How To Create Good Sequences

What Are Sequences?

A sequence is a set of stars within a given field of view for which we have accurate, standardized photometry. Sequences are usually used to calibrate differential photometry (photometry with respect to stars within a field). For visual observers, such sequences might be visual magnitudes of 3-4 stars per magnitude step, covering the entire range of variability of a given star. This permits the differential interpolation of the variable's magnitude with respect to the sequence. For CCD observers, a sequence might be a set of stars that span the magnitude range of the variable, and also bracket the variable's color. They would be used to set the zero point for the magnitudes in a given field, and optionally can be used as comparison stars for the variable. Sequences that would be used for determining the transformation coefficients for a given CCD/filter set are best when the stars have the same magnitude but span a wide range in color. In other words, the concept of a sequence remains the same, but how you create the sequence will depend on its intended use.

All-Sky Photometry

Sequences are created by performing All-Sky Photometry, a generalized form of differential photometry. There are many standard stars in the Johnson/Cousins system that are located primarily near the Celestial Equator. The set that most observers use was published by Arlo Landolt. To create a sequence, you want to perform differential photometry with respect to these standards. Simple, right? The complication arises when the desired sequence field does not include any of the Landolt standards. You must determine the nightly instrumental zero point (and transformation coefficients) using the Landolt standards, plus determine the atmospheric extinction coefficients, and then use these parameters to adjust your instrumental photometry in the desired field so that standardized magnitudes come out of the reductions. This procedure is called all-sky photometry, and has been discussed in another tutorial.

Sequences require more care than does differential photometry. Not only do you have to account for extinction, but the color range of stars in the field usually require careful attention to the transformation coefficients so that red stars are not given systematically different magnitudes than blue stars. Likewise, the stars will appear randomly over the field of view of the CCD, and you need to make sure you have good flat fields so that there isn't a systmatic difference between, say, the center and the edge of a field.

Preparing for Sequence Photometry

Say that you need to do sequences for three fields: one that transits near the zenith, one on the Celestial Equator, and one that is far south for your location. How would you go about obtaining sequences in each of these fields?

First, you should not attempt all-sky photometry (necessary for the sequences) on any night that does not look 'photometric.' What I mean by this is that the sky should be clear, with no hint of clouds, no dust or smog or smoke should be present, and the seeing should be reasonably good and stable. For many sites (such as the Midwestern U.S.), this will limit you to only a handful of nights per quarter. For other sites, such as in the Southwestern U.S., interior Australia, or South Africa, you may have seasons during which any night is likely to be photometric. You make do with what you have! You won't know for sure until you reduce your data on a later date whether the night in question was truely photometric, but you ought to 'stack the deck' as much in your favor as possible.

Next, I usually go to the Digital Sky Survey web site and pull down photographic images of the desired fields. Then I draw on each image the field of view of my CCD, and move it around to maximize the number of stars and magnitude range that I want to cover. These charts will be used at the telescope later.

Then I check to see when the fields will transit. You want to minimize the refraction and extinction effects to get the best possible results, and that is done when the fields pass across your meridian. I mark those transit times down on a schedule.

You then need to include observations of Landolt standards. There are about two dozen fields where many Landolt standards will fall within the field of view of your CCD, and I usually make sure I observe those fields. An example is SA110-503 where 8 Landolt standards are within a 5arcmin circle. I try to set up my schedule so that I observe two such fields very early in a night, and then observe at least one field per hour the rest of the night, ending my night with another Landolt field observation. It sounds like a lot of extraneous images beyond the program fields, but when doing sequence work, you really, really want to understand what the sky is doing. The first two fields are observed at large and small airmasses, and then I alternate between large and small airmasses the rest of the night. You should try to observe Landolt standards at larger airmass than any program field, and with a range of colors that extends beyond the colors expected in your program fields.

You also need to insert any calibration frames that are necessary for your CCD camera. For example, if you normally take sky flats, they will occupy the first and/or last half-hour of a night; bias or dark frames might best be done when it is dark out and a program field is not scheduled; dome flats also require darkness.

The end result is a schedule of observations, where you have program fields interspersed with standard star observations.

Observing Sequences

Even though a given field might have some bright stars, I almost never try to measure a star brighter than V=10. The Tycho satellite has observed all of these stars in B \& V, and you would just be duplicating that effort. After throwing these stars out, I take a test exposure and see what the next brightest star magnitude is. I set the final exposure such that the brightest star I want to measure does not saturate the detector, but the faintest star still has enough signal-to-noise to provide a good measure. For most CCDs, this dynamic range is between 4 and 5 magnitudes. A field that contains a V=11 star will give reasonable photometry down to V=16 or so in a single exposure. If you need more dynamic range than this, you will need to take at least two frames with differing exposure times. In that case, I try to adjust the exposure times so that there is at least a magnitude overlap between exposures, or a maximum of a factor of 40 between the exposure times. An exposure of 10 seconds could then be taken alongside an exposure of 400 seconds to give an 8-magnitude range. CCDs are very linear detectors, and this 8-magnitude range is likely to be very close to the truth.

You need to use at least two filters so that you can get a color to be used in the transformation process. If you used, say, just a V filter, then a red star in the frame would be systematically brighter or fainter than a blue star. These two frames with differing filters will most likely have to have different exposure times as well. I find B \& V to be the best filters for most sequence work as there is a very good relationship between V, (B-V) and (V-R) such that you can estimate (V-R) easily. Another good choice is V \& I and is often easier to do with a front-illuminated CCD such as found in many amateur cameras. A poor choice is V \& R.

Processing the Data

I usually extract every star from every frame I take. For your purposes, you may not want to be so complete. In that case, I'd be sure to extract every Landolt standard from their frames, and then pick and choose within your program frames a reasonable set of stars based on your purpose. Avoid any star that has a neighbor, as the neighbor can affect the calculated brightness of the desired star.

I usually use aperture photometry unless the field is particularly crowded. The best aperture size depends on the seeing and your pixelization, but in general I select an aperture that closely matches the photoelectric aperture Landolt used (10-14arcsec) so that I include any faint stars in my digital aperture that Landolt would have included in his original observations. For simplicity, keep this same aperture for all of your reductions on a given night. There are methods to compensate for different apertures on different frames, or for using point-spread-function fitting, but for the purposes of this simple tutorial, I will avoid them.

Do your normal transformation and extinction calculations so that you can compare your resultant standard star observations with the published values. The standard deviation of these measures tells you how stable the night was and how well you can determine the magnitude of a particular star for that night. Actually, the error for a given star is the square root of the sum of the nightly standard deviation (squared) plus the Poisson error (squared) of the star in question. That means the measurement error is a combination of both effects. My criteria for a photometric night is a standard error of less than 0.02mag; a marginal night might have errors up to 0.03mag; anythihg greater than this is considered non-photometric and the observations are retaken. You may have different criteria. For example, sequences for visual observers need only have 0.1mag accuracy, and sequences for calibrating photographic plates might only need 0.2mag accuracy.

I find that you can have systematic effects in a single night's observation of a program field. In the Southwestern U.S., we get cloud bands that can mess up measures to the far north, yet not affect the equatorial measures of the Landolt standards. I find the only way to get a good handle on these errors is to measure each field on at least two, and preferably three or more, nights and compare the results. By measuring on different nights, you also can check for variability of stars within the field, especially if the nights in question are spaced several nights apart. By taking the average of the measures, you get a better determination of the field magnitudes, and by taking the standard deviation of the measures of a given object, you get a better understanding of the real errors in the measurement of that object. It would be great if Landolt standards were placed everywhere in the sky, and even better if a standard star appeared in every frame you take with your CCD, but that is not currently the case. I usually use a minimum of three nights if I am really serious about getting the best results for a field. If on careful inspection, one of those three nights is significantly different than the other two, I discard that night and take data on a fourth night. Taking data on several nights also guards against the inclusion of possible variable stars in your sequence. You should always take at least two measures of a field, and use a single measure of a field only as a last resort.

I mentioned earlier that you need to get good flatfields. If there is a gradient in your flats, you may be able to get a good transformation since the Landolt standards appear in the center of your image, yet get very poor results for objects that appear in the edges of your program field. You can guard against that by rastering a Landolt standard across your CCD field of view and seeing if there are systematic differences in its magnitude estimate depending on location in the field of view. This requires either a very stable night, or else having the entire field filled with 'standards' on a single frame (such as looking at a well-studied open cluster such at M67).

Inspecting Your Data

For a given field, I calculate the mean magnitude and standard deviation for each object measured. Then I plot magnitude vs. error (a scatter plot) to check the data quality. A plot of this type for a field centered on suspected variable NSV528 is shown below.

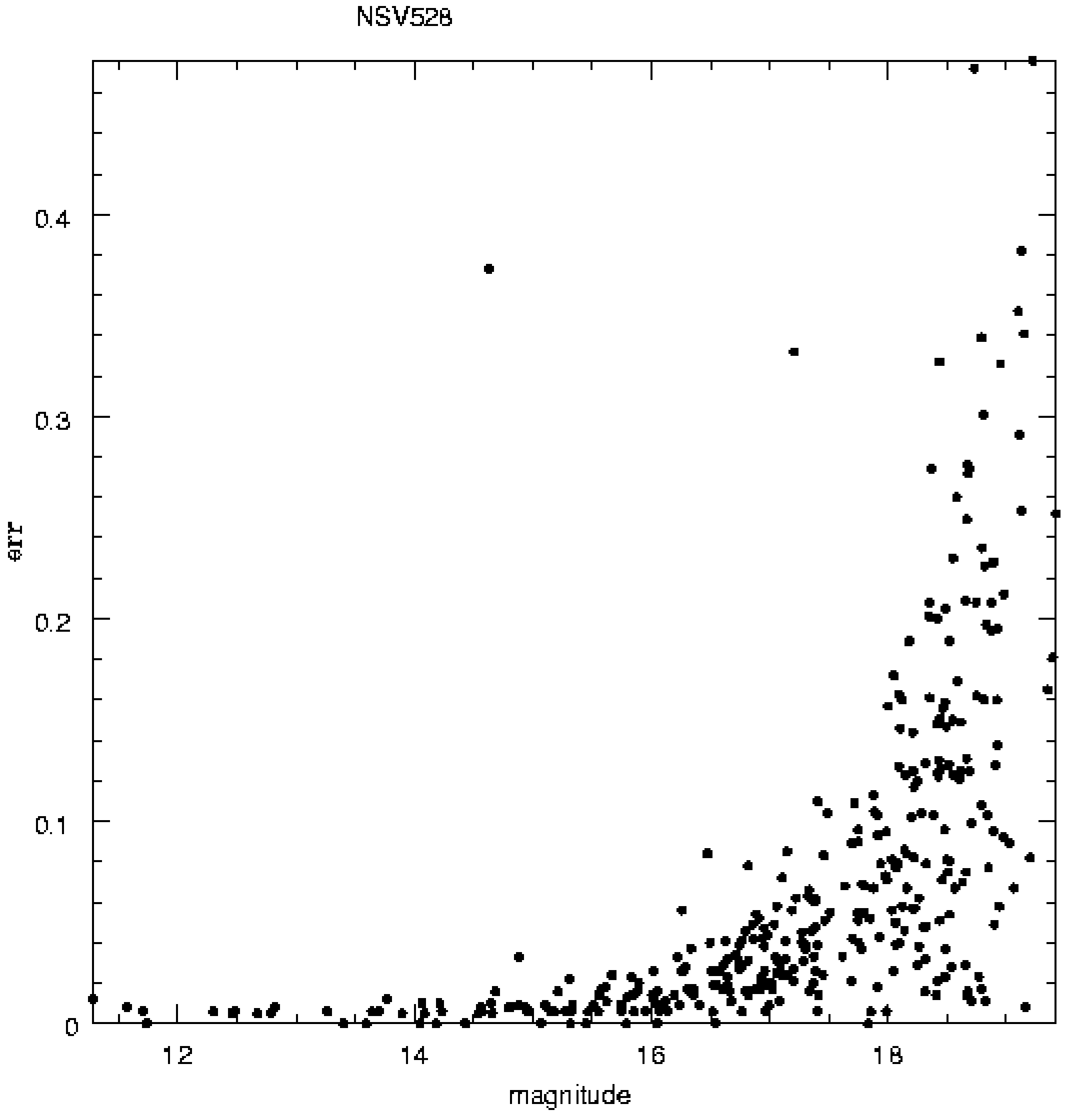

This plot shows the typical ridge line of Poisson photometric error: the errors are reasonably small to V=16, then exponentially get worse towards fainter magnitudes. Note the stars with large error around V=14.5 and V=17; these are both real variables. This kind of plot finds poor nights quickly (if you have 3 or more nights represented in the plot) since the smallest error (that is, for the bright stars) will be increased and the plot will look like there is an added offset. Likewise, if you have some saturated stars on the bright end, they will have larger error than some stars slightly fainter since the bright stars will have varying amount of saturation depending on the seeing, etc. There is usually about a 4-magnitude 'sweet spot' on a CCD for which the Poisson errors are small, and you need to adjust your exposures so that the sweet spot covers the desired sequence magnitude range.

Some of the problems typically found with all-sky photometry are shown on the scatter plot for TX Cam below, based on four nights of data:

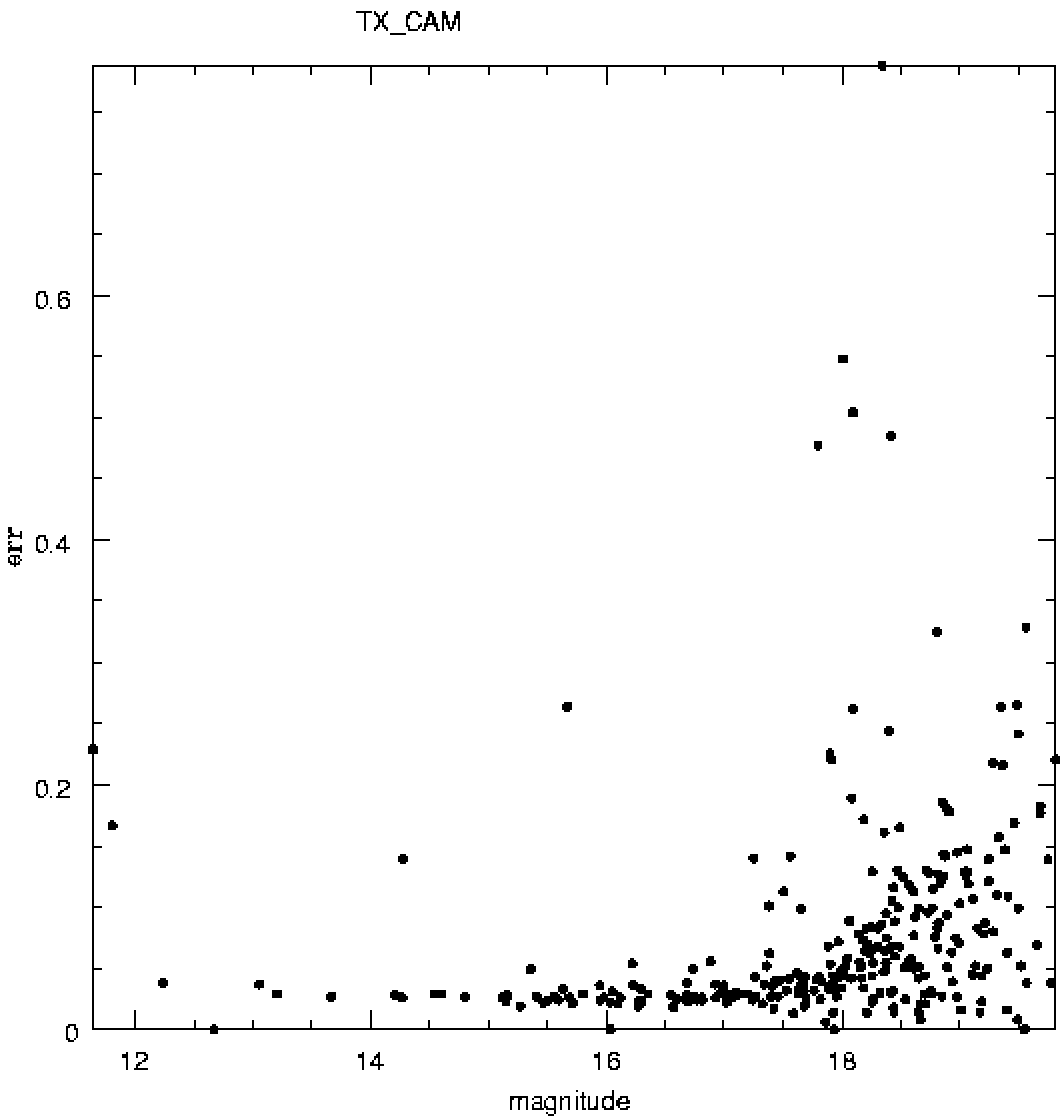

(this is not a great example, but I've eliminated all of the poor data from most of my normal photometry!). There are several problems with this dataset. Note that after magnitude V=17 or so, there are a number of stars that fall well above the ridge line. This field is somewhat crowded, and these faint stars happen to lie in the wings of brighter stars, or are only detected on a couple of the nights. A more important problem is an offset in the y direction (error axis), where the well-exposed stars fall on a horizontal line at about err=0.03. These stars have Poisson errors of less than 0.01mag, so the fact that they don't fall closer to err=0 on the plot indicates that one or more of the nights was slightly non-photometric. If the night was really non-photometric, you would see an offset of a tenth of a magnitude or more for the well-exposed stars. There are four stars that fall above the line and are brighter than V=17. The point at V=15.67, err=0.264 is a star that is on the edge of my CCD frame and so only appears on two of the nights, and even for those two the data is marginal. The point at V=14.276, err=0.140 is only a couple of arcsec away from the variable TX Cam and therefore highly contaminated. The The point at V=11.806, err=0.167 is a saturated star, and therefore has higher error than its Poisson error would indicate. If you had more bright stars in this field, you would see most of them having larger error and actually looking like a mirror image of the exponential rise of error found for faint stars. The point at V=11.643, err=0.264 is the variable TX Cam itself. While some of the scatter is due to the inherent variability of this star, the majority of the error is just saturation. I know it is the variable because it has the proper coordinates, magnitude, and color for this Mira.

Other Techniques

I've just told you how to take sequences on photometric nights. If you need a sequence quickly, and the weather does not look promising, you can sometimes cheat and do your program field bracketed by standard fields on either side, either in time or in space. This can yield reasonable results even if you only have, say, a cloudfree hour during the night. For such quick-and-dirty determinations, you can even use nearby sequences, such as those around a variable star, that have been published by others. The results will not be as good as if you use Landolt standards, but are often good enough.

Other methods include overlapping fields from a known set of standards or a nearby sequence, taking frames until you reach the field in question. You can then tie the frames together by setting the zero point of each frame based on its overlap with the previous frame. For bright sequences and wide field cameras, you can often include several Tycho stars in your field and work your way down in magnitude steps so that you overlap in brightness instead of space.

Conclusion

You can see that getting good sequences is not hard in theory, but in practise you need to guard against anything that causes a systematic error. This means you need to understand your hardware and do good flats, and you need to take many standard star observations through the night to understand what the sky is doing. I would start with a known field, such as M67 or NGC7790, and attempt to recreate the magnitude sequences known there with the techniques given above. Once you know you can get the same results as others have in a known field, then you can move on to the much scarier task of calibrating the unknown fields. Good sequences are necessary to many aspects of astronomy, and people who can create good sequences are in high demand (and regard!) by the community.

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