From www.studio.adobe.com

Advanced pixel blending

Blending pixel layers in Adobe® Photoshop® CS2 lets you do things that you simply can't do by adjusting a single image. One situation where we blend pixel layers is when we want to extend the apparent dynamic range of digital captures. Of course, the best way to do this is to take several bracketed exposures and merge them to a 32-bit floating point HDR (High Dynamic Range) image. But that requires planning, a hefty tripod, and a subject that moves little if at all.

It may seem that it's pointless to blend different renderings of a single digital capture. After all, all the data is in the capture, so why not simply edit it to get the results you want? We feel that there are two reasons:

- It's extremely difficult to get everything right from a single rendering of the image.
- Most cameras capture 12 bits of data. By blending multiple renderings of the 12-bit capture, you can populate more levels in 16-bit space than you can by stretching and squeezing the original 12 bits.

The image below shows an "impossible" image produced just this way.

The sun is just outside the frame at the top of the image—we wouldn't even have attempted this with film!

We chose this completely over-the-top example to make it clear just how far you can take this approach, but the techniques are equally applicable to interior shots with bright window light, or to any other kind of wide-dynamic-range scene.

We started by creating three different renderings of the same camera raw image, one for highlights, one for midtones, one for shadows, using the Camera Raw plug-in.

The image, rendered for shadows from Camera Raw.

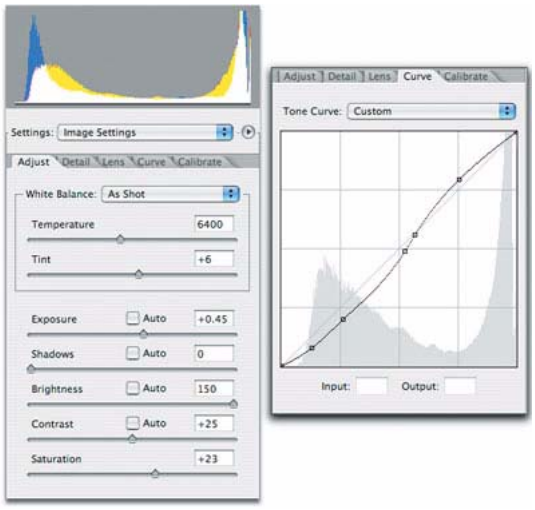

The Camera Raw settings and the Curve tab with the image rendered for shadows.

The image, rendered for midtones from Camera Raw.

The image, rendered for highlights from Camera Raw.

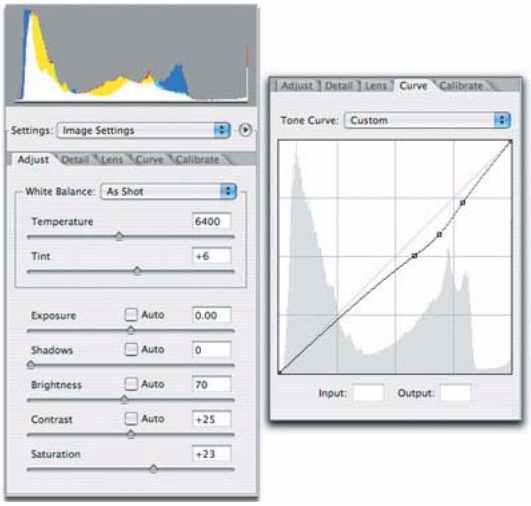

The Camera Raw settings and the Curve tab with the image rendered for midtones.

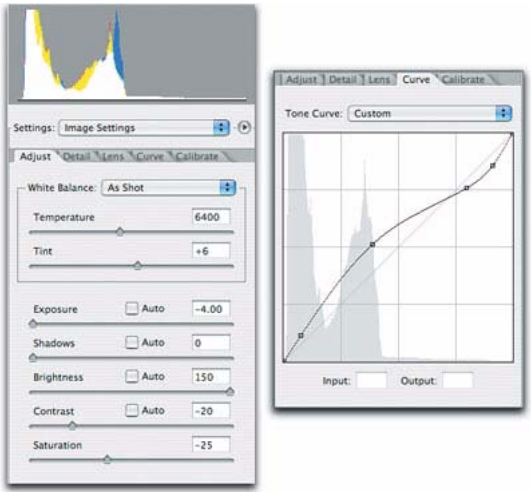

The Camera Raw settings and the Curve tab with the image rendered for highlights.

In many cases, you can get away with only two renderings, but this is an extreme example! Our first rendering is for the shadows, using significant positive Exposure and Brightness adjustments, with fine-tuning applied using Camera Raw's Curve tab. The second rendering concentrates on the sky and the distant hills. The third rendering makes use of Camera Raw's extended highlight recovery, with a major negative Exposure correction further aided by the Curve tab to squeeze every last drop of highlight detail out of the image.

Combining the pixel layers

We always start with the lightest layer and add the darker ones, because it makes the masking much easier. Why? Because a very good starting point is to load the luminosity of the lightest layer as a selection, and use it as the layer mask for the darker layer. That way, the shadow detail in the lightest layer gets preserved, while the blown-out areas are replaced by the darker layer.

We start by opening the first rendering in Photoshop. Then we open the second rendering, select all, copy, and paste it into the first rendering, where it appears as a new layer.

The image before layer masking.

Loading the luminosity as a selection

Our next step is to mask the pasted layer. As we just noted, the best starting point for the mask is to load the Luminosity of the Background layer as a selection, using either of the following methods:

- Hide all layers except the Background and then Command-click on the RGB channel tile in the Channels palette.
- Hide all layers except the Background and then press Command-Option-~(tilde). Mac users take note that OS 10.4 and later hijacks this useful keyboard shortcut to "Move Focus to Window Drawer," but you can reassign it or turn it off in System Preferences.

A key point in either case is that the selection that gets loaded is the combined luminosity of all visible layers—that's why you have to "solo" the background layer, which you can do quickly by Option-clicking the Background layer's eyeball icon in the layers palette. The image below shows the layers before and after masking.

Left: *Option-click to "solo" the Background layer, then*
Center: *Command+click the RGB tile or press Command+Option+~(tilde).*
Right: *Click the Add layer mask icon to load the selection as a layer mask.*

The image after layer masking.

With less-challenging images, two layers masked this way may be all you need, but in this case, we need to add a third layer for the highlight detail. The procedure for adding the third layer and masking it is almost identical. We open the third rendering of the image, the one adjusted for highlight detail, Select All, Copy, and paste it into our layered document, where it becomes the third, topmost layer. (That's assuming that the top layer was selected when you pasted—if it wasn't, and the new layer comes in between the two existing ones, just drag it to the top of the layer stack.)

To make the mask for the third layer, click its eyeball icon in the layers palette, then load the combined Luminosity of the first two layers either by Command-clicking the RGB Channel thumbnail in the Channels palette, or by pressing Command-Option-~(tilde)—remember, the selection that gets loaded is the luminosity of all visible layers. Then, select the third, topmost layer, click the eyeball to make it visible, and click the Add layer mask icon. The selection is added as a layer mask to the third (highlight) layer.

Adding and masking the third layer: Before layer masking.

***Left:** Click to hide the new layer, then*
***Center:** Command-click the RGB tile or press Command+Option+~(tilde).*
***Right:** Click the Add layer mask icon to load the selection as a layer mask.*

Adding and masking the third layer: After layer masking.

The resulting image is pretty flat, but all the essential components are now in place. We can improve matters somewhat by tweaking the layer opacities.

Adjusting the layer opacities as shown in the Layers palettes.

The adjusted result.

Producing the final image

To get from here to the result shown back at the beginning of this tutorial, we used masked layers set to Multiply to darken, Screen to lighten, and Soft Light to add contrast. We left all the layers at 100 percent opacity and used the layer masks to control them.

On this image, the process was iterative, bouncing back and forth between the layers, and in all honesty it's unlikely that we'd be able to produce absolutely identical results twice in a row, though we'd come pretty close. The image below shows the almost-final image, with layers applied to add contrast (Soft Light), lighten (Screen), and darken (Multiply).

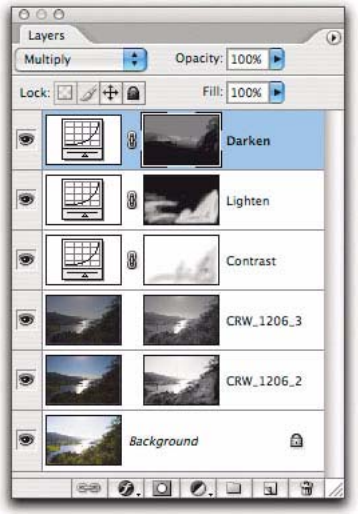

Layers set to Soft Light, Screen, and Multiply.

The adjusted image.

We started by adding a layer set to Soft Light to increase contrast. We left the layer mask white while we added two more layers, one set to Screen, the second to Multiply. On those layers, we inverted the layer mask (target the mask and press Command-I) to hide the effect, then we painted the lightening and darkening into the image by painting on the masks with a soft-edged brush at opacities between 10 and 20 percent. Then we returned to the Soft Light layer, and painted out some of the contrast in the foreground trees.

While we didn't in this case resort to doing so, we'd be remiss if we failed to mention that it's also possible to edit the layer masks using Levels or Curves. To lessen the strength of a layer, you can use the white Output Levels slider in Levels to turn the white (fully revealing) areas of the mask to a light gray, or to let just a little of the effect show through fully masked areas, you can use the black Output Levels slider to turn the black (fully concealing) areas of the mask to a dark gray.

We finished off the image by adding a fully masked Soft Light layer, then painting the additional contrast in with a 5-percent-opacity brush to get the final result shown in the initial image at the start of the tutorial.

Alternative workflows

In this example, we created three separate DNG files by saving them from Camera Raw with their respective settings, then opened them in Photoshop and copied-and-pasted to get them all into the same document. But there other useful ways to handle the task of combining multiple renderings of the same raw image.

- **Edit and Open in Camera Raw without saving settings**. If the edits needed to produce the different renderings are simple enough, it's probably not worth saving them as separate settings. Instead, you can open the raw file in Camera Raw hosted by Photoshop (this doesn't work when Camera Raw is hosted by Bridge), and Option-click the Open button (it changes to Open a Copy when Option is pressed) to open a copy of the image without updating the Camera Raw settings.

This technique is useful when you want to produce a few different renderings quickly without altering the "master" settings for a raw file. Once the images are open in Photoshop, the techniques for combining them are the same as in the example we've just covered.

- **Place Smart Objects**. Another alternative is to place the raw file multiple times into a Photoshop document as Smart Objects. That way, you can edit the settings for each rendering by simply double-clicking the layer thumbnail for the raw-place-as-Smart Object—it opens the image up in Camera Raw and lets you edit the settings, which are applied only to that specific Smart Object.

Our enthusiasm for this workflow is tempered by two considerations, though. You have to create a Photoshop document with the correct dimensions before you start placing the Smart Objects, and while you can edit the images in place, you don't see them in the context of the other layers. So it makes changing the individual renderings a little quicker than going back to the raw file, editing it, and opening a new version, but it doesn't make getting the settings right any easier.

One of the most appealing aspects of layer-based editing is its nondestructive nature—the edits aren't committed until you flatten the file—but sometimes, you can take nondestructive editing so far that you create extremely complex files that are both large and hard to understand. So don't be overly afraid to mix a little destructive editing in with the nondestructive stuff. Remember that you can edit pixel layers directly, and sometimes, it's a good idea to do so. If you're nervous about making edits without an escape route, there's one more Photoshop feature that provides a handy fallback position—the History palette.

Excerpted from "Real World Adobe Photoshop CS2" by Bruce Fraser and David Blatner © 2005. Used with the permission of Pearson Education, Inc. and Peachpit. To purchase this book, please visit www.adobepress.com.

Adobe, the Adobe logo, and Photoshop are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States and/or other countries. Mac is a trademark of Apple Computer, Inc., registered in the United States and other countries.
© 2006 Adobe Systems Incorporated. All rights reserved.