
Color Balancing Techniques

Written by Jonathan Sachs
Copyright © 1996-2007 Digital Light & Color

Introduction



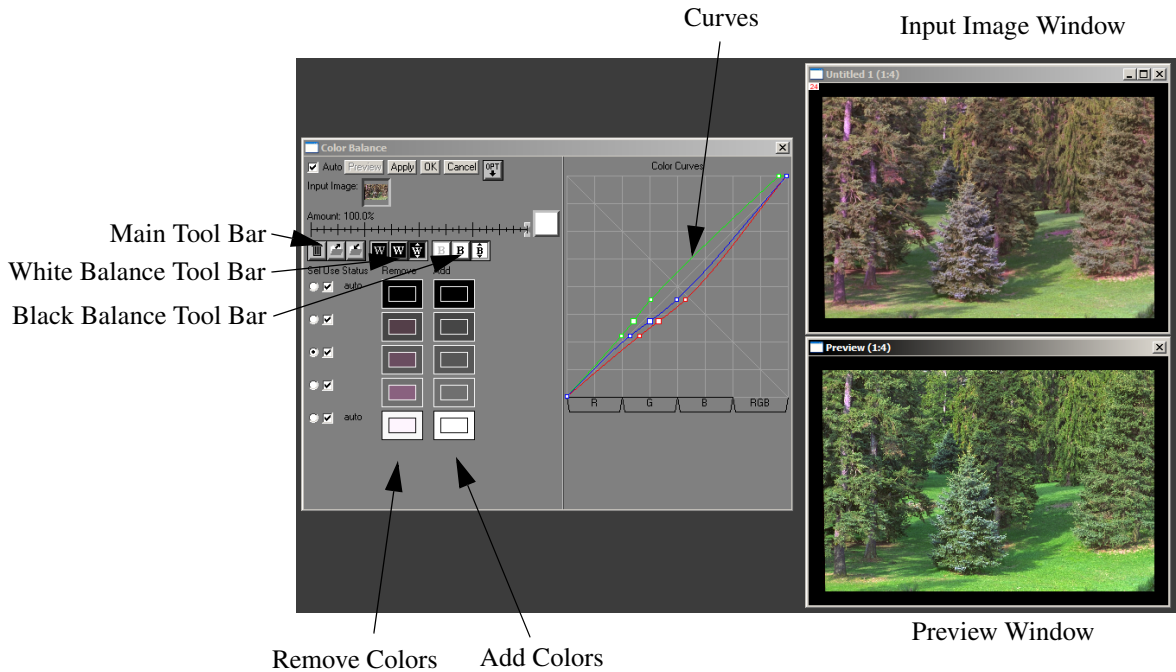
Color balancing refers to the process of removing an overall color bias from an image. For example, if an image appears too red, it is said to have a red cast. Removing this red cast brings the image back into balance. Color casts can arise from many different causes: the way the original scene was illuminated, filters you may have used, and how the image was captured and processed. Since it is difficult or impossible to control all the factors that can create a color imbalance, it is often easiest to correct the problem at the end of the process.

The Color Balance Transformation

A perfectly color balanced image renders the full range of neutral grays, including black and white, without any tinge of color. Sometimes this may be effect you are trying to create, and sometimes it may not. For example, photographers often take pictures in early morning or evening light because of the dramatic shadows and the yellow-orange color of the light that imparts a warm glow to everything it illuminates—even white objects.

The Color Balance transformation is based on removing and adding colors to an image. To remove a color cast, you specify the color of the cast you want to remove as the Remove color. To add a color cast, you specify the color of the cast you want to add as the Add color. The net result is to replace occurrences of the Remove color in the input image with the corresponding Add color to create the result image.

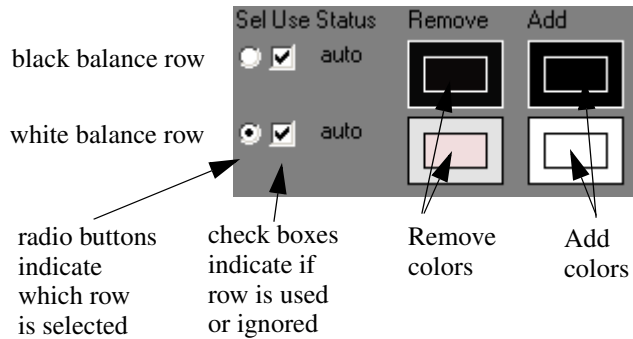

For an overview of the layout of the Color Balance transformation dialog box, see below:



This section details a series of operations you can perform using the Color Balance Transformation starting the simplest automatic procedures and proceeding by stages to increasingly advanced techniques.

Automatic White and Black Color Balancing

When the Color Balance transformation starts up, it examines the entire input image to determine the average color of its brightest and darkest parts. If you select automatic white or black balancing (see below), the white or black Remove colors are set to the values Picture Window found in this initial analysis.



White or black auto balancing without range extension sets the Add color to a neutral version of the Remove color based on its luminance (apparent brightness). If you select range extension, the Add colors are set to pure black or pure white. This has the effect of expanding the dynamic range of the result image by making highlights as bright as possible and shadows as dark as possible. Without range extension, the overall brightness and contrast of the input image is preserved. Whether range extension is a desirable or not depends on the image and what effect you are trying to achieve.

Below is an example image corrected using auto white and black balancing with and without range extension:



Original Image



Auto White& Black



Auto White& Black with Range Extension


Balancing by probing the input image

Sometimes Picture Window's algorithm for automatic color balancing locks onto the right color casts, in which case you're done. However if for some reason the automatic black and or white colors don't achieve the desired effect, you may have to set the Remove colors manually. This can easily occur if there is no pure black or pure white object in the original scene. Or, even if there are pure highlights and shadows, you may not want the color casts completely removed.

The easiest way to select Remove colors manually is by probing the input image (i.e. by clicking on a selected point in the input image window). This method

depends on being able to find something in the input image you know is neutral. For portraits, you can often use the whites of the eyes, white walls or furniture, or articles of white clothing. For landscapes, clouds, white flowers, waterfalls, or tree trunks are often a neutral white or light gray.

First click on the radio button to select the white balance row if it is not already selected.



Then go to the input image window and click on a lighter part of the image that you wish to appear neutral in the result image. This sets the Remove color to the color you clicked on and sets the Add color to a neutral version of the Remove color of the same luminance. If you want range extension or if you want to add a color cast back into the image, click on the white balance Add color (see above). This brings up a color picker—adjust the color to change the color cast or adjust the brightness slider to change the dynamic range.

Similarly, you can override automatic black balance by clicking on the radio button for the black balance row and then clicking on a darker part of the input image.

When black and white are not enough

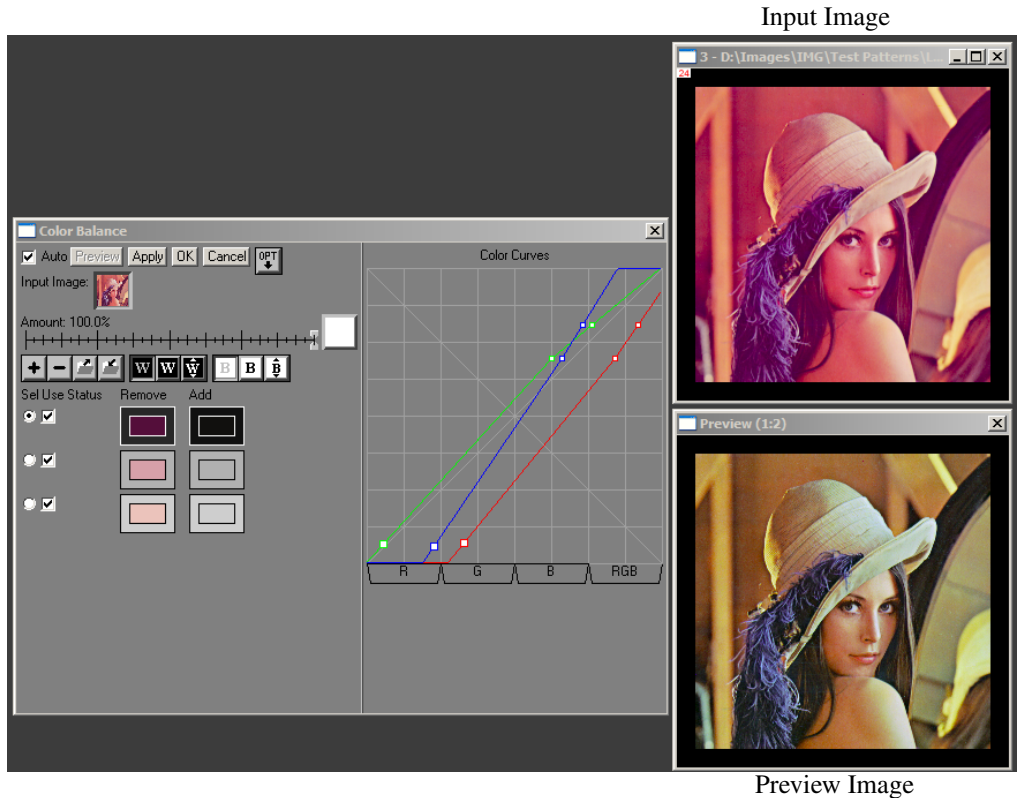
There are times when simple white and/or black balancing is too simple to get the job done. For example, consider overexposed parts of an image such as specular reflections (reflections of the sun from shiny objects) or other burned out spots in the original scene. Regardless of any color cast in the overall image, these highlights look white or nearly white because one or more of the channels are clipped. Automatic white balancing may lock onto the white highlights and not remove the main color cast, while removing the main color cast will add a color cast to the highlights. Dealing with this problem requires removing different colors from the highlights than from the rest of the image.

To deal with this and other situations, the Color Balance transformation lets you specify as many as seven different sets of colors to remove and add to the input image. The easiest way to add a new row of color patches is the Shift-Click on an

area of the input image. This creates a new row and sets the Remove color to the probed color and the Add color to a neutral version of the Remove color with the same luminance



The highlights in the input image (top left) are white but there is an overall purple cast to the image due to the artificial lighting. Adding an extra row (the middle row) by probing the light midtones on one of the white uniforms fixes the problem (see Preview window top right) without adding a cast to the specular reflections.





In this example, neither the auto white nor the auto black produces a good result. The two lighter Remove colors were obtained by probing the whites of the two eyes, one of which is a little darker than the other. The darker Remove color came from the mirror frame. The corresponding Add color was entered manually as a dark brown using the color picker.

If necessary you can add more rows by probing (shift-clicking on) other parts of the input image, but three rows are sufficient most of the time.

Matching colors between two images

You can use Color Balance to match colors between two images. Whenever you remove one color and add another, any occurrences of the Remove color are changed to the Add color. Thus to get white (or any other color) from one image to match another image, you simply remove the color from the input image and add the color you want to match from the other image.

Applying consistent color balance changes to several images

You can save the current Color Balance settings in a file and then reload them later for use with another image. This is done via the   buttons on the main tool bar.

Changing the default settings


If you save the current settings in a file with the name *default.cb* located in the same folder in which Picture Window is installed, then those settings become the default settings for new sessions. If *Remember Settings* under *File/Preferences* is set to *Yes*, then each time within a session that you use Color Balance, it will remember and start up with the settings from the last transformation. You can revert to the defaults by holding down the Shift key while you select Color Balance from the Transformation menu or you can set *Remember Settings* to *No* to always start with the defaults.

Color Balancing for Different Light Sources

Different films are balanced for different light sources. Similarly, many digital cameras have white balance settings for light sources such as Tungsten and Daylight. The purpose of these settings is to ensure that white objects appear white in the final image. If the color of the light used to illuminate the object is different from the color for which the film or camera was set, white objects will come out colored.

You can create your own custom color balance settings to correct for this kind of problem as follows:

1. Photograph a pure white sheet of paper under the illumination conditions and using the film or camera settings for which you wish to correct.

2. In the Color Balance transformation, click on the image of the white paper to identify the highlight color to remove. This setting will cause the image of the white paper to come out perfectly neutral in the output image. You can then use this setting to correct other photographs taken in the same light with the same film or camera settings.
3. Save the color balance settings by clicking the  button in the main tool bar and entering a file name. The file you create in this manner can be reloaded later when you wish to make a color correction to another image.

Color balancing using a reference target

For a more definitive version of the procedure outlined above, you can photograph a reference target that has a graduated series of gray patches instead of just a white sheet of paper. By shift-clicking on several of the gray patches in the reference image you can create a set of curves that precisely remove casts from the entire tonal range. If you then save these settings, you can use them to correct subsequent images take under the same lighting conditions.



In the above example, a standard Kodak Q-13 gray step wedge was photographed under tungsten lighting on a copy stand with the camera set for daylight white balance. The resulting image is, as you might expect, very orange. A set of correction

curves was created from this test target by setting the black and white Remove colors to the colors of the darkest and lightest patches, and then shift-clicking on a series of patches in the image (#3, #6, #9 and #12) to create additional rows for intermediate gray levels. These settings can then be used to accurately correct other photographs taken under the same conditions.

If you don't have a Q-13, Macbeth ColorChecker, IT8 or other standard target that includes a series of neutral grays, you can print one on your own printer. Just make sure you print it using only black ink so it will come out truly neutral and that you print it on white paper. The step wedge file *Gray Wedge 15.tif* is available for download from our web site for you to print.

Color curves

The right side of the Color Balance dialog box displays a curve for each of the color channels, or you can view all three curves at once. You select the one you want to view using the tabs below the curves (R/G/B/RGB). If you select a single channel, its histogram is displayed for reference purposes in the curve background.

Points are marked on each curve according to the current values of the color patches. The x-coordinates of the red curve points are determined by the red channel values of the Remove colors; their y-coordinates are determined by the red channel values of the Add colors, and similarly for the other two channels. The point corresponding to the current row is identified by having larger markers. Unless you have RGB selected, you can directly manipulate the curves by dragging any of the points.

Useful Tip: In most cases smooth curves produce the best results. Usually if any of the curves has a kink in it (such as a sudden change in slope), the end result will look strange. The more Remove colors you are using, the easier it is to accidentally introduce kinks in the curves so it is a good idea to check them from time to time.

To avoid “foldover”, the sequence of x-coordinates for each of the three curves must always be in ascending order—if the selected colors do not have each channel in strictly ascending order, the out-of-order points are flagged with an *err* status and are temporarily disabled. This condition is unlikely to occur unless two adjacent Remove colors are very close to each other or one or more of the Remove colors are highly saturated. To correct the problem, delete one or more of the inconsistent rows.

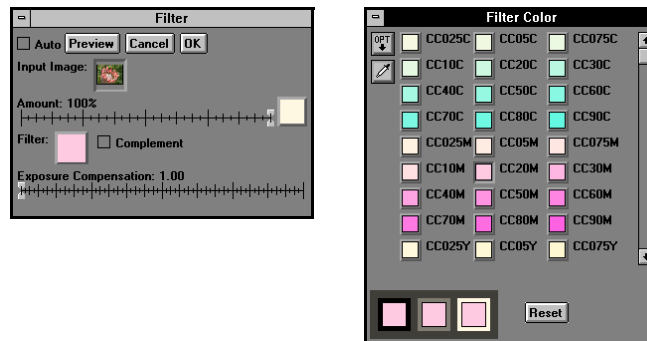
The way the color balance transformation works is to simply apply the curves derived from the Remove and Add colors to each channel of the input image to produce the output image (possibly modified by the Amount control). Because of the way the curves are constructed, any occurrence of a Remove color in the input image will be changed to the corresponding Add color. Colors intermediate in brightness between two adjacent Remove colors are computed as a smoothly blended average of the two Add colors. Thus, if the Add colors are all neutral and the Remove colors correspond to parts of the input image that should be neutral but appear with a color cast, the curves will remove that color cast.

Color correction using CC filters with the Filter Transformation


The traditional method of controlling color balance when printing color images in the darkroom is to use one or more color correcting (CC) filters. CC Filters are a graduated set of standard color filters available in each of the primary and secondary colors (red, green, blue, cyan, magenta, yellow). The filters are designated by their density and color. For example, a CC10M filter is a magenta filter with a density of 0.10. Filters with higher density transmit less light and are therefore darker and more intensely colored. Color balancing an image using CC filters has two steps. First you must identify the color cast of the image, and this takes a bit of practice. Then, to remove that color from the image, you must apply a filter whose color is the complement of the color you want to remove. For example, to remove the magenta cast in the image above, you would need to use a green filter (since green is the complement of magenta). In addition, the density of the CC filter you use must be proportional to the strength of the color cast in the original image.

This method of color balancing is fully supported by Picture Window with the Filter transformation and setting the filter color using the color picker's filter bank option. You can select any combination of standard CC filters and adjust for the filter factor using the Exposure Compensation slider. The procedure is as follows:

1. Open the input image and click on its window to select it.
2. Select Transformation/Filter from the main menu. The Filter dialog box is displayed.



3. Click on the large square button just to the right of the caption **Filter:** in the dialog box. Then select **Solid Color** from the menu that pops up. A color picker is displayed.

4. Click on the  button in the upper left corner of the color picker dialog box. Then select **Filter Bank** from the menu that pops up. This replaces the standard color picker color wheel with a set of colored buttons. Each button corresponds to a standard CC or Wratten filter and they are labeled with their filter designation. If necessary, you can use the scroll bar to bring more buttons into view.

Each time you click on a filter button, the corresponding filter color is selected. You can combine as many different filters as you like. To deselect a filter button, simply click on it a second time. To deselect all the filter buttons, use the **Reset** button. If a filter button is selected, it is displayed as depressed.

5. Click the Preview button in the Filter dialog box to see the effect of applying the filter, and adjust the filter color as necessary until the desired result is obtained. If you select the Auto preview check box at the top of the Filter dialog box, then each time you click on a new point in the color wheel, Picture Window will automatically rebalance the image and display a preview of the result for you to evaluate. Comparing the corrected version with the original by arranging the input and preview windows side by side is usually very helpful.

If a filter makes the image too dark, you can adjust the Exposure Compensation slider to lighten the image. To select the complement of the color you have specified, click the Complement check box. For example, to remove magenta from an image, you could select a magenta filter and then click the Complement check box to create a green filter that reduces magenta in the input image.

The Color Correction transformation

Both the Color Balance and the Filter transformations apply corrections uniformly to the entire image. In many cases, this works well, but sometimes the situation is a little more complicated. For example, Fuji Velvia has a tendency to produce highly saturated reds while reproducing other colors more realistically. While this works wonders in the red rock country of Utah, it can wreak havoc on skin tones or even pine needles. If you try to tone back the reds, you will then start boosting the greens and blues in the image as well (see below). This phenomenon is sometimes referred to as *crossover* and cannot be corrected with a color correcting filter.



Original



Filtered to reduce red

While some problems of this type can be fixed using the midtone balancing feature of the Color Balance transformation, the Color Correction transformation (available only in the Pro version of Picture Window) is the ultimate tool for making just this kind of selective color adjustment. It works by creating custom color lookup tables that apply different corrections to different parts of the color space.

For example, this lets you adjust just the reds in the sample image while leaving all other colors in the image unchanged as illustrated below.

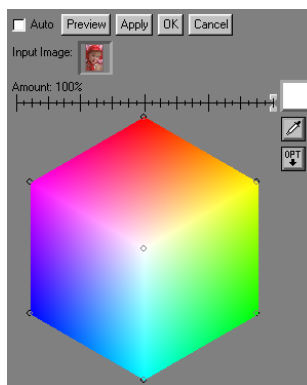


Original



Reds reduced using Color Correction

The Color Correction transformation is based on the same color wheel Picture Window uses in its color picker. The color wheel is a hexagon with white in the center and the fully saturated colors around its perimeter. Every possible combination of hue and saturation in the HSV color space corresponds to a different point on the color wheel. The higher the saturation, the further the point is from the center; variations in hue corresponding to rotations about the center. All the colors in the color wheel correspond to the maximum possible brightness ($V = 100\%$); every possible color can either be found somewhere on the color wheel or can be generated by darkening one of the colors on the color wheel by reducing each of its components by the same fraction.



To specify which colors are preserved and which are changed, the Color Correction transformation uses *control points*. Each control point consists of two parts: an input color and an output color. When the transformation is applied, the input image is examined one pixel at a time. If the hue and saturation of the pixel (regardless of its value or brightness) matches the input color of any of the control points, the pixel will be changed to the output color with the same brightness in the result image. For colors between control points, the output color is interpolated using the nearest available control points so there are no abrupt transitions between colors.

To preserve colors in one region of the color wheel, you create a control point whose input and output colors are identical. To make one color (and those in its immediate neighborhood) change to another, you create a control point whose input color is the one you want to change and whose output color is to one you want to change it to.

Control points are displayed on the color wheel. If the input and output colors of a control point are the same, it is displayed as a small circle at the corresponding location on the color wheel. If the output color differs from the input color, then an arrow is extended from the input color to the output color, indicating that the input color will turn into the output color.

Initially, there are seven control points, each of which have the same input and output color. One is located at the center of the color wheel (white) and one at each vertex of the hexagon (red, yellow, green, cyan, blue, and magenta). Nailing down the central white point guarantees that white (and consequently all neutral grays) will be left unchanged. Similarly, nailing down the fully saturated primary and secondary colors at the vertices of the hexagon keeps them from being changed. Running the transformation with these settings will, however, leave the input image unchanged as all the output colors are the same as the input colors.


While you can change either the input or output colors of any of these initial seven control points, you can also add new control points or remove existing ones. There are two ways to add a new control point:

1. You can shift-click on the color wheel to add a control point at a specific location.
2. You can enable the Probe (by depressing the eyedropper icon just to the right of the color wheel) and then click on a point in the input image window that has a color you want to change or preserve. This adds a new control point at the location on the color wheel corresponding to the color of the image at the point you clicked.

To change the output color of a control point, just click and drag it -- this causes an arrow to grow out of the control point whose tail is the input color and whose head is the output color. Once you separate the two, you can change the input color by dragging the little circle or the output color by dragging the head of the arrow.

To remove a control point, control-click on it.


As an illustration, let's see how to use this transformation to change reds in an image without changing anything else.



Leaving the initial seven control points in place, one additional control point was added by simply clicking on the eyedropper icon and then clicking on a sample of the red pine needles in the input image window. The control point was then dragged towards the center of the color wheel (see the arrow above) to more or less halve its saturation while leaving its hue the same, yielding the much more natural effect at the right. Note that the foliage, sky, and rocks have not been changed at all, nor

have the whites in the image been contaminated. This kind of very selective color balancing is very difficult to perform any other way.

In the following example, the more or less neutral waterfall can be made more blue without deadening the yellows, reds, and greens in the foliage by dragging just the white point:



If the same color shift was introduced into the image using the Filter or Color Balance transformation, the results would look like this:



Note the loss of vibrance in the foliage above the waterfall.

The Color Correction transformation can also be used with color printers. Some day, precision color management will be seamlessly integrated into all Windows printer drivers and images will all look the same on paper as they did on the monitor. Until then, the situation is less ideal. If your printer consistently renders certain colors incorrectly, you can use the Color Correction transformation to create (and save for later reuse using the Opt button) a set of lookup tables that compensates for your printer's idiosyncrasies. If a particular color is coming out with the wrong saturation or hue, simply create a control point and drag the color in the opposite direction. With some trial and error, you should be able to correct at least the more obvious problems.

Use Masks for Even More Control

For an extra level of control, the Filter, Color Balance, or Color Correction transformations can be used with a mask.

If only one section of an image needs correction, you can create a mask that isolates just that part of the image and restrict the color correction to that region.

Density masks are useful when you want to make color corrections just to the shadows, midtones, or highlights of an image. A density mask is simply a mask whose transparency at any given point is determined by the brightness of the image at that point.

For more information on creating and using masks, please see the Picture Window manual, the help file, or the white paper on Creating and Using Masks.