

Chapter 5: Photometry – measuring images

Now that you have a set of carefully calibrated CCD images, it is time to measure the brightness of the stars you have captured. This is the process known as *photometry*. As with image acquisition and calibration, there is software available to do most of the hard work, but it is important that you understand it and use it properly or your results may not be scientifically useful.

Since there are a variety of software packages available, including the AAVSO's own photometry program – VPhot, this guide will not attempt to delve into specifics of how to use any particular program. Instead it will focus on concepts and techniques common to all of them, which will help you to produce good data.

What is differential photometry?

There are two kinds of photometry that are commonly done in astronomy:

- differential photometry – in which the magnitude you derive for the variable star is compared to the magnitude you derive for stars of known brightness in the near field at the same time, so that a “standardized magnitude” for the variable can be determined.
- all-sky photometry – a more complicated procedure in which the star magnitudes are derived directly using the results of nightly calibration of your system and current atmospheric conditions using a set of standard stars outside the field-of-view.

Only differential photometry will be covered in this guide because it is far easier and yields excellent results. It is also much more forgiving when observing conditions are not ideal. For example, if a thin cloud passes through your field of view when you are taking an image, chances are good that it will affect the magnitude of the comparison stars you measure as much as it will affect the magnitude of the target star. The magnitude *difference* between them will therefore be nearly the same and your results may be unaffected.

Here are the steps involved with performing differential photometry on your images:

1. Check your images
2. Identify the stars
3. Set the aperture
4. Choose the check and comparison stars
5. Measure the magnitudes
6. Determine the uncertainty

1. Check your images

Although you might have done this before, a visual inspection of each image can save a lot of time and frustration. Look for clouds, airplane or satellite trails or cosmic ray hits that could contaminate any of the stars (both the target and comparisons) you wish to measure. If you've taken a set of time-series images of the same field, you can examine them all in sequence to look for changes over time.

Double-check all of the stars you are measuring to be sure that none of them are saturated. Remember that just because you may not see blooming from a star in your image, doesn't mean that it can't be saturated. One way to see if a star image is saturated or not is to examine a point spread function (psf) plot of the star's brightness profile (see sidebar). If it looks like the top of the curve is flat, chances are good that the star has saturated the detector and there will be no way to derive a good magnitude for it. If you have not yet determined the linearity of your camera it would be a good idea to do so (see Chapter 3, page 16). With practice, you will get a feel for the best exposure time to use for your images based on a star's magnitude and the filter you are using.

Examples of some of the problems you might see when inspecting your images are shown on page 40.

2. Identify the stars

Study your images carefully — especially in crowded fields or in cases where the stars you wish to measure are very faint. It is not uncommon for a close companion or nearby star to be confused with the variable star you wish to measure, particularly when the companion is brighter. A large-scale (zoomed in) chart should always be consulted when you are imaging a field that is new to you so you can make sure that there are no hidden surprises, and you observe and analyze the right star.

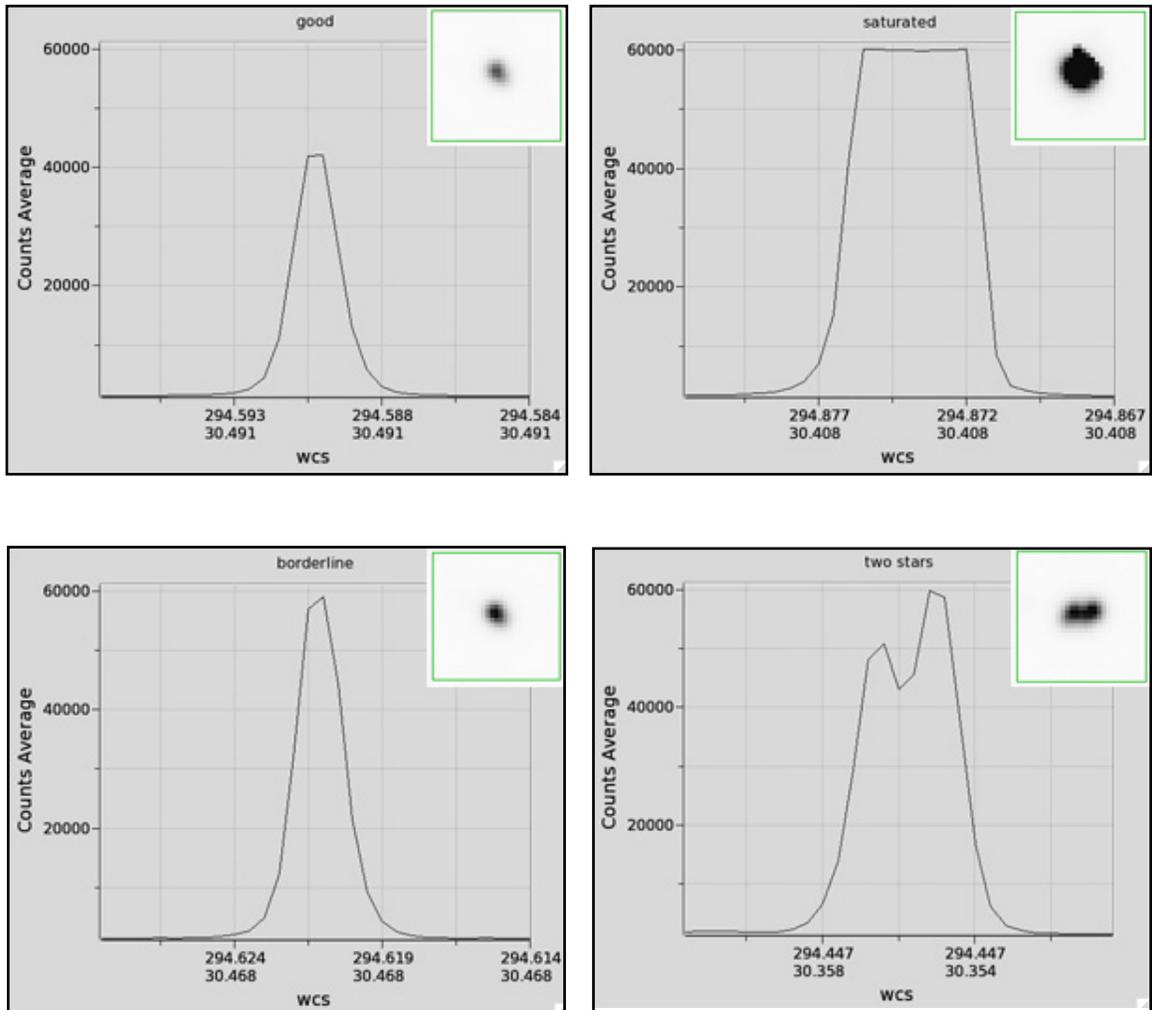
Depending on which software package you use, star identification will either be done automatically or you will have to do it yourself using your charts. In either case, it is important to check to be sure that the variable and comparison stars are correctly identified. Astrometry software is good but not perfect! It can be thrown off by defects in your images or misidentify stars with close companions.

If your software does not import comparison star sequence information from the AAVSO, you will have to do this yourself. The best way to get the information you need is to use the AAVSO Variable Star Plotter (VSP) to make a chart and get a Photometry Table. Using the chart, you can identify the comparison stars and record the published magnitudes for each filter color in the appropriate places. Using a DSS image in your chart can also be very helpful.

InfoBox 5.1 – The PSF Plot

Your photometry software should provide a way for you to make a point spread function (psf) plot of a selected star from your image. Generally this will be a two- or three- dimensional plot of the ADU count for each pixel versus a cross-section or radial cut through the star as seen on your image.

Such a plot can be very useful in determining whether a star in your image is saturated or perhaps blended with another star. Below are some sample psf plots (created using DS9) along with a close-up of the star being measured from the image.



3. Set the aperture

Strictly speaking, photometry is simply the measurement of the amount of light energy received per unit time. In this guide, we will concern ourselves only with the method known as *aperture photometry*, so named because we measure the strength of light in little circles or apertures, centered on individual stars in our image.

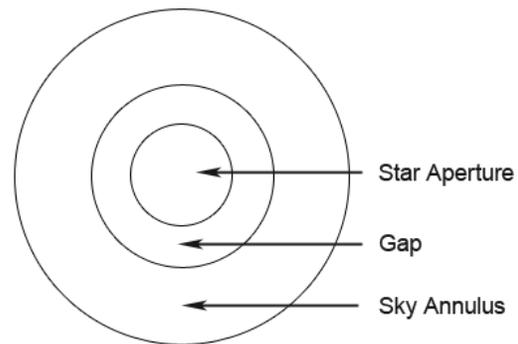
Two other ways in which photometry can be performed include *point spread function (PSF) fitting* and *image subtraction*. These techniques are useful for making measurements in very crowded fields but since both are very complicated and are rarely included in commercial CCD software packages, they will not be covered here.

The aperture is comprised of three parts as seen in the diagram:

Star aperture (or Measuring aperture) – this is the innermost circle, which surrounds the star you are measuring.

Gap – this is simply a space between the signal circle and the sky annulus.

Sky annulus – the outer ring that is used to capture information about the sky background.



The software package you use will probably create these circles automatically as soon as your image is loaded. However, you should have some control over the size of the each ring and may need to make small adjustments to suit your image or avoid problems. One important rule to remember is that *you must use the same sized set of rings for every star in the same image*.

Here are some other suggestions and guidelines regarding the size of the aperture rings:

- The diameter of the star aperture should be 3 to 4 times the rough average FWHM of all the stars you wish to measure. Your software should provide a way for you to determine the FWHM. (FWHM, or “full-width at half-maximum” is defined in Chapter 3, page 17.)
- Make sure it looks like the brightest star you are planning to measure fits completely within the star aperture. If the aperture is too small it won’t measure the star completely. If the aperture is too large, you increase the chance of including other faint stars in it.
- The diameter of the inner circle of the sky annulus should be about 5 times the average FWHM (or about 10 pixels across).
- Adjust the outer ring of the sky annulus if necessary. A bigger sky annulus yields better signal-to-noise ratio (SNR) but it is good to avoid field stars if you can.

- If there is no way to avoid “contamination” from field stars in the sky annulus, don’t panic! Your software may be able to remove their contribution automatically; consult your software manual to see if and how this is done.

4. Choose the check and comparison stars to use

This is a very important step because you will get different results depending on which comparison stars you use. In general, the more comparison stars you use, the better, since any errors or slight variability will be averaged out. However, it is important that you inspect the comparison stars you plan to use and select them with care to be sure that you have eliminated the ones that will give you worse results.

If at all possible, please use AAVSO comparison star sequences. Many software packages will allow you to load them automatically. If not, you can find the recommended comparison stars for each field by using the AAVSO chart plotting tool (VSP) and requesting output in the form of a “photometry table”. The table will give you the position of each comparison star along with its magnitude and the magnitude error in each bandpass.

AAVSO sequences have been carefully designed to use stars for which magnitudes have been determined very accurately, are known not to vary or have close companions, and are of a color similar to the variable. The other advantage is that by using a standard set of comparison stars, your results should compare more favorably with those of other AAVSO observers when your data are combined in the AAVSO International Database. Researchers using your data will like that.

Here are some guidelines to follow when choosing which comparison stars to use:

- Try to select comparison stars close to the target and not near the edges of the image where they could be distorted.
- The comparison stars should be similar in color to each other, but not necessarily to the target star.
- Don’t use red stars (many of which are themselves variable) or very blue stars. A good rule of thumb is to pick sequence stars that have (B-V) colors between +0.3 and +1.0, with (B-V) of +0.7 being a good mean value. But do realize you will be limited to whatever stars appear in the field, and you may not have much of a choice.
- Pick comparison stars that are similar in magnitude to the target star.
- Be sure that none of the stars you select have companions.
- Choose comparison stars with a signal-to-noise ratio (SNR) of at least 100.
- Choose stars with similar magnitude errors, preferably all less than .01 – .02
- Ensure that none of the comparison stars you choose are near the saturation point in your image.

Check stars are important in that they can be used to determine if any of your comparison stars are varying or if other problems may exist with your image. A check star is simply a star of known brightness that doesn't vary which can be treated in the same way as you treat your target star. You should be able to compare the magnitude you determine for it with its published magnitude (in the same color) and the results should be very close. The check star should be as similar in color and magnitude as the variable as possible and it can be chosen from the list of available comparison stars in the same field as the target.

If you are processing several or many images taken of the same field on the same night (time series) it is a good idea to plot the magnitude of the check star versus time. If all goes well, the result should be a straight horizontal line. If your check star's magnitude varies, then something is wrong. Could a cloud have passed by when you weren't watching?

5. Measure the magnitude

In most modern software, with the entry or acquisition of comparison star data, a click of your mouse will provide you with the magnitude of your target. It is good, however, to understand what is going on within your software to accomplish this (especially if you have older software that is not as automatic).

The first step the software takes is to measure the *instrumental magnitude*. This is simply a number related to the count of the photons (or ADUs) captured within the aperture. By subtracting the instrumental magnitude of a comparison star from the instrumental magnitude of the target star, you get what is known as a *differential magnitude*. Here is the formula:

$$\Delta v = v_{\text{measured}} - c_{\text{measured}}$$

Where Δv is the differential magnitude, v_{measured} is the instrumental magnitude of the variable star and c_{measured} is the instrumental magnitude you just measured of the comparison star.

In order to make your observations more useful to the scientific community, you now need to convert from a differential magnitude to a *standardized magnitude* by adding to it the published magnitude of the comparison star like this:

$$V = \Delta v + C_{\text{published}}$$

Nearly all software packages available today, will allow you to perform what is known as *ensemble photometry*. What this does is to compare one-by-one the variable star with each comparison star you selected. Using the equations above it will then compute the standardized magnitude of the target star based on each comparison star and return the results as a weighted average of all these

values. You are left with one standardized magnitude for your target star, which is generally less error-prone than it would be if only one comparison star was used. If it seems like one of the comparison stars in the ensemble is noisy or has a problem which is adversely affecting your results, try removing it from the ensemble and re-compute the average again.

Please note that we are using the convention that lower-case letters stand for instrumental magnitudes, upper-case, italicized letters (like V) are for standardized magnitudes, and non-italicized upper-case letters are for magnitudes which have been *transformed*. We'll explain transformation in the next section, but to summarize quickly: you may have taken an image with a standard Johnson V-filter, but you must perform some additional calculations to place your measured "v" magnitude onto the Johnson V system with the highest accuracy. We'll show you how in Chapter 6.

InfoBox 5.2 – A note about magnitudes

The magnitude system dates from the second century BCE, and is attributed to the Greek astronomer Hipparchus. It is a logarithmic system where brighter stars are assigned smaller magnitudes. The system was developed to classify naked-eye stars, but has been adapted in the telescopic age to measure optical brightness for many kinds of astronomical objects. There is a direct relation between magnitudes and fluxes: five magnitudes difference in brightness corresponds to a multiplicative factor of 100 difference in flux, meaning that each magnitude corresponds to a factor of approximately 2.5 in flux. Because the magnitude scale is logarithmic, ratios of fluxes can be expressed as differences in magnitudes. The relative difference in magnitudes between two objects with different measured fluxes can be obtained with the following equation:

$$\text{mag}_1 - \text{mag}_2 = -2.5 \log_{10} (\text{flux}_1 / \text{flux}_2)$$

For a longer discussion, see the AAVSO website: <http://www.aavso.org/magnitude>

Your software will probably convert measured fluxes (number of ADU within your measurement aperture) to instrumental magnitudes for you, but be aware that it may use arbitrary zero-points for these instrumental magnitudes. This may lead to strange-looking (but otherwise perfectly legitimate) instrumental magnitudes like "-12.567". Such instrumental magnitudes are fine as long as *all of the stars are measured with the same instrumental zero-point*. This is because the zero points cancel each other out when differential magnitudes are calculated.

6. Determine the uncertainty

The magnitudes that you measure only provide part of the information of your observation. Every legitimate piece of scientific data comes not only with a measurement, but also with an *uncertainty*, which tells the researcher who uses your data how well constrained your measurement is. Therefore, it's important that you accurately calculate and submit the uncertainty in your magnitudes along with the magnitude itself.

Your measurement uncertainty will contain both a *random* component and a *systematic* one. Random noise includes things such as photon noise (which is proportional to the square root of the number of photons your camera receives), and thermal noise in your CCD detector. These noise sources need to be characterized, but very little can be done to reduce them, and they put a lower limit on your uncertainty. Systematic uncertainties are related to your instrumentation, and can include things such as the way that your measurement apertures influence your output magnitudes, or whether you have uncertainties or errors in your flat fields or in the magnitudes you use in your comparison star values. We will not go into a detailed discussion of the theory of uncertainties here, but we recommend the AAVSO's CHOICE course *Uncertainty about Uncertainties* and the accompanying notes by Aaron Price for further discussion. We'll simply limit ourselves to *how* to do this.

The easiest but not ideal way is to let your CCD software do the work. Most CCD software will either return an uncertainty in magnitudes or will tell you the signal-to-noise ratio (SNR or S/N). A handy approximation is to assume that the uncertainty in magnitudes is $1/\text{SNR}$, so that a stated SNR of 50 yields an uncertainty of 0.02 magnitudes. The reason we say this is not ideal is (a) the SNR will be calculated just for each image that you measure, and will not tell you anything about noise from non-photometric conditions for example, and (b) you have to trust that the software is doing this correctly. Most software *now* does a reasonable job of doing this, but historically that was not always the case for all software. As always, look at your results and see if they make sense.

Beyond that first method, there is no one best way to calculate uncertainties, but it depends on what and how you plan to observe. If you're making multiple measurements of a star during a single night (e.g. a time series run), you can use the variations observed in either your variable or your comparison and check stars to estimate the total photometric uncertainty. There are two choices here. If you know the variable isn't changing in brightness on short timescales (a Mira star, for example), then you can calculate the magnitude of the variable on each frame, and then calculate the standard deviation of those measures of the variable to give you the uncertainty. (Note: ideally, for a slowly varying star, you would go a step further and combine all of the measures of the variable made on a single night into one magnitude instead of submitting the entire time series.) If the variable does change on short timescales (a cataclysmic variable, for example), then you should

instead obtain the uncertainty from multiple measures of your comparison or check stars instead. In all cases, you compute the uncertainty using the equation for the standard deviation, σ :

$$\sigma = \sqrt{\sum(x_i - x)^2 / (N-1)}$$

where x_i are the individual magnitudes, x is the average magnitude, and N is the total number of measures being averaged. You would then report σ as your uncertainty. Note that if you are using the standard deviation of a comparison or check star for this test, you should use one that has similar brightness to the variable.

If instead you only take one image per filter of a given field, you're limited to calculating uncertainties based on the information contained in this one image. For the case of a faint star, *you must use the CCD equation*:

$$S/N = \frac{N_{star}}{\sqrt{N_{star} + n(N_{sky} + N_{dark} + (N_{readnoise})^2)}}$$

where N is the number of photons received from each of star, sky, dark current, and the readnoise of the CCD, and n is the number of pixels in your measurement aperture. Although this may look complicated, it is simply a modification of the case where you're measuring the uncertainty due just to photon noise. To see this, imagine that N_{star} is much larger than any of the other terms. In that case, the CCD equation approaches the limit of the square root of the number of photons received.

Note two things here. First, note that N in the above equation is the number of photons, rather than the number of ADU, which is what your CCD measures. This introduces a slight modification to the equation for ADU which includes the *gain*, G :

$$S/N = \frac{N_{ADU} \times G}{\sqrt{(N_{ADU} \times G) + n_{pix} \times ((N_{ADU,sky} \times G) + N_{dark} + (N_{r.n.})^2)}}$$

Second, note conveniently that *you can use the SNR value from your software instead of the full CCD equation* to estimate the uncertainty in the case of a star whose brightness is well above both the sky background and readnoise.

The next best option with single-image photometry is the case where you have multiple comparison stars available in the frame. In this case, you can measure all of the comparison stars along with the variable, calculate the magnitude of the variable obtained using each of the comparisons, and then calculate the standard deviation of all of these magnitudes. This will take into account the intrinsic uncertainties in both the variable and the comparison stars.

The CCD equation is universal, but is also somewhat involved to calculate, since you have to measure all of these things individually, and it doesn't provide information about other sources of uncertainty beyond what was present in that specific image, such as sky conditions. However, with a single image, it's the best you can do and you should use it, especially in the case where you're working with faint stars and lower S/N.