

Combining Images


by Dennis Wilkins

Picture Window and Picture Window Pro provides a very flexible set of functions for combining images that can be used to make panoramas. The **Composite** transformation has a number of alignment features that enable “splicing” images that do not line up perfectly. When making a panorama, the rotation of the camera causes a change in image perspective. There are a number of automated panorama programs on the market that attempt to correct the perspective and blend each successive image to make a seamless whole. I own a couple of these programs, and have tested a few more, and they generally work pretty well if you rotate the camera through the nodal point of its lens. But I have found in some cases, even using nodal rotation, they cannot seem to make a clean blend. And since they typically use a “lap-dissolve” with the processed image fading gradually from one input image to the next, the results can be less than desirable.

Picture Window and Picture Window Pro will not automatically stitch two or more images together, but the extensive control provided for alignment and even perspective control when combining images makes stitching panoramas easy and more precise than most automated stitching programs.


Stitching provides more than just making wider angle shots . . . it also increases the pixel count of an image. Stitching only two images together from a three megapixel camera can make a five megapixel image with about 20% overlap. And Picture Window doesn't really require any overlap, although you need enough to see how to align the two images.

The shots below were made in the eastern Colorado plains, of a landmark called the West Pawnee Butte. This butte was a major signpost for the early pioneers traveling west. These shots were made right after sunrise one morning, with the camera set to manual exposure so that there would be no shift in density from one image to the next.

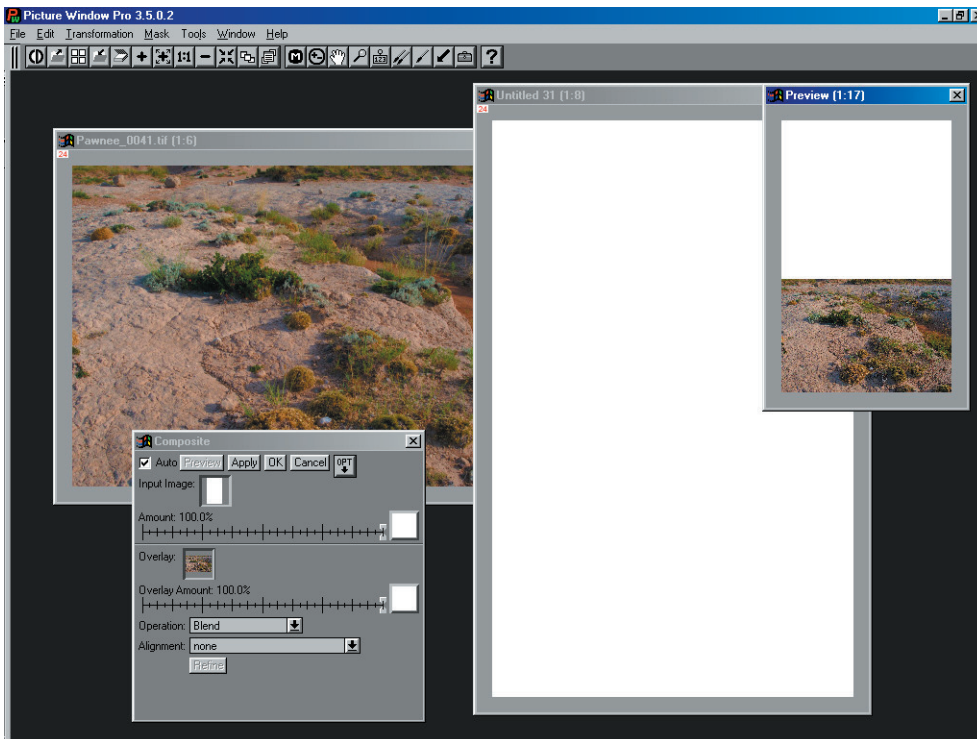


This is not the usual “left to right pan” but is a vertical pan. Before stitching these images together, we need an image window that has a space large enough for the combined three images. The width is known (3008 pixels in each case), but the resulting height will be less than three times the 2000 pixel height of each image because there is considerable overlap. Overlap was used here in order to enable a couple of automatic panorama programs to process this image – neither one produced an image as “clean” as Picture Window, and both required a lot of manual adjustments and clean-up work. Overall, they took as long to complete the panorama as it took me using Picture Window.


The first order of business is to create a new window, 3008 pixels wide by about 4800 pixels tall, more than enough to include the height. This is easy to do . . .



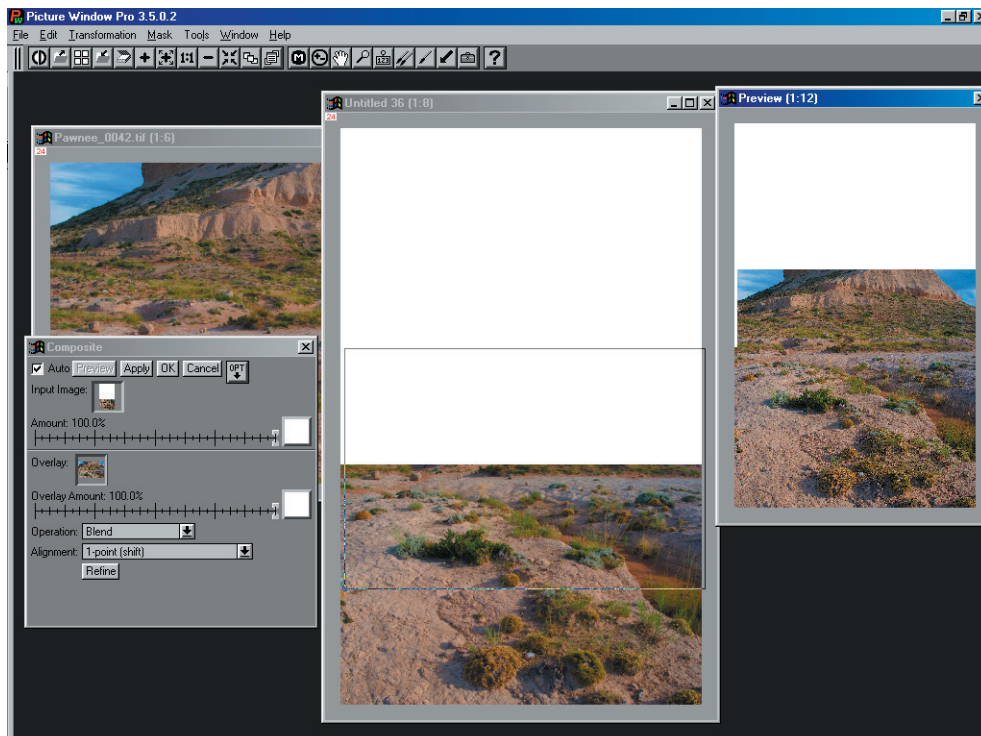
Here is a window ready to start adding the three images. It's easiest if you keep in mind that Picture Window references composite images to the lower left corner, so adding the lowest view first makes sense. Below we have the new tall window selected as the Input Image and the lowest view selected as the Overlay image. Alignment can be set to "none" since this will place the overlay image at the bottom of the Input Image window as shown in the Preview window.



Clicking **OK** creates a new window . . .



The new window has the lowest image in position and ready to add the middle image . . .




In this case, you need to use the newly created window as the Input Image, and the middle view image as the Overlay image. You must be able to shift the overlay image up to a position that fits

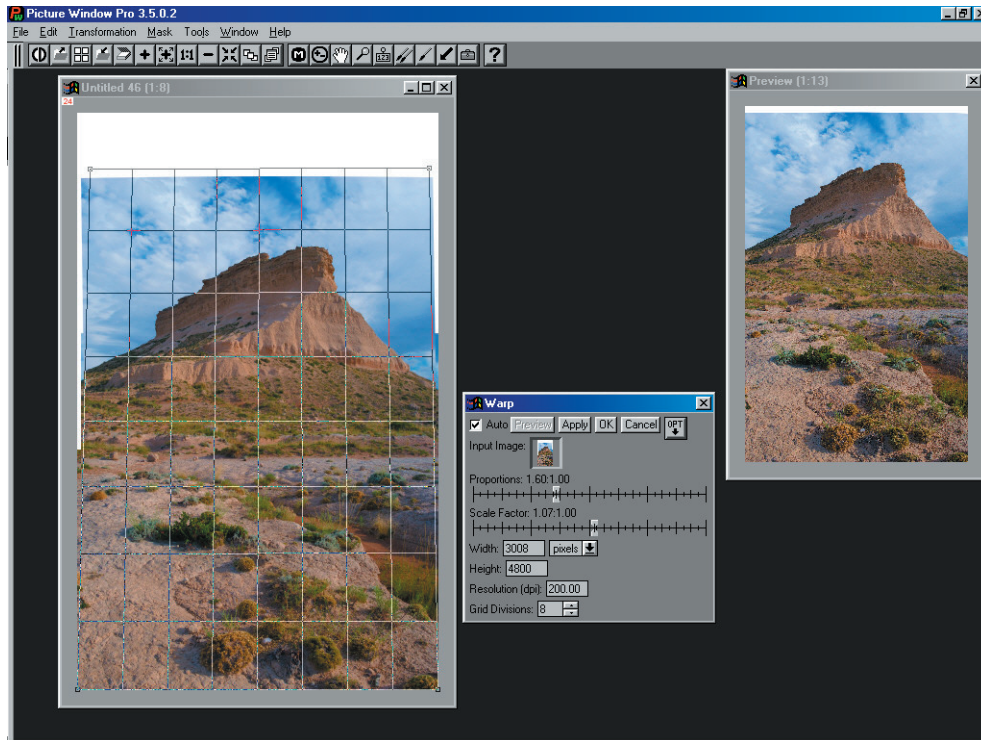
the lower image. In this case I found the 1-point (shift) alignment works fine, but in some images, especially with strong linear visual elements, you may need to use one of the other alignment options such as 4-point (warp) in order to adjust differing perspectives between two joined images. You can see that the two images were not aligned precisely, right-to-left, and the alignment frame is shifted a little to the right as well as raised. You can see the result in the Preview window – in actually making the alignment it is best to enlarge the alignment window to 1:1 for the best precision.

Once the alignment is set, clicking **OK** creates a new window with the composite image. Then using this new image as the Input Image, you can add the next (high view) image.

In this case there was a problem using just the 1-point (shift) alignment. Due to rotation, and possibly some barrel distortion, the lower edge of the upper image could not be aligned across the whole width. The left and right sides could be perfectly aligned, but the middle was off. This indicates we need more than the 4-point (warp) alignment. We need the full “Monty”, the Multi-point alignment. Adding a single mid-point on the lower edge of the alignment frame enabled pulling the center down a little, which corrected the alignment problem.

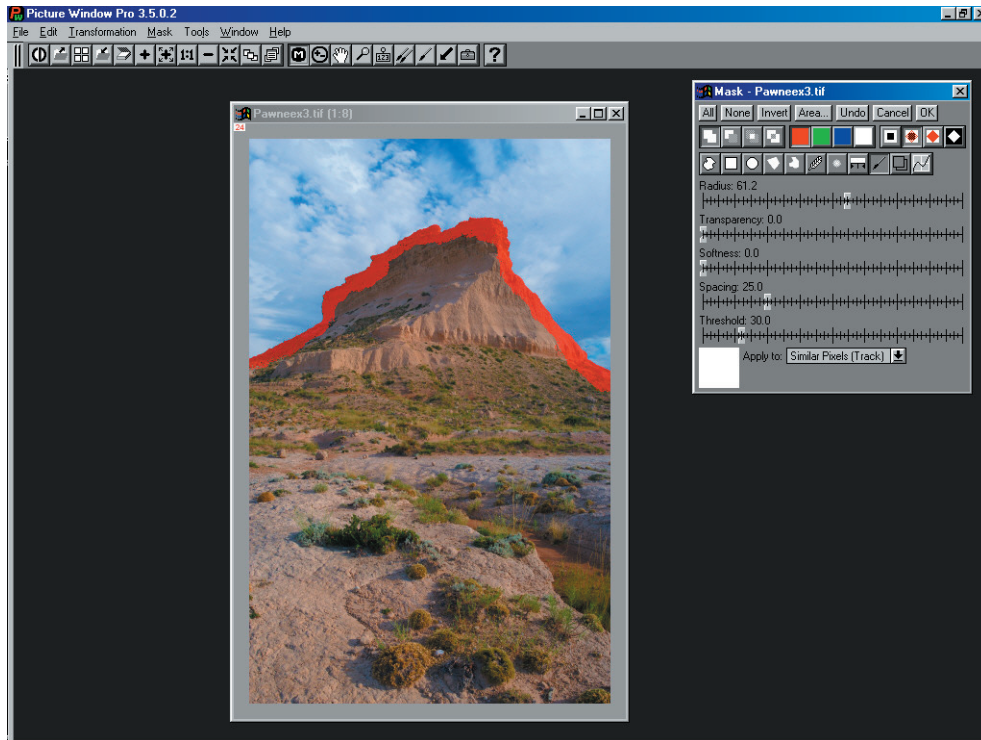


You can see in the Preview window that there are some edges not fully covered by the middle and upper images. This can be taken care of by cropping, cloning, or warping. Since the top image has some perspective convergence and “squashing” due to the camera being angled up, using a warp tool makes sense. In this case, it’s not as obvious as a picture of a building where you can see the walls tilting in, but the convergence is there.

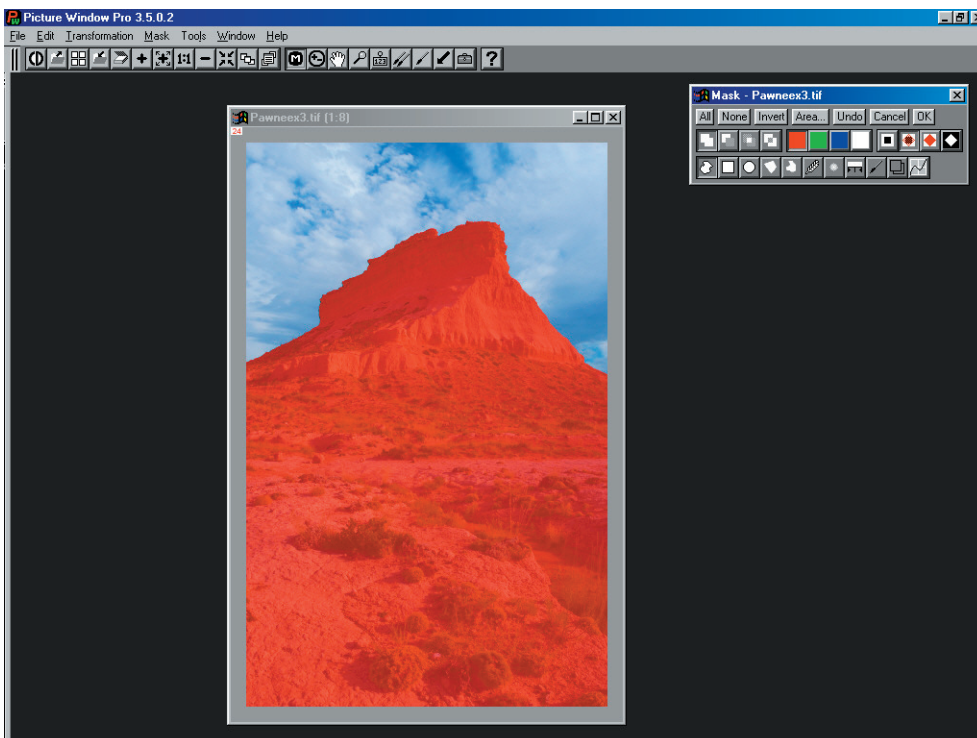


Since the frame is taller than needed, the top of the Warp transformation grid can be used to stretch the image up, or the image can be cropped. I did both. The warp process used some of the extra height of the window, then I cropped the rest.

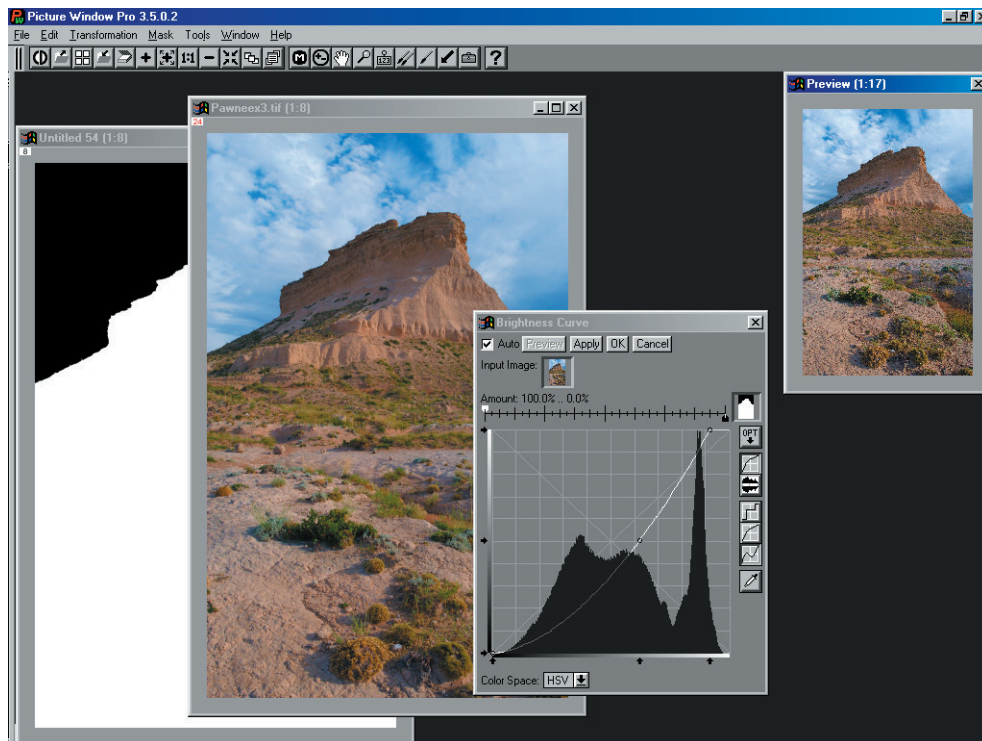
Finally, the tones need a little work. This was shortly after sunrise, the sky was dark blue (altitude at the butte is about one mile above sea level), and the butte and ground were sharply lighted. So one tonal adjustment to the whole image will not work. It must be time for a mask . . .




I used a Paint tool in the mask toolbox, with about a 60 pixel Radius, no Transparency, No Softness and a Threshold of 30 with Similar Pixels (Track). The upper edge of the butte was painted first, as shown above. The image was enlarged to 1:1 so the paint tool could be used right to the edge of the rock. Then, the Freehand Outline tool was used to fill in the rest of the land. The mask was saved so it could be used for a couple of transformations.



The mask was used with a **Brightness Curve** to deepen the sky tones first.



Note the sliders were adjusted so the sky was processed, but not the land (black Amount slider at 100%, white at 0%). The Brightness Curve transformation was bowed down in the middle with the upper right corner pulled a little to the left to lighten the brightest part of the clouds. Then the land image was processed using the same mask but opposite Amount slider settings and a different transformation curve.



The resulting image looks like early morning on the plains, and a view of the Pawnee Buttes that promised the early settlers traveling west by wagon train that Denver was only a day or two away . . .

