vein watch out for any disturbing occlusions, e.g., a tree trunk near you partly occluding another one a little further away. Very often all that is needed is a small change of position to the side to achieve separation and a much more satisfying composition. More subtly watch out for objects that have similar colour and tone that can appear merged in your photograph. For example look at Figure 7. This photograph took many visits to get right. The central problem is the fact that the top of the main pine tree is black (it

Figure 6: Carnethy Hill.

is dead) and the main face of the mountain is also very dark (snow doesn't seem to settle there). It might not be completely clear from the reproduction but the light, which is at an angle and behind the tree, gives its top a luminance that separates it from the mountain.

Always remember that a photograph is flat so if you want to convey depth you need to employ the familiar devices of lead in lines and recession (perspective and/or tonal). When looking at a scene it is often helpful to close one eye. This removes the depth information we gain from binocular vision. It isn't a perfect way to judge things since we cannot alter the focal length of our eyes, but it is a very effective learning aid. Of course at times you might want to convey flatness but this had better be deliberate rather than an oversight.

The danger of unwanted highlights was mentioned above and problems with generally bright areas were discussed in §3.1. In examining a possible composition it is good practice to half close your eyes. Any problem areas caused by contrast will be easier to spot. Once you have set up it is also a good idea to see the image dimmed (this is a more precise version of half closing the eyes, we look only at the composition). For this you need to stop down the diaphragm and close it; most lenses leave the diaphragm wide open till exposure so that you have a bright viewfinder image (normally a good idea). Unfortunately there is an increasing trend towards leaving this capability out. There is also a widespread misunderstanding that the main use of this tool is to judge sharpness, how that can be

Figure 7: Glen Etive with An Grianan.

done with a dim image and a relatively small viewfinder image (on smaller formats) is never explained. In any case we now have two reasons for this very useful function: examining the contrast range and placing grads. In some cases the problems of contrast are too great in which case returning at a better time is the only solution.

In this discussion we should not forget the role of the lens, or rather its focal length (relative to the chosen format). This is one of photography's most significant contributions to how we are enabled to view the world. Choose a main subject, let's say a rock. Now half fill your fame with it using a wide angle lens, then do the same with a standard one followed by a moderate telephoto. Watch how the surrounding context of your subject changes. Look also at the effect on the shape of your subject. To put it briefly, wide angle lenses allow us to expand perspective while telephoto lenses compress it. The temptation to fill the frame by zooming in, or to get a wide vista by zooming out, all from the same position should be resisted at all costs.

Finally, your compositions will improve if you make a habit of examining all forms of visual representation and working out why something is satisfying or not. For example look at the astonishing painting L'absinthe by Edgar Degas and the way that the picture space is organized (a web search will take you to various places with digital versions).

4 Equipment

4.1 Cameras

Digital or film? Couldn't care less. What does it lets you do is the point. Does it give you full control without driving you round the bend or does it assume that it knows best and all you have to do is point and shoot[15]? We have seen above that for landscape photography it is best to have:

1. Easy control over exposure. This means being able to set the aperture and shutter speed conveniently and quickly.

2. Spot meter function.

3. Control over focus, auto focus is pointless here.

4. Easy way to make exposures over a long time (so some form of cable release is needed).

5. Ability to attach filters reliably.

6. Ability to stop the lens diaphragm down.

None of these dictates an expensive camera[16]. Of course you also need good lenses to go with it. Zoom lenses are OK provided you resist the temptation to stand in one place as discussed in §3.5. In some cases the front element of the lens rotates as you focus, this is best avoided as it makes it more difficult to use filters. When I started in landscape photography I used a 28-135mm zoom lens with my 35mm SLR. It didn't take me long to decide that my work would improve with the use of prime (fixed focal length) lenses. Switching to them was a liberation for me but it might not be so for you, the point is to think about how you work and what helps or impedes you. Naturally financial considerations are also a factor. Second hand equipment can be an excellent purchase but you do need to be sure that it is all in good order (a knowledgeable friend or reliable dealer are obviously helpful). In the intervening years since I first wrote this article I have bought some lenses on ebay and have not been disappointed. Clearly some care has to be taken, a Leica 50mm f/2 Summicron in good condition is never going to be cheap.

So if you have an automatic compact digital or film camera do you have to discard it for landscapes? No, is the short answer. If you cannot afford extra equipment as described above you can still derive a great deal of enjoyment and make good photographs. The

[15] I must stress that I am not against automation or even the point and shoot approach. It has allowed very many people to record their memories, either of loved ones or places. The problem comes in when this is pushed as the best approach for all situations.

[16] For SLR users we could also add the ability to put the mirror up before exposure in order to reduce vibrations. The effect is at its most pronounced for exposures of around $1/4''$ to say around $1/8''$, for long exposures there is relatively little effect but obviously it is best to get rid of vibrations altogether.

points explained here should help you to improve your work and also see why you cannot achieve certain results. It is tempting to say that good photographs are not a function of the equipment but to be honest if your camera limits you to a moderate wide angle (as happens with a lot of compacts) then there are things you cannot do[17]. All the same there is plenty you *can* do. Even if you have, or at some point buy, the extra equipment your compact camera comes in very handy as a visual notebook for ideas. I must admit that I do not myself do this, even though I own a compact digital camera. I did try it but found it a complete turn off, still many do find it very useful so give it a try. What you should *not* do is buy a lot of expensive equipment in one go at least not till you are really sure of what will work for you.

4.2 Tripods

Landscape photography is possible without a tripod but many opportunities are closed off. We have discussed the factors that affect our choice of aperture and shutter speed. These often dictate a slow speed. Even when a fast speed is available we are likely to prefer a slow one, as discussed in §3.3. Furthermore the use of a tripod helps with precise composition; see also the discussion in §3.1.

Advice on tripods is often phrased in terms of their weight. However this is not really the issue. The important thing is that the leg sections should be tightly coupled and dissipate vibrations rather than amplify them. A sturdy tripod is more likely to do this but a poor heavy tripod is worse than a high quality light one. If this was not the case then carbon fibre tripods would never have succeeded. A good tripod will not have a head with it since the correct one will depend on intended usage. For landscape photography geared heads (such as the Manfrotto 410) are a very good choice because they provide very precise control. Many use a ball head, some of which cost a fortune, but I prefer the ability to control the head in three independent planes since this makes minute precise adjustments very easy.

Another factor to bear in mind is how low you can set the tripod (in a fuss free fashion). This is very useful if there is a strong wind or if you want to get a dramatic low angle or a close up shot. At the other extreme think about how far you can extend the legs. If you intend to photograph in very uneven terrain (my favourite) the longer the better. This is useful, e.g., on a very steep bank. Of course this is only worth having if the tripod remains sturdy at full extension. It is clear from all these considerations that a centre column is not necessary and indeed undesirable for landscape photography.

Unfortunately the very best tripods and heads are also very expensive but there are

[17] There is a deplorable trend to advise beginners that for a wider view they can just stitch images together (another 19th century practice by the way). There are at least two problems here: firstly the perspective is not the same as that of a single shot made with a wider angle lens, i.e., it is a different interpretation not an equivalent one as is implied. Secondly, wide landscapes are often at their best in fast changing light and stitching will not work with this, at least not convincingly.

good choices for more moderate prices. Before buying, set up the tripod at full extension and check how sturdy it is. The best tripods just have three independently moving legs and a good collar for the head. Carbon fibre reduces the weight but pushes up the price, nice if you can afford it but not essential.

Finally, a small spirit level is invaluable in keeping horizons level. Good heads will have one or more spirit levels (mine has three, one for each plane) but they can be difficult to use. The simplest solution is to use a small spirit level on the camera (these are sold by photographic shops) and adjust the head appropriately for my 10×8 camera I use a small carpenter's level bought from B&Q for very little money.

4.3 Filters

We have already discussed the use of ND graduate filters in detail. These are the most essential ones for landscape work. Others that are worth having are the 81 series warm up ones. There are several of these with increasing effect. A good first purchase is the 81b. This is strong enough to make a difference but not so strong as to be overly imposing. The stronger ones have specific uses, especially in correcting colour shifts with some transparency film or for conditions when the light is very bluish. It is true that you can add warmth later (e.g., in photoshop) but I prefer to visualize my end result at the time of making it. One undesirable feature of the 81 series is that if used with a blue sky they can make it go murky. This is not really a problem with 81b or weaker (81, 81a) but is best avoided for stronger ones (81c, 81d, 81ef).

A very popular and often misused filter is the polarizer. This can be used to eliminate reflections or saturate the blue colour of the sky. The problem is that if overdone this can turn the sky almost black (you rotate the filter to alter the effect). Furthermore the polarization effect is not even over the sky so that with wide angle lenses you can end up with a very uneven and unnatural result. If the sky is largely clear of clouds then don't use this filter. If it has some clouds that you want to stand out it can be a good choice but use it in moderation. The filter works most strongly when the sun is at right angles to it and not at all if it is aligned with the filter. Another, and perhaps better use, is to eliminate reflections from scattered light in order to saturate the colour of autumn leaves etc. I do own a polarizer but hardly ever use it. You will see references to circular or linear filters. This does not refer to their shape but rather the way in which they achieve the effect: visually there is no difference but auto focus cameras require the circular type.

There are very many other filters but most are not desirable for landscape work, best to avoid any "special effects" abominations. There are various useful ones but they are rather specialized and will not be discussed here with the exception of coral grads. These provide an alternative way of adding warmth to the landscape. Whereas the 81 series are amber coloured these are reddish. As their name suggests they have a coloured portion that fades to clear. They are used inverted (colour portion at the bottom) so that the landscape colours are shifted towards the red but the sky is unaltered. Care has to be taken not to overlap the coloured part with the sky. Used with care they can be extremely effective especially when the sky has some blue and the land is predominately reddish brown. They should not be used if there is clear water (e.g., a stream) in the landscape, reddish water seems a bit odd. Like the 81 series coral ones come in various strengths but for landscape work no. 1 and no. 2 are the most useful.

Filters are made by various manufacturers with widely varying prices. Amongst the very best are those from Lee Filters and, as you have already guessed, they are not the cheapest. You do get what you pay for though, e.g., cheaper ND grads are not long enough to allow arbitrary placement of the transition zone and tend to have a colour cast.

Above all remember the golden rule for using filters. They should be used for subtle enhancement rather than wholesale change.

Finally always shade your filters if there is any trace of direct light. Even if the light is from behind it can strike the edge of a filter and be refracted down to cause flare. A common practice is for the photographer to stand between the sun and the filters or use the palm of the hand. You can also buy special holders that can be useful under various conditions. Working in the field you are bound to get small specks of dirt on the filters. This is not necessarily a problem but always start your session with clean filters (there are various liquids and soft lint free cloths for this).

5 The weather etc.

Perhaps the worst advice I have ever seen is to check the weather forecast before going out to avoid disappointment. I ignore the weather forecast except for when intending to climb a mountain when I must know the conditions to be properly prepared for safety (for excellent forecasts visit `http://www.mwis.org.uk`). Some of my very best photographs have been made in drizzle or even heavy rain. Look across a vista on a day of soft rain and marvel at the astonishing recession as visibility diminishes gently (of course some rainy days are just awful, but there are surprisingly few of these, even in Scotland!). There is much more to landscape photography than sunrises and sunsets. Rather than thinking about great light, turn the question round: given *this* light what would be great to do with it. The most depressing situation for me is a day of unrelieved sunshine in spring and especially summer; even days like this have their uses (e.g., location searching or just keeping fit).

Naturally if you are wet, cold and miserable you will not enjoy yourself and more importantly your photography will suffer. Thus you need the best outdoor clothing you can afford. Wickable fabrics move perspiration away from your body so that it evaporates. Very good quality waterproofs (e.g., Gore-Tex XCR) mean that you can stay dry and comfortable even in a downpour. You do need to check how waterproof your camera is especially if it relies on electricity (I use a wooden one, trees don't mind getting wet). Good boots are essential. Wellington boots are also very handy either for coastal work or standing in the middle of a river. A good walking pole is very handy especially on uneven terrain, many a fall has been averted by my one!

If you intend to work on the coast then make sure you know the tide times. For the UK you can get free 7 day predictions from the United Kingdom Hydrographic Office at `http://easytide.ukho.gov.uk/EasyTide` and there are various apps also available (good ones are not free but not expensive). Take great care on the coast that you do not get sand in your camera or lenses, this can be very destructive. Also ensure that you get rid of any salty water as soon as possible after a session.

There are sites for sunrise and sunset times and position (or you can buy a special sun compass). I have never used these, preferring to get to know my location well enough to work this information out for myself. Admittedly if you only have a short time in a place a special sun compass can be handy (I now have a free one from a magazine). A great alternative is TPE (The Photographer's Ephemeris). The online version is free while the apps cost a little; I have now invested in the iOS version.

6 Bibliography

Studying the work of others is a great source of inspiration. Take care however to draw a distinction between the work that you admire and the work that is best for you to do. The strongest work comes from an inner urge coupled with determination; if it is just an imitation then it shows. There is a lot you can learn by trying to reproduce a well known composition but this is best viewed as a technical exercise, just as art students make copies of well known works as part of their studies.

Anyway here is a short list of some of the books that I have enjoyed. Probably they will not all be to your taste but I hope that some of them will give you a start. There is also a short list of books that I have found infuriating but best not to go into details here.

1. John Blakemore, *John Blakemore's Black and White Photography Workshop*, David and Charles, 2005.

2. Joe Cornish, *First Light*, Argentum, 2002.

3. Joe Cornish, *Scotland's Coast*, Aurum Press, 2005.

4. Peter Dombrovskis and Jamie Kirkpatrick, *In the Forrest*, West Wind Press, 2001. (Difficult to get, any books by Peter Dombrovskis are great if you can find them.)

5. E.H. Gombrich, *The Story of Art*, Phaidon, 1995.

6. E.H. Gombrich, *Art & Illusion*, Phaidon, 1991.

7. E.H. Gombrich, *The image & the Eye*, Phaidon, 2002.

8. Michael Hiley, *Frank Sutcliffe, Photographer of Whitby*, Phillimore, 2005.

9. Richard Oveden, *John Thompson*, National Library of Scotland, 1997.
10. Sara Stevenson, *The personal Art of David Octavius Hill*, Yale University Press, 2002.
11. David Ward, *Landscape Within*, Argentum, 2004.
12. Charlie Waite, *The Making of Landscape Photographs*, Collins and Brown, 1992.