Tricks for High-Precision CCD Photometry

Definition

I'm going out on a limb, but will define high-precision photometry as any observation with 0.01mag or better internal accuracy. Every CCD system I know of is capable of this kind of photometry; you just have to understand your hardware limitations and take reasonable care in making your observations.

High-precision photometry is easier to do with differential photometry. All-sky photometry (where you measure program objects with respect to standard stars elsewhere in the sky) is more difficult and will be left to the reader to study in textbooks like "CCD Photometry". I'm going to assume that you have already been doing simple photometry, and are familiar with terms like 'fwhm'.

Pixel Size

You don't often have the luxury of changing your pixel size after acquiring a telescope and CCD system! However, if you are contemplating the purchase of a new system and want to do high-precision photometry, then try to optimize your pixel size. For most front-illuminated chips, you want 3 pixels per fwhm seeing size. For example, if your typical seeing is 2arcsec, then you would want 0.67arcsec pixels. This oversampling smooths out the sub-pixel variations caused by the overlying gate structure of a CCD. Of course, using more pixels to sample the image means you reduce your field. Compromises always have to be made! For a back-illuminated chip, you can get by with 2 pixels per fwhm.

Other ways to avoid the errors of undersampling your images are to only observe in the red (the gates become transparent there), or to take multiple exposures and dither between each exposure so that the image gets moved to different pixels.

I don't recommend the use of field-crunchers, as they cause distortions at the edges of the field, usually reduce the image scale so that you are undersampled, and may not be well color-corrected in the near-IR. If you want the big field to find comparison stars, why not find some new program objects that have nearby comparison stars? If you need the big field because your telescope mount can't point well enough to find a smaller field, then improve your mount or pointing model. Sorry, but that is life!

Binning

Most CCD cameras give you the ability to 'bin' your pixels to make superpixels in 2x2, 3x3 or 4x4 mode. We don't recommend doing this for photometry unless you are greatly oversampled (rare on amateur telescopes). While such binning decreases the read noise (you are reading one pixel instead of 4, 9, or 16), it makes it much more difficult to center your measuring aperture and actually decreases your dynamic range (assuming your CCD is set up so that full well is pretty near full range on your ADC). The exception is when you are trying to measure an extended object, where apertures would be huge and the surface brightness is low and read noise is important. That is beyond the scope of this article.

Anti-Blooming

Antiblooming gates decrease the blooming that occurs on bright stars in the CCD field. This is great for making pretty pictures of extended objects, but causes problems for photometry. The A/B gates decrease the effective size of the pixel and therefore decrease the sensitivity of the chip. They also decrease the full-well of the pixel because they siphon off charge from bright objects, thereby making the response nonlinear. I recommend either not using a CCD with antiblooming capability, or if possible, turn off the antiblooming feature in your camera control software. If you can't do this, then try to select a comparison star that is nearly identical in brightness to your program object and keep the signal level below 1/2 range on your ADC.

Filters

While you can certainly do high-precision photometry with an unfiltered CCD, I recommend the use of wide-band filters. These give you some color information about an object and permit the comparison of your measures with other observers. The most common filter set is the Johnson-Cousins UBVRI. These filters are not cheap, costing about $150 per filter, but this is a necessary evil if you want to really get into photometry.

Flatfielding

To get the highest accuracy, you need to flatfield your images. This removes the high frequency pixel-to-pixel variations in sensitivity for your CCD, as well as lower frequency wavelength sensitivity variations and telescope vignetting. Flatfielding is the single most important calibration step for your CCD. More details about flatfielding can be found in any textbook.

The best flatfields are of a uniform light source of about 5000K. Neither the blank sky nor a dome screen illuminated by a tungsten light source matches these criteria, so you always have to make concessions. Sky flats can be either made 'tracked' or 'untracked', are made normally at dusk/dawn but can also be formed from a median of all frames taken during a night ('master sky flat'). I usually take tracked dusk sky flats, taking 3-5 flats per filter with small (30arcsec) moves between frames for star rejection. I then combine the flats by (1) subtracting any additive component; (2) normalizing the flats to a common mean value; (3) summing the flats with rejection of 1 high pixel to remove cosmic rays and stars; (4) normalizing the result so that the mean of the flat is 1.000. Sky flats can only be made during a 20-minute interval at dusk and dawn, and are difficult when clouds are present, and may be hard with some open-truss telescope tubes (scattered light). Untracked sky flats also work, but you are then usually limited to the signal/noise in a single image and may have to take several before any star trail is of minimal amplitude.

Dome flats are similar. You illuminate a white screen with a lamp, and take multiple flats per filter for cosmic ray rejection. Here you may not need to normalize the flats before rejection since the lamp should be fairly constant. Dome flats require a dark dome (wasting night-time hours), need a special paint to get good blue/UV reflectivity, and require filtration on the lamp to mimic a bluer blackbody.

For any flat, expose to about 1/2 full well. You want enough signal to have high signal/noise in each pixel, but not so much that you run the risk of saturation or nonlinearity. Use long enough exposures that any shutter vignetting is unimportant. Watch out for scattered light. This is an additive component that should be subtracted out before making a flat.

Shutters

You should never use exposures less than 3 seconds duration. Shorter than this, the open/close time of a shutter causes vignetting in the field, affecting the photometry. If you use shorter exposures, you need to generate a shutter vignetting correction table. IRAF has a discussion of this, as well as the discussion in CCD Photometry.

Also, your estimate of the shutter open duration is based on knowledge of when the computer sends the 'open' command and when it sends the 'close' command. These timings may or may not correspond closely to the actual open time, so longer exposures lessen the impact of any timing error. While such errors do not directly affect differential photometry (since all objects on the frame have the same exposure time), they do affect any all-sky calibration of the field, or any case of traditional photometry where the comparison star does not fall on the same image as the program object.

If your program object is so bright that you need to use shorter exposures, then find a different program object! Realistically, you have a couple of solutions to the problem. First, you can make the shutter correction as mentioned above. You can also purposely defocus your images to spread the peak light over several pixels. With a refractor, this works well. With a reflector, you start seeing the central obstruction if you get far out of focus, and the images get definitely non-gaussian. I defocus to perhaps 2x fwhm with a reflector, and only use aperture photometry. Note that defocussing changes the flatfielding, so again don't defocus much. Another alternative is to wait for cloudy weather. That cuts the light down so that you can work on brighter fields. Finally, you can always reduce your effective telescope aperture with aperture masks, neutral density filters, using blue filters to hit the poorer part of the CCD sensitivity, or even use narrow-band filters like H-alpha.

Scintillation and Image Motion

Scintillation modulates the incoming light from an object and can be considered an additive noise source like Poisson noise. It has a very high frequency, and so can be averaged out by using longer exposures. My rule of thumb is to keep my exposures longer than 10 seconds to lessen the effect of scintillation. Scintillation is also larger the closer you get to the horizon, so I rarely take important frames at airmasses greater than 2.0.

Image motion is the basic component of seeing. For small telescopes, you will see lots of image motion, which can be removed to a large extent by the new 'Adaptive Optics' tip/tilt mirror systems. If you don't use an AO system, then your best bet is to use longer exposure times to let the image motion 'fill in' the seeing profile. Again, I would use 10-second exposures for averaging out the effects. Image motion is greater towards the horizon, so again keep your airmass less than 2.0. Image motion is less important with bigger apertures, as the motion is integrated over the mirror surface. Likewise, tip/tilt has less effect for bigger telescopes, and you have to go to higher-order terms.

Differential Refraction

As you get closer to the horizon, starlight spreads out into a spectrum. This causes all kinds of problems. Any autoguider has a specific bandpass, and any CCD observations at a different bandpass will have trailed star images since the autoguider is not tracking the same 'image'. The spread-out star images are quite different than those at the zenith, so doing all-sky photometry gets difficult. Within a single CCD field, you can get differential effects since the amount of refraction is different at the top of the frame than it is at the bottom. Again, the solution is to stay away from the horizon.

Aperture Size

You have a definite advantage over the older photoelectric systems because you have an infinite number of apertures that can be used to extract your data, and you can do it after-the-fact. At the same time, you need some basic rules as to what aperture to use to get the maximum signal/noise in your images.

The most basic rule is to use an aperture that is 5x the typical fwhm seeing profile. For example, if your seeing is 2 arcsec, you would want to use a 10arcsec diameter aperture. Larger apertures are ok for bright objects or uncrowded fields.

For faint objects, or for crowded fields, you may wish to use a smaller aperture. This is ok, but you should use the same aperture for all objects in a single frame. Smaller apertures do throw away a larger fraction of the light in a stellar image, but also include a smaller amount of sky and so you get increased contrast. Centering becomes important, as any error gets magnified as you use smaller apertures. Faint companions to any comparison star may be included in larger apertures but excluded in smaller ones, so the basic concept of constancy for a comparison star may be compromised if you use a large aperture one night and a smaller one on a different night.

For standardizing your field by the use of Landolt standards and all-sky photometry, I highly recommend using a 12arcsec aperture, no matter what the seeing is. Landolt used a photoelectric system with this size aperture, and you need to use the same size aperture to include any faint background stars that Landolt did in his measurements.

Mixing aperture sizes from image to image or within a given image can be done, but you then need to make aperture corrections. That is beyond this simple document, so I recommend reading the appropriate section in CCD Photometry or else the article by Steve Howell in the PASP Conference Series volume 23.

Comparison Stars

Use comparison stars that are roughly the same brightness as your variable. Somewhat brighter is certainly ok, but be sure that any star you want to measure is nowhere near the saturation point of your CCD. Select comparison stars that are also similar in color to your program object to lessen transformation errors.

A neat trick is to use multiple comparison stars in the same field. This gives you a handle on how good the frame is (form magnitude differences between the comparisons). Also, by averaging the comparison star magnitudes to form a 'mean' comparison, you improve the signal to noise in the measurement. Then, the major error in your measurement is the Poisson noise in the program object itself. Using multiple comparison stars to form a mean is also a method to use fainter comparison stars in those fields without good bright comparisons. In fact, I usually measure every star in the field, rejecting those with signal/noise less than 30, but averaging all the remaining stars. Why not -- such measurements are free and give the added advantage that you might discover new variables (about 2 percent of all stars are easily found to be variable). Also, multiple comparison stars average out any spatial errors (such as gradients in your flatfield) since they tend to be spread randomly over your image.

Differential vs. All-Sky Photometry

For the highest precision, you need to perform differential photometry. Keep the comparison star(s) and the program object on the same frame. This eliminates almost all sky effects such as clouds. In fact, I've done 5mmag photometry with 4 magnitudes of cloud extinction. Keep your fields centered the same way night-to-night. This removes much of any flatfielding error since you are using the same part of the CCD for the measurement of each star.

For some fields, where there are no comparison stars, and for the times when the sky is photometric and you want to standardize the comparison stars in your fields, you need to perform all-sky photometry. This is an intermixing of program field observations with observations of Landolt fields. You need to observe several Landolt standards with both red and blue colors, and at varying airmasses from as close to 1.0 as possible to 2-3 airmasses. Some observers, such as Berdnikov (MSU), never do differential photometry, but instead do all-sky observations of their program objects. This requires lots of extremely good photometric nights, a commodity not common at most sites in the U.S. Don't do it unless you are an expert! More details about all-sky photometry can be found in CCD Photometry. Good photometric nights in the Southwest will yield 0.01-0.02mag precision with all-sky photometry, so inferior to differential photometry even at the best sites.

Summary

You can do high-precision photometry with almost any CCD system, but you need to be careful. Follow the guidelines above as much as possible.

Arne Henden / USNO / aah@nofs.navy.mil