

Chapter 1

The DSLR Revolution

A few years ago, I said that if somebody would manufacture a digital SLR camera (DSLR) that would sell for under \$1000 and would work as well as film for astrophotography, I'd have to buy one.

That happened in 2004. The Canon Digital Rebel and Nikon D70 took the world by storm, not only for daytime photography but also for astronomy. Within two years, many other low-cost DSLRs appeared on the market, and film astrophotographers switched to DSLRs en masse.

There had been DSLRs since 1995 or so, but Canon's and Nikon's 2004 models were the first that worked well for astronomical photography. Earlier digital cameras produced noisy, speckled images in long exposures of celestial objects. Current DSLRs work so well that, for non-critical work, you almost don't need any digital image processing at all — just use the picture as it comes out of the camera (Fig. 1.1). The results aren't perfect, but they're better than we often got with film.

As you move past the beginner stage, you can do just as much computer control and image enhancement with a DSLR as with an astronomical CCD camera. Some hobbyists bring a laptop computer into the field and run their DSLR under continuous computer control. Others, including me, prefer to use the camera without a computer and do all the computer work indoors later.

1.2. CHOOSING A DSLR

1.2.2 Shopping strategy

Because of rapid technological progress, you generally want the *newest* DSLR that works well for astrophotography, not the most ruggedly built one. It's better to buy a low-end DSLR today and another one in three years with a new, improved sensor, rather than sink all your money into a professional-grade camera that will commit you to using today's technology for a decade.

Of course, if you can justify the expense for other reasons, go ahead and enjoy your Canon EOS 5D or Nikon D200; these cameras have big, bright viewfinders and are a joy to use. Apart from price, one disadvantage of pro-grade DSLRs is that they are heavy enough to unbalance a medium-sized amateur telescope. Another is that pro-grade cameras are more complicated to operate, and that can be a problem in the dark. Generally, professional cameras are designed for people who use them all the time and can easily remember a large number of controls. Entry-level cameras with simpler controls are easier to use, even for advanced work, as long as they have the features needed.

You don't have to have the best camera on the market in order to get good pictures. You just have to have a camera that is *good enough*. All astronomical instruments, including big observatory telescopes, have measurable limitations. We work so close to the limits of the laws of physics that perfection is unobtainable.

Also, buying the *very* newest camera has some drawbacks. This week's hot new DSLR may not yet be supported by your software, although updates usually come quickly. It also may not have fully debugged firmware (the software inside the camera); watch the manufacturer's web site for firmware upgrades.

One last warning. Unreliable camera vendors are common on the Internet, and they often advertise impossibly low prices. Before dealing with a stranger, do some searching and find out whether others have had good experiences with the same merchant. You can find out the reputation of a vendor from www.epinions.com and www.resellerratings.com. Remember that even the best large-volume dealer has a few unsatisfied customers, and a "perfect" score may mean simply that not many customers have been surveyed.

Highly reliable camera dealers include B&H in New York

2.1. IMAGE FILES

with large DSLR images. Regardless of what software you use, you'll need a computer with ample RAM (at least 1 to 2 GB recommended) and an up-to-date operating system. Smaller astronomical CCDs work well with older, smaller laptops, but DSLRs do not.

2.1.2 Raw vs. compressed files

Most digital photographs are saved in JPEG compressed format. For astronomy, JPEG is not ideal because it is a "lossy" form of compression; low-contrast detail is discarded, and that may be exactly the detail you want to preserve and enhance. What's more, operations such as dark-frame subtraction rely on images that have not been disrupted by lossy compression.

For that reason, astrophotographers normally set the camera to produce raw images which record exactly the bits recorded by the image sensor (or nearly so; some in-camera corrections are performed before the raw image is saved).

The term *raw* is not an abbreviation and need not be written in all capital letters; it simply means "uncooked" (unprocessed). Filename extensions for raw images include .CRW and .CR2 (Canon Raw) and .NEF (Nikon Electronic Format). Adobe (the maker of *Photoshop*) has proposed a standard raw format called .DNG (Digital Negative) and has started distributing a free software tool to convert other raw formats into it.

Raw images are compressed — their size varies with the complexity of the image — but the compression is lossless or nearly so; the exact value of every pixel in the original image is recovered when the file is decoded in the computer. Canon and Nikon raw images occupy about 1 megabyte per megapixel, varying with the complexity of the image. Uncompressed, a 12-bit-deep color digital image would occupy 4.5 megabytes per megapixel.

2.1.3 "Digital film" and camera software

DSLRs record images on flash memory cards, sometimes called "digital film." Unlike real film, the choice of "digital film" doesn't affect the picture quality at all; flash cards differ in capacity and speed, but they all record exactly the same data and are error-checked during writing and reading.

3.1. TAKING A PICTURE MANUALLY




Figure 3.1: Display panel of a Canon XTi (400D) set up for deep-sky photography. *F00* means the lens is not electronic and the camera cannot determine the *f*-stop. *C.Fn* means custom functions are set.



Figure 3.2: For manual control of shutter and aperture, set mode dial to M.

3.4. MENU SETTINGS

High ISO NR: Normal. This is a Nikon setting that apparently improves picture quality but does not affect the “star eater.”

Magnified View: Image Review and Playback. This is a Canon XTi (400D) setting that makes it easier for you evaluate focus right after taking a picture. With this option enabled, you don't have to press \triangleright (“Play”) to magnify the picture. You can press “Print” and “Magnify” together to start magnifying the picture while it is still being displayed for review.

3.4.2 Settings for an astrophotography session

Picture Quality: Raw + JPEG. If the memory card has room, I like to have the camera save each picture in both formats. Each file serves as a backup of the other, and the JPEG file contains exposure data that can be read by numerous software packages. This takes only 20% more space than storing raw images alone.

Review: Off. At an observing session with other people, you'll want to minimize the amount of light emitted by your camera.

LCD Brightness: Low. In the dark, even the lowest setting will seem very bright.

LCD Display When Power On: Off. On the Canon XTi (400D) and similar cameras that use their LCD screen to display camera settings, you probably don't want the display shining continuously at night. Instead, switch it on with the DISP button when you want to see it.

Auto Power Off: Never. You don't want the camera turning itself off during lulls in the action; at least, I don't.

Long Exposure Noise Reduction: Your decision. If turned on, this feature will eliminate the hot pixels in your images by taking a dark frame immediately after each picture and automatically subtracting it. This takes

5.3. OPTICAL PARAMETERS




Figure 5.7: One way to use a Meade or Celestron focal reducer with a telescope that lacks a threaded rear cell, or with a 2-inch-diameter electric focuser.

5.3 Optical parameters

5.3.1 Focal length

The *focal length* of a telescope or camera lens is a parameter that determines the size of objects in the image (Fig. 5.8). In conjunction with the film or sensor size, the focal length also determines the field of view.²

If your telescope is a refractor or Newtonian, the focal length is also the length of the tube (the distance from the lens or mirror to the sensor or film). Technically, the focal length is the distance at which a simple lens forms an image of an infinitely distant object, such as a star.

²If you're new to optical diagrams, you may wonder why Fig. 5.8 shows rays of light spreading apart while Fig. 5.1 shows them coming together. The answer is that Fig. 5.1 shows two rays from the same point on a distant object, but Fig. 5.8 shows one ray from each of two points some distance apart.

7.5. LENS MOUNT ADAPTERS



Figure 7.8: Result of using crosshairs in Fig. 7.7. Canon Digital Rebel (300D) at ISO 400, Sigma 105-mm $f/2.8$ lens wide open. Stack of five 3-minute exposures of the North America Nebula, dark-frame subtracted, cropped.

7.5 Lens mount adapters

Besides the ubiquitous T-ring, originally designed for a line of cheap telephoto lenses, there are other adapters for putting one kind of lens onto another kind of camera body.

Obviously, any adapter has to contain a receptacle to fit the lens and a flange to fit the camera body. The trouble is, all of this takes up space, placing the lens farther from the film or sensor than it was designed to be. There are three ways to deal with this problem:

- Add a glass element (typically a $1.2\times$ Barlow lens) so that the lens will still focus correctly. This is a common practice but degrades the image quality.
- Leave out the glass element and give up the ability to focus on infinity. The adapter is now a short extension tube and the lens can

7.5. LENS MOUNT ADAPTERS

only take close-ups.

- If the camera body is shallower than the one the lens was designed for, then make the adapter exactly thick enough to take up the difference. In this case, the lens focuses perfectly, without a glass element.

The third kind of adapter is the only one that interests us, and it's only possible if the camera body is shallower, front to back, than the body the lens was designed for. In that case, the adapter makes up the difference. For example, a Nikon lens can fit onto a Canon EOS body with an adapter 2.5 mm thick.

Olympus Four Thirds System DSLRs are much smaller than 35-mm film cameras, almost any film SLR lens should be adaptable to fit them. The Canon EOS is one of the shallowest full-size SLR bodies, and adapters exist to put Nikon, Contax-Yashica, Olympus, M42, and even Exakta lenses on it. The one thing you *can't* do is put older Canon FD lenses on an EOS body, because the FD body was even shallower.

Nikon bodies generally can't take non-Nikon lenses because the Nikon body is one of the deepest in the business. Only the Leicaflex is deeper.

One curious zero-thickness adapter does exist. Pentax's adapter to put screw-mount lenses on K-mount cameras simply wraps around the screw-mount threads, taking advantage of the K-mount's larger diameter. More commonly, though, screw-mount-to-K-mount adapters are about 1 mm thick and do not preserve infinity focus.

7.5.1 Adapter quality

Not all lens mount adapters are equally well made. Novoflex (www.novoflex.com) markets top-quality adapters through camera stores, but the prices are high, presumably to allow the dealer a good markup on a slow-selling item. A respected supplier with more competitive prices is Fotodiox (www.fotodiox.com). On eBay you can buy adapters directly from the machinists in China who make them.

Fig. 7.9 shows what to look for. Good adapters are usually made of chrome-plated brass or bronze and often include some stainless steel. Some are made of high-grade aluminum. For higher prices you get more

8.2. LCD FOCUSING




Figure 8.3: Confirming focus on Canon 300D LCD screen. Compare star images to dot that indicates magnification (marked by arrow here).

8.2 LCD focusing

8.2.1 Confirmation by magnified playback


No matter how well you can focus optically, *you should always confirm the focus electronically*. The easiest way to do so is to play back the image on the LCD screen and view it with maximum magnification (Fig. 8.3). On a Canon DSLR, it is easy to compare the star images to the dot that indicates the portion of the picture being displayed. As you press the + button repeatedly, the dot gets smaller and the star images get bigger. If, at maximum magnification, the star images are still smaller than the dot, they're in focus.

Your test image can be a much shorter exposure than what you eventually plan to make. All you need is to capture a few stars. In most situations, a 5-second exposure is enough.

If your DSLR provides magnified live focusing, you don't even have to take a test exposure. Just aim the camera at a star and adjust the focus while watching the image. This feature was introduced on the Canon EOS 20Da (now discontinued) and is gradually spreading through the world of DSLRs.

9.1. TWO WAYS TO TRACK THE STARS

TRACKING WITH ALTAZIMUTH MOUNT Image rotates; long exposure photographs are not possible



TRACKING WITH EQUATORIAL MOUNT Telescope rotates with image; long exposures work as intended




Figure 9.2: Tracking the stars with an altazimuth mount causes field rotation, which can be overcome by taking very short exposures and doing a rotate-and-stack. (From *How to Use a Computerized Telescope.*)

9.2. THE RULES HAVE CHANGED

9.2 The rules have changed

The rules of the tracking and guiding game are not what they used to be. In the film era, the telescope had to track perfectly for 20 or 30 minutes at a time. Only equatorial mounts could be used because an altazimuth mount can only go a minute or two without excessive field rotation. Guiding corrections had to be made constantly, either by an autoguider or by a human being constantly watching a star and pressing buttons to keep it centered on the crosshairs. One slip and the whole exposure was ruined.

It was also important to guard against flexure and mirror shift. During a half-hour exposure, the telescope and its mount could bend appreciably. Also, notoriously, the movable mirror of a Schmidt-Cassegrain telescope would shift slightly. For both of these reasons, guiding was usually done by sampling an off-axis portion of the image through the main telescope.

Today, we commonly take 3- to 5-minute exposures and combine them digitally. That makes a big difference. Tracking errors that would be intolerable over half an hour are likely to be negligible. If there is a sudden jump, we can simply throw away one of the short exposures and combine the rest. We can even work with very short exposures on an altazimuth mount, using software to rotate as well as shift the images so that they combine properly.

9.3 Setting up an equatorial mount

9.3.1 Using a wedge

An equatorial mount is simply an altazimuth mount tilted toward the celestial pole. The typical fork-mounted amateur telescope becomes equatorial when mounted on a wedge of the proper inclination, pointed north (or south in the Southern Hemisphere).

Well-made wedges are sturdy, easy to adjust, and *heavy*. To save weight and to take advantage of the vibration-damping properties of wood, I built the wooden wedge shown in Fig. 9.4. Its obvious drawback is that there are no adjustments; the only way to adjust it is by moving the tripod legs or altering their length, and it only works at latitudes close to 34° north.

Would-be wedge builders should note that the mounting bolts of a Meade

10.2. CAMERA CONTROL




Figure 10.4: Pins on Canon N3 connector (used on 20D and higher models) have same function as Tip, Ring, and Sleeve on Digital Rebel series. Do not short Tip and Ring together; they have different voltages.

Digital Rebel, making it compatible with serial port cables, custom timing devices, and whatever other accessories you may have built.

Fig. 10.4 identifies the pins on the camera. You can trace the connections into the RS-80N3 and add a short piece of cable with a female phone jack to match the Digital Rebel, connected to the corresponding wires, so that whatever you plug in will operate in parallel with the existing switches. On the wires entering the RS-80N3, Tip is red, Ring is yellow or white, and Sleeve is the copper shield. These are the same colors used on stereo cables, so if you get your 2.5-mm phone jack by taking apart a stereo extension cable, you can simply match colors — but always verify the connections with an ohmmeter.

11.3. NEBULAE, RED RESPONSE, AND FILTER MODIFICATION

It takes about a million pixels to make a pleasing full-page picture. This means the output of a 10-megapixel DSLR can be binned 3×3 with pleasing results.

The easiest way to accomplish binning is simply to resize the image to 1/2 or 1/3 of its original linear size using *Photoshop* or another photo editing program. Many astronomical software packages offer binning as an explicit operation.

11.3 Nebulae, red response, and filter modification

11.3.1 DSLR spectral response

Figure 11.5 sums up a lot of information about the wavelengths of light to which cameras respond. Reading from the top, the first thing you'll notice is that color film has a gap in its response between green and blue, and DSLRs don't. This is important because the strong hydrogen-beta and oxygen-III lines from emission nebulae fall in that gap. That's one reason nebulae that look red on film may come out blue with a DSLR.

DSLR makers don't publish spectral response curves, so the curves at the top of Fig. 11.5 are estimated from Sony CCD data sheets plus a number of published tests. What is most uncertain is the effect of the infrared-blocking filter in front of the sensor. We know that with this filter removed, the sensor has strong response up to 700 nm or further.

11.3.2 Filter modification

There are two reasons astrophotographers often have their DSLRs modified for extended red response. One is the strong hydrogen-alpha ($H\alpha$) emission from nebulae at 656.3 nm. The other is that working in deep red light, with filters blocking out the rest of the spectrum, is a good way to overcome skyglow in cities.

Canon's EOS 20Da (now discontinued) and some Fuji DSLRs have been manufactured with extended red response. Usually, though, astrophotographers rely on third parties to modify their cameras. Reputable purveyors of this service include Hutech (www.hutech.com), Hap Griffin (www.hapg.org), and LifePixel (www.lifepixel.com). For those who

Chapter 12

Overview of image processing

This chapter will tell you how to start with raw image files from your camera, perform dark-frame correction, decode the color matrix, combine multiple images into one, and carry out final adjustments.

Vita brevis, ars longa. Digital image processing is a big subject, and I don't plan to cover all of it here. In particular, in this and the following chapters I'm going to skip almost all of the mathematics. To learn how the computations are actually done, see *Astrophotography for the Amateur* (1999), chapter 12, and other reference books listed on p. 263.

This is also not a software manual. For concreteness, I'm going to give some specific procedures for using *MaxDSLR* (including its big brother *MaxIm DL*) and, in the next chapter, Adobe *Photoshop*, but in general, it's up the makers of software to tell you how to use it. My job is to help you understand what you're trying to accomplish. *Many different software packages will do the same job equally well*, and new software is coming out every day.

12.1 How to avoid all this work

Before proceeding I should tell you that you don't have to do all this work. A much simpler procedure is to let the camera do most of it for you. Here's how:

- Turn on long-exposure noise reduction in your camera. That way,

12.3. DETAILED PROCEDURE WITH MAXDSLRL

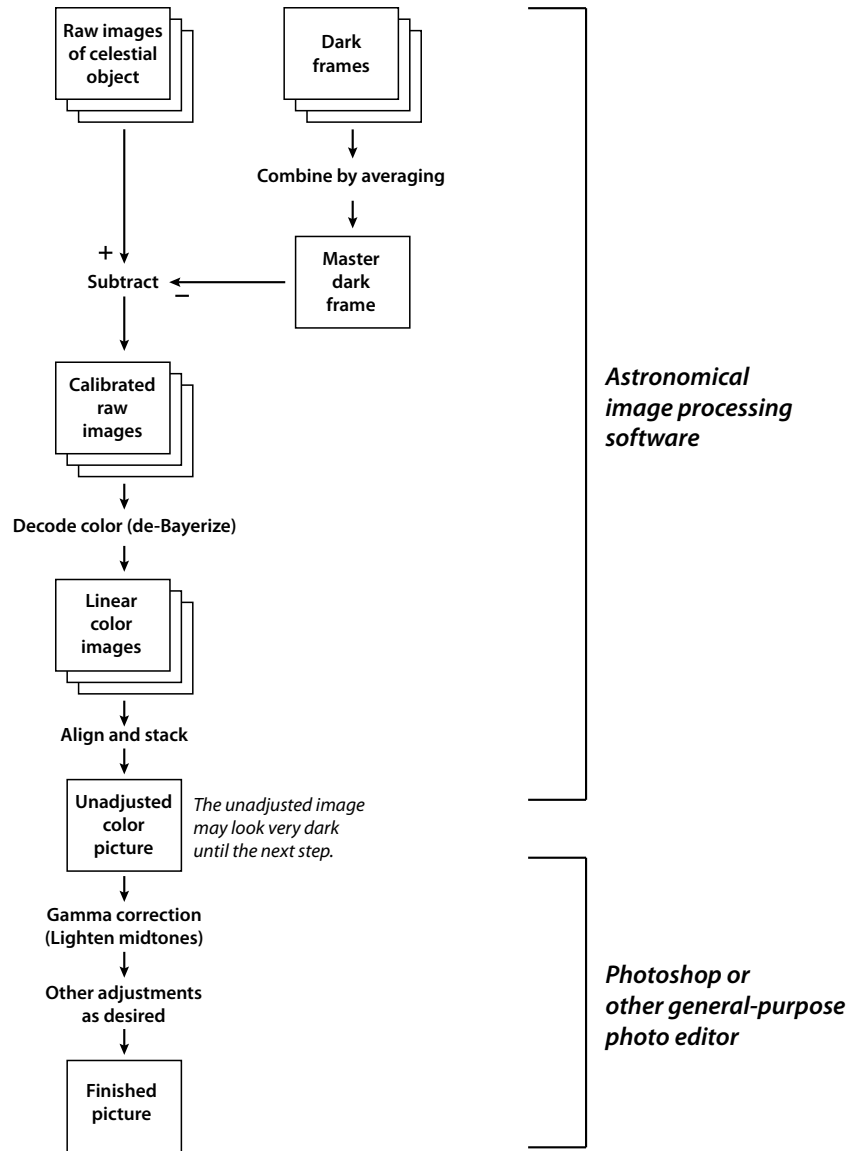


Figure 12.1: The path to a high-quality digital image from raw image files and dark frames. Good software automates much of the process.

12.3. DETAILED PROCEDURE WITH MAXDSLRL

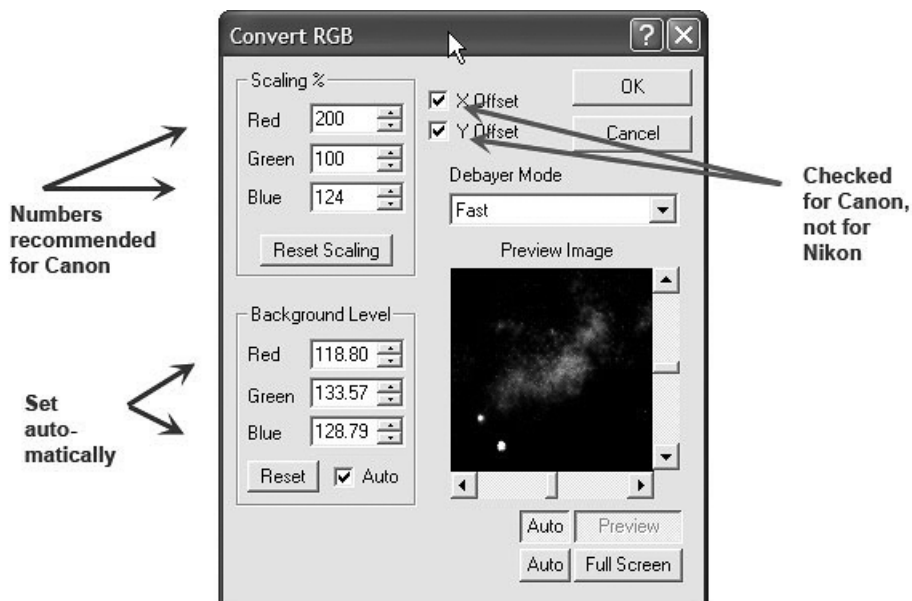


Figure 12.7: Settings to de-Bayerize (demosaic) the raw images and convert them to color.

offset refer to. Other software packages automatically use the right offset for each type of camera, but for maximum versatility, *MaxDSLR* leaves it up to you.

To check the offset, *I urge you to process a daytime picture or two* (skipping the dark-frame step, of course). It's OK if your pictures have a color cast, but they should not swap colors, rendering yellow as blue or the like. Only a picture of a familiar, colorful daytime object can confirm this.

12.3.4 Combining images

12.3.4.1 How to align and stack

Now it's time to stack (combine) your color images. *MaxDSLR* can usually do this automatically. But you have some decisions to make:

13.5. SHARPENING




Figure 13.6: The concept of unsharp masking: (a) original image; (b) unsharp mask; (c) result of stacking them; (d) contrast stretched to full range.

13.5. SHARPENING



Figure 13.9: A wavelet (one of many types). A complex signal can be described as the sum of wavelets.

more appropriate for long-lasting waveforms, such as sound waves; wavelets are more appropriate for nonrepetitive waveforms, such as images.

After analyzing an image into wavelets, you can selectively strengthen the wavelets of a particular width, thereby strengthening details of a particular size but not those that are larger or smaller. This is how *RegiStax* brings out fine detail on planets without strengthening the film grain, which is even finer (see p. 272).

Image processing that involves separating the image into different frequency components is often called *multiscale* processing.

13.5.5 Deconvolution

Suppose you have a blurred image, and you know the exact nature of the blur, and the effects of the blur have been preserved in full. Then it ought to be possible to undo the blur by computer, right?

Indeed it is, but the process is tricky. It's called *deconvolution* and has two pitfalls. First, you never have a perfectly accurate copy of the blurred image to work with; there is always some noise, and if there were no noise there would still be the inherent imprecision caused by quantization.

Second, deconvolution is what mathematicians call an *ill-posed problem* — it does not have a unique solution. There is always the possibility that the image contained even more fine detail which was completely hidden from view.

For both of these reasons, deconvolution has to be guided by some criterion of what the finished image ought to look like. The most popular

13.5. SHARPENING




Figure 13.10: Deconvolution shrinks star images and brings out fine detail. This example is slightly overdone, as shown by dark doughnuts around stars in front of nebula.

13.6. COLOR CONTROL

criterion is *maximum entropy*, which means maximum simplicity and smoothness. (Never reconstruct two stars if one will do; never reconstruct low-level fluctuation if a smooth surface will do; and so on.) Variations include the Richardson-Lucy and van Cittert algorithms.

Deconvolution has long been a specialty of *MaxIm DL* (not *MaxDSLR*), although several other software packages now offer it. In Fig. 13.10 you see the result of doing it. Fine detail pops into view. The star images shrink and become rounder; even irregular star images (from bad tracking or miscollimation) can be restored to their proper round shape. Unfortunately, if the parameters are not set exactly right, stars are often surrounded by dark doughnut shapes.

Because deconvolution is tricky to set up and requires lots of CPU time, I seldom use it. My preferred methods of bringing out fine detail are unsharp masking and wavelet-based filtering.

13.6 Color control

13.6.1 Gamut

Like film, computer monitors and printers reproduce colors, not by regenerating the spectrum of the original light source, but simply by mixing three primary colors. This works only because the human eye has three types of color receptors. Creatures could exist — and indeed some human beings *do* exist — for whom the three-primary-color system does not work.³

By mixing primaries, your computer screen can stimulate the eye's color receptors in any combination, but not at full strength. That is, it has a limited *color gamut*. Colors outside the gamut can only be reproduced at lower saturation, as if they were mixed with white or gray. Nothing on your computer screen will ever look quite as red as a ruby or as green as an emerald.

The gamut of an inkjet printer is also limited, more so than the screen

³Severely color-blind people have only two primary colors. There are also humans with normal color vision whose primary red is not at the usual wavelength, and it is speculated that a person who inherits that system from one parent and the normal system from the other parent could end up with a working four-color system.

14.1. COMBINING IMAGES




Figure 14.2: Airplane trail disappears when image is median-combined with two more images. If the images had been averaged, the trail would have remained visible at reduced contrast.

14.1.1.4 Sigma clipping

The average gives the best smoothing, but the median has the ability to reject abnormal pixels. Is there a way to combine the two?

That is the idea behind *sigma clipping* — averaging with abnormal values excluded from the mix. To judge which values are abnormal, the program computes the average and standard deviation (σ , sigma), then rejects any pixels that differ from the mean by more than $k\sigma$ where k is a factor you specify, and computes a new average omitting the rejected pixels.

14.1. COMBINING IMAGES

14.1.2 Stacking images in *Photoshop*

So far we have only discussed image stacking done automatically by astronomical software. Sometimes you'll want to stack images under full manual control using *Photoshop* or *Photoshop Elements*. Here is the procedure. It introduces an important concept, *layers*.

1. Open both images in *Photoshop*.
2. In the toolbar, choose the Marquee Tool (dotted box).
3. Right-click on one of the images and choose *Select All*.
4. In the main menu, choose *Edit, Copy*.
5. Click on the title bar of the other image.
6. In the main menu, choose *Edit, Paste*.

You have now pasted the second image on top of the first one. However, your goal was to mix them 50-50, so you're not finished yet. Proceed as follows.

- (7) In the main menu, choose *Window, Layers* to make the *Layers* window visible.
- (8) In the *Layers* window, set the opacity of Layer 1 (the added image) to 50% (Fig. 14.3).
- (9) Hit Ctrl-+ several times to magnify the image so that you can see single pixels.
- (10) Choose the Move tool, and use the mouse or the arrow keys to move one image relative to the other (Fig. 14.4) until they are perfectly aligned.
- (11) In the main menu, choose *Layer, Flatten Image*. This combines the layers into one.

Your image is now ready for further processing.

While aligning the images, you may want to set the blending mode (*Normal* in Fig. 14.3) temporarily to *Difference*. Then, when the images are

14.2. CALIBRATION FRAMES




Figure 14.6: The effect of flat-fielding on a wide-field view of Orion with a Canon EOS 20Da and $H\alpha$ filter. (a) Image calibrated with dark frames but not flat fields. (b) A flat field taken through the same lens on the same evening. Note slight vignetting and prominent dust spot. (c) Image calibrated with dark frames and flat fields. (Slightly undercorrected; see text.)

14.5. THE EXTREME BRIGHTNESS RANGE OF NEBULAE

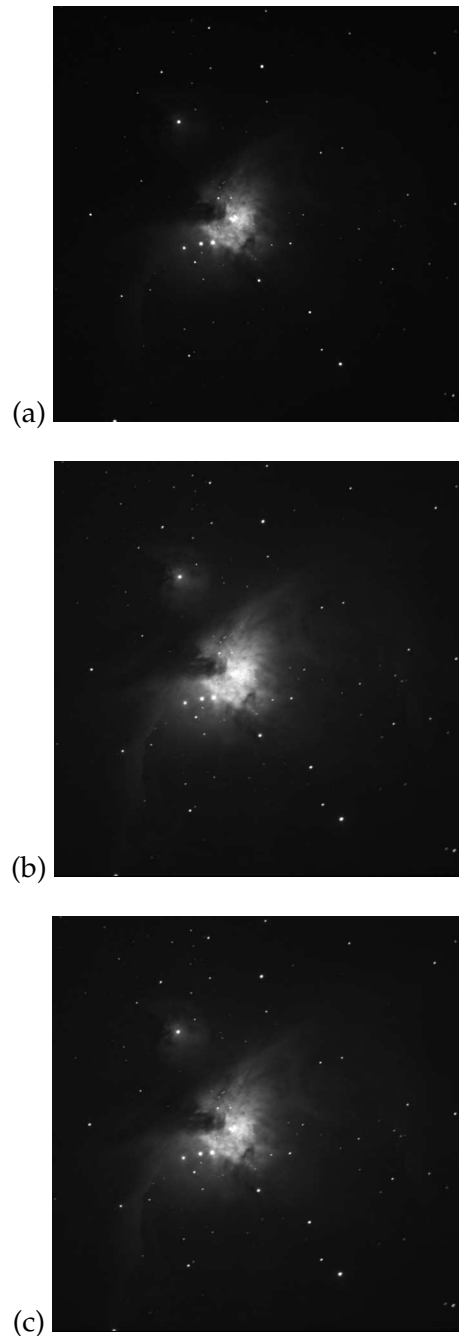


Figure 14.10: Combining different exposures with layer masking. (a) Orion Nebula, 1 minute, 8-inch telescope at $f/6.3$, Canon XTi, ISO 400. (b) Same, 2 minutes. (c) Combined, using a layer mask to take the central region from the shorter exposure.

14.5. THE EXTREME BRIGHTNESS RANGE OF NEBULAE

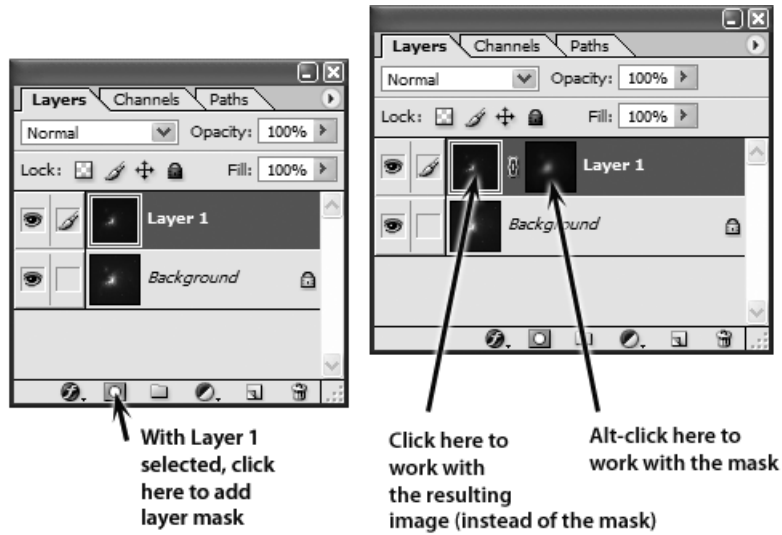


Figure 14.11: How to add a mask to a layer in *Photoshop* and how to work with it.



Figure 14.12: The layer mask used in Fig. 14.10.

14.5. THE EXTREME BRIGHTNESS RANGE OF NEBULAE



Figure 14.13: The Orion Nebula (M42) in exquisite detail — a combination of 29 5-minute, 5 1-minute, and 5 10-second exposures at ISO 800 using a Canon XT (350D) and an Orion ED80 8-cm telescope with a focal reducer giving $f/4.7$. Wire crosshairs in front of the telescope produced diffraction. Processing included layer masking, Richardson-Lucy deconvolution, and noise removal with *Neat Image*. (Hap Griffin)