Constructing High-Resolution Lunar Mosaics

by Frank Barrett

May 30, 2004

Constructing High-Resolution Lunar Mosaics

by Frank Barrett

Why do Lunar Mosaics?

Consider the Moon. Whether you're a visual observer or a deep sky imager it always seems to be in the way. It's always filling the night sky with its sun-reflected light pollution. But on those crystal clear nights when the Moon is "ruining" the show have you ever stopped to observe those photons? They really are putting on quite a show all by themselves! The Moon offers some truly spectacular views and its face changes every night. It is the beauty of the lunar landscape that lures many into capturing its natural aesthetics in high-resolution digital images.

One problem with imaging the Moon is that it is so large, about 30 arc minutes in diameter. To capture all the detail the Moon has to offer requires imaging at a fairly long focal length - the longer the better. The limitation comes, of course, from the *seeing conditions*, that is to say how steady the atmosphere is at your location. The steadier the seeing, the longer the focal length you can get away with before your images start to look like a blurry mess. But with longer focal length comes the challenge of field of view. Depending on the chip size in your imaging device you will not likely be able to take a full phase image of a large object like the Moon in one frame unless you use smaller focal length optics.

It would seem that image detail and field of view are trade offs. If you want a large field of view (at least 30 arc minutes) you must compromise on image detail, and if you desire image detail, you must compromise on field of view. Actually there is a solution that allows you to maximize your focal length for the seeing conditions and capture the full content of the current lunar phase. The solution is to take many overlapping frames of different areas of the Moon and stitch them together into a single seamless image. This result of this process is called a ***mosaic***.

This article discusses a proven technique for constructing high-resolution lunar mosaics. Much of the work in developing this technique was done using a CCD camera and so the article is slanted with a CCD bias. Nevertheless, with a bit of thoughtful creativity and application, most, if not all, of the techniques presented here should be applicable to imaging with a webcam, digital camera, and/or film.

Preparing for the Mosaic

The first question to be answered is "How many images are needed to complete the mosaic?" The answer to this question helps you gauge how much work is going to be required in terms of both image acquisition and image processing. If the number is too high (say over 50) you might want to consider a shorter focal length (or a larger camera). On the other hand if you are up to the task there really is no limit to the number of lunar images that can be stitched together into a mosaic. The number of images needed is a function of:

1. the effective focal length of the optics used
2. the image detector size
3. the amount of image overlap allowed between image segments

You can calculate the field of view of your imaging device using this formula:

```
FieldOfView_X = ((detector_size_X * 57.296) / focal_length) * 60

FieldOfView_Y = ((detector_size_Y * 57.296) / focal_length) * 60
```

The detector size and focal length are specified in millimeters and the field of view is given in arc minutes. Finally, take into consideration the amount of image overlap. To ensure no gaps are present in the completed mosaic you should overlap your image segments in amounts of 15% to 25%. It is better to err on the side of too much than too little. Subtract this percentage from the field of view. Assuming a 20% overlap and a full lunar disk, the number of images in the x and y dimension then becomes a simple calculation:

```
Images_X = 30 / (FieldOfView_X * 0.80)

Images_Y = 30 / (FieldOfView_Y * 0.80)
```


Round these numbers up to the next whole number and multiply together to get a good starting estimate. As an example, consider an SBIG ST7 CCD camera on a Celestron C8. That's a focal length of about 2000mm and an image size of 6.9mm by 4.6mm. If you do the math you will see that it takes 20 images to capture the entire disk of the full Moon.

Acquiring the Image Segments

The process of acquiring the image segments is relatively straightforward. There are plenty of techniques that will work. Common sense should be the best guideline; just be consistent. Some key points to keep in mind:

1. Be careful not to leave any gaps; you want complete coverage! One way to be completely sure you have complete coverage is to construct a rough mosaic as you go. Also, if you have a mount capable of begin driven by a computer, note that some planetary programs, such as Software Bisque's *The Sky*, have a mosaic tool to assist in acquiring mosaic segments.
2. Ensure that the camera does not rotate at all during the acquisition process. Any rotation will complicate the stitching process!
3. Use a filter. The Moon is quite bright. A Moon filter or some sort of neutral density filter is almost always required to prevent the images from saturating. The Orion Moon filter works well as do variable polarizing filters.
4. Line up your camera such that the x and y axis are parallel with the RA and Dec axis of your mount. It would seem natural to align the camera on the Moon's terminator, but you will find that it just complicates the acquisition process. Lining up the camera with the compass directions simplifies the process since movement in just one direction is all that is needed to acquire the next segment.
5. Work hard on good focus. You want a high-resolution image, right? If the temperature is changing rapidly you may need to refocus occasionally.
6. Check the collimation of your optics. If you are having a particularly hard time focusing, it could be that your optics are out of alignment.

7. Use the shortest exposure time possible. You are trying to catch a moment when the air is still. A shorter exposure time increases the chances of that happening.
8. Take your images when the Moon is high in the sky. You are shooting through less atmosphere and the seeing effects will be minimized.
9. Take multiple images of each segment. Depending on the seeing conditions you may want to take between 5 to 10 images (or more) of each segment. Having multiple images of each segment allows you to evaluate and select the crispest image.

Preparing the Image Segments

If you take multiple images of each segment, you can either combine these images or just select the crispest image of the series. The decision should be based on the signal-to-noise ratio of the image set. If you image with a CCD camera that is thermoelectrically cooled and have properly collected ample dark and flat frames, you can probably get away with using a single image for each segment. Otherwise, you may need to align and combine your images to get a reasonably good signal-to-noise ratio. A high signal-to-noise ratio will allow more aggressive sharpening to the completed mosaic and will also result in a more natural blending of individual segments.

After you have processed the images for each segment down to a single image the next step is to normalize the histograms for the image set. The goal here is to set the dark point and white point in exactly the same spot for each image segment. In other words, we want to determine a common dark point and white point such that the image is not clipped on either the dark or the light end of the histogram for any image in the set. If you are unfamiliar with some of this terminology, perhaps some definitions will help. The image ***histogram*** is a representation of the number of pixels in an image for a given brightness value range. For example, in Software Bisque's CCDSoft program, select the image and press **Ctrl-H** to display the histogram. The ***Dark Point*** is the point in the histogram below which all pixel values are displayed as pure black and conversely the ***White Point*** is the point in the histogram above which all pixel values are displayed as pure white. Any image data below the black point or above the white point is lost since it is rendered as pure black or pure white. This is called ***clipping*** and is perhaps the most misunderstood and abused aspect of digital image processing.

Since we are dealing with multiple image segments we want to be sure that the histogram of the *combined* mosaic is not clipped on either end. The first step is to set the dark point and white point for each individual segment such that the image histogram is not clipped (see **Figure 1**). Then tab through each image and take note of the lowest black point and the highest white point values. The minimum black point and maximum white point will most likely not be from the same image segment. Now go back through the all the images and set each image's black point to the minimum value and the white point to the maximum value. Don't become too concerned if the images start looking too bright or too dark. What you are doing with this process is capturing the mosaic's total tonal dynamic range. You may be surprised in the end just how much real information is hidden in all those individual segments! As a final step save the image segments in TIFF format.

Normalizing the histograms is a vital step and is perhaps one that many imagers constructing mosaics overlook. Failure to properly execute this step may result in some areas of the finished mosaic to be burned out and other faint areas to be lost into the darkness. Once the information is clipped it can never be retrieved, so be careful!

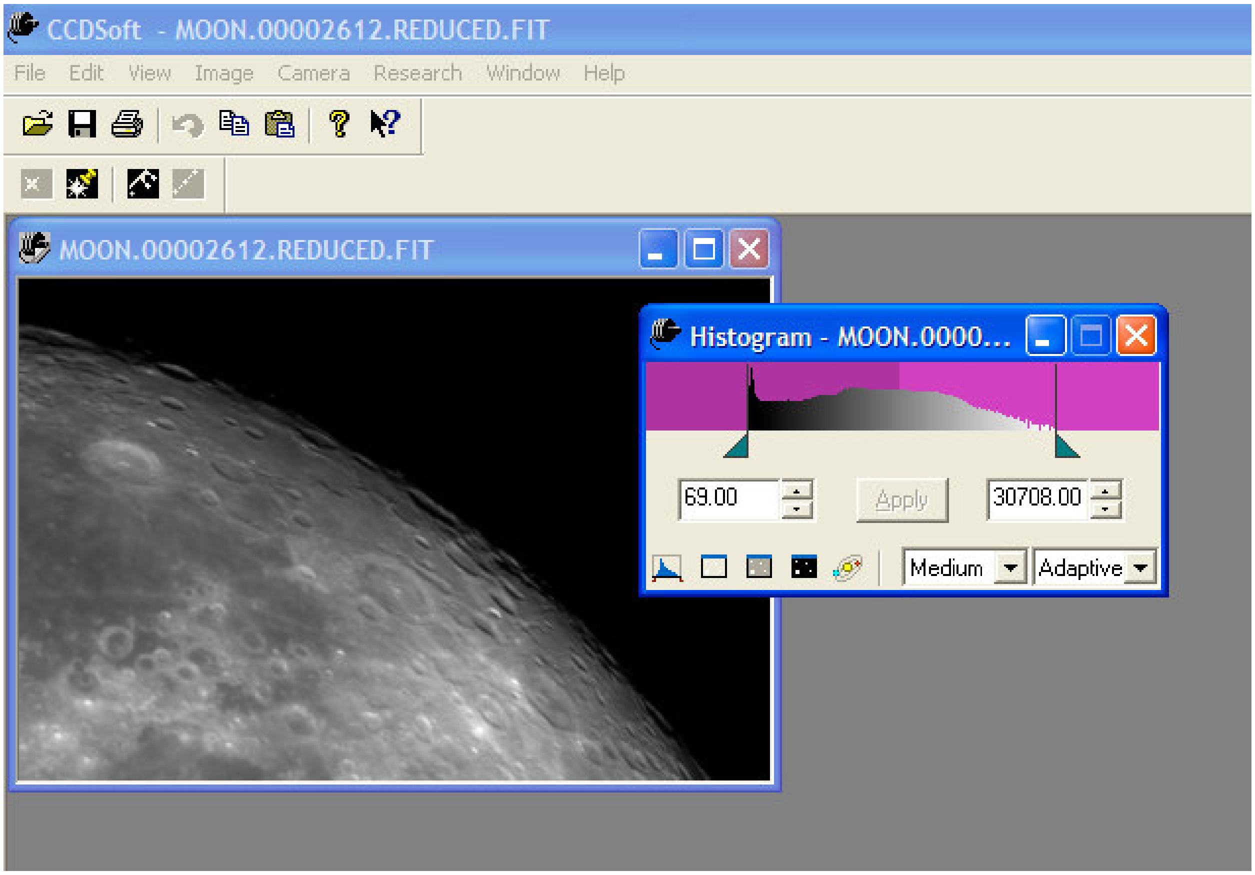

Figure 1. Adjust the histogram for each image segment so that image data is not clipped on either end. As you do this take note of the smallest dark point value and highest white point value. Set the histogram of all image segments to these two values so that the resultant mosaic will not be clipped on either end.

Stitching the Mosaic Together

Now comes the fun part, putting all those segments together. This discussion assumes the processing is performed in Adobe Photoshop. The first step is to create a canvas large enough to contain all the segments (**File** | **New...**). Calculate this from the maximum number of segments in each axis and the pixel dimensions of your camera. It is better to err on the side of being too large. This isn't critical though as Photoshop will allow you to enlarge the canvas in any direction needed (via **Image** | **Canvas Size...**), but it's good to get this right from the beginning. When creating the image it's a good idea to set your background color to pure black initially fill the new canvas.

Now let's add our segments. It's generally a good idea to follow the same order in which the image segments were acquired. As long as you are laying down segments that are adjacent to one another alignment should be pretty easy. Open your first image segment and copy it to the clipboard. In Photoshop this is accomplished by pressing **Ctrl-A** to select all (or **Select** | **All**), and then **Ctrl-C** to copy (or **Edit** | **Copy**). Then click on the mosaic canvas and paste this image in, **Ctrl-V** (or **Edit** | **Paste**). Notice that Photoshop creates a new layer. Use the Move tool to place this image in the approximate correct location. Open the second image segment and repeat the process. Use the Move tool to align the segment as closely as possible to its correct position over the previously pasted segment. Zooming in on the overlapping region can do this fairly accurately. For fine adjustments use the keyboard cursor arrow keys to move the layer in single pixel increments.

Here's a trick that can virtually guarantee perfect pixel alignment of the two layers. In the Layers view change the blending mode (it's the drop down menu that typically says 'Normal') of the layer on top to **Difference**. You will see the overlapped area between the two image segments become dark, perhaps totally black (see **Figure 2**). Using the cursor keys and the Move tool, move the segment up and down and left and right until the dark area looks completely black (see **Figure 3**). It may be that this area won't go completely black. That's quite all right; the images can be offset by as much as a half a pixel width relative to each other and it may not be possible to achieve perfect alignment. Just use the setting that results in the darkest area. Then change the blending mode of the layer back to **Normal**. A real purist could achieve perfect alignment by resampling each segment by 2x or more and performing the alignment on the enlarged images. Generally speaking this is not necessary. When we blend the seams perhaps you'll see why.

Continue this process with all remaining segments. When you align a layer over multiple layers below you may discover that you cannot get everything perfectly aligned. For example, moving a pixel left makes one area darker and moving to the right makes another area darker. Again, don't get too hung up on this. Use the setting which gives the darkest overall area and you will achieve a good result.

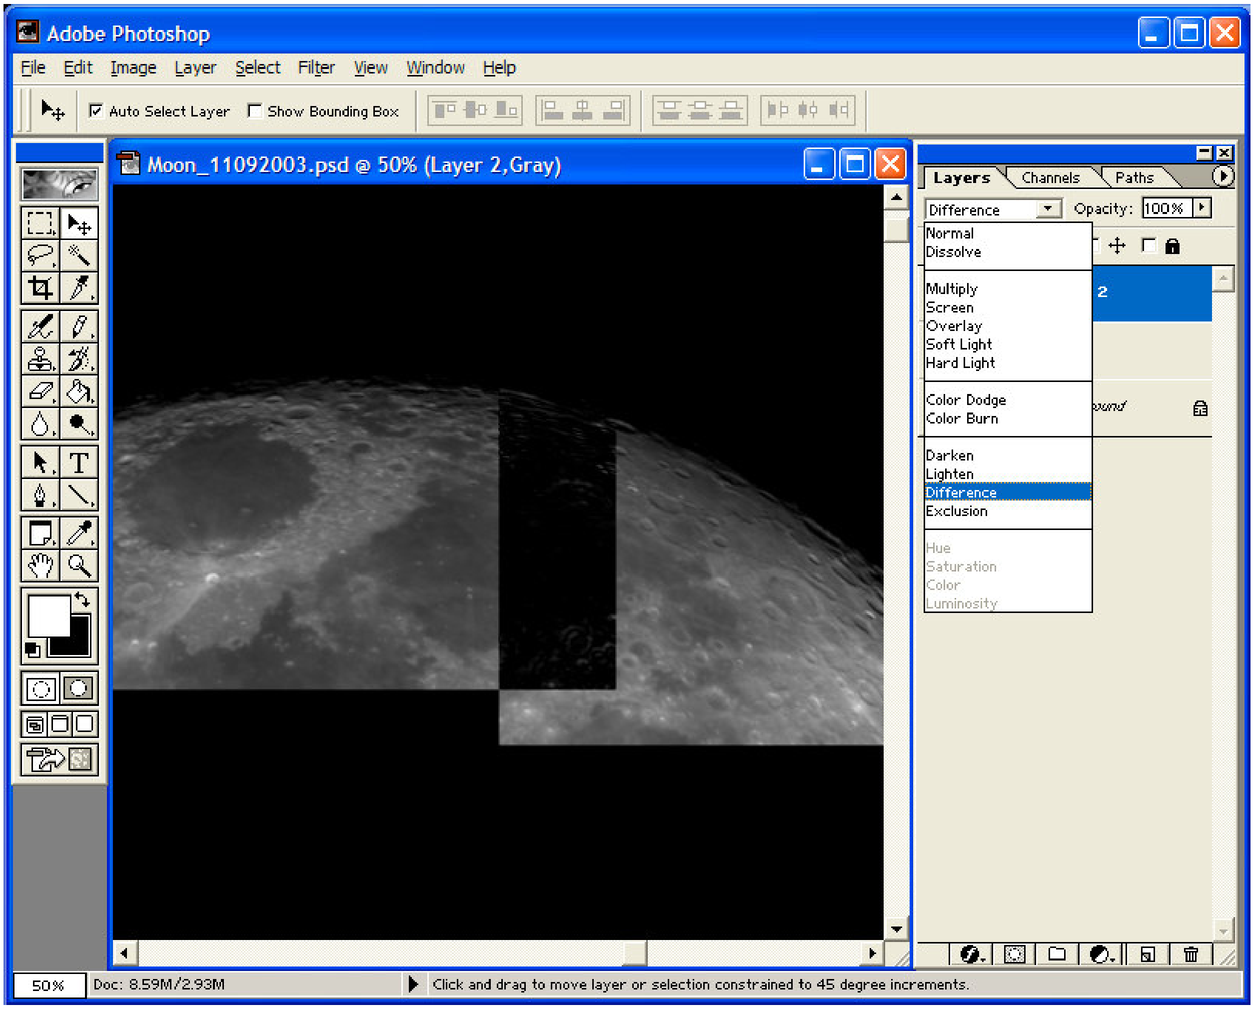

Figure 2. Set the layer blending mode to "Difference". This shows where the two images are not properly aligned.

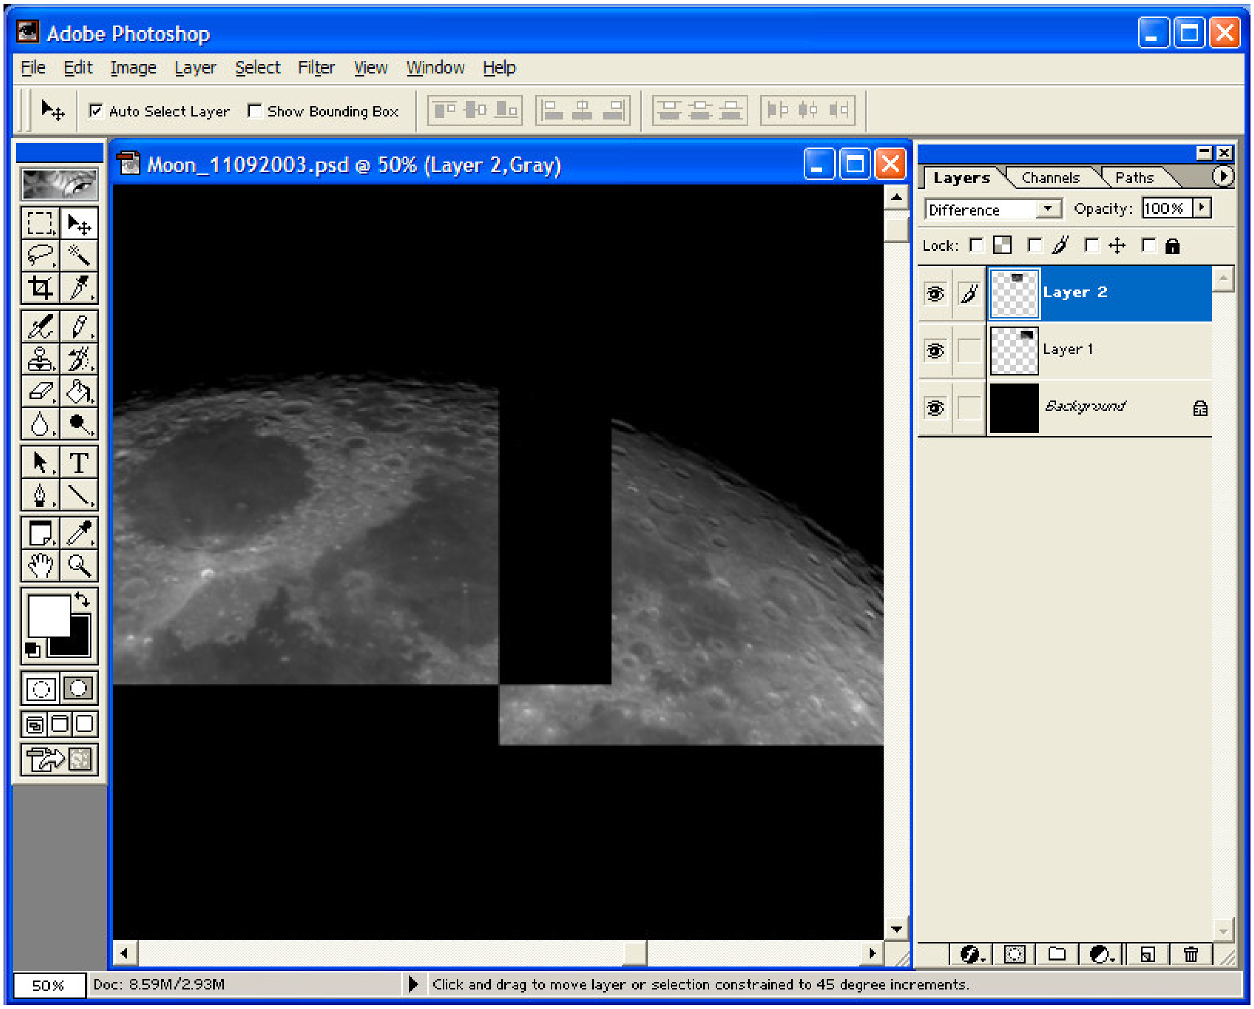

Figure 3. Using the Move tool, nudge the layer until the dark area is completely black. Then set the layer blending mode back to "Normal".

Blending the Seams

Before proceeding it's a good idea to save your work with a different filename just for insurance (**File | Save As...**). If an error is made later in the process you can return here without having to go through the alignment process again. If you have acquired a good set of images and done a good job on the alignment you probably are not noticing many seams. If you do happen to notice a few seams take heart as these are easily hidden using a simple technique.

First you should evaluate the overall brightness levels. Does it look like a particular segment stands out as too bright or too dark compared to those around it? If so, select that layer and adjust the brightness and contrast to bring it in line with its neighbors (**Image** | **Adjust** | **Brightness/Contrast**). If it looks like you are able to fix one part of the segment, but another portion is incorrect, use the Lasso tool and select the offending portion, feather it a generous amount (**Select** | **Feather...**), and adjust the brightness and contrast of that portion by itself. With enough tweaking you should be able to achieve a perfect blend.

It should be pointed out that if you are using a CCD camera and have properly calibrated your images and normalized the histograms as noted earlier, you should not have to adjust the brightness of any of your segments. You may, however, have a few minor alignment problems to clean up. This is done by using the Eraser tool with a fairly large radius and soft brush to erase the edge of each of the overlapped image layers. This will soften the edges and blend the layers nicely. But how can we see this edge to erase it?

Here's the trick. Select the top layer. Add an Adjustment Layer for Brightness and Contrast (**Layer** | **New Adjustment Layer** | **Brightness/Contrast...**). Be sure to choose "Group With Previous Layer" (see **Figure 4**). Increase the brightness of this layer by about 50% (see **Figure 5**). Don't worry about how the image looks; we're going to delete this layer in a moment anyway. What this layer does is show us the image segment's edges. Use the Eraser tool with a soft brush and fairly decent size radius. For starters, try a radius of 45. Select the layer with the image segment. It will be the layer just below the adjustment layer we just added. Note that Photoshop will leave the adjustment layer selected by default so don't forget this step! Carefully erase the edge that overlaps the image below. You'll see the edge softening as you erase. That's the effect we're after (see **Figure 6**). Take care to only erase the edge. Do not penetrate too far into the image or you could erase too much. You need not soften the edges that extend out into the black sky. When you are happy with the result, delete the adjustment layer...it has served its purpose. Be sure the adjustment layer is selected, and then delete the layer (**Layer** | **Delete Layer**).

Repeat this procedure with the remaining layers. You will notice as you work your way down that the adjustment layer will show you hard edges and soft edges (see **Figure 7**). The soft edges are edges you have already softened in the layers above. These do not require erasing. Doing so may result in eliminating image data. Only erase the hard edges! If at any time you see a completely black area as the result of your erasing it means you have erased too much data. Use undo immediately (**Edit** | **Undo**)! You will notice when you get to the bottom image segment that it has no hard edges. That's because in a perfectly overlapped mosaic the bottom image segment has no internal edges visible from the top and thus will not need any softening. Save off a copy of the layer stack at this point. Now flatten the image and save again with a unique filename.

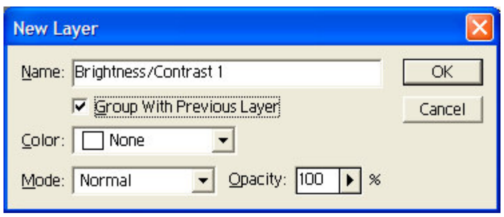

Figure 4. Adding a Brightness/Contrast Adjustment Layer makes it easy to see the image segment's edges. Be sure to check "Group With Previous Layer".

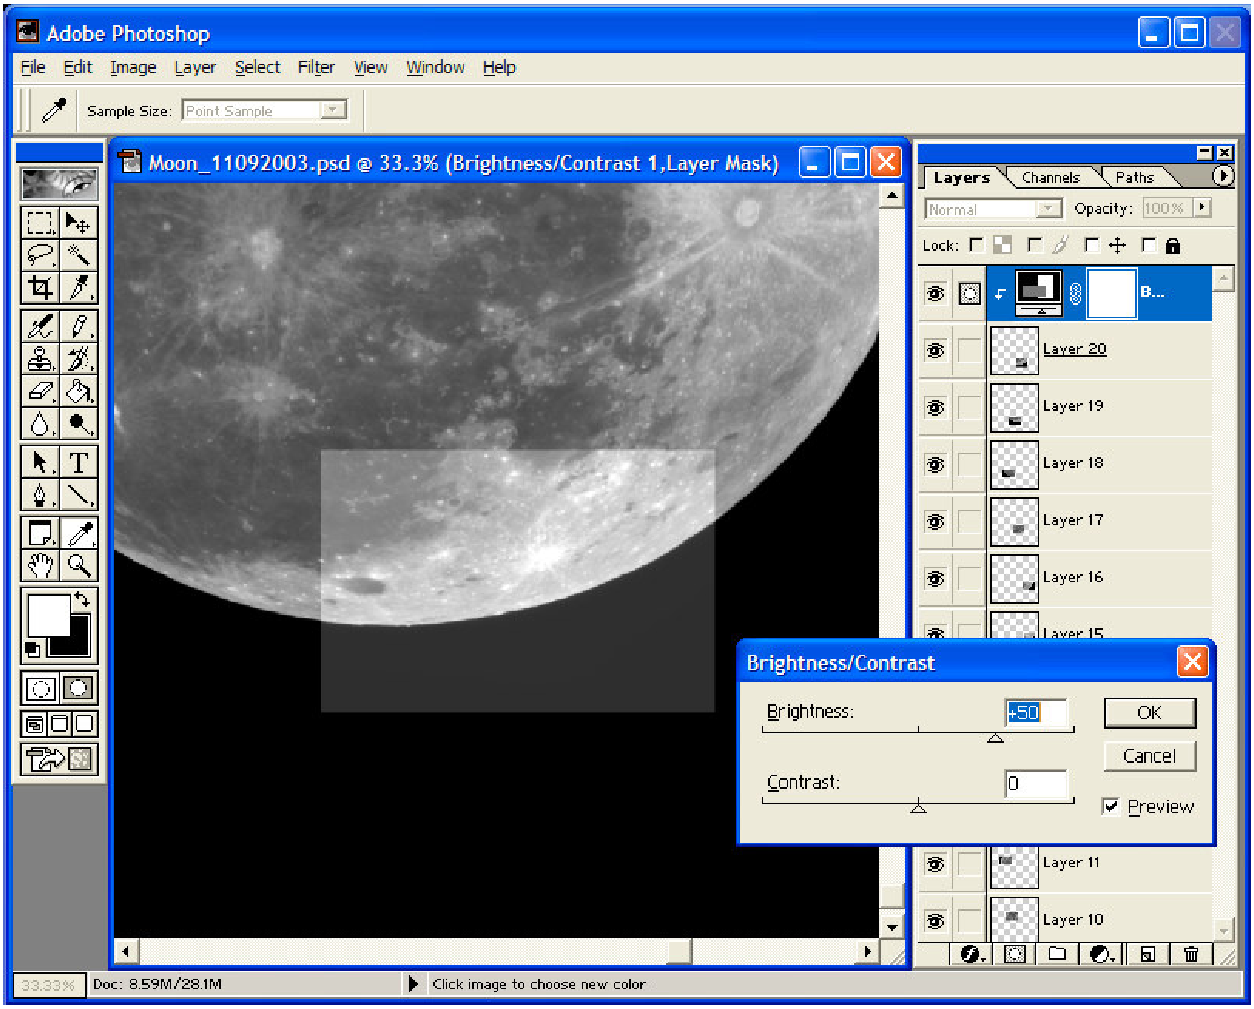

Figure 5. Adjust the Brightness up by 50 percent. Notice how this shows us the edges of the image segment.

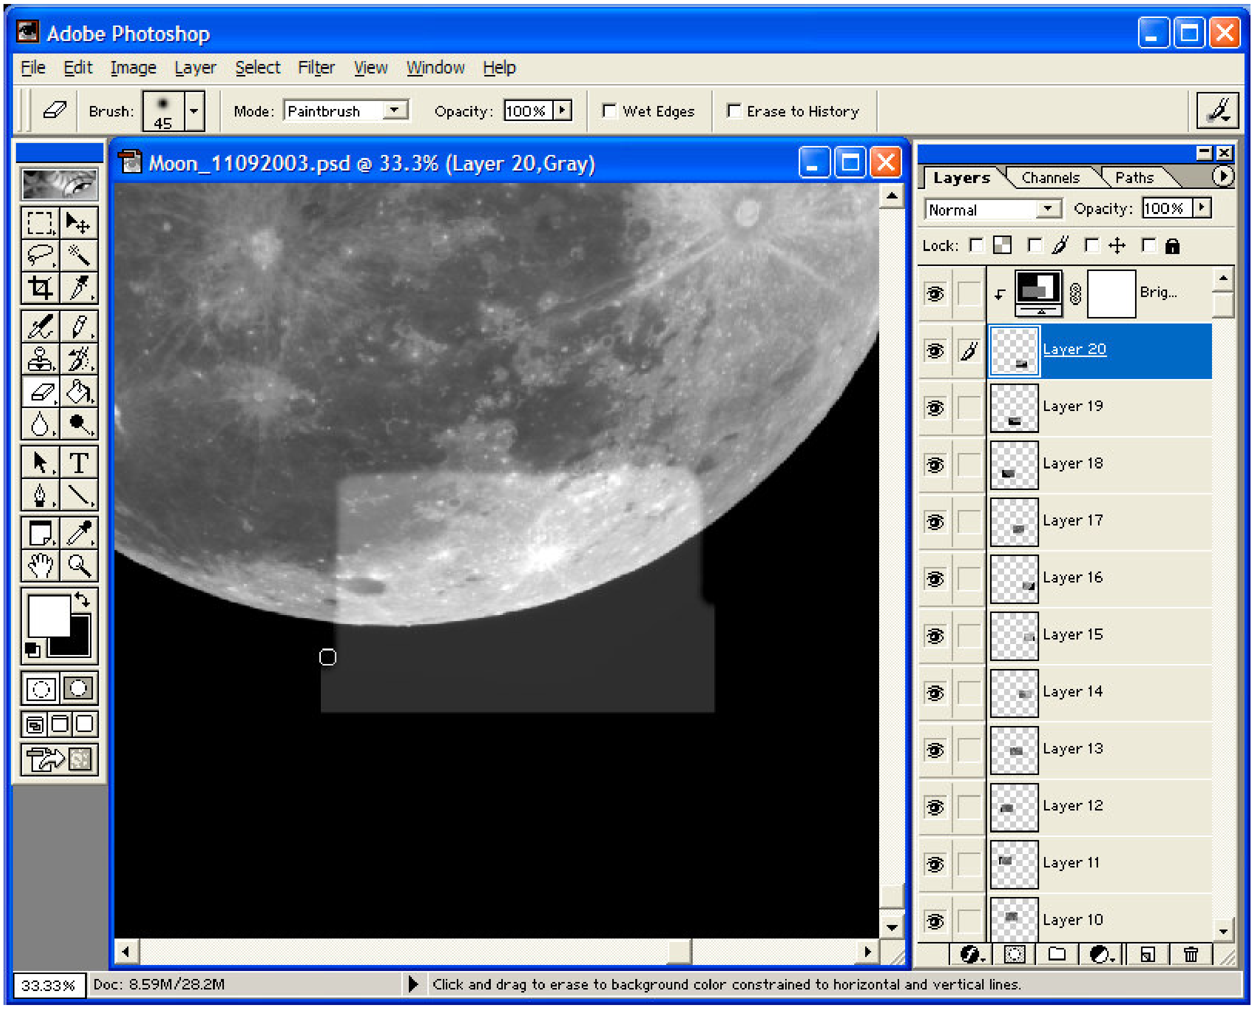

Figure 6. Use the Eraser tool with a soft brush to soften the edges. Be sure to select the image layer before erasing and don't erase too much!

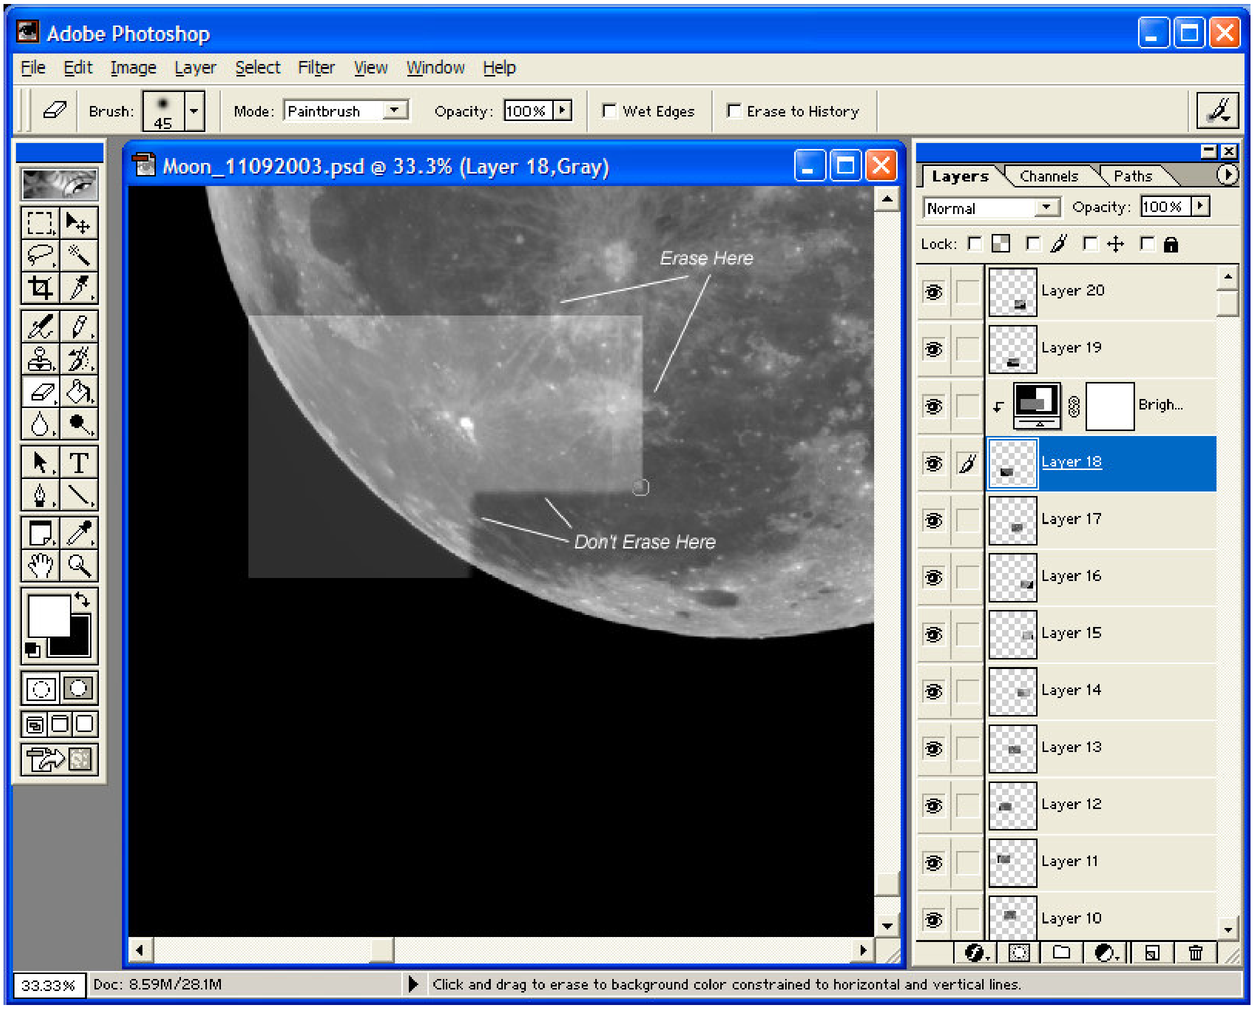

Figure 7. When softening edges in lower layers, be sure to erase *only* the hard edges. If you start to see the black canvas underneath, undo immediately!

Sharpening and Deconvolution

High-resolution images almost always require some level of sharpening. It's the price we pay for having to image through a turbulent atmosphere. There are many ways to improve the resolution of your mosaic. The nuances of the many image sharpening and deconvolution techniques are way beyond the scope of this article. However, if you have Photoshop you need not look any further than the Unsharp Mask (**Filters** | **Sharpen** | **Unsharp Mask**). The best way to learn sharpening techniques is to experiment. Try using several iterations of the Unsharp Mask (see **Figure 8**). First use a large radius to sharpen large features and then use progressively smaller radii to sharpen smaller features. Each time you apply the

Unsharp Mask duplicate the top layer (**Layer** | **Duplicate Layer…**) and use the layer properties of the previous layer to document the parameters used for that layer (**Layer** | **Layer Properties…**). For example, "Unsharp 135%/0.9/0." That way if you ever decide you went too far or not far enough, you know what you did earlier, and you have the option of returning to the original image data and starting over.

Use moderation. When sharpening, it is easy to go too far. You know when you've gone too far when the image starts to look garish or noise and graininess become noticeable. When you get to a point you like, it is a good idea to back off just a bit. It's always better to have a slightly soft image than one with too many artifacts.

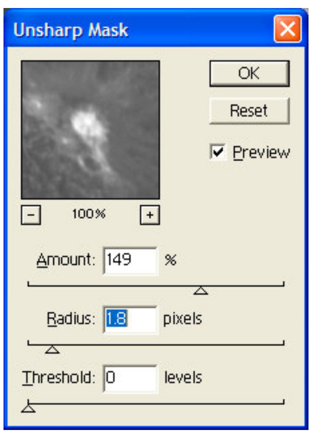

Figure 8. Despite its name, the Unsharp Mask is very useful for sharpening the image. Experiment with Amount and Radius until you get a good result. Try toggling Preview on and off to see the before and after effects.

Final Tweaks

At this point you should be pretty happy with the image. Sometimes it helps to add a Levels and/or Curves Adjustment layer to tweak the image contrast. Be careful not to go too far that you burn out highlights or lose dim shadow areas. As an optional step you may wish to rotate the image so that the poles and/or terminator line up vertically (**Image** | **Rotate Canvas** | **Arbitrary...**). In all likelihood your mosaic did not end up perfectly centered on the canvas. The best way to fix this is to simply use the image crop tool to cut away excess canvas area. Save these intermediate steps with unique filenames.

As a final step flatten the layers (**Layer** | **Flatten Image**) and save off as a JPEG file to share with other folks on the internet (**File** | **Save As...**). If your image is large you may want to resample the image to smaller dimensions so that it will fit in a browser window (**Image** | **Image Size…**). You may also want to post a link to the larger image as well. See **Figure 9** for a final processed image.

Figure 9. The final processed image. The full resolution image can be found in the author's gallery at **http://celestialwonders.com**.

Conclusion

Creating lunar mosaics is fun. It is, perhaps, a bit more work than some are willing to undertake, but there is a feeling of great satisfaction when the work is complete. The technique works quite well. It is by no means the only way to go about constructing a lunar mosaic. Try the technique and see what kind of results you can achieve. Experiment. Above all, don't allow fear of failure or laziness rob you of the satisfaction of creating high-resolution images of our beautiful Moon! Have fun with it!

You can find more examples of lunar mosaics constructed with this technique at the author's gallery at ***http://celestialwonders.com***.

Biography

Frank Barrett (frankb@celestialwonders.com) works by day as a Senior Software Engineer. He has been an Amateur Astronomer for over 30 years and has been most keenly interested in astrophotography for the last two years.