

THE BOOK OF™ GIMP

A COMPLETE GUIDE
TO NEARLY EVERYTHING

OLIVIER LECARME AND KARINE DELVARE



5

Composite Photography

In Chapter 2, we retouched existing images. In this chapter, we show you how to create unique pictures, using multiple photographs to build composite images.

5.1 Tutorial: Building a Composite Portrait

A *composite portrait* is an image that combines several different portraits to build a new portrait. For example, you could merge several artists' renderings of eyewitness accounts to create the most accurate possible likeness of a dangerous criminal. You could also use these techniques to guess what the child of two people would look like or to invent a strange hybrid animal.

In this example, we create a composite of several young women who kindly agreed to let us use their pictures for this demonstration.

Each portrait was taken in front of the same neutral background, with the same lighting, and with the girls in approximately the same position. But differences in skin hue or in the tilt of their head will create some interesting challenges. Using photographs taken under disparate conditions would be much more difficult,

though. Imagine trying to copy a nose lit from the left onto a face lit from the right. Of course, you could flip the first image vertically, but that could lead to undesirable effects elsewhere in the image. Scaling a feature because the source photograph is much larger—or much smaller—than the target one would lead to differences in sharpness. And putting the chin of a middle-age man on the face of a young woman, or vice versa, would, for many reasons, look odd. For this example, we chose conditions that make the project feasible and simple while still being instructive.

Begin with the four 900×1200 portraits shown in Figure 5.1. The four portraits are arranged as layers within a single image. Our composite will have the eyes from portrait 1, the hair from portrait 2, the nose from portrait 3, and the mouth and chin from portrait 4.

The first thing you need to do is correct the lighting so it's the same across all four portraits. Use the **Image: Colors > Levels** tool on each layer separately: Select the layer in the Layers dialog and then the Levels tool, and click **AUTO**. The result is shown in Figure 5.2. The skin hues still vary widely but less than they did before the adjustment.



Figure 5.1 The four source portraits



Figure 5.2 After lighting correction

Taking Features from Different Portraits

Of all the features, the girls' hair is the most difficult to copy. No selection tool could capture it precisely, and correcting it with the Quick Mask tool would be tedious—and extremely difficult. The best solution is, therefore, to copy the other features onto the face with the hair you want (for this example, portrait 2). The composite girl will be wearing the black shirt from this portrait, but all of the girls are dressed in black anyway, so the shirt doesn't make much of a difference.

To create the new image that will be the composite portrait, simply click and drag the layer of portrait 2 from the Layers dialog to the Toolbox. Immediately save this image (**SHIFT+CTRL+S**), and from time to time, save it again (**CTRL+S**), so you don't risk losing all your work.

Next, copy the features from the other portraits. You can do this in two ways:

- Select the feature in the source image, copy it, paste it into the target image, and create a new layer (**SHIFT+CTRL+N**) with the resulting floating selection.
- Use the Clone tool to copy from the source image to the target one, preferably into a new transparent layer.

Each technique has advantages and drawbacks, so we show you both techniques and let you decide which one works best for you. Creating a new layer for each new feature is important because it allows you to change one feature without having to rebuild the whole portrait. Containing the features in different layers also lets you make adjustments to the color,

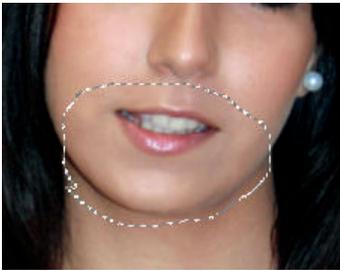


Figure 5.3 Selecting the mouth and chin

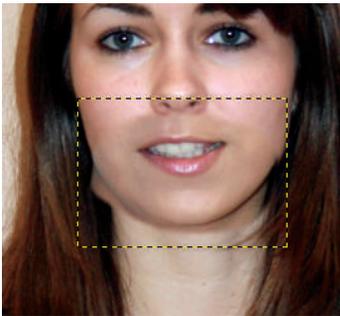


Figure 5.4 Pasting the mouth and chin

orientation, or size of one feature without changing the rest of the portrait.

Once you're satisfied with the placement, size, and coloration of the copied features, you can flatten the image and then do some minor retouching to better incorporate the new features into the face.

Zoom the initial portraits and your new portrait to the same zoom factor, for example 50%. Make sure portrait 4 is the active layer. Select the Free Select tool (**F**). Check the FEATHER EDGES box and choose a RADIUS of 10. Build a tight selection around the mouth and chin of portrait 4, as shown in Figure 5.3. If you copy too large an area, you might end up with discrepancies in skin color.

Copy the selection made in portrait 4, and paste it into the composite picture. Press **(SHIFT+CTRL+N)** to create a new layer with the floating selection. The copy appears in the center of the image; move the copy to its proper place with the Move tool. Because the girl in



Figure 5.5 Selecting the eyes

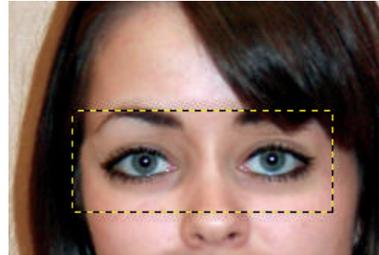


Figure 5.6 After pasting the eyes

portrait 4 tilted her head, the new mouth and chin are crooked. To correct this, flip the new layer horizontally with **Image: Layer > Transform > Flip Horizontally** and then rotate it (**(SHIFT+R)**) into the proper position. The mouth and chin are now aligned, but the color is uneven, as shown in Figure 5.4. We'll correct this later. Double-click the new layer in the Layers dialog and change the name to mouth.

Make sure portrait 1 is visible and active and then copy the eyes, again using the Free Select tool. To avoid selecting the eyebrows and the bridge of the nose, build the selection in two parts. First, draw a circle around one eye, press and hold the **(SHIFT)** key, and draw a second circle around the other eye. See Figure 5.5. You could also change the MODE in the tool's options, rather than pressing **(SHIFT)**.

Again, copy the selection, paste it into the composite picture, and create a new layer from the floating selection. Name this layer eyes. Move the layer into place, and rotate it slightly (**(SHIFT+R)**). The result, shown in Figure 5.6, isn't bad, although you could improve the skin color transition.

For the nose, use the Clone tool (**(C)**). Create a new transparent layer and call it nose. Select



Figure 5.7 Copying the nose

the Clone tool and select the Hardness 075 brush. Set the Clone tool's size option to 20, and select ALIGNED as your ALIGNMENT option. Using a tablet stylus if you have one, or the mouse if you don't, **[CTRL]**-click the tip of the nose in portrait 3 and then paint on the composite picture, beginning at the tip of the nose. The resulting image has two main problems: First, the skin color is inconsistent, and second, the tilt of the nose isn't quite right. The nose tilt can be corrected by rotating the new layer slightly. One possible result appears in Figure 5.7.

Smoothing Transitions

Next, smooth the features to create a more coherent composite portrait. Here, we demonstrate a few of the ways you can do this.

For the mouth and chin, use the Levels tool. Select it, and with the mouth layer active, move the right and middle triangles slightly for each of the three color channels. You should get a skin color similar to that of the original portrait. The area under the chin is still much too dark, however. Because the chin is on a separate layer, you can simply erase the part you don't need. Select the Eraser tool (**[SHIFT+E]**), and choose the brush and size that you used with the Clone tool. Zoom in to work more precisely, and if you have one, use a tablet pen. You'll still see a hue discrepancy between the chin and the neck, but don't worry about that.



Figure 5.8 The final composite picture

The main problem now is the color of the nose. Correct it using the Levels tool. Move the middle triangle in the Value channel liberally toward the left. Try somewhere near 1.8.

Now you've done all you can with a multilayer image. To complete the composite portrait, you need to flatten the image. As a precaution, save the multilayer image and then duplicate it (**[CTRL+D]**), and work on the copy. Merge all layers (**Layers: right-click > Flatten Image**), so you can use the Smudge, Dodge/Burn, or Clone tools on the image as a whole.

Several parts of the composite portrait require delicate touch-ups. Use the Smudge tool (**[S]**) to smooth the transitions between the various components of the final portrait, especially around the nose and on the upper part of the mouth and chin fragment copied from portrait 4. You can also use the Healing tool (**[H]**) to correct slight irregularities on the nose bridge or on the cheekbones.

Finally, select the whole face, but not the hair, and apply a very light Gaussian blur with a radius of 1 or 2 pixels. This makes the face smoother

and helps to hide any visible borders that remain between the features and the face. Figure 5.8 shows the final result.

5.2 Selections, Overlaying, and Blending Modes

This chapter focuses on the construction of composite images using elements from several different photographs. The three main tools for this process are selections, overlaying, and blending modes.

Masks and Selections

As the preceding tutorial demonstrates, making selections is one of the most important, and trickiest, tasks involved in photo manipulation. So we'll spend some time clarifying the functionality of the selection tools available in GIMP.

When you build a selection with the Rectangle Select tool, you get an area bounded by a simple geometric form. The simplicity makes it easy to tell whether pixels are within the selection. If you choose to feather the edges of the selection, however, the boundary is less clear: Some pixels are clearly out of the selection and some are deep within it, but what about those pixels on the feathered border?

The situation is even more complicated when you use the Fuzzy Select tool or the Select by Color tool. The selections that these tools build are usually not defined by clear and simple outlines. This shortcoming is obvious when you turn on Quick Mask, which displays the unselected pixels in red (by default). If the selection is feathered, you see that some pixels are completely red (unselected); others aren't red at all (selected); and some pixels are, well, reddish. Figure 5.9 shows a feathered selection made with the Select by Color tool, whereas Figure 5.10 shows the same selection with the Quick Mask tool active. Both images are shown with a zoom factor of 800%.



Figure 5.9 A feathered selection made by the Select by Color tool



Figure 5.10 The selection as a Quick Mask

As you may have guessed, those reddish pixels are partially selected. A selection and a mask are actually the same thing. They are both grayscale images that specify to what degree each pixel is selected. If the pixel value of the mask is 0 (i.e., black), the corresponding pixel in the image is not selected at all. If this value is 255 (i.e., white), the corresponding pixel is fully selected. A partially selected pixel has a value somewhere between 0 and 255.

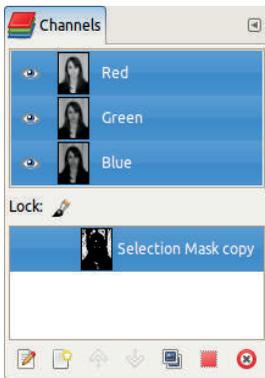


Figure 5.11 A selection saved in the Channels dialog

Along the feathered edge of a selection, pixels are partially selected. All actions made to the selection affect these partially selected pixels to a lesser degree: Painting is semitransparent; burning is partial; cutting leaves some of the pixels in place; and so on. Partial selections are what make selections so useful in photo manipulation. They allow you to build composite pictures that don't look like clumsy collages.

Because a selection and a mask are the same thing, and a mask is a grayscale image, you can manipulate it with the same tools that you would use on an image. You can paint on it with the painting tools demonstrated in Chapter 3. You can select and edit specific areas. And you can save it and use it later with **Image: Select > Save to Channel**. The mask appears as an additional channel in the Channels dialog, as shown in Figure 5.11. In this figure, the channel is selected (its line is emphasized), so if you paint on the image now, you'll be painting on the mask. To paint on the image itself, switch to the Layers tab, select the layer you want to modify, and begin painting.

You can also use this property to change the mask in the same way we use the Quick Mask tool. To do this, open the Channels dialog and click the box to the left of the mask so the eye is visible. The mask now appears over the image (in gray by default). You can change the mask's display color by selecting **Channels: right-click >**



Figure 5.12 The Edit Channel Attributes dialog

Edit Channel Attributes to open the dialog shown in Figure 5.12. You can change the channel's name and its opacity, as well as the color of the mask. Click the large button on the right to open the Color chooser, where you can pick any color you want. But keep in mind that using a color that contrasts with the image is generally best.

When a selection (a mask) is saved in a channel, you can use it again later by simply clicking the red square button at the bottom of the Channels dialog. If you hover the pointer over the button, you'll see that the standard key combinations allow you to activate different selection modes. You can save an unlimited number of selections as masks in the Channels dialog.

A mask used to determine the transparency of pixels in a layer is known as a *layer mask*. A layer mask is a part of a specific layer, and its pixels specify the transparency of the corresponding pixels in that layer. If a mask pixel is white, the corresponding layer pixel is opaque and thus visible in the image if not hidden by another pixel located in an upper layer. If a mask pixel is black, the corresponding layer pixel is transparent and thus invisible. Intermediate values for the mask pixels produce intermediate transparency levels for the layer pixels. Note, however, that the layer itself is unchanged.

You can add a layer mask by applying Add Layer Mask from the **Image: Layer > Mask** menu. Once added, you can access its options via the same menu. You can also add or manipulate a layer mask by right-clicking in the Layers dialog. When you add a mask to a layer, you can choose among several options: a completely



Figure 5.13 *The first photograph*



Figure 5.14 *The second photograph*

white or completely black mask, the current selection, a specified channel, and others. Once you've built the mask, you can change it by painting on it, for example, with Quick Mask. You can change some of its properties, such as how it displays on the image. Finally, you can apply it to the layer, making its effect final.

We already demonstrated some of these options in previous chapters, and we come back to them later in this chapter.

Overlaying Images

Simply overlaying opaque images isn't very interesting. Modifying the opacity of the top layer sometimes yields a satisfying result, but doing so doesn't give us much control over it. Once you build a selection, however, you can create your own unique and imaginative images.

For example, take the photographs shown in Figures 5.13 and 5.14. Open the first one, and add the second one as a layer (**CTRL+ALT+O**). Then, with the Fuzzy Select and Select by Color



Figure 5.15 *Selecting the sky*



Figure 5.16 *The final result*

tools, select the sky in the top layer, which requires about a dozen clicks in several areas of the sky. Recall that pressing **[SHIFT]** allows you to add to the selection, whereas pressing **[CTRL]** allows you to subtract from the selection. The completed selection is shown as a Quick Mask in Figure 5.15.

Our goal is to make the sky transparent so the building from the other photograph is visible behind the tropical vegetation. If we simply cut the selection, however, the sky is replaced by opaque white because the layer doesn't contain any transparency (indicated by the boldface layer name) and so any cut pixels are replaced by the background color. To add transparency, add an Alpha channel to the top layer: **Layers: right-click > Add Alpha Channel**. Now if you cut the selection, you get Figure 5.16.

You can also use a selection to copy something from one photograph and paste it into another one. You did this earlier on a smaller scale when you created the composite portrait at the beginning of the chapter. Moving larger



Figure 5.17 The first photograph



Figure 5.18 The second photograph

Credit: Vincent Lecarme



Figure 5.19 Pasting the man

objects from one photograph to another presents its own challenges, as you'll soon see.

As you saw previously, newly pasted objects are added to an image as a *floating selection*. They appear in the Layers dialog (which you should always keep visible in your workspace) and act somewhat like a layer located on the top of the stack. But as long as a floating selection exists, you can't change anything else in the image. The other layers are inactivated, just as the unselected regions of an image are inactive as long as a selection exists.

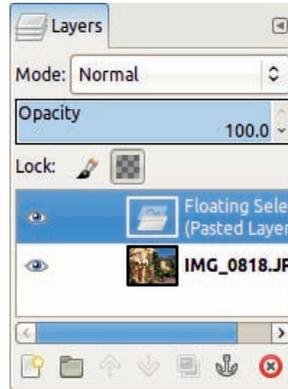


Figure 5.20 The new layer as a floating selection



Figure 5.21 The final result

With a floating selection, you can do the following:

- Create a new layer to contain the object you've just pasted.
- *Anchor* the pasted content to whatever was active just before you created the floating selection. This layer might be the active layer (highlighted in the Layers dialog) or the active layer mask (see the previous section), but not the active channel.

Let's say you want to combine elements from the photographs shown in Figures 5.17 and 5.18. You've decided to place an overdressed man in front of a sunny villa on the French Riviera. Select the man on the right, as shown in Figure 5.18. Make the selection with the Free Select tool and complete it by painting with Quick Mask. Don't very careful when selecting his legs, however, because they won't be visible.



Figure 5.22 The upper layer and the lower layer

Copy the selection from Figure 5.18 and paste it onto the image in Figure 5.17. The man appears in the center of the image, as shown in Figure 5.19, as a floating selection (see Figure 5.20).

Create a new layer for this floating selection and move the man to the bottom-right corner of the sunny scene. He's darker than his surroundings, so select the Levels tool and add some light and contrast. Also apply **Image: Layer > Transform > Flip Horizontally** to flip him around. You can see the result in Figure 5.21.

Using Blending Modes

GIMP has 21 blending modes, which you can select from the Layers dialog. The **MODE** option contains a drop-down menu with a list of blending modes that will act on the active layer. Each blending mode uses a different mathematical model to compute a new pixel value from the active layer and the layer below. Although we won't go into any detail regarding the mathematics of blending modes, we will introduce some of the modes most useful for overlaying images.

Begin with the photographs from Figure 5.22. Put them into the same image with the first one as the upper layer. The blending modes are arranged in a logical order, as explained in Chapter 12. In the following list, we briefly introduce

these modes, in a different order, so we can run through them more quickly.

- If the mode is set to Normal, the upper layer completely hides the lower one, except if its opacity is less than 100%. Figure 5.23 (left) shows an image with a top layer opacity of 50%. Each new pixel contains half of the upper pixel and half of the lower one.
- Dissolve mode looks the same as Normal mode until you reduce the opacity of the top layer. As you decrease the opacity of the top layer, the layers are blended using *dithering*, and the resulting pixels are a random mixture of the upper and lower pixels. Opacity corresponds to the likelihood that the new pixel value will be taken from the upper layer. Figure 5.23 (middle) shows the effect of Dissolve mode with the top layer at 50% opacity.
- In Multiply mode, the pixels' values are multiplied and then *normalized*. The resulting pixels are darker than those of the initial layers. Figure 5.23 (right) shows the result of Multiply mode: black pixels whenever one of the pixels is black and light pixels only when both pixels are light.
- Divide mode divides and normalizes the pixel values. The result depends largely on the upper pixel. If the upper pixel is black, the



Figure 5.23 Normal mode and Dissolve mode, 50% opacity; Multiply mode



Figure 5.24 Divide mode, Screen mode, and Overlay mode

result is almost white, as shown in Figure 5.24 (left).

- The mathematics of Screen mode are more complicated. As Figure 5.24 (middle) shows, the dark areas in the top layer are more transparent than the light ones. If you want to learn more about Screen mode, flip ahead to Chapter 12.
- Overlay mode combines Multiply and Screen modes and results in a washed-out upper layer. As Figure 5.24 (right) shows, the upper layer is no more than a ghost in the picture.

- Dodge and Burn are the traditional techniques used in analog photo development to make areas in a picture lighter or darker. As you can see in Figure 5.25 (left and middle), these modes operate according to the same principles: The upper layer is used to dodge or burn the lower one, and the intensity of the action corresponds to the value of the upper layer.

- Hard light and Soft light are similar combinations of Multiply and Screen. As shown in Figure 5.25 (right), Hard light mode isn't very



Figure 5.25 Dodge mode, Burn mode, and Hard light mode



Figure 5.26 Soft light mode, Grain extract mode, and Grain merge mode

useful in this example because the upper layer almost completely hides the lower one. As Figure 5.26 (left) shows, Soft light mode produces a more interesting result, although in this example, the result is very similar to Overlay mode (see Figure 5.24, right).

- Grain extract and Grain merge are another pair of sibling modes. Grain extract is similar to Difference mode, whereas Grain merge is similar to Addition mode (which we explain next). These modes are supposed to mimic film grain, but sometimes the similarity isn't

apparent. Figure 5.26 (middle) illustrates how Grain extract uses the light areas from the upper layer to exaggerate the light areas in the lower one. Figure 5.26 (right) shows the more subtle effects of Grain merge, which uses the texture of the upper layer to change the lower layer.

- Difference, Addition, and Subtract are also sibling modes. Difference mode, shown in Figure 5.27 (left), subtracts the lower pixel from the upper one and uses the absolute value of the result. The resulting image is



Figure 5.27 *Difference mode, Addition mode, and Subtract mode*

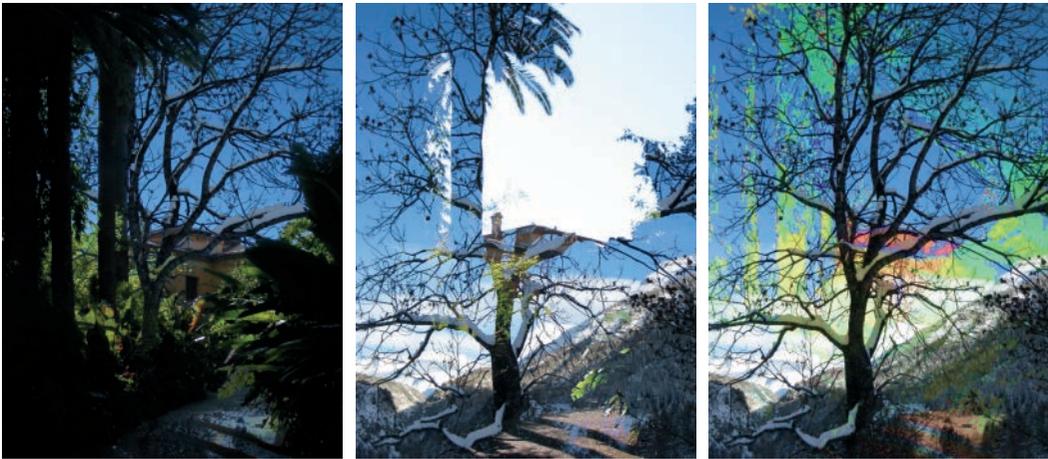


Figure 5.28 *Darken only mode, Lighten only mode, and Hue mode*

difficult to predict. Addition mode, shown in Figure 5.27 (middle), adds pixel values, and values greater than 255 appear white. The resulting image is much lighter. Subtract mode, shown in Figure 5.27 (right), subtracts the pixel values, and values that would be negative appear black.

- Darken only and Lighten only, shown in Figure 5.28 (left and middle), do what their names imply: The darkest or the lightest pixel is selected from the two layers.
- The last four modes are related: The new pixel gets some combination of HSV components from the initial layers. In Hue mode (Figure 5.28, right), the upper layer colorizes the lower one. Saturation mode (Figure 5.29, left) takes the hue and value components from the lower layer. Color mode (Figure 5.29, middle) takes only the value from the lower layer. Value mode (Figure 5.29, right) is the exact opposite: It takes only the value from the upper layer.



Figure 5.29 Saturation mode, Color mode, and Value mode

5.3 Building a Panorama

With a digital camera, you can easily take several pictures of the same subject and capture a view from multiple vantage points. You can also build *panoramic images* (i.e., wide pictures that cover a larger area than you can see at any one time).

Software tools that specialize in building panoramas are available. One of them is Hugin (see <http://hugin.sourceforge.net/>), which is free and works on GNU/Linux, Mac OS X, and Windows. It uses a powerful algorithm called *SIFT*, which was developed at the University of British Columbia. It builds the panorama in a completely automated way. The algorithm finds a collection of *control points* in the different photographs. The result is spectacular, but the algorithm is processor intensive, especially if the panorama is being built from a lot of large photographs.

We also use Pandora, a GIMP plug-in developed by Akkana Peck, at the end of this section. After you install it, Pandora is accessible in **Image: Filters > Combine**.

Taking the Pictures

Some digital cameras offer a panorama mode designed to ensure the images overlap of each picture.

sufficiently. Don't worry if your camera lacks that feature; taking the pictures you need for your panorama, using normal settings, is easy. Just follow this advice:

- Keep your feet firmly planted as you shoot. Rotate the top of your body to capture each successive shot, but try to keep the camera at the same height as you move across the landscape so each shot has the same amount of sky visible. Too much vertical motion can reduce your final panorama to a very narrow strip. If feasible, use a tripod for greater stability.
- Do not change the exposure as you move: The settings must be the same for all pictures, even if the lighting conditions are not the same. Choose the optimal settings for some part of the intended panorama (the center, for example), fix them, and then take all the pictures. If you're not sure how to do this, refer to your camera's manual.
- Take more pictures than necessary. To build a panorama, you need a lot of overlap. A camera's optical system generally distorts the corners, and sometimes the sides, of every image. Only the center of an image isn't distorted in any perceptible way. This is why you build your panorama using mainly the center of each picture.

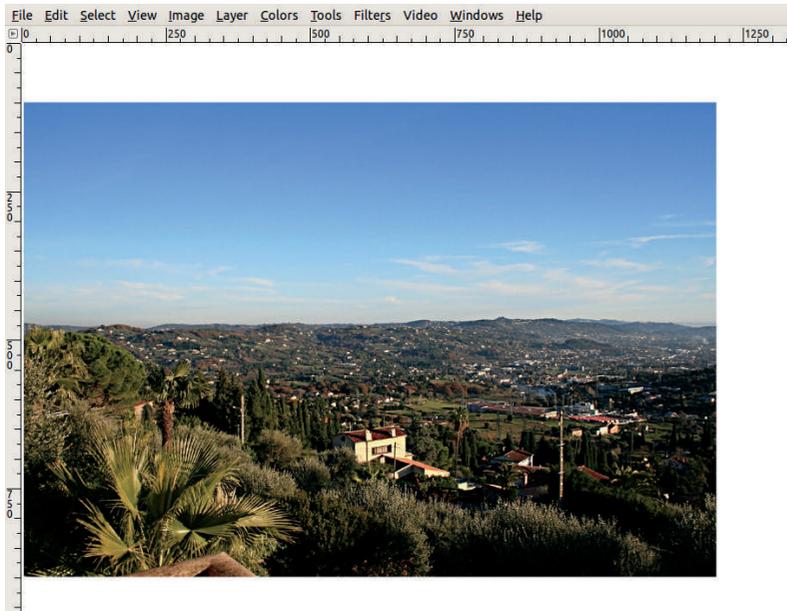


Figure 5.30 Adding the first photograph

Perhaps your camera can take photos with a really high pixel count. For example, Olivier’s camera has a maximum image size of 3888×2592 . When loaded in GIMP, an image of that size occupies 90.9MB. If you load six of them, probably the minimum number for a panorama, they occupy 545.4MB, which is too much for most computers to handle. And as soon as you begin working on an image, its size increases because of the stored information that allows you to undo. Basically, unless your computer is an extraordinary beast with a huge amount of memory, building an enormous panorama causes it to crash.

Generally, really large panoramas are impractical anyway. A panorama intended for a web page shouldn’t be much wider than 1000 pixels. If you want to print the image, a width of 4000 pixels results in a print size of 13.3 inches or 33.9 cm when printed at 300 ppi.

You can take the pictures for your panorama using the highest setting available on your camera and then scale them down to a usable

size. Be careful not to scale down too much, however, or you’ll lose precision and sharpness.

Superimposing the Pictures

Olivier took eight photographs of the view from his balcony, following the advice in the previous section. The images are 3888×2592 , but we scaled them down to 1200×800 . When scaling an image, choose the Sinc (Lanczos) algorithm, which, although somewhat processor intensive, does the best job. After scaling down, we applied the **Image: Filters > Enhance > Unsharp Mask** filter with its default parameters to compensate for the loss of sharpness caused by scaling.

Because these photographs overlapped a lot, as they should, the complete panorama will not be more than 2800 pixels wide. As keeping the camera at exactly the same height for all shots was difficult, we created a new image that was 2800×1000 , which happened to give us some wiggle room. This image isn’t extremely large, but it will have to be scaled down for a web page. On the other hand, this image is rather small

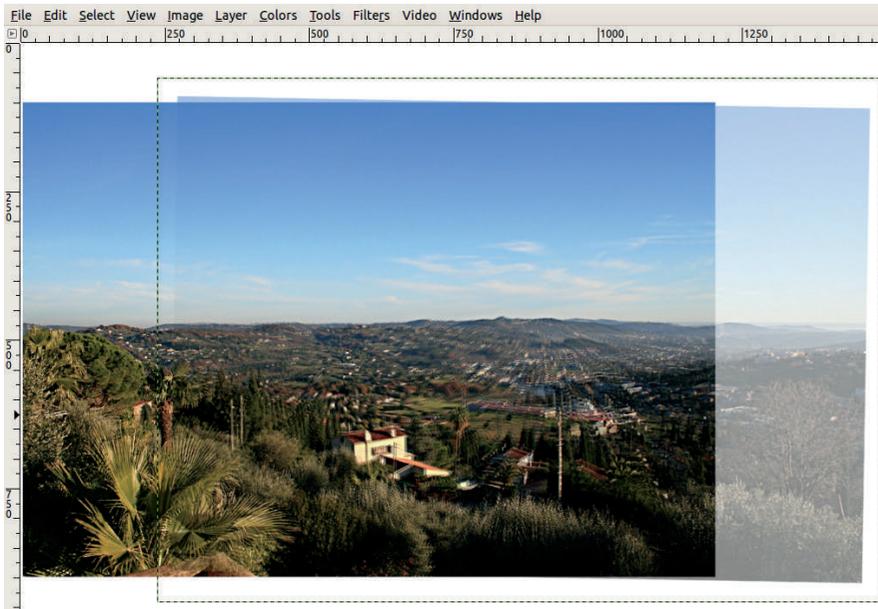


Figure 5.31 Loading the second photograph

for printing. Here's how we would build the panorama.

1. Load the first photograph into this new image with **[CTRL+ALT+O]**. The new layer appears in the center of the image; move it completely to the left, as shown in Figure 5.30.
2. Load the second photograph as a new layer and change its opacity to 50% so you can see whether it's positioned correctly. Move it so the overlapping areas of the two photographs align. Focus on superimposing the house in the middle so it's as sharp as possible. Use the arrow keys to tap it into place. As Figure 5.31 shows, a perfect superimposition is impossible because of the lens distortion. Here, we are able to get the house in sharp focus, but the surrounding hills are blurry and the electric pole on the left is doubled. We can tell from the horizon that the camera was slightly tilted when the second picture was taken.
3. To improve the transition from one photo to the other, add a white layer mask to the top layer. Select the Blend tool and choose a linear gradient from black to white without repetition. You want the top layer to transition from transparent on the left to opaque on the right. To create a perfectly horizontal gradient, press **[CTRL]** while dragging from left to right, beginning somewhat to the right of the left edge and ending at the house.
4. Once the gradient is built, click the layer thumbnail to select the layer itself. Set the opacity of this layer back to 100% and move it very slightly (with the arrow keys) so the house is as sharp as possible. You also need to rotate the top layer a tiny bit to compensate for the camera's tilt. This is tricky because partial transparency disappears while you are using the rotation tool. Figure 5.32 shows the result so far as well as the Layers dialog.

Touching Up the Panorama

Repeat this process for the remaining photographs: Load the photograph, set transparency to 50%, position it, remove the

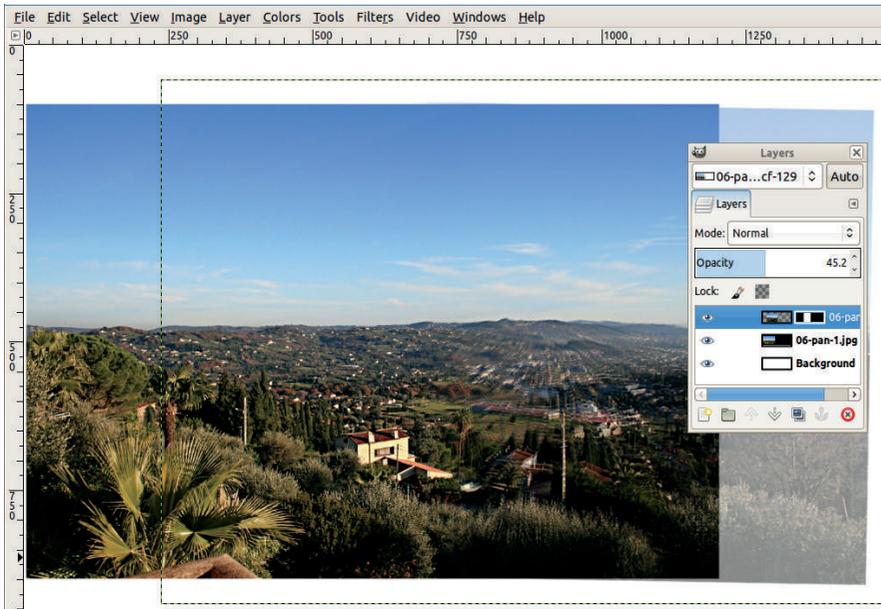


Figure 5.32 Positioning the second photograph



Figure 5.33 All the photographs are positioned.

transparency, add a layer mask, draw a gradient on this mask, and finish positioning the photograph. To avoid creating weird artifacts in the scenery, be careful not to make the gradient too wide and avoid placing it in an area with a lot of fine detail. As you add layers, the size of the image increases and, in this case, is 123.4MB after we've added all eight photographs.

Sometimes, if a detail from the lower layer appears out of place, changing the gradient is helpful. You can adjust the gradient by painting on

the layer mask, or you can add another gradient over the problem area—above the existing gradient. Set the second gradient to Multiply mode (from the MODE menu in the tool options dialog). To see the layer mask itself, **[ALT]**-click its thumbnail or right-click in the Layers dialog.

The current result appears in Figure 5.33. The color of the sky is uneven at a couple of the junctions between photographs. Because the sky contains clouds and a light haze on the extreme right, replacing it completely with some



Figure 5.34 *The final panorama*

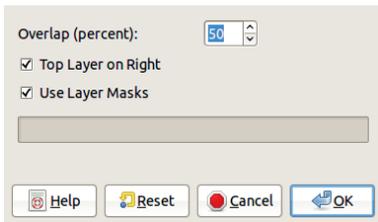


Figure 5.35 *The Pandora filter dialog*

gradient between two different shades of blue would detract from the image. We prefer to leave the panorama as is, although some careful (and tedious) work with the Dodge/Burn tool (in Dodge mode) or the Clone tool could improve the result.

With all the pictures in place, all that's left is to crop and flatten the image. Our final panorama appears in Figure 5.34.

The Pandora Plug-in

The Pandora plug-in, mentioned at the beginning of this section, can also be used to build a panorama. Here's how you use it:

1. Load all the photographs as layers in a new image. Do this quickly using `CTRL+ALT+O`.
2. Select the **Image: Filters > Combine > Spread out layers** filter. In the dialog that appears (Figure 5.35), select the **OVERLAP** quantity between images (in this case, 80 is better than 50) and decide whether the top layer is on the



Figure 5.36 *The original image*

left or right. Note that the last image loaded becomes the top layer.

3. The finishing touches are left up to you. Move the layers and fine-tune the layer masks to get a smooth panorama.

5.4 Image Overlaying

In the previous section, we considered a very specific application of image overlaying. Now, we show you techniques for more general use.

Digital Collage

Begin with the photograph shown in Figure 5.36. We want to replace the seascape, seen through the window, with Figure 5.37. To do this, you need to make the seascape transparent. You could try selecting it with the Select by Color tool or the Fuzzy Select tool, but this time try a brand new method instead.



Figure 5.37 A new background

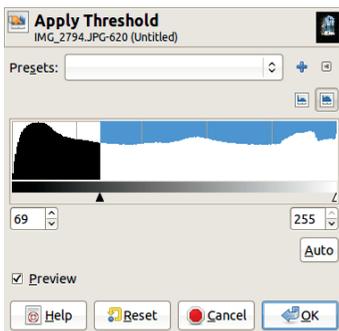


Figure 5.38 The Apply Threshold dialog

Duplicate the layer, and choose the **Image: Colors > Threshold** tool, which opens the dialog shown in Figure 5.38. Move the black triangle to the left until you see a silhouette of the window only. Invert it with **Image: Colors > Invert** to get the result in Figure 5.39. Use this as a mask to make the window fully transparent.

Now do the following:

1. Select the top layer and copy it.
2. Open the image shown in Figure 5.37 as a new layer by selecting **Image: File > Open as Layers** or pressing **[CTRL+ALT+O]**.
3. Duplicate the bottom layer, and move it to the top of the layer stack.
4. Add to it a layer mask (**Layers: right-click > Add Layer Mask**).
5. Paste the copy made earlier, and anchor it to the layer mask (**[CTRL+V]** and **[CTRL+H]**).
6. Arrange the layers as shown in Figure 5.40.



Figure 5.39 Threshold result after inversion

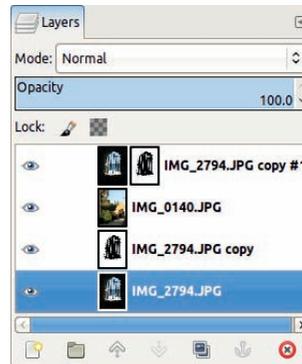


Figure 5.40 The Layers dialog for Figure 5.41



Figure 5.41 With the new background added

The resulting image appears in Figure 5.41. You could stop here. But the wall is so dark that its stones are almost invisible. Fortunately you can adjust this without altering the landscape seen through the window: In the Layers dialog, click the thumbnail in the top layer to select the layer and not the layer mask. Use the Levels tool (**Image: Colors > Levels**) to adjust the Gamma triangle until you're satisfied. Only the layer is affected, not the layer mask or the other layers in the stack. If, on the other hand, you want



Figure 5.42 *The final result*



Figure 5.43 *The frame*

to adjust the layer mask, you could select its thumbnail and then use the Brush tool to add to or remove from the mask.

The final image appears in Figure 5.42. Note the initial image is still present and unchanged in the bottom layer.

Next, use the same principles to build an image that's slightly more complicated. The photograph shown in Figure 5.43 will serve as a frame for the portrait shown in Figure 5.44. You want this portrait to appear in the open French windows, larger than life, but with the landscape still visible behind her. You first need to remove the beige background that's currently behind the subject, which is easy enough. Once you've done that, however, you can't simply copy the portrait and paste it in because her



Figure 5.44 *The portrait*

shoulders would extend beyond the window frame.

To position the portrait properly in its unconventional frame, follow these steps:

1. Open the image shown in Figure 5.43. Add an Alpha channel (**Image: Layer > Transparency > Add an Alpha Channel** or **Layers: right-click > Add an Alpha Channel**).
2. Using the Free Select tool, make a selection by clicking the four corners of the French windows.
3. When the selection is complete, cut it and then paste it back into the image as a floating selection. Create a new layer from the floating selection (click the leftmost button at the bottom of the Layers dialog or press **[SHIFT+CTRL+N]**). Move the new layer to the bottom of the layer stack. The image looks the same, but it's divided into two different layers.
4. Open the photograph shown in Figure 5.44 as a new layer. Add an Alpha channel to this layer. Select the Fuzzy Select tool and select the beige background. Decrease the **THRESHOLD** in the tool options, and as usual, press

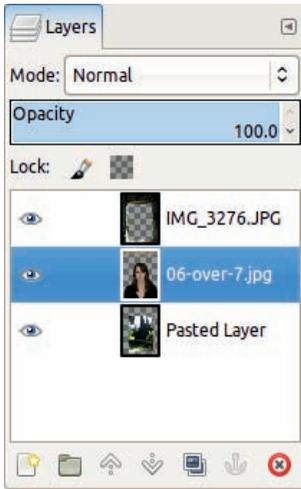


Figure 5.45 The Layers dialog for Figure 5.46



Figure 5.46 The final result

- [SHIFT] to add to, or [CTRL] to remove from, the selection as you build. Take care when selecting the area around her hair. When you're finished, cut the selection ([CTRL+X]) and discard it ([SHIFT+CTRL+A]).
5. Select the Scale tool ([SHIFT+T]), choose NUMBER OF LINES in the GUIDES options, and click in the portrait layer. Enlarge the portrait, and move it into the window opening.



Credit: Vincent Lecarme

Figure 5.47 The initial photograph



Figure 5.48 Self-compositing in Hard light mode



Figure 5.49 Duplicating the top layer twice more

6. Finally, arrange the layers so the French window is on top, the portrait layer is in the middle, and the cutout landscape is at the bottom, as shown in Figure 5.45.

The final result should resemble Figure 5.46.

Self-Compositing

Here, begin with a photograph of a winter scene, shown in Figure 5.47, that is almost uniformly gray and flat. To transform this picture into the image shown in Figure 5.48, simply duplicate the



Figure 5.50 *The initial photograph*



Figure 5.51 *Self-compositing in Screen mode*

layer and change the mode of the top layer to Hard light. By repeating this process, you can increase the contrast further, but this adds a strong blue cast to the image, as shown in Figure 5.49. Correct this effect using the Levels tool.

The photograph shown in Figure 5.50 (taken from an airplane) is much too dark. To lighten it, begin by duplicating the layer, as you did in the previous example, but this time use Screen mode.

Figure 5.51 shows the result after two duplications. Now the color and detail in the foreground are visible, but some details in the lighter areas of the background have been lost.

To bring back some of that background detail, select the light areas in the image (i.e., the sky and the clouds) with the Fuzzy Select tool, but this time, forget about pressing **[SHIFT]** or **[CTRL]**. Instead, click somewhere in the sky, and then drag the cursor down or to the right. As you do, the tool threshold increases, which increases the selected area. If you drag the cursor



Figure 5.52 *Selecting the light areas*



Figure 5.53 *Adding a layer mask to shield the lighter parts*

up or to the left, the selected area decreases. The result appears in Figure 5.52.

Invert the selection (**[CTRL+I]**), add a layer mask to the top layer, and in the dialog that appears, choose to create the mask from the current selection. Do this for the top layer only, however, to avoid creating an unnatural effect on the lighter areas of the mountains. Figure 5.53 shows the result. The image could be improved more by adjusting the Gamma pointer in the Levels tool and by removing a few minor, age-related imperfections from the photograph, as you saw in Chapter 2.

Self-compositing photographs can also be used to create some artistic, filter-like effects. Open the photograph shown in Figure 5.54 and duplicate the layer. If you set the top layer's mode to Difference, the resulting image will be completely black because the difference between two equal quantities is always zero. But if you use the Move tool to move the top layer a few pixels up and to the left, you get the



Figure 5.54 *The initial photograph*



Figure 5.57 *Using Divide mode*



Figure 5.55 *Using Difference mode*

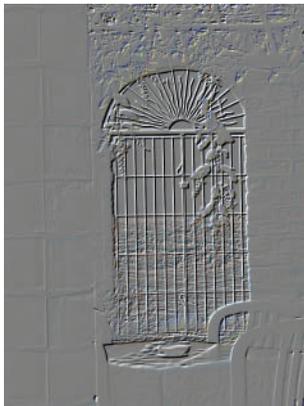


Figure 5.56 *Using Grain extract mode*



Figure 5.58 *The initial photograph*

image shown in Figure 5.55, which is a bit more interesting.

Change the top layer's mode to Grain extract to get Figure 5.56. Then change it to Divide mode to get Figure 5.57. You can create variations on these effects by changing the amount and direction of offset or the top layer's opacity or by adding new layers in different modes, and so on.

A Composite Photography Project

The photograph shown in Figure 5.58 was taken early in the morning when clouds partly filled in the valleys in front of Olivier's house. He planned to use that image in a composite photography project, but he snapped the photo in a rush and the image doesn't truly capture what he saw.

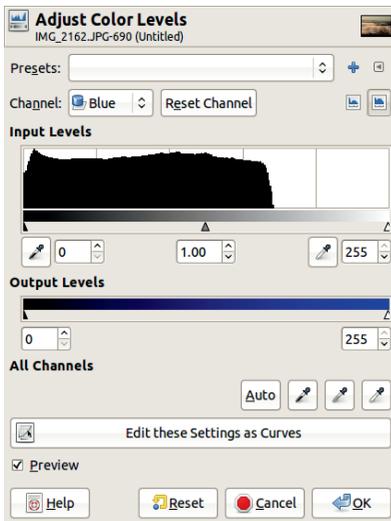


Figure 5.59 The Blue channel's histogram



Figure 5.60 After extending the levels

If you call the Levels tool and look at the color histograms, you see the Green and Blue channels are truncated on the right. The Blue channel, which appears in Figure 5.59, is the worst. Simply clicking the **AUTO** button corrects this, and you get the result shown in Figure 5.60. This landscape is now very much as Olivier remembers, so he decided to use it in the following creative project, shown in Figure 5.61.

Here's how to create this image:

1. Add the image shown in Figure 5.36 on page 131 to Olivier's landscape as a new layer.
2. Scale this layer to the same width as the underlying image (**Image: Layer > Scale Layer**).
3. Set the top layer to Darken only mode.
4. Crop the image.



Figure 5.61 Compositing with another image in Darken only mode



Figure 5.62 Another image to be composited



Figure 5.63 Compositing with another image in Burn mode

Similarly, add the photograph shown in Figure 5.62 to the image as another layer, and select **Burn** mode. You'll get the image shown in Figure 5.63.

In the following examples, we duplicated an image layer and then coupled various blending modes with a Gaussian blur of radius 10 that we applied to the top layer. Figure 5.64 shows the result when the top layer is set to **Divide** mode.



Figure 5.64 A blurred layer in Divide mode



Figure 5.65 A blurred layer in Difference mode



Figure 5.66 The top layer in Burn mode

To create Figure 5.65, we chose Difference mode, and we also adjusted the Levels by moving the Gamma triangle to the left.

In Figure 5.66, we removed the Gaussian blur and set the layer to Burn mode. Note how this enhances the clouds in the foreground.

You can add depth by moving the top layer 10 pixels up and to the left. The offset images in Figures 5.67 and 5.68 were cropped to illustrate the effect.



Figure 5.67 An offset layer in Divide mode



Figure 5.68 An offset layer in Difference mode

5.5 Exercises

Exercise 5.1. Finding a single set of images that can be used to accurately demonstrate all blending modes is a difficult task. The choice made in “Using Blending Modes” on page 123 is not ideal. Find a better set of images for this demonstration.

Exercise 5.2. Creating a uniform sky is often the most difficult part of panorama construction. Explore solutions to the problems seen in Figure 5.34 on page 131. Adjust the layer mask before merging the layers or touch up the final result using the Smudge tool, for example.

Exercise 5.3. The photograph shown in Figure 5.43 on page 133 contains two potted plants with large leaves, which might look nice in front of the subject instead of behind her. Devise a method to do this.

Exercise 5.4. Using the concepts demonstrated in “Self-Compositing” on page 134, create a unique work of art with your own digital photos.