
Using Picture Window in Astrophotography

Written by Jonathan Sachs
Copyright © 1998,1999 Digital Light & Color

Introduction

Picture Window can be used to combine multiple photographs of the same region of the sky to create enhanced images. Several common situations that come up in astrophotography are addressed in this document:

- Combining two or more exposures to obtain a single image equivalent to one long exposure.
- Combining black and white photos taken through red, green, and blue filters to produce a color image.
- Using a short exposure image to bring out highlight detail washed out in a long exposure image (e.g. in a photo of a globular cluster).
- Creating a mosaic from two or more overlapping images.

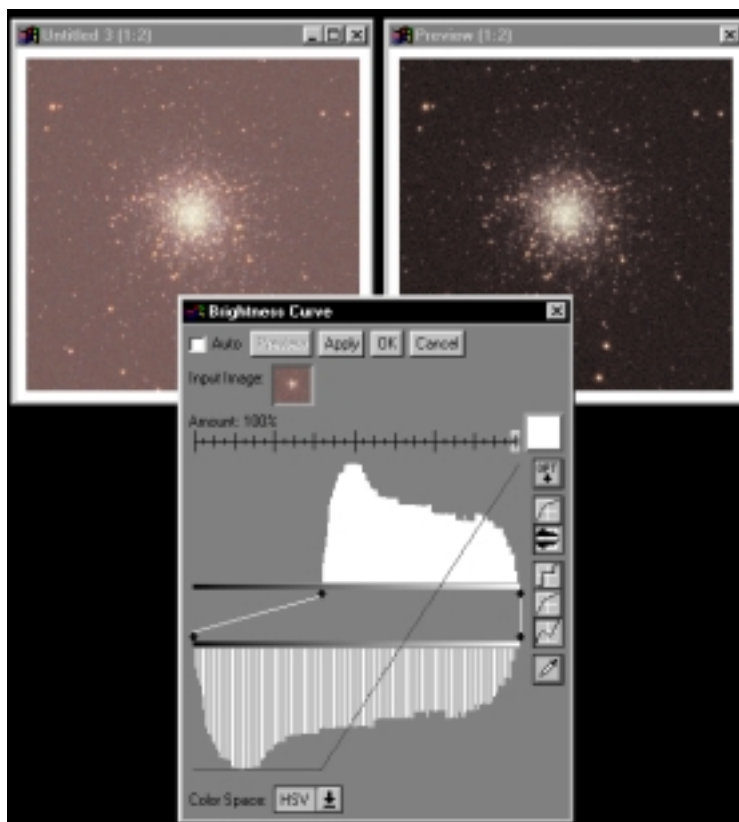
Finally, techniques for overall image enhancement are also covered.

Preparing images for use

Before images can be combined, they must be normalized so that they cover a consistent range of brightness values. There are two ways to do this using Picture Window:

The Manual Method

For more precise control of the process you can use the Gray/Brightness Curve transformation. The basic idea is so make the background sky black and the brightest objects in the photo white. This gives you the maximum usable range of data to work with.



This example illustrates a typical raw scan of a negative in the upper left corner. The Brightness Curve dialog box has been set to Log Histogram (via the Opt button just to the right of the curve) and Histogram mode. The histogram of the image illustrates that the darkest data in the image is much too light. By setting the upper left marker at the level corresponding to the darkest image data and the lower left marker at black (all the way to the left), this creates a curve that stretches the contrast of the image as illustrated in the preview in the upper right. In this particular example, the brightest image data is already white and needs no further correction, but if it were not, you could make it white by setting the upper right marker to the brightest level in the image and moving the lower right marker all the way to the right.

After adjusting the dynamic of the range of the image, you may also need to correct its color balance. This adjustment is normally made by assuming that the brightest spots on the image should be a pure white. If they are some color other than white, the image can be corrected using the Color Balance transformation. Color Balance is also capable of performing an automatic adjustment of the dynamic range, but for the greatest control, the method described above is the more accurate.

To color balance an image, click on its window and select Transformation/Color/Balance from the main menu. The default correction will make the lightest colors in the image white, so you can just click OK to accept these settings. To expand the dynamic range as well, click on the Opt key at the top of the Color Balance dialog box and select Full Range. This will adjust the settings so the dynamic range runs all the way from black to white and the lightest areas of the image will be corrected to white if necessary. For a full description of other color balancing options and techniques, please see the manual or help file.

The Automatic Method

A quicker and simpler way to expand an image to use the full dynamic range is to use the Gray/Levels and Color transformation if the image is black and white or the Color/Balance transformation if the image is color. Then click the Opt button in the upper right corner of the dialog box and select the Full Range option. Finally, click OK. For color images, this not only makes the image use the full range of pixel values but it also makes the image run from pure black to pure white by removing and color from the brightest objects and the sky background.

The automatic method works by scanning the image to locate its brightest and darkest areas and their colors. For this to work properly, it is important that the image is properly cropped so as not to include any of the film around the edge of the image

or any other extraneous image data. Rarely, the algorithm may fail, in which case you can use the manual method described above.

A word about numerical accuracy

When combining multiple images to lengthen the effective exposure time, you will get more accurate results if you scan and work with 16-bit black and white (or 48-bit color) images. This lets you accumulate image data without having to worry about truncation errors that occur in the blending process when using 8-bit images.

Registering images

Regardless of the task, the first step is to accurately register all the images so their data can be combined pixel for pixel. Registering two images simply means accurately aligning one photo so it matches the other. One image must be arbitrarily designated as the base image and the other as the overlay image. Registering the two images causes the overlay image to be warped to align it with the base image.

The primary workhorse for registering images with Picture Window is the Composite transformation. Here is the basic procedure:

Step 1: Open the images and select Composite

Starting with each image in its own window, click on the window of the file you want to use as the base image and then select the Transformation/Composite command. This pops up the Composite dialog box. Leave the Amount slider and the Overlay Amount slider both set to 100%.

Step 2: Select the overlay image

Go down to the Overlay control and click on the small white square. This pops up a menu of image windows; select the overlay image from the menu.

Step 3: Select the alignment option

The simplest case is one in which two images can be registered simply by shifting and rotating the overlay image to match the base image. This case is described first;

the case where additional image distortion needs to be corrected will be dealt with below in the section on multipoint registration.

To combine two images using only shifting and rotation, select the *2-point (shift/rotate)* Alignment option. As soon as you do this, Picture Window displays two numbered alignment points over the base image window and the overlay image window. The initial positions of these alignment points are in the lower left and lower right corners of each image.

Step 4: Place the alignment points

The general idea is to position alignment points on the same features (typically stars) in each image.

Assuming 2-point shift/rotate alignment, when Picture Window composites the images (using more than 2 points is discussed later), it shifts the overlay image to make alignment point 1 in the overlay image coincide with alignment point 1 in the base image, and then rotates the overlay image so that the line between alignment point 1 and alignment point 2 in the overlay image points in the same direction as the line between alignment point 1 and alignment point 2 in the base image. For the greatest accuracy, try to position the two alignment points as far apart as possible.

Picture Window pays attention to the precise location of each alignment point, even down to a fraction of a pixel. By zooming the images beyond 1:1, you can enlarge the pixels on the screen sufficiently to permit placing alignment points within a pixel.

As a convenience in setting alignment points when zoomed in, you can use the keyboard shortcut of pressing *1* or *2* on the keyboard to recenter the display of the base and overlay image windows on the corresponding alignment point.

Step 5: Refining the alignment point locations

Once you have gotten the alignment points close to the same feature in each image, you can use Picture Window's automatic refinement feature to compute the optimal locations of the alignment points down to a tiny fraction of a pixel. Simply click the Refine button (located just below the Alignment control) and Picture Window will search for the location for each alignment point in the overlay image that results in the best match with the base image. For this to work, the initial locations of each alignment point must be within 5 pixels of the optimal location, and the alignment

points in the base image must be set on prominent features (as opposed to featureless areas that are all more or less the same color).

Step 6: Checking your results

To check how well the two images are aligned, set the Operation to Absolute Difference and then click Preview. This displays the absolute value of the difference between the two images in the preview image. The better the match between the two images, the closer to black the preview image will be. You can zoom in on the preview image to see fine detail as necessary. If the images are misregistered, double star images should be visible as you zoom in. In most cases there will be some residual difference due to statistical variations, differences in exposure time or filtration, etc. There will also be stars visible anywhere in the preview image that the base image is not covered by the overlay image (i.e. if the overlay image does not completely cover the entire field of view of the base image).

Step 7: Generating the final registered image

To produce the final, registered version of the overlay image, set the Operation control to Register and then click OK.

Step 8: Sharpen the registered image

A side effect of the bilinear interpolation of the overlay image performed when it is warped is to soften the result image slightly. The Sharpen transformation can restore most of the sharpness lost by resampling.

Combining two images to simulate a long exposure

Once you have your registered images, combining them is relatively easy. To combine two exposures of the same object to obtain the effect of a single, long exposure, you simply composite the two image using no alignment points, and the Blend operation with the Overlay Amount slider set to 50% (assuming the exposures are of equal duration). If the exposures were of different lengths, the slider should be set accordingly to blend the two image in proportion to the exposure times. A value of 100% gives all the weight to the overlay image; a value of 0% gives all the weight to the base image. For example, if the base image had an exposure time of 2 hours and the overlay image had an exposure time of 1 hour; set the slider to 33.3% (the percentage of the total exposure time due to the overlay image).

Combining more than two images

If you have more than two exposures to combine, the best strategy is to combine them pairwise. For example if you have four images, combine the first and second and then the third and fourth; then combine the two results. If you have an odd number of images, reduce the weighting of the last one to reflect that it represents a shorter exposure than the others. For example, to combine three images of equal duration, combine the first two at 50% and then combine the result image with the third at 33.3%.

Combining red, green, and blue images to make a color image

To combine three black and white images of the same part of the sky taken through red, green, and blue filters into a single color image, use the Transformation/Gray/Combine Channels transformation. Please refer to the manual or help file for details on using this transformation.

A beneficial side effect of registering the three color channels is to reduce chromatic aberration which results from the red, green, and blue images having slightly different magnifications. When the images are registered, they are automatically scaled to the same size, thus reducing or eliminating color fringes around the edges of the image.

Using multipoint alignment

In an ideal world, two astrophotographs could be aligned by simply shifting and rotating the overlay image with respect to the base image. This assumes negligible optical distortion in the telescope, negligible nonlinearity in the scanner used to digitize the images, and high dimensional stability of the film. Any of these factors can result in small distortions which degrade the result image if they are not compensated for.

In some cases 3-point or 4-point alignment may be enough to get a good match. The real solution for hard-to-match images, however, is multipoint alignment.

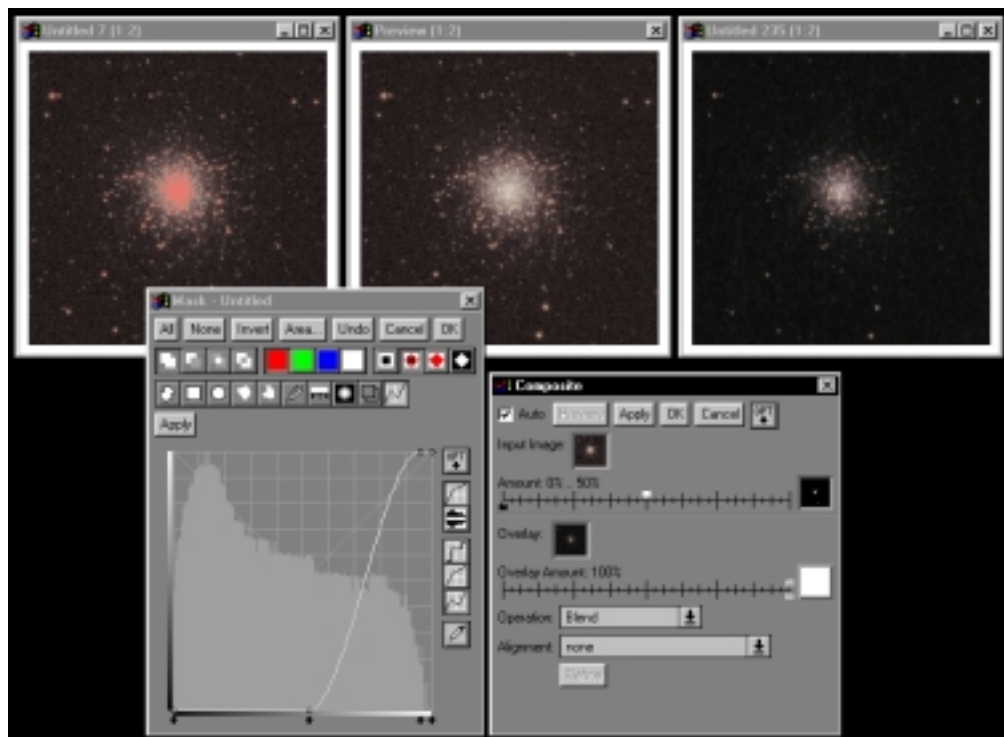
When you select the Multipoint alignment option, Picture Window initially displays three alignment points. These first three points are special as they are used to define a baseline affine (shift/rotate/ scale/shear) warping of the overlay image. Additional alignment points are then used to make local corrections to this baseline warp to allow the alignment of multiple features in each image. Four different interpolation schemes are available for determining the local corrections to the overlay image inbetween the alignment points. Use the options menu (accessed by clicking the Opt button in the upper right corner of the dialog box) to select the desired method. Depending on the nature of the distortion, you may get better results from one of these than with the others.

In multipoint alignment mode, Picture Window also displays a small tool bar just to the right of the Refine button. To add a new pair of alignment points, click the button labeled with a plus sign and then click on the input image window. To remove an alignment point pair, click on the minus sign and then click on the alignment point in the input image window.

Using a short exposure image to restore highlights to a washed out long exposure image

This procedure is a trickier as it requires using a mask to limit the application of short exposure image to the highlights of the long exposure image. For the purposes of this description, assume that the long exposure image will be used as the base image and the short exposure image as the overlay image. Before compositing the two images, you need to create a mask that isolates the highlights of the long exposure image. One of the easiest ways to do this is to create a density mask or, in other words, a mask whose transparency is based on the brightness of the image. This can

be done using the Curve tool in the Mask dialog box and selecting a curve like the following:



Since astrophotographs have highly skewed histograms, it is often necessary to use the Log Histogram option as shown in this example. This option can be selected by clicking on the Opt button just to the right of the curve. The curve illustrated above generates a mask that is black where the base image is dark or midtoned and white in its highlights. You can experiment with the transition location and drop-off rate to get the desired effect. Remember to click the None button in the Mask dialog box before trying new curve settings—otherwise the new mask will be combined with the previous mask. The easiest and most flexible way to make this work is to invoke the Mask tool on the base image before starting the Composite operation. This will automatically select the mask into the Amount control and, as you alter the mask, the effects will ripple through into the Composite operation via its Preview feature. You can further modulate the effect by moving the white slider in the Overlay Amount control -- this will alter the blending percentage for those areas where the mask is white.

Creating a Mosaic

When using the Composite transformation, the dimensions of the output image are set to match the dimensions of the input image. For this reason, if you want the output image to extend beyond the edges of the input image, you must first extend the input image to make room to drop in the overlay image. The easiest way to do this with the Geometry/Add Border transformation. This simple transformation lets you add pixels above, below, to the left, and to the right of the image in any solid color.

The next step is to create a mask for the overlay image so that it is blended smoothly with the input image in the area of overlap. This is best done with the Gradient transformation. Simply create a linear gradient that runs from black to white across the overlap between the two images and is white over the rest of the overlay image. To ensure that the gradient image is not a color image, be careful how you select the gradient colors. Start by clicking on one of the black, gray, or white patches at the bottom of the color picker and only use the brightness slider to set the gray level.

Once you have extended the input image and created a blending mask for the overlay image, you are ready to composite. Click on the extended input image to select it, select the Composite transformation from the main menu, and select the overlay image into the Overlay control and the overlay mask image into the Overlay Amount control. Finally, you need to align the two images, so select 2-point alignment and set the alignment points on two stars in the area of overlap, as far apart as possible. Use the Refine button to zero in on the best match and then click OK to composite.

If you are compositing more than two images, simply repeat the above steps for each additional image.

Final Image Enhancements

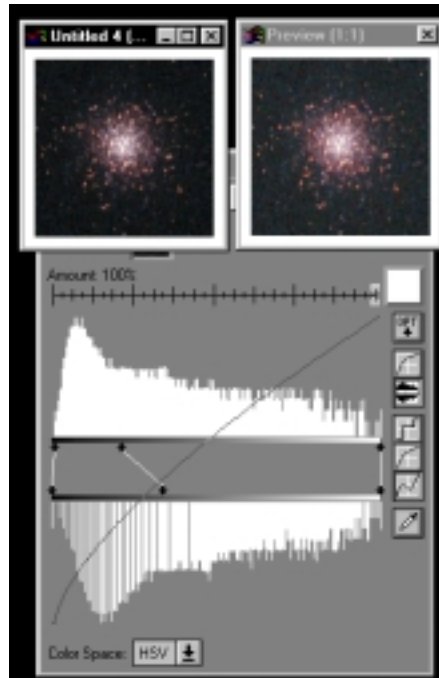
After generating a final composited image, there are a number of enhancements to consider in the interest of improving its overall appearance. They should probably be applied in the order in which they are described.

Sharpening

Judicious application of Unsharp Masking using the Sharpen Transformation can make the final image crisper. Sharpening should always be done after compositing since it amplifies noise in the image.

Gray Curve Adjustment

You can emphasize faint detail by applying a curve that increases the contrast in the shadow areas (at the expense of losing a little contrast in the midtone and highlight areas).



Saturation Enhancement

Increasing the saturation of color images using the Saturation or Saturation Curve transformation will make them more colorful.