75th birthday series

The Secret Life of Masks in Photoshop

Douglas A. Kerr P.E. (Ret.)

Issue 2
May 19, 2011

ABSTRACT AND INTRODUCTION

In digital image processing, a *mask* is a map that determines, for each pixel location, to what extent (if any) the pixel there is susceptible to various “happenings”.

The image editing application **Adobe Photoshop** includes three principal kinds of *masks*. These masks can serve to isolate parts of the image that can then be included in overall image makeup, deleted, moved, or modified—or even “partially included”.

In this article, we explain the nature of these three masks, what they influence and how, and how we can construct and modify them.

THE CONCEPT OF A MASK

Simplistically, in digital image processing, a *mask* is a map of the image frame in which each point (each pixel location, for pixel-oriented images) has a property that, for openers, we will consider as being either transparent or opaque. (We will see later that in many cases there can be intermediate values of that property, which we think of as being "semi-transparent".)

Basically, where the mask is transparent, various happenings can "occur". Where it is opaque, the happenings are blocked. The kinds of happenings, and the significance of their “getting through” or not, vary with the mask and how it is used.

CAVEAT

Photoshop has an enormously rich repertoire of features, often with many ways to access them. This article does not pretend to describe everything that is known about masks in Photoshop, nor to be a treatise on their effective use. The latter is better left to those with more experience than I.

Rather, here I attempt to accurately describe, in considerable detail, the principal kinds of masks used in Photoshop from a functional point of view.

PLATFORM

Specific references in this article are to Photoshop CS5, with user interface details for operation on a Windows XP platform.

THE PRINCIPAL PHOTOSHOP MASKS

There are three principal masks in Photoshop (the names here are those used in the Photoshop documentation):

- Selection
- The Layer Map or Pixel Map
- The Vector Map

SELECTION

Notation

The Photoshop functionality "Selection" operates by way of a *mask*[1] and we must recognize this to grasp all its properties and to fit it into the overall story. I will call this mask the "Selection mask" (a term we will, however, not find in the Photoshop documentation).

One way of looking at its operation (as with all the masks we will study) is that it allows us to **select** certain portions of the image such that only they will be affected by various things we can do.

Basic properties

The Selection mask is a *global* mask, by which I mean it is not associated with a layer. Only a single Selection mask can be in existence at any time, and it influences any layer that contains image content.

It is a pixel-oriented mask—it has a value for each pixel location of the image. For the moment, we will recognize two possible values, *transparent* and *opaque*.

Basic generation

If we "make a selection" (using the Ellipse, Rectangle, or Row Marquee tools, or either of the Quick Selection tools), a Selection mask comes into existence. The interior of the closed figure that is

[1] Some say "not so". I discuss this at the end of the article.

"drawn" becomes a *transparent* region in the mask, and the remainder of the mask is *opaque*.[2]

The selection

It can be useful to speak of the transparent region as the *selection*.

Seeing the mask

The boundaries between the transparent and opaque regions of the mask are visible on the screen with a *marquee* ("marching ants"); it in effect shows the boundary of the selection.

Dismissal

The Selection mask can be dismissed in various ways (including the keyboard shortcut Ctrl + D).

Inversion

The Selection mask can be inverted (opaque and transparent regions interchanged).

Procedure: (Menu) Select > Inverse;
(KB) Shift + Ctrl + I;
Right click anywhere; choose Select Inverse on the context menu.

Select All

If we do Edit > Select All (Ctrl + A), a Selection mask is created that is everywhere transparent.

Adding or deleting transparent region area

Once a Selection mask has been created, by making an initial transparent region ("selection"), we can add other transparent areas (disjoint from the original one or overlapping it), delete parts of the transparent region, or draw a figure whose intersection with the existing region (their common area) becomes the transparent region.

Again, the marquee tool is used, and we select the specific operation from icons on the tool ribbon: Add to selection, Subtract from selection, or Intersect with selection.

---

[2] Regardless of whether the tool (at this first "hit") is set to "new selection", "add to selection", "subtract from selection", or "intersect with selection".

Transformation of the mask

We can move the entire Selection mask.

Procedure: With any tool except Move selected, right click anywhere. In the context menu, click on Transform Selection. Drag anywhere except on one of the little square handles that appear on the corners of the visible boundary of the transparent area. The entire mask will move.

When done, Hit Enter to finalize the transformation.

Or, to abandon the change, hit Esc.

The entire original mask (covering the entire frame) moves. The area "cleared" by the original mask moving becomes opaque in the resulting new mask.

Because of this, if the original mask had an opaque "background", it will seem as if only the transparent area moved.

We can rotate the entire Selection mask.

Procedure: With any tool except Move selected, right click anywhere. In the context menu, click on Transform Selection.

Drag anywhere outside the bounding box embracing all the transparent areas. The entire mask will rotate about a pivot point shown with a little "target". The pivot point can be moved by just dragging it.

Use Shift-drag to constrain the rotation to multiples of 15°.

When done, Hit Enter to finalize the transformation.

Or, to abandon the change, hit Esc.

As with moving the mask, the area "cleared" by the original mask rotating becomes opaque in the resulting new mask.

Because of this, if the original mask had an opaque "background", it will seem as if only the transparent area moved.

We can scale or skew the combination of all the transparent areas.

Procedure: With any tool except Move selected, right click anywhere. In the context menu, click on Transform Selection. Little square handles will appear on the corners of the "bounding box" of the overall transparent region and at the centers of its sides. (If the region is only a single rectangle, the bounding box will be the same as the region boundary.)

To stretch or compress the collective transparent region in width, drag one of the handles on a vertical side of the bounding box.

To stretch or compress the collective transparent region in height, drag one of the handles on a horizontal side of the bounding box.

To change both the height and the width, drag one of the handles at a corner of the bounding box. Use Shift-drag to maintain the aspect ratio.

To reshape the collective transparent area (perhaps in a "trapezoidal" way), use Ctrl-drag on one of the corner handles. To constrain the movement of the corner to either a vertical or horizontal direction, use Shift + Ctrl-drag

To skew the collective transparent region, Ctrl-drag one of the handles on an edge of the bounding box.

When done, Hit Enter to finalize the transformation.

Or, to abandon the changes, hit Esc.

THE QUICK MASK EDIT MODE

This functionality is an alternate way to make or modify the Selection mask.

If no Selection mask already exists

If we click on the Edit in Quick Mask Mode icon (or press Q), the system enters the Quick Mask Edit mode. What Photoshop calls a *Temporary Mask* is created, initially everywhere transparent.

Assuming the "default" version of this mode, if we paint (with a Brush tool, for example) using black "paint", where we paint becomes opaque; the rest of the mask remains transparent.

If we paint with white over an area that we had earlier painted black, this restores that area to transparent.

Seeing the mask

While in the Quick Mask Edit mode, the Temporary Mask is shown to us with a semitransparent red overlay (reminiscent of the "Rubylith" masks used in black and white film photography). In the default form of the Quick Mask Edit mode, the red shows us the opaque regions.

“Depositing” the resulting mask

When we exit the Quick Mask Edit mode, **if we had painted any with black**, a Selection mask is created, identical to the Temporary Mask (and the Temporary Mask disappears). Even if we had overpainted all our original black painting with white, this will still happen, creating a Selection mask that is everywhere transparent.

When we exit the Quick Mask Edit mode, **if we had never painted with black**, whether or not we painted any with white, no Selection mask is created.

It might seem that having no Selection mask is no different from having a Selection mask that is everywhere transparent, but there is a subtle difference.

If there is a Selection mask in effect and we do Copy, all of the image pixels that are in transparent regions of the mask are copied to the Clipboard, and can be Pasted (they will go to a newly-created layer). If

that Selection mask is everywhere transparent, that will be all of the pixels. (It is as if we had done Select All.)

However, if there is **no** Selection mask in effect, and we do Copy, nothing is copied to the Clipboard, and it retains whatever was put on it last (perhaps from another image, or another day, or even another application). If we Paste, that other content is what will be pasted into this image, on a new layer. (Could be a shock!)

If a Selection mask already exists

If there is an existing Selection mask and we invoke the Quick Mask Edit mode, the Temporary Mask will initially be identical to the Selection mask. The overlay will show us its current structure (the overlay color by default indicating the opaque regions).

We can then modify the mask by "painting" (or the equivalent). When the Quick Mask Edit Mode is exited, the Selection mask will be updated to reflect these modifications.

Transformation of the mask

When in the Quick Mask Edit mode, we can transform the mask in all the same ways we can when in the normal mode, with one difference (shown in green):

Procedure: (For all transforms) With any tool except Move selected, right click anywhere. In the context menu, click on Free Transform.

Etc. etc.

Alternate Quick Mask Edit mode

When using the Quick Mask Edit mode to establish or modify a Selection mask, we can choose an "inverted" form in which the red overlay shows the transparent (not opaque) regions, painting with black creates a transparent (not opaque) area, and painting with white creates an opaque (not transparent) area.

Nevertheless, the implications of the resulting opaque, transparent, and semitransparent regions are the same as always. It is just that the paint colors, and the overlay indication, work upside down.

Procedure: Double click on the Quick Mask Edit Mode icon; set "Color Indicates: Selected Areas" in dialog

Overlay color options

In addition to the two options as to the "sense" of the color overlay, we can vary its color and its opacity.

With an opacity of 100%, wherever the indication shows red (or whatever color), we cannot see the actual image pixels; with a lower opacity, we can see the image through the overlay.

Procedure: Double click on the Quick Mask Edit Mode icon; set Color and Opacity

"Painting" via paths

If, when in the Quick Mask Edit mode, we stroke or fill a *path*, this has the same effect as painting; that is, it can create a transparent or opaque region of the Selection mask.

Not an additional mask

It is important to recognize that the Temporary Mask on which we work in the Quick Mask Edit mode is not a mask that is ever used for any "mask" purpose in Photoshop operations. It is only a "Selection mask in process".

It is interesting that Photoshop chooses to call this construct in process a "mask", while not calling the construct that it becomes when finished a "mask".

WHAT THE SELECTION MASK DOES

Based only on what we have spoken of so far (to keep things reasonably simple), the impact of a Selection mask includes these major aspects:

- If we put the mouse pointer anywhere in the selection (the transparent region) and drag, all the pixels in the transparent region are dragged to another location on the same layer. The original locations of these pixels now carry "transparent" (impotent) pixels.

- If we put the mouse pointer anywhere in the selection and Alt-drag, copies of all the pixels in the selection are dragged to another location on the same layer. The original locations of these pixels retain their original values.

- If we do Delete, all pixels in the selection are replaced by transparent (impotent) pixels.

- If we do Copy (Ctrl + C), a copy of all pixels in the selection go to the Clipboard.

- If we do Cut (Ctrl + X), all pixels in the selection go to the Clipboard. The original locations of these pixels now carry transparent (impotent) pixels.

- If we apply to the layer an adjustment (such as a change in contrast or lightness), it only affects the pixels in the selection.

- If we apply a fill to the layer, only pixels within the selection are overlaid by the fill color or pattern. (Since the fill pixels are normally fully-opaque, in effect the existing pixels are replaced.)

What we will learn next will add some complication to the above. We'll revisit it later to take these complications into account.

SEMITRANSPARENCY

Introduction

If, while in the Quick Mask Edit mode, we paint with a color that is not black or white but rather some in-between gray, we create a region in the Selection mask that is "semitransparent". Such a region will "let stuff through partially", a concept whose further explanation will come later. (The basic notion seems obvious, but exactly what it means in each situation is a bit tricky.)

We will introduce the notion of a transparency parameter, t, for the mask at each point. In a (fully) transparent region, $t = 1$. In a (fully) opaque region, $t = 0$. In a semitransparent region, t has some value in between.

If in some region $t = 0.45$, we can think of this as meaning that in that region, "stuff is let 45% through". Again, exactly what this means for different kinds of stuff will have to come later.

It turns out that if we "paint" a certain mask region with a gray color for which R,G,B = C,C,C, then for that region:

$$t = \left(\frac{C}{256} \right)^{1.25} \qquad (1)$$

The selection

We can (cautiously) consider all regions in which $t > 0$ to be the *selection*.

Seeing the mask

In "normal" mode (not the Quick Mask Edit mode), we will see outlined by "marching ants" all regions in which $t > 0.5$. I call this the "indicated selection".

The meaning of *t*

If, over a certain region of our mask, $t = 0.45$, what does this mean? With regard to "letting through" some image adjustment, I seems to mean this: the pixel RGB values at a certain point (in that region), after the modification, are the sum of 0.45 times the what would be the fully modified RGB values plus 0.55 (1-0.45) times the unmodified RGB values. That is, “45% of the change happens”. Maybe.

If we drag on any spot in the **indicated selection** (where $t > 0.5$), all pixels that are in the **selection** (where $t > 0$) will be dragged to a new location. Or if we do Cut or Copy, all pixels in the selection will go to the Clipboard. and can then be Pasted onto (for example) a new layer of the image.

In either such case, if $t = 1$, then the dragged or Pasted pixels arrive in their original state. But what if t is, for example, 0.45?

Recall that in most situations a pixel has both a *color* and an *opacity*. The lower the opacity, the more the pixel is "transparent". Normally, pixel opacity is 1. The influence of a pixel's color comes from the pixel's opacity. If the pixel’s opacity is zero (so it is fully transparent), it is impotent and thus invisible.

If we have some image pixels with color R,G,B = 80,128,92 (and 100% opacity, the normal situation), and the mask there has $t = 0.45$, then when the pixels arrive at their new locations, they still have (for themselves) color R,G,B = 80,128,92 but opacity 45% (they are only "45% potent).

Then if, for example, the existing "background" at their new location is black (RGB = 0), the resulting "composite" color at the new location of the pixels is now R,G,B = 36,57,41.[3]

Note that in this regard, pixel opacity ultimately works as a factor on the (nonlinear) R, G, and B values, not on the underlying linear values (often indicated as r, g, and b). The same is true with regard to mask transparency, t.

Semitransparency in other masks

We will find this identical concept in the other two masks we will study.

---

[3] Assuming the “normal” blend mode.

The red overlay revisited

In the Quick Mask Edit mode, in its default operation, a red overlay shows the opaque area. For semitransparent areas, the red overlay still appears, but with reduced opacity. Essentially, its opacity is scaled by 1-*t*. The limit is reached for transparent areas (*t* = 0), where the overlay opacity becomes zero—that is, we see no overlay (just as we earlier learned).

WHAT THE SELECTION MASK DOES—UPDATE

Now that we have made the acquaintance of semitransparency, we need to update our summary of the impact of the Selection mask:

- If we put the mouse pointer anywhere in the indicated selection and drag, a copy of all pixels in the selection will be dragged to another location on the same layer. Their opacity ("potency") at their new location is scaled by the transparency value (*t*) in the mask for their original locations. Left behind is the "remainder" of each pixel"—that is, the original pixels but with their opacity scaled by 1-*t*. From a visible standpoint they are lying atop a background of the checkerboard pattern indicating "nullity", its own visible density diluted by the presence of the pixel remainder lying above it.

- If we put the mouse pointer anywhere in the indicated selection and Alt-drag, all pixels in the selection will be dragged to another location on the same layer. Their opacity ("potency") at their new location is scaled by the *t* for their original locations. Left behind are the pixels with their original opacity.

- If we do Delete, the fraction *t* of the opacity of all pixels in the selection disappears, leaving the "remainders" of themselves, with opacity 1-*t*. From a visible standpoint they are lying atop a background of the checkerboard pattern indicating nullity, its own visible density diluted by the presence of the semitransparent pixels lying above it.

- If we do Copy (Ctrl + C), all pixels in the selection are copied to the Clipboard with their opacity scaled by the mask transparency value at their location (*t*). If they are Pasted, their opacity at their new location is that scaled value. The source pixels are not affected.

- If we do Cut (Ctrl + X), all pixels in the selection are copied to the Clipboard with their opacity scaled by *t*. If they are Pasted, their opacity at their new location is that scaled value. At the source locations, we have the "remainder" of each pixel, with opacity 1-*t*. From a visible standpoint they are lying atop a background of the

checkerboard pattern indicating nullity, its own visible density diluted by the presence of the pixel remainder lying above it.

- If we apply to the layer an adjustment (such as a change in contrast or lightness), it only affects pixels in the selection. The degree of the change is scaled by the value of *t* in the mask for each location.

- If we apply a fill to the layer, only pixels in the selection are overlaid by the fill color or pattern. The opacity of the new pixels is scaled by the value of *t* in the mask for each location. The opacity of the existing pixels is, in effect, reduced to 1-*t* of their original value by the semitransparent "fill" pixels above them.

No effect on visibility

Note that, although the Selection mask influences what pixels can be "exported", and on whether or not they are affected by various operations, and on their opacity (and thus appearance) in new lives they may receive, it has no influence on the visibility of pixels in place. (That is, the appearance of pixels in place is not one of the kinds of "stuff" that this mask either lets through, blocks, or partially lets through.)

Another wrinkle

If we have not yet "made a selection", is a "default Selection mask" in place (perhaps one that is everywhere transparent)? In one sense yes: every place on the image is accessible to modification. In another sense no: no part of the image is enabled to be affected by a drag, delete, Cut, or Copy operation.

Tactical use

Typically, a Selection mask is used in a tactical sense: it controls some operation, and is then dismissed.

LAYER-ASSOCIATED MASKS

The Selection mask is global: its impact is on whatever layer is active. Operations governed by it in general only affect the active layer.

In contrast, the two other principal masks, the Pixel Mask/Layer Mask and the Vector Mask, are associated with a specific layer.

Very often, these masks are persistent. That is, they are often left in place for an extended period, often to govern what portions of the basic material placed in a layer are of any effect on the overall image being "constructed", or what portion of the image is affected by an adjustment layer.

We will discuss these two masks in the following sections.

THE PIXEL MASK (LAYER MASK)

This mask is associated with a particular layer, and is pixel oriented. It is sometimes called in the Photoshop documentation the “Pixel Mask”, and sometimes (almost randomly) the “Layer Mask”. For consistency, I will call this the Pixel Mask.

Basic generation

Assume for the moment that there is no Selection mask in effect.

If, with a certain layer selected, in the Masks panel we click on the "Add a pixel mask" icon, a Pixel Mask is created, associated with that layer. By default, it is initially everywhere transparent.

If, with the pixel mask icon on the layer control bar selected, if we paint (with a brush tool, for example) with black "paint", we create an opaque region on the mask. If we paint with white "paint", we create a transparent region (actually meaning taking away from an opaque region we already drew—it would mean nothing to paint “transparency” in a region that is already transparent).

Seeing the mask

We can "see" the mask in one of these three ways:

- By its effect. For example, if the lower layers are made invisible, then, for the regions in which our mask is opaque, we will see the “nullity” pattern (usually a gray checkerboard).[4]

- As a grayscale map, with black indicating the opaque regions and white indicating the transparent regions.

  Procedure: Alt + Click on the layer mask icon in the layer control bar.

- With a red overlay, red for the opaque regions.

  Procedure: Alt + Shift + Click on the layer mask icon in the layer control bar;
  Or click on the layer mask icon on the layer control bar and key “\”.

There are some other visibility modes. If we click on the on the pixel mask thumbnail in the layer control bar and then right-click on it, and

---

[4] This is usually described as the pattern that “indicates pixel transparency”. In this context (and especially when we deal with multiple layers), a more insightful way to think of it is that, if we have fully transparent pixels in a region (with nothing below), we see right through them to this visible indicator of “nothingness”. If we have partly transparent pixels in a region, we partly see through them to this indicator (which will then appear “faded”).

select Refine Mask, a dialog will open that includes various tools for refining the boundaries of the mask, a matter that is beyond the scope of those article.

But this dialog also includes a dropdown menu that offers several different ways to see the mask. These include the three noted above, plus these:

- Reveal Layer. We see only the layer of interest, with no effect of the mask.
- Marching Ants. Here, the boundaries of the mask are indicated with a "marching ants" marquee, just as we commonly have for the Selection mask. The mask is in effect.
- On White. Here, we see the material admitted by the transparent regions on a white background.
- On Black. Here, we see the material admitted by the transparent regions on a white background.

Inheritance from a Selection mask

If a Selection mask is in effect when, in the Masks panel (with a layer selected), we click on the "Add a pixel mask" icon, a Pixel Mask is created, associated with that layer, matching the Selection mask. The Selection mask itself is dismissed.

Semitransparency

Just as with the Selection mask (when using the Quick Mask Edit mode), with the Pixel Mask of a layer selected, we can paint with a gray color and create a semitransparent region.

As with the Selection mask, if the gray color we paint with is R,G,B = C, C, C, then, for the region we create, the transparency, *t*, is set to:

$$t = \left(\frac{C}{256} \right)^{1.25} \qquad [1]$$

The implications of varying degrees of transparency in the mask are very comparable to those for the Selection mask, as applicable to the various kinds of stuff.

Transparency scaling

We can make all the black-painted regions semitransparent rather than opaque, and gray-painted regions less opaque, over the entire mask,

by setting the Density on the Masks panel (with the mask selected) to a value less than 100%.

Feathered edges

We can set a Feather property for a Pixel Mask. This makes for a "not completely sharp" transition between transparent and opaque or semitransparent regions. This is done on the Masks panel with the mask selected.

Alternate initialization

If we create a Pixel Mask by using Alt+Click on the Add a Pixel Mask icon on the Masks panel (rather than just Click), the mask that is created is initially everywhere opaque. The effects of painting on the mask, with black, white, or gray paint, is no different than before.

Inversion

A Pixel Mask can be inverted (opaque and transparent regions interchanged).

Procedure: Select the mask
On the Masks panel, click on Invert.

No transformation

There is no way to shift or otherwise to transform a Pixel Mask.

If we transform the layer, the mask is transformed with the content (continues to enclose the same image features).

But there is a workaround.

- Create a Selection mask from the Pixel Mask.

  Procedure: With the Pixel Mask selected, in the Masks panel, click on Load Selection from Mask (the little dotted circles at the lower left of the panel).

- Delete the Pixel Mask.

  Procedure: Select the Pixel Mask.

  In the Masks panel, Click on Delete Mask (the garbage can at the lower right of the panel).

- Modify the Selection mask in whatever way is needed, either in normal mode or in the Quick Mask Edit mode.

- Create a new Pixel Mask from the Selection mask.

  Procedure: With the layer selected, in the Masks panel, click on Add a Pixel Mask.

THE VECTOR MASK

Introduction

A Vector Mask is associated with a particular layer. It is not a map with a value for every pixel location in the image frame. Rather, it comprises regions (transparent and opaque or semitransparent) separated by boundaries that have mathematical descriptions. These boundaries are *paths*, as this term is used in Photoshop, and are composed of (straight) line segments and cubic Bézier curves.

If, in the Mask panel for a certain layer, we click on the "Add a vector mask" icon, a Vector Mask is created, connected to that layer. It is initially transparent everywhere.

If, with the vector mask icon on the layer control bar selected, we draw a path (using a shape tool or one of the pen tools), with “Add to Path” selected for the tool, the region enclosed by that path remains transparent, and the rest of the mask becomes opaque. If instead we have selected “Subtract from Path” for the tool, the region enclosed by the path becomes opaque, and the remainder of the mask remains transparent.

That is:

- “Add to Path” means “make a transparent region”. (We might think of it as a “selection”.)
- “Subtract from Path” means “make an opaque region”
- Whatever the “sense” of the first hit, the remainder of the mask is made (or left in) the opposite sense.

We can "see" the mask by its effect. If the lower layers are made invisible, then, for the regions in which our mask is opaque, we will see the “nullity” pattern (usually a gray checkerboard).

The paths that serve as the boundaries between transparent and opaque regions can be edited using the various path editing tools, including the pen tools and the Direct Selection and Path Selection tools. The details of this are beyond the scope of this article.

No inheritance from a Selection mask

A Vector mask cannot be made directly from a Selection mask.

Semitransparency

As before, we can think of the transparent regions of a Vector Mask as having transparency, t, equal to 1, and the opaque regions having t = 0.

We can make all the opaque regions semitransparent by setting the Density to a value less than 100% (using a slider on the Mask control panel with the mask selected).

If we set Density to 55%, we have set *t* for all originally-opaque regions to 0.45.

The implications of the various values of *t* in this mask are almost identical to the implications of the various values of *t* in a Pixel Mask.

Feathered edges

In a similar way, we can set a Feather property for a Vector Mask. This makes for a "not completely sharp" transition between transparent and opaque or semitransparent regions.

Alternate initialization

If we create a Vector Mask by using Alt + Click on the Add a Vector Mask icon on the Masks panel (rather than just Click), the mask that is created is initially everywhere opaque. The effects of painting on the mask, with black. white, or gray paint, is no different than before.

Note that if our first "hit" is with black paint, this creates an opaque region and the reminder of the mask becomes transparent, the same result as with the default initialization. Similarly if our first "hit" is with white paint, this creates a transparent region and the reminder of the mask remains opaque, the same result as with the default initialization.

Thus there is really no point to using the alternate initialization unless our objective is to make a mask that is everywhere opaque.

No inversion

A Vector Mask cannot be directly inverted (transparent and opaque/semitransparent regions interchanged).

No transformation

There is no way to directly shift or otherwise transform a Vector Mask.

If we transform the layer, the mask is transformed with the content (continues to enclose the same image features).

EFFECT OF A LAYER MASK

Introduction

By "Layer Mask" I mean either the Layer Pixel Mask or the Vector Mask. This is a different use of the term than in the Photoshop documentation, where "Layer Mask" is an alternate name for the Pixel Mask.

The effect of a layer map is different whether it is associated with an image layer, a fill layer, or an adjustment layer.

For an image layer

- Wherever the Layer Mask for an image layer is transparent ($t = 1$), the pixels in the layer content contribute to the overall image buildup "at par". If we only have this layer visible, we see these pixels "as they are".

- Wherever the Layer Mask for an image layer is opaque ($t = 0$), the pixels in the layer content do not contribute to the overall image buildup at all. If we only have this layer visible, we do not see these pixels at all, but rather see the checkerboard pattern that indicates nullity.

- Wherever the Layer Mask for an image layer is semitransparent ($0 < t < 1$), each pixel in the layer content contributes to the overall image buildup with its opacity ("potency") scaled by the mask transparency, t, at that point. If we only have this layer visible, we see these pixels "diluted", lying atop the checkerboard pattern that indicates nullity (its own opacity diluted by the semitransparent pixels on top).

If there are both a Layer Pixel Map and a Vector map for a layer, then the effect is as described above, with the map transparency being the product of the transparency of the two masks at each point (I think).

For a fill layer

The effect of a Layer Mask on a fill layer is very similar to that on an image layer.

For an adjustment layer

- Wherever the Layer Mask for an adjustment layer is transparent ($t = 1$), the effect of the specified adjustment on the layer(s)[5] below is "at par".

- Wherever the Layer Mask for an adjustment layer is opaque ($t = 0$), there is no effect of the specified adjustment on the layer(s) below.

- Wherever the Layer Mask for an adjustment layer is semitransparent ($0 < t < 1$), the effect of the specified adjustment on the layer(s) below is scaled by the transparency value (t) of the semitransparent region.

Not a "selection" for drag, Copy, or Cut

Unlike the case with the Selection mask, the pixels in non-opaque regions of a Layer Mask of either kind are not uniquely susceptible of being dragged, Cut, or Copied.

APPLYING A LAYER MASK

Typically, we leave a layer mask in place for an extended period. It may, for example, define the portion of the overall content of its image layer that will be part of the overall image buildup.

This can be advantageous, compared to, for example, "cropping out" only the desired portion of the layer content, since we can, at a late stage of the process, change exactly the portion of that content to be included in the overall image makeup.

However, if we are certain we have identified the appropriate part of the content, we can permanently dispose of the unwanted part by "applying" the mask. This essentially "crops to the mask" the layer content.

On an image or fill layer, we can do this with respect to either a Pixel Mask or a Vector Mask.

Procedure: Right Click on the mask Icon, choose Apply Mask from the context menu

Having done so, the layer mask as such has vanished.

---

[5] An adjustment layer may be set so that its effect is only on the next image layer below or on all image layers below.

ADVANCED MASK MANAGEMENT

There are other ways to make these three kinds of mask than are described here, including ways to make one from another (directly or indirectly), turn masks into paths to be used for other purposes, make masks from paths initially made for other purposes, make masks from "objects" in the image, and the like. These are beyond the scope of this article.

IS THE SELECTION MASK REALLY A MASK?

An assertion

Some assert that the "thingy" that implements the Selection functionality is not properly thought of as a *mask*. The argument typically goes something like this:

- A *mask* blocks part of an image, **excluding** it from some activity.
- The Selection thingy, however, selects part of an image, **including** it in some activity. Thus is it not a mask—rather the opposite of that.

Analysis

In fact, for each of the three masks discussed in this article, the **inclusion** and **exclusion** aspects coexist, always in complement. If we put aside the complicating matter of semitransparency, then any mask:

- **excludes** certain regions of the image, thus **including** the remainder, or, if we prefer to look at it another way,
- **includes** certain regions of the image, thus **excluding** the remainder.

Now, it may well be that what we most often use a particular kind of mask for, or the technique we most often use to construct it, will sometimes seem to emphasize one or the other of these two always-coexisting aspects. But that does not mean that there are two conceptually different kinds of constructs.

An example

Here's an illuminating example. We have a busy shot of a scene in a basketball game, and we wish to highlight the ball by brightening it (so the viewer can easily find it). We might proceed in two ways:

- We select the basketball, and then apply a Brightness adjustment to the layer. The Selection thingy **allows** the adjustment to affect

the basketball, which we want. It also **blocks** the adjustment from affecting the rest of the frame, also what we want.

- We add a Brightness adjustment layer above the image layer. We provide it with a Pixel Mask, and make its transparent region embrace only the basketball. The Pixel Mask **allows** the adjustment to affect the basketball, which we want. It also **blocks** the adjustment from affecting the rest of the frame, also what we want.

These don't sound like opposing concepts, do they?

The language of Photoshop

Why does Adobe not speak of the Selection construct in Photoshop as a *mask*? Probably because that is not necessary to explain its basic uses ("select the basketball so we can brighten it", "select the goat so we can Copy it to the clipboard"). And this practice was doubtless established at a time when the broader notion of masks had not become familiar.

But note that when we wish to construct, or modify, the Selection construct in the Quick Mask Edit mode, while the thingy is in "on the workbench" it is indeed called a mask. Why? Perhaps because in this context, we deal with the thingy in a more complicated way (perhaps now including the concept of semitransparency), and its true maskiness is manifest.

Whether or not we think it would help to tell the basic user that the Selection construct is a *mask*, it is one.

In another application

By way of comparison, in Picture Publisher (the image editor I mostly use), the functionality that is directly comparable to Selection in Photoshop is spoken of as *masking* (there is only that one).

The discussion of the mask in the documentation primarily reflects the **inclusive** aspect, but also discusses the **exclusive** aspect.

A CHANCE FOR MISUNDERSTANDING

I recently posted a note on an important digital photography forum summarizing the steps to conveniently make a layer pixel mask by "painting" it with the image content all remaining visible for guidance.

An esteemed colleague, a skilled Photoshop operator, reported that he was surprised by my description, in which I said that where the mask was "painted white", it was transparent, with the result that pixels in

that area were “admitted” to the image buildup. He said that, to the contrary, normally the “painted white” part of a layer pixel mask “blocked” pixels.

It took a little dialog to sort this out. It turns out that what he was speaking of was this: If we have a mask on a layer, then where that mask is “white”, the pixels on a layer beneath were blocked.

Here is how that happens. Where the mask is “white”, it is transparent, and therefore “admits” the pixels on its layer (which is what it controls) to the image buildup. Assuming that the blending mode is “normal”, then such pixels, residing on an upper layer, and being opaque (as pixels normally are unless their opacity has been diluted by some effect) are “in front of” the pixels of the lower layer, which are thereby blocked from view.

But it is inappropriate to say, because of this two-stage phenomenon, that the transparent (”white”) portion of the mask “blocks pixels”.

I have been assured by my colleague that, in discussion of Photoshop technique, this is the recognized convention: in discussing the impact of a “white” or “black” region of a mask, we speak not of its effect on the pixels it governs (in its layer) but rather the secondary effect of those pixels on a lower layer (assuming there is one, of course—often there isn’t). I hope that really isn’t so.

In any event, I discourage descriptions based on this outlook.

#