

Color Filter Array (CFA) Processing

Color cameras do not require filter wheels to produce color but require some unique image processing techniques.

An observation made of the first edition was its lack of emphasis on using conventional digital (color) cameras. It is easy to get hung up on these things but the camera is just one part of a vast array of other equipment, software and processes that are common, regardless of the choice. Whether or not a camera is a DSLR (or mirror-less model) or a dedicated CCD, it is still connected to a telescope, power and a USB cable. There are some differences though; some snobbery and also misleading claims about the inefficiency of color cameras: For an object that is not monochromatic, there is little difference in the received photon count via a color camera or a sensor behind a filter wheel for any given session. Unfortunately, conventional digital cameras alter the RAW image data and the (CMOS)output is not as linear as that from a dedicated CCD camera. Monochrome sensors do have other advantages; there is a slight increase in image resolution, the external red and green filters reject Sodium-yellow light pollution and more efficient for narrowband use. As a rule, dedicated cameras are more easily cooled too. These are the principal reasons both book editions major on monochrome CCD cameras, using an external filter wheel to create colored images and unmolested linear sensor data. This chapter addresses the important omissions concerning linearity and color formation in conventional color cameras.

Both one shot color (OSC) CCDs and conventional digital cameras have a Color Filter Array (CFA) directly in front of the sensor that require unique processing steps. In addition, photographic camera's RAW formats are not an unprocessed representation of the sensor's photosite values and require attention during image calibration and linear processing. The most common CFA is the Bayer Filter Array, with two green, one red and one blue-filtered sensor element in any 2x2 area, but there are others, notably those sensors from Fuji and Sigma. Consumer cameras suitable for astrophotography output RAW files, but this moniker is misleading. These cameras manipulate sensor data by various undocumented means before outputting a RAW file. For instance, it is apparent from conventional dark frame analysis, at different exposure conditions, that some form of dark current adjustment kicks in. At some point, a long dark frame has lower minimum pixel values than a bias frame. Each RAW file format is potentially different: In some cases the output is not linear, especially at the tonal

extremes, which affects conventional calibration, rejection and integration processes. In others, most notably older Nikon DSLRs, the RAW file had noise reduction applied to it, potentially confusing stars for hot pixels.

The traditional workstreams that transfer a RAW camera file to say Photoshop, go through several translation, interpolation and non-linear stretching operations behind the scenes to make the image appear "right". This is a problem to astrophotographers who are mostly interested in the darkest image tones. These are at most risk from well-meaning manipulation designed for traditional imaging. So how do we meet the two principal challenges posed by non-linearity on image calibration and color conversion, using PixInsight tools?

Image Calibration

First, why does it matter? Well, all the math behind calibration, integration and deconvolution assume an image is linear. If it is not, these do not operate at their best. It is important that we calibrate an image before it is converted into a non-linear form for an RGB color image. Second, it is never going to be perfect: RAW file specifications are seldom published and are deliberately the intellectual property of the camera manufacturers. To some extent, we are working blind during these initial processing steps. For instance, if you measure RAW file dark current at different temperatures and exposure lengths, you will discover the relationship is not linear (as in a dedicated CCD camera) but the calibration process assumes they are. In addition, as soon as a RAW file is put into another program, the programs themselves make assumptions on the linearity of the image data during color profile translations. Many applications (including PixInsight) employ the open-source utility DCRAW to translate a RAW file into a manipulatable image format when it is opened. Over the years, this utility has accumulated considerable insight into the unique RAW file formats. In the case of most photo editing programs, these additionally automatically stretch the RAW image so it looks natural.

Each of the various popular image file formats, JPEG, TIFF, FITS, PSD and the new XISF, have a number of options: bit depth, signed/unsigned integers, floating point, with and without color profiles and so on. When PixInsight loads one of the myriad RAW file formats, it

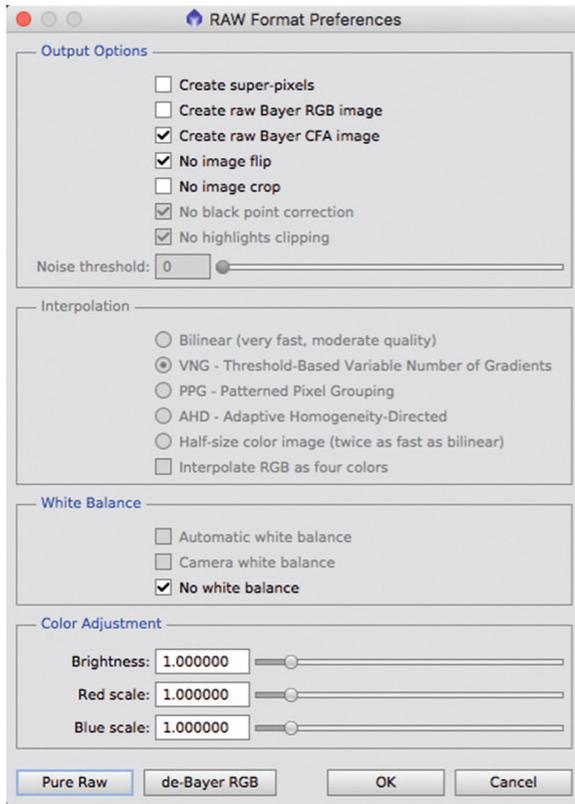


fig.1 It is essential to set the right output options in the RAW Format Preferences for DSLR_RAW (found in the Format Explorer tab). Here it is set to convert to a monochrome CFA image (Pure Raw) without any interpolation (DeBayering).

converts it into an internal format called DSLR_RAW. This has several flavors too, under full control in the Format Explorer tab. The options allow one to retain the individual pixel values or convert into a conventional color image. A third option falls in-between and produces a colored matrix of individual pixel values (fig.2).

These options do not, however, change the tonality of the image mid tones (gamma adjustment). For example,

if you compare an original scene with any of the DSLR_RAW image versions, it is immediately apparent the image is very dark. This is more than a standard 2.2 gamma adjustment can correct. (2.2 is the gamma setting of sRGB and Adobe 1998 color profiles.) The reason is that the sensor data is a 14-bit value in a 16-bit format. This is confirmed by a little experimentation; if one opens a raw file with the HistogramTransformation tool and moves the highlight point to about 0.25, to render a clipped highlight as white, and adjusts the gamma from 1.0 to 2.2 it restores normality in the RAW files' mid-tones and the image looks natural (similar to if you had imported it into Photoshop).

Calibration Woes

When you look up close at a RAW file (fig.2) and consider the calibration process, one quickly realizes that calibration is a pixel by pixel process, in so much that the bias, lights and darks are compared and manipulated on the same pixel position in each image. To create a color image it is necessary to interpolate (also called de-mosaic or DeBayer) an image, which replaces each pixel in the image file by a combination of the surrounding pixels, affecting color and resolution (fig.2). Interpolation ruins the opportunity to calibrate each pixel position and is the reason to keep the pixels discrete during the bias, dark and flat calibration processes. To do this we avoid the interpolated formats and choose either the Bayer CFA format or Bayer RGB option in the DSLR_RAW settings (fig.1). These settings are used by any PI tool that opens a RAW file. The Bayer RGB version, however, occupies three times more file space and separates the color information into three channels. This has some minor quality advantages during image calibration but is computationally more demanding. (You might also find some older workflows use 16-bit monochrome TIFF to store calibration files. When the



fig.2 From the left are three highly magnified (and screen-stretched) output formats (see fig.1) of the same Canon CR2 RAW file; raw Bayer RGB, Bayer CFA and DeBayered RGB file (using the VNG option). You can see how the DeBayer interpolation has smeared the effect of the bright pixel. The green cast is removed later on, after registration and integration, during the color calibration processes.

Debayer tool is applied to them, they magically become color images.)

The latest calibration tools in PixInsight work equally well with image formats that preserve the CFA values, as does the BatchPreprocessing (BPP) Script, providing the tool knows these are not de-mosaiced RGB color images. It is important to note that image capture programs store camera files in different formats. Sequence Generator Pro gives the option to save in Bayer CFA or in the original camera RAW format (as one would have on the memory card). Nebulosity stores in Bayer CFA too, but crops slightly, which then requires the user to take all their calibration and image files with Nebulosity. After a little consideration one also realizes that binning exposures is not a good idea. The binning occurs in the camera, on the RAW image before it is DeBayered and the process corrupts the Bayer pattern.

When the file has a RAW file extension, for instance. CR2 for Canon EOS, PI knows to convert it using the DSLR_RAW settings. When the file is already in a Bayer CFA format, PI tools need to be told. To do this, in the BPP Script, check the CFA images box in the Global Options and in the case of the separate calibration tools, enter "RAW CFA" as the input hint in the Format Hints section of the tool. When dark and light frames are taken at different temperatures and/or exposure times, a dark frame is traditionally linearly-scaled during the calibration process, as it assumes a linear dark current. Any non-linearity degrades the outcome of a conventional calibration.

Fortunately, the dark frame subtraction feature of PixInsight's ImageCalibration tool optimizes the image noise by using the image data itself, rather than the exposure data in the image header, to determine the best scaling factor. As mentioned earlier, while both CFA formats calibrate well, of the two, the Bayer RGB format is potentially more flexible with manual calibration. The BPP script produces the exact same noise standard deviation per channel (with the same calibration settings) but when the color information is split into three channels and using separate tools, it is possible to optimize the settings for each to maximize image quality.

Color Conversion

The CFA formats from either photographic cameras or one-shot color CCDs retain the individual adjacent sensor element values, each individually filtered by red, green or blue filters. In contrast, a conventional RGB color image has three channels, with red, green and blue values for a single pixel position in the image. The conversion between the two is generically called de-mosaicing or more typically DeBayering and the next step, registration, in our linear processing workflow requires DeBayered images. (Integration uses registered light frames in the same vein but note that the ImageIntegration tool can also integrate Bayer RGB/CFA files, for instance, to generate master calibration files.)

The BPP Script DeBayers automatically prior to registration when its CFA option is checked. If you are

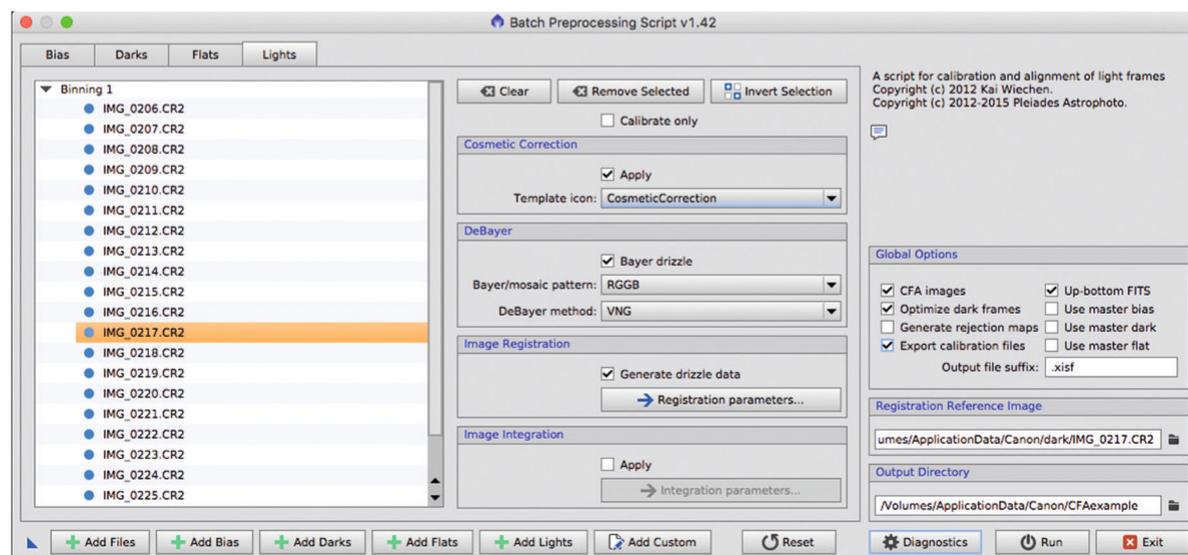


fig.3 The Batch Preprocessing Script set up to work with CFA files. In this specific case, rather than use a DeBayer interpolation to generate RGB color files, it has been set up for Bayer Drizzle. As well as the normal directories with calibrated and registered images, it additionally generates drizzle data that are used by the ImageIntegration and DrizzleIntegration tools, to generate color files at a higher resolution and approaching the optical limitation. Make sure to use the same file format for bias, dark, flat and light.

registering your images with the StarAlignment tool, however, you need to apply the BatchDeBayer or BatchFormatConversion script to your calibrated image files before registering and integrating them. One thing to note; when using the BPP script, if the Export Calibration File option is enabled, the calibration master files and calibrated images are always saved in the mono Bayer CFA format but when integrating calibration files with the ImageIntegration tool, the final image is only displayed on screen and can be saved in any format, including Bayer CFA or Bayer RGB. The trick is to make a note of the settings that work for you.

DeBayering is a form of interpolation and combines adjacent pixels into one pixel, degrading both color and resolution (fig.2). Since the original Bryce Bayer patent in 1976, there have been several alternative pixel patterns and ways to combine pixels to different effect. PixInsight offers several, of which SuperPixel, Bilinear and VNG are the most common and that interpolate 2x2, 3x3 or 5x5 spatially-separate sensor elements into a single color “pixel”. These various methods have different pros and cons that are also more or less suited to different image types. Most choose VNG over the Bilinear option for astrophotography since it is better at preserving edges, creates fewer color artefacts and has less noise. The SuperPixel method is included too for those images that are significantly over-sampled. This speedy option halves the angular resolution and reduces artefacts. I stress the word over-sampled since the effective resolution for any particular color is less than the sensor resolution. For images that are under-sampled (and which have 20+ frames) there is also an interesting alternative to DeBayering called Bayer Drizzle with some useful properties.

Registration and Integration

Star alignment works on an RGB color image, in which each pixel is an interpolated value of its neighbors. After registration, integration does likewise, outputting an RGB file. During integration, some users disable pixel rejection if the file originates from a photographic camera but otherwise enable them for dedicated OSC cameras. That, of course, leaves the door open to cosmic ray hits, satellites and aircraft trails. It certainly is worth experimenting with both approaches and compare the output rejection maps to tune the settings.

Bayer Drizzle

As the name implies, this uses the resolution-enhancing drizzle process on Bayer CFA images to produce color images. Drizzle is a technique famously used to enhance the resolution of the Hubble Space Telescope’s images, by

combining many under-sampled images taken at slightly different target positions. This technique can recover much of an optical system’s resolution that is lost by a sensor with coarse pixel spacing. For drizzle to be effective, however, it requires a small image shift between exposures (by a non-integer number of pixels too) that is normally achieved using dither. Most autoguiding programs have a dither option and many users already use it to assist in the statistical removal of hot pixels during integration.

Bayer Drizzle cleverly avoids employing a DeBayer interpolation since, for any position in the object, a slight change in camera position between exposures enables an image to be formed from a blend of signals from different sensor elements (and that are differently filtered). In this case, the resolution recovery is not compensating for lost optical resolution but the loss in spatial resolution that occurs due to the non-adjacent spacing of individual colors in the Bayer array. Thinking this through, one can see how wide-field shots may benefit from this technique as the angular resolution of the sensor is considerably less than the optical resolution.

As with image calibration and integration, one can either use several tools consecutively to achieve the final image stack or rely on the BPP Script to accomplish the task more conveniently, by enabling the Bayer drizzle option in the DeBayer section and the Drizzle option in the Image Registration section. With these settings, the BPP script generates calibrated and registered files, as normal, in a folder of the same name. In addition, it generates drizzle data in a subdirectory of the registered folder, using a special drizzle file format (.drz). Bayer drizzle requires a two-stage image registration and integration process:

- 1 Add the **registered** fits images into the ImageIntegration tool (using Add files button).
- 2 Add the drizzle .drz files in the **registered/bayer** sub-folder using Add Drizzle Files button. This should automatically enable the Generate Drizzle data option.
- 3 Perform image integration as normal, maximizing the SNR improvement and at the same time just excluding image defects.
- 4 In the DrizzleIntegration tool select the updated .drz files in the **registered/bayer** folder. Set the scale and drop shrink to 1.0.

If you have many images, say 40 or more, it may be possible to increase the scale beyond 1.0. Considering the actual distribution of colored filters in the Bayer array, a scale of 1.0 is already a significant improvement on the effective resolution of any particular color. In this process the initial ImageIntegration does not actually integrate the

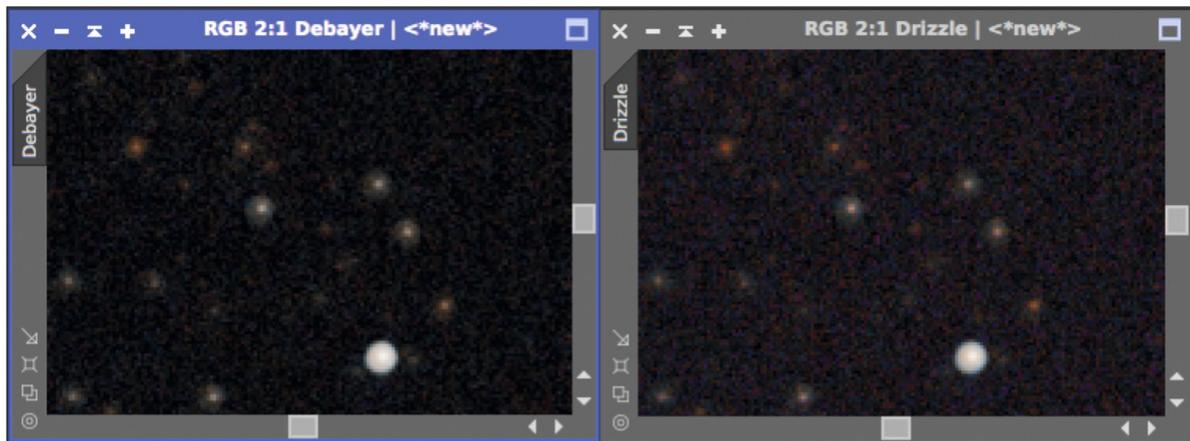


fig.4 This 2:1 magnified comparison of DeBayered registration and Bayer Drizzle process on 40 under-sampled subframes. The Bayer Drizzle process produces marginally tighter stars with more saturated color (and chroma noise). You have to look hard though!

images but simply uses the registered files to work out the normalization and rejection parameters and updates the drizzle (.drz) files with these values. The DrizzleIntegration tool uses these updated drizzle files to complete the image integration. The proof of the pudding is in the eating and fig.4 compares the result of 40 registered and integrated wide-field exposures, taken with a 135 mm f/2.8 lens on the EOS through a standard DeBayered workflow and through the Bayer Drizzle process. To compare the results, we need the sensor to be the limiting factor, or the lens resolution and seeing conditions may mask the outcome. In this case, the wide angle lens has a theoretical diffraction limited resolution of approximately 2.6 arc second and appears poor until one realizes that the sensor resolution is over 6.5 arc seconds/pixel and under-sampled to use for comparison purposes. (In practice, seeing noise and tracking errors probably put the effective resolution on a par with the sensor resolution.)

Post Integration Workflow

The practical workflows in the First Light Assignment section assume a common starting point using separate stacked images for each color filter. In these, the luminance and color information are processed separately before being combined later on. In the case of CFA images, we have just one RGB image. These workflows are still perfectly valid though. The luminance information (L) is extracted from the RGB file using the ChannelExtraction tool or the Extract CIE L* component toolbar button and follows its normal course. For the color file, the various background equalization, background neutralization and color calibration are simply applied to the RGB file. Once the RGB file has been processed to improve color saturation, noise and is lightly stretched, it is mated with its now deconvoluted,

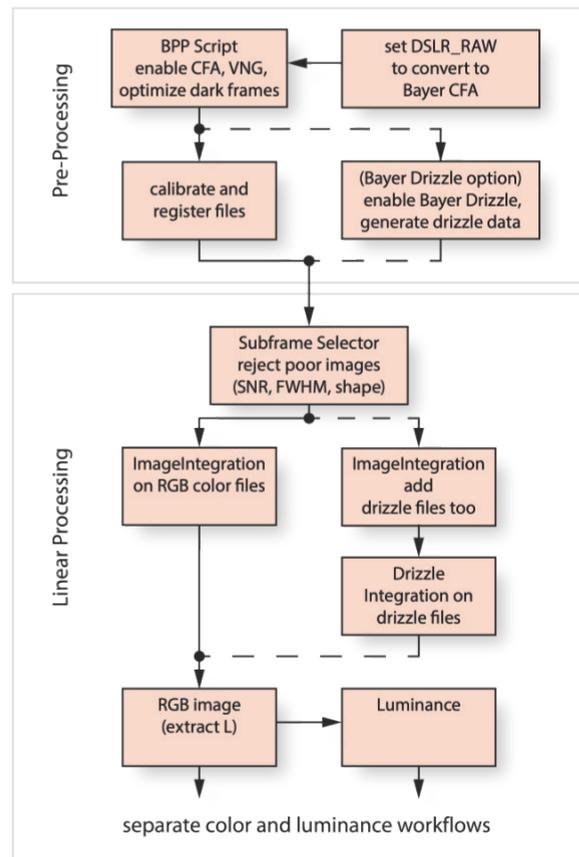


fig.5 The CFA workflow using standard DeBayered registration or Bayer Drizzle process, through to the start of the separate processing of the color and luminance data (the starting point for many of the practical workflows throughout the book). The BPP Script can be replaced by the separate integration (of bias, darks and lights), calibration and registration tools if one feels the urge.