
Registering and Distorting Images

Written by Jonathan Sachs
Copyright © 1999 Digital Light & Color

Introduction to Image Registration

The process of getting two different photographs of the same subject to line up with each other is called *image registration*. Even if two pictures are taken from the same location, the corresponding scanned images often will not line up because of small variations in the placement of the film or print in the scanner. Film curvature, scanner nonlinearities, and other distortions can make precise registration even more difficult.

Some applications that depend on accurate image registration are listed below:

- Extending the latitude of slide film by combining over- and under-exposed images of the same subject.
- Combining two or more astrophotographs of the same area of the sky to create a single, enhanced image that captures more detail and faint stars than either of the two images individually.
- Combining three black and white images of the same subject taken through red, green, and blue filters to create a single RGB color image.
- Creating a composite of two faces.
- Overlaying an aerial photograph of a parcel of land with a survey or map of the same property.

The Composite Transformation

While Picture Window's Composite transformation has many uses, it is at heart a very powerful tool for registering images. The Composite transformation accepts an input image (also called the *base image*). When you select an image into the Overlay control, this image (called the *overlay image*) will be registered with and superimposed over the base image. The base image determines the dimensions of the result image and remains fixed; the overlay image is shifted, rotated, scaled, or distorted as necessary to align it with the base image.

What are Alignment Points?

You make the overlay image line up with the base image by setting *alignment points* in both images. The type of geometric transformation applied to the overlay image is determined by the alignment method you select. As soon as you select a

method that uses one or more alignment points, Picture Window displays numbered alignment points in both the base image window and the overlay image window. While these alignment points are initially located in the corners of the images, you can drag them anywhere in the image. In addition, a rectangle is displayed over the base image indicating where the corners of the overlay image will fall in the final composite (see below for examples).

When the Composite transformation creates its result image, it does whatever is necessary to the overlay image to bring each of its alignment points into coincidence with each of the corresponding alignment points in the base image.

Alignment Methods

You select an alignment method with the Alignment control in the Composite dialog box. Each additional pair of alignment points you specify gives you more control over the way the overlay image is transformed. The available alignment methods are:

None

When you select this option, Picture Window simply positions the overlay image over the base image so as to align their lower left corners; the overlay image is neither scaled nor rotated.



1-point (shift)

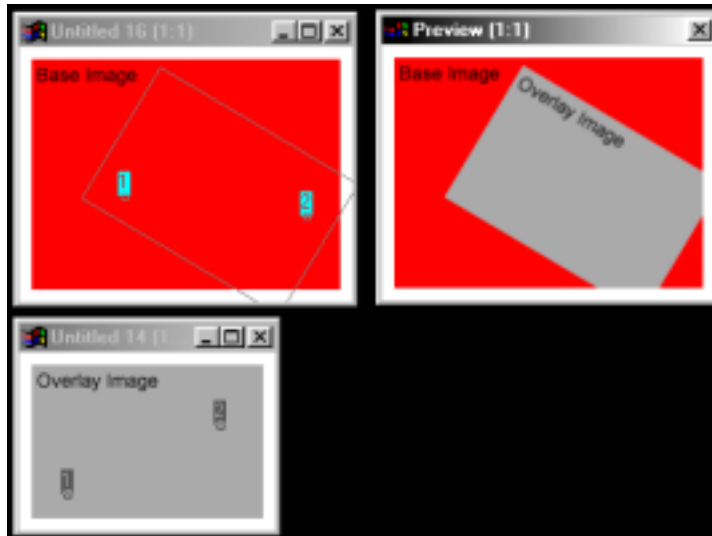
When you select this option, the Composite transformation works with a single alignment point. The overlay image is shifted with respect to the base image so as to bring the two alignment points into coincidence; the overlay image is neither scaled nor rotated.



2-point (shift/rotate)

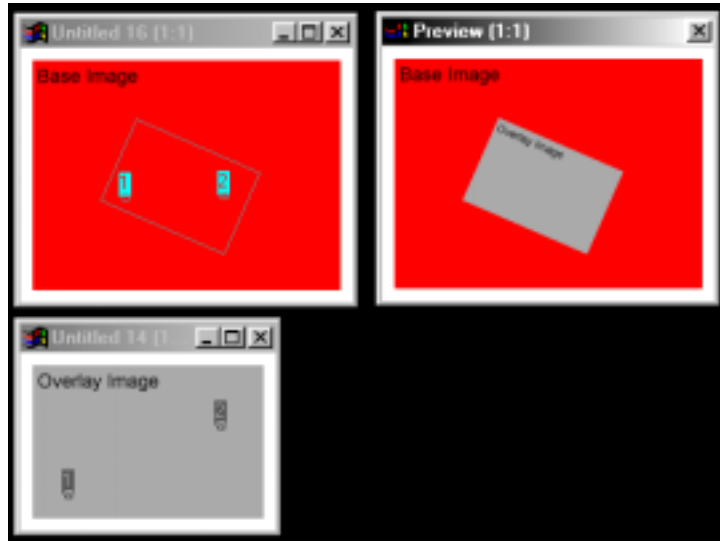
When you select this option, the Composite transformation works with two alignment points. The overlay image is shifted and rotated with respect to the base image so as to bring alignment point 1 in both images into coincidence; the overlay image is then rotated (but not scaled) so the line between alignment points 1 and 2 in the overlay image points in the same direction as the line between alignment points 1 and 2 in the base image.

Use this option when you know the two images have the same scale and can thus be brought into coincidence by a simple shift and rotate. For example, you should be able to register two images of the same subject taken from the same location with this method.



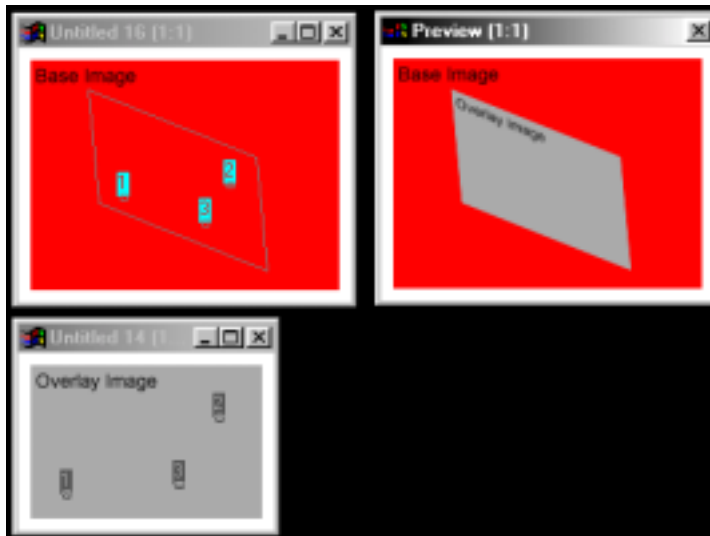
2-point (shift/rotate/scale)

When you select this option, the Composite transformation works with two alignment points. The overlay image is shifted, rotated, and scaled with respect to the base image so as to bring both alignment points both images into coincidence.



3-point (shift/rotate/scale/shear)

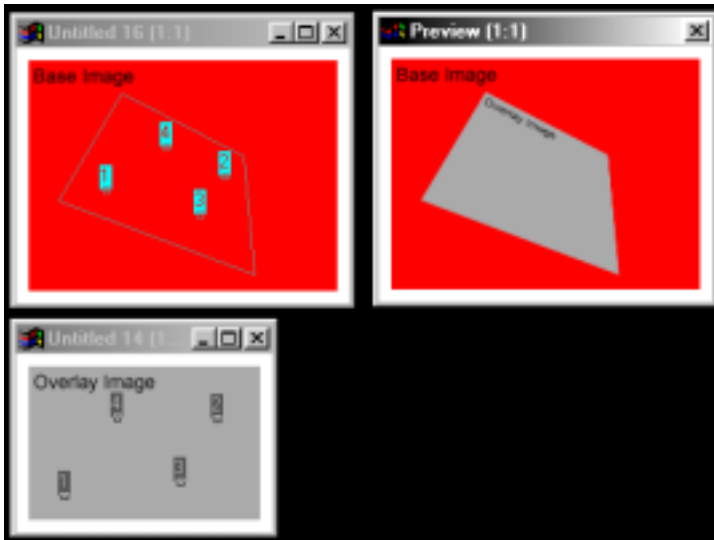
When you select this option, the Composite transformation works with three alignment points. The overlay image is shifted, rotated, scaled, and sheared with respect to the base image so as to bring all three alignment points both images into coincidence. This kind of geometric operation is called an *affine transformation*.



4-point (perspective warp)

When you select this option, the Composite transformation works with four alignment points. A perspective transformation is applied to the overlay image so as to bring all four alignment points both images into coincidence

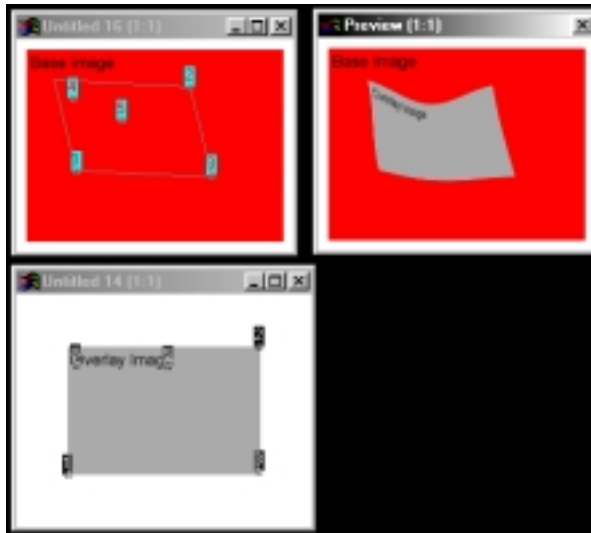
This alignment method can be used, for example, to drop an image into a photo of a room that has a picture frame on one of the walls. The overlay image can be warped to make it fit inside the picture frame in the same plane as the wall.



Multi-point

Note: multi-point alignment is available only in the Pro version of Picture Window, not in the standard version.

When you select this option, the Composite transformation works with anywhere from a minimum of 3 up to a maximum of 62 alignment points. Using one of several mathematical techniques, the overlay image is distorted so as to bring all pairs of alignment points into coincidence.



This is by far the most powerful of the alignment methods, and can be used to correct for a variety of small distortions that can make the simpler methods described above fail to accurately register the overlay image with the base image. Or it can be used to add complex, free-form distortions to images.

Using the Operation Control

The Operation control in the Composite dialog box determines how the two images are combined once they have been registered. The most common settings are:

Blend

Selecting this option causes the overlay image to be blended with the base image based on the Overlay Amount slider. A setting of 0% makes the overlay image totally transparent, leaving the underlying base image completely visible. A setting of 50% blends the two images in equal proportions, similar to

making a double exposure. A setting of 100% makes the overlay image opaque, completely obscuring those parts of the base image that are covered by the overlay image. Areas of the base image that are not covered by the overlay image will still be visible around the edges.

Register

This option is identical to Blend, with the exception that areas of the base image that are not covered by the overlay image show up black in the final image. For image registration, use this setting and leave the Overlay Amount slider set to 100%.

Absolute Difference

Selecting this option causes Picture Window to compute the absolute value of the difference between the base and the overlay image. The absolute difference image is black (zero) where the two images are identical; where the images differ, the absolute difference increases and the result image is brighter. The absolute difference image is then blended with the base image according to the Overlay Amount slider as described above.

This setting is mostly useful during the preview process to check how well the current alignment method and alignment point locations are registering the two images. The better the match, the darker the preview image.

Step-by-Step Registration Procedure

This section describes the step-by-step procedures to register images using the Composite transformation and multi-point alignment.

Step 1 -- Prepare the images.

The first step is to collect the images you will be registering. You must then decide which one you will use as the base image and which one will be the overlay image. Remember, the base image remains fixed and the overlay image is distorted to match the base image.

The size of the final image produced by the Composite transformation is inherited from the base image. This means that any part of the overlay image that lies outside the input image will be cut off. To avoid this, use the Geometry/Add Border transformation to create enough blank space around the base image to accommodate the overlay image.

Step 2 -- Getting Started with Composite

First click on the base image window to make it the current window and then select Transformation/Composite from the main menu. This pops up the Composite dialog box with the base image selected as the Input image. Next, click on the small white square labeled Overlay. This should pop up a menu of image windows -- use this to select the overlay image. Finally, use the Alignment control to select an alignment method. For simple compositing where no scaling or distortion is required, you can use *2-point (shift/rotate)*; for more complicated situations, use *multi-point*.

Step 3 -- Position the Alignment Points

As soon as you set the alignment method, labeled alignment points will appear in the base and overlay image windows. Initially, these points are located in the corners of the images. If you selected 2-point alignment, the alignment points should appear in the lower left and lower right corners; if you selected multi-point alignment, a third point should appear in the upper left corners of each image.

You can move an individual alignment point by dragging it to a new location; you can move all the alignment points in a window at once by holding down the Shift key while dragging anywhere in the window.

Alignment points are generally used to register images by placing them on the same feature in each image. Thus when the alignment points line up, so do the underlying images. Small but distinct features that can be located very precisely make the best alignment points. Try to avoid setting alignment points in blurry or featureless areas as these are much harder to match. For the greatest accuracy, alignment points should also be set as far apart as possible; if the corners of the images match up, it is likely that the centers of the images will as well.

By zooming in on the base and overlay images beyond 1:1, you can even position alignment point locations to within a fraction of a pixel. But you don't need to position them by hand that accurately; you can click the Refine button at the bottom of

the Composite dialog box and let Picture Window jiggle each alignment point to find the best local match between the two images. For the Refine button to work, the initial alignment point locations only need to be within about 5 pixels of their optimal locations, but it does require a distinct feature to lock onto.

Step 4 -- Multi-point Alignment

If you selected multi-point alignment, the first three alignment points (labeled 1, 2, and 3) are treated differently from the rest. These first three points define a baseline affine transformation (see 3-point alignment above) which is applied to the overlay image to align it with the base image. Additional alignment points, if any, are used to create local distortions in the overlay image so you can get specific image features to line up if the baseline transformation is not accurate enough.

To create an additional pair of alignment points, use the miniature tool bar at the bottom of the Composite dialog box and click the + icon. Then click on the base image at the approximate location where you want the alignment point to be located. A corresponding alignment point will be automatically created in the overlay image using the baseline affine transformation defined by the first three alignment points. Once the new alignment points have been created, you can drag them anywhere in the image. You can repeat this procedure to create as many alignment points as necessary (Note: you must click the + button once for each new point)

When placing alignment points or reviewing their locations, it is often useful to place the base and overlay images side by side and zoom way in on each one so you can see exactly where the alignment points are placed. This means you can't see the entire image at once, and it is easy to lose track of where individual alignment points are placed. As an aid in this situation, you can use a handy keyboard shortcut: pressing 1 on the keyboard automatically recenters both the base and overlay image windows on alignment point 1. Similarly for points 2..9, a..z, and A..Z.

The strategy for placing alignment points is really two-fold; they can be used to equal advantage either to specify that some part of the overlay image should be moved to a new location or that some part of the overlay image should be left where it is. This second strategy can be useful if you find that parts of the image you did not want to change are being affected.

As if all this flexibility weren't enough, you have more options to play with. First, you can select the mathematical method used to interpolate between the alignment points. When you set alignment points, you are specifying how specific points in the overlay image will be moved to their corresponding locations in the final image.

Interpolation techniques are used to determine how points between the alignment points will be moved, and there are an infinite number of different ways to transform the overlay image that meet the constraints dictated by the alignment points.

To change interpolation methods, click the Opt button in the upper right corner of the Composite dialog box and choose Shepard, Hardy, Thin Plate, or Gaussian. The smoothest interpolations are usually produced by the Hardy or Thin Plate methods. The Shepard or Gaussian methods are good for making adjustments that are more localized. The Hardy and Gaussian methods can be further customized by adjusting the R slider that appears at the bottom of the Composite dialog box. Smaller values of R make the distortions more localized; larger values make the distortions smoother by increasing the radius of influence of each alignment point. As you make the distortions radical and more local, you increase the probability of *foldover*, a phenomenon where the same image data can appear in two or more locations in the final image (a little like a funhouse mirror).

Step 5 -- Check your Progress

To see how you are doing, set the Operation control in the Composite dialog box to Absolute Difference and click the Preview button. This will show you the difference between the base image and the distorted overlay image. Where the two images are identical, the preview image will be black; the more they differ, the brighter the preview image will be. This technique works best when you have two more or less identical photographs of the same subject; if the images are basically different (for example over- and under-exposed photos of the same subject), the absolute difference will never be totally black, no matter how well you register the images.

Step 6 -- Create the Final Image

Once you are satisfied with the number of alignment points and their placement, set the Operation control to Register and click OK. This should create the final image, namely a version of the overlay image that now lines up with the base image.

Step 7 -- Sharpen the Image

The registration process involves resampling the overlay image. Picture Window uses a technique called bilinear interpolation to do this and a side effect of this interpolation technique is to blur the image very slightly. You can compensate for this slight blurring if you wish by using the Sharpen transformation.

Using Multi-point Alignment to Distort Images

Ironically, the same mathematical techniques used to precisely register two photographs can also be used to create subtle or wild free-form image distortions. The trick is to select the same image as both the base image and the overlay image. When the base and overlay images are the same, both sets of alignment points are displayed on the same image. To avoid total confusion, Picture Window displays alignment points a little differently in this situation. If both alignment points are located in the same place, the display is normal. When you move one of the alignment points, it is shown as a line with one end labeled and the other end marked with a small circle. When the composite operation is performed, the labeled alignment point is moved to the other end of the line.

Step-by-Step Distortion Procedure

Step 1-- Prepare the Image

The size of the final image produced by the Composite transformation will be the same as the image you start with. This means that any part of the distorted image that ends up outside the original image will be cut off. If you want to avoid this, use the Geometry/Add Border transformation to create additional blank space around the image.

Step 2 -- Getting Started with Composite

First click on the image window to make it the current window and then select Transformation/Composite from the main menu. This pops up the Composite dialog box with the image selected as the Input image. Next, click on the small white square labeled Overlay. This should pop up a menu of image windows -- use this to select the same image as the Overlay image. Finally, set the Operation control to Register and the Alignment control to Multi-point.

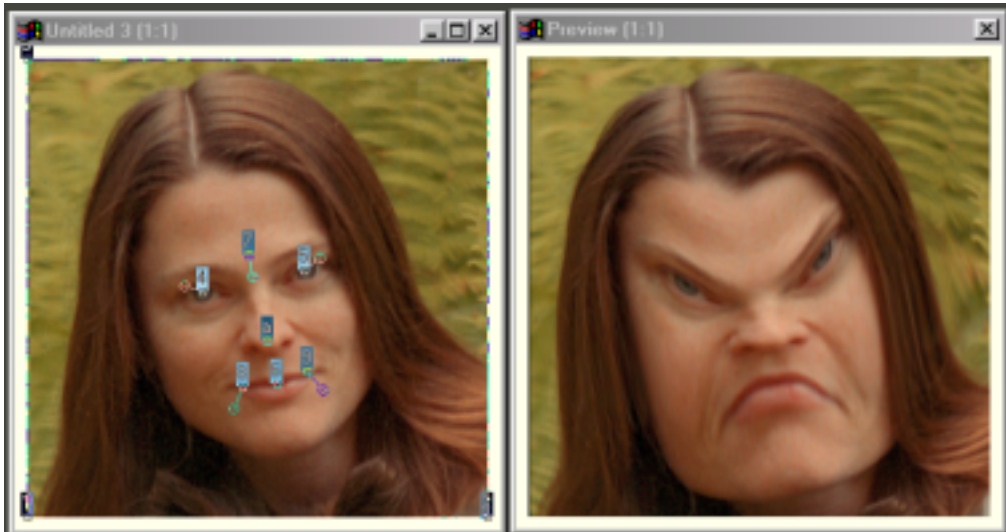
Step 3 -- Position the Alignment Points

As soon as you set the alignment method, three labeled alignment points will appear in the base and overlay image windows. Initially, these points are located in the lower left, lower right, and upper left corners of the image.

You can move an individual alignment point simply by dragging it to a new location; you can move all the alignment points in a window at once by holding down the Shift key while dragging anywhere in the window.

The first three alignment points are special -- they define a baseline 3-point affine transformation (see above). Additional alignment points create local distortions of this baseline transformation. For most image distortions, leaving the first three alignment points at their initial locations is a good starting point as this indicates that you want to baseline transformation to have no effect.

To start distorting the image, click on the + icon on the small tool bar at the bottom of the Composite dialog box and then click on the image window to add a new alignment point. When you drag the new alignment point, it separates into two points connected by a line. The labeled end defines where the image data will be moved from; the end with the small circle defines where the image data will be moved to. When you click Preview, Picture Window shows you the effects of the distortion you have specified. The more alignment points you add, the more control you can exercise over the image geometry..



As with image registration, there is a dual strategy for placing alignment points. It is equally useful to use an alignment point to specify displacements for certain key features as it is to add an alignment points with no displacement to nail down other features and prevent them from being affected.

For image distortion, the Gaussian method (selected via the Opt button) can be a good choice as its effects tend to be local and can be further controlled with the R slider.