Chapter 7: Photometry and science

The first six chapters of this guide give you everything you need to make variable star observations with a CCD that may be useful for science. Most of the requirements, procedures, observing and analysis techniques are outlined there, and you are ready to start observing. This chapter is intended to give you some additional astronomical background that will aid you in planning and executing observations on your own that are more likely to yield scientifically useful results. In many cases, an observing campaign requested by the AAVSO or other organization will tell you exactly what observations they want and why; here we want to give you background on general principles that should guide your observing techniques. You can consider Chapter 7 an "extra", but you should at least read through it to see how we at the AAVSO think your observations should be made. In particular we want to focus on two things: (1) why filtered and transformed observations are useful, and (2) what to think about when forming an observing plan for specific classes of variables, including filter use, observing cadence, and exposure times.

Before we go further, "step zero" of your observing process should be to consult the AAVSO website to see what resources we have for observers, and what stars we're asking for data on. As an example, the AAVSO (and several other variable star organizations) run Observing Campaigns where data are requested on specific stars at specific times. There are also many perennial targets for which data are always needed, so there will be no shortage of targets for you. We won't cover which specific stars to observe here in the *Guide*, because there are too many that are worthy of observers' time — it's worth a book just for that topic alone. Just keep in mind that you can be selective of what targets you explore to improve how likely it is that your data will be used by researchers. The exception is when you yourself are the researcher, and you have a well–defined, novel research question that you want to answer with your observations, but that is also a topic for an entirely different guide.

Photometry and filters

Before you start, you may want to read Appendices A and B of this *Guide* that cover some physical background on light, and how stars radiate. The simplest thing to take away from that discussion is that starlight contains more information than how much of it arrives at your telescope at a given moment, and that you can learn more by making observations with standardized filters than by simply taking an unfiltered image. Photometric filters have well–defined wavelength cutoffs and transmission properties, and they were designed to closely approximate a standard system, such as the Johnson–Cousins or Sloan system bandpasses. If you measure starlight through one of these filters, you are making a measurement not of the total amount of light that comes in, but the total amount of light within a wavelength range defined by the bandpass of the filter.

Filtered photometry may provide very useful astrophysical information. Stars with different physical properties (like temperature or chemical composition) will have unique spectral characteristics as measured in each of these filter systems. For example, a star of spectral type "A" will have a spectrum such that if you obtain calibrated measures of the star in Johnson B and V, the difference in those calibrated magnitudes will be close to 0.0. Stated in a more familiar way, the (B-V) color of an A–star is close to zero. That was set by definition — it was how the magnitude systems were defined in the first place in the Johnson system. The (B-V) color of a G–type star, cooler than an A–type star, will be somewhere around +0.7 — the calibrated B–band magnitude of that star will be 0.7 magnitudes fainter than the V–band magnitude. Spectral types for stars are based in large part on their temperatures, which in turn are reflected in how their spectra appear. More importantly, if you obtain a set of calibrated photometry for a given star, you can then compare those colors against known spectral calibrations to determine the approximate spectral types of your stars. Precise spectral typing is more complicated (and usually involves taking spectra), but photometric colors can give you some useful information about the properties of stars. One obvious example that we won't go into here is the color–magnitude diagram, where the magnitudes and colors of stars in clusters lie on very well–defined locations on this diagram, and these locations correspond to different evolutionary stages like the main sequence and red giant branch.

Things get even more interesting for variable stars, because their colors can change while their overall light varies. Remember that colors may correspond in part to the temperature of a star. We also know that some stars change temperature during the course of their variations. A pulsating star like a Cepheid or RR Lyrae can change by 1,000 K or more during a pulsation cycle, and it so happens that this temperature shift results in a substantial change in color, especially in (B-V). So you'll see a few things if you perform calibrated multifilter photometry of a Cepheid. First, you'll see the V–band light curve will have a different amplitude than the B–band curve (and may even have a slightly different shape and phase). Second, because of the difference between V and B, you'll see that the color curve — a plot of (B-V) versus time — is also variable. This is useful information in Cepheids, because it's a good way of showing (for example) during what part of the light curve the star is hottest. You'll find similar examples in other classes of variables whose temperature changes during their variation, dwarf novae being a good example; they go into outburst because their accretion disks transition to a hot, bright state that temporarily overwhelms the light coming from the cooler, redder secondary star. There are also some other physical processes that can cause color changes, obscuration by dust being one example. Dust preferentially scatters bluer wavelengths of light out of the line of sight, making the underlying star appear redder than it otherwise would. Dust is one reason some long–period variables and R Coronae Borealis stars appear very red.

So why is all of this relevant to variable star photometry? Note that we used the word "calibrated" many times in the discussion above. When spectral standards were created, they were done so using very well–defined filters and equipment whose properties are measured and understood. They were also established in such a way that atmospheric extinction was calibrated and removed from the measurements. Your filters, your equipment, and your observing conditions will almost never match those of the observers who created the spectral standards that defined the various properties of stars. Thus, if you obtain a "V magnitude" and a "B magnitude" for a star without calibrating your filters and equipment or determining the atmospheric extinction, they will be different than those of the known standards. You might measure the (B-V) color of the G–type star mentioned above and find that it's +0.8 instead of +0.7, and that of the A–type star is +0.05 instead of 0.0. That's why you have to determine your transformation coefficients using well–defined standards: you're determining the corrections that you need to apply to your data so that your measurements are on the same system as those of agreed–upon standards. In that way, your magnitudes can be most easily compared to everyone else's magnitudes. It isn't that your magnitudes are "wrong" — it's that they're different. But the problem is then how to understand data from many different observers, all of whom are different. Ultimately your data will be a lot more useful if you can minimize the differences between your magnitudes and standard magnitudes. That is why we spend so much time asking people to transform their data.

Time considerations: variability timescales, exposure times, and cadence

If you've been a variable star observer for awhile, you're probably aware that different stars vary in different ways. Some stars can vary with timescales of seconds or minutes (like some cataclysmic variables) while others may change over weeks, months, or years. Some stars may even show both kinds of variability. This is something you need to keep in mind when deciding how to observe a given star. If you have many different variable types in your observing program, you almost certainly don't want to use the same method for every star. The three primary things to keep in mind are:

1. You must be able to obtain useful signal to noise with an exposure time that's less than the timescale of variation.

2. You will need to average multiple observations of bright stars where the integration time is very short (ten seconds or less) due to scintillation.

3. You should not over–observe a star whose timescale of variation is very long, nor under–observe a star whose timescale is very short.

Point (1) above is mainly a concern for stars that have very fast variations and are intrinsically faint. The classic example of this is the orbital light curve or superhump of a short–period cataclysmic variable. There are a number of CVs whose orbital periods are 90 minutes or less, but which are also very faint. The trick is to figure out how to balance signal–to–noise requirements with the requirement that your exposure time doesn't smear out any interesting rapid variations.

Point (2) is a common concern for those instrumental observers working at the bright end, brighter than 7th or 8th magnitude for many typical SCT+CCD systems. Scintillation is a rapid change in intensity of starlight caused by inhomogeneities in Earth's atmosphere. There's nothing you can do to avoid it, only to average out its effects. The atmospheric eddies responsible for scintillation have a broad distribution in sizes, and is worst (a) with small apertures and (b) short timescales. We'll assume that you can't arbitrarily increase your aperture size, so the only corrective method you can use is to make multiple measures and average them. You'll likely see rms errors on the order of a few to several hundredths of a magnitude when your exposure times are ten seconds or less. If the stars you're observing vary on much longer timescales than your exposure time (Miras and other bright giants being classic examples), then you should absolutely take several exposures, measure the magnitudes, and submit the averaged magnitudes as your result. Submitting every frame's magnitude serves no scientific purpose.

This leads naturally into point (3) on optimizing observing cadence. Different classes of variable stars vary on different timescales, from milliseconds to millennia. Your observations should be optimized to the type of variability you want to search for, and you should also realize that some kinds of variability may be beyond the reach of your equipment.

As an example, take the case of a slowly–varying star and lots of photons. Bright Miras in the AAVSO program are examples of these. Nearly all of the well–observed Miras in the AAVSO archives are easily measurable by CCD observers (with filters) throughout almost their entire range of variation; there are hundreds of Miras that spend most of their time brighter than V=14–15. The question then is how often to observe? The simple advice we give to visual observers — no more than once every 1–2 weeks — is equally good for CCD observers. A somewhat more sophisticated answer would be to take a few sets of observations — 3 or 4 exposures in each of your filters — on a single night, and then average together the resulting magnitudes in each filter. You'd then submit the averages rather than the individual magnitudes, and you'd submit them as groups of magnitudes so that a researcher would have not just magnitudes but colors. How often you should do that depends on the star, but in general for periodic stars it is good to have between 20 and 50 observations equally–spaced throughout the period of variation of the star. If the period is 500 days, that's one night every 10 days at most. If the period is 100 days, that's no more than once every two days (and should really never be more than once every 4-5 days).

Some observers don't do this, and there are some egregious examples in the AAVSO International Database where observers were doing intensive time series of a Mira as if it were a rapid variable. Those data are not technically wrong, but they are largely wasted effort, and for the most part aren't useful for researchers in that form. [The only possible use of such data would be to look for rapid variations not typical of such stars, as might be caused by accretion onto an unseen companion.] Usually, an observer can make a more useful contribution if they take a few observation sets of one star, then move to take similar data on several other stars. There are plenty of variables in need of coverage, and a conscientious CCD observer could potentially create some wonderfully useful data sets for lots of stars.

Sometimes, you may encounter the exact opposite case — you may have a faint object that varies rapidly, and you're starved for photons (unless you have an enormous telescope). As an example of this case, look at one night's observations of the eclipsing polar CSS 081231:071126+440405 by AAVSO observer Arto Oksanen:

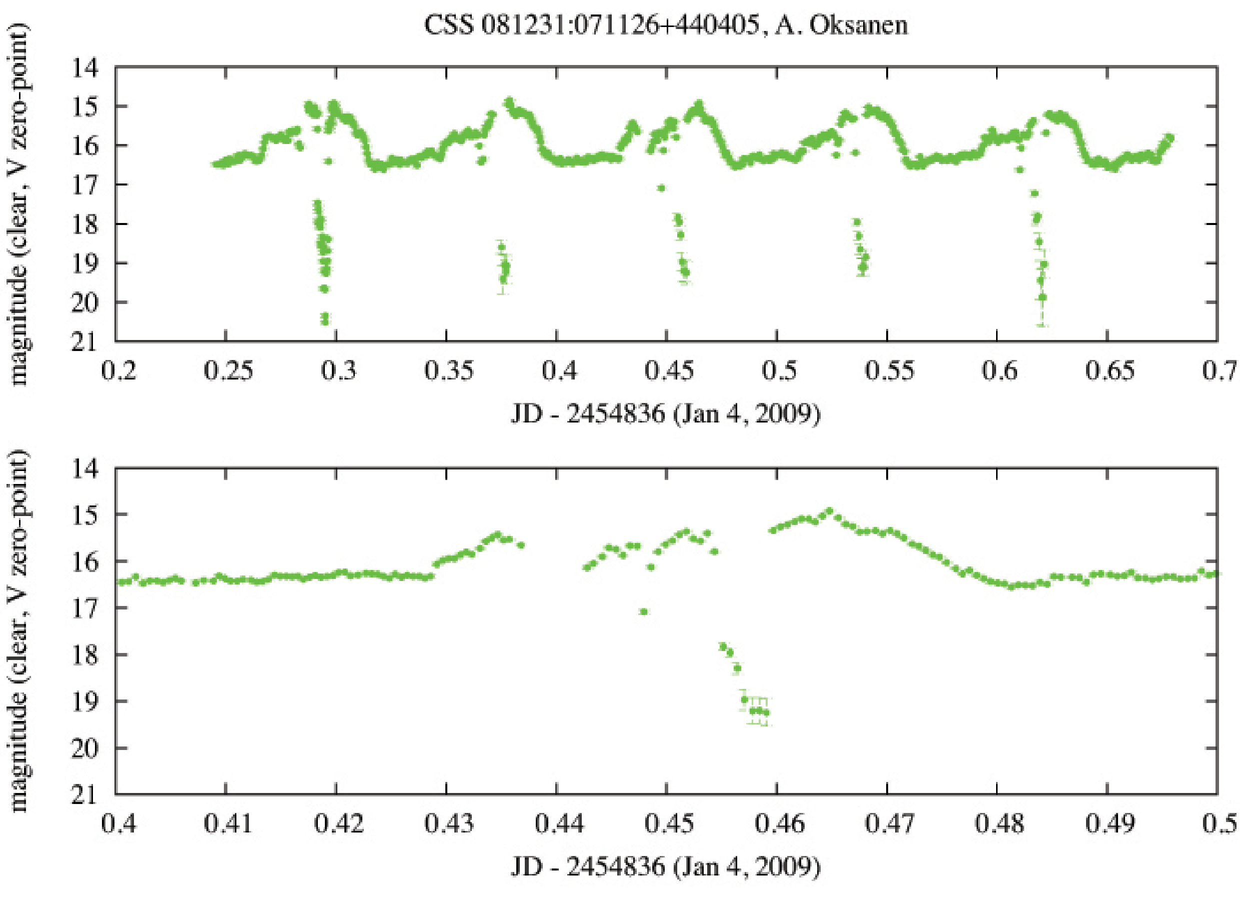

Figure 7.1 – Unfiltered time series of an eclipsing AM Herculis–type cataclysmic variable. Note that the error bars are very small, and note also the number of times observations are made. The observing cadence is approximately one observation per minute, including both exposure and chip readout.

These data were taken through a clear filter using a 0.4–meter (16–inch) telescope. When the star is between 15 and 17, the photometric uncertainties are around 0.015 to 0.02 magnitudes, which is well below the overall amplitude. Equally important is that the observing cadence is around one observation per minute. The orbital period of the star is just over 117 minutes, and so the observing cadence provides ample coverage throughout the orbital cycle. The result is that most of the orbital variations of this star are very well measured, and the overall light curve looks great.

The only time when it starts to be problematic is during the extremely short, deep eclipse, when the star goes below magnitude 20. First, the eclipse entry is extremely sharp — only a few seconds — so it isn't possible for an observing cadence of 1/minute to resolve that feature. Second, the eclipse is very deep — more than three magnitudes — so the eclipse causes the added problem of losing signal to noise. Uncertainties on the eclipse magnitudes approach 0.3 magnitudes, more than ten times larger than during the bright part of the orbit.

In this case, there's really nothing you can do to improve either temporal resolution or signal–to–noise during the eclipse — you're limited by the aperture of your telescope and the number of photons you're detecting, and there's no astrophysical reason to either shorten or lengthen exposure times. Shortening the exposures to improve temporal resolution would make the photometry too noisy to be useful, while longer exposures would simply smear out the eclipse leaving you with only a few data points during that interesting feature. This is an extreme case, but the number of interesting, faint stars like this is only going to increase as large–scale surveys like LSST begin finding new stars. For the more general case where you might have some options, simply be aware of the kind of variability you might see, and think ahead of time what your exposure times and cadence should be.

This is also a good example to raise the question of whether you are better off observing without a filter. Although we covered filters separately, they're relevant here to a discussion of timing because filters all lower your overall signal, and thus impact your exposure times and signal–to–noise; some filters may lower your signal so much that you can't make useful observations with them using your equipment. There are two principles to remember here:

1.) If the target is bright and you can get good signal to noise with an appropriate exposure time, you should always use filters. (Note that "good" will be defined by your project goals, but > 20 is a reasonable value.)

2.) If the target has very red colors, you must use filters unless there is some overriding reason where unfiltered photometry is useful (e.g. transient searches and gamma–ray burst afterglows). If you cannot use a filter on a known red target, you are better off observing a different target.

In this case, the object is very faint at times (with eclipses below magnitude 20), so you are definitely photon–starved. The variations are also relatively rapid, so you want to keep them as short as possible. But the most important reason you can forgo using a filter is that this star is very blue like most cataclysmic variables. If you were to take a spectrum of this star, you'd find the continuum is relatively flat, and doesn't change very much with wavelength. In this case, broadband variations match variations measured through filters reasonably well, and unfiltered observations are a good compromise that gets you slightly higher signal to noise and/or shorter exposure times at the expense of spectral information that, in this case, isn't as important as the other information you get.

Exceptions

Every rule has exceptions, and the guidelines for observing cadence and exposure time are no different. The most important thing to remember from the discussion above is that your exposure times have to be sufficient to detect the behavior you're searching for, and your observing cadence also has to match the timescales you want to cover. There may be research projects that look for behavior different from what is normally expected for a given variable star class. One example could be the discovery of an extrasolar planet transit in a longer period variable like an M– or K–giant. You might normally observe such a star once every several days, but a transit might vary on timescales of minutes to hours. You have to make observations with a much faster cadence in that case. In general such cases are rare, and usually happen when a star is already known to be special in some way (for example a Mira variable in a symbiotic system). You can certainly take high–cadence data to go exploring for interesting phenomena yourself, but realize those data will rarely be used as–is. You should consider examining your high–cadence data yourself offline, then averaging them and submitting the averaged data to the AAVSO archives rather than the individual points.

One more caution about Mira stars: do not make unfiltered observations of Miras, semiregulars, or other red variables in general. Unfiltered observations are really only suitable for "blue" stars (with B-V around 0.0). For red variables, your CCD is likely sensitive in the near–infrared, and red stars will be much brighter than you might expect them to be. You'll probably find occasional examples of someone reporting "CV" magnitudes for a Mira or semiregular star that are two or three magnitudes brighter than both visual data and filtered CCD data. Such observations really are wrong since the "CV" bandpass is very misleading to researchers. You might be tempted to observe very faint Mira stars without a filter in order to provide coverage at minimum, but the spectral properties of such data are so poorly constrained that they will not provide researchers with much useful information, and may actually cause more confusion than anything else. If you don't have filters for your CCD camera, you should avoid nearly all types of red variables, and restrict your work primarily to cataclysmic variables. Again, exceptions might be very faint transients like gamma–ray bursts.