Digital SLR Astrophotography

by Michael A. Covington

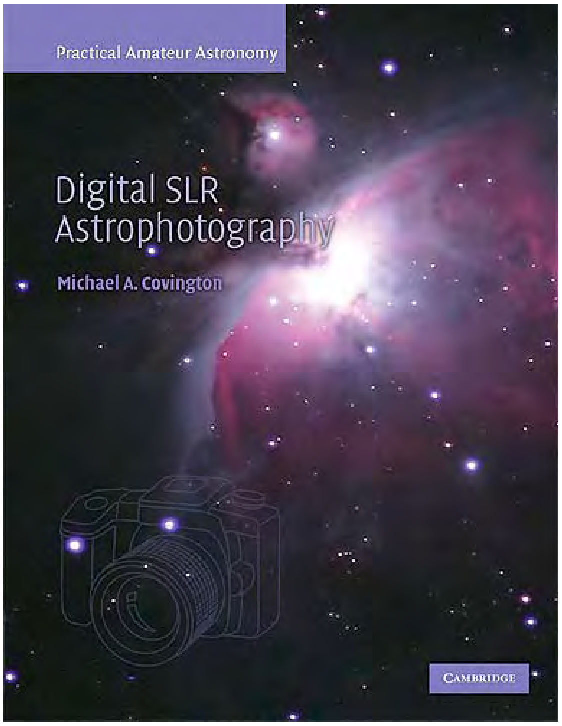

Sample Pages

For more information see www.dslrbook.com.

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 (Figure 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.

1.1 What is a DSLR?

A DSLR is a digital camera that is built like a film SLR (single-lens reflex) and has the same ability to interchange lenses. You can attach a DSLR to anything that will form an image, whether it's a modern camera lens, an old lens you have adapted, or a telescope, microscope, or other instrument.

Unlike other digital cameras, a DSLR does not normally show you a continuous electronic preview of the image. Instead, the viewfinder of a DSLR uses a mirror and a focusing screen to capture the image optically so that you can view

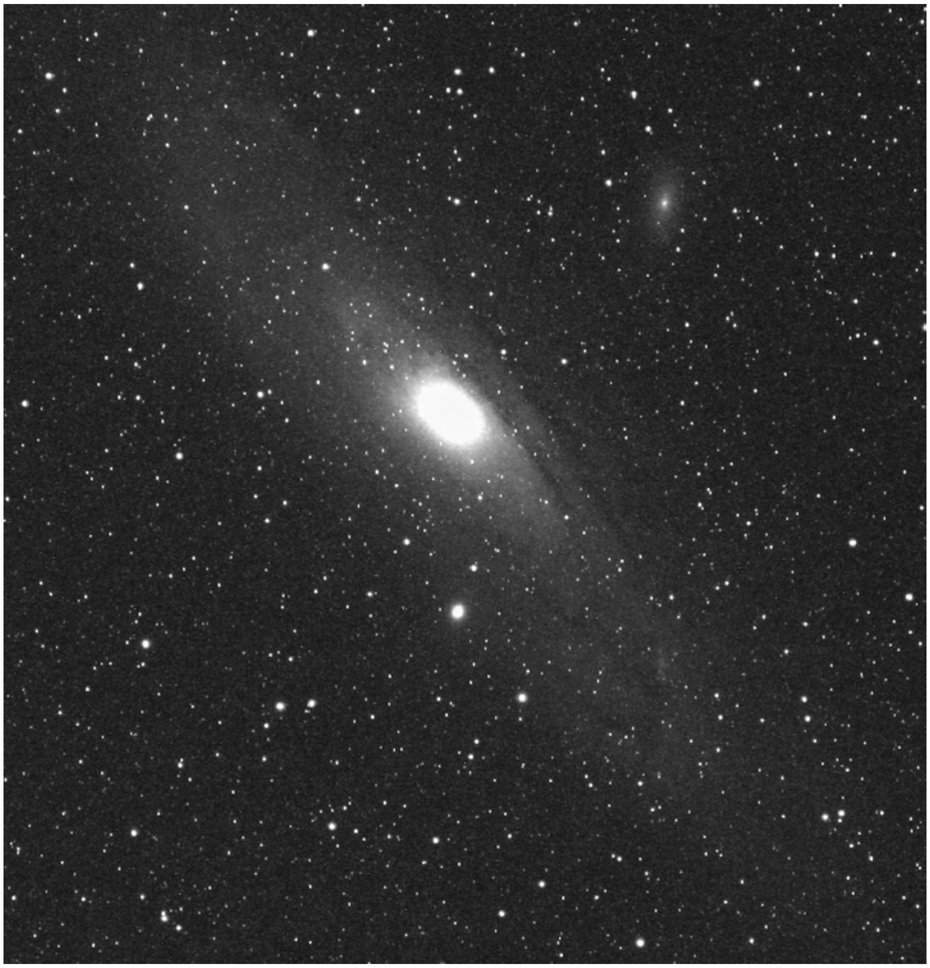

Figure 1.1. The galaxy M31 as the image came from the camera, with no processing except adjustment of brightness and contrast. Canon Digital Rebel (300D); single 6-minute exposure through a 300-mm lens at $f/5.6$, captured as JPEG. Some noise specks are present which newer cameras would eliminate with automatic dark-frame subtraction.

and focus through an eyepiece. When you take the picture, the mirror flips up, the image sensor is turned on, and the shutter opens.

The reason a DSLR doesn't show the electronic image continuously is that its sensor is much larger than the one in a compact digital camera. Big sensors are good because they produce much less noise (speckle), especially in long exposures, but operating a big sensor all the time would run down the battery. It would also cause the sensor to warm up, raising its noise level. That's why you normally view through the mirror, focusing screen, and eyepiece.

Some DSLRs *do* offer "live focusing" or "live previewing" for up to 30 seconds at a time. The Canon EOS 20Da, marketed to astronomers in 2005, was the first. Live previewing enables you to focus much more precisely than by looking

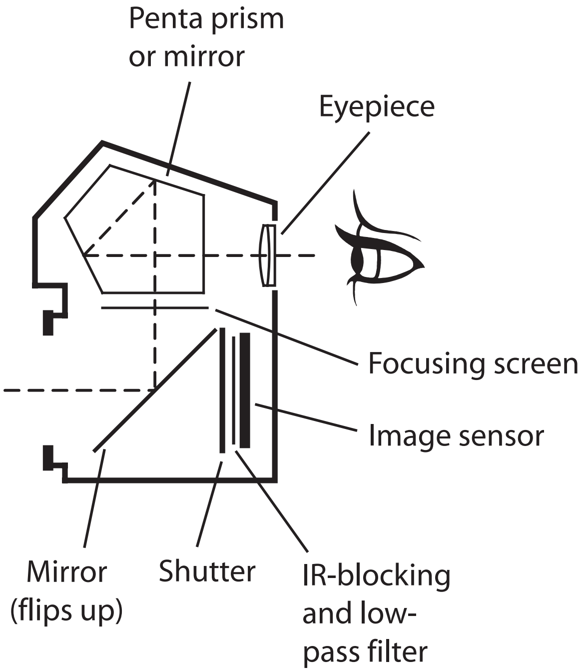

Figure 1.2. A DSLR is a single-lens reflex with a digital image sensor. Mirror and eyepiece allow you to view the image that will fall on the sensor when the mirror flips up and the shutter opens.

Figure 1.3. A more elaborate view of what's inside a DSLR. Note computer circuitry ("DIGIC II") at right. (Canon USA.)

Figure 3.6. Dots punched out of red vinyl tape make good covers for the power-on LED and self-timer indicator.

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 ▷ ("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.

Chapter 4
Four simple projects

After all this, you're probably itching to take a picture with your DSLR. This chapter outlines four simple ways to take an astronomical photograph. Each of them will result in an image that requires only the simplest subsequent processing by computer.

All of the projects in this chapter can be carried out with your camera set to output JPEG images (not raw), as in daytime photography. The images can be viewed and further processed with any picture processing program.

4.1 Telephoto Moon

Even though the Moon is not ultimately the most rewarding object to photograph with a DSLR, it's a good first target.

Put your camera on a sturdy tripod and attach a telephoto lens with a focal length of at least 200 and preferably 300 mm. Take aim at the Moon. Initial exposure settings are ISO 400, $f/5.6$, 1/125 second (crescent), 1/500 second (quarter moon), 1/1000 (gibbous), or 1/2000 (full); or simply take a spot meter reading of the illuminated face of the Moon. An averaging meter will overexpose the picture because of the dark background.

If the camera has mirror lock (Canon) or exposure delay (Nikon), turn that feature on. Let the camera autofocus and take a picture using the self-timer or cable release. View the picture at maximum magnification on the LCD display and evaluate its sharpness. Switch to manual focus and try again, varying the focus slightly until you find the best setting. Also adjust the exposure for best results. If you have mirror lock or prefire, you can stop down to $f/8$ and use a slower shutter speed.

Figures 4.1 and 4.2 show what you can achieve this way. Images of the Moon benefit greatly from unsharp masking in *Photoshop* or *RegiStax*; you'll be surprised how much more detail you can bring out.

To make the face of the Moon fill the sensor, you'll need a focal length of about 1000–1500 mm. In the next example, we'll achieve that in a very simple way.

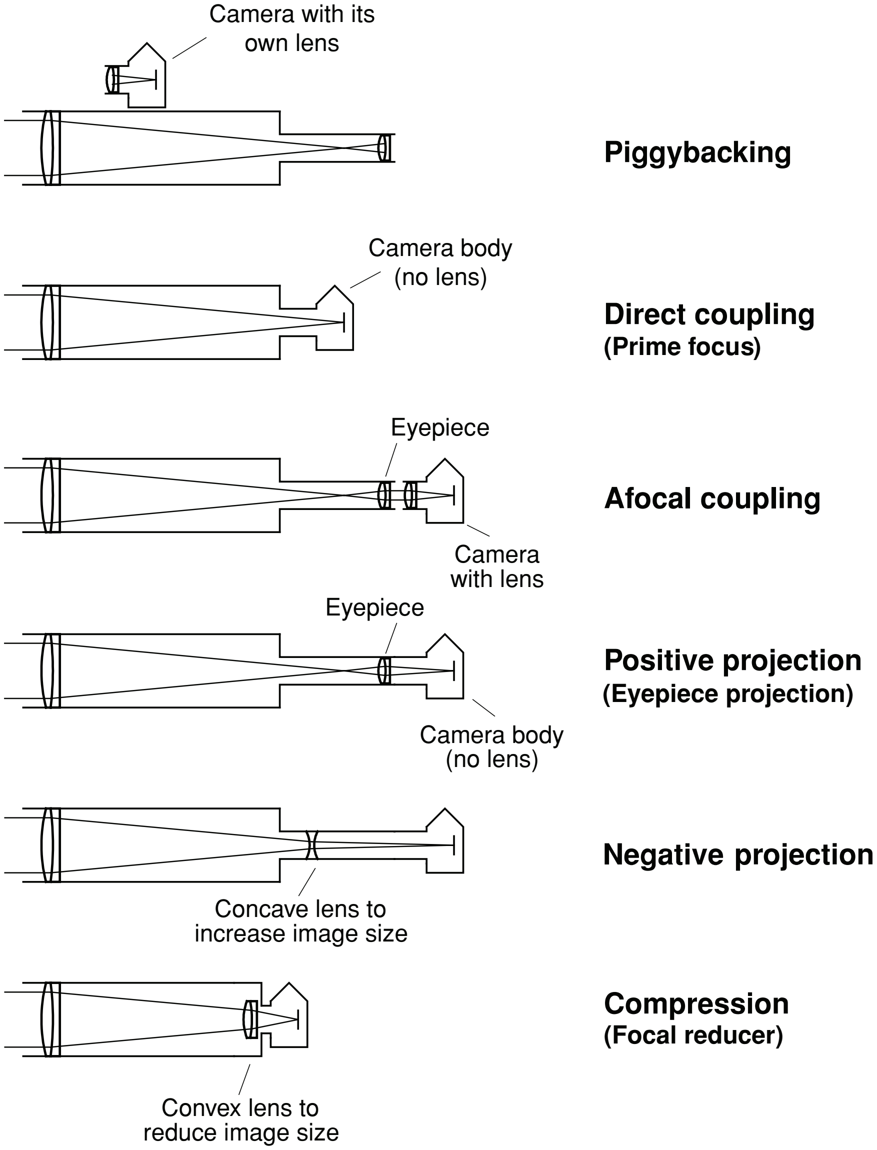

Figure 5.2. Ways of coupling cameras to telescopes. Piggybacking, direct coupling, and compression are main modes for deep-sky work.

Not all of these modes work equally well, for several reasons. First, DSLRs excel at deep-sky work, not lunar and planetary imaging. Accordingly, we want a bright, wide-field image. That means we normally leave the focal length and f-ratio of the telescope unchanged (with direct coupling) or reduce them (with compression). The modes that magnify the image and make it dimmer – positive and negative projection and, usually, afocal coupling – are of less interest.

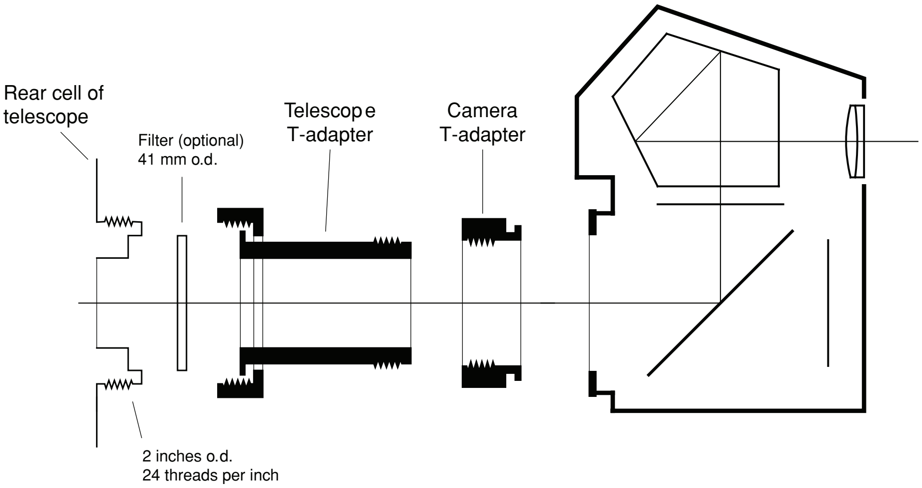

Figure 5.5. Schmidt–Cassegrains have a threaded rear cell and accept a matching T-adapter.

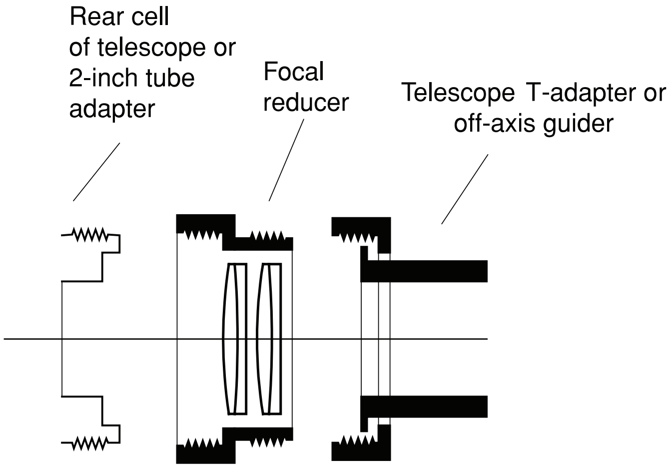

Figure 5.6. Meade or Celestron focal reducer screws onto rear cell of telescope.

You can get a simple eyepiece-tube adapter that screws into the T-ring (Figure 5.4) or a telescope T-adapter for other types of telescopes (Figure 5.5). Here you'll encounter the other screw coupling that is common with telescopes, the classic Celestron rear cell, which is 2 inches in diameter and has 24 threads per inch. Meade adopted the same system, and telescope accessories have been made with this type of threading for over 30 years. In particular, that's how the Meade and Celestron focal reducers attach to the telescope (Figure 5.6).

If your telescope doesn't have Celestron-style threads on its rear cell, or you're using an electric focuser that makes them inaccessible, then Meade's and also

$$\frac{22.2}{3888} = 0.00571 \text{ mm horizontally}$$

Now find the field of view of a single pixel; that is, pretend your sensor is 0.00571 mm square. Suppose the telescope is a common 20-cm (8-inch) $f/10$ Schmidt–Cassegrain with a focal length of 2000 mm. Then:

$$\text{Field of one pixel} = 57.3^\circ \times \frac{0.00571 \text{ mm}}{2000 \text{ mm}} = 0.000\,016\,359^\circ = 0.59''$$

As with astronomical CCD cameras, the pixel size works out to be close to the resolution limit of the telescope.

5.3.6 "What is the magnification of this picture?"

Non-astronomers seeing an astronomical photograph often ask what magnification or "power" it was taken at.

In astronomy, such a question almost has no answer. With microscopes, it does. If you photograph a 1-mm-long insect and make its image 100 mm long on the print, then clearly, the magnification of the picture is × 100.

But when you tell people that your picture of the Moon is 1/35 000 000 the size of the real Moon, somehow that's not what they wanted to hear. Usually, what they mean is, "How does the image in the picture compare to what I would see in the sky with my eyes, with no magnification?" The exact answer depends, of course, on how close the picture is to the viewer's face.

But a rough answer can be computed as follows:

$$\text{"Magnification" of picture} = \frac{45^\circ}{\text{Field of view of picture}}$$

That is: If you looked through a telescope at this magnification, you'd see something like the picture. Here 45° is the size of the central part of the human visual field, the part we usually pay attention to, and is also the apparent field of a typical (not super-wide) eyepiece.

So a picture that spans the Moon (half a degree) has a "magnification" of 90. Deep-sky piggyback images often have rather low "magnification" computed this way, anywhere from 10 down to 2 or less.

5.4 Vignetting and edge-of-field quality

As *Astrophotography for the Amateur* explains at length, few telescopes produce a sharp, fully illuminated image over an entire 35-mm film frame or even an entire DSLR sensor. The reason is that telescopes are designed to work with eyepieces, and with an eyepiece, what you want is maximum sharpness at the very center of the field.

Figure 7.9. Mid-grade and high-grade Nikon to EOS adapters. Avoid cheap imitations.

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.

Figure 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 accurate machining, more elaborate mechanisms to lock and unlock the lens, and more blackening of parts that could reflect light.

M42 screw mount to Canon adapters are the simplest, and my experience is that they always work well. I have also had good results with an inexpensive

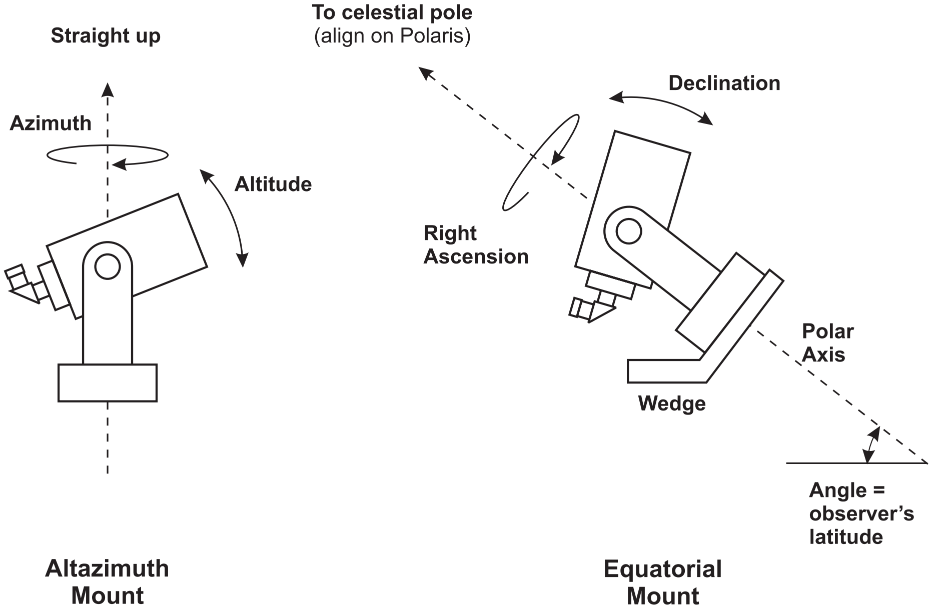

Figure 9.1. Two ways to track the stars. The altazimuth mount requires computer-controlled motors on both axes; equatorial needs only one motor, no computer. (From *How to Use a Computerized Telescope*.)

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.

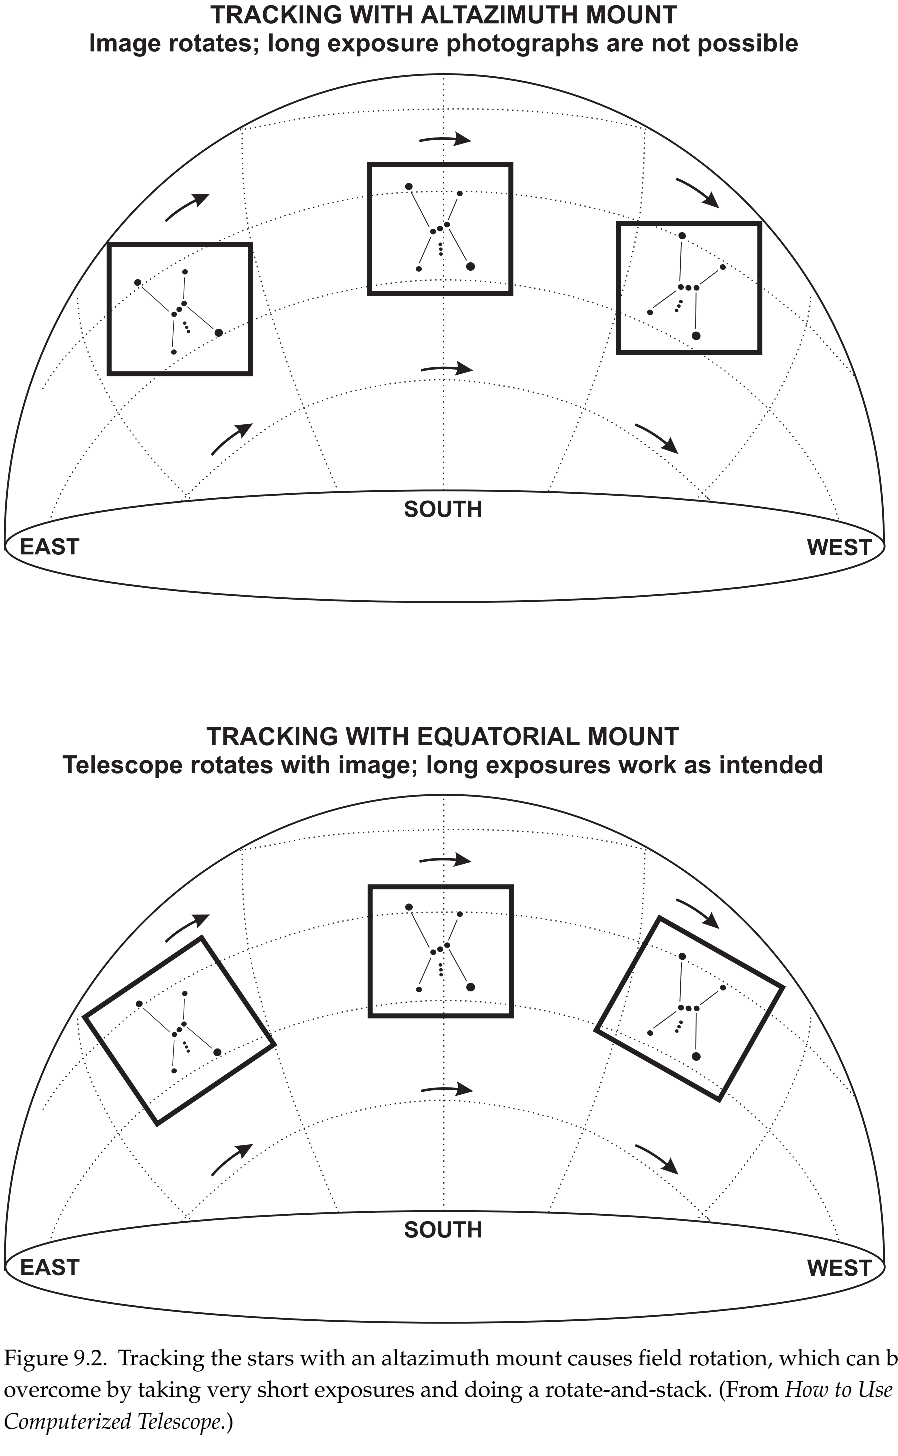

Figure 9.2. Tracking the stars with an altazimuth mount causes field rotation, which can be overcome by taking very short exposures and doing a rotate-and-stack. (From *How to Use a Computerized Telescope.*)

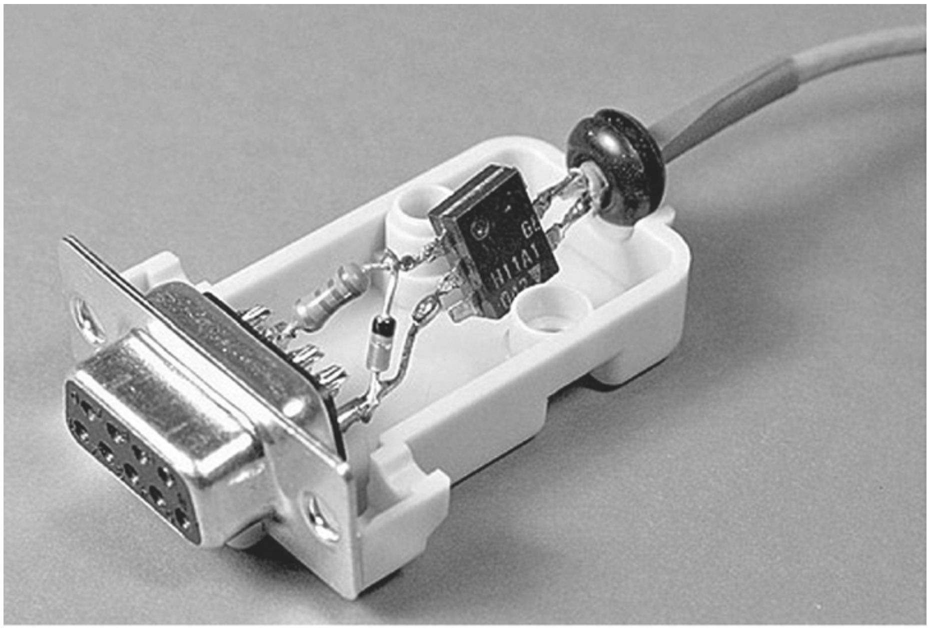

Figure 10.6. The serial cable circuit can be assembled in a large nine-pin connector, then connected directly to the computer or to a USB serial cable.

For details of the cable needed, consult the instructions that accompany your software; there are many variations. Figure 10.5 shows one popular type.

The software interface is very simple. Instead of outputting data to the serial or parallel port, the software simply sets one pin positive instead of negative. On a serial port, this is usually the RTS line (pin 7; pin 5 is ground); on a parallel port, pin 2 (and pin 18 is ground).

The value of the resistor can vary widely; erring on the side of caution, I chose 10 kΩ, but some published circuits go as high as 47 kΩ. The second version of the circuit uses an optoisolator to eliminate a ground loop, but if the camera and computer are also tied together by their USB ports, this is a moot point.

If your DSLR isn't one of the Canons with the 2.5-mm phone plug, Figure 10.7 shows what to do. You'll need a cable release or infrared controller for your camera. Find the switch in it that trips the shutter; use a voltmeter to identify the positive and negative sides of it; and wire the transistor across it.

10.3 Networking everything together

Now it's time to sum up. The typical modern DSLR astrophotographer uses a laptop computer for control, a webcam for autoguiding, and a computerized telescope. Here are all the connections that have to be taken care of:

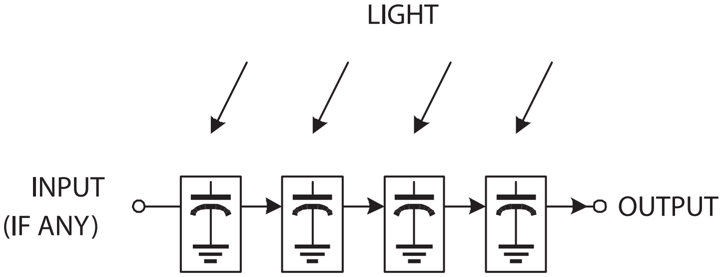

Figure 11.1. Charge-coupled device (CCD) consists of cells (pixels) in which electrons can be stored, then shifted from cell to cell and retrieved. Light makes electrons enter the cells and slowly raise the voltage during an exposure. (From *Astrophotography for the Amateur.*)

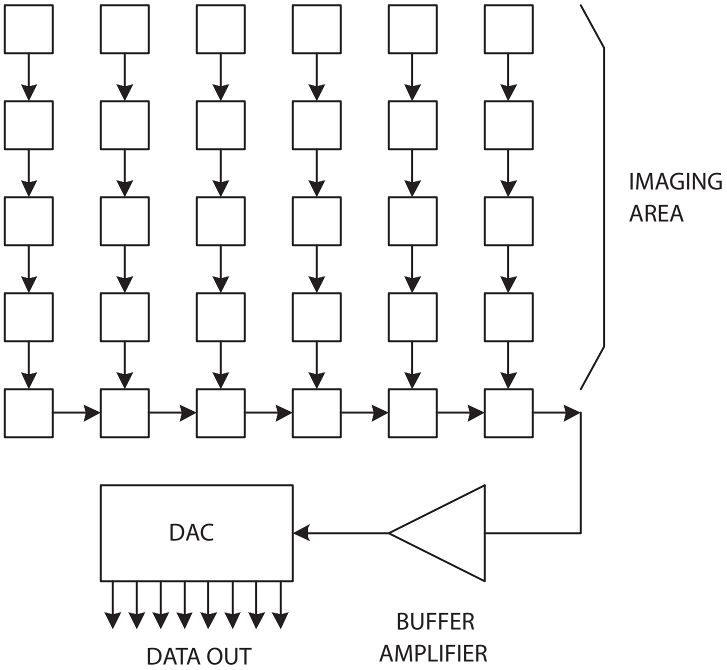

Figure 11.2. CCD array shifts the contents of each cell, one after another, through other cells to the output for readout. CMOS arrays have an amplifier for each pixel, with row and column connections to read each pixel directly. (From *Astrophotography for the Amateur.*)

Today, however, CCD and CMOS sensors compete on equal ground. The CMOS amplifiers are under the pixels, not between them, and the quantum efficiency is comparable to CCDs. So are all other aspects of performance, including noise. Canon DSLRs use CMOS sensors (to which Canon has devoted a lot of work), most others use CCDs, and neither one has a consistent advantage over the other. Nikon adopted a Sony CMOS sensor for the high-end D2X camera.

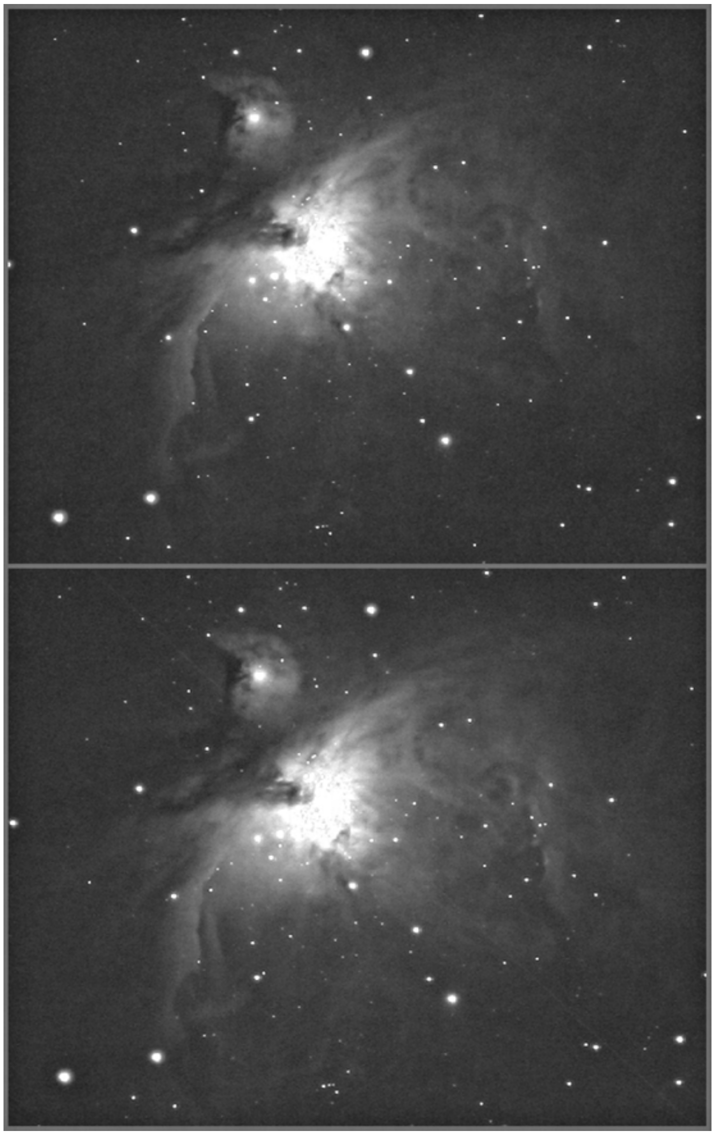

Figure 11.6. Effect of filter modification. *Top:* unmodified Canon XTi (400D). *Bottom:* Canon 20Da with extended hydrogen-alpha sensitivity. Each is a single 3-minute exposure of the Orion Nebula (M42) with 420-mm $f/5.6$ lens, processed with *Photoshop* to give similar color balance and overall contrast. See the back cover of this book for the same pictures in color.

Other modified cameras show somewhat more of an increase. In the picture, the right-hand edge of the nebula is representative of thinner, redder hydrogen nebulae elsewhere.

What is striking is how well the *un*modified camera works. Yes, the modification helps – but it's not indispensable. The unmodified camera shows the nebula reasonably well.

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. 195.

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, whenever you take a celestial photograph, the camera will automatically take a dark frame and subtract it.
- Tell the camera to save the images as JPEG (not raw).
- Open the resulting image files in *Photoshop* or any photo editor and adjust the brightness, contrast, and color balance to suit you.

Why don't we always take the easy way out? For several reasons.

First, we usually want to combine multiple images. With digital technology, ten 1-minute exposures really *are* as good as one 10-minute exposure – almost. They're certainly a lot better than *one* 1-minute exposure. Combining

Chapter 13

Digital imaging principles

This chapter is a crash course in the principles of digital image processing. For more about most of these concepts, see *Astrophotography for the Amateur* (1999), Chapter 12.

13.1 What is a digital image?

A digital image is fundamentally an array of numbers that represent levels of brightness (Figure 13.1).

13.1.1 Bit depth

Depending on the *bit depth* of the image, the numbers may range from 0 to 255 (8 bits), 0 to 65 535 (16 bits), or some other range.

The eye cannot distinguish even 256 levels, so 8-bit graphics are sufficient for finished pictures. The reason for wanting more levels during manipulation is that we may not be using the full range at all stages of processing. For instance, a badly underexposed 16-bit image might use only levels 0 to 1000, which are still enough distinct levels to provide smooth tones. An 8-bit image underexposed to the same degree would be unusable.

For greatest versatility, some software supports floating-point data, so that levels can be scaled with no loss of precision; in a floating-point system, you can divide 65 535 by 100 and get 655.35. You can also use large numbers without going out of range; if you do something that produces a value greater than 65 535, it will not be clipped to maximum white.

Note that *Photoshop* always reports brightness levels on a scale of 0 to 255, regardless of the actual bit depth of the image. This is to help artists match colors.

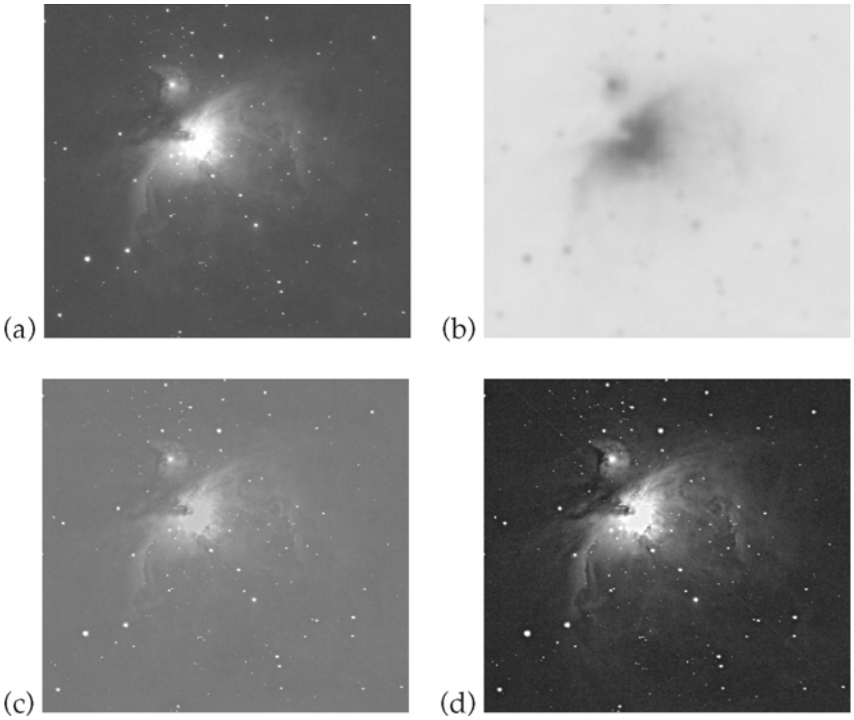

Figure 13.6. The concept of unsharp masking: (a) original image; (b) unsharp mask; (c) result of stacking them; (d) contrast stretched to full range.

13.5.3 Digital development

Digital development processing (DDP) is an algorithm invented by astrophotographer Kunihiko Okano that combines gamma stretching with unsharp masking.[2] It is particularly good at preserving the visibility of stars against bright nebulae. Many astrophotographers like the way it combines several steps of processing into one (Figure 13.7).

Some care is needed setting the parameters because implementations of digital development that are optimized for smaller CCDs will bring out grain in DSLR images. The unsharp masking radius needs to be much larger.

In *MaxDSLR,* the unsharp masking part of digital development can be turned off by setting a filter matrix that is all zeroes except for a 1 in the center; digital development then becomes a kind of gamma stretch.

13.5.4 Spatial frequency and wavelet transforms

Another way to sharpen an image is to analyze it into frequency components and strengthen the high frequencies.

---

[2] Web: *http://www.asahi-net.or.jp/~rt6k-okn*.

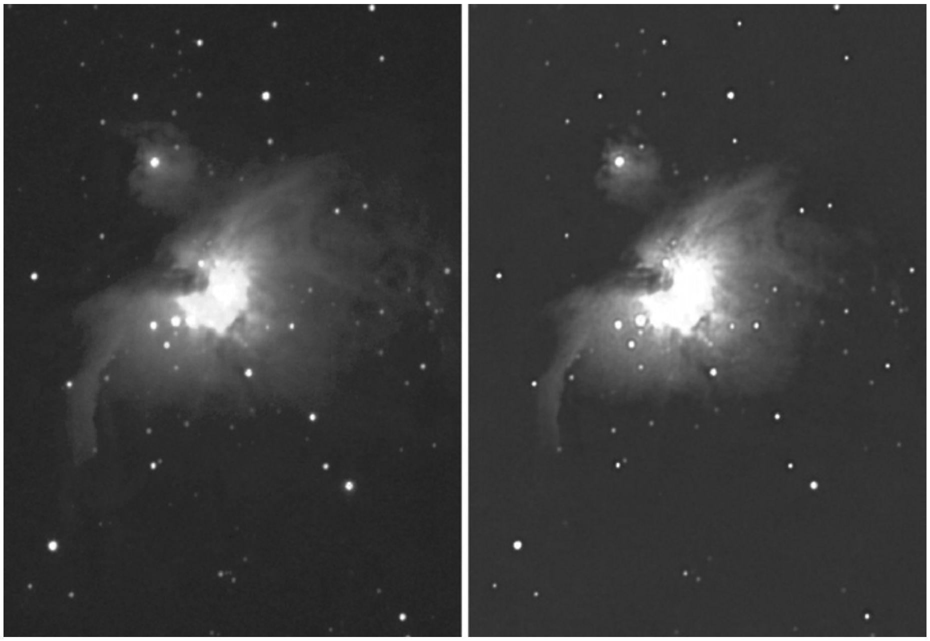

Figure 13.10. Deconvolution shrinks star images and brings out fine detail. This example is slightly overdone, as shown by dark doughnuts around stars in front of the nebula.

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.

---

[3] Severely 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.

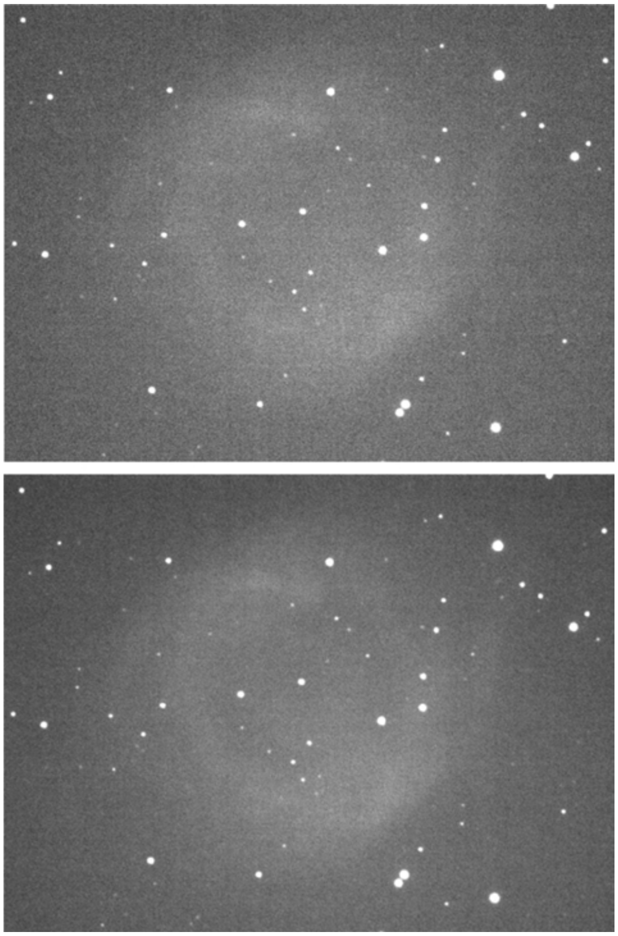

Figure 14.1. Combining images reduces noise. *Left:* Single 6-minute exposure of the Helix Nebula (NGC 7293) through an 8-inch (20-cm) telescope with $f/6.3$ compressor, with Canon Digital Rebel (300D) at ISO 400. *Right:* Average of three such exposures. Random noise is reduced by a factor of $\sqrt{3} = 1.7$.

That may not matter if the only bright spots in the image are star images and you don't mind having them all end up maximum white. In that situation, summing is a good way to strengthen faint objects such as nebulosity.

Average (mean)

The average (the mean) is the sum divided by the number of images. The resulting pixel values are always in the same range as the originals. This is actually the most common way of combining images. As with summing, every image contributes equally to the finished product. In fact, taking the average is exactly equivalent to summing the images and then scaling the pixel values back to the original range.

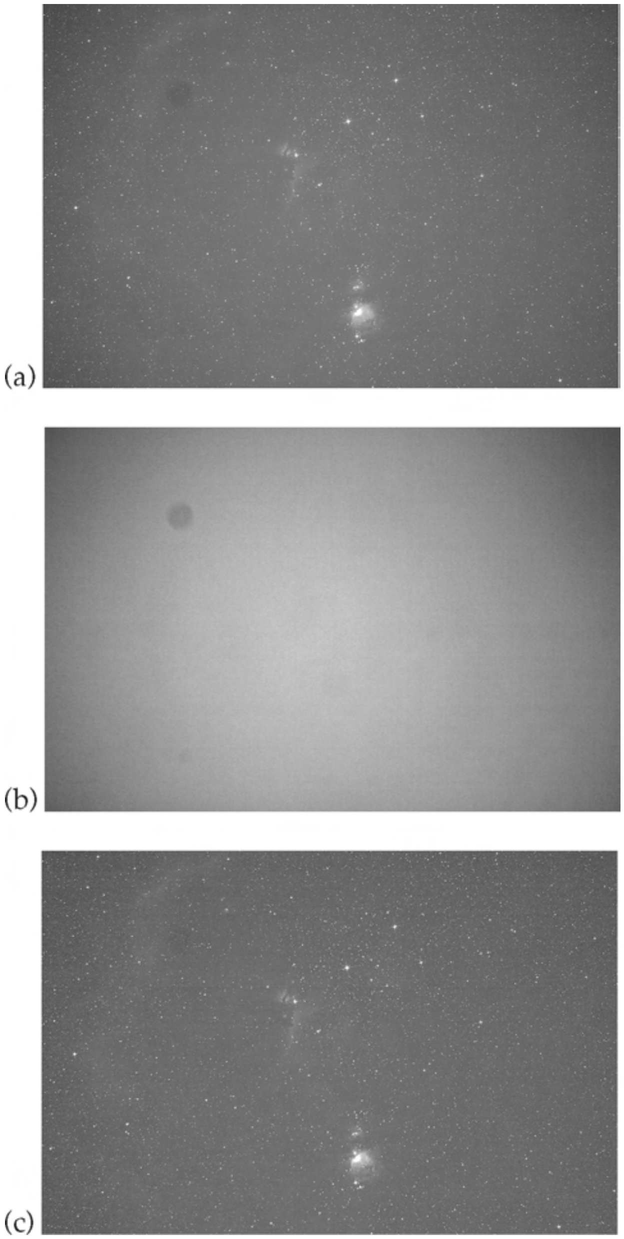

Figure 14.6. The effect of flat-fielding on a wide-field view of Orion with a Canon EOS 20Da and Hα filter. (a) Image calibrated with dark frames but not flat fields. (b) A flat field taken through the same lens on the same evening. Note the slight vignetting and prominent dust spot. (c) Image calibrated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

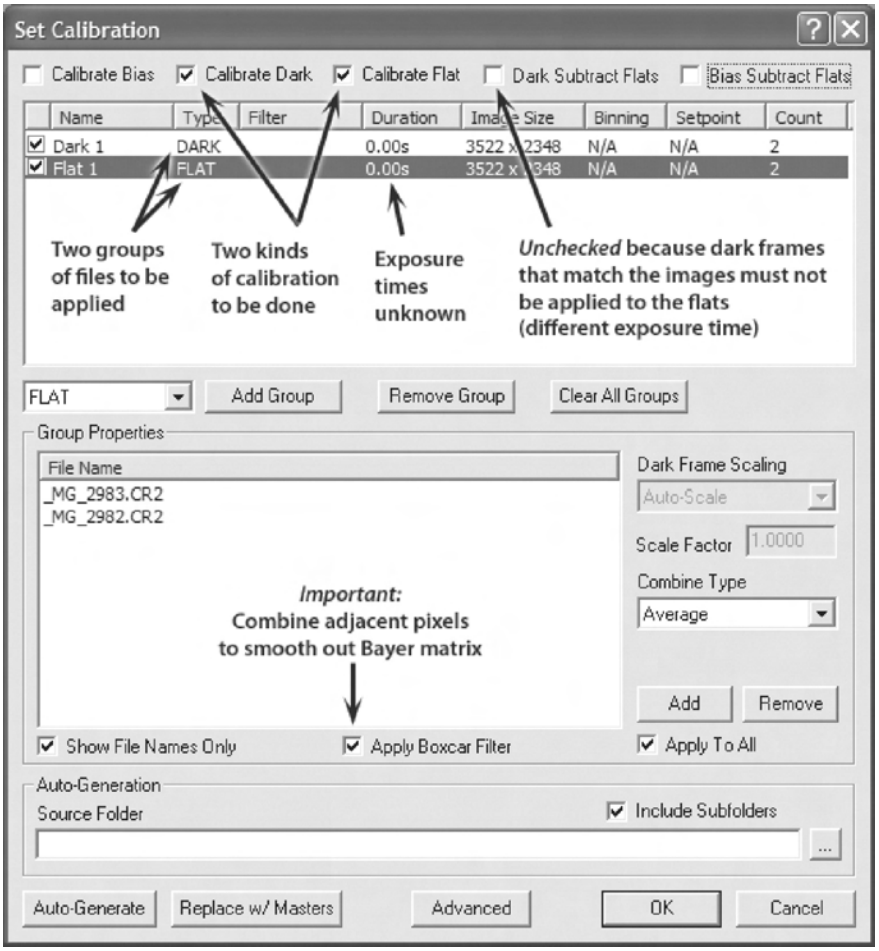

☐ Calibrate Bias ☑ Calibrate Dark ☑ Calibrate Flat ☐ Dark Subtract Flats ☐ Bias Subtract Flats

	Name	Type	Filter	Duration	Image Size	Binning	Setpoint	Count
☑	Dark 1	DARK		0.00s	3522 x 2348	N/A	N/A	2
☑	Flat 1	FLAT		0.00s	3522 x 2348	N/A	N/A	2

Two groups of files to be applied

Two kinds of calibration to be done

Exposure times unknown

***Unchecked* because dark frames that match the images must not be applied to the flats (different exposure time)**

FLAT

File Name
_MG_2983.CR2
_MG_2982.CR2

Dark Frame Scaling: Auto-Scale

Scale Factor: 1.0000

Combine Type: Average

Important:
Combine adjacent pixels to smooth out Bayer matrix

☑ Show File Names Only ☑ Apply Boxcar Filter ☑ Apply To All

☑ Include Subfolders

Figure 14.7. Flat-field and dark-frame calibration can be done in a single step in *MaxDSLR*.

The computation is very sensitive to errors. That's why it's so important to match the original ISO setting and optical conditions.

Better yet, make a *calibrated flat field*. That is, along with your flat fields, take some dark frames that match them. Subtract the dark frames from the flat fields, then average the flat fields (*without* aligning or de-Bayerizing them). Save the calibrated flat field and use it to calibrate your astronomical images.

Another reason for making a calibrated flat field is that you can adjust it. If your images come out undercorrected (with vignetting and dust spots still visible), scale up the contrast of the corrected flat field just a bit, and try again. If they are overcorrected, with reverse vignetting (edges brighter than the center), do the opposite.

Figure 14.7 shows the *Set Calibration* window in *MaxDSLR* when both darks and flats are to be applied. Don't make the mistake of subtracting, from your flats, the dark frames that go with the images; the exposure times are very

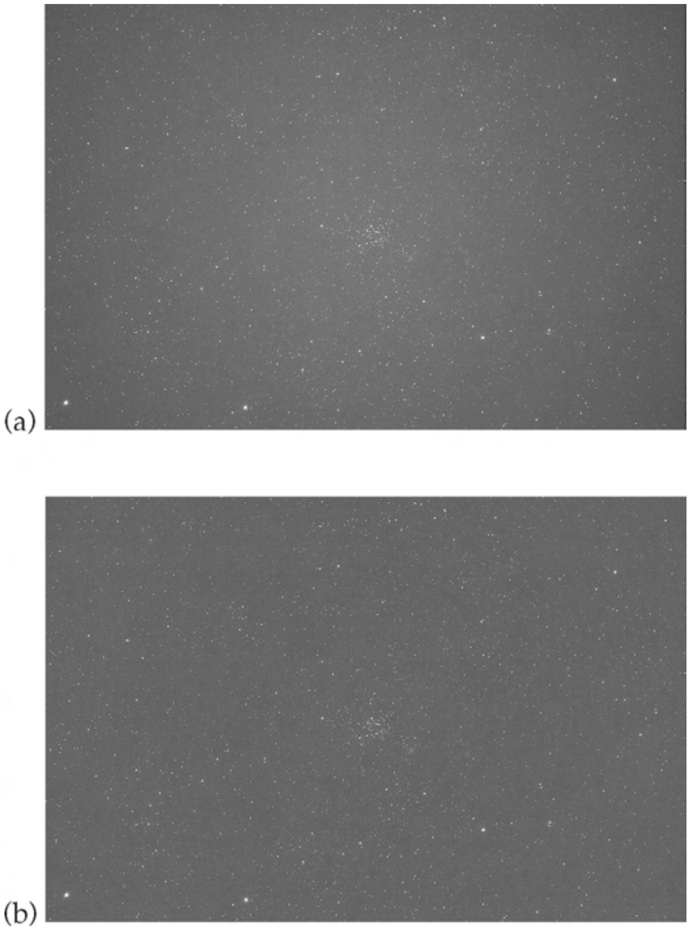

Figure 14.8. Removal of vignetting. (a) Mildly vignetted wide-field view of star cluster M35 (single 2-minute exposure, Nikon 200-mm $f/4$ lens wide open, unmodified Nikon D50 at ISO 400 in Mode 2). (b) After applying *Auto Flatten Background* in *MaxDSLR*.

14.4 Removing grain and low-level noise

Because no two pixels are alike, and because photons and electrons are discrete particles, digital images of faint objects are grainy. To a considerable extent, grain can be removed by software without obscuring the desired detail in the image. After all, if you can tell the difference between the grain and the image, then, in principle, so can the computer.

All computer graphics textbooks say that you can reduce noise by blurring the image – that is, by averaging each pixel with some of its neighbors. That is true but unsatisfying for two reasons. First, blurring the image makes stars disappear. Second, if you blur a random grain pattern without reducing the contrast, often what you end up with is just a bigger random pattern.

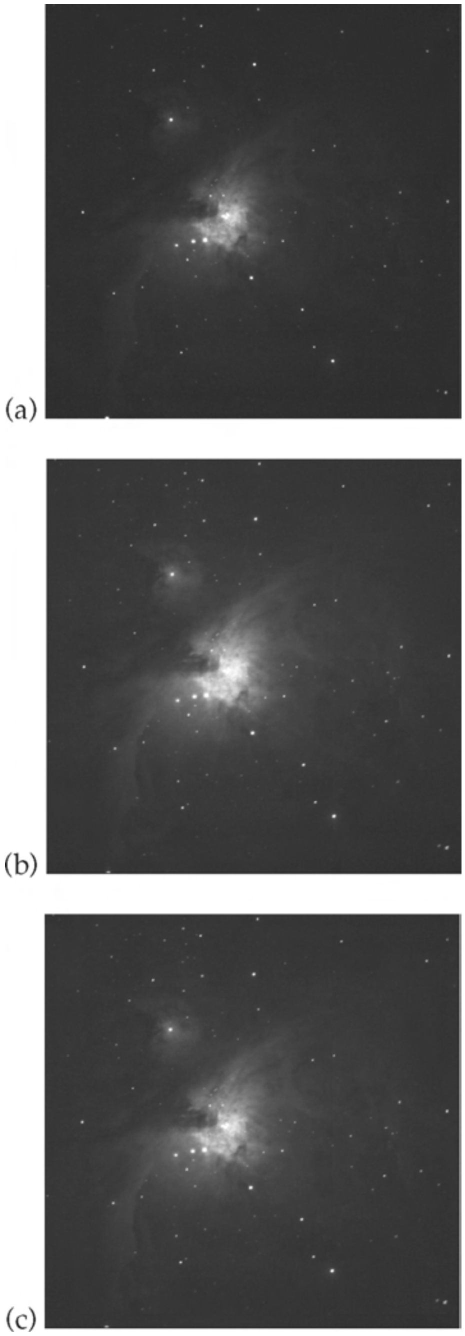

Figure 14.10. Combining different exposures with layer masking. (a) Orion Nebula, 1 minute, 8-inch telescope at $f/6.3$, Canon XTi, ISO 400. (b) Same, 2 minutes. (c) Combined, using a layer mask to take the central region from the shorter exposure.

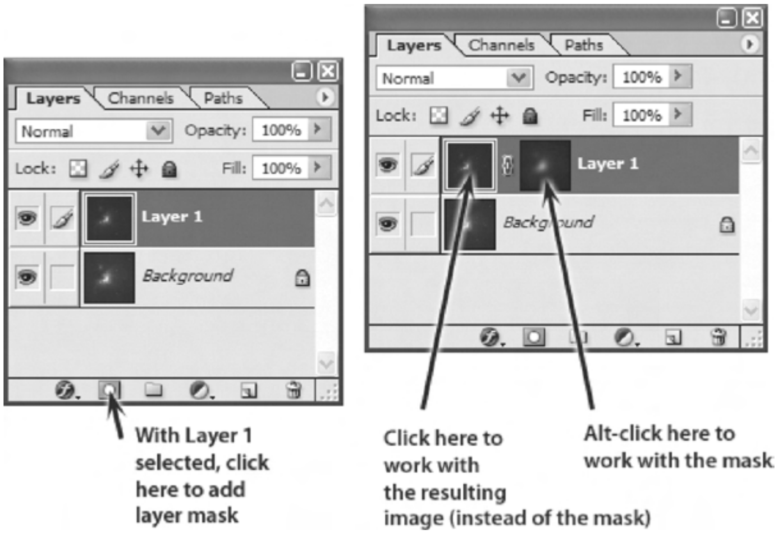

With Layer 1 selected, click here to add layer mask

Click here to work with the resulting image (instead of the mask)

Alt-click here to work with the mask

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

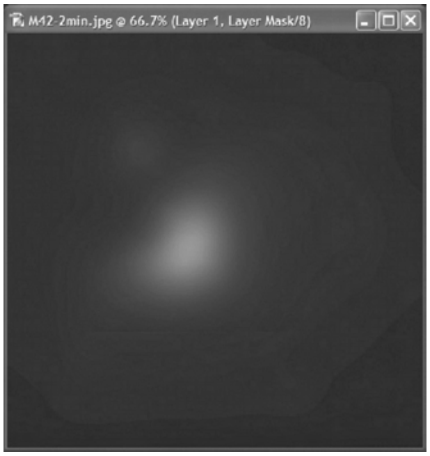

Figure 14.12. The layer mask used in Figure 14.10.

(10) Finally, flatten the image and make other adjustments, such as levels and unsharp masking.

You can combine any number of images in this way, and the results can be impressive (Figure 14.13).

14.6 Other *Photoshop* techniques

In the *Photoshop* main menu, *Image, Adjustments, Auto Levels* will automatically stretch the image, in each of three colors, so that it uses the full brightness range.

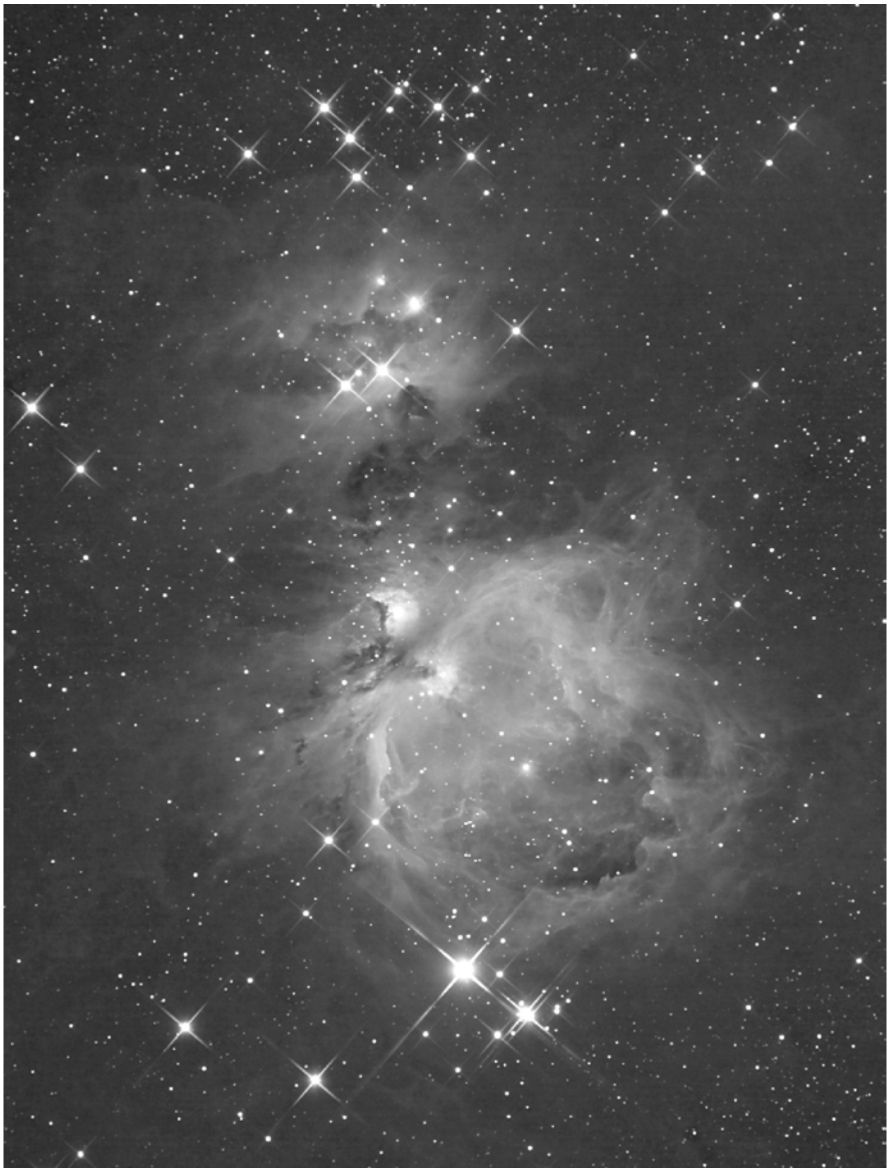

Figure 14.13. The Orion Nebula (M42) in exquisite detail – a combination of 29 5-minute, five 1-minute, and five 10-second exposures at ISO 800 using a Canon XT (350D) and an Orion ED80 8-cm telescope with a focal reducer giving $f/4.7$. Wire crosshairs in front of the telescope produced diffraction.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