
Creating and Using Masks

Written by Jonathan Sachs
Copyright © 1996-2000 Digital Light & Color

Introduction

A mask is an image used to identify part of another image. Masks are one of the most powerful features of Picture Window -- they can be used to serve several important functions:

- To localize the effects of a transformation to part of an image. For example, to lighten or darken just one part of an image using the Brightness transformation, you must first create a mask that isolates the part you want to modify.
- To protect part of an image from being modified by a tool like clone or paint. You can use masks the way artists sometimes use masking tape or frisket, to protect part of an image from being affected by one of the freehand painting tools.
- To identify those parts of one image to be composited with another. If you want to place a copy of an object from one image into another, you must first create a mask that defines the boundary of the object you want to composite.

A mask is just an ordinary black and white image -- what makes a mask special is the way it is used. There is however one very important rule about using masks that you must keep in mind at all times: *A mask must be a black and white image and it must have the exactly the same width and height in pixels as the image with which it is used.*

Each pixel in a mask image can be anywhere from totally opaque (white) to totally transparent (black). By default, a mask will apply at full strength where it is opaque and will have no effect where it is transparent, although you will see later how to alter this using Picture Window's double slider. You can use a semitransparent mask to partially mask a region. For example, when using a mask with the Brightness transformation to lighten or darken part of an image, it usually works best to make the edges of the region you want to lighten or darken semitransparent so that the lightening or darkening effect will not have too obvious a boundary. This process of softening the edges of a mask is called *feathering*.

Creating and Modifying Masks

Picture Window has many powerful tools for creating different kinds of masks. While any black and white image of the correct size can be used as a mask, the most common ways to create a mask are with the Mask tool or the Gradient transformation. The Mask tool is the primary method for creating new masks, so we'll start there.

Once you have created and saved your mask, you can use any of Picture Window's freehand tools or transformations to modify it since it is simply a black and white image. Just remember that a mask always has to be exactly the same size as the image it is used with, so if you crop or resize the image, the mask will cease to be usable unless you apply an identical operation to the mask.

How the Mask Tool works

The Mask Tool displays masks as a colored overlay on an image. By default, masks are displayed as a semitransparent red overlay, but you can use the controls in the Mask dialog box to select alternate overlay colors and/or transparencies. The areas where the red overlay is displayed correspond to regions of the mask image that are white (opaque) -- the area where the mask is black (transparent) are clear.

Along the bottom of the Mask dialog box is a tool bar that lets you select individual masking tools. Each masking tool provides a different way to create or modify a new mask region. For example, you can use the freehand outline tool and simply draw a curve around the part of the image you want to mask.


The Mask dialog box has another tool bar which lets you select a combine mode. By default, the mode is set to Add, so each new region you mask is added to the

current mask. For example if you want to create a mask that contains several regions, you can use the freehand outline tool to circle them one at a time. Each new region you draw will be added to the mask you have created so far. But what if you make a mistake and accidentally mask part of the image you didn't want to include? Then you can switch to Subtract mode and each new region you mask will be removed from the current mask, so you can just outline the part you want to remove and it will be unmasked.

For a detailed description of all the masking tools, please see the Picture Window manual and help file. There is additional information in the General Masking Tips and Techniques section at the end of this document. Once you have mastered the basic masking techniques, you can go back and experiment with some of the more advanced masking tools.

To create a new mask for an image

First select the image you want to mask by clicking on it. Then execute the Mask

New command from the main menu or click on the  icon on the tool bar. This creates a new and initially transparent mask overlay for the current image and pops up the Mask dialog box. This dialog box provides a number of tools and options for modifying the mask overlay.

To convert the current mask overlay to a black and white image

When you click the OK button on the Mask dialog box, the current mask overlay is “peeled off” the base image and converted to a separate black and white image in its own window. The mask image is white where the mask is in full effect, black where the overlay is transparent, and takes on intermediate gray levels where the overlay is semitransparent. The resulting black and white image is no different from any other black and white image. You can save a mask image for later use or apply any of Picture Window's transformations or tools to modify it further.

To convert a black and white image to a mask overlay

Use this procedure when you want to rework a saved mask. Remember that a mask must be a black and white image and it must have exactly the same width and height in pixels as the image to which it is applied. First make sure that both the image and the mask are open in separate windows. Select the image you want to mask by clicking on it and then execute the Mask command from the main menu. Select the name of the mask image you want to use from the list of available masks

at the bottom of the Mask submenu. This makes the mask image into an overlay for the original image and pops up the Mask dialog box. If the name of the mask window does not appear at the bottom of the mask menu, double check (using the Window/Info command) to make sure that the mask image is a black and white (and not a color) image and that it has exactly the same size, in pixels, as the image you want to mask.

Creating a mask with the Gradient transformation

The gradient transformation can be used to create a number of useful masks, colored filters, and backgrounds. When creating a mask with the gradient transformation, remember that all of the colors on the color line must be neutral shades of gray -- otherwise the resulting gradient image will not be black and white. To guarantee that the color line contains only grays, click on the white patch in the lower right corner of the color picker for each control point on the color line and then use the value slider to adjust the gray levels.

For example, to lighten or darken the upper part of an image, you can create a linear gradient that is light at the top of the image and dark at the bottom. This gradient image can then be used as a mask in conjunction with the Brightness transformation to control which parts of the image to lighten or darken. Similarly gradients can be used with the Composite transformation to blend one image into another in a controlled manner.

Other transformations you can use to create or modify masks

To create masks from text, use the Text transformation and place white text on a black background (or black text on a white background). A text mask can then be used to spell out a message on an image by altering its color or brightness.

To roughen a mask you can use the Add Noise transformation or create a texture using the Texture transformation and then apply the texture to the mask. Other special effects like Emboss or High Contrast can be used to create interesting masking effects.

Subtle control over semitransparent masks is often best obtained by using the Brightness Curve transformation.

Masks can be softened using the Gaussian Blur option of the Blur transformation or sharpened using the Sharpen transformation.

Using Masks with Transformations

Many of Picture Window's transformations can be used with masks -- the mask restricts the effect of the transformation to a particular region of the input image. For example, you can create a mask that consists of a black circle on a white background and use this mask to brighten or darken just that circular region of an image. Masks can also contain shades of gray, so you can control exactly how much of a particular transformation to apply to each part of an image. You can select a mask either before initiating a transformation or after.

To select a mask before running a transformation


If you initiate a transformation while the Mask tool is active and a mask overlay is displayed over the input image, then that mask will be automatically selected into the amount control (if any) for that transformation. In addition, the mask tool will remain active and you can alter the mask while the transformation is active.

Remember to click on the image window before going to the Transformation menu so Picture Window will know which image you want to operate on. If you try to start up a transformation while the Mask dialog box is the current window, most of the transformation menu items will be grayed out.

If the colored mask overlay interferes with your use of the transformation by making it hard to see the differences between the input image and the preview image, you can use the mask display tool bar in the Mask dialog box to temporarily make the mask overlay transparent.

To select a mask after running a transformation

First make sure the mask is open in a window. Then select the transformation from the Transformation menu. If the transformation accepts an amount mask, you will see a slider with a large button to its right:




Click on the button to pop up a menu of available masks and select the one you want to use by name. If the name of the mask window does not appear in the list, double check (using the Window Info command) to make sure that the mask image

is a black and white image and that it has exactly the same size, in pixels, as the input image.

Using the double slider

Once you have selected a mask into an amount control, its slider changes from a single to a double slider; one black and the other white. A thumbnail image of the mask also appears on the button face.

The white slider controls what percent of the transformation is applied to the input image in regions where the mask image is white (opaque). The black slider controls how heavily the transformation is applied where the mask is black (transparent). Gray parts of the mask produce an intermediate effect.




Initially the white slider is set to 100% and the black slider to 0%. This means that the parts of the input image where the mask is white receive the full force or the transformation while black areas remain untouched. At the other extreme, when you reverse the positions of the black and white sliders, you invert the effect of the mask, applying the full transformation to the parts of the input image where the mask is black and leaving the white areas untouched.


A Simple Example -- Creating and Using a Mask

Let's put all the pieces together and see how masking works in practice. First, we'll start by opening a file containing an image we want to work on and clicking on it to make it the current window.



Next, click on the  button on the tool bar to bring up the Mask dialog box. Using the default settings, position the cursor over the input image and, holding down the left button, crudely outline the group of two flowers on the left. When you release the mouse button, Picture Window should display a red overlay over the area you just outlined. Next, outline the flower on the right. Since the combine mode is set to add, this should add a second region to the current mask:



Now, to soften the edges of the mask, select the  (feather) button from the tool bar at the bottom of the Mask dialog box and set the slider to roughly 25 pixels. Then click the Apply button:



Next, leaving the mask overlay on the image, click once on the input image and then select Transformation/Gray/Brightness from the main menu and go to the Amount control in the Brightness dialog box (just below the Input Image control). The Amount control should have a double slider (white on top and black on the bottom) with both sliders initially set to zero (in the middle). There should also be a miniature representation of the mask image in the small square to the right of the slider.

Move the black slider to the left until the readout shows roughly -25% and move the white slider to the right until the readout registers about 10% . Then click Preview to see the effect.




By setting the white slider to 10% , you are specifying that you want the image to be made 10% brighter in those areas where the mask is white. By setting the black slider to -25% , you are specifying that you want the image 25% darker in those areas where the mask is black. Where the mask has an intermediate gray value, the image is lightened or darkened by an intermediate amount. By experimenting with the sliders you can create many different effects, some subtle and some extreme.

Until you click OK or select a new input image, both the Mask tool and the Brightness transformation are both active. This means you can also modify the mask using the Mask tool and then click the Preview button in the Brightness dialog box to see the effect on the transformation -- in effect the mask becomes another setting in the transformation dialog box that you can alter. However, only one Mask tool can be active at a time.

Because the same masking techniques work with most of Picture Window's transformations, you can use the either the same mask or a different one to perform many other operations. For example, the following two examples show the result of

selective color balancing and selective embossing using the same input image and mask as above.



Using Masks with Freehand Retouching Tools

Picture Window's clone, paint, and miscellaneous tools can be used with masks -- the mask prevents part of the input image from being affected by the tool. For example, if you create a mask that consists of a white circle on a black background, you can use this mask to protect that circular region of an image while you paint around it. Since masks can also contain shades of gray, you can control exactly how much of a particular tool is applied to each part of an image.


Procedure:


- Create a mask overlay on the input image before selecting the tool.
- Select and apply the desired retouching tool.

Masked regions are fully protected from the action of the tool; unmasked regions are totally unprotected. Partially masked regions are partially susceptible to the tool.



Image with circular mask overlay

You may find it helpful to temporarily suppress the mask overlay while using the tools so you can see the input image better. To do this, click on the  button in the

Mask dialog box. To restore the semitransparent colored overlay, click on the  button.



Mask protects image from paint tool



Red overlay temporarily suppressed

Using Masks with the Composite Transformation

Picture Window's Composite transformation lets you perform a number of different functions, ranging from simple color overlays to creating complex collages. This section will explore the full range of capabilities of this powerful transformation and show you how to achieve a number of different effects.

The basic idea behind the Composite transformation is to start with a base image (the input image) and to overlay over that base image either a solid color or an overlay image. By using masks, you can control what part of the base image is covered by the overlay and what part of the overlay image is placed over the base image. The Composite transformation also allows you to control the size and orientation of the overlay image by positioning alignment points. This feature lets you line up features in the overlay image with specific locations in the base image to provide precise placement of the overlay.

The following sections will illustrate progressively more complex operations.

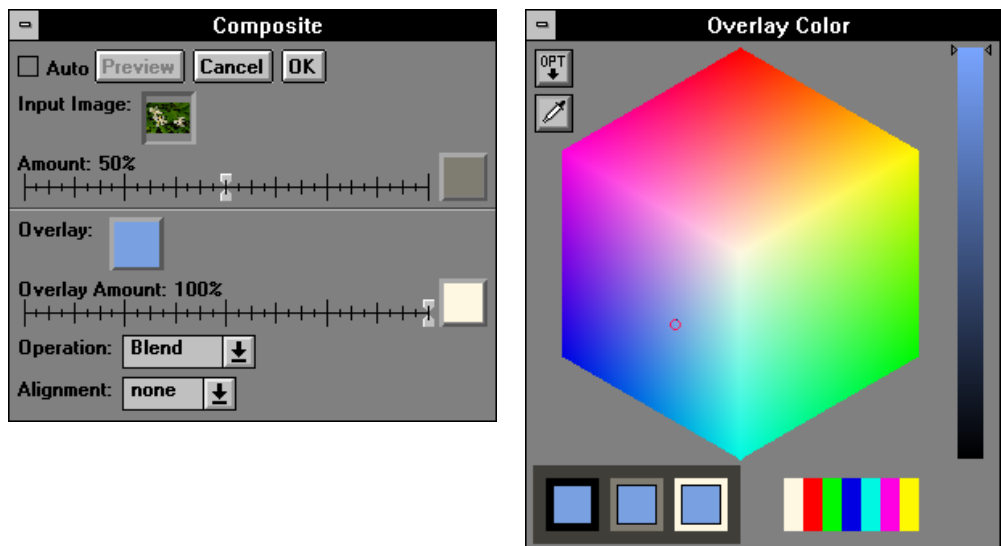
Applying a Solid Color Overlay

The simplest form of the Composite transformation lets you apply a solid color overlay to all or part of the base image.

You control the transparency of the overlay with the Amount slider (the one just below the Input Image control). Setting the Amount slider to 0% produces a totally transparent overlay which lets the original image show through in its entirety. A 50% transparency blends the input image and the overlay color in equal amounts. Setting the transparency to 100% makes the overlay color totally opaque and hides the input image completely.


You control the color of the overlay with the Overlay control (the one just below the Amount control). Just click on the square button just to the right of the caption “Overlay:”. This will display a menu -- select Solid Color (the first item) from the menu and a color picker will be displayed. Then use the color picker to select the desired overlay color.

When overlaying with a solid color, the Overlay Amount and Alignment controls are not used.




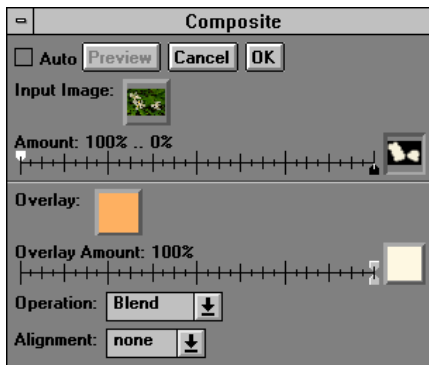

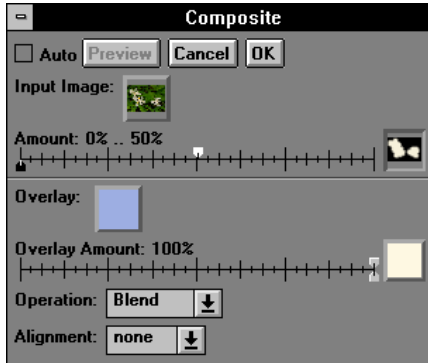
Once you have selected a color and transparency, click the Preview button to get an idea what kind of effect you have selected. Using Composite in this way is a little like taking a double exposure of the image and a full frame of solid color. If neces-

sary, adjust the overlay color and transparency and preview again until you achieve the desired results -- then click OK to finish up.


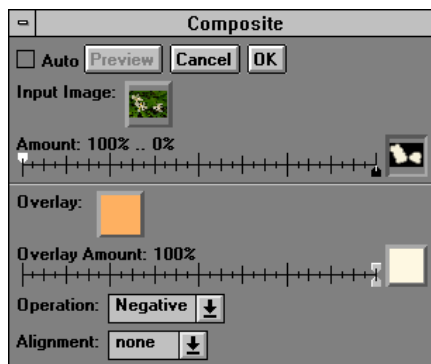


As illustrated earlier, you can work with a mask over the input image by either initiating the Composite transformation with a mask overlay in place or by selecting the name of the mask window from the menu that appears when you click the white square to the right of the amount slider. If you do not use a mask, the overlay color will be layered over the entire input image with a transparency controlled by the Amount slider. If you do use a mask, the transparency of the overlay color will vary according to the opacity of the mask and the settings of the black and white amount

sliders. Using the same mask we created before, we can get a variety of effects by playing with the overlay color and the amount sliders:



You can experiment with additional variations by changing the Operation setting. The default value, Blend, produces straight transparency effects as shown above. The Lighten (and Darken) settings are similar, except they limit the overlay to only lightening (or darkening) the base image. The Filter setting is similar to the blend setting but rather than averaging the overlay image with the base image, the overlay image is applied to the base image as a filter. The difference is like the difference between a double exposure (50% blend) and sandwiching two transparencies (100% filter). The Negative setting combines the input and base images using an XOR operation which creates some interesting special effects:



Overlaying One Image Over Another

The next step beyond overlaying the base image with a solid color is to overlay the base image with an overlay image. The simplest form of this is to combine two images of the same size to produce an effect like a double exposure. To achieve this effect, open the overlay image in a window and click on the square button of the Overlay control. Next, select the name of the window containing the overlay image from the menu that pops up. This will display a miniature of the overlay image in the square to indicate which image you have selected.

If the two images are not the same size, it does matter which one you choose as the base and which one is the overlay. The result image is always the same size as the base image. If the overlay image is larger than the base image, any excess that extends beyond the edges of the base image will be discarded. Similarly, if the over-

lay image is smaller than the base image, the parts of the base image that are not covered by the overlay image will show through in the result image. If you want the result to include all of both images, you will have to use the Crop transformation to add a border to the base image before compositing.

You control the transparency of the overlay image with either the Amount slider (the one just below the Input Image control) or the Overlay Amount slider (more on the difference between these two later). Setting the Amount slider to 0% produces a totally transparent overlay image which lets the base image show through in its entirety. A 50% transparency blends the input image and the overlay image in equal amounts. Setting the transparency to 100% makes the overlay image totally opaque and hides the input image completely.

Overlay of one entire image with another

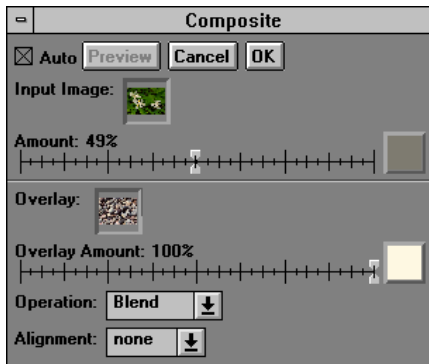
If you do not use a mask for the input image or the overlay image, the overlay image will be applied to the entire input image with a transparency specified by the Amount slider or the Overlay Amount slider:



Input Image



Overlay Image



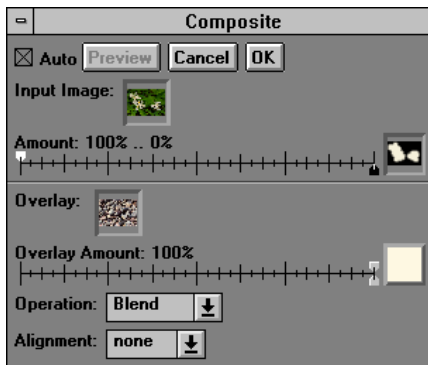
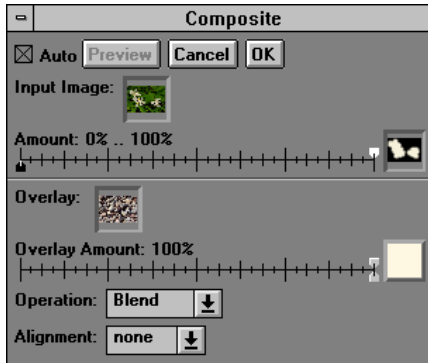
Composite Dialog Box



Preview Image

Selective image overlay using a mask

As before, you can work with a mask over the input image by either initiating the Composite transformation with a mask overlay in place or by selecting the name of the mask window from the menu that appears when you click the white square to the right of the amount slider. If you do not use a mask, the overlay image will be layered over the entire input image with a transparency controlled by the Amount slider. If you do use a mask, the transparency of the overlay image will vary according to the opacity of the mask and the settings of the black and white amount sliders. Using the same images and mask we created before, we can get a variety of effects by playing with the amount sliders:



The above examples illustrate using the Amount double slider to invert the mask. In the first example, the overlay image is applied to the base image where the mask is opaque (white). In the second example, by reversing the positions of the

sliders, we have effectively inverted the mask and caused the overlay to be applied to the parts of the base image where the mask is black. Intermediate settings of the two sliders can be used to create various partial blending effects.

Different operations for different effects

The Operation setting in the Composite dialog box lets you vary the way in which the overlay image is applied to the base image.

Blend


The default setting is Blend which computes a weighted average of the two images. The higher the overlay amount slider setting, the more of the overlay image is used and the less of the base image. At a 50% setting equal amounts of the two images are combined to give an effect like a double exposure. At low settings, the output image is primarily made up of the base image with only a small amount of the overlay image mixed in.




Base Image




25% Blend



50% Blend



75% Blend




Overlay Image

Lighten


If you select Lighten, a blend is performed but its result is applied to the base image only where it is lighter than the base image. This preserves areas where the base image is darker.




Base Image



75% Lighten



100% Lighten




Overlay Image

Darken


When you select Darken, the darker parts of the base image are replaced and the lighter parts preserved:




Base Image



75% Darken



100% Darken



Overlay Image

Filter


If you select Filter, the overlay image is applied to the base image as a filter. At 100%, the effect is like sandwiching two transparencies and viewing one through the other. Reducing the percentage is like using a paler filter:



Base Image



50% Filter




100% Filter




Overlay Image

Subtractive Filter


If you select Subtractive Filter, the overlay image is applied to the base image as a subtractive filter. This is computed as the negative of filtering the negative of the base image with the negative of the overlay image:




Base Image



50% Sub. Filter



100% Sub. Filter



Overlay Image

Negative


If you select Negative, the two images are combined using a mathematical technique called “exclusive or” sometimes abbreviated XOR which can produce interesting visual effects but is difficult to predict:



Base Image



100% Negative



Overlay Image

Positioning the overlay image using alignment points

The previous examples have assumed that we already have two images which have been scaled and oriented compatibly. Often, however, we need to reposition, resize, rotate, skew, or alter the perspective of the overlay image to make it conform to the base image. The key to doing this is to use alignment points.

If you use the default Alignment setting of None, then Picture Window simply aligns the base and overlay images by lining up their lower left corners. This setting is primarily useful if the two images are already properly aligned and have the same size in pixels.

Picture Window also supports 1-point, 2-point, 3-point, or 4-point alignment. These options let you specify one, two, three, or four pairs of alignment points to make the overlay image line up with the base image. Each pair of points includes one point in the base image and one point in the overlay image. The overlay image is adjusted as necessary to bring each pair of alignment points into alignment in the result image. The more alignment points you specify, the more ways Picture Window must be able to distort the overlay image to guarantee that all the points line up.

When you specify 1-point alignment, Picture Window merely slides the overlay image over the base image until the two alignment points (one in the base image and one in the overlay image) are coincident.

When you specify 2-point alignment, Picture Window first slides the overlay image to align the first pair of points and then scales and rotates the overlay image as necessary to align the second pair of points as well. There is also a variant of 2-point alignment that lets you shift and rotate the overlay image only without affecting its size.

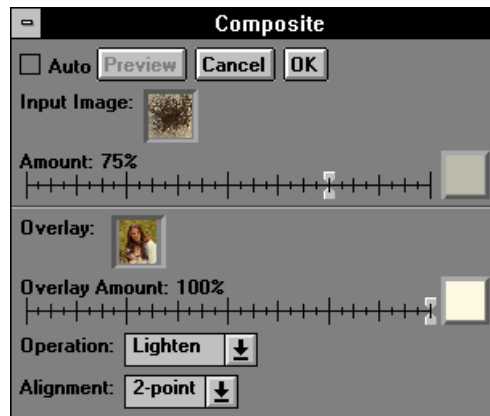
When you specify 3-point alignment, Picture Window skews the overlay image to get the third pair of points into alignment using what mathematicians call an *affine transformation*.

When you specify 4-point alignment, Picture Window applies a perspective warp to the overlay image to make all four pairs of points match up.


Picture Window Pro also supports an advanced method called *multi-point alignment* which is described in detail elsewhere.

As soon as you modify the Alignment control to specify one or more alignment points, numbered control points are displayed as an overlay over both the input and the overlay image. Initially, these alignment points are placed in the corners of the base and overlay images, but you can drag them to any location you desire. Each alignment point is numbered so you can tell which points in the base image correspond to which points in the overlay image.

Let's look at an example to see how this works:



Composite Dialog Box -- Note Amount set to 75%, Operation set to Lighten, Alignment set to 2-point




Input Image



Overlay Image

Note the two numbered alignment points displayed over each image. Point 1 causes the center of the nose to appear in the center of the dark area in the result image. Point 2 determines the size and angle of the head by specifying where the center of

the hair line will fall.



Result Image

Note that face has been adjusted to bring both of the overlay image's alignment points into coincidence with the corresponding alignment points in the base image. By selecting the Lighten operation, only those pixels of the overlay image that are lighter than the corresponding pixels in the base image are incorporated into the result image, creating the lattice-like effect.

Stitching Two Images Together

The Composite transformation can be used to join two images together. For example, you can create a custom map that spans two adjacent topographical maps by scanning the appropriate sections of both and then combining them into a single image. Or you may have two overlapping landscape photos taken from the same location that you want to join into a composite.

The one thing to remember about the Composite transformation when joining images together is that the result image is always the same size as the base image. If you want the result image to extend beyond the edge of the base image, and this is normally the case when joining images, then you must first create room for the second image by adding a border to the first image. The easiest way to do this is with the Crop transformation (under Geometry). This transformation can add a border of any size or color to an image along any or all of its four sides.

Once you have created space for the second image, create a mask for the first image that identifies the region you want to cover with the second image. This is normally best done by selecting a rectangle or quadrilateral mask. If you want to blend the two images smoothly, soften the edge of the mask; to butt the two images together with a sharp edge, leave a hard edge on the mask.

Next, select the first image and start up the Composite transformation. Select your mask as the Amount mask for the base image and select the second image as the Overlay. Select 2-point alignment and choose alignment points over the same objects in both images to get the two images to align perfectly. Click Preview to check your progress -- make any necessary refinements in the alignment or mask and then click OK and you're done.

To composite two wide angle shots into a single panoramic photo that really looks natural, you will need to use 4-point alignment and a blending mask to hide the seam between the two images. This technique only works for rectilinear lenses (a rectilinear lens is one that reproduces straight lines as straight lines. It will not work for fisheye lenses or other lenses that distort straight lines into curves).

Here are two overlapping photos taken from the same location using a 20mm lens on a 35mm camera body:



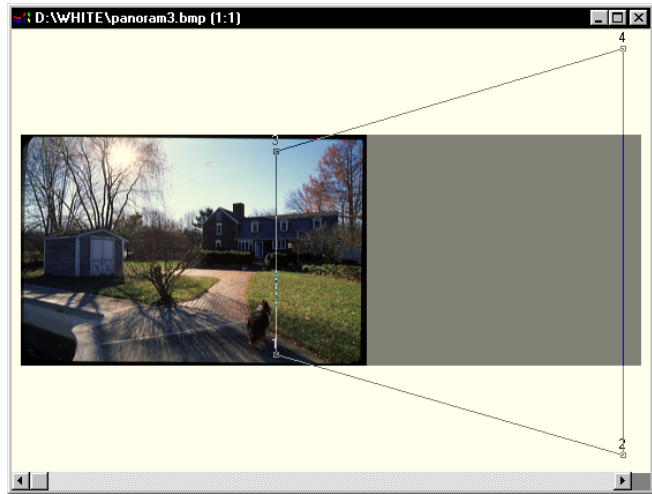
The first step in creating a composite is to add a border to the base images to make a place to drop the overlay image onto:



Next, create a blending mask for the overlay image which is all white except along the left edge where it fades gradually to black. When you create the composite image (see next step) select this mask image into the Overlay Amount control:




The final step is to composite the overlay image with the base image using 4-point alignment to perspective warp the overlay image to match the base image.



Alignment points positioned over base image.
(Alignment points left in corners of overlay image)



Result image

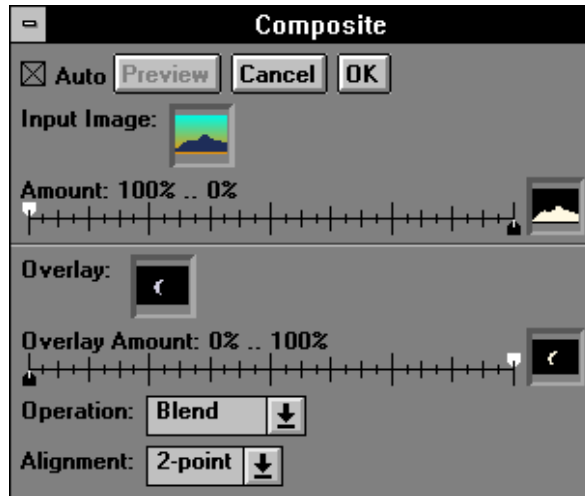


Result image trimmed to remove excess


Using both an Amount mask and an Overlay Amount mask

As soon as you use 1-point, 2-point, 3-point, or 4-point alignment, there is a difference between the Amount mask (a mask image selected into the Amount control) and the Overlay Amount mask (a mask image selected into the Overlay Amount control). The Amount mask always remains aligned with the input image while the Overlay Amount mask is moved, scaled, rotated, and skewed along with the overlay image. Thus if you want to identify an object in the overlay image to be placed over the base image, you should create an appropriate mask and select it into the Overlay Amount control. To protect part of the base image from being covered by the overlay image, you should create a mask to cover the area you want to protect and select it into the Amount control.


The following example shows the different functions of the two masks -- the images have been simplified for purposes of illustration. The base image contains a dark mountain outlined against a lighter sky, and a mask has been constructed for the base image that isolates the mountain from the background sky. The overlay image is a light blue crescent moon against a black sky background, and the overlay mask isolates the moon from its background. The goal of this exercise is to place the moon into the sky, behind the mountain:



Composite Dialog Box -- Note the reversed positions of Amount sliders to allow overlay of sky area only.



Input Image




Amount Mask Image



Overlay Image



Overlay Amount Mask Image



Result Image

The result image shows the desired effect. For convenience, alignment points were placed on each of the two tips of the crescent moon in the overlay image. Then, by positioning the corresponding points in the base image, the moon could be placed accurately in any location. If an Amount mask were not used, the moon would cover the part of the mountain that it overlaps. If an Overlay Amount mask were not used, then the black background sky from the moon image would be overlaid over the base image.

General Masking Tips and Techniques

Seeing the mask

In the process of creating a mask, you may find it useful to reverse the masked and unmasked regions with the Invert button to help visualize the edges of the masked region. It may also be useful to vary the color of the mask overlay and the mask display controls. Sometimes making the mask opaque or displaying only the mask helps you see exactly what part of the image the mask covers or to detect pinholes in the mask. Zooming in on the mask is also a useful technique for detecting pinholes too small to show up at lower magnification factors.

Working from rough outline to fine detail

To mask a complex region, start by making a crude mask of the entire region with the outline tool, then zoom in and refine the mask by adding or removing small amounts at a time around the edges.

Using the polygon tool

The polygon tool lets you quickly create masks for a buildings or other objects with hard, straight edges. To create a complex outline, you start by clicking and dragging the polygon tool to create a rectangle that more or less covers the area of the object you want to mask. To add vertices to the polygon, hold down either Shift key while clicking with the cursor at the desired location. This will replace the nearest edge of the existing polygon with two new edges that pass through the new vertex. Unlike the freehand lasso tool, you can reposition polygon vertices by clicking and dragging them, so you don't have to be too careful how you place them.

Polygon edges are antialiased so you can get really precise, crisp edges -- this also means you can position vertices with an accuracy of fraction of a pixel if you zoom in on the image to a high magnification.

Using the spline tool

The spline tool is very similar to the polygon tool with the exception that instead of connecting the vertices with straight lines, it passes a smooth curve through the vertices. This is useful when you need to create a mask for an object that has smooth edges that are curved such as the edge of a leaf or the outline of a arched doorway.

You add vertices the same way as for the polygon tool -- to control the shape of the curve, simply add more vertices. If you need the curve to change direction sharply to follow a corner or sharp point, simply insert two vertices at the same location.

Using the color range mask tool

The color range mask tool is very useful for creating a mask based on the color of a region of the input image.

The normal technique for using this tool is to click on a pixel in the image you want to mask to define one of the colors you want the mask to cover. Next, click the Contract button to shrink the range down to just that color. Finally, hold the Shift key down while passing the cursor over other parts of the image that you want to include in the mask. As you do this, the range will be automatically extended as necessary to include all the colors you pass over.

To soften the edges of the mask, you can drag the lower set of markers for each of the selected channels to partially mask colors that are a little outside of the selected range.

Filling pinholes in a color range mask

When you create a mask with the range tool, it may have small pinholes. You can fill these quickly by painting over them with the mask paint tool using an opaque, hard-edged brush.

Optimizing performance

The range tool works faster if you first isolate a region with the outline tool and then switch to overlap mode. This lets Picture Window ignore areas outside of the outlined region when applying the range tool.

Using the brush tool

The mask brush tool lets you paint directly on the mask with a brush of adjustable radius, transparency and softness. A large hard-edged brush is good for masking out large areas quickly. A smaller hard-edged brush is good for filling pinholes left behind by the color range tool.

Masking around the edges of complex objects

One of the most difficult masking situations arises when you need to mask an object with complicated edges against an irregular background. One technique is to select *Similar Pixels (Fixed)* or *Similar Pixels (Track)* in the *Apply To* option. Then, as you paint around the outside of an object the effects of the brush will be limited to colors that match the “background” color at the center of the brush within an adjustable tolerance (the lower the tolerance, the closer pixels have to match the background color to be included in the mask).

If you use the *Similar Pixels (Fixed)* option, the brush tool samples the background color once when you start a new stroke -- this is normally the most useful option. If you use the *Similar Pixels (Track)* option, the tool samples the background color each time you move the cursor, automatically tracking variations in the background. When you use either of these options, make sure you keep your brush just outside the edge so you don’t accidentally pick up colors from the object.


Combining the color range with the brush tool

Another technique for making the mask brush tool more intelligent is to select *Pixels in Color Range* in the *Apply To* option. Now when you apply the brush, only pixels that fall within the current color range will be affected. You can flip back and forth between the color range and brush tools to change the color range as you work with the brush.

Creating density masks

The brightness curve mask tool lets you create *density masks*. A density mask is simply a mask whose transparency is a function of the brightness of the image. Where this comes in handy is solving problems like increasing the saturation of the midtones in an image. If you create a mask which is white where the image has midlevel brightness and fades to black in the both shadow and highlight areas (see the middle example below), you can then use this with the Saturation transforma-

tion to adjust the saturation of just the midtones. By adjusting the location, slope and shape of the transitions from white to black, you can fine tune the results. Here are some examples of simple curves and their effects:



Feathering

Feathering is very important especially when using masks with transformations. If you don't feather the edges of a mask, then the boundary between the transformed and unmodified regions of the input image will often be very obvious. Feathering blends the effects of the transformation around the edges and creates a more natural result. The larger the feather radius, the more the edges of the mask are softened.

You may find it useful to save a copy of your mask before feathering. This lets you go back and refeather the mask after trying it out.

For most transformations that adjust the brightness and color of an image, gradual feathering that falls off out from the regions you are modifying usually creates the best results (4-10 pixel radius).

When using the composite transformation, a smaller feather radius (1 to 4 pixels) usually works better.

When using the blur transformation to remove grain from the sky, try using the range tool to select the sky region and then feather with a negative radius and sudden falloff to avoid blurring the horizon line or any foreground elements silhouetted against the sky.

To feather a mask more than is possible using the feathering tool, you can use the mask blur tool. When you use this technique, the mask will shrink by approxi-

mately the radius of the blur, so you need to start with a slightly larger mask to allow for this erosion.