AAVSO Guide to CCD/CMOS Photometry

with Monochrome Cameras

Version 1.0 (July 2022)

AAVSO
185 Alewife Brook Parkway, Suite 410
Cambridge, MA 02138

Phone: +1 617 354-0484
Email: aavso@aavso.org

ISBN 978-1-939538-62-8

Preface

This *AAVSO Guide to CCD/CMOS Photometry with Monochrome Cameras* has existed in a number of different forms, including *The AAVSO Guide to CCD Photometry*, since AAVSO observers first began using CCD[1] cameras in the 1990s. Although CCD cameras continue to be the major technology tool for photometry, manufacturers are shifting from CCD to CMOS technology. These cameras have provided a dramatic increase in the amount of image-derived data, now accounting for more than 80 percent of all data submitted to the AAVSO International Database (AID). The decreasing cost and increasing usability of consumer-grade camera systems, particularly CMOS cameras, will only increase the future amount of image-derived data.

The ease with which data can be obtained and extracted from a camera does not necessarily indicate the ease with which science can be done with that data. This version of the *Guide* is an update that broadens its scope to include both CCD and CMOS cameras and addresses some of the needs of beginning photometrists (including a Quick Start Guide in Appendix F), while maintaining the AAVSO's overarching mission to enable science by amateur observers and not just to generate data. While still covering elementary use of CCD and CMOS cameras and reduction of data, this *Guide* is aimed to help you generate scientifically useful data by generating precise and accurate standard magnitudes. The AAVSO emphasis on the scientific value of our data is reflected in this *Guide's* guidance for new observers to actively learn, self-assess, and improve. Ultimately, the scientific utility of your data matters far more than how much of it you collect.

This *Guide* is intended to serve beginner and intermediate imaging observers who want to use their equipment to obtain variable star photometry of the highest quality. It is possible to create photometric data with a small telescope and camera that equal the quality of data from many professional observatories, and there is in principle no difference between data from an amateur observer and data from a professional. We aim to help you obtain the best data possible. We'll tell you how to get data out the back end of your system, and we'll also explain why and how to do this the right way so that your data provides researchers with useful information.

The *Guide* will always be a work in progress, and we depend on the community to help us develop and document best photometric practices. You may find things in this document that are out of date or unclear. Please give us feedback as to what works for you and what doesn't.

Please send any feedback or suggestions to aavso@aavso.org. Clear skies,
Sara Beck, *AAVSO Technical Assistant, Science Team*
Arne Henden, *AAVSO Director Emeritus*
Matthew Templeton, *AAVSO Science Director 2010-2015*
Mark Munkacsy, *Coordinator, CCD/CMOS Photometry Guide Development*

1 Underlined words are defined in the Glossary, beginning on page 130. When a glossary word is used multiple times, only the first occurrence on each page is underlined.

Acknowledgements

The AAVSO and the authors of this *Guide* would like to express their sincere thanks and appreciation to the following people who have contributed in some way to this publication:

For their time and programming skills in creating the transformation tools:

Gordon Myers
Richard Sabo
George Silvis

For help, encouragement, and valuable feedback:

Tom Calderwood
Dave Cowall
Tim Crawford
Blake Crosby
Thibault de France
Allan Hollander
Robert Jenkins
Ken Menzies
John O'Neill
Sebastián Otero
Miguel Rodrigues
Phil Sullivan
Gary Walker
Ed Wiley

Much of the CCD content in this *Guide* was originally published as *The AAVSO Guide To CCD Photometry* (Version 1.0) in September 2014.

Revisions to that Guide are:

Version 1.1: February 2015

- March 3, 2015: update to Chapter 6
- January 11, 2016: minor change to "Deciding how many images to make", Chapter 4.
- February 9, 2016: minor change to "Submitting Observations to the AAVSO", Appendix C.
- February 26, 2016: corrected and replaced plots in "InfoBox 3.1 - How to determine the linearity of your camera", Chapter 3.
- March 2, 2017: modified Chapter 4 relating to calibration images, and reformatted equations in Chapter 5.

The AAVSO Guide To CCD/CMOS Photometry With Monochrome Cameras includes new content on CMOS cameras as well as revisions to pre-existing CCD content. While it could be considered as Version 2.0 of the original *AAVSO Guide To CCD Photometry*, it is presented as Version 1.0 of a new publication, published July 2022.

Table of Contents

Chapter 1: Becoming a photometrist

If you own or have access to a telescope with a digital monochrome astronomical camera, whether CCD or CMOS, you can obtain scientifically useful variable star data. Similarly, a Photoelectric Photometer (PEP) can be used, but it is different enough in practice that it has its own AAVSO Observers' Guide. It is also possible to use digital color cameras (DSLR or one-shot-color (OSC)) to obtain useful variable star data, as described in a separate AAVSO Observers' Guide.

The AAVSO supports several observing modes, with digital observing and visual observing (either aided or unaided eye) being the two most popular. Both kinds of observing have strengths and weaknesses, and each has its place in variable star astronomy. This *Guide* can help a novice observer become a better imaging photometrist[2] using cooled monochrome CCD or CMOS cameras. This is critical for our mission, because the quality of data we receive from observers impacts the quality of science that researchers will do with it. A digital astronomical camera is capable of obtaining very good variable star data, and like most scientific instruments, it is also capable of obtaining very bad data. We want to help you aim for and produce good data.

Our imaging observer community is drawn from a number of different populations. Some former (and current) visual observers made the leap to camera-based observing. Some people who used astro-cameras for astroimaging wanted to do more than astrophotography. Some people may use remote or shared facilities to obtain astronomical observations and want to maximize their value. Some people may have come across an article on variable star observing and thought *I want to try that!* They may have taken the leap directly into digital observing.

For the sake of simplicity, this *Guide* assumes you have a passing knowledge of astronomy — you should know, for example, how stars move across the sky during the night, what astronomical coordinates are (Right Ascension and Declination), and what the magnitude of a star means. We also assume that you already know how to set up and operate your telescope, how to connect your camera to a computer, and how to use the camera and telescope software to operate them. You should be able to turn on your telescope, point it to a star field or have the telescope point itself, focus, and take an image with the camera. If you've taken a picture of a star field, cluster, nebula, or galaxy with your telescope that you're reasonably satisfied with, you're all set with what you need to know. If you're just starting out with a new instrument, learn the basics of how to operate it and have some fun first. Get a good feel for the telescope and camera, and especially how to use it to take images that track properly.

Along with this you should be comfortable with the software that came with your telescope and camera, or at least have a copy of the software manuals. Most commercial imaging software will have

2 Underlined words are defined in the Glossary, beginning on page 130. When a glossary word is used multiple times, only the first occurrence on each page is underlined.

everything you need to prepare your images for scientific usability. Later in the *Guide* we'll talk about how to extract data from those images, something that can be done with many commercial software packages or with the AAVSO's own *VPhot* software. More on that later.

In general, you do not need to be a mathematician, engineer, or astrophysicist to obtain good data. Some background knowledge of mathematics (including algebra) *will* be assumed; perhaps more valuable, though, is self-discipline and attention to detail. While many of the calculations required for photometry can be automated within a spreadsheet, you need to understand what goes into that spreadsheet and what comes out. You will need to develop the habit of examining (and questioning) your results carefully, assessing whether they make sense every time you submit an observation.

Finally, we'll assume that you have an interest in both variable stars and creating good quality scientific data. Familiarity with variable stars before you start doing photometric observing would be great, at least at the level of knowing what a variable star is in comparison to a non-variable star, but you can learn as you go along, and we'll cover the basics of what variable star photometry is and why we do it in the next chapter. Many of our best imaging observers got their start as visual observers, and we encourage everyone to pick up this guide's sister publication, the *AAVSO Manual for Visual Observing of Variable Stars.*

Note that obtaining good data may involve making some mistakes and (crucially) learning from them. Taking **very** good data is complicated, and it requires effort and discipline. It's easy to get bad data from a camera; it's fairly easy (or at least straightforward) to get good data. It's harder to get **great** data whether you're an amateur or a professional, but we're confident that you **can** do it if the circumstances allow; otherwise, we wouldn't be writing this. It's okay to make mistakes, and if you learn from them, you'll be on your way to collecting good data.

1.1 Scope of This *Guide*

This *Guide* is written for users of dedicated digital, cooled, monochrome astro-cameras, based on CCD or CMOS technology. These cameras are normally sold without a lens kit, so you will be mating the camera to a telescope. If you are using a DSLR camera (traditional or mirrorless), we point you to the AAVSO *DSLR Observing Manual*, which uses terminology better aligned to those cameras.

1.2 What Is Photometry?

When we "observe" a variable star, we're measuring the amount of light that the star appears to give off at a given moment. We repeat that measurement as often as needed to completely track all of the light variations. If our individual measurements are repeatable (precise) and representative of the true value (accurate), we can then make physical models that try to explain why the brightness changed in that manner. Your task as a variable star observer is to make good measurements so that researchers can make good models. The better your data, the better their models. The process of measuring the light from a star is called photometry, and a person who does this is a photometrist. We're hoping you'll become one, and a good one at that, once you work through this *Guide*.

There are a number of details about how you make that measurement that can improve researchers' chances of making realistic models. Sometimes you can generate excellent data for some stars just by pointing your telescope at the target, taking one or more images, and processing the images with simple methods. Sometimes you will spend many hours a night on just one star, taking images over and over again as quickly as you can. You will often use one or more filters to measure light within well-defined wavelength ranges (bandpass). You will even spend time measuring specially selected non-variable stars to better calibrate your observations. All of these and more are involved in turning your observations into useful data.

Photometry is the measurement of starlight intensity (i.e., magnitude or flux) by any means. While this *Guide* will teach you how to do photometry with a digital camera, that isn't the only instrument capable of doing this, and your ultimate goal isn't to be a "good imaging observer", it's to be a good *photometrist* who is using an imaging camera. There's a difference. Nearly everyone can saw a piece of wood in half, but that doesn't make them a carpenter. A digital camera will produce numbers that get turned into another set of numbers inside your computer, and perhaps another set of numbers in your analysis software, spreadsheet, and so on. Those numbers aren't photometry unless the process is correct. Don't focus on the technology, focus on the purpose. Your goal isn't to produce *numbers*; it's to produce *knowledge* that may lead to understanding. We'll show you why and how, starting now.

1.3 A Word About Terminology

This *Guide* sometimes walks a fine line between using words that are precisely correct and words that will be most understandable to readers. This is particularly true in discussions about error, accuracy, sources of error, measurement precision, and reporting uncertainty. In most cases, the precisely correct term is "uncertainty." The *Guide's* use of multiple terms is not intended to be misleading; the precise reader should consider all of the different terms and phrases we use to be synonyms of "uncertainty."

1.4 How To Use This *Guide*

- Chapter 2 (and Appendices A and B) provides background on the science of photometry and establishes the fundamentals of differential photometry, the measurement of star brightness by comparison with a sequence of stars that have known brightness.
- Chapter 3 talks about hardware and software and will help you identify the decisions you will need to make about *how* to employ your photometry system to make the science

measurements you're pursuing.

- Chapter 4 explores what happens at the telescope – the images you need to capture, particularly including images needed for calibration.
- Chapter 5 (paired with Appendix C) explains how to extract measurements from your images.
- Chapter 6 talks about transforming your measures into what is called the "standard system" to correct for differences between observers due to different filters, cameras, and telescopes.
- Chapter 7 (supported by Appendix E) is about *self assessment*, how you determine the accuracy and precision (repeatability) of your measurements.
- Chapter 8 then revisits some of the topics of Chapter 2, exploring how the equipment and processes of the prior chapters can be tailored to obtain specific scientific data.
- Appendix F offers a step-by-step primer on getting started, for those who have never done photometry before.
- The Glossary defines words that are underlined in this *Guide*.

Chapter 2: An Introduction to Photometry

2.1 What Photometry Measures

We start with an image, a completely hypothetical image, but perhaps something that looks similar to the following:

And let's assume that this is a very well studied patch of sky, where every star visible in this image has a well-known and documented brightness.

Further, just for a few pages, let's talk about brightness in terms of luminous flux, instead of using the more common term, magnitude. Luminous flux is a measure of brightness that describes how many photons per second are arriving from that star. It doesn't matter how spread out that light is: our measure of luminous flux (or just "flux" for short) is the sum across all the pixels capturing light from that particular star.

So we now choose two stars in that image; let's call them star "V" and star "C". Even though we know the brightness of all of the stars in the image, we're going to pretend that we don't know the brightness of the star we've chosen to call V (short for "variable"). Using this image, we want to determine how bright V is. And since we're pretending, let's pretend that we know that C is a constant star with a brightness that does not vary.

We start by dropping fixed-size circles around V and C:

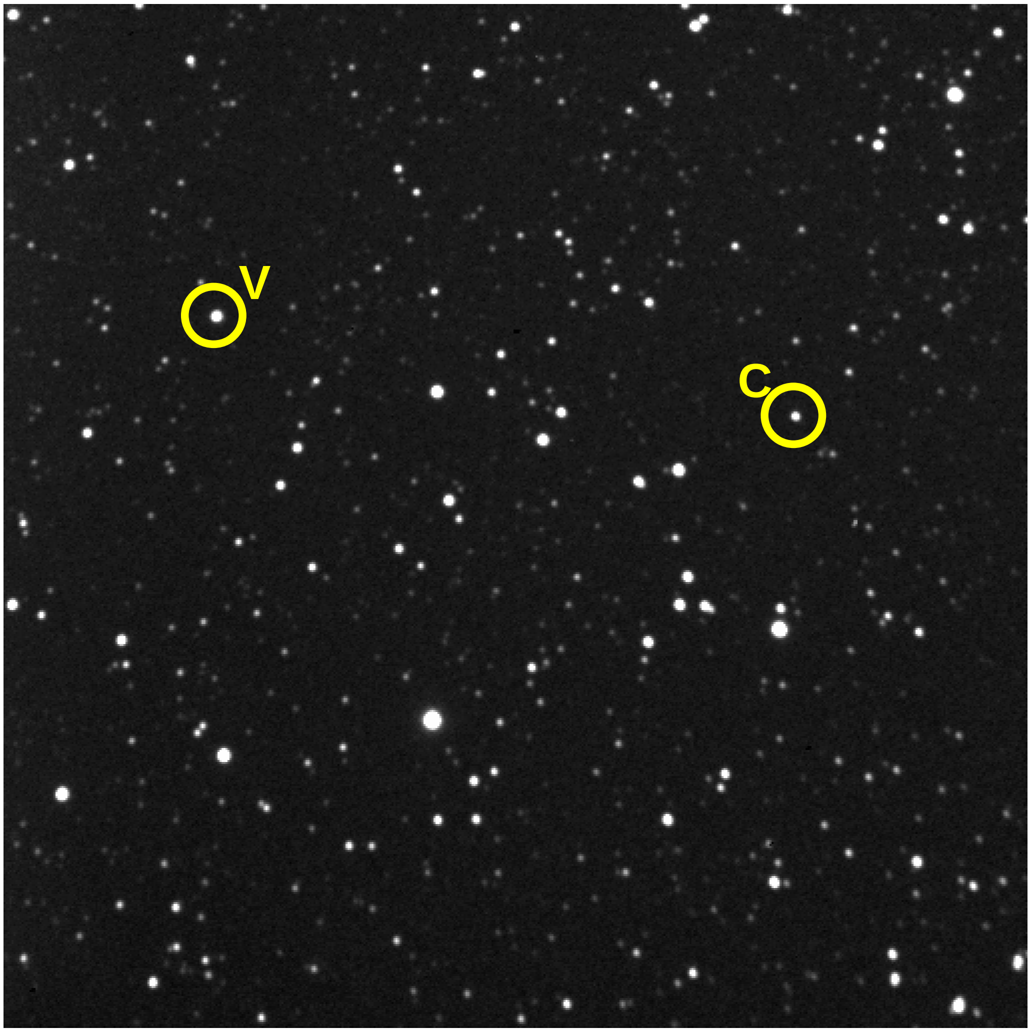

The image consists of pixels, and each pixel has a "count" reflecting the amount of light gathered during the exposure. If we look at the pixels around V, we find the following pixel values (the numbers are in units that we call Analog-to-Digital Units (ADU)), and we've superimposed the yellow circle around V onto the numbers:

487	492	494	496	487	488	493	502	492	489	494	495	490	496	494	480	495	494	480	481
490	493	491	492	485	489	495	495	491	489	485	495	488	495	491	495	494	497	491	490
492	491	489	491	491	485	490	498	491	496	495	496	497	494	497	492	495	493	492	497
494	485	486	491	500	494	494	493	498	493	491	496	494	495	492	492	492	500	489	493
493	493	494	492	483	494	494	[illegible]	504	502	501	[illegible]	501	490	491	491	489	491	495	495
489	490	496	494	493	[illegible]	501	505	512	517	515	518	510	[illegible]	492	494	495	493	498	496
487	502	500	499	498	505	505	536	582	605	579	545	517	510	495	500	492	494	496	496
489	490	498	495	[illegible]	500	533	672	967	1137	943	692	555	529	[illegible]	494	492	491	498	494
493	494	497	495	496	513	596	1041	2281	3144	2173	1036	628	541	512	498	493	492	492	490
494	497	492	496	508	524	633	1301	3619	5529	3576	1321	692	547	510	499	499	493	486	492
490	483	494	499	502	510	585	1039	2646	4124	2831	1131	642	533	501	492	490	489	493	490
492	487	492	493	[illegible]	502	537	676	1178	1616	1299	766	575	522	[illegible]	494	490	492	492	491
486	491	492	492	[illegible]	500	509	540	614	688	659	580	526	510	501	491	489	490	495	494
494	493	494	494	498	[illegible]	492	507	513	521	529	515	503	[illegible]	494	494	494	496	492	493
490	494	497	493	498	501	[illegible]	[illegible]	498	502	506	[illegible]	[illegible]	487	491	490	490	499	490	492
485	482	483	496	495	491	496	499	495	495	499	496	498	502	494	495	494	499	496	498
487	489	491	497	497	497	489	489	492	495	498	488	491	500	497	495	495	490	500	500
495	492	493	487	486	500	494	487	485	489	488	493	486	492	493	488	500	491	493	500
490	486	484	492	493	491	494	491	499	494	490	496	488	496	497	497	502	497	495	492
493	486	491	492	494	495	490	495	500	497	497	493	492	496	494	497	499	496	498	496

Using these counts and the known brightness of C, we can establish the brightness of V.

The light from each star extends over more than one pixel, and we make the circles big enough to encompass very nearly all of that light, while excluding the vast majority of the image from consideration. The circles cover some fixed number of pixels, and for both V and C we add up all the counts in their respective circles, and we'll call those numbers "sum_of_C_circle_counts" and "sum_of_V_circle_counts." These counts represent the total amount of light received, flux times the exposure time.

Flux depends on a lot of things: the size of our telescope, the thickness of the atmosphere we're imaging through (affected by both our telescope's altitude above sea level and the viewing angle in the sky), the amount of dust, smog, and moisture in the air, the amount of dirt on our telescope's optics, and the efficiency of our camera in collecting photons. But, since we're pretending, let's pretend that all those factors are identical for all the stars in this one image.

We will use a simple ratio to relate the brightness (flux) of C and V:

$$\frac{flux_V}{flux_C} = \frac{\text{sum\_of\_V\_circle\_counts}}{\text{sum\_of\_C\_circle\_counts}}$$

Because the brightness of C is known and we were able to measure the pixel counts in our image for stars C and V, we can use this equation to describe the brightness (flux) of V in terms of the brightness of C. For example, if the sum of the V aperture counts is half of the sum of the C aperture counts, we would say that the brightness of V is 0.5 times the brightness of C. In this (pretend) case, we know the brightness of all the stars, so we know the true ratio $\frac{flux_V}{flux_C}$, so we can compare the ratio we computed with what is already known in order to assess how well we did.

Why jump through these hoops to indirectly determine the magnitude of V? Well, a direct (absolute) measurement of V's flux in any meaningful way requires that we know (and correct for) all those things we mentioned earlier: air thickness, telescope size, dirt on the optics, atmospheric absorption, and so on. That's actually rather difficult. What we have done is called differential measurement, and it is capable of very good results with a whole lot less work. Almost all AAVSO photometry is gathered this way.

2.1.1 Complication 1: Repeatability (Precision)

Let's make lots of images of these stars. When we do this, though, we discover very quickly that no matter how perfectly identical we make a series of images of this field, the number of counts that we find in our circular apertures is *different* each time. Compared to the total number that's in our aperture sum, the variation isn't huge, but it's definitely there.

Since this is a pretend experiment (and since we do know the actual brightness of the star V), we can look at how well our measured ratio compares with what we know to be the true ratio. And we realize that if we average together all the ratios that we calculated from all the images we took, the average comes closer to what we know is correct than the ratio that we see on any single image.

It turns out that this variation is caused by the same effect that causes one hundred tosses of a coin to result in something a little bit different than 50/50. Sometimes we see 49 heads, sometimes 52, sometimes something else – but it averages fairly close to 50/50. The more times we toss the coin, the closer to 1:1 becomes the ratio of heads to tails.

The root cause of this is that we are actually capturing individual photons of light in our camera, and the arrival and capture of a photon is a discrete event that follows very specific statistics, sometimes called Poisson noise or shot noise.

When we measure brightness, we find that our results become more repeatable (have greater precision) when we capture more light. We can do this through longer exposures, a larger telescope, more exposures (stacked images), or by making multiple measurements and averaging them together.

2.1.2 Complication 2: The Background Isn't Zero

When we look in detail at our images, we find that the background isn't completely black: the background counts aren't equal to zero. Were we to perfectly cover up the aperture of our telescope and take an exposure, every pixel would have some counts. (The background in our image earlier was around 490 counts.)

These counts would mess up our magnitude computations but they are relatively easy to eliminate. They mess things up because:

$$\frac{flux_V}{flux_C} \neq \frac{flux_V + background}{flux_C + background}$$

If we performed more experiments, we would discover several things about the background:

- The longer the exposure, the more background counts we see.
- However, even zero-length exposures have nonzero background counts.
- The background counts are temperature-dependent, higher in the summer and lower in the winter.
- Even if we lock the telescope and camera into a light-tight box, there is still background being recorded.

We identify two distinct sources of those background counts. (Later on, we will identify a third.) One we call bias: the background counts that show up even for zero-length exposures. The other we call dark current: the background counts that scale with exposure time and with temperature.

That first source, bias, is not affected by the exposure length – it is purely a function of reading the data – and the amount of bias will vary somewhat from pixel to pixel.

The dark current is due to thermal effects. Most cameras that we use for photometry incorporate a built-in thermoelectric cooler to reduce this effect. But, even a well-cooled camera has residual heat in the imaging chip. That heat causes charge to accumulate in the pixels, and the longer the chip runs the greater the charge that will pile up. This extra extra charge will also vary from pixel to pixel.

We deal with this complication by creating what we call Bias and Dark Images. In these special images, we allow no starlight, so our camera captures only the bias and dark signals. Once we have these images, we can subtract the bias and dark image values, pixel-by-pixel, from the values in our image of C and V. This effectively subtracts the background level. We call this process calibration, and you will do it on *every* photometry image that you create.

2.1.3 Complication 3: Light Pollution and Skyglow

The sky in which our stars are embedded is not perfectly dark (this is the third background source). Our pixels pick up light from the stars combined with background light. This light comes from chemical reactions in the earth's atmosphere (called airglow), and from very faint stars near V and C. And that's only if you're lucky: most of us live with light pollution that scatters off of air molecules and aerosols and then enters our telescopes. Like dark current and bias, the background contaminates the pixels and throws off our measurements. We need to get rid of it.

In the diagram above, the yellow circles we placed around V and C are called "apertures." These apertures will gather some amount of background in addition to the light from the star. If we could measure the amount of background light we could subtract it. We estimate the background in the aperture by placing an annulus around our circle:

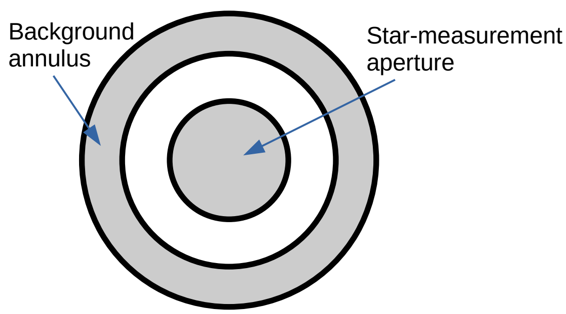

It surrounds the aperture but does not overlap. (There may or may not be a gap between the aperture and the background annulus, but there should never be pixels shared between both.) Modern photometry algorithms will detect any stars in the annulus, and will exclude their pixels from the computation of the background. We take the counts in the annulus and divide them by the number of pixels therein. This gives us an average background expressed as counts-per-pixel. We then subtract that amount from the counts in each aperture pixel. This is how we compensate for background skyglow and light pollution.

2.1.4 Complication 4: Aperture Size Matters

Out of curiosity, we do some experiments and discover to our dismay that changing the size of the yellow circle (the star measurement aperture) changes our results. The problem isn't that we get totally outlandish flux ratios for C and V; instead, we find that if our apertures are too small, then measurement repeatability seems to suffer. Similarly, if we make the aperture too big, we see the same thing: rising fluctuations in the measurement of our flux ratio.

Further, for some choices of star C and V, as the aperture diameter gets bigger, the ratio seems to consistently diverge from what we know is the true ratio. Sometimes the ratio becomes too big and sometimes too small.

Further experiments narrow down this behavior to two separate effects. First, when we make the measurement aperture too big, we start picking up light from adjacent stars in addition to the light from V or C. This shifts our ratio, and we use the phrase faint star contamination to refer to the cause.

Second, even when there is no contamination, making the aperture too big or too small degrades what we call our signal-to-noise ratio (SNR). This is a measure of how much starlight we've captured in the aperture relative to the amount of pixel noise. (More precisely, SNR is the ratio of the total number of counts in the measurement aperture from the star to the amount of random noise in our measurement. A large SNR is good (100 would be excellent) and a low SNR (below 10-20) provides poor results.) We find that when we make the diameter of the star aperture 3-4 times the apparent width of the star's image we get the best SNR. (Later on (Chapter 5), we'll provide more information on how to measure the width of a star's image.)

2.1.5 Complication 5: Location in the Frame Matters

Another upsetting discovery is that the location of C and V in our image seems to matter. We notice that the closer either star gets to an edge or corner, the bigger the discrepancy between our measured flux ratio and what we expect. But sometimes even just moving the stars a few pixels will noticeably shift the flux ratio.

After a few more experiments, we come to an unsettling conclusion: all pixels are not created equal. As we travel across the sensor, we find that the pixels have slightly different sensitivities to light. Those differences may be small, but enough to throw off the measurements we want to make. And the problem gets worse: the telescope optics do not uniformly illuminate the whole camera sensor. Pixels near the center of the camera are receiving the most light. Pixels in the corners of the camera receive substantially less. Pixels in-between get an in-between amount of light. The effect is called vignetting, and it is especially apparent in old portrait photographs, where the corners of the picture are actually black from the lack of light. Vignetting is frequently a significant problem when a focal reducer is used on the telescope. Also, dust on telescope optics will rob some pixels of some light.

Ideally, we want a camera with uniform sensitivity and a telescope with uniform illumination across the whole sensor. But since these don't exist, we must compensate with our real-world equipment using a process called flat-fielding.

The process of flat-fielding (usually considered a part of the calibration process we described earlier using bias and dark images) starts with a measurement of how much the combined sensitivity and illumination affects each pixel. This is easier than it sounds. We do it by taking a special exposure of a target that has perfectly even light across the whole field-of-view. This is usually either a white screen inside the observatory or the sky at twilight. The point is to image a field that ought to give the same counts in every pixel.

We can look at that Flat Image and compare the brightness of every pixel to the average pixel's brightness. Some will be brighter, and some will be dimmer. By dividing each pixel's actual value by the value of the average pixel (a process called *normalization*), we get a number (hopefully, somewhat close to 1.0) that describes how *this* pixel varies from an average pixel. We can now take our image with stars in it, and adjust the value in every pixel by dividing it by that number that was close to 1.0. Pixels that are somewhat too bright will be divided by a number a little bigger than 1.0 (reducing that pixel's value) and pixels that are somewhat too dim will be divided by a number a little smaller than 1.0 (raising that pixel's value).

If we perform this flat-field calibration carefully, we'll find that our flux ratio of V to C becomes much more independent of the stars' locations within the frame.

2.1.6 Complication 6: Bright Stars and Faint Stars

Not surprisingly, we discover that we have problems with very bright stars and very dim stars.

As stars get dimmer, we have more and more trouble finding the star's center in order to properly place the measurement aperture. Further, we see that the ADU counts at the star's center don't change much from the sky background ADU. Earlier, we found the concept of signal-to-noise ratio useful. This is another situation where it helps. Once the SNR drops below 10 or 20, the star begins to get lost in the general background noise, and our flux ratio changes significantly with every image. We find that we can correct for this (by improving the SNR) by stacking multiple images or by using longer exposure times.

But we find that we also have to be cautious about too-bright stars. Beyond a certain brightness, the flux ratio begins to deviate dramatically from the correct value as a result of saturation within either the camera sensor's pixels or within the supporting electronics. We can see that it doesn't matter which star is too bright: the ratio falls apart if *either* star C or V is the bright one. Further, if we have a CMOS camera with adjustable gain, the camera's gain setting seems to have a dramatic effect on the point when the ratio begins to fall apart.

This concept of staying in the camera's linear region turns out to be critically important (much more important for photometry than for general astrophotography). As you work through the rest of this *Guide*, you will find lots of advice on how to determine where the limits are for your camera and how to recognize when you've exceeded the linearity limit.

2.1.7 Complication 7: Color

Our final complication comes from the realization that star color is affecting the accuracy of our flux ratio. And the effect is sometimes dramatic, with the flux ratio being off by a factor of two for some telescope/camera systems and some star combinations.

After some experimentation, we realize the following:

- When stars V and C have the same color, the flux ratio is closest to correct.
- By using colored filters that restrict the wavelengths of light reaching our camera's sensor, we can somewhat reduce the flux ratio errors when the colors don't match.
- With filters in place, the remaining flux ratio errors are not random, but instead seem very repeatable; indeed, we can come up with mathematical corrections based on the difference in color between stars C and V to reduce much of the remaining flux ratio errors.

2.1.8 Summarizing the Complications

This section has walked through the basics of differential photometry. It has covered the fundamental process of measuring star brightness within circular apertures after first calibrating our images to remove variations in pixel sensitivity and background levels. It introduced the idea of signal-to-noise ratio (SNR above 100 is good; less than 10 is poor), the importance of staying within the camera's linear range (avoiding saturation), and the importance of color.

This chapter continues with a more detailed look at color, since color-related considerations are ubiquitous within this *Guide*, affecting equipment, observation planning, the selection of appropriate comparison stars, and post-analysis color correction.

2.2 Color and Photometry

The study of color is an integral part of becoming a good photometrist. There really are three reasons for this:

- Comparing the relative brightness of stars that have different colors can be complex, because their relative brightness can be different across different swaths of the spectrum,
- Star color can change over a variable star's cycle; measurement of that color can provide important physics information about what's going on, and
- Measuring brightness through multiple filters can give coarse insight into a star's spectrum by measuring brightness at a few different portions of the spectrum.

We start with the first reason: the challenge of comparing brightness between two stars. Because all stars radiate at multiple wavelengths (see Appendices A and B), making a statement about the brightness of a star means we need to agree on what portions of the spectrum we are including in our measurement. In our case, because we are talking about differential photometry, the question we wrestle with is: How does color affect our measurement of the relative brightness of two stars?

Stars emit energy across the entire electromagnetic spectrum. Your telescope and camera can only capture a small slice of that total. Further, your equipment captures light of different colors with different efficiency. The overall *color response* of your camera, filters (if any), telescope, and local atmosphere combine to determine the overall color response of your system, which profoundly affects the apparent color of the things you look at. Photometrists often use color response graphs to describe the color response of the parts of their system as well as the response of the system as a whole. Figure 2-1 is an example of a color response graph; this shows the color response (sometimes called *spectral response*) of one particular camera.

Although graphs like Figure 2-1 are commonly available for filters and for cameras, they are rarely made available for the telescope itself. Further, as telescopes age and mirror/lens coatings slowly degrade, the spectral response curve for a telescope will change over time. The atmosphere itself introduces another time-varying component of our overall color response curve, since water vapor and particulates have a large effect on color-related light absorption in air. Thus, every photometry system will have a different spectral response.

Does this matter? Well, let's think our way through it, starting with a case that's so simple that it doesn't happen in real life:

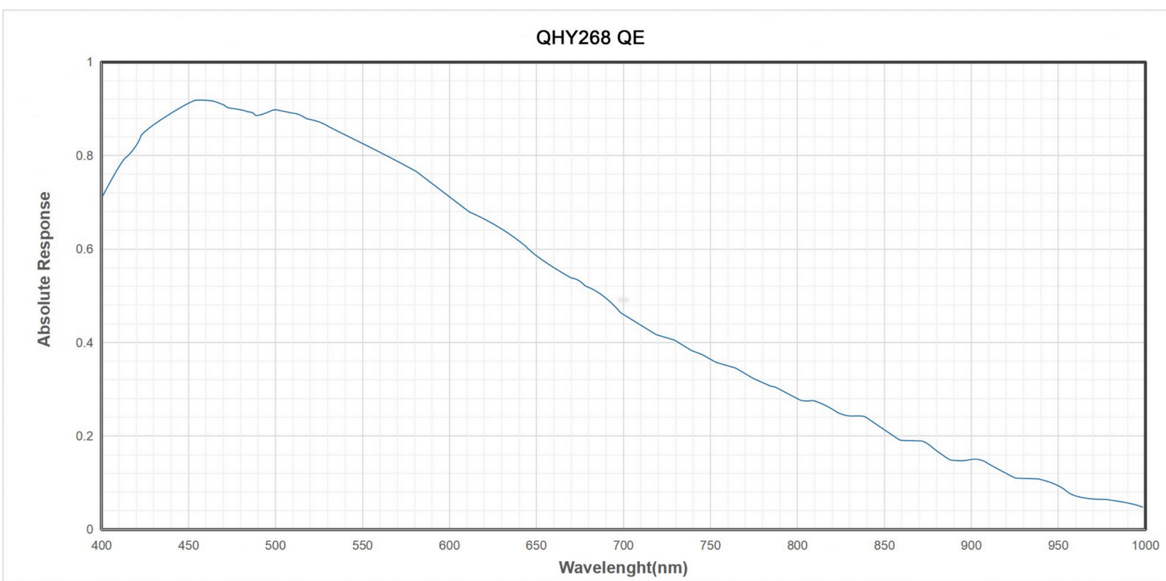

Figure 2-1 — Spectral response of the QHY268M camera. (Graph by QHYCCD, used by permission)

When observing two stars with *identical* colors (spectra), observers will always agree on relative brightness, regardless of the color response of their systems.

Although it's often difficult (impossible?) to find two stars with identical color, an important consequence of this simplification is that we can minimize color-related inaccuracy by using comparison stars that are close in color to our variable star. (Section 6.3 will introduce the concept of a color index and what it means to compare the colors of two stars.)

So, since stars come in a variety of colors (spectral profiles), let's assert something that's easy to say, but difficult to do:

If two stars *don't* have identical colors (spectra), then observers need to match the color response of their systems in order to agree on the stars' relative brightness.

This is at the crux of the photometrists' challenge: how to match the color response of multiple observing systems. Over the years, astronomers have defined a series (now counted in the dozens) of standard color systems that specify an overall system color response curve. With these standard color systems, a photometrist aligns their system to one of the standard systems, and then informs anyone using their data that this data conforms to the response curve of that standard color system. This process of aligning your system is called *transformation* and is covered in Chapter 6.

As a first step toward alignment with a standard color system, you can purchase and install a color filter designed to match a standard response curve. Multiple commercial vendors offer photometric filters that align to common standard color systems, such as Sloan, Cousins, and Johnson. However:

Alignment with a standard color system is not all-or-nothing. Proper filter selection can get your system's overall color response close to that of a standard color system, but applying calculated correction factors (color transformations) will get your data closer to the standard. The color transformation factors are unique to the color response of your particular system and will change as your equipment ages. The closer you are to a standard color system, the smaller your systematic measurement error (bias/accuracy).

Effectively, there are three different levels of correction to align with a standard photometric color system:

1. **No correction at all:** measurements made without filtering, with a clear filter, or with a completely non-standard filter (e.g., a blue-blocking exoplanet filter). Although these measurements can't be directly combined with measurements from other observers, there are several cases where unfiltered observations are useful: when the source is known to have a

neutral color (where all wavelengths are equally bright, as is typical in hotter objects like cataclysmic variables in outburst), when the object is very faint and simply detecting the source has great value (as in gamma ray bursts), or where timing or period determination is the overriding scientific goal (as with eclipsing binary stars). In fact, clear filters are commonly used for situations where temporal variation is of primary interest and where targets are so faint that any loss of flux through standard filters is unacceptable.

2. **Untransformed standard filters:** measurements made with filters that do align with a standard photometric color system, but have not been adjusted using color transformation coefficients. Residual errors due to color can be as high as several tenths of a magnitude.
3. **Transformed into a standard color system:** measurements made with filters that align with a standard and with subsequent adjustments applied to the measurements based on transformation coefficients calculated against standard calibration star fields. With care, residual errors due to color can be reduced below about 0.03 magnitudes.

This *Guide* encourages you to pursue this third level of color correction; it's an important part of being a good photometrist. However, the AAVSO also recognizes that important science can be conducted (with careful planning) with the other levels of correction and that the cost of photometric filters may be prohibitive for some observers. We encourage *all* observers to work toward the third level as they pursue their self-assessment and improvement programs (Chapter 7).

2.3 Why We Perform Photometry

The brightness of a variable star changes measurably due to physical processes inside, on, or around the star. There are many variable star classes, and each represents a different way that a star can vary. Stars may change in size, shape or temperature over time (pulsators), they may undergo rapid changes in light due to physical processes around the star (accretors and eruptives), or they may be eclipsed by stars or planets in orbit around them (binaries and exoplanets). The key is that something is physically happening to the star itself or in its immediate vicinity. You may see a star twinkle in the sky (scintillation), but that variation is due to the Earth's atmosphere. Variable stars change all on their own, independent of anything happening here on Earth.

Different kinds of stars vary on different timescales. Some stars may take weeks, months, or years to undergo changes that we can detect. Others take days, hours, minutes, seconds or much less. Some stars vary regularly, and in their variations we can see patterns that repeat over time. Other stars undergo chaotic change that we can never predict exactly. Some stars vary the same way for centuries, while others, like supernovae, flare up briefly and then disappear, never to be seen again. And some stars do not seem to vary at all, at least within the precision of our equipment!

Variable stars also exist with a range of apparent brightness (how bright they appear to us) as well as a range of intrinsic luminosities (how much light they actually give off). A star may be intrinsically luminous, but if it is thousands of light years away, it will appear faint. Variables also have a range of amplitudes — how much their light changes over time. Some variable stars can vary by ten magnitudes or more, which is a factor of ten thousand in flux, a huge change! Some variable stars vary by a millimagnitude (0.001 magnitude) or even less, and their variations may be impossible for you to detect. There are innumerable variables in between, and there's no shortage of targets that you'll be able to do productive work on, regardless of the size of your telescope.

Variable stars are interesting for a number of different reasons, but ultimately, we study them because they're like physics laboratories. We can't go and touch a star or change it to study it, but if we can understand how the light from a variable star changes, we can learn more about how stars work. The same fundamental physical processes that operate here on Earth (e.g., gravity, fluid mechanics, electromagnetism, light and heat, chemistry, and nuclear physics) operate exactly the same all over the universe. By watching *how* stars change over time, we can learn *why* they change. Your observations provide the raw material that powers scientific inquiry. Scientists can speculate endlessly about why things appear and behave the way they do, but ultimately those hypotheses have to be tested in order to productively advance our scientific understanding. That's where observing comes in, and it's where you can make a valuable contribution to variable star science. If you give researchers valid and accurate data, they can make accurate models of the universe, and our understanding increases and improves. Conversely, if they have bad data, those scientists may make bad models, misleading us and hindering progress in the field. Figure 2-2, on the next page, is a good example of a light curve built with AAVSO data and informing the refinement of astronomical models. This is just one of many light curves that you can contribute to and study during your observing activities.

Variable stars are particularly interesting because they often tell us more about themselves than what we can get from a constant star. They can tell us something about the circumstances under which they formed, how they spend their lives, and how they eventually evolve and die. Learning more about what stars are and why they behave as they do gives us a more complete picture of our universe, both in the present and over cosmic timescales, providing insights into everything from planets and stars to galaxies and beyond. That's ultimately what all of variable star astronomy is about.

We repeat that the goal of photometry isn't the numbers that come out of the camera and your data processing, it's the *science* that you can do with those numbers. In order to do science, your results have to represent something physically meaningful, and have to be presented in a way that is useful for rigorous scientific analysis. That's our goal, and that's where we're aiming with this *Guide*.

2.4 The Photometry Workflow

In order to perform photometry, you will be assembling your own personal photometry workflow.

Photometrists speak of their workflow as the sequence of tasks that moves them from start to finish along the path of making a photometric measurement, as well as the tools (particularly software tools) they use along the way. Although different photometrists will subdivide their workflow in different ways, a complete photometry workflow will contain the following elements:

- **Planning Phase:** The activity that takes observing goals and decomposes them into specific targets on specific dates. It also includes exposure planning: deciding how many exposures will be needed (and how long each will be), along with an observing cadence. The planning phase verifies suitability of the field to the available equipment and to the comparison star sequence. The output of the planning phase is, not unsurprisingly, an observing plan.
- **Calibration Phase:** In Chapter 3 we will explore all the elements of the calibration process. Many calibration steps can be performed under cloudy skies (or even during daylight). The observing plan is used during calibration to choose camera settings and exposure times for the various calibration frames that are exposed and then combined. The output of the calibration phase is a set of master calibration frames to be used during the analysis phase.

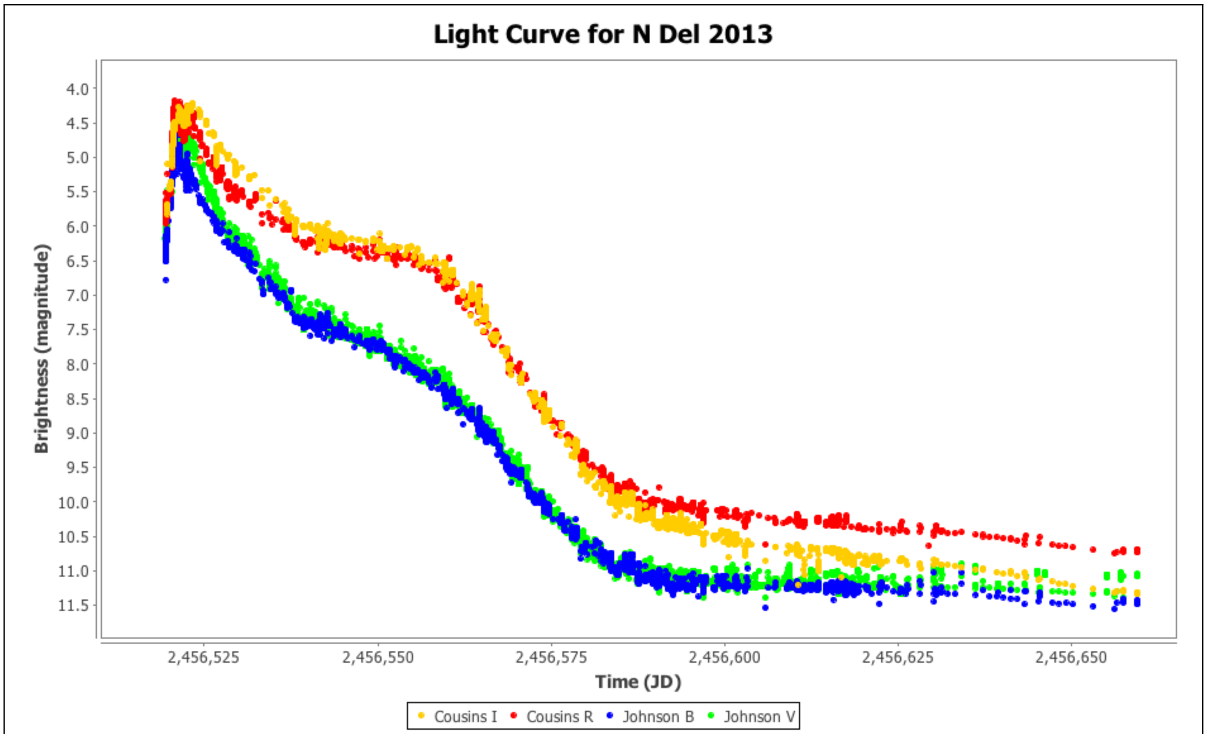

Figure 2-2 — Light curve of Nova Del 2013 (V339 Del) as plotted with VStar. Note how the overall brightness changes, but the relative brightness of each band also changes as different physical processes dominate in the nova evolution.

- **Science Exposure Phase:** Chapter 4 concentrates on what is done in the dark during an observing run. The observing plan is used as the master guide for the observing session, describing exposure times, the order in which filters are used, and image counts. Supporting observation activities such as guiding, focusing, pointing, and image framing are all performed during this phase to support the ultimate product of the phase: a series of images that are passed along to the analysis phase.
- **Analysis Phase:** Chapter 5 describes the analysis phase, which typically involves the manipulation of data in the form of images and numbers derived from those images. Photometrists strive to interconnect the software tools they use during this phase to minimize the amount of human "touch" involved in passing numbers from one analysis activity to the next. Each human interaction introduces the opportunity for human error; all those potential error points must be double-checked at some point in the workflow to detect the inevitable mistakes and omissions that happen when humans do anything.
- **Self-Assessment Phase:** This phase is somewhat unusual because it operates within two distinctly different time frames. One timeline concentrates on a single observing session, providing an independent check on the prior phases of the workflow. It both detects mistakes within the workflow and also estimates photometric uncertainty for the session. The uncertainty estimate provides yet another check against unreliable or incorrect measurement by alerting the photometrist to an observing run with measurements inconsistent with what's been seen on other observing sessions. The other timeline focuses on long-term photometric performance trends. This helps the photometrist identify improvements that are needed to the workflow itself, to observing technique, and to the parts of the telescope/camera systems.
- **Reporting Phase:** This phase takes the results of the analysis phase, formats them into the AAVSO Extended File Format, and submits them to the AAVSO. A final quality check is used to detect submission errors.

Section 3.5 will address the software that photometrists use, and will go into more depth about the flow of information through the process.

2.5 Photometry Quality and Observing Programs

AAVSO observers use a diverse collection of telescopes, filters, cameras, and software to perform photometry. The sky offers an even wider population of variable stars, varying over different timescales, exhibiting brightness ranges that span at least 4 orders of magnitude, and colors that range from ordinary to extreme. Some photometry systems are better suited to observing some variable star types than others.

This *Guide* encourages you to learn photometry: to understand *all* of the steps and the rationale for each. As you proceed on this journey, we will explain how the options at each step affect photometric quality (the accuracy, uncertainty, and precision of your measurements). Different observing programs have different photometric quality needs; this *Guide* will help you:

- Understand how to match your photometric system with your observing program.
- Understand the importance of knowing your photometric quality for various equipment configurations available to you.
- Create your own self-assessment program (Chapter 7) to build confidence in your assessment of photometric quality.
- Choose your observing programs in a deliberate way. You can't obtain good data for every variable star in the sky with any single system, regardless of its size or cost. However, there will be many objects that can be observed easily and effectively no matter what you have — just realize you should figure out what those objects are before you get to the telescope.
- Match your observing cadence to your target stars' behavior. Over-observing can be as seriously detrimental as under-observing. Section 4.6.3 will explore the relationship between the number of images that your capture, your reporting cadence, and the underlying behavior of the star your are observing.

The AAVSO International Database (AID) merges data from thousands of observers. Its value to researchers is directly influenced by the quality of the data it contains. Researchers turn to the AID because of its reputation for quality. *Your* observations have a direct bearing on the AID's reputation for quality. A single poor observation that results from improper technique, an awkwardly-placed cosmic ray hit, a numerical transposition error, vignetting uncorrected by a master flat, or a stuck filter wheel has a much more profound influence on the astronomical community than just a single bad data point; it chips away at the reputation of the entire AID. *Your* knowledge matters, *your* technique matters, *your* processing workflow matters, *your* personal checks and balances matter. And with a little bit of care and hard work, *your* photometry matters and *your* contributions can help improve the quality and reputation of the AID and its ability to support new discoveries and insights.

In the next chapter, we'll cover the very basics of the equipment and software you'll need to start doing image-based photometry. Every telescope (from a small refractor to a large reflector), every mount (from a German equatorial to a fork to a simple alt-azimuth), and every camera (from a 12-bit CMOS to a 16-bit CCD) will have its own peculiarities, but we'll cover the essentials that you should have when you go out to the observatory for a night of variable star photometry.

Chapter 3: Equipment and Software Overview

Since you are using this *Guide*, it is assumed that you already have a telescope, mount, camera, and all the associated equipment needed to do photometry. Therefore, there is no point in describing what equipment you should get, but rather how to make the most of the system you have. There are many different types of telescopes, cameras, and software packages. In this chapter, we mainly want to explain the things that all will have in common, and what will generally be required to get good data from any system. Consider this chapter as being less about doing photometry, and more about the critical step of preparing for photometry before you get out to your observatory to start observing.

3.1 Telescope and mount

Most telescopes can work well with digital cameras. Smaller telescopes, like a 70mm refractor, are good for imaging brighter stars. Larger diameter telescopes, like a 300mm reflector, help with the fainter variables where increased light-gathering is needed. Each telescope system may be subject to various optical aberrations such as coma, but today's telescopes, from the simplest telephoto lens to a modest refractor or to the largest and most sophisticated 0.5m reflector, can all support useful photometry. Cost will almost certainly be a key criterion for your choice, but keep in mind that virtually all of today's available telescope systems will permit you to start in photometry and collect useful science data.

Figure 3-1 – Two of the original AAVSOnet telescopes: BSM–Hamren, 65mm Astro–Tech AT–65EDQ (left) and Coker 30, a 30cm Meade LX–200GPS (right). Both have subsequently been updated.

One of the difficulties of using a large telescope with a small camera is that the field of view can be much smaller than what you may be used to through an eyepiece. In general, the smaller the focal length of the telescope, the larger the field of view, which makes finding the field and trying to capture all the comparison star candidates in the same frame a little easier. However, very fast systems (<f/3) may cause some optical aberrations and difficulties with coated interference photometric filters.

It is important that you try to reduce stray light from entering the system. This is generally more of a problem with reflectors. Remove your camera and look through the end of your telescope at the night sky. Look for reflections or glints of light off any internal surfaces. If you see anything more than the stars out the end of the telescope, your camera will pick that up as well, affecting your images. Consider trying to eliminate that stray light with paint or flocking material inside the tube.

Having a good mount for your telescope is very important for success, and observers often describe the mount as the most critical element of their system. A good mount brings simplicity and efficiency to the entire imaging process. Equatorials are generally preferred, because during medium and long exposures alt-azimuth mounts cause field rotation that is impossible to remove without a rotator. A fork mount will avoid meridian flips and can be deployed on an equatorial wedge. Whether you choose a German equatorial mount (GEM), an equatorial fork mount, or an alt-azimuth mount is a case of personal preference. All can work well with appropriate effort, although it is recommended to buy as high a quality mount as you can afford. It is important that the mount be well-aligned and track accurately. It will save time and frustration if the mount also helps find the target field with GoTo controls or allows your computer to take you to the field. Auto-guiders are not essential, but helpful both for longer exposures and time-series runs.

Finally, consider using an observatory to house your equipment. Although not absolutely essential for getting good data, some sort of permanent mount (with a way to protect it from the elements) will save time and frustration setting up and breaking down your equipment each night. Even a good, sturdy watertight box on rollers that you put over your mount will save hours of setup and alignment each session. With a more substantial structure, you will feel comfortable leaving your camera and computer attached and ready for use. There are many solutions, and they don’t have to be expensive.

As a warning, though, some observatories can raise camera and mount internal temperatures far above ambient air temperature. Direct solar radiation can heat the roof; inadequate ventilation can cause temperatures inside the observatory to soar during the day. The electronics in cameras, mounts, and computers can be damaged if power is applied while circuit board temperatures are excessive. After the observatory is opened, allow sufficient time for electronics *internal* temperatures to moderate *before* power is applied to the equipment. An inexpensive USB temperature monitor can be used to measure equipment temperature during the day to track and record peak temperatures.

3.2 Optical Coverage

Your camera and telescope work together to define how much of the sky you capture in a single

Roll-off roof shed *BSM—HQ's housing*

Figure 3-2 – Examples of observatories providing weather protection for permanently-mounted telescopes.

image and the size of the smallest sky detail that can be detected.

3.2.1 Field of View

The field of view (FOV) of your system is how much of the sky you capture in each image. It is important for you to know and understand this number and to design an observing program consistent with your FOV. It is a good idea to check your FOV against a star chart or your planetarium software to see if your field is indeed large enough to contain the variable star you wish to image as well as all the comparison stars you will need for photometry at the same time.

To calculate the FOV, use the telescope's focal length and your detector size in millimeters:

FOV = (57.3 × width/focal length) by (57.3 × height/focal length)

(FOV in degrees, focal length in mm, height & width of the chip in mm)

Below are two examples of systems using the same camera :

Camera: SBIG ST402 (KAF-0402), Chip Size = 765 × 510 pixels (6.9mm x 4.6mm)

Example 1: *Telescope:* Takahashi refractor, focal length = 530mm
FOV = Height: 57.3 * (4.6mm/530mm) = 0.5º = 30 arcmin;
Width: 57.3 * (6.9mm/530mm) = 0.746º = 44 arcmin
44' × 30'

Example 2: T*elescope:* Celestron 11" SCT, focal length = 2800mm
FOV = Height: 57.3*(4.6mm/2800mm) = 0.094º = 5.6 arcmin
Width: 57.3 * (6.9mm/2800mm) = 0.141º = 8.5 arcmin
8.5' × 5.6'

There is no "best" FOV. However, experience indicates that a FOV smaller than 15 arcminutes across can create problems finding good comparison stars in the same field as the variable. An excessively large FOV (more than about 45 arcminutes across) can challenge the photometrist by requiring differential extinction calculations to compensate for atmospheric differences across the large image. This challenge is exacerbated by observing through high airmass at low elevation angles.

As sensor sizes have grown (and camera costs have dropped), some photometrists have chosen to limit their working FOV to something *smaller* than what the sensor is capable of, using only the central portion of each camera frame. Examples of this include recent cameras using monochrome APS-C sensors that are a little bit too large to be fully illuminated with 1.25" photometric filters, leaving the corners of each frame unusable. Instead of replacing your existing 1.25" filters with larger filters, you can choose instead to crop to a subframe for your photometry.

Sensors with rectangular (not square) dimensions are common. Most photometrists prefer orienting the camera so the long axis spans the east-west direction. Because camera hardware and software defines the long axis as "width," this tends to be more convenient when working with VSP charts and the VPHOT displays, both of which keep east-west in the horizontal direction.

3.2.2 Image Scale (Pixel Scale or Resolution)

Another critical piece of information about your system is image scale or resolution. This is how much of the sky is captured by each individual camera pixel. The image scale of your system can be computed using this equation:

Image scale = (pixel size/focal length) × 206.265

(image scale in arcsec/pixel, pixel size in microns, focal length in millimeters)

You can get the pixel size from the manufacturer's specifications on your camera. The focal length of your telescope can also be expressed as f/ratio times the aperture in millimeters.

Knowing the image scale of your system is critical to matching your camera to the typical seeing conditions in your location. Simply use this equation:

Seeing = Image scale * FWHM

(seeing in arcsec, image scale in arcsec/pixel, full width half maximum (FWHM) in pixels)

In most locations, seeing averages between 3 and 4 arcsec, but it certainly varies from location to location and could be better or worse on any given night.

3.2.3 Seeing and Sampling

Ideal images of point sources made by optics have an intensity pattern called an *Airy disk*. However, in practice the light from stars (generally considered to be a point source) has to pass through the Earth's atmosphere which diffuses and expands the pattern into the *seeing disk* captured by your camera, which is significantly larger than the Airy disk.

When you inspect the intensity profile of a star in your image, you will see that it is spread across a group of pixels, with some brighter ones near the center surrounded by some dimmer ones. The edge of the star is indistinct (and, technically, extends to infinity). Collectively, the set of pixels making up the star image from your camera is called a seeing disk because atmospheric seeing conditions have a profound effect on the intensity of the light. To measure the size of a disk that doesn't have sharp edges, photometrists use the term Full Width, Half Maximum (FWHM). This is defined as the number of pixels that are filled to one-half the brightness level between the background and the brightest pixel in the star's image. In order to get the best photometry results, you should strive to sample so that the FWHM of your seeing disk is spread across two to three pixels. This will help optimize the signal-to-noise ratio (SNR) and improve the precision of your measurements.

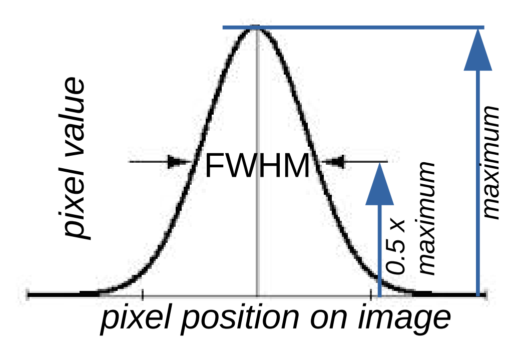

Figure 3-3 — Full Width Half Maximum point spread function

So how do you know if your system gives proper sampling of the seeing disk? The answer is simple: you measure it directly. Just use any of your well-focused images close to the zenith. Most camera control software has a tool for measuring the shape of a single star, including the size (FWHM) of the seeing disk expressed in pixels. This is your system's sampling of the star image. Measure several stars around the center of the image that have a good SNR but are not saturated. The FWHM may vary a bit across the image because of seeing effects and optical aberrations. It may also change over time as the seeing (scintillation) changes. However, the FWHM will be similar for all stars, bright or faint, in the image. (Your eye will often perceive the brighter stars to be bigger than fainter stars, but measure carefully and you will find that *all* stars have roughly the same FWHM.) You are just looking for an approximate number of 2-3 pixels per FWHM. (But, realize that the *full* profile of each star is far larger than this and spread across many more pixels.)

Often, achieving this goal won't be feasible or even possible, given that it is highly dependent on the seeing conditions and limitations of your equipment, but you may be able to tweak it a little. If you are averaging a FWHM of less than 2 pixels, you are probably *undersampling*. If the FWHM of your seeing disk is more than 3 pixels in diameter, you may be *oversampling*. Either situation could affect the accuracy of your photometry, although undersampling is much worse than oversampling. Fortunately, there are things you can do to remedy the situation.

3.2.3.1 What should I do if my system is undersampling?

The goal is to try to increase the size of the seeing disks on your images. One option is to defocus your telescope a bit, then increase the exposure time. If you choose to defocus, be very careful that other nearby stars are not close enough to blend together and affect the photometry. Also try to create flat frames (see 4.5) that are defocused the same amount, and take all your images with the same amount of defocusing (which can be very tricky!). Chronic undersampling suggests a bad match between your sensor and your telescope; you might find that adding a good-quality focal extender or Barlow could also help the situation somewhat if you can tolerate the smaller FOV. Or, consider purchasing a new camera with smaller pixels.

3.2.3.2 What should I do if my system is oversampling?

First of all, check the focus and make sure that the seeing disks are as small as possible. If the FWHM is greater than six pixels, you might consider using a focal reducer. Not only would it reduce your oversampling by increasing the image scale, but it would also give you a larger field of view. Binning is another option, and is explained in a later section.

Oversampling is not nearly as serious as undersampling, because with undersampling you actually lose information about the star by concentrating the star into too few pixels – there is nothing you can do during post-processing that will restore that missing information. Oversampling, though, is more about optimization of the system SNR; fine photometry is possible when somewhat oversampled.

3.3 Camera

Cameras range widely in quality, complexity, and cost, but most can be used quite successfully for photometry. The important thing is that you should get to know your camera well in order to get the most out of it. Then you can use that knowledge to set up your observing program appropriately.

Figure 3-4 provides a conceptual overview of what happens within the camera. During each exposure, incoming photons of light knock electrons "loose" within each of the sensor's pixels. The sensor's *quantum efficiency (QE)* describes how perfectly that happens. Some of today's mass-market cameras achieve QE exceeding 80% in certain wavelengths. The QE is wavelength-dependent, typically peaking within the visual spectrum and falling off at both ultraviolet and infrared wavelengths. Some cameras have better blue QE than red, while others perform better at red wavelengths.

When each pixel's number is read from the camera (the pixel's "ADU count"), the ratio of the number

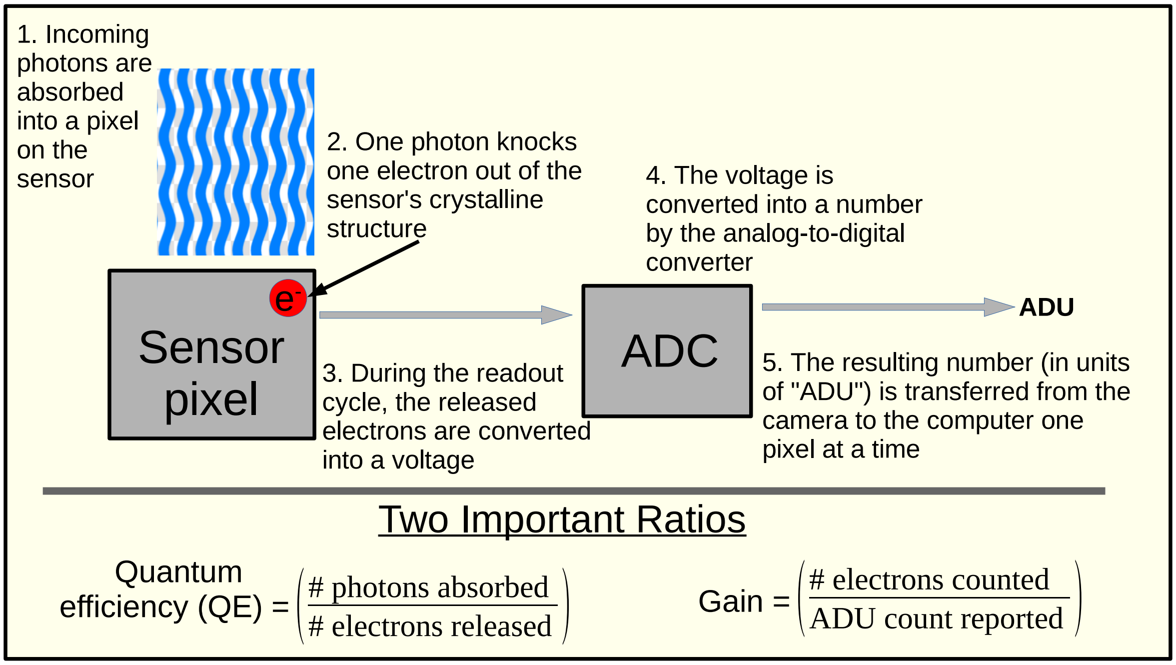

Figure 3-4 — Quantum Efficiency (QE) and System Gain

of electrons to the reported ADU count is the camera's gain. Some cameras provide a gain of one or less, meaning that *every single electron* has been counted. Some cameras provide a gain greater than one, meaning that the total number of electrons has been scaled down to create the number reported by the camera for that pixel. In a perfect camera, every pixel has the same gain and the amount of incoming light has no effect on gain. In real cameras, several things can cause gain to fluctuate with the amount of incoming light; collectively, we call this *nonlinearity*, discussed more below.

Two different underlying technologies are used in today's astronomical cameras: charge-coupled devices (CCD) and complementary metal oxide semiconductors (CMOS). CCD technology is older, and most semiconductor manufacturers have shifted to fabrication of CMOS sensors exclusively. Both technologies have been used very successfully for photometry; neither technology has inherently superior photometric performance. However, the supporting circuitry is quite different between CCD and CMOS cameras, and these differences result in different camera characteristics, summarized in Table 3-1.

The remainder of this section addresses those camera characteristics and settings common to both technologies, followed by the items that are unique to CMOS cameras.

Table 3-1 – Comparison of CMOS and CCD Camera Characteristics

Feature	***CMOS***	***CCD***
Gain	Adjustable	Fixed
Offset	Adjustable	Fixed
Readout Mode	Sometimes selectable	N/A (single mode)
Saturation behavior	Nonlinearity	Nonlinearity + blooming
Thermoelectric cooling	Regulated	Regulated
Shutter	Usually electronic	Often electromechanical
Amplifier glow	Sometimes an issue; absent in most recent designs	Non-issue (electronics are not on sensor chip)
Support for live video	Common	Rare
Binning	In software, after the ADC	Typically within the sensor, before the ADC

3.3.1 Thermoelectric Cooling

In addition to incoming light knocking loose an electron in the sensor's pixels, electrons can also be randomly knocked loose by thermal energy (heat). This adds undesirable thermal signal and noise to the final image. By cooling the sensor, cameras can reduce that thermal noise. This cooling is done with a thermoelectric cooler that uses an electric current to create a temperature differential. Different cameras use different cooling designs (and some cameras have no cooler), but most cameras provide an electronic thermostat to maintain the sensor at a fixed temperature. Since the number of thermal electrons is very temperature-sensitive, this temperature regulation reduces the thermal noise level and holds it constant across an entire evening's observations.

Rapidly changing temperatures cause structural stress in the sensor itself, so best practice is to change sensor temperature slowly. It is important to control both the cool-down cycle at the start of the observing session as well as the warm-up cycle when observing is finished.

The sensor temperature should be set as cold as possible to reduce dark current. You should set the temperature to the coldest the camera can reach using a power level not exceeding 80% (so that there is still a little cooling reserve power if needed). In practice, temperatures of -20C to -30C reduce thermal noise to such a low level that temperatures below this aren't needed. Give the camera about 30 minutes to stabilize before you start taking images.

Your calibration images should be created using the same temperature setting as your science images. During the summer you may have to operate your camera at a warmer temperature. Many photometrists choose a relatively small set of temperatures and create a set of calibration images for each of those temperatures. At night, they cool their camera to the coldest temperature they can reach for which they have calibration data.

3.3.2 Monochrome vs. One-Shot-Color

This *Guide* covers only monochrome cooled cameras: all pixels are the same, and the camera's sensor contains no color filtering (i.e., there is no Bayer matrix). The sensor designer usually sets a goal to have the sensor respond to as wide a range of light wavelengths as possible. You separately purchase colored photometric filters that you can insert into the optical path to define and reduce the bandpass.

Although one-shot-color cameras often provide something the manufacturer calls "monochrome mode," photometry using monochrome mode images from color cameras is usually inferior to photometry based on separate RGB images from the same camera. Again, if you are using a one-shot-color camera, we refer you to the AAVSO *DSLR Observing Manual* instead of this *Guide*.

3.3.3 Camera Shutter

Camera shutters come in several different types:

- **Electromechanical shutter:** Electromechanical shutters typically involve mechanical vanes that are moved by a small motor to either cover the sensor or expose the sensor to starlight. Slight differences in exposure time across the sensor are common, becoming more significant as total exposure time becomes short. Different camera vendors have used different shutter designs; some significantly reduce this effect, and some can be very problematic. (You can experiment using your flat-field system configuration to measure shuttter-caused gradients.) Exposure time gradients of 10% across the sensor are not unusual with exposure times of 0.5 second. Short exposures (< 1 second) should generally not be used with an electromechanical shutter, although some cameras are known to need corrections for any exposures less than about 5 seconds. Electromechanical shutters simplify the process of making dark exposures (see 4.4), because when an exposure is made with the shutter closed, no light will fall onto the sensor.
- **Electronic shutter (global):** A global electronic shutter begins an exposure in all pixels simultaneously and ends the exposure simultaneously across all pixels. This is done electronically, with no mechanical shutter to block the light landing on the sensor. A global electronic shutter makes very short exposure times possible (down into the milli- and micro-second range). However, making a dark exposure requires you to somehow block light from entering the camera. (This is typically done by either covering the end of the telescope tube or by inserting a light-blocking "black filter" into the light path.)
- **Electronic shutter (rolling):** With a rolling electronic shutter, pixels do *not* all start an exposure at the same time. Instead, the start of the exposure "rolls" across the pixels, taking a significant amount of time to move from the first pixel to the last. Similarly, the end of the exposure "rolls" across the pixels in the same order. Each pixel contributes the same exposure

time to the overall image, but the pixels' exposure times are not synchronized. During typical long exposure times used in photometry, the "rolling" of the shutter doesn't matter and doesn't affect photometric accuracy. However, if a video is being created with the sensor (planetary imaging, perhaps), the rolling of the shutter for short exposures can interact with atmospheric seeing fluctuations to create artificial image artifacts. Again, this isn't an issue for typical exposures for photometric use. Creating dark exposures requires the same techniques used for global electronic shutter cameras.

3.3.4 ADC Bit Depth

The analog-to-digital converter (ADC) shown earlier in Figure 3-4 creates its digital number in a binary counter that contains a specific number of binary digits (bits). Dedicated astro-cameras today use 12-, 14-, or 16-bit converters. The number of bits in the ADC establishes a hard upper limit on the largest ADU count that can be generated by a pixel, according to the following table:

ADC Bit Depth	Largest ADU per Pixel
12-bit	4,095
14-bit	16,383
16-bit	65,535

The ADC is also not capable of reporting negative numbers. If the voltage presented to the ADC is slightly negative due to noise, the ADC will truncate the value to zero. If the voltage presented to the ADC would result in a number bigger than the number listed in the table above, the ADC will report only as high as the number in the table.

3.3.5 Saturation

When too much light lands on a pixel (either because the exposure is too long or because the light intensity is too high), an effect called saturation will prevent the overexposed pixels from reporting an ADU value that reflects the true number of photons absorbed. This particular form of nonlinearity can be caused by either of two different effects:

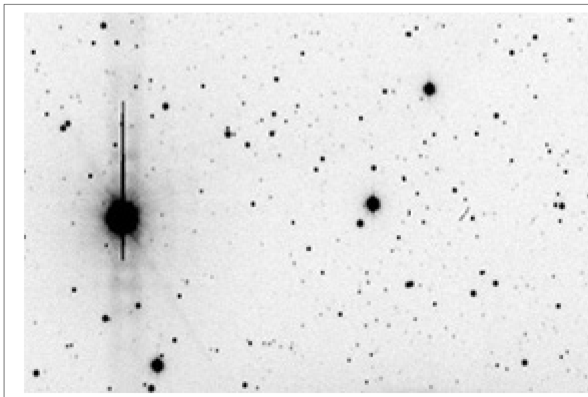

Figure 3-5 — Reversed (negative) image showing blooming.

- "Full Well Depth" exceeded: When a sensor is designed, the size and internal structure of each pixel establishes an upper limit on the number of electrons that can be knocked loose and held for subsequent reporting when the image is read out. If bright starlight releases more electrons than this, the sensor will generally behave in decidedly nonlinear ways; in some sensors, the excess electrons will affect adjacent pixels. (Figure 3-5 shows an example of this with an image artifact called *blooming*.) Exceeding full well depth tends to be a *soft limit,* with pixel behavior starting to change at some electron count somewhat below the pixel's so-called "Full Well Depth" (which is specified by the manufacturer in electrons).

- ADC clipping: Depending on the sensor's gain, the full well depth limit can be reached either *before* or *after* the ADC reaches the limit of what it can report (e.g., 65,535 for a 16-bit ADC). In a sensor with a fixed gain, the sensor's design team will try to make the two limits coincide, so that the ADC reaches its largest possible count at the same time that the pixel response starts to become nonlinear. However, in a sensor with a variable gain, you will commonly see *either* of these two effects to be limiting, depending on the gain you choose.

Saturation is fatal for photometry. As a general rule, if *any* pixel in a star's image saturates (due to *either* effect), measurement of that's star's brightness will be too faint and should not be used. This is true whether the star is a variable star, a comparison star, or a check star.

Detecting saturation is sometimes simple, and sometimes more difficult. It's simple if saturation is due to ADC clipping: if a pixel holds the maximum value for your sensor's ADC bit depth (the table above), that pixel has saturated. A plot of the star's profile will appear flat at the top. However, if saturation is due to nonlinearity because the maximum well depth has been exceeded, saturation is more difficult to detect. See Info Box 3-B for a procedure you can follow to determine at what ADU level your pixels begin to show full well depth limit nonlinearity. *Note that the procedure of Info Box 3-B must be repeated separately for each different gain, binning, or readout mode setting that you use in a variable-gain camera. In a fixed-gain camera, the procedure need only be done once.*

Binning (described below in 3.3.7) combines the values of multiple camera pixels into a smaller number of reported pixel values. Binning can hide saturation and is further discussed below.

Most CCD cameras are designed with an “anti-blooming gate” (ABG) to prevent blooming from happening when full well depth is exceeded by siphoning off the spill-over electrons before they contaminate adjacent pixels. This is great for keeping unsightly spikes out of your pretty galaxy photos, but it reduces the linear range of the camera. CCD cameras without the anti-blooming gate (NABG) are often preferred for photometry, although not required.

When you use a CCD camera with ABG you must measure the camera's limitations and take care not to exceed that range. Whether your camera uses an ABG or not, you have to know what the camera's linearity limit is so you can prevent nonlinearity from affecting your target or comparison stars.

InfoBox 3-B — How to determine the linearity and saturation limits of your camera

1. Set up a light source by illuminating a white screen. (It does not to be perfectly uniform, just stable).
2. Point the telescope at the screen and adjust the brightness until a 10 second exposure will result in a mean central region ADU count of 10,000.
3. Take a series of images where the exposure time increases in 10 second increments (i.e. 10, 20, 30, 40, etc.) until it obviously saturates.
4. Plot the exposure time versus mean central region ADU count.
5. Take one or two more exposures between each 10-second one through straight sections of the plot, then at even more frequent intervals in places of interest of the plot — e.g. where it begins to curve at either end or in any other non-linear section(s).

From this plot you should be able to determine at what count your camera saturates and if there is any non-linear behavior along the way.

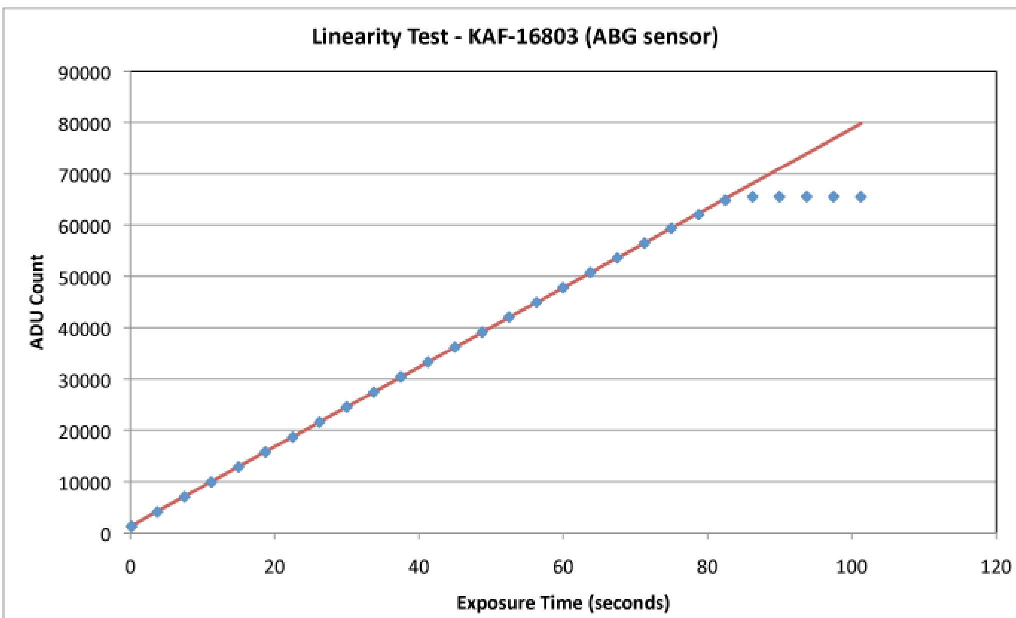

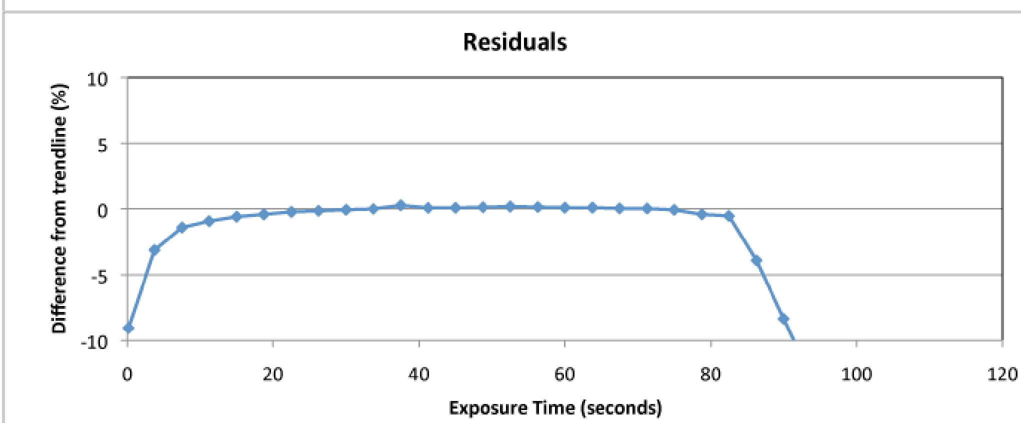

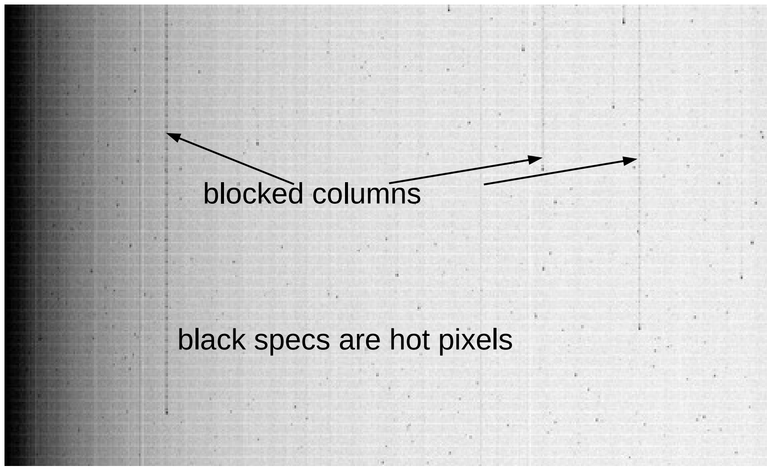

Figure 3-6 — Reversed image of a bias frame showing blocked columns and hot pixels.

3.3.6 Sensor Defects

The sensor itself may have (or develop over time) problems such as “hot pixels”, “blocked columns” or other defects. Most defects are not a problem and will not affect the quality of your photometry as long as you adjust the framing of your images so that critical stars don't fall on a defect.

One way to avoid problems caused by chip defects is to inspect a few of your images carefully and note what you see (Figure 3-6). You can draw a rough sketch of the defects on the image and include the pixel coordinates for each. Also, since sensors degrade with time, it would be a good idea to repeat this exercise at least once each year. Having that information at your fingertips as you point your telescope at the star field you are imaging will better help you to avoid using any bad areas to measure stars of interest to you.

3.3.7 Binning

Binning is something you can do to increase your effective pixel size by grouping pixels together. Your software can be set up to sample (or bin) a group of 2 pixels by 2 pixels (or more) to make those four pixels act as one. There is a tradeoff, however. Resolution will be lost, so you have to be sure that star images have not blurred together with other nearby stars.

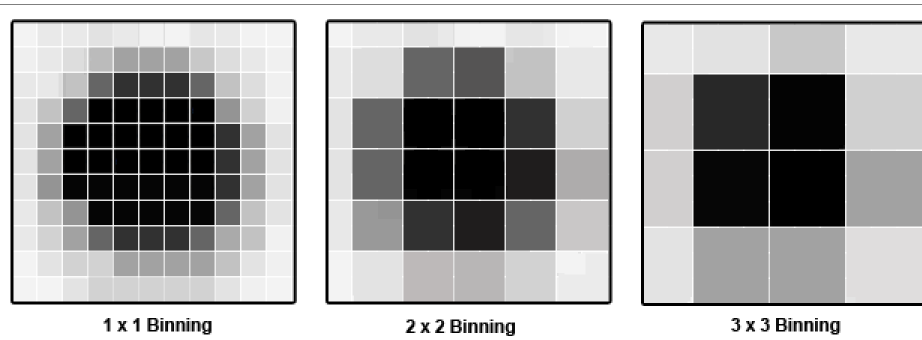

Figure 3-7 – Examples of the same star, captured with three different binning settings.

There are two techniques used for binning; one of the two is unique to CCD cameras. The two techniques are differentiated by whether the pixel combining is done *before* the pixel charge enters the ADC or occurs *after* the A/D conversion. Binning before the ADC is only possible with CCD cameras, and reduces the number of A/D conversion cycles, reducing the amount of readnoise in the binned pixel. Binning after the ADC causes the noise from multiple A/D conversions to be combined into the final, binned pixel. Post-ADC binning offers no noise advantage over unbinned images. (But it does reduce total image size.)

Post-ADC binning can also subtly affect saturation and the nonlinearity threshold. If one of the raw pixels in the binning group is saturated, the accuracy of the photometry will suffer. If you do a linearity test (InfoBox 3-B) you should ensure that you run this test using the same level of binning that you will use for your science images. Your calibration frames must also be binned to the same degree. Analysis and experience consistently suggest that keeping your peak ADU levels (binned or unbinned) to no more than 60-70% of the limit you measure during your linearity checks will keep you from experiencing saturation. But be particularly aware that if you are using CMOS-style binning, you cannot perform reliable photometry if your ADU levels are at 90% (or above) your measured linearity limit.

3.3.8 Image Compression

Different software packages and camera vendors each have "preferred" image file formats. For photometry we use only the Flexible Image Transport System (FITS) standard for storing your images. This format is something of a "least common denominator" that most serious software packages support and that is supported by all software used in serious photometry. (There are many other image file formats that take less disk space (for example, .JPG), but do not preserve the camera's inherent linear response or result in a loss of data. These formats must be avoided at all costs, since even a transient use of this format along the way to getting your image into the FITS standard will make any subsequent photometry from that image useless.)

It is possible to compress FITS files using general-purpose file compression tools (e.g., zip, bz2, compress). These work fine and can save a significant amount of computer disk space, although the amount of compression is typically less than what is achieved using one of the specialized lossless compression algorithms tailored to FITS (such as Tiled Rice compression).

3.3.9 Filter Wheel

Many camera systems will include some option to place filters of various kinds into the beam path between the telescope and the sensor. A filter wheel provides a rotating carousel; you install your color filters into the carousel. Most filter wheels are motorized under control of either the camera or your computer. Once you've installed your filters, you tell the associated software which filter color is in which carousel position. Whenever you want a different filter used, the computer will run the filter wheel's motor to position the desired filter in place.

During calibration you will create flat images that capture (among other things) the size and location of dust specks on your optics, including on your filters. It is important for your filter wheel to accurately reposition your filters to the same location every time you put a particular filter in place. Positioning errors of just a few microns are sufficient to move the effect of a dust mote onto a new set of sensor pixels and degrade the quality of your flat image.

Filter wheels are only needed if you own more than one filter.

3.3.10 Camera Characteristics and Settings Unique to CMOS Cameras

3.3.10.1 Offset

Offset is a fixed bias amount that's added to the voltage derived from the electrons knocked free by photons. The offset is usually added during the ADC process. CCD cameras (with rare exceptions) add a fixed bias amount; CMOS cameras give you control over the bias amount. Setting an appropriate bias amount for photometry is somewhat more important than setting bias for making pretty pictures, but (luckily) there are few downsides to setting the offset a little too high. Typical offset values are around 100 ADU.

Offset matters because the ADC won't report negative numbers, and photometric software often makes an implicit assumption that background image noise is distributed evenly above and below the average background level. Info Box 3-D goes into this in more detail. Your goal is to set offset high enough that the entire background histogram is in the positive region. A larger-than-strictly-necessary offset *will* reduce the system's dynamic range. However, for real-life cameras, the actual dynamic range reduction due to this "extra" offset is less than 0.1 dB (out of a typical dynamic range of 70 or 80 dB), and does not reduce photometric usability.

3.3.10.2 Readout Mode

Some cameras provide you with a choice of what the vendor calls "readout mode." Many cameras, though, don't give you this choice or hide the choice as something that changes "under the hood" between two specific values of a user-selectable gain setting.

Different camera vendors package different things into their readout modes, but they commonly include changes in the pixel readout process that change:

- Read noise,
- Readout time (the amount of time needed to process all the pixels in an image),
- Range of gain settings that can be used,
- Handling of high dynamic range images.

For example, some cameras provide one readout mode intended for use with video recording (fast processing with higher inherent image noise), and a different readout mode intended for use with

InfoBox 3-D — Why Offset Matters

If we create a short-exposure image with little or no light entering the camera, we can take all the ADU pixel values and create a histogram of those values. It will look something like the graph on the left, below. This entire graph can be shifted left or right by changing a fixed *offset* that is incorporated into the sensor's ADC. With a large positive offset, the graph shifts to the graph on the right, below.

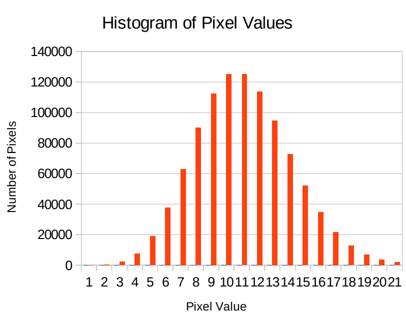

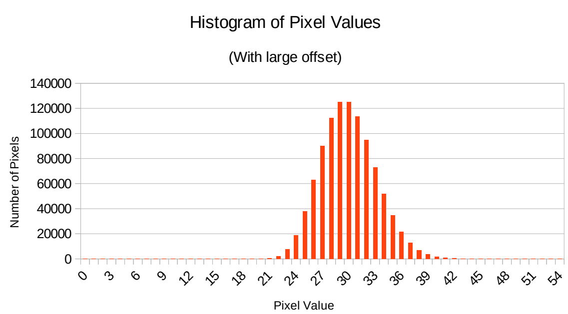

But if the offset is too small, the histogram can move far enough to the left that the left edge of the distribution is clipped to zero (seen at right, below). Notice that the symmetry that the distribution had with larger offset values has been completely lost. Some of the mathematical algorithms found in image processing software assume that this distribution of background pixel values is well-behaved; in particular, some algorithms assume that the only peak coincides with the average background noise level. In this graph, there are two peaks: one at the back-ground average and another one at a pixel value of zero. These problems can lead to analysis software generating *incorrect* values for background noise level, which can propagate into the calculation of measurement accuracy.

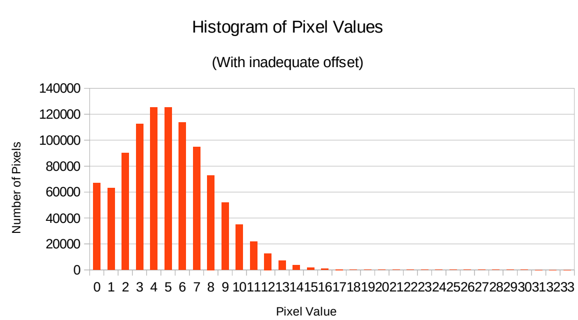

By moving the distribution far enough to the right, this clipping around zero is eliminated.

single-exposure astrophotos (slow processing, higher gain, reduced noise). In general, the readout mode that provides highest dynamic range is recommended for photometry.

3.3.10.3 Gain

Most CMOS cameras have adjustable gain settings, but camera vendors rarely recommend any particular setting, and *never* recommend a setting for photometry. What to choose? The best place to start the selection process is with the camera manufacturer's website. Look for a specification curve that provides the camera's dynamic range. Figure 3-8 is an excerpt from the QHY manual for the QHY268M camera. (Note that QHY originally provided three readout modes (what they simply call "modes", see 3.3.10.2). When the fourth mode was added, some of their graphs did not get re-titled (e.g., the dynamic range graph below, which has four curves, but a title that only acknowledges the first three.))

The highest dynamic range (DR) is with readout mode 1 (blue curve), which shows two "peaks," one at a gain setting of 0 and the other at a gain setting of 56. When multiple gain settings provide the

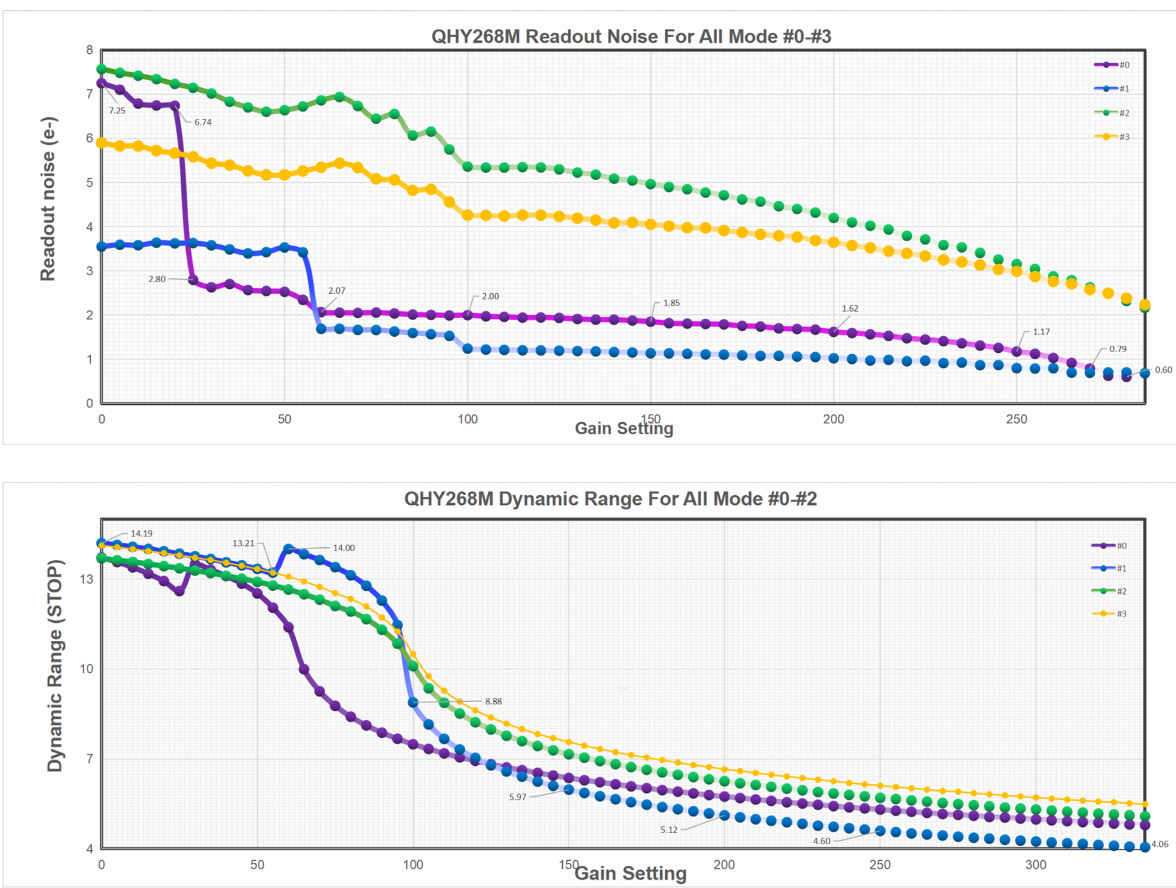

Figure 3-8 Camera Performance Graphs for QHY268M Monochrome Camera. (Graphs by QHYCCD, used by permission.)

same (or similar) dynamic range, you should choose the *lower* gain setting. (This is because the majority of AAVSO observers battle light pollution and moonglow as the dominant noise sources affecting SNR. The lower gain setting allows longer exposure times, more captured photons, and better SNR than a higher gain setting for identical conditions.) For this camera, the preferred gain setting would be 0 with readout mode 1.

[Note also that the *Gain Setting* of the camera is different from the *System Gain* (measured in electrons per ADU). Although the Gain Setting affects System Gain, the units for the Gain Setting are often arbitrary; using a Gain Setting of 100 does not give you a System Gain of 100 electrons/ADU. The manufacturer usually provides a graph translating the Gain Setting into the System Gain.]

3.3.10.4 Overscan and Optic Black Regions

Some CMOS cameras have special areas in the image that do not respond to light: overscan and optic black regions. If your camera has this feature, you want to eliminate these areas from your images. Most camera drivers give you the option to exclude these areas from your downloaded images.

3.4 Filters

In photometry, filters limit the wavelength range of light coming into the detector at a given time. This gives you the ability to measure the rough spectrum of a source within well-defined wavelength bands, providing more physical information about the starlight, and bringing your measurements into alignment with a standard photometric color system (Chapter 6). Filtered photometry provides additional physical information about the target star, and in general, increases the usefulness of your observations. Using several filters provides valuable astrophysical information and is recommended.

Different filters pass different total amounts of light, requiring different exposures. For example, Johnson B filters often require exposures that are twice as long as exposures for Johnson V filters.

If you use only one filter, the best choice is Johnson V. Magnitudes from images made with this filter most closely mimic observations made visually. If you want a second filter, the next most useful is Johnson B followed by Cousins I, and Cousins R in that order. Johnson U is rarely used and is for advanced photometrists. “Johnson” and “Cousins” denote standard filter band passes developed by Harold Johnson and Alan Cousins respectively.

Some colored-glass filters tend to degrade over time, so it is important to inspect your filters annually and make new calibration images frequently (see next chapter). Coated interference filters are generally more durable than colored-glass filters.

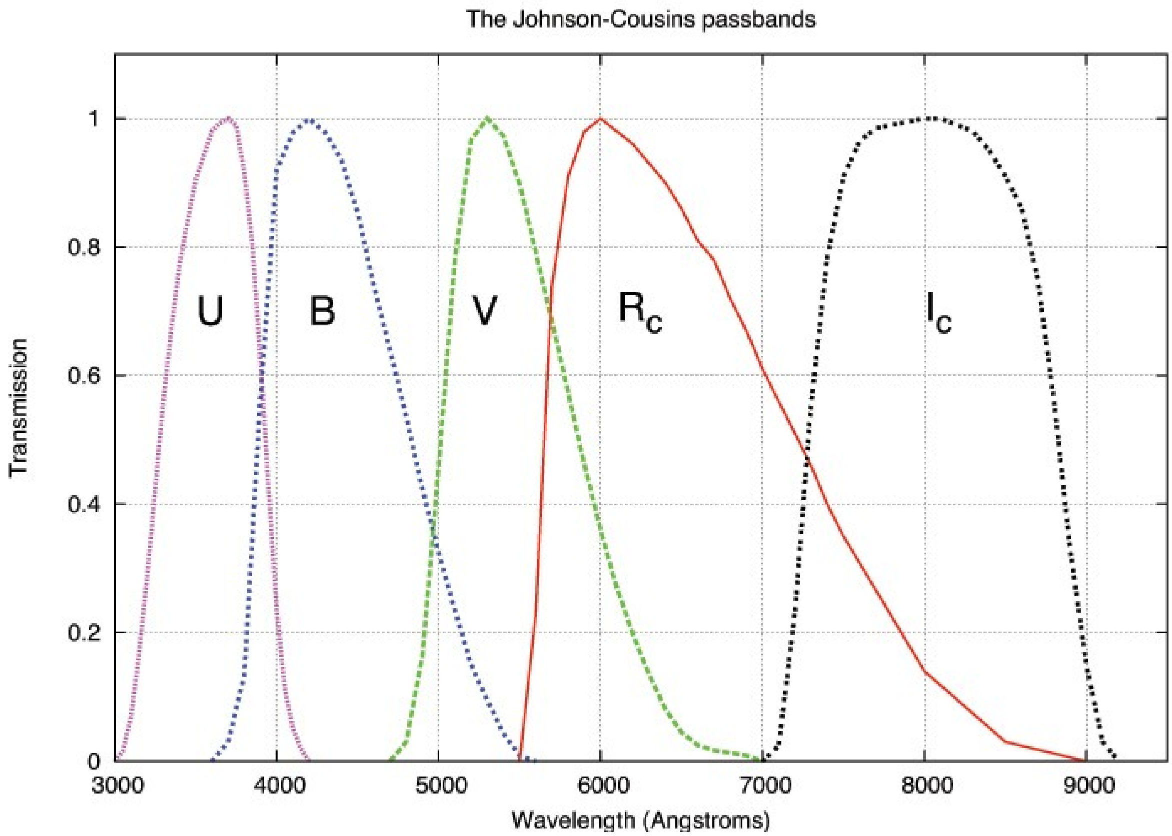

Figure 3-9 — Plot of transmission versus wavelength in each of Johnson–Cousins filters. (courtesy of Michael Richmond, RIT)

3.5 Computer and software

Your computer software performs a variety of functions:

Camera interface — controlling the camera itself, selecting filters, making exposures, perhaps focusing. Often, your camera will come with its own imaging software.
Mount interface – controlling the mount, pointing the telescope.
Calibration – Combining raw bias, dark, and flat images into masters; calibrating science images with master calibration images.
Astrometry — also known as "plate solving" to find the RA and Dec position of each of your target stars; measuring FWHM; locating star centers; identifying stars using existing catalogs.
Photometry — Performing differential brightness measurements using comparison stars; assessing precision and accuracy.
Transforms – Finding color transform coefficients (a few times a year) and applying those transforms every night.
Reporting – create an AAVSO report in the proper format.

This *Guide* does not recommend any specific software; there are simply too many offerings to list them all. Further, each available software package provides its own unique set of functions, and the various software packages overlap in some complex ways. Each photometrist must assemble, learn, and test their own software processing workflow.

However, we *do* single out the tools provided by the AAVSO: VPhot, the TransformGenerator, the TransformApplier (embedded within VPhot), and WebObs. These AAVSO software tools *are* consistent with everything in this *Guide*. Obviously, though, they are not complete. They won't provide control of your camera or mount, and they don't provide you with a way to apply darks and flats during calibration, among other things. However, even if you choose not to base your processing workflow on the AAVSO tools, you should consider using them to test and validate the tool suite that you do assemble. Run a set of test images through VPhot and the TransformApplier. If you get the identical AAVSO report from that process as you do from your workflow, you can feel confident with the workflow you assembled. However, if the results differ by more than a few millimagnitudes, that could be a sign of something amiss.

Be particularly sensitive to software that primarily caters to the astrophotography community. With extreme care, it is possible to successfully incorporate these tools into a photometry workflow. However, astrophotography typically relies on nonlinear processing techniques that will ruin otherwise excellent raw photometric images. Astrophotography software must be carefully configured to ensure that *only* algorithms that preserve linearity are allowed within the photometry workflow.

Obviously, you will need to have a computer available to run all this software. There are no fixed requirements, but Windows is the most commonly used operating system. Some photometry software packages will only run on Windows machines, and there may not be a Mac or Linux version available. (However, there are photometrists who do not use Windows at all in their workflow, if you are so inclined.) It is necessary to carefully track operating system updates for potential compatibility problems with the other software you have installed. (Being able to back out a Windows update is sometimes necessary to avoid losing a night's worth of data collection to an unexpected compatibility problem. A good strategy is to set your system to make updates during periods of the day your system is typically inactive.)

The images you create with your camera will be saved in the FITS file format. FITS (Flexible Image Transport System) is the standard method of storing scientific images into computer-readable files, and is supported by all software packages. A useful feature of the FITS format is that information about the image (such as target name, time of exposure, filter used, etc.) can be stored in a human-readable format in the FITS header, a part of the FITS file in addition to the image itself.

Another necessary computer function is keeping accurate time. The computer controlling your camera

InfoBox 3-E – Assembling a Photometry System

It would be really convenient if this *Guide* could provide a checklist of software and hardware components to acquire and combine to create a good photometry system. But experience has shown that this is actually a poor way to go about assembling a system for several reasons:

- No one starts from scratch — everyone leverages some collection of mount, telescope, camera, filters, and electronics that are already available.
- Compatibility is important — to some degree, investment in one component can offset cost-savings in another. A good example is the interplay between the mount and the guide system: a superb guide system can provide good results if tuned well to a less-than-superb mount.
- Different observing programs have different needs — investment in a good system for timing exoplanet transits will be different than investment in a good system for performing BVRI photometry on long-period Mira variables.
- Computer operating system preferences can run fanatically deep — a photometrist with an aversion to Windows will certainly choose different system software, and in turn that can influence equipment choices.

Instead, this Guide suggests that you start simple (with what you already have, if at all possible), gain experience, faithfully self-assess using the techniques of Appendix E, and dynamically focus on the "weakest link" of your photometry system as revealed by the "root cause assessment" of Chapter 7.

What we *can* provide here is a checklist of equipment and functions that need to be considered. Go through this list and convince yourself that you've thought through each item. Then start performing photometry with your system and tuning the system based on your own self-assessment.

✔ Telescope	✔ Mount	✔ Computer
✔ Camera	✔ Pier/tripod	✔ Guide camera
✔ Filters	✔ Observatory/weather protection	✔ Photometry software
✔ Focuser(s)	✔ Dew control	✔ Data/image storage
✔ Filter wheel	✔ Flat light source	✔ Computer networking
✔ Focal reducer	✔ Guider optics	✔ Guider software
✔ Image acquisition software	✔ Light block for making darks (shutter, lens cap, ...)	✔ Electrical power at the telescope

needs to have Internet access in order to keep its clock synchronized to existing special purpose Internet time servers. Different operating systems use different software products to keep their computer clocks synchronized, but a good software package for Windows computers is Dimension 4. It is important that your computer's clock stays synchronized since that clock provides the time which eventually ends up in the FITS header of your images. Without synchronization (perhaps because you lack Internet access), your computer's built-in timekeeper could be off by several seconds (if not more) in a very short period of time. This may not seem like much, but for measuring short-term variations in some stars or doing transit/eclipse timing work, it could make a critical difference in the correctness of your data.

The other important function of a computer is data storage and archiving. You will quickly find yourself accruing lots of images that consume a lot of storage space. You should decide how you will handle this *before* you start. Everyone makes mistakes or misses problems with images once in awhile, and it is not uncommon for observers to find a processing error, a comparison star sequence change, or some other reason to retrieve and re-process images from the past. Therefore, it is essential that your files are complete and organized so you can find what you need.

Your approach to data storage and archiving should also consider software versioning and system/data backup. You should have a backup plan (and *test* your ability to restore from a backup).

These are the things you should keep in your files:

- Nightly logs containing notes on what is being observed, weather, the moon phase, etc.
- Calibration images
- Nightly raw images
- Calibrated images (flat field and dark subtracted images)
- Logs of observations
- Notes regarding processing

3.6 Charts

Using proper variable star charts is an important part of any variable star observing program and the AAVSO has created an online tool to make this easy for you. You can find the "Variable Star Plotter" (VSP) along with links to help pages on the AAVSO website at https://www.aavso.org/vsp.

Some of the options you might find helpful for photometry include the following:

Choose a chart orientation — selecting the "CCD" option will create a chart with North up and East to the left, much as your camera software should display it.

Do you want a chart or list of field photometry? — You may choose to plot a chart or simply a table of field photometry. You need both. The photometry table identifies candidate comparison stars since it gives position, color, and magnitude information using different filters.

It is also important that you plot a correctly oriented chart of the part of the sky you are imaging so you can use it to check to see that you have identified the star field correctly. Inspect the chart very carefully and if necessary, create a large scale (zoomed in) chart, so that you can use it to check for close companion stars near the variable or any of the comparison stars you plan to use.

The AAVSO comparison star sequences have been carefully chosen and calibrated so please use them! Using non-AAVSO sequences can make your data worse than useless. Systematic biases of as much as 0.1 magnitudes exist between some catalogs. A casual review of AAVSO light curves quickly shows scatter that may be caused by such systematic errors.

Many software packages (e.g., VPhot) already include AAVSO comparison star information so you won't need to type it in, but you should still check to be sure it is not out-of-date. Revisions, updates and new sequences are being produced all the time, largely as a result of requests from observers.

Would you like to display a DSS image *on the chart?* — This option will overlay an image from the Sloan Digital Sky Survey on your chart. This can also help with field identification as it shows the stars in a way that more closely resembles what you will see on your images.

Would you like a standard field chart? — You may find this option useful when you are preparing to image a standard field for the purpose of computing your transformation coefficients. Selection of this option means that comparison star labels will be omitted from all but the "standard stars". See Appendix D for more on transformation.

Figure 3-10 — The AAVSO's Variable Star Plotter (VSP) with the CCD-specific options enlarged.

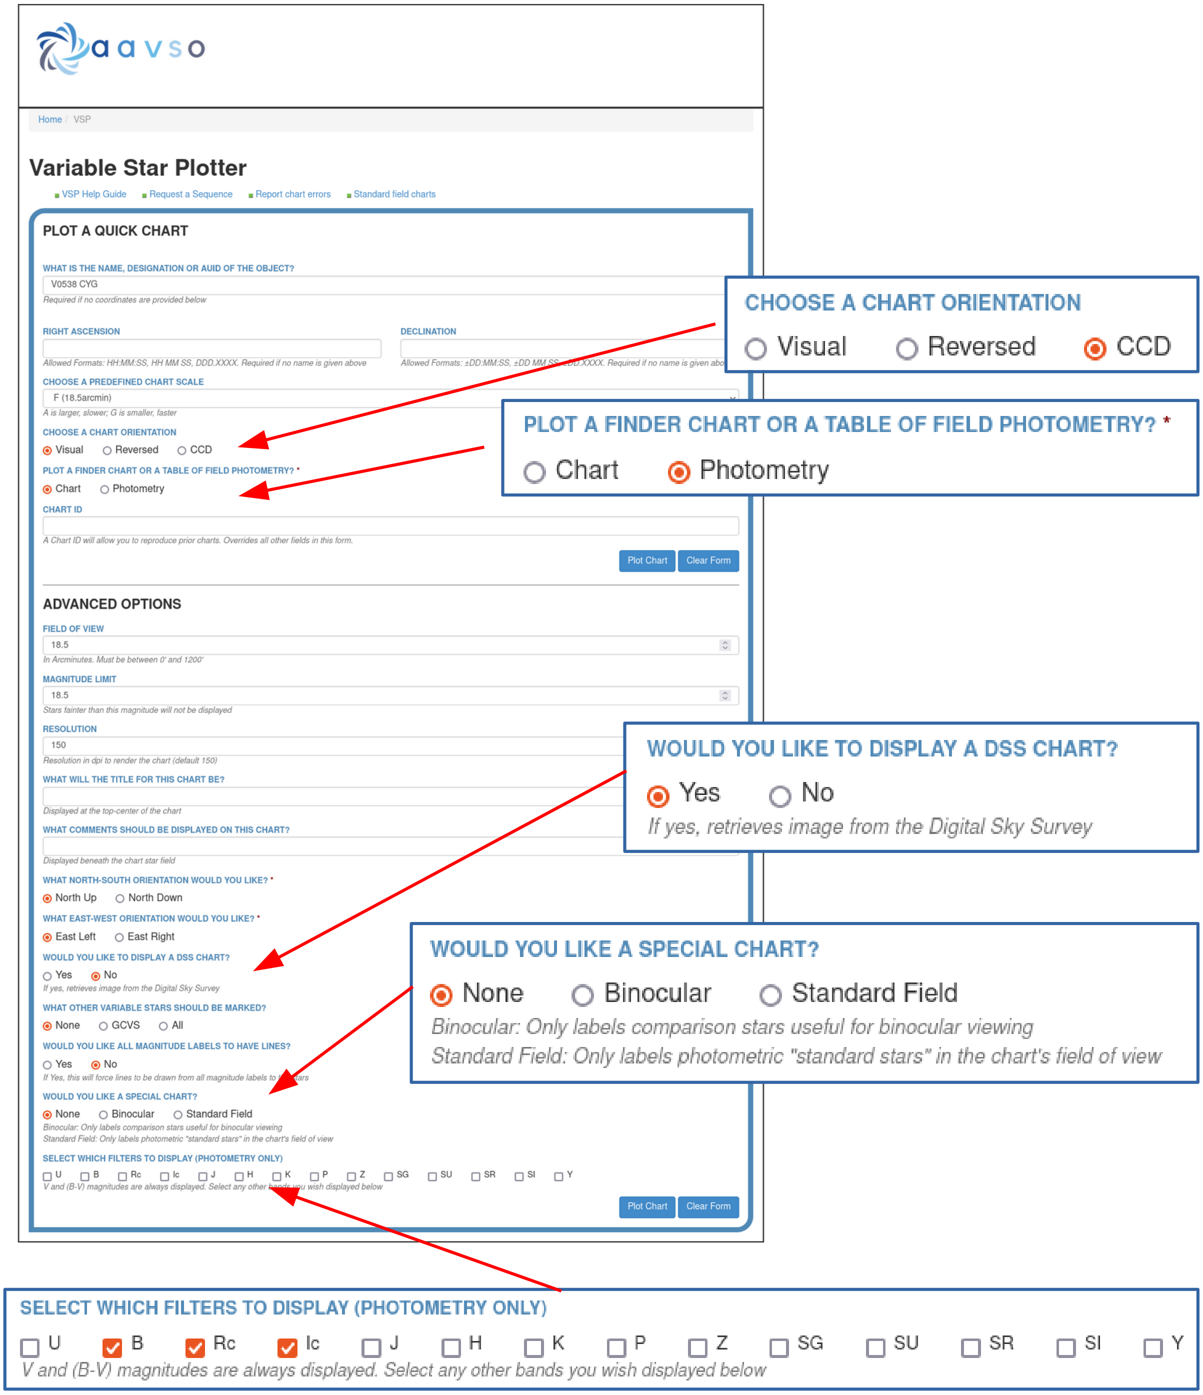

Chapter 4: Image acquisition & processing

4.1 The Nature of Photometric Images

When we image one of our target stars (along with nearby sequence stars that become comparison star candidates), the ADU counts that we get from each pixel contain a variety of information:

- Light from the stars themselves,
- Background skyglow (both man-made light pollution as well as naturally occurring background light),
- Thermal (dark current) electrons, perhaps including amplifier glow,
- A background (bias) level that shows up in every pixel just due to the readout process.

In addition, we see noise: pixel variations across the image that are themselves due to a variety of effects:

- Pixel-specific variations in the background bias level,
- So called hot pixels, that respond more aggressively to thermal (temperature-related) effects than the average pixel,
- Cold pixels, that are less sensitive than their neighbors,
- … and others.

Photometrists use a process called calibration to remove as many of these signals as possible. Image calibration fundamentally boils down to using a series of calibration images to perform pixel-by-pixel adjustments to our photometry images. The adjustments are performed by either subtracting a calibration image pixel-by-pixel from our "science image" or by scaling each of the pixels in our science image based on the value of the corresponding pixel in a calibration image.

Through the calibration process we remove degrading signal from our science images; with care, we can succeed in doing this without simultaneously adding (much) new noise. We do this with careful attention to technique, which helps keep signal-to-noise ratio as high as possible in all of our images, including the calibration frames.

Each of our images contain a variety of signals. Which signals are in which image types is determined by how the image was created. The different kinds of signals that can appear in these images include:

1. Light signal (ADU counts resulting from photons entering the camera),
2. Bias signal (ADU counts – sometimes called an offset – that results from reading out an image with no thermal or photon signal),

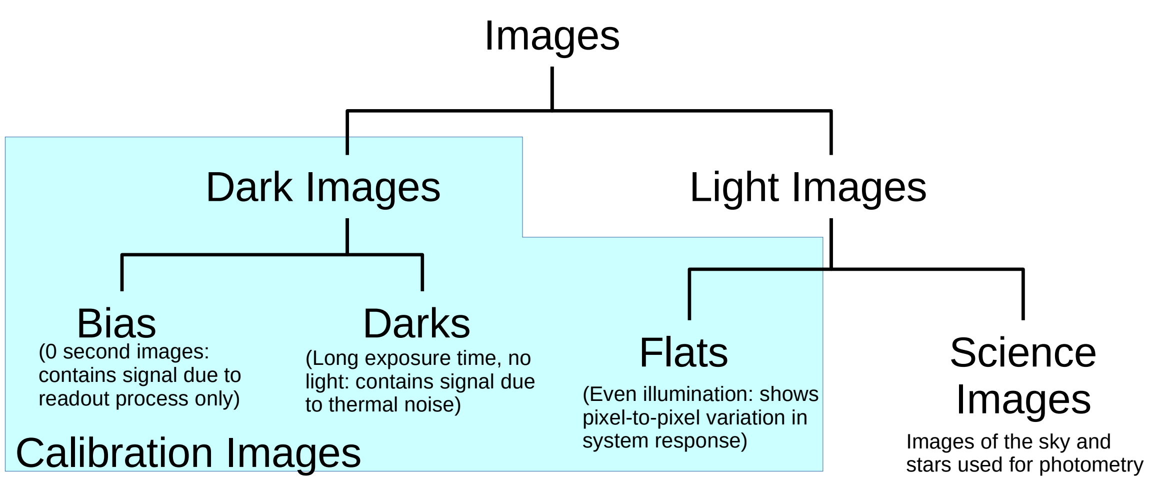

Figure 4-1 – Taxonomy of Image Types. Science images are calibrated using three types of calibration images.

3. Dark signal (ADU counts that result from random thermal electrons),
4. Field flatness signal (ADU counts that vary across the image due to pixel-to-pixel sensor differences, uneven illumination caused by vignetting, and dust specks)

Figure 4-1 summarizes the types of images that photometrists create. Three of the four types are calibration images; the fourth type is the science image.

4.1.1 The Nature of Noise

When we illuminate our cameras with a flat, even source of light, we find that there is an average number of photons that are received by each pixel along with random pixel-to-pixel variations in the number of received photons. The random variations are called shot noise or Poisson noise, and are caused by the inherent unpredictability in the timing of when discrete, random events (the arrival of a photon that kicks loose an electron) will actually occur. Poisson noise has the inherent characteristic that the expected variation in the number of events is equal to the square root of the number of events. (See Figure 4-2.)

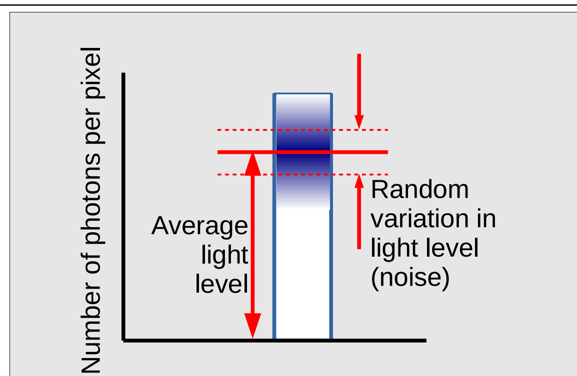

$$Random\ variation = \sqrt{(Total\ number\ of\ events)}$$

$$SNR = \frac{Total\ number\ of\ events}{Random\ variation} = \sqrt{(Total\ number\ of\ events)}$$

Figure 4-2 – The inherent unpredictability of photon arrival times creates image noise.

All of the discrete events that are captured by our pixels have this behavior, regardless of their source; skyglow (light pollution) photons, dark current electrons, and photons from the stars themselves exhibit this shot noise superimposed on the actual count of the number of events from that source. When we talk about these sources of signal in our photometry systems, this *Guide* will be careful to distinguish between the *level* of the signal and the amount of *noise* superimposed on the signal.

During the calibration process we remove the *level* of some of these interfering signals. (For example, we subtract dark images to remove the level of dark current.) However, once a *noise* source has been introduced, we are stuck with it for good; there's no way to subtract the random noise from a pixel (because it's random). Further, if the dark image that we are subtracting has inherent noise (and it will), that noise actually combines with the noise in our science image, *adding* to the total noise in our final image. For this reason, light pollution is bad; it creates measurement noise even though we subtract the background sky brightness during our photometric measurements.

How do we reduce the effects of noise in our work? Usually by increasing exposure time (but without causing saturation!) and by combining multiple images (stacking or averaging). Notice from Figure 4-2 that SNR improves as the total number of events increases. This is the key to improving SNR and reducing the relative contribution of noise to our measurements. When we stack images, the number of events in each stacked pixel increases, causing a rise in the corresponding SNR.

Here are some key things to remember:

1. You *can't* eliminate noise – even the starlight we measure comes with its own shot noise.
2. You *can* subtract unwanted signal levels (thermal dark current, skyglow), but that won't get rid of the noise that was originally associated with the unwanted signal.
3. Noise in your calibration images will combine with noise in the science image — it is important to understand the sources of noise in your calibration images and work to minimize that noise (often by combining many calibration images into a single master — averaging is a good way to reduce noise).
4. Noise is caused by *events*, and multiplying the number of events by a constant doesn't change the underlying number of events. The ADU counts that we read out from our camera only matches the number of events if the system gain is 1.0. When calculating noise, the relationship between the number of events and the ADU value has to be remembered.

We cannot prevent noise in the images, but we can minimize the noise contribution by keeping the signal-to-noise ratio (SNR) high for each of the four kinds of images. Further, by combining multiple calibration images we further increase the SNR and smooth the noise; we need to take a sufficient number of these calibration images to insure that the result (master flats, bias and darks) minimizes the noise introduced into the science image.

4.2 Making Calibration Images

One of the keys to collecting scientifically useful data is calibrating your images properly. The data in the science images needs to accurately represent the signal from the stars. We accomplish this with calibration frames. Each of the types of calibration frames introduced above are described in detail in subsequent sections:

1. **Bias Frames:** These are exposures with a length of zero seconds (or as short as you can), with no light entering the sensor. They measure how the camera's pixels respond to the readout process in the absence of any signal.
2. **Dark Frames:** These are timed exposures with no light entering the camera. They measure how the pixels respond over time in the absence of light; they record the dark current due to thermal electron activity. (These calibration images will also include the bias signal.)
3. **Flat Frames:** These are light exposures made by exposing the telescope, filters, and camera to a diffuse light source that evenly illuminates the entire field of view. The even illumination means the flat frame captures pixel-to-pixel variations across the sensor, including variation in the pixel itself as well as any reduction in light reaching the pixel (possibly caused by dirt specks or uneven illumination patterns), and associated noise. And, like the science frame, it contains both bias and dark signal (and their associated noise).

You may find that your imaging software does most of the work for you. Just be sure to specify which kind of calibration frame you are taking so that your software will know what to do with it later when they are being combined. The other decisions you will have to make as you set up your calibration sequence relate to exposure times, the number of images to make, and what filter to use.

Your software will also make it easy for you to combine (stack) images together and apply your calibration frames to your science frames. Depending on which software package you are using, the steps of combining frames together or subtracting or dividing frames can be automated.

All your calibration and science images must be taken with the same temperature setting (per section 3.3.1). For CMOS cameras you should also use the same gain, offset, binning, and readout mode as your science images.

Photometrists have developed multiple ways to perform image calibration for photometry; those different calibration techniques fall into two broad categories:

1. Those built upon the way that many sensors' response to dark currently grows linearly over time (called **Scaled Darks**), and

2. Those that don't (called **Unscaled Darks**).

These two variations affect several parts of the calibration process and are explained further as calibration is detailed later in this chapter. Both variations can be used by beginners. However, it isn't recommended that you mix different variations in any one observing session.

When using scaled darks, bias frames become important, because the dark frames contain both bias

InfoBox 4-A – Combining Images

In order to reduce the amount of noise in calibration images, this *Guide* advises you to make many calibration frames of each type and to combine them. Some of today's software packages offer a multitude of ways to combine images. What do these different techniques do, and what do we recommend? Here is a list of commonly-encountered combination algorithms:

- **Average** – The value of a pixel in the combined image is set equal to the average of the values of that particular pixel found in all the images being combined. *(Good noise reduction, but individual image problems such as cosmic ray hits do appear in the output*.)

- **Median** – Each pixel in the combined image is set equal to the median of the values of that pixel found in all the images being combined. Half of the images have pixel values larger than the median and half have pixel values smaller than the median. *(Reduces sensitivity to individual image problems, but poorer noise reduction than averaging.)*

- **Min/Max clip** – Similar to averaging, but the brightest value and the dimmest value are omitted prior to averaging (similar to judging in some Olympic events). *(Reduces sensitivity to individual image problems, but poorer noise reduction than averaging because some data is simply not used.)*

- **Sigma clip** – Similar to min/max clipping, but instead of dropping the largest and smallest values, each value is compared to the standard deviation of all the values for that pixel. Those outside the standard deviation band around the average are dropped. *(Very good at rejecting individual image problems, but requires more images than simpler methods.)*

- **Standard Deviation masking** – Shifts back and forth between simple averaging and simple median on a pixel-by-pixel basis based on the standard deviation of that pixel. High standard deviation triggers use of the median. *(Very good at rejecting individual image problems, but requires more images than simpler methods.)*

This *Guide* recommends making a relatively large number of calibration frame exposures that are combined using Min/Max clipping, Sigma clipping, or Standard Deviation masking, if one of those three is available. Otherwise, use median combining.

and dark current signals, and only the dark current scales linearly with time. Having bias frames permits us to separate the linearly-time-varying part from the constant part. We can then create a dark frame of any arbitrary exposure duration by scaling the time-varying part (as long as we are creating a dark frame for an exposure time less than or equal to the exposure time of our original dark frames) and adding the fixed part (which we get from the bias frames).

With unscaled darks, we make no such assumptions about the behavior of dark current and avoid scaling the dark frames. Instead, we create a library of multiple dark frames, one for each exposure time that will need to be calibrated. With unscaled darks, we can leave our dark frames with both their bias signal and their dark signal combined, reducing or eliminating the need to create bias frames.

The two techniques require different quantities and types of calibration frames, and thus differ in the amount of time you need to set aside to create your calibration frames.

Photometrists regularly disagree over which technique yields better photometry. In some cases, the camera mandates the use of unscaled darks; some sensors exhibit enough nonlinear behavior that scaled darks simply don't work well for that particular camera model. Photometrists have found, though, that excellent photometry is possible (in general) with either technique.

4.3 Bias Frames

Your camera and its electronics add intrinsic signals to every image, regardless of the exposure time. Bias frames are used to compensate for readout noise, interference from your computer, and other electronic noise. They also remove any constant signal put in your images by the camera's hardware or software drivers.

Raw bias frames are created by taking zero-second exposures (or the shortest exposure possible with your system) without allowing any light into your camera and at the same binning (1x1 or 2x2), gain, offset, and readout mode as the science images. Since the bias frames you take will be combined to create a "Master Bias", it is necessary to take a lot of them so that any random noise is smoothed out. Using only a small number of noisy bias frames could actually introduce more error into your science images than it removes.

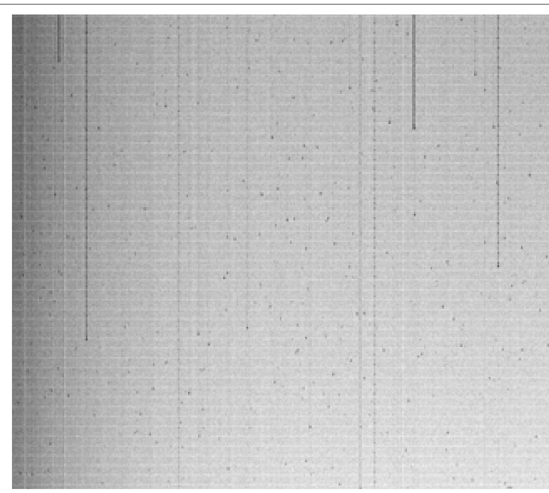

Figure 4-3 — Negative image of a bias frame

You must make and use bias frames if you are using the scaled dark calibration technique; if you are using unscaled darks, bias frames are optional. New bias frames must be taken if you change the camera's temperature setpoint or change a camera setting (such as gain, on CMOS cameras).

4.4 Dark Frames

The thermal energy of electrons within the sensor slowly generate an electron signal in proportion to the temperature and exposure time. This happens not because the pixel is exposed to optical light, but because these thermal electrons have a chance to build up in each pixel over time. *Dark frames* quantify the dark current (the thermal signal) in the sensor so it can be subtracted from the data images. Hot pixels can also generally be controlled with good temperature regulation and will diminish as the chip temperature goes down.

To make raw dark frames, ensure that there is no light entering your camera, and take images having the same binning-mode (1x1 or 2x2) and the same or longer exposure time than you will need for your science images. (Again, keep gain, offset, and readout mode constant if they are adjustable.)

If using unscaled darks, you create a “raw Master Dark” by simply combining all raw darks of the same exposure length. It is used by photometrists who take their science images at one or a few exposures (all at 120 seconds, for example). Since the raw Master Dark contains both dark signal and bias signal, the science image can be dark- and bias-calibrated by simply subtracting the raw Master Dark from the science image. This option is simple, requires no bias frames, and guarantees strict linearity, but requires a library of raw Masters of as many exposure times as used for the science images.

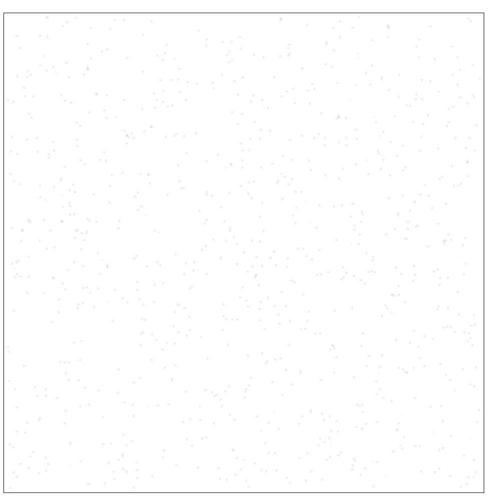

10-second dark frame (negative image)

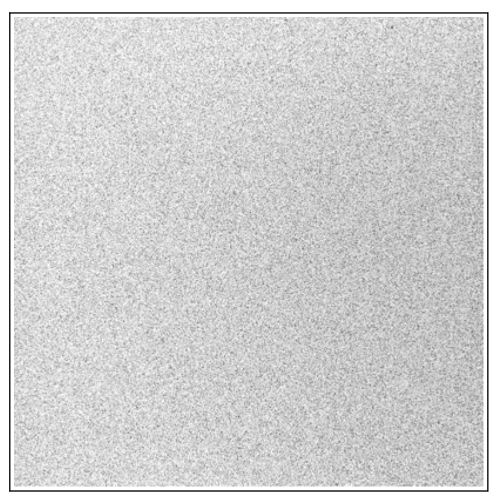

300-second dark frame (negative image)

Figure 4-4 – Examples of dark frames with different exposure times.

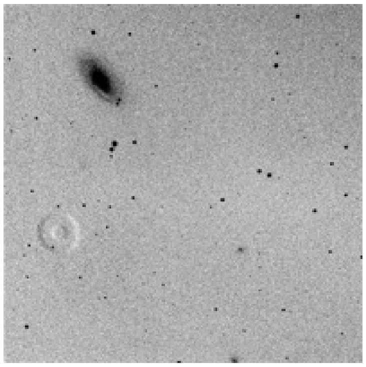

Improperly flatted image showing effect of dust speck

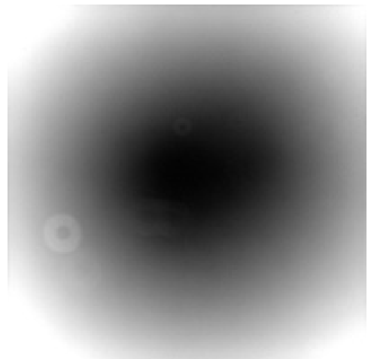

V-Flat with dust

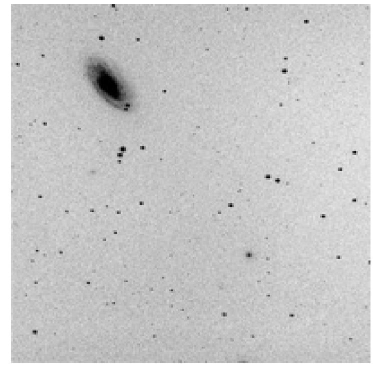

Image with flat applied

If you're using unscaled darks, spend a few minutes now thinking about how this will affect the planning of your observing session. Creating 10 dark frames with a 120 second exposure time takes at least 20 minutes. If you are using multiple exposure times, you need to allocate time to create enough darks for *each* exposure time that you will use.

If using scaled darks, you create a Master Dark by subtracting the Master Bias from each raw darkframe. This Master Dark contains no bias signal and, when used in combination with the Master Bias, can be used with science images that have exposure times equal to or less than the exposure time of the raw darks used to make the Master Dark. The software scales the pixel values of the Master Dark to equal the exposure time of the science image. Because you're not exposing a separate set of raw dark frames for each exposure time, scaled darks can be convenient. Just as with raw bias frames, you must take new raw darks whenever something changes within your system (such as using a new computer, changing sensor temperature, different wiring, etc.).

4.5 Flat Frames

A flat frame is a light image used to adjust your science image to compensate for variations in pixel response and problems in the light path from telescope to sensor. Dust on optical surfaces, reflections from baffles, and poorly aligned optics can all cause light distribution to vary in your system. By taking images of a uniform light source, many of these gradients can be recorded, quantified, and then removed from the science image just as the Master Bias and Master Dark frames remove other kinds of signal.

The critical part about making raw flats is finding a uniform light source. Many people use commercial or homebuilt light boxes or a uniformly illuminated white surface against the wall of their observatory. Another popular procedure is to use the sky itself at morning or evening twilight (see

InfoBox 4-B). In either case, it is important that the source be uniform, otherwise the images taken will not accurately reflect the problems in your light path, but the problems in your light source!

To take raw flat frames, ensure that the temperature of your camera is stable and the same as the temperature used for your raw bias and raw dark frames. The focus should be set to that used for your raw science images, otherwise your "dust donuts" will not match what is affecting the science images. In addition, you must take flats for the binning (1x1 or 2x2), gain, offset, and readout mode you are using for your science images. Exposure times will vary with each filter unless you can adjust the brightness of your light source to compensate for the differences. (If you are using an electromechanical shutter, beware of the minimum exposure time guidelines from section 3.3.3.) The goal is to expose your camera to about 50-70 percent of the full well depth of the pixels (explained in section 3.3.5).

Take at least 5-7 images for each filter. If your light source is the twilight sky, ask your software to "median combine" your raw flats together for each filter to remove any stars that may have been included; otherwise average them. This makes a "raw Master flat." A Master Flat for each filter will be created when the Master Dark and Master Bias are subtracted (scaled darks), or raw Master Dark (unscaled darks). If scaling, use a Master Dark with an exposure time that is equal to or longer than that of the flat to permit scaling the Master Dark to the exposure time of the raw Master flats.

InfoBox 4-B — Taking Twilight Flats

Using the sky itself is the easiest (and least expensive) way to create good flat frames. However, it is not fool proof. By following the suggestions below, you should be able to avoid the major pitfalls.

- Use the approximately 20-30 minute window, starting when the sun is 5°-7° below the horizon in the evening or ending when the sun is 5°-7° below the horizon in the morning.
- Point your telescope toward the zenith.
- **Move your telescope between frames so that stars don't wind up in the same place on any two frames.** Consider placing a white T-shirt over the end of your telescope to further diffuse light from any stars that get imaged.
- Avoid imaging the Milky Way region because too many stars will be captured.
- Don't take flats when there is a bright moon or clouds in the sky.
- Pick an exposure time for each filter which will result in ½ full well but not less than 3 seconds or more than 30 seconds.
- Make flats for the B filter (if used) during the brightest period and the rest of the filters when it is a bit darker.

You can use the set of Master Flats just created for more than one observing session, but it is best practice to take flats each night you observe. Flats can be taken at a time when science images cannot be collected (due to clouds or twilight), so collecting them does not reduce your science image collection. If you make twilight flats (where it is not possible to get more than 4-5 images in each filter due to your time constraints as the sun sets), you can build a running average taken over several

InfoBox 4-C —— Quick Guide to Making Calibration Images

All calibration frames should be captured at the same temperature as that of the science images. Allow the camera cooler to run for ~ ½ hour to reach thermal equilibrium before taking images.

Bias Frames

- Should be done in the dark with shutter closed and/or lens cap on.
- Exposure time should be zero seconds (or shortest possible).
- Take 100 images and average them together to create a Master Bias.

Dark Frames

- Should be done in the dark with shutter closed and/or lens cap on.
- Exposures should be the same (or longer, if using scaled darks) as for your science images.
- Take 20-30 images.
- If combining into an unscaled raw Master Dark use this only with science images of the same exposure and do not use the Master Bias.
- If combining into a Master Dark, subtract the Master Bias from each, then combine them all together to create a Master Dark for use with science images of equal or shorter exposure. Use this with the Master Bias in calibration.

Flat Frames

- Take images of a uniform light source or the twilight sky.
- Ensure that focus is good and the same as that of science images.
- Exposure time should result in flat images with pixels containing ADU counts of about half of the pixel's full well depth and no greater than the linearity limit.
- Take 11-21 images for each filter, combine them together, then subtract a Master Dark and Master Bias to create a Master Flat. Note: your software may perform the dark subtraction automatically for you; see what options are available.
- If using scaled darks, when preparing your Master Dark for flats, use raw darks that are no longer than the longest exposure times for your set of raw flats. Master Darks prepared from raw darks of long integration times (300 seconds for example) may contain hot pixels that do not properly scale to flats taken at much shorter exposures.

nights. How long a set of master flats can be used depends very much on observatory conditions. Dust has a way of getting into everything, no matter how hard you try to keep it out! If anything changes in your optical path (such as adding a focal reducer, cleaning or replacing a filter, or removing or rotating your camera) you *must* create new Master Flats. We recommend dividing a new Master Flat with your last Master Flat to look for new dust or other effects. If you consistently see new features, then you may need to take flats more often.

4.6 Science Image Acquisition

Now that you have a set of calibration frames to work with, it's time to start collecting images of actual variable stars. There are several factors to consider as you create these images.

4.6.1 Use of filters

Your observing plan should include your decisions on which filters to use for your observations. Your software should embed filter information into the FITS files that are created for each exposure.

4.6.2 Choosing exposure times

The exposure time you select for each image depends on a number of factors including the brightness of the variable at the time, which filter you are using, the quality of your telescope's drive mechanism, and whether or not you are guiding. In general, you should use the longest exposure time appropriate for both overall brightness and timescale of the variation you wish to measure. The most critical aspect of choosing an appropriate exposure time for a given filter is not to saturate the image of the variable or any of the comparison stars. Doing so will give you a false reading of the star's brightness which will result in worthless data. (InfoBox 4-D explains how you can determine if saturation has happened.)

To avoid this problem, you should measure your camera's linearity (see section 3.3.5) at the gain, offset, binning, and readout mode that you've chosen, and use the resulting curve to establish the maximum and minimum "safe" exposure time for each magnitude star you are likely to image. You can then save your findings as exposure time versus star magnitude for each filter in a table for future reference. This will save you a lot of time and possible frustration in the future.

Keep in mind that a star image can saturate long before it "blooms" (i.e. you see vertical spikes coming out of it)!

Here are some other useful tips related to choosing exposure times:

- If you are uncertain as to the exposure time to use on a new target, always err on the side of

InfoBox 4-D — 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.

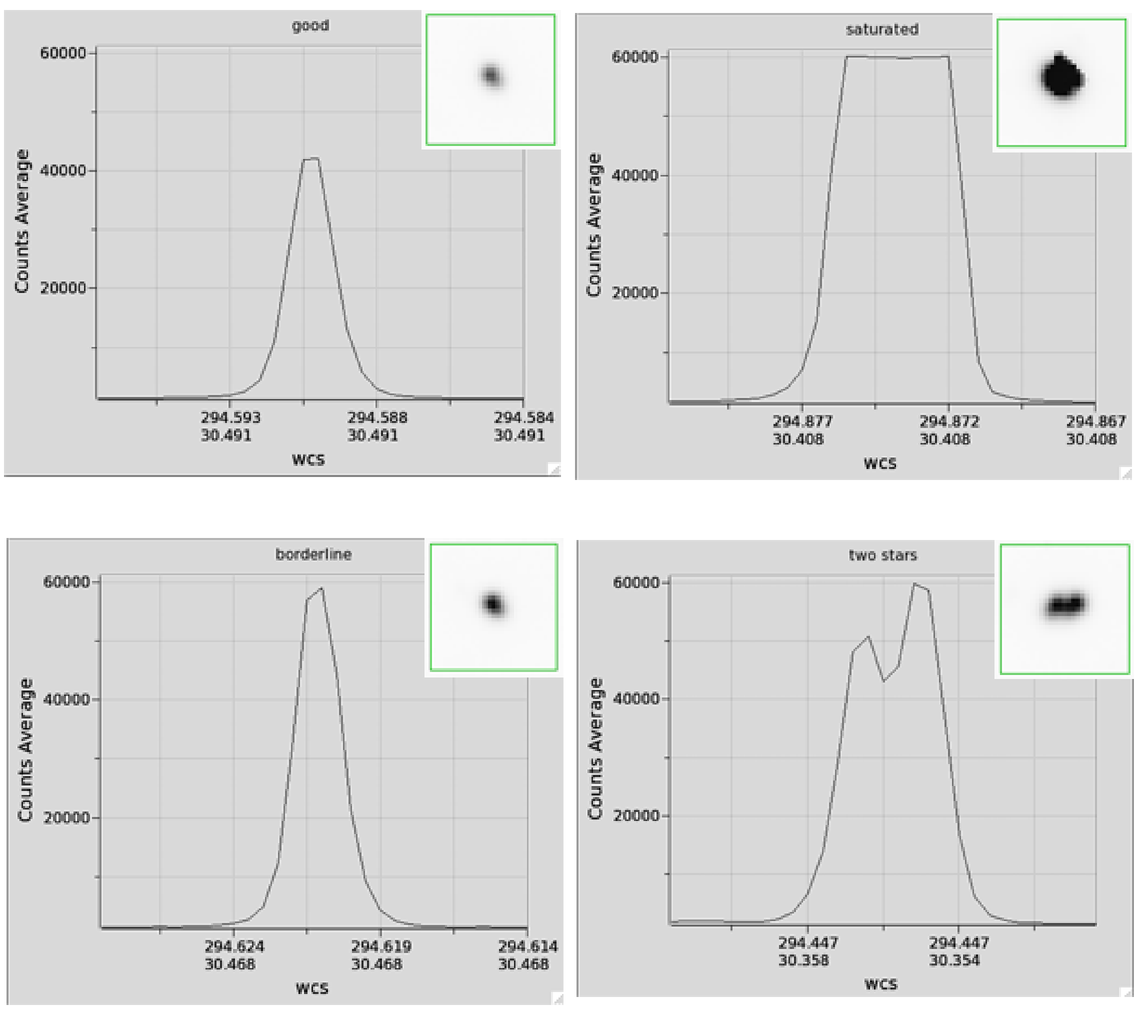

a shorter exposure or take a series of test images (and use your software to measure SNR).

- Very long exposures are best broken into several shorter exposures. The longer the exposure, the more chance there is that your image could be spoiled by drive abnormalities, a passing satellite, cosmic ray hits, passing cloud, etc. The shorter images can be stacked to improve the SNR.
- If your camera has an electromechanical shutter you should avoid exposure times less than 10 seconds and never take exposures of less than 3 seconds. Anything shorter will cause the shutter opening and closing to expose different regions of the sensor for significantly different times. This is not a problem for cameras with electronic shutters.
- Different filters will nearly always require different exposure times, not only because of filter and sensor response, but because the star may emit much less light in one band than another. This is especially true of bluer filters, particularly when observing red stars.

4.6.3 Deciding how many images to make

This section distinguishes between the number of measurements that you *make* (using photometry to extract a magnitude) and the number of measurements that you *report* (an addition to the AAVSO International Database). These numbers will often not match if, for example, you average multiple measurements into a single submission (to both improve quality through averaging and to estimate uncertainty by measuring the standard deviation of the contributing measurements).

Four interrelated considerations influence how many exposures you need to make:

Observing Cadence: In most cases, the needed observing cadence derives from the behavior of the star you are observing. For now, we only distinguish between a cadence that calls for a single *reported* measurement for the night (as you would for a long-period variable) and one that requires a time series, providing multiple reported measurements over the course of a single night. A time series observing run in which hundreds of measurements are made of one target star over the course of a single night should be reserved for stars which are actually doing something in the astrophysical sense over that short a time scale (e.g., short-period variables or eclipsing binaries). More information on this subject is covered in the section of this *Guide* on “Photometry and Science” (see section 8.2). The key is that the cadence must be *appropriate* for the science. Too many observations of some kinds of stars in too short a time distort the light curve, contaminate the database, and waste your time. Too few observations of other stars can leave your data incomplete and potentially meaningless.

SNR: The science that you are supporting will dictate the minimum SNR that you can tolerate. For instance, if your science goal is to distinguish whether an eruptive star is in its active or quiescent state, a relatively low SNR may suffice. On the other hand, if you are looking for color shifts associated with an exoplanet transit, you probably need a relatively high SNR. Remember, also, that your SNR goal must be considered for both the target and for the comparison star(s). Having a good target SNR but a poor comparison star SNR leads to higher measurement uncertainty than having a good SNR for both.

Exposure Time: In the previous section (4.6.2), you decided what your exposure time would be, based on a series of factors, one of which is the SNR you need. In some cases, the exposure time that you want will be sufficiently long that you choose to break up the total integration time into a series of exposures that you stack into a single image that will be analyzed to create a single measurement.

Measurement Uncertainty: This *Guide* describes several ways of estimating the uncertainty of your measurements; one (good) way is to make multiple measurements and measure the standard deviation of those individual measurements about their mean. Three measurements are the minimum number that offers a reasonable mean and standard deviation.

Use these four considerations to make several decisions:

- How many reported values will be submitted for the session? If more than one, with what cadence?
- How many photometric measurements will contribute to each reported value?
- How many images will be combined (stacked) for each photometric measurement?

As a general rule, the answers to the first two questions will be the same for all the filters that you use for that target star on that night. The answer to the third question may well be filter-dependent, affected by both the target star's color and by the sensitivity of your system in each filter passband.

The answers to these three questions will determine both the number of images that you need as well as whether you elect to interleave images with different filters.

4.7 Special cases and other issues

4.7.1 Bright stars

Bright stars pose a special problem for photometrists. In order to avoid saturation of your star image, you will want to use a short exposure time. However, in addition to possible issues caused by the shutter opening and closing, very short exposure images can suffer more from scintillation effects than longer ones where the “twinkling” is averaged out over a longer period of time. To avoid such problems, it is recommended that you typically take exposures of longer than 10 seconds duration. When you reach the point where you cannot take a short enough exposure to avoid saturation, you may wish to try one or more of the following techniques:

- Use an aperture mask on the end of your telescope to reduce the amount of incoming light getting to your camera. (Note that you will need to retake flats if you do!)
- Try using a B filter instead of a V filter. The filter itself reduces the amount of light reaching your camera, plus some cameras are less sensitive to blue light than to the longer wavelengths of the V, Rc, or Ic bands.
- Defocus the image a little. This spreads the light out over several pixels, allowing you to increase the exposure time before saturation occurs.

In any case, if you are forced into very short exposure times (< 10 seconds) to avoid saturation, take multiple images and combine them into a single measurement if the star varies slowly enough. This will help lessen the impact of scintillation, which can incorrectly yield false magnitude variations of more than 0.1 magnitudes! A useful thumb rule is to stack enough short exposures to yield a total image time of at least 20 seconds.

InfoBox 4-E — Scintillation

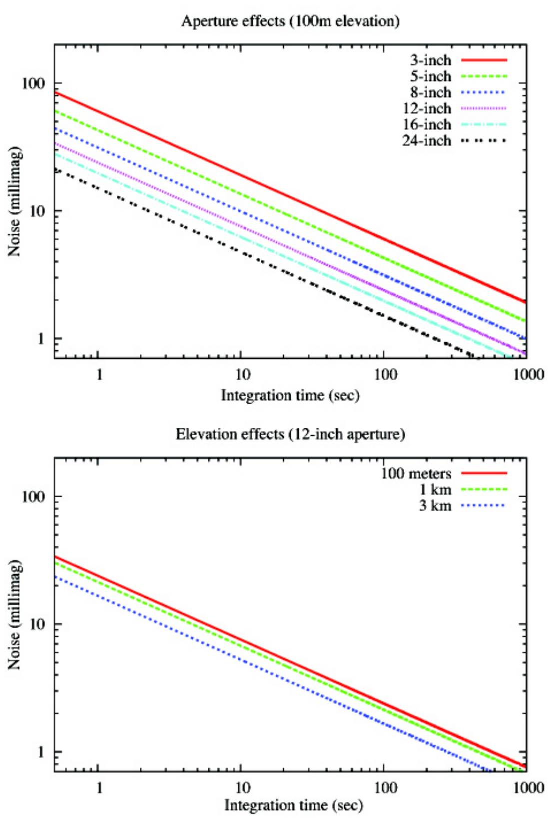

Scintillation is caused by refraction of starlight by individual turbulent cells in the atmosphere. The stars scintillate on both short and long timescales, but the amplitudes of changes on short timescales are larger. Scintillation has been measured experimentally (see *Young 1967*) and the noise effects on a signal can be approximated as a function of the telescope aperture, the exposure time, the airmass, and the elevation of the telescope site. This graph shows the effects of aperture (top) and site elevation (bottom) on the scintillation noise as a function of exposure time using Young's equation (assuming S_0 =0.09, airmass=1.5). Larger telescope apertures serve to average over more small turbulent cells, so the noise effects in large aperture telescopes are greatly reduced. Radu Corlan's website has useful tables of scintillation effects, available at: *http://astro.corlan.net/gcx/scint.txt*.

4.7.2 Crowded fields

Inexperienced observers should avoid imaging fields in which the stars are very close together. The reason is that it is very difficult to do accurate photometry when stars are touching or overlapping each other. Data that accidentally combines light from two stars is generally of very little use. In order to separate two closely spaced stars, you could use a larger telescope or use mathematical techniques such as point spread function (PSF) fitting, which is beyond the scope of this *Guide*.

The one exception to this guideline is when the nearby star has 1% or less of the counts of the target star throughout the range of the variable. In this case, you could use the combined light of the variable and the nearby star. However, in crowded fields, this is rarely the case. Worse, variables with large ranges (like Miras) may be much brighter than the nearby star at maximum, but fainter at minimum. This case often leads to observers confusing the identities of the two stars, and the AAVSO archives have a number of "flat-bottomed" light curves as a result.

4.7.3 Near horizon

Observations made low on the horizon should also be avoided. Observe objects only when the airmass is less than 2.5 (or altitude > ~23°). When light from a star has to pass through a thicker cross-section of the earth's atmosphere, its brightness is diminished. This is known as *attenuation* or *atmospheric extinction*. It is possible to apply corrections to your data to make up for this, but it gets complicated since the rate of attenuation changes rapidly as you near the horizon. The effect also differs depending on the color of the stars you are measuring. At some point, you will need to apply different amounts of extinction to every star even in the same field of view. The seeing also gets worse as you get closer to the horizon.

The thickness of the atmosphere is quantified in terms of airmass. Airmass is defined as the length of the light path as it passes through the atmosphere, referenced to the length of the shortest possible path — straight up. Thus, the airmass for an object directly overhead is 1.0 and the airmass for something on the horizon is very large.

When submitting your data to the AAVSO, you should include the airmass for each observation. If your photometry software does not calculate it or you cannot get airmass from your planetarium software, you can estimate the zenith angle of your target and compute it yourself (see InfoBox 4-E).

4.8 Image Inspection

During the observation session, it is important that you perform at least one round of quality control by inspecting the images. You will become aware of potential problems with your system or procedures as well as conditions outside of your control which may affect your final results (e.g., *InfoBox 4-F — Estimating airmass*

Airmass (X) at sea level can be approximated using this formula:

$$X = 1/\cos(\theta)$$

Where θ is the zenith angle (measured from 0º at the zenith to 90º at the horizon)

Altitude (angle above horizon)	Zenith angle (angle from overhead)	Formula Airmass	Airmass (actual)
90°	0°	1.00	1.00
60°	30°	1.15	1.154
30°	60°	2.00	1.995
25°	65°	2.37	2.357
20°	70°	2.92	2.904
10°	80°	5.76	5.60

satellite trails). In some cases, you can still use the images, but in others you cannot. Either way, it will save you a lot of trouble later on when you try to figure out why one observation is so different from the rest.

The next few pages contain a list of common image problems and how they manifest themselves. Examples of images with these problems can be found on page 71.

4.8.1 Saturation

Stars that are much too bright for the exposure time often display blooming. Again, it is important to note that a star's image can be saturated well before you see blooming. To see if a star has saturated, check its ADU count in the brightest center of the star. Do this for the target star as well as for the check star and all the comparison stars you plan to use. If the ADU count for any of them gets close to or exceeds the linearity limit of your camera, then that star is saturated and should not be included in any measurements. Alternatively, look at the star profile to determine if the profile is "flat-topped." It is perfectly fine to use other unsaturated stars in the field as long as they aren't affected by blooming spikes from any star that is saturated.

4.8.2 Filter problems

Your filter wheel is a fairly delicate piece of equipment. Sometimes the filter wheel can stick, causing it either to not turn at all or to rotate only halfway into position. A filter stuck part-way will often obscure the stars in part of your image. If the filter wheel does not turn at all, you may think you are imaging in a certain color when you aren't. This may be harder to detect until after you perform your photometry and see how the magnitudes of the stars you measured compare with the magnitudes you derived from another color filter. If something doesn't make sense, go back and check it!

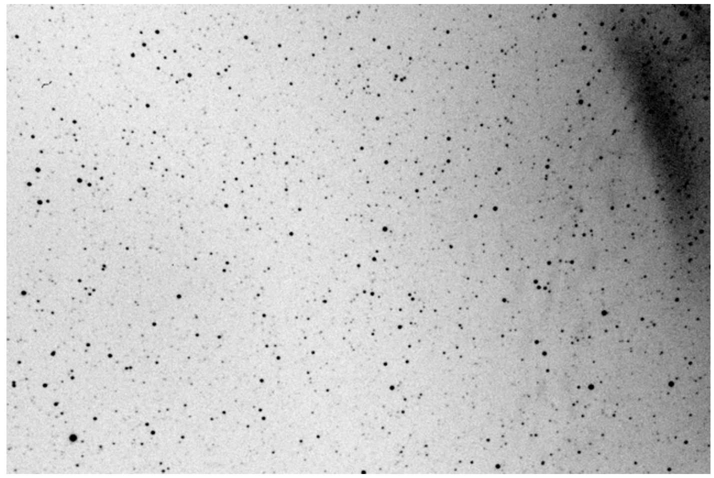

Negative image showing the effect of scattered light from the moon in the upper right corner.

4.8.3 Scattered light

Reflections off the inside of your telescope tube or other optical elements can cause bright areas, rings or double star images that could affect your results. This is particularly evident when the moon is up or there are bright stars or planets near the field you are imaging.

4.8.4 .Ghosts (residual bulk images or RBIs)

Due to the way the sensor works in your camera, if you image something bright you may see a "ghost" of that same object on the next image. You can tell it is a ghost if it looks like a fuzzy patch and gradually fades with each subsequent image. Generally these artifacts are not a problem unless they interfere with a star you are trying to measure or confuse you as you identify the field. They are more prevalent with images taken using a red (e.g. Rc or Ic) filter. To avoid them, try warming up your camera and waiting few minutes for the image to "bleed". When you cool your camera down again, it should be gone. Another possible option is to keep bright objects near the edge of your field of view so the ghost is unlikely to affect anything.

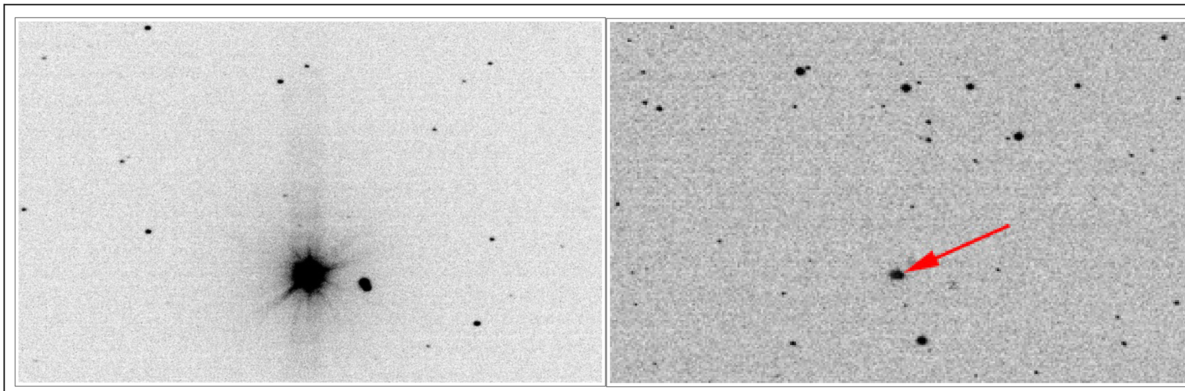

Figure 4-5 – The left image is a negative image of the bright star DY Eri. The following image (shown on the right) shows a "ghost" of DY Eri in a different field.

4.8.5 Atmospheric problems

When you are setting up your equipment for a night of imaging, take a few moments to study the sky! Record what you see, especially if there are clouds about, and make notes about the seeing conditions and transparency. As it is difficult to see thin cloud in a very dark sky, you should consider recording what you see when it is still twilight or during dawn.

It isn't always easy to detect the effect of thin cloud in your images, but later on as you study the results of your photometry and suspect that something could be wrong, your notes may come in very handy. In rare cases, a thin, uniform cloud may affect your target star and the comparison stars you are using to the same degree, and due to the way differential photometry works, the effect will cancel out. However, this is rarely the case so you should regard your measurements during questionable weather conditions with a great deal of skepticism.

Figure 4-6, below, shows how clouds can affect your images (on the left) and how the resulting poor photometry can cause large variation in your photometric measurements (on the right).

4.8.6 Cosmic rays

It is not unusual for you to see the effect of cosmic ray hits on your images especially if you are observing from a higher-altitude location. These will manifest themselves as small streaks, curls or small, sharp (1-3 pixel) bright spots on your images. They are random and generally do not pose a problem. If however, one should happen to land in the signal circle or the sky annulus of a star you are measuring, the effect might be noticeable.

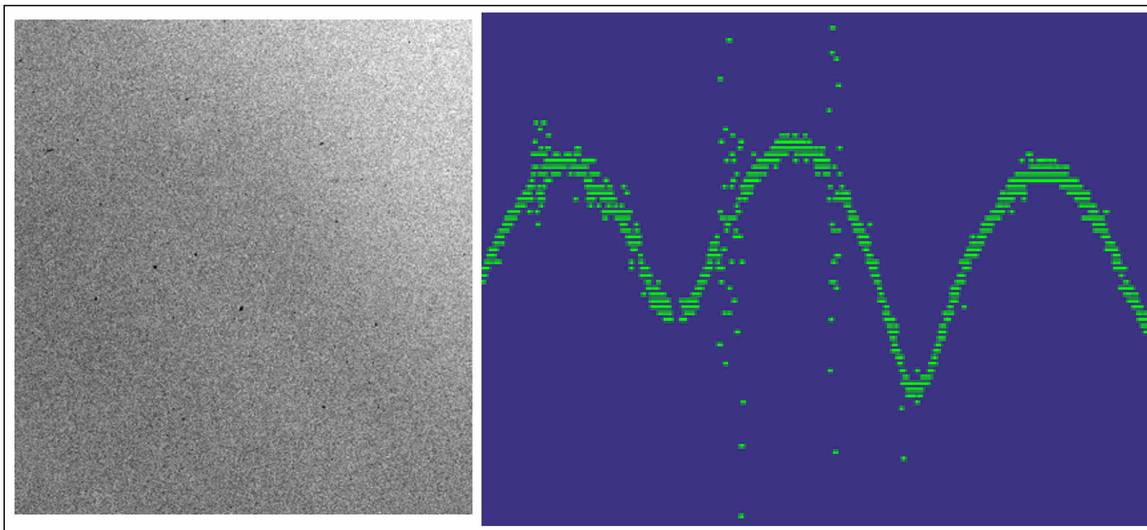

Figure 4-6 – On the left is a negative image showing light gradients caused by dense clouds. The lightcurve on the right shows the effects of intermittent clouds on a lightcurve of VW Cep.

4.8.7 Airplanes/satellites/meteors

Much as with cosmic rays, airplane, meteor and satellite trails which pass through your image are not a problem as long as they aren't too close to a star you are measuring. If you are unlucky enough for this to happen, you may have to choose other comparison stars or skip using this image altogether.

More potential image problems:

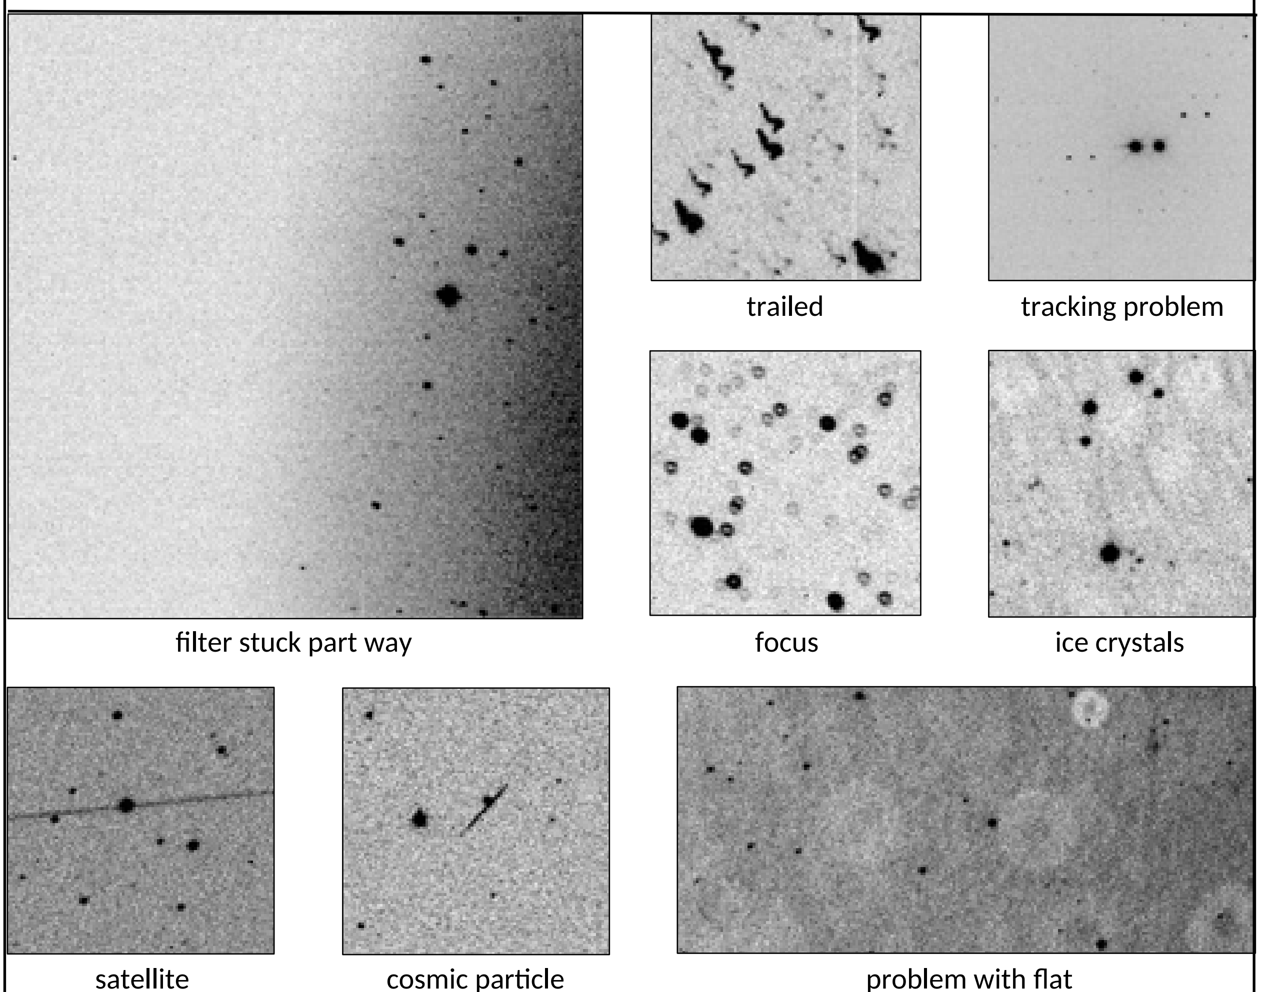

Chapter 5: Photometry — Measuring Images

Now that you have a set of carefully calibrated images, it's time to measure the brightness of the stars you have captured. 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. This *Guide* does not replace your software's user manual; you need to study and understand your software's documentation. However, this *Guide* should not be contradicting that documentation. If you find something inconsistent, don't blindly assume that one or the other is correct – reach out to one of the AAVSO forums and seek guidance.

5.1 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 a star of known brightness in the same field of view 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 is covered in this *Guide* because it is far easier, normally yields excellent results, and is 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 star 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. This is because we assume that the sky conditions are the same throughout the image.

The analysis phase of the workflow introduced in section 2.4 includes the following steps:

- Check your images
- Identify the stars
- Set the aperture
- Choose the check and comparison stars

- Measure the magnitudes
- Determine the uncertainty

5.2 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 comparison) 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 InfoBox 4-C). 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. It is vital that you have already determined the linearity of your camera (InfoBox 3-B) and that you don't exceed 50-70% of that limit (particularly if you are using CMOS-style binning). 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.

5.3 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. Further, when some software encounters faint stars adjacent to bright ones, it will "recenter" your cursor onto the bright star regardless of which you tried to select. 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 is either done automatically or done by you using charts. In either case, it is important to check to be sure that the variable, comparison, and check stars are correctly identified. Astrometry software is good but not perfect! It can be thrown off by defects in your images and can 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 candidate 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.

If your target star is not in the AAVSO database you can request it be assigned a unique identifier (via the VSX search page), and ask the AAVSO Sequence Team to provide a comparison star sequence.

5.4 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 aperture photometry, so named because we measure the strength of light in circular 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 software packages, they will not be covered here.

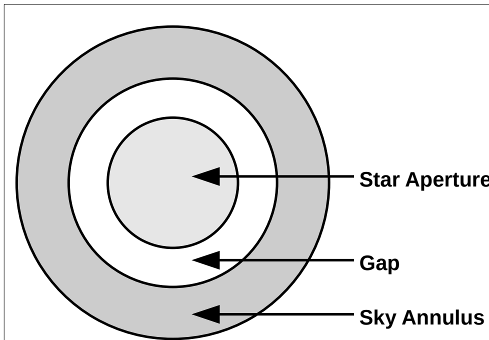

Photometry aperture and sky annulus

The aperture has 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 brightness of the sky background. The annulus has both an inner radius and an outer radius.

The software package you use will probably give you some control over the size of the each aperture; you 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 average FWHM of all the stars you wish to measure. (Hence, aperture *radius* will be 1.5-2x the FWHM.) 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 32.) Be careful with your software, since some report the star's radius and others report the star's diameter. Also be careful that both measurements are in pixels and that you aren't mixing arcseconds and pixels.
- Don't worry if the brighter stars in your image seem to exceed the limits of the star aperture.

Even though all the stars have roughly equal FWHM, the brighter stars may *appear* bigger than the fainter stars.

- The diameter of the inner circle of the sky annulus should be about 5 times the average FWHM (typically about 10 pixels across).
- Adjust the outer ring of the sky annulus. A bigger sky annulus yields better signal-to-noise ratio (SNR) of the sky background. Stars in the sky annulus do not typically affect the measure of the sky background because most photometry software uses techniques to exclude stars falling within the sky annulus. Ideally, the sky annulus should contain about 10x as many pixels as the star measurement aperture, which suggests an outer diameter about 9 times the average FWHM.

5.5 Choose the check and comparison stars to use

AAVSO sequences have been carefully designed to use stars for which magnitudes have been determined very accurately and are known not to vary or have close companions. By using a standard comparison star, your results will compare more favorably with those of other AAVSO observers when your data are combined in the AAVSO International Database. Researchers using your data will appreciate that.

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

- You can choose either a single comparison star or multiple stars, called a comparison star ensemble. There are strengths and weaknesses to each approach. *Do not use ensembles* until you are an accomplished photometrist and understand and can compensate for the complications they carry. Although some software encourages ensembles as a default, all software packages support the use of single comparison star photometry.

	Single Comparison Star	Comparison Star Ensemble
Strengths	• Easy to apply color transforms using VPhot • Aligns well with AAVSO Extended File Format, enabling later corrections if sequence is adjusted	• Reduces uncertainty of the comparison star reference by averaging across the ensemble
Weaknesses	• Uncertainty in magnitude of the comparison star directly affects accuracy of the variable's magnitude	• Requires an alternate approach to apply color transforms in VPhot or alternate algorithm specifically designed for use with ensembles • Any mistakes in your computations are extremely hard to detect and/or correct after submission

- Select a comparison star located close to the target and not near the edges of the image where it could be distorted. It is useful to select comps within the center half of the image, if possible. In very large fields of view (i.e., degrees), it is essential to select comps within 30 arcminutes of the image center (and target star) to minimize uncorrected extinction effects.
- Use a comparison star similar in color to the target star. If the color is identical to that the of the target star, the estimated magnitude does not require later color transformation. Unfortunately, comps of identical color are rarely found. Similar, but not identical, color comps are sometimes available, minimizing inaccuracy due to color effects.
- 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 a comparison star that is similar in magnitude to the target star.
- If using an ensemble, bracket the target star's magnitude by selecting comp stars with magnitudes that are within one magnitude of the target. Note that as the target magnitude changes, the best comps may need to be re-selected. In fact, as long as selected comparison stars are sufficiently bright, (SNR>20), all comps could be used to calculate the target magnitude under all conditions (bright or faint). It just requires more effort to confirm those comps which yield good agreement with the average target magnitude.
- Try to select a comp star that has no nearby companions that could contaminate the apertures, especially the measurement aperture. The impact of extraneous stars in the sky background aperture may be statistically removed in many photometry software including VPhot.
- Choose a comparison star with a signal-to-noise ratio (SNR) of at least 100. (Ideally, the target star will also have an SNR above 100.) When target stars are very faint, it may become necessary to include comp stars with fainter SNR, perhaps as low as SNR>20.
- Choose a comp star with small magnitude errors, preferably less than 0.01 - 0.02.
- Ensure that the comparison, check, and target stars are all below 50-70% of the saturation point in your image.

A check star is an important quality check. It will detect any variation in your comparison stars or other problems that may exist with your image. A check star magnitude is required in AAVSO Extended File Format Report. A check star is simply a comparison star of known brightness that doesn't vary which can be treated in the same way as you treat your target star. The selected check star magnitude must be similar to the target magnitude so it will best represent the potential accuracy and precision of the target. It is critical that you compare the magnitude you determine for it with its published magnitude (in the same bandpass), and the agreement should be very close (<0.05 magnitude). If the check star magnitude agreement is poor, effort should be spent checking image quality and double-checking the analysis procedure for problems that may similarly degrade the accuracy and precision of the variable star target. The check star is 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?

5.6 Measure the magnitude

The magnitude system dates from the second century BCE, and is attributed to the Greek astronomer Hipparchus. It is logarithmic, and assigns smaller magnitude numbers to the brighter stars. Originally developed to classify naked-eye stars, it has been adapted in the telescopic age to measure optical brightness for many kinds of astronomical objects. There is a direct relation between magnitude and flux: 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.

Here in this chapter, we will work with two different kinds of magnitude measures: instrumental magnitude and standard magnitude. Although different, the two are interrelated and it's possible to convert from one kind to the other.

Instrumental magnitude is defined by the equation:

$$mag = -2.5 \log_{10}\left(\frac{integrated\ flux}{exposure\ time}\right)$$

Standard magnitude is defined as a relationship between two stars. There are standard stars across the sky whose standard magnitudes have been carefully measured. The standard magnitude of any star can be expressed relative to one of those standard stars using the equation:

$$(mag\,1 - mag\,2) = -2.5 \log_{10}\left(\frac{flux\,1}{flux\,2}\right)$$

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

Because instrumental magnitudes are not tied to any particular reference brightness, many factors affect the measured instrumental brightness. Telescope size, camera efficiency, atmospheric transparency, airmass (and, thus, observing altitude above sea level), and wavelength passband all affect the measured instrumental magnitude. For any given observation, there is an offset between instrumental magnitude and standard magnitude; that offset is called the zero point. When performing differential photometry, we simultaneously image our target and at least one comparison star of known standard magnitude. That enables us to calculate the zero point, which in turn allows us to convert *any* instrumental magnitude from that image into its corresponding standard magnitude. We can even average together the zero points from multiple images to find an averaged zero point with less uncertainty than any one zero point estimation from any one image.

Your software will probably convert flux (number of ADU within your measurement aperture) to instrumental magnitudes for you, but be aware that it may use arbitrary reference 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 reference point.* This is because the instrumental reference points cancel each other out when differential magnitudes are calculated. 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 the software, because the AAVSO Extended File Format requires reporting of some check and comparison star data using instrumental magnitudes.

The whole concept of brightness is intertwined with how we look at a star's spectrum. Brightness is measured across a part of the spectrum, with photometrists using color filters as a primary tool in determining which parts of the spectrum will be included in a particular measurement of brightness. Hence, when talking about the brightness of a star (its magnitude), we often include a statement of which standard filter was used to measure that brightness. So, if you look up the brightness of a star you will find multiple magnitudes listed, with each one labeled as to its standard measurement filter.

The software will measure the flux used in the equations above using the apertures discussed back in section 5.4. The ADU counts in the pixels contained in the inner measurement aperture are summed; this sum includes both the light of the central part of the star along with background light. The software will look at the ADU counts in the pixels contained in the outer annulus to determine a background brightness that is then subtracted from the inner aperture sum. That background-adjusted inner aperture sum is used to calculate the star's instrumental magnitude.

The known standard magnitudes of the comparison star (or ensemble of comparison stars) will then be combined with the measured instrumental magnitude of the comparison star (or ensemble of stars) to calculate the zero point (or averaged together to give a zero point). If a comparison star ensemble is used, you receive one mean estimated magnitude for your target star, which is generally more accurate than if only one comparison star was used. If it seems like one comparison star in the ensemble is noisy or has a problem that adversely affects your results, remove it from the ensemble and recompute the average. Again, though, postpone your use of ensembles until you've become proficient in photometry and estimation of your measurement uncertainties.

Once the zero point is established, it is added to the instrumental magnitude of every measured star in the image to transform those instrumental magnitudes to standard magnitudes.

Many software packages today keep this process hidden, and you may never have to deal with an instrumental magnitude. However, some software packages use a process that gives you responsibility for calculating the zero point and adjusting the instrumental magnitudes. Read the reference material available to you for your software package.

5.7 Measure the uncertainty

Your star magnitudes 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 to accurately calculate and submit magnitude uncertainty along with the magnitude itself. Uncertainty comes in two forms. Precision is a measure of the random error associated with identically repeated measurements (repeatability). Accuracy is a measure of the systematic error between your measurements and known (true) values.

Random uncertainty includes things such as photon noise (proportional to the square root of the number of photons your camera receives), and both read noise and thermal noise in your sensor. These noise sources need to be characterized. We attempt to reduce random uncertainty and improve precision by increasing SNR.

Systematic uncertainties measure how far your observed (calculated) magnitudes are from the true, actual magnitude of the star. Systematic uncertainties measure accuracy, not precision. For a slowly-varying star, making multiple measurements of its brightness lets you directly measure the amount of measurement scatter (precision) and calculate the mean magnitude. The difference between that mean and the actual magnitude of the star is the systematic uncertainty. A useful measure of accuracy is the difference between the observed and known magnitude of the check star.

All current measures of uncertainty in photometry are precision uncertainties. Accuracy uncertainties can only be judged by examining the fit of the result to a model (e.g, a light curve model) under the assumption that the model is robust (if not true) or by comparing the difference between your result and the accepted magnitude of a standard star (with due regard to random errors of both).

Most software either returns a precision uncertainty in magnitudes or provides the signal-to-noise ratio (SNR). A handy approximation is that the uncertainty in magnitudes is 1/SNR, so an SNR of 50 yields an uncertainty of 0.02 magnitudes. But this is not ideal because (a) the SNR is calculated directly from each image and doesn't 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, 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.

You have two choices. If you know the variable changes only slowly (a Mira star, for example), then you can calculate the magnitude of the variable on each frame, and calculate the standard deviation of those measures to give you uncertainty. (Remember that for a slowly varying star, you will go a step further and combine all the measures made on a single night into one magnitude instead of submitting a multiple-measurement 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 star instead. In all cases, you compute uncertainty using the standard deviation equation:

$$\sigma=\sqrt{\left(\frac{\sum (x_i-\bar{x})^2}{N-1}\right)}$$

where x_i are the individual magnitudes, $\bar{x}$ 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 in order to have a representative error value.

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*:

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

where N is the number of photons received from each of star, sky, dark current, and the readnoise of the sensor, 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 to just 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, N in the above equation is the number of photons, rather than the number of ADU, which is what your camera provides. This introduces a slight modification to the equation for ADU which includes the gain, G:

$$SNR = \frac{(N_{ADU})G}{\sqrt{((N_{ADU})G + n_{pix}((N_{ADU,sky} + N_{ADU,dark})G + (N_{readnoise})^2))}}$$

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.

Another option with single-image photometry (for advanced photometrists) uses multiple comparison stars (an ensemble) in the frame. In this case, you measure all of the comparison stars along with the variable, calculate the magnitude of the variable obtained using each of the comparison stars, and then calculate the mean and 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 low SNR.

Chapter 6: Transforming your data

6.1 Why is transformation necessary?

The AAVSO International Database is composed of data collected from many different observers, at different times, from around the globe. The beauty of such a system is that it allows all interested observers to contribute to the archive, thus it has a great deal of potential to expand the duration and breadth of coverage for the target stars. Unlike data collected through surveys, which can experience coverage gaps due to bad weather conditions, equipment failures or discontinuation of funding, the AAVSO approach reduces the effect of such problems. On the other hand, the fact that each observer is using different equipment and procedures can introduce offsets which are difficult to reconcile from one observer to another.

Assuming that the procedures outlined in this *Guide* have been followed carefully, and no mistakes have been made along the way, the largest remaining differences between measurements reported by two different observers looking at the same star with the same filter at the same time is likely caused by differences in the color response of each observer’s equipment. Each telescope, filter, and camera combination has its own unique characteristics, which, depending on the color of the star being measured and the filters used, can result in magnitude differences of anywhere from several hundredths of a magnitude to a few tenths of a magnitude from one observer to another. Even two photometric filters purchased from the same vendor may have a slightly different spectral response that will affect your measurements!

By transforming your data to a standard color system, these differences can be greatly reduced. This will both bring your observations more in line with those of other observers who have transformed their data and also make the whole database more scientifically useful. It is the goal of the AAVSO to get all photometrists to transform their data as a matter of course.

6.2 The Two Steps of Transformation

There are two parts to the process of transforming your data:

- **Determining transformation coefficients:** This step is performed relatively rarely (a few times a year) to calculate the numerical coefficients that capture the difference between your system's response and the standard color system. You will do this by measuring the brightness and color of many stars of known brightness and color (a standard field) and comparing your measurements (through each of your filters) against the known brightness of each star.
- **Applying transformation coefficients:** This step is performed after each observing run as an integral part of processing each observation. You use the coefficients that you found earlier to adjust each of your measurements to align it with the standard color system.

6.3 General overview and assumptions

For consistency with the rest of this *Guide*, the explanation that follows assumes you are performing differential photometry and that you have at least two filters (the minimum required for a normal transformation).

In astronomy, the color of a star (or color index) is expressed as the difference in magnitude measured with two different filter bandpasses. So, for example, the color index B-V is obtained by subtracting the V magnitude from the B magnitude. By convention, the shorter wavelength filter letter is always given first. Some color indices are used more often than others. The most widely used measure is B-V (i.e. the magnitude measured using a Johnson B filter minus the magnitude measured using a Johnson V filter). The smaller the B-V color index, the bluer the star. Color indices can be positive or negative numbers. The bluest stars are generally associated with the most negative color indices.

Color indices are important throughout astronomy. They give a hint at the nature of each star's spectrum by describing the relative strength of starlight over wavelength regions. So, for example, a large positive B-V color says the spectrum is "humped" toward the red and we are describing a cool, red star. Further, because color indices measure the *difference* between magnitudes, color indices tend to be somewhat less sensitive to zero-point errors.

Physically, there is a color index for every filter *pair* that you have. If you have two filters (e.g., B and V), you are dealing with a single color index (one pair: BV). If you have three filters (e.g., B, V, and R), there are three pairs (BV, VR, and BR). With four filters there are six pairs, and with five filters there are ten pairs. Each color index is a difference in magnitude, and we give each color index a name that shows which two magnitudes were subtracted for that color index. However, not all filter pairs have the same astronomical significance. Some color indices simply don't get used much (for example, B-I) and are usually ignored and not even computed.

When you make brightness measurements with multiple filters, in addition to measuring a brightness (magnitude) for each filter, you also effectively measure a color index for each filter pair. Similarly, there are two kinds of transformation coefficients that you will calculate: one set of transformation coefficients that is used to correct your measured color indices, and one set that is used to correct measurements of individual filter magnitudes.

And at this point, you may very well be scratching your head, since adjusting individual filter magnitudes will automatically change the color indices that you calculate (when you subtract the magnitudes of two different colors), and vice versa (since adjusting the color indices implies a change to the individual magnitudes). Indeed, magnitude and color are intertwined and tangled. They are just different views of the same data. This interdependence adds a complication to the mathematics of transformation that will be addressed later.

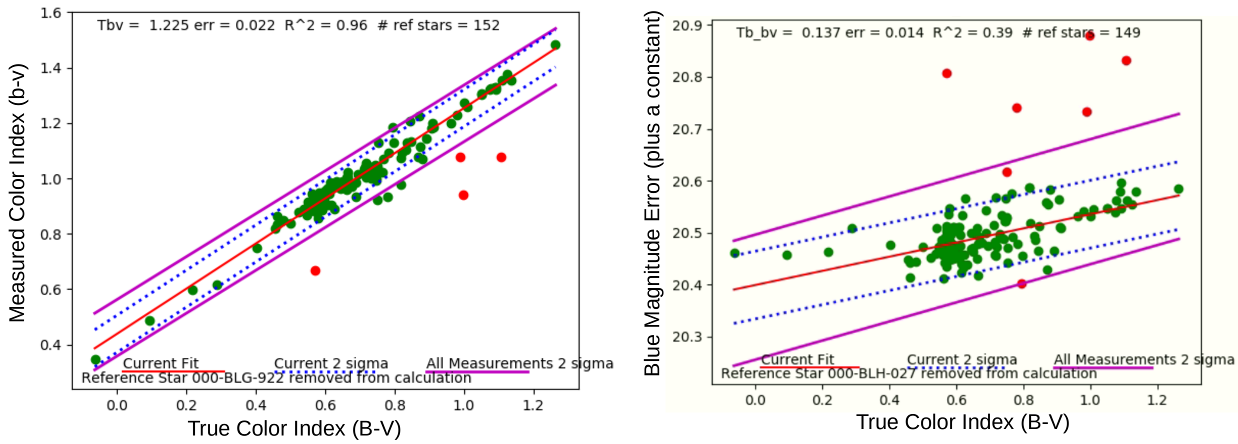

Figure 6-1 — Examples of the two transformation graphs. The graph on the left is for a color transformation coefficient and plots the measured color against the known (true, standard) color. The graph on the right is for a magnitude transformation coefficient and plots magnitude error (plus a constant) against the known color. In each case, the transformation coefficient itself is derived from the slope of the best-fit (red) line through the points. The purple lines show the standard deviation of the points around the best-fit lines. (Illustration courtesy George Silvis)

6.4 The Transformation Coefficient Graphs

The process of determining your transformation coefficients begins with an observing session where you image a standard field (of stars with known brightness) using all of your photometric filters. For each star, you will generate a set of brightness measurements, and from those magnitudes you will calculate your observed color indices for each pair of your filters.

You'll then create two kinds of graphs as seen in Figure 6-1, one that plots your measured color indices against the known color indices for each star and one that plots the difference between your measured magnitude and the known magnitude against the known color index.

If your system is already perfectly aligned with the standard color system, the first set of plots (measured color index against known color index) will have a slope of exactly 1.0. Similarly, if your system is already perfectly aligned with the standard color system, the second set of plots will have a slope of 0.0, since error will be completely independent of color.

The AAVSO Transform Generator (TG) tool will help you measure these slopes. It will take your brightness measurements, compute your measured color indices, look up the known color indices, and generate all the plots for you. You will interact with the TG tool to look at each plot, assess the plot quality, and delete any points that are very inconsistent with the others (such as the red points in the graphs above). Once you have finished reviewing and editing each plot (and there can easily be a dozen plots if you have four or more filters), the TG tool will compute the slope of each line to give you the value of the corresponding coefficient; it will also use data scatter in the plots to estimate the precision (uncertainty) of the coefficient. The resulting set of coefficients will be put into a data file for storage.

The TG tool will allow you to compare coefficients measured on different dates and to average coefficients from different observing sessions. Your final set of coefficients will be formatted and stored for use in your daily observing sessions. Documentation for both the TG tool and the Transform Applier tool are available at https://www.aavso.org/transforms-everything-you-need-transform-your-ccd-observations.

6.5 Naming the Coefficients

Recall that there are two types of transformation coefficients: ones that adjust *color* and ones that adjust *magnitude*. The naming convention for transformation coefficients is different for each type:

- For color transformation coefficients, the name used is Txy, where x and y are the one- or two-letter names of the filters. For example, the color transformation coefficient that relates the measured b-v color index to the known (true) B-V color index is known as Tbv. (Notice, also, that we use upper case filter letters to refer to standard system measurements and lower case filter letters to refer to our untransformed (instrumental magnitude) measurements.) Sometimes the filter codes require two letters each, so Tsgsr is the color index associated with the difference between Sloan green and Sloan red magnitudes.
- For magnitude transformation coefficients, the name used is Tx_yz, where x is the one- or two-letter name of the filtered magnitude that is being adjusted and yz is the 2-4 letter name of the color index used as the basis for the adjustment. For example, the magnitude transformation coefficient that relates the instrumental b magnitude to the known (true) B-V color index is known as Tb_bv.

6.6 Applying the transformation coefficients

Although you *compute* your transformation coefficients relatively rarely, you will *apply* them every night. You apply these coefficients to transform your measured target data to the standard system.

The series of graphs that gave you information about how your system's colors correlate to standard colors and how your magnitudes need to be adjusted based on standard color provide all the information that's needed to perform the inverse calculation: adjusting your measured colors and magnitudes to match the standard. Actually, you've got *too many* parameters; you only need to calculate as many corrections as you have filters (one correction per filtered magnitude), but you probably have at least double that number of transformation coefficients.

The application of the transformation coefficients can be complex for at least three reasons:

- The equations that are used need to be solved for a different set of unknowns than when the equations are in the form that generates the transformation coefficients.
- There are more equations available (and more transformation coefficients) than unknowns that you are solving for (i.e., the system is over-constrained).
- The equations are based on color differences between pairs of stars, but if you're using a comparison star ensemble, you have more than just the single pair of stars that you deal with when there's one target star and one comparison star. (This is one of the reasons that you should stay away from ensembles as a beginner.)

Hence, this part of the process can involve a daunting amount of arithmetic if you do it by hand. (The equations are provided in Appendix D to this *Guide*.) However, both the VPhot-TransformApplier package and LesvePhotometry can be trusted to apply your transformation coefficients correctly as a normal part of completing analysis.

(Note, though, that VPhot currently (2022) provides two different methods for applying transformation coefficients, only one of which works with multiple comparison stars (a comparison ensemble). The two methods do not use identical equations, and apply slightly different transformation corrections. One method is called "TransformApplier" and works extremely well when there is a single comparison star. The other method is called "Two Color Transform" and works with ensembles of comparison stars, but requires some non-intuitive application to use if more than two filters have been used.)

Chapter 7: Assessment and Improvement

Whether you are a beginner or an experienced photometrist, you will eventually find that the relentless pursuit of self-assessment and improvement becomes second nature. The aim is always twofold: to understand the sources (and amounts) of error contributions to your photometry and to create a roadmap that leads to meaningful reduction of those errors. Photometrists find that a critical step to improve accuracy is to transform your data to a standard color system.

Three activities are fundamental to this process:

- Assessment of your measurement error
- Establishing root cause
- Improvement

All three activities are needed, and aren't necessarily completed in sequential order because of some interaction across the three. Further, you will find yourself iterating through the activities repeatedly; as one error source is reduced, other error sources will grow in importance and require reassessment (and finding the new root cause).

Because your techniques will evolve as your experience grows, we start this chapter with a set of basic techniques that don't require any special analysis techniques compared with what you are already doing as part of your normal analysis and reporting cycle.

7.1 Basic Analysis Techniques

For every photometric measurement you make of a star (variable, comparison, or check star), at the moment of your measurement, that star has some actual brightness that would be recorded by a perfect photometrist using a perfect telescope and camera system. The difference between your measurement and that actual brightness is the error in your measurement. Since the brightness of the variable is unknown to you and since you've already used the comparison stars as the basis for making your measurement, your check star measurements provide a great estimate of measurement accuracy.

Six straightforward techniques for assessing (and improving) your photometric performance are described in the following paragraphs.

7.1.1 Learn

There are several resources that make a great starting point for improving your photometry:

- CHOICE course: CCD Photometry

- CHOICE course: Photometry with VPhot
- CHOICE course: Fundamental Statistics for Photometry

The more you know, the better you'll be able to assess your own data and recognize the problems that will inevitably occur. Look for training opportunities for the software used in your personal photometry workflow.

7.1.2 Make multiple measurements of your target star

If you are observing a target with a time series in order to capture short-term variations, then you are already making multiple measurements of your target over the course of the session. If not, then plan your exposure sequences so that you can make multiple transformed measurements; you can then compare those measurements and get a good estimate of measurement uncertainty (precision) by looking at the standard deviation of those measurements.

7.1.3 Use your check star

Your check star should already be in the AAVSO sequence for your target. That sequence gives you an expected magnitude for the check. How closely does your measurement for the check star compare with the value in the AAVSO sequence? This gives another estimate of uncertainty.

7.1.4 Add standard fields to your observing plan

The standard fields (discussed further below) are spread across the sky; you can always find at least one field convenient to your location, date, and time. The standard fields take no longer to acquire photometric images than do AAVSO program stars; add a standard field to your observing plan regularly. From your standard field images you can easily measure both precision and accuracy by looking at how your measurements change from image to image and by comparing your measurements to the known standard magnitudes for those standard fields (downloaded easily from the AAVSO VSP).

7.1.5 Transform!

If you aren't already doing it, transform your data (Chapter 6). Photometrists find that this is a relatively easy step that consistently improves your measurements.

7.1.6 Compare against the light curve

Make it a standard part of your process: compare every measurement you make against the AAVSO light curve for that target. Highlight your own observations and look for systemic offsets between your data and that of other observers.

7.2 Expanding Your Program: Estimating Accuracy and Error

There are at least two different sets of error values that you're looking for: an *accuracy* value and a *precision* value. The precision (*repeatability*) is far easier to measure: make a series of measurements (images) of the same star under steady-state conditions and see how much the measurements differ. (This variability is defined as the standard deviation of the measurements.) Be aware that your precision measurement is unique to that particular night, field, and equipment configuration. Changing any of the following changes your precision:

- Different exposure times (longer exposure times will improve SNR),
- Different color filters (affects the efficiency of light collection and SNR),
- Different sky background (light pollution or moonglow levels change noise levels),
- Camera settings (gain, readout mode, offset).

One way to assess your photometric accuracy is to compare your measurement with the measurements of other people. There are three ways to do this, and you can (and should) embrace all three:

- **Expand your observing program to include photometry of standard stars in standard fields (e.g., Landolt fields – see below)**; these standard stars have had their brightness carefully measured and published so that you can easily compare your results.

- **Compare your measurements of a check star** during everyday photometry with the photometric sequences distributed with AAVSO charts. There is usually more uncertainty in the published magnitudes of these sequence stars than in the data for standard stars, so as your photometry skills improve, at some point your accuracy may be better than the published accuracy of the check stars.

- **You can compare your variable star photometric results with the results of other observers** using the AAVSO light curve generator and data download tools.

7.2.1 Standard Fields

Several hundred non-variable stars have very well-known brightness. These particular stars have been measured dozens or hundreds of times by experienced photometrists using well-calibrated systems. The AAVSO maintains charts for these standard fields that include this accurate photometry. A list of the standard fields directly supported by the Variable Star Plotter is at https://www.aavso.org/standard-stars-vsp. (There are two lists there, one for northern observers and one for southern observers.)

Using your equipment, you can measure the brightness of these standards and compare your results to the known standards. These standard fields are an excellent way to assess your own total error; the AAVSO recommends that you incorporate photometry of standard fields into your day-to-day observing cycles.

There are several types of standard fields. Some are standard clusters, and others are field stars, not located in clusters. Many of these latter stars use photometry from Arlo Landolt, and are known as Landolt standards, or Landolt fields.

As you choose the standard fields that you will use regularly, be aware that some of the standard clusters have typical star separations measured in arcseconds near the cluster center. The presence of multiple stars within your photometry measurement aperture becomes a source of error. The Landolt standards cover a relatively wide range of star magnitudes; avoid standard stars that are outside the range appropriate for your equipment.

7.2.2 Comparing Photometry With That Of Other Photometrists

This can be done either before submitting your data to the AAVSO or after you've done so. The best way to compare is to launch the light curve generator (LCG2) for the star you'd like to compare. If you are comparing data you've already submitted, you can highlight your own data by going to the "Contributors to this Plot" area under the graph and selecting your observer code. Note that this is accuracy relative to a model and not a direct assessment of accuracy relative to truth. However, conforming to a model is the best we can do with variables.

Some questions you can ask yourself as you assess your data:

1. Does my own data show more scatter than photometric data from other observers? (Exclude visual observations as you make this assessment.)
2. Does my own data show any systematic trends compared to other observers' photometry? Is my data consistently brighter or fainter than others? (If so, this is often an indicator of a comparison star inconsistency or a filter/transformation issue.) (As an example, see the lightcurve fragment for V Aur on the next page.)
3. Are any of my data points wildly inconsistent with the community consensus? (Possible causes: stuck filter wheel, photometry aperture shifted off of correct star center, star misidentification) (As an example, see the lightcurve for RZ Cyg below.)

If you've done this step correctly, you now have some estimates of accuracy (bias) and precision (repeatability).

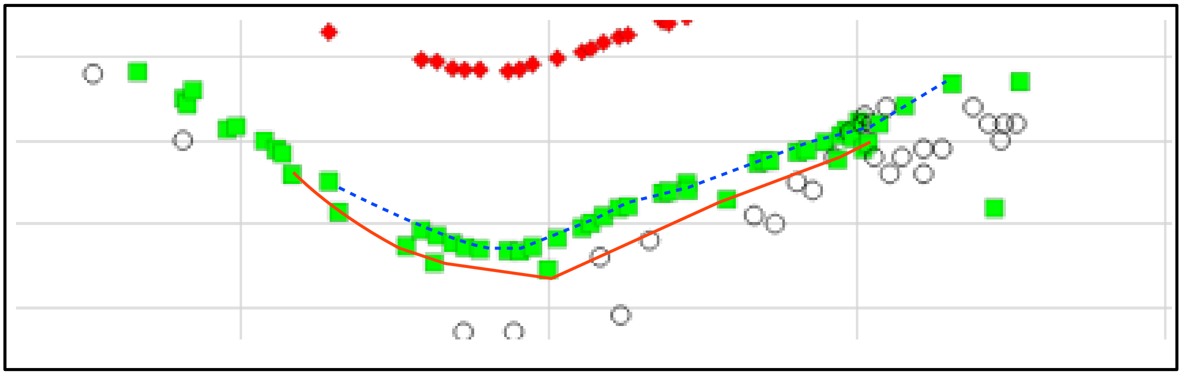

Figure 7-1 — V-band photometry from two different observers for V Aur shows a distinct, systemic offset between the two observers of several tenths of a magnitude.

7.3 Establishing Root Cause

Unfortunately, your accuracy and precision values alone won't be sufficient for you to know what to do next. You will have to go through an analysis and experimentation process to determine the root cause of the errors.

Errors never come from a single source. There are many contributors to photometric error, and they generally never go away completely. At any given moment, though, one contributor is bigger than any other contributor. You fix the big one, it shrinks, and a new error contributor becomes the biggest.

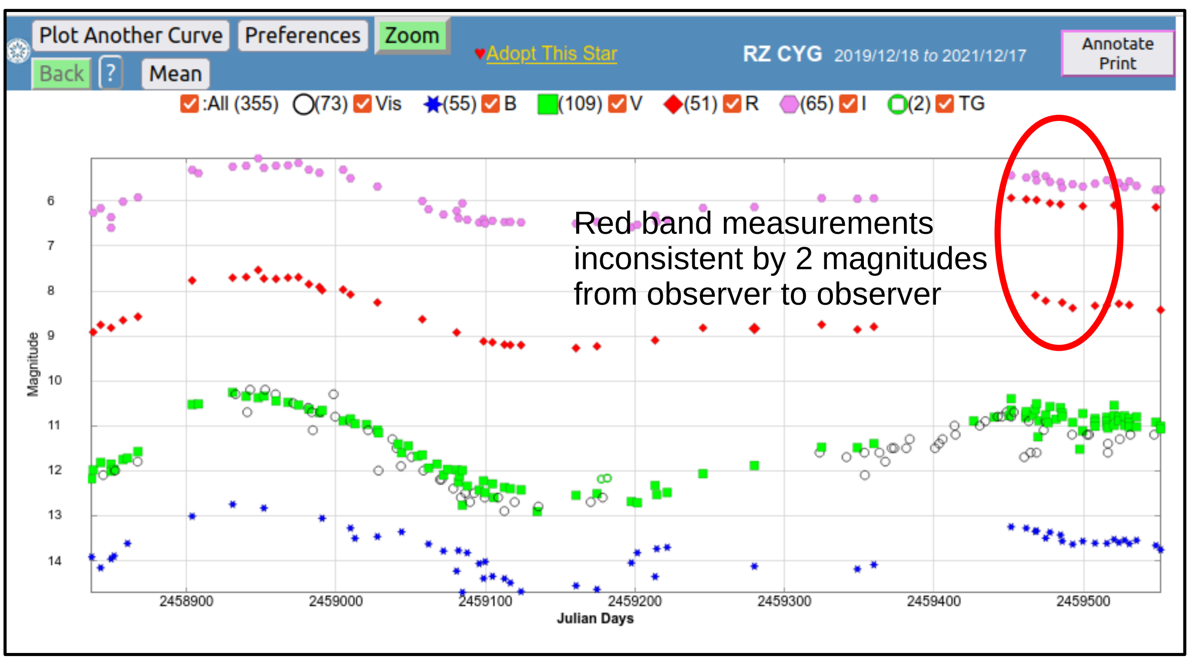

Figure 7-2 – RZ Cyg lightcurve showing inconsistency suggesting mislabeling of images or use of incorrect filter.

7.3.1 Learning From Standard Field Measurements

1. Are your errors the same on a given star every night? If so, your error contributions due to poor guiding, poor focus, wind, and low SNR are not as great as errors from other sources.
2. Are any trends apparent? Are faint stars typically measured as brighter than the standard magnitude? Dimmer? Are stars near the edges of your images measured as brighter than they should? If so, the particular trend will give you a hint as to the corresponding error source (see InfoBox 7-A).
3. If some stars are consistently off by more than others, look at the colors of those stars. Errors that correlate to color are a potential indicator that color transformation coefficients are either not being used or have been miscalculated.

InfoBox 7-A – Establishing Root Cause: Coma and Field Flatness

A good (and relatively easy) experiment to explore root cause creates a plot of FWHM as a function of distance from the center of your images. Here is a typical plot for an image made by a 14-inch Schmidt-Cassegrain with no coma corrector or flattener in the optical path.

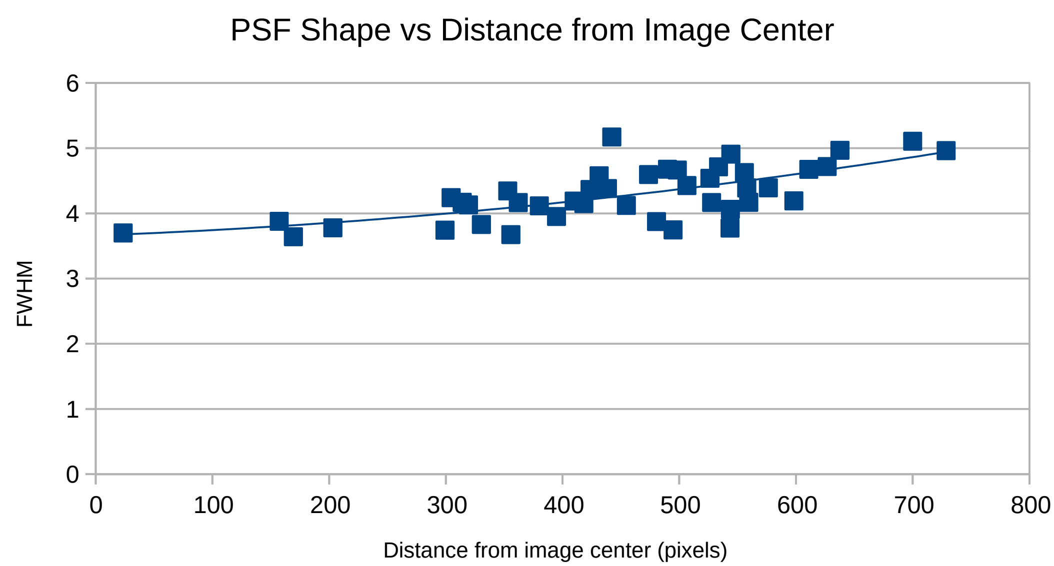

This systematic extra fuzziness on stars displaced from the image center is a cause of photometric error. The error can be minimized by either restricting photometry to the center of the image or by adding a coma corrector and field flattener into the optical path (which was the chosen solution, which in turn added more vignetting (another source of photometric error) that had to be addressed through the flat-fielding process).

4. Make plots of error vs. star SNR and random error (repeatability) vs. SNR. Random error should increase as SNR drops and should be somewhat consistent with the photometry software's estimate of measurement precision.

In all probability, the dominant error source will not be obvious. As a next step, consider giving yourself a periodic report card. Assign a grade for each category, perhaps using the report card template shown in Appendix E. Save these report cards. Fill out a new report card after each major system change (after you've optimized your procedures for the new configuration), or at least annually. Based on your photometry program, you may want to delete some topics from the report card or add new ones. Keep the report card meaningful and honest. Use it to help collect your thoughts about the most meaningful changes you can take to improve photometric quality.

The process of filling out the report card may trigger questions in your mind. One useful part of the process of completing the report card is creating experiments specifically to isolate or rule out specific error sources. For example, if you suspect a detector nonlinearity may be degrading your photometry, you may choose to perform an experiment to measure linearity using your usual configuration for making flats.

Remember that being able to rule out error sources is an important step toward establishing the underlying root cause(s) of your errors.

7.4 Improvement

Sometimes improvement is easy once the dominant error source is known, but sometimes it's hard. For example, if you can determine that the root cause for grossly incorrect measurements belongs with a filter wheel that sticks in cold weather, getting a filter wheel tuneup may completely solve that problem. Or maybe you realize that your photometric errors are dominated by poor target star SNR, and all you need to do is increase exposure times. Your Appendix E "report card" is your friend in this process; it should help make clear what isn't a candidate (for the time being) for improvement.

But sometimes the process includes the challenge of reconciling household financial budget with your photometric error budget: a complete set of photometric filters can be expensive.

This is where we encourage experimentation. Take small, incremental steps, if at all possible. Gain confidence that you have correctly identified the root cause. Process images using a completely independent software suite and see how your results change. Process someone else's images and see how your measurement error compares.

And above all else, take the long view. Never stop pursuing improvement. Never stop assessing.

Chapter 8: Photometry and science

The first seven chapters of this *Guide* give you everything you need to make variable star observations 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 this Chapter 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 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.

8.1 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 were designed to closely approximate a standard system, such as the Johnson-Cousins or Sloan system band passes. 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 band pass of the filter.

Filtered photometry provides 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 magnitudes were initially defined in the Johnson system. The B-V color index of a G-type star, cooler than an A-type star, will be somewhere around +0.7, meaning 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 magnitude-color (Hertzsprung Russell) 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 brightness varies. Remember that colors may correspond in part to the temperature of a star. We also know that some stars change color and temperature during the course of their variations. A pulsating star like a Cepheid or RR Lyrae can change by 1,000ºK 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 B and V, 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 (among other things) 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 index 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 others' 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.

8.2 Time considerations: variability timescales, exposure times, and cadence

If you've been a variable star observer for a while, 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 exposure times and cadence for every star. The four 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 neither over-observe a star whose timescale of variation is very long, nor under-observe a star whose timescale is very short.
4. It's appropriate to collect about 100 data points over the full period of each variable star to provide a reasonably well-defined light curve.

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 measuring bright stars, brighter than 7th or 8th magnitude for many typical systems using moderately-sized telescopes. Pay attention to scintillation (section 4.7.1).

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.

Point (4) is a recommended rule of thumb that yields measurements which identify light curve changes of about 1%.

As an example, consider a slowly varying target star. 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 digital camera 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 photometrists. 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 observer could potentially create some wonderfully useful data sets for lots of stars.

Sometimes, you may encounter the exact opposite case: you 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 (Figure 8-1).

These data were taken through a clear filter using a 0.4-meter (16-inch) telescope. When the star is between magnitude 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 one exposure per 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.

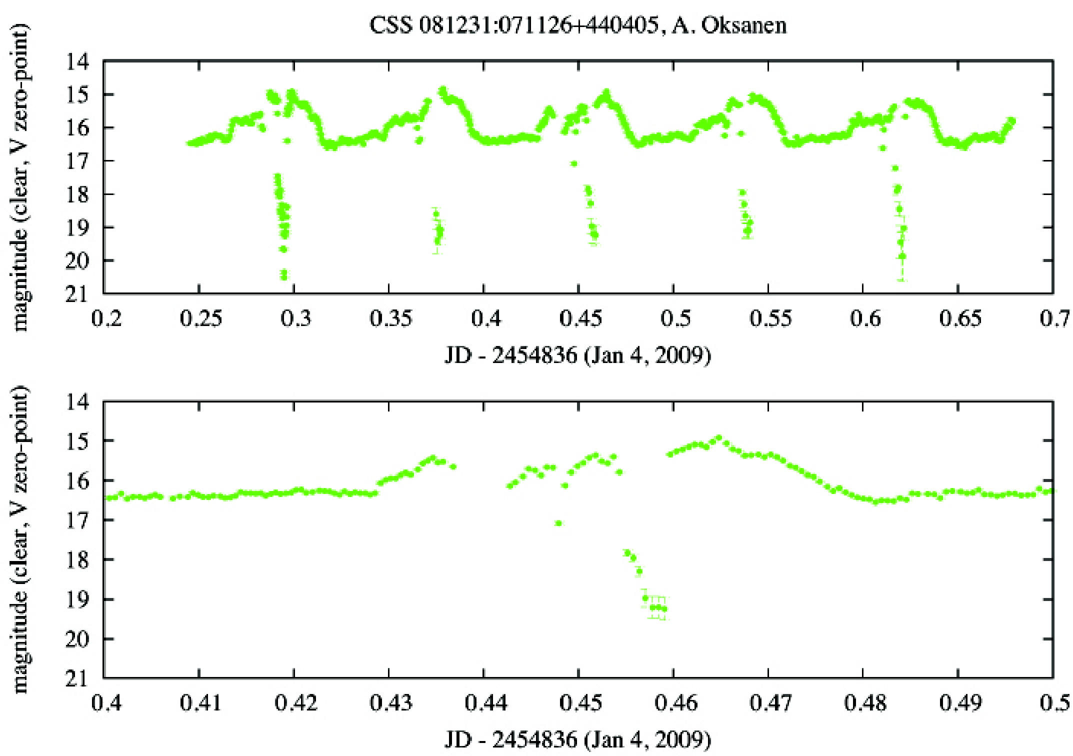

Figure 8-1 — Unfiltered time series of an eclipsing AM Herculis-type cataclysmic variable. Note that the error bars are very small, and note also the number of times observations are made. The observing cadence is approximately one observation per minute, including both exposure and chip readout.

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:

- If the target is bright and you can get good SNR with an appropriate exposure time, you should always use filters if you have them. (Note that "good" will be defined by your project goals, but >20 is a reasonable SNR value.)
- 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). Its spectrum 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.

8.3 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 camera is likely sensitive in the far red, 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 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.

Introduction To The Appendices

Readability was emphasized during the writing of this *Guide*, and each chapter has been kept focused on a set of core, fundamental concepts. To maintain the flow of these chapters, a number of concepts that are relevant and useful have been gathered into a series of appendices to the *Guide*. In each case, you can find references to these appendices within the body of the *Guide*.

In some cases, these appendices provide additional detail for the curious reader (Appendices A and B on the nature of starlight). In other cases, the appendices provide detailed instructions (Appendices C (submitting observations), E (self-assessment report card), and F (getting started)). Appendix D explores color transforms in more detail; depending on the software you use, the details in this appendix may not be needed by a beginner. And Appendix G (resources and links) and H (glossary) can help clarify what you've read on these pages.

Appendix A: What is starlight?

There's much more information in starlight than how much of it there is and when you measure it. We ask observers to use standard filters when doing photometry because filters allow you to measure both the amount of light, and its spectral distribution. The key physical property of light relevant here is the wavelength. Light is composed of *photons*, which are small bundles of electric and magnetic fields that travel through space at the same speed — the speed of light, **c**. These small bundles behave both like particles and waves, and since they are waves, they have a characteristic wavelength.

In optical light, the different *colors* you see correspond to light with different wavelengths. Red light has longer wavelengths than yellow light, which has longer wavelengths than green light, which has longer wavelengths than blue and violet light. All of the different colors of light observed together are called a *spectrum*. The visual spectrum is roughly composed of all light having wavelengths between 300 and 700 nanometers, from the violet to the red. There's more light beyond that range, too. Beyond the violet toward shorter wavelengths lie the ultraviolet, X-ray, and gamma-ray regions of the electromagnetic spectrum. Beyond the red toward longer wavelengths lie the infrared, microwave, and radio regions. We only define the visual spectrum this way because that's what the human eye is capable of seeing — our eyes aren't sensitive to light outside that range. Most normal stars emit the bulk of their light in the optical and infrared, and our own Sun emits the greatest amount of light around 500 nanometers, which appears green to our eyes.

A related quantity for each photon is its energy, which is also a function of wavelength. Specifically, the energy carried by a photon is inversely proportional to wavelength:

$$E = hc/\lambda$$

where h is Planck's constant, c is the speed of light, and λ is the wavelength. Note the inverse relationship with wavelength: shorter wavelength, blue photons have more energy than longer wavelength yellow photons, which have more energy than even longer wavelength red photons. The wavelengths of light that astrophysical sources emit are related to the total energy density of the system that's doing the emitting. A relatively cool star is unlikely to emit high-energy radiation *unless there are special sources of energy within the system*. Conversely, a hot star may be capable of emitting higher energy radiation, *but it also emit photons with lower energy*. (More on that in Appendix B.)

There is another property of light that we won't go into detail in this *Guide*, and that is its *polarization.* Photons are bundles of electromagnetic radiation, where each particle consists of an oscillating electric and magnetic field. All photons received from a single source may be assumed traveling in parallel when they get to your detector, but the oscillation axes of each photon will be different. The fields may oscillate in a single direction perpendicular to the direction of motion but with random orientation, or they may have a circular component to the oscillation (i.e., the photon is elliptically or circularly polarized). If the emitting source is polarized or if the light passes through a polarizing medium (like a dust cloud), there will be a preferential orientation for most photons you see. Circularly polarized light can also be created in environments or physical processes having strong magnetic fields.

Polarization can be measured with special filters, but it is a time-consuming process. We won't discuss it further, but be aware that it is another fundamental property of light that you observe.

Appendix B contains a brief discussion of radiative processes common in stellar astronomy, and how these can be described or explored using photometry.

Appendix B: Why and how stars radiate

Both the amount of light generated and the wavelength spectrum of light that an object like a star emits will depend on the physical properties of what's emitting the light. The spectrum of starlight is generally very complex on close examination, but the physics responsible for it can be broadly generalized into two processes: continuum emission, and line emission and absorption.

Continuum emission is any physical process that emits photons with a broad range of different wavelengths. As an example, think of what you see when you hold a prism in sunlight — you see several bands of color with red, orange, yellow, blue, indigo, and violet. All of those colors are present in sunlight at the same time, but you don't see them individually — the Sun simply looks white.

B.1 Blackbody radiation

A special kind of continuum emission is blackbody radiation, emitted by all objects — any objects — with temperatures above absolute zero. The amount of light and the wavelength distribution of photons in the blackbody spectrum depend on one parameter: the temperature. The key things to remember are, for two identically-sized stars, if one star is hotter than another, (1) it will emit more light overall, and (2) the spectrum of light it emits will have more light at shorter wavelengths. If you have two stars whose physical sizes are the same and are the same distance away from us but one is at 10,000 K and the other at 5,000 K, the hotter star will be brighter (more light), and bluer (more emission at shorter wavelengths). Thus you can use starlight to take the temperature of a star without touching it — a neat trick! The equations describing blackbody radiation were worked out by Max Planck early in the 20th century, and you'll often see blackbody radiation referred to as *Planck radiation*.

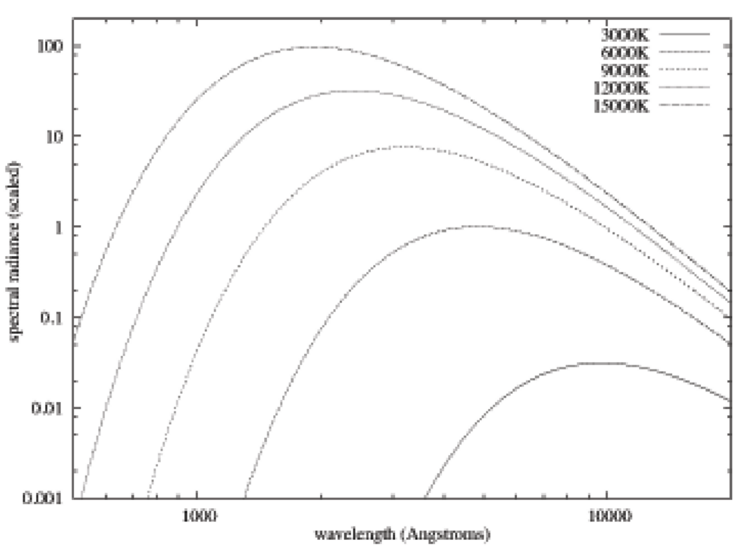

Figure B-1 — Black body spectra scaled to the peak spectral radiance of a blackbody at 6000 Kelvin. The Sun's effective temperature is about 5774 Kelvin. That of an A0 star is about 10,000 Kelvin, while that of an M star is below about 4000 Kelvin. Compare the bandpasses of filters shown in Figure 3-1 with the curves shown here.

There are a few concepts related to blackbody radiation that are very useful in stellar astrophysics. First, Wien's Law is a simple equation that gives you the wavelength at which a black body emits the most light (i.e., the peak of the blackbody spectrum):

$$\lambda_{max} = b/T$$

where λ is the wavelength, T is the temperature of the blackbody, and b is a constant known as *Wien's displacement constant*. You can derive this by taking the equation of a blackbody and determining where the curve is maximum: you determine the temperature and wavelength at which the derivative is zero. This is a really handy equation, because it lets you roughly estimate the temperature of any blackbody-like object by simply measuring the wavelength of the peak of its spectrum. Many stars behave so similarly to blackbodies that this is straightforward to measure; where it breaks down are for stars that have such strong atomic or molecular absorption that their optical spectra don't match a blackbody very well. (This often happens for M stars whose spectra peak in the near-infrared anyway.)

Another relation is the Stefan-Boltzmann Law, which provides a simple relationship between the energy flux per unit area from the surface of a black body and its temperature:

$$f_{bol} = \sigma T^4$$

where f_{bol} is the total energy flux per unit area, T is the temperature, and σ is a constant (the Stefan-Boltzmann constant). The hotter a blackbody gets, the more total energy it emits. Again, this yields another interesting astrophysical application. You may be able to estimate a star's effective temperature by some means (photometric, or spectroscopic). The total luminosity (the light emitted in all directions) by a blackbody is simply this quantity f_{bol} times the total surface area: $4\pi R^2$. Combining these two things, you get the interesting equation

$$L_{bol} = 4\pi R^2 \sigma T^4$$

There are a few potentially interesting quantities there, namely the luminosity (which can be tied into the distance to the star) and the radius of the star. This is important astrophysically; the luminosity of a star is proportional to both its effective temperature and to its radius. Spectral types also include *luminosity classes* from *dwarf* to *supergiant.* A star might have an effective temperature of 4000 K, but there will be a huge difference in luminosity depending on whether its radius is that of a dwarf star or a supergiant.

B.2 Line emission and absorption

The second process, line emission and absorption, are two things caused by the same physical process: the emission or absorption of individual photons by atoms. Atoms are composed of nuclei (protons and neutrons) surrounded by electrons having very specific orbits. The orbits of these electrons correspond to specific energy levels. If an electron transitions from a higher energy level to a lower one, it will release the resulting energy difference as a photon with that energy. Since wavelength corresponds to energy, these electron transitions correspond to specific wavelengths of light. These wavelengths (or combinations of wavelengths) are unique to each atomic species. If you have a sample of hydrogen gas and excite it (say in a fluorescent tube), it will emit light at several discrete wavelengths corresponding to the electron energy levels of a hydrogen atom. Likewise if you have a sample of nitrogen, sodium, or neon gas (all common in fluorescent bulbs) they'll have different spectra. (This is why "neon signs" have different colors — they use different gases.)

The inverse of emission is absorption: if you have a photon of the right wavelength to excite an atom that has an allowed electron transition with just the right energy, the atom will absorb the photon. If you have a source of continuum emission (like the photosphere of a star) along with some gas that can absorb energy (like hydrogen, calcium, iron, or other elements in a star's atmosphere), the star's spectrum would look like a blackbody with some specific wavelengths reduced or missing. So when you take a spectrum of a star, you'll see mostly a continuum of light, but with dark bands appearing along the dispersion axis. The amount of absorption that you see depends on many different factors including the abundances of different atomic and molecular species, and the temperature of the star. A-type stars for example are defined as having the strongest absorption lines of hydrogen in their spectra. As another example, molecular absorption occurs in cool, M-type stars, and the kind of absorption you see depends on whether the star's atmosphere is richer in oxygen or carbon.

The astrophysics of radiation and radiative transfer is a very rich subject. Much of what was discussed above was laid out even before the golden age of quantum mechanics by the 19th century physicist Gustav Kirchhoff, and are summarized by Kirchhoff's three laws of radiation:

1. Hot, solid (or optically thick) objects emit a continuous spectrum.
2. A hot, optically thin gas emits light at discrete wavelengths characteristic of the chemical composition of the gas.
3. A continuous spectrum passing through a cool, optically thin gas will show absorption lines characteristic of the chemical composition of the gas (and at identical wavelengths to the emission lines that would appear if the gas were hot).

Kirchhoff outlined these rules in the 19th century, before atomic physics and quantum mechanics were understood. But for many cases of interest in variable star astronomy, these rules broadly describe everything you'll see, and the mathematical models of how light is created and how it propagates in a physical system are rooted in Kirchhoff's laws.

We won't cover spectral analysis in this *Guide*, but it is possible to use observation and measurement of the strengths of spectral lines in a star to figure out what the star is made of. Atomic line measurement in the laboratory was and still is a major field in laboratory astrophysics. Absorption and emission lines will change their appearance in a complicated way that depends on relative abundances in the plasma, the temperature (and temperature structure when looking through a thin gas), and pressure. Some lines and groups of lines are so strong and prominent that they serve as proxies for the overall "metal abundance" (i.e. the abundance of everything except hydrogen and helium). In some cases, these can be so strong that they can even be detected in broadband light, and thus can be detected with filtered photometry rather than spectroscopy.

B.3 Other processes

There are other sources of radiation, including magnetic fields (especially important in active stars that generate X-rays), nuclear reactions, and radioactive decay (which power the interiors of stars and are also responsible for the energy that powers supernovae and their light evolution). Many variable stars will have multiple sources for radiation and absorption. As an example, the UV Ceti stars are low-mass, young, dwarf M stars, usually very cool. These objects are generally very faint since their cool temperatures mean they radiate a relatively small amount of light, mostly in the red and infrared. However, they can also emit enormous amounts of blue, ultraviolet, X-ray, and even gamma-ray radiation in very short bursts due to magnetic reconnection events in their atmospheres analogous to solar flares on our own Sun. These stars are naturally very faint in blue, so when large flares occur, they may have enormous amplitudes in blue light, but relatively little in red. A bright flare may have a B-band amplitude of 3 or 4 magnitudes, but much less than a magnitude in R- or I-band.

The physics of radiation is one of the earliest courses a student in astronomy would take, and while it isn't required to be an observational astronomer, knowledge of radiative processes may provide you with some insight into what you're observing. One particularly useful book on the topic is George Rybicki and Alan Lightman's *Radiative Processes in Astrophysics*. A detailed reference on spectral lines and stellar spectra is David Gray's *The Observation and Analysis of Stellar Photospheres.*

Appendix C: Submitting observations to the AAVSO

Submitting observations to the AAVSO — whether obtained visually, by using a digital camera, a photoelectric photometer, a DSLR, or in some other way — is all done through use of the online tool WebObs (*https://www.aavso.org/webobs*).

You must choose whether you wish to "Submit observations individually" or "Upload a file of observations". If you have just a small number of observations then the individual option may be the easiest for you. If on the other hand, you are submitting a large number of photometry observations (either time-series or for many different stars), creating a file in "AAVSO Extended File Format" is definitely the better way to go. Fortunately, many of the photometry software packages in use today come with an option to export your results in the form of an AAVSO report — you simply need to upload it through WebObs. Should you have to create or tweak your own report, however, it is essential that you follow the format outlined in this Appendix. Even if you submit individual observations, you may find some of the field descriptions in the "Data" section helpful.

C.1 General information

The "Extended Format" file must be a plain text (ASCII) type file. It is not case sensitive. There are two parts to the file; *Parameters* (or header information) and *Data*. There are restrictions on the name of the file: the only acceptable filename extensions are .txt, .csv, and .tsv. A full description of the Extended Format is at *https://www.aavso.org/aavso-extended-file-format*.

C.2 Parameters

The Parameters are specified at the top of the file and are used to describe the data that follows. Parameters must begin with a pound/hash sign (#) at the start of the line. There are six specific parameters that the AAVSO requires you to include at the top of the file. Personal comments may also be added as long as they follow a pound/hash sign (#). These comments will be ignored by the software and will not be loaded into the database. However, they will be retained when the complete file is stored in the AAVSO permanent archives.

The six parameters that are required are:

#TYPE=Extended
#OBSCODE=
#SOFTWARE=
#DELIM=
#DATE=
#OBSTYPE=

Here is an explanation of each:

- **TYPE:** Should always say “Extended” for this format.
- **OBSCODE:** The official AAVSO Observer Code for the observer, which was previously assigned by the AAVSO. To request an observer code, log in to https://aavso.org and go to your profile tab under **`"My Account."`** There you will find a **`Request Observer Code`** button.
- **SOFTWARE:** Name and version of software used to create the format. If it is private software, put some type of description here. For example: “#SOFTWARE=AIP4Win Version 2.2”. This is limited to 30 characters.
- **DELIM:** The delimiter used to separate fields in the report. Any ASCII character or UNI- CODE number that corresponds to ASCII code 32-126 is acceptable as long as it is not used in any field. Suggested delimiters are: comma (,) semicolon (;), exclamation point (!), and pipe (|). The only character that cannot be used are the pound/hash sign (#) and the space. If you want to use a tab, use the word “tab” instead of an actual tab character. Note: Excel users who want to use a comma will have to type the word “comma” here instead of a “,”. Otherwise, Excel will export the field incorrectly.
- **DATE:** The format of the date used in the report. Times specify the midpoint of the observation. Convert all times from UT to one of the following formats:
    - JD: Julian Day (Ex: 2454101.7563)
    - HJD: Heliocentric Julian Day
    - EXCEL: the format created by Excel’s NOW() function (e.g. 12/31/2007 12:59:59 a.m.)
- **OBSTYPE:** The type of observation in the data file. It can be CCD (for either CCD or CMOS cameras), DSLR (for either DSLR or one-shot-color), or PEP (for photoelectric photometry). If no obstype is specified, CCD is assumed. If you use a CMOS camera, please report it as CCD.

The OBSCODE and DATE parameters may also be included elsewhere in the data. Our data processing software will read these parameters and will expect all following data to adhere to them. (For example, you can add “#OBSCODE=TEST” to the report and all subsequent observations will be attributed to observer TEST.)

If you want to put a blank line between your parameter records and your data records, be sure to comment the line out with the pound/hash sign (#). WebObs will not accept a file with blank lines that are not commented out.

C.3 Data

After the parameters, come the actual variable star observations. There should be one observation per line and the fields should be separated by the same character that is defined in the DELIM parameter field. If you do not have data for one of the optional fields, you must use "na" (not applicable) as a place holder. The list of fields are:

- **STARID:** The star's identifier. It can be the AAVSO Designation, the AAVSO Name, or the AAVSO Unique Identifier (AUID), but ***not*** more than one of these. *(25 character limit)*
- **DATE:** The date and time of the observation, in the format specified in the DATE parameter. The AAVSO requires that you report the midpoint of the exposure time. If you stack images, this becomes more complicated so please add a note about how you have computed the exposure time in the NOTES field.
- **MAGNITUDE:** The magnitude of the observation. Prepend with < if a fainter-than. A decimal point is required (e.g. "9.0" rather than "9").
- **MAGERR:** Photometric uncertainty associated with the variable star magnitude. If not available put "na".
- **FILTER:** The filter used for the observation. This can be one of the following letters (in bold):
    - **U:** Johnson U
    - **B:** Johnson B
    - **V:** Johnson V
    - **R**: Cousins R (or Rc)
    - **I:** Cousins I (or Ic)
    - **J:** NIR 1.2 micron
    - **H:** NIR 1.6 micron
    - **K:** NIR 2.2 micron
    - **TG:** Green Filter (or Tri-color green). This is commonly known as the "green channel" in a DSLR or color CCD camera. These observations use V-band comp star magnitudes.
    - **TB:** Blue Filter (or Tri-color blue). This is commonly known as the "blue channel" in a DSLR or color CCD camera. These observations use B-band comp star magnitudes.
    - **TR**: Red Filter (or Tri-color red). This is commonly known as the "red channel" in a DSLR or color CCD camera. These observations use R-band comp star magnitudes.
    - **CV:** Clear (unfiltered) using V-band comp star magnitudes (this is more common than CR)
    - **CR:** Clear (unfiltered) using R-band comp star magnitudes
    - **SZ:** Sloan z
    - **SU:** Sloan u
    - **SG:** Sloan g
    - **SR:** Sloan r

- **SI:** Sloan i
- **STU:** Stromgren u
- **STV:** Stromgren v
- **STB:** Stromgren b
- **STY:** Stromgren y
- **STHBW:** Stromgren Hbw
- **STHBN:** Stromgren Hbn
- **MA:** Optec Wing A
- **MB:** Optec Wing B
- **MI:** Optec Wing C

Please note: There are a few other (rarely used but legitimate) filters, which can be specified. If you are using a filter that is not listed here, please contact AAVSO HQ with as much information as possible about what you are using and we will let you know how to report it.

- **TRANS:** YES if transformed using the Landolt Standards or those fields that contain secondary standards as discussed in Appendix C, or NO if not.
- **MTYPE:** Magnitude type. STD if standardized by utilizing the published magnitudes of the comparison stars or DIF if differential (uncommon). Differential means that the published magnitudes of the comparison stars were not used and only instrumental magnitudes are being reported. DIF requires the use of CNAME. ***Please note that use of the word "differential" in this case is not the same as saying you are doing "differential photometry".***
- **CNAME:** Comparison star name or label such as the chart label or the AUID for the comparison star used. If not present, use "na". *(20 character limit)*
- **CMAG:** Instrumental magnitude of the comparison star. If not present, use "na".
- **KNAME:** Check star name or label such as the chart label or AUID for the check star. If not present, use "na". *(20 character limit)*
- **KMAG:** Instrumental magnitude of the check star. If not present, use "na".
- **AIRMASS:** Airmass of observation. If not present, use "na".
- **GROUP:** Grouping identifier (maximum 5 characters). It is used for grouping multiple observations together — usually an observation set that was taken through multiple filters. It makes it easier to retrieve all magnitudes from a given set in the database in case the researcher wanted to form color indices such as (B–V) with them. If you are just doing time series, or using the same filter for multiple stars, etc., set GROUP to "na." For cases where you want to group observations, GROUP should be an integer, identical for all observations in a group, and unique for a given observer for a given star on a given Julian Date.
- **CHART:** Please use the sequence ID you will find in red at the bottom of the photometry table. If a non-AAVSO sequence was used, please describe it as clearly as possible. *(20 character limit).*
- **NOTES:** Comments or notes about the observation. If no comments, use "na". This field has a maximum length of several thousand characters, so you can be as descriptive as necessary. The convention for including a lot of information as concisely as possible is to use subfields after

any freeform comment you wish to make. The subfield format is **`|A=B`**; the **`|`** character is the separator, **`A`** is a keyword name and **`B`** is its value. To make it possible to easily parse this information, use keywords taken from this list:

- VMAGINS, CMAGINS, KMAGINS are the instrumental magnitudes of target, single comp, and check star
- CREFMAG and KREFMAG are the reference magnitudes of comp and check
- CREFERR and KREFERR are the errors of the reference magnitudes
- VX, CX and KX are the airmass values for target, comp and check
- Transform coefficients can also be documented here. See the example below.

  Not all the values are necessary. But using this mechanism you can document your submission in much better detail. Here is an example of a notes field created by TransformApplier:

  `5 records aggregated|VMAGINS=-7.244|VERR=0.006|CREFMAG=13.793|CREFERR=0.026|KREFMAG=14.448|KREFERR=0.021|VX=1.1501|CX=1.1505|KX=1.1500|Tv_bv=0.0090|Tv_bvErr=0.0100|TAver=2.47`

C.4 Examples

Here is a simple report with multiple stars (the data used are not necessarily realistic!):

```
#TYPE=EXTENDED
#OBSCODE=TST01
#SOFTWARE=MAXIM DL 6.0
#DELIM=,
#DATE=JD
#OBSTYPE=CCD
#NAME,DATE,MAG,MERR,FILT,TRANS,MTYPE,CNAME,CMAG,KNAME,KMAG,AMASS,GROUP,CHART,
NOTES
SS CYG,2450702.1234,8.235,0.003,V,NO,STD,105,10.593,110,11.090,1.561,na,13577KCZ,outburst
V1668 CYG,2450702.1254,18.135,0.0180,V,NO,STD,105,10.594,110,10.994,1.563,na,3577KCZ,na
WY CYG,2450702.1274,14.258,0.004,V,NO,STD,105,10.594,110,10.896,1.564,na,13577KCZ,na
SS CYG,2450722.1294,10.935,0.006,V,NO,STD,105,10.592,110,10.793,1.567,na,13577KCZ,na
```

Note the existence of the #NAME, DATE... line in the above format. Since it is prepended with a #, it will be ignored by our software. Feel free to do this if it makes writing and reading the format easier for you.

Reporting ensemble photometry is permitted under this format. You need to pick one star (the check star) in addition to the target to be measured by the technique. The check star should not be included in the comparison star ensemble. This star's calculated magnitude should be put in the KMAG field, so that if the true magnitude of the check star is found to be different at a later date, a simple zeropoint offset can be added to your ensemble value. If ensemble is used, CNAME should be set to ENSEMBLE and CMAG should be set to "na", as shown below.

```
#TYPE=EXTENDED
#OBSCODE=TST01
#SOFTWARE=IRAF 12.4
#DELIM=,
#DATE=JD
#NAME,DATE,MAG,MERR,FILT,TRANS,MTYPE,CNAME,CMAG,KNAME,KMAG,AMASS,GROUP,CHART,
NOTES
SS CYG,2450702.1234,11.235,0.003,B,NO,STD,ENSEMBLE,na,105,10.593,1.561,1,070613,na
SS CYG,2450702.1254,11.135,0.003,V,NO,STD,ENSEMBLE,na,105,10.492,1.563,1,070613,na
SS CYG,2450702.1274,11.035,0.003,R,NO,STD,ENSEMBLE,na,105,10.398,1.564,1,070613,na
SS CYG,2450702.1294,10.935,0.003,I,NO,STD,ENSEMBLE,na,105,10.295,1.567,1,070613,na
SS CYG,2450702.2234,11.244,0.003,B,NO,STD,ENSEMBLE,na,105,10.590,1.661,2,070613,na
SS CYG,2450702.2254,11.166,0.003,V,NO,STD,ENSEMBLE,na,105,10.497,1.663,2,070613,na
SS CYG,2450702.2274,11.030,0.003,R,NO,STD,ENSEMBLE,na,105,10.402,1.664,2,070613,na
SS CYG,2450702.2294,10.927,0.003,I,NO,STD,ENSEMBLE,na,105,10.292,1.667,2,070613,na
```

In this report, the ensemble solution gave 11.235, 11.135, 11.035 and 10.935 for the B, V, Rc, and Ic (respectively) magnitudes of SS Cyg for the first group, and 11.244, 11.116, 11.030 and 10.927 for the second group. The ensemble solution also gave 10.593, 10.492, 10.398, and 10.295 for the BVRcIc magnitudes of the check star for the first group.

C.5 After Submission

Once you have submitted your observations to the AAVSO database, it is a good idea to take a look at the light curves of the stars you have observed using the Light Curve Generator (https://www.aavso.org/LCGv2) or VStar (https://www.aavso.org/vstar) and see if you think that your data makes sense. If you find that your observations seem to be very different from those of other observers using similar equipment, it is important that you go back and check things against your observing notes or original images. Your observations may be correct while those of another observer or observers could have problems, but if you see a discrepancy, you should start by checking your own data again.

It is not uncommon for observers to make typographical errors resulting in the mislabeling of a star, reporting the wrong date or time, and mixing up the reported bands. If your report seems correct, go back and review your images. Could you have misidentified any of the stars, included a close companion in the aperture or saturated the target or any of the comparison stars?

If you do find a problem, you have the power to fix it. One of the other options available to you through WebObs is “Search for observations”. Using this search tool you should be able to narrow your search so that you can isolate the observation or observations with problems. Then you can either delete the observations and resubmit the corrected ones or edit the erroneous observation. Which option you choose depends on how many observations you have and the nature of the error.

One thing to note about the WebObs Search tool use is that by clicking the little unlabeled box in the left corner of the header row of the “Results” page, you can select all of the observations on that page which makes it much easier to delete a large group of observations rather than clicking on them one-by-one.

If you discover a problem with your data that would be very time consuming to correct, please do not hesitate to contact AAVSO Headquarters to ask for help. Alternatively, if you see something suspicious about another observer’s observations, you can report that "discrepant" magnitude to AAVSO HQ either through use of VStar, Zapper or an email describing what you see.

Appendix D: Supporting Calculations for Color Transformations

Recall from Chapter 6 that there are two types of transformation coefficients: the color transformation coefficients and the filter magnitude transformation coefficients.

D.1 Definition of Color Transformation Coefficients

The color transformation coefficients describe the difference between the measured color indices of stars and the standard color indices of those stars. These coefficients are named Txy, where x and y denote the filter colors used for the index. The value of the coefficients is commonly derived from a plot of measured color index vs standard color index as the reciprocal of the slope of the best-fit line. So, for example, in the graph shown in Figure D-1, the slope of the red line is 0.816 and the value of Tbv is $\frac{1}{0.816}=1.225$.

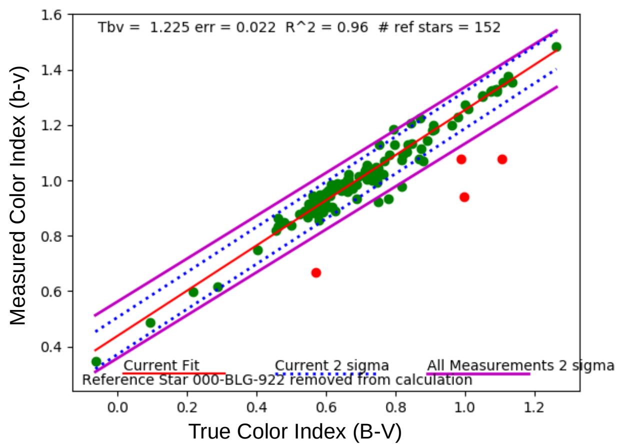

Figure D-1 – typical color transformation graph

From the definition of slope as the change in y values divided by the change in x values for two points, this establishes an expectation that for any two stars observed with this system:

$$Tbv=1.225=\frac{(B-V)_{star\,1}-(B-V)_{star\,2}}{(b-v)_{star\,1}-(b-v)_{star\,2}} ,$$

where b and v are observed (instrumental) magnitudes for the respective stars in Johnson blue and visual filters, and B and V are standard (known) magnitudes in Johnson blue and visual bands for the respective stars.

D.2 Definition of Filter Magnitude Transformation Coefficients

On the other hand, the filter magnitude transformation coefficients describe how the error between the standard magnitude and the measured magnitude changes as a function of standard color. These coefficients are named Tx_yz, where *x* denotes the color associated with the error and *yz* denotes the color index used as the independent variable. So for example, in the graph shown in

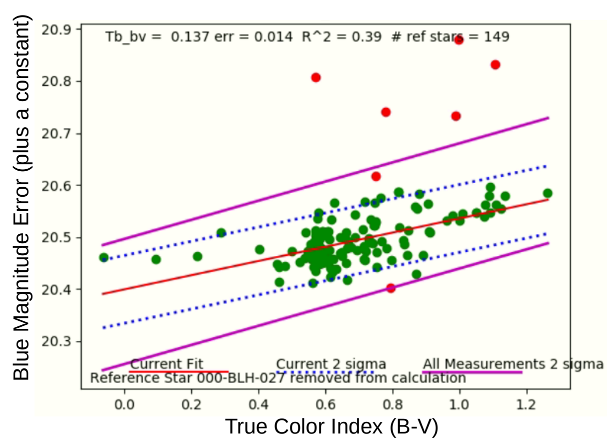

Figure D-2 – Typical magnitude transformation graph

Figure D-2, the slope of the red line is 0.137, which is the value of Tb_bv. This establishes an expectation that for any two stars observed with this system:

$$\mathrm{Tb\_bv} = 0.137 = \frac{(B-b)_{star\,1} - (B-b)_{star\,2}}{(B-V)_{star\,1} - (B-V)_{star\,2}}$$, again using the definition of the slope of a line.

D.3 Available Coefficients

Depending on the filters available to you, there may be a surprising number of coefficients defined for your system. Table D-1 provides a few examples. Note that some magnitude transformation coefficients that could physically be defined are not listed, because they have little or no astronomical value.

Table D-1 — List of Coefficients for Typical Filter Combinations

Available Filters	B, V	B, V, I	B, V, R, I
Color Transformation Coefficients	Tbv	Tbv Tvi Tbi	Tbv Tri Tbr Tvr Tvi Tbi
Magnitude Transformation Coefficients	Tb_bv Tv_bv	Tb_bv Tv_bv Tb_bi Tv_vi Ti_bi Ti_vi	Tb_bv Tb_br Tb_bi Tv_bv Tv_vr Tr_vr Tr_ri Ti_ri Tv_vi Ti_vi Tr_vi

D.4 Calculation of the Coefficients

Once a set of calibrated images has been obtained of a standard field using all available filters, the coefficients are calculated from linear regression best-fit lines through the corresponding plots. The color transformation coefficients are equal to the reciprocal of the corresponding slope, while the magnitude transformation coefficients are equal to the slope itself.

D.5 Measuring Your Transformation Coefficients

D.5.1 Step 1 — Image a standard field and calibrate the images

The first step in determining your transformation coefficients is to image a "Standard Field" using each of your filters. Standard fields are star fields for which the magnitudes of selected stars are known very accurately in several colors. For your convenience, the AAVSO has prepared standard sequences for six star clusters that were selected on the basis of several factors including their range of colors and quantity of stars that conveniently fit into one image (see Table D-2).

Table D-2 — Standard clusters

Name	RA	Dec	Mag Range	Diameter (arc min)	Star Count
NGC 1252	03:10:49	-57:46:00	8 - 15	300+	36
M67	08:51:18	+11:48:00	7 - 16	74	211
NGC 3532	11:05:39	-58:45:12	8 - 13.5	30	288
Coma Star Cluster	12:22:30	+25:51:00	5 - 10	450	92
M11	18:51:05	-06:16:12	8.5 - 17	20	416
NGC 7790	23:58:23	+61:12:25	10 - 20	7	218

You can obtain charts for these fields with the AAVSO Variable Star Plotter (VSP) by typing in the RA and Dec of the field you want and selecting the FOV and limiting magnitude appropriate to your system. Be sure to select "Yes" to the question, "Would you like a standard field chart?" You should also print out the associated Photometry Table containing the published magnitudes of all the Standard stars, which will be needed if your software does not load the comparison star photometry for you.

Now use normal good photometry practices to image the standard field. Try to image the field when it is high in the sky, and use an exposure time that provides as many ADU as reasonable without saturating the brighter stars. Take several images with each filter and calibrate them with bias, dark, and flat frames. Stack the images to increase SNR.

To minimize the effects of spurious problems or atmospheric effects, it is a good idea to repeat the entire process of imaging the standard field and computing your coefficients over several nights. Your results from each of the nights can then be averaged together to get a better set of coefficients.

D.5.2 Step 2 — Measure the images to obtain magnitudes

Using your photometry software, measure as many stars as you can to obtain their magnitudes. There is no need to select a specific target star or check star. As with any crowded field, be careful not to measure any stars that are so close together that their images blend or overlap. Also be very careful with star identification, and in the case of multiple stars with the same identifier in VPhot or on the VSP chart, double check the RA and Dec to be sure you know which is which.

D.5.3 Step 3 — Compute the transformation coefficients

AAVSO volunteers developed software tools to assist you with both computing transformation coefficients (e.g., TransformGenerator – TG) and applying coefficients to transform your data (e.g., TransformApplier in VPhot) (see https:// www.aavso.org/transform to download the programs and read useful information about them).

Here is the process you would use to generate transformation coefficients from within VPhot:

1. Collect your images of M67 in all your filters and calibrate them with bias, darks, and flats.
2. Upload the images to VPhot.
3. Select all the images (i.e., all filters), and select **`Time Series`**. (This now works with multiple filters.)
4. For sequence you should select **`com M67 AUID`**. This sequence labels the stars with AUIDs and designates one star as a target star (in this case, HW Cnc). This isn't used in the traditional sense of a target star; it simply keeps the internal Time Series process happy.
5. Click on **`Start the Analysis`**.
6. Once done, select **`General Export`**.
7. At the bottom of the page, you will find a button **`Save to AIP fmt for TG`**. Clicking this button will pack all your M67 photometry into one file in a format that TG will accept.
8. Open the TG tool. Select your scope, **`M67`**, and **`AIP/Maxim format`**.
9. The TG tool uses the output of VPhot to automatically create your transformation graphs. You review each of the graphs and delete outliers (shown in red in Figure 6-1). TG will also permit you to average this data set with other data sets that you've generated.
10. When finished, output the transformation coefficients into a INI file that you can then provide to VPhot and import into your telescope profile.

The final result is a set of transformation coefficients (both the color transformation coefficients and the filter magnitude transformation coefficients) that you will apply every night as an integral part of your analysis process.

Transformation coefficients should be computed at least once a year, but if you change anything in your optical train (replacing a filter, adding a field flattener, etc) you will have to compute your coefficients again.

D.6 Application of the Coefficients

Applying the transformation coefficients to your observations is more complex than the computation of the coefficients. To understand why, consider the simplest case: an observation using only blue and visual filters. We then have three coefficients (Tbv, Tb_bv, and Tv_bv). If we observe two stars, designating one as the target and one as the comparison, we then have three equations available (in these equations, values shown in green are known and values shown in red are unknown).

$$Tbv = \frac{(B-V)_{comp} - (B-V)_{target}}{(b-v)_{comp} - (b-v)_{target}} \qquad (1)$$

$$\text{Tb\_bv} = \frac{(B-b)_{comp} - (B-b)_{target}}{(B-V)_{comp} - (B-V)_{target}} \qquad (2)$$

$$\text{Tv\_bv} = \frac{(V-v)_{comp} - (V-v)_{target}}{(B-V)_{comp} - (B-V)_{target}} \qquad (3)$$

Thus, we have three equations in two unknowns. This is an overconstrained system, and given the way the coefficients were derived empirically from data with inherent measurement error, it is likely that the three equations have no self-consistent, single solution. However, it is fairly straightforward to take either equations (1) and (2) and solve for the unknowns, or to take (1) and (3) and solve for the unknowns. Just for illustration, we take (1) and (3). Rearranging (1), we get:

$$(B-V)_{target} = (Tbv)((b-v)_{target} - (b-v)_{comp}) + (B-V)_{comp} \qquad (1a)$$

Substituting that into (2) gives:

$$\text{Tb\_bv} = \frac{(B-b)_{comp} - (B-b)_{target}}{(Tbv)((b-v)_{comp} - (b-v)_{target})} \qquad (4)$$

This can clearly be solved for *B*. The value of *B* can then be substituted into (1a) to calculate a value of *V*.

Yes, the algebra is somewhat messy, but is straightforward. Similarly, (1) and (3) can be used to calculate *B* and *V*. The final values will be slightly different than with (1) and (2), but should be close.

As the number of filters increases, the algebra becomes more messy and the system of equations becomes even more overconstrained. As was mentioned earlier, some color indices have more physical significance than others (e.g., B-V has proved a better measure of overall star color than B-I). This preference for certain color indices has established a preferred set of transformation coefficients for each filter combination.

Available Filters	**To solve for**				
	U	**B**	**V**	**R**	**I**
UBVRI	Tu_ub	Tb_bv	Tv_bv	Tr_vi	Ti_vi
UBVI	Tu_ub	Tb_bv	Tv_bv		Ti_vi
BVRI		Tb_bv	Tv_bv	Tr_vi	Ti_vi
BVR		Tb_bv	Tv_bv	Tvr	
BVI		Tb_bv	Tv_bv		Ti_vi
VRI			Tv_vi	Tr_vi	Tvi
BV		Tb_bv	Tv_bv		
VR			Tv_vr	Tr_vr	
VI			Tv_vi		Tvi

Applying color transformations to a comparison star ensemble is potentially complex, depending on how the ensemble is combined (remember our recommendation for beginners to avoid ensembles). VPhot has two color transformation tools: the TransformApplier and the Two-Color-Transform tool. Only the Two-Color-Transform tool can be used with ensembles, and it forces a different mix of equations than is recommended in the earlier table on this page. LesvePhotometry, on the other hand, *will* apply color transformations for all three of its ensemble modes. Review the documentation for the software you consider using, and be careful to keep your transformation workflow aligned with the recommendations of this *Guide*.

Clearly, you can only use transformation coefficients that use data from filters that you used for your observation. If you were unable to obtain images with the Johnson B filter (perhaps the star was red and faint and was not visible in your B images), then you can't use Tv_bv to adjust your V images and will want to shift to Tv_vi.

There's another, more subtle issue associated with the quantity of images you obtain in what is called an observation "group." The reporting format provides a way to identify the observations in a group so that subsequent researchers can understand how you derived your data. The measurement sets that go into the TransformApplier should be balanced across colors: three sets of blue photometry, three sets of visual, and three sets of red. These will be matched to each other during the transformation process. If you were to add a fourth red photometry set, it would not be transformed since there are no visual and blue data to go with it, and using a single data set more than once seriously taints your final result.

Appendix E: Self-Assessment Report Card

Photometry Self-Assessment Report Card		Date: ______
	Grade (A-F)	Notes
Image Quality		
Focus & Plate scale/sampling		
Tracking		
Exposure/Saturation, adequate SNR		
Skyglow/Background		
Field Uniformity (e.g., coma, vignetting)		
Image defects and artifacts		
Calibration Procedure (Bias, Dark, Flat)		
Color		
Filters		
Color Transformations		
Photometric Reduction		
Aperture selection/placement		
Comparison stars		
Check stars		
Reporting		
Assessment & Improvement		
Use of standard fields		
Comparison with other observers		
My current hardware/software/procedural baseline includes:		
My next improvement initiative will be:		

This template can be used as the basis for periodic self-assessment. Topics should be tailored out of the template depending on the needs of your observing program. Topics that should be included in each assessment category include:

Focus and Plate Scale/Sampling:

- Do my images consistently have a FWHM between 2 and 3 pixels?
- If I make multi-hour time sequences, how much does star size vary over time? (Is focus drifting?)

Tracking:

- Are my stars usually round? How much does FWHM measured vertically differ from FWHM measured horizontally?
- Do I ever get "twin stars"? (Multiple, side-by-side images of every star in an image?)
- Does poor tracking limit the SNR I can achieve on target or comparison stars by limiting my exposure times?

Exposure/Saturation, Adequate SNR:

- Do I have a linearity curve for my camera that's recent? Do I use it to exclude stars that approach (or exceed) my camera's linearity limit?
- If binning, how am I detecting saturation in binned pixels?
- Is my star SNR consistently aligned with the accuracy needed for my particular observing program? (For example, if I need 0.02 magnitude precision, is my SNR *at least* 50?)
- Are the answers to the first three questions consistent across all the filters that I use? Do I have a filter where I struggle to get sufficient SNR?

Skyglow/Background:

- Is background sky brightness flat (even) across my images? If it isn't flat, how am I deciding whether to exclude that image from my photometry?
- Does skyglow limit my photometric precision by limiting my SNR? Is there anything I can do to reduce skyglow?
- How does skyglow change from one filter to another?

Image Defects and Artifacts:

- Do I know what a cosmic ray looks like? Do I scan every image for cosmic ray hits before believing the photometry I extract from that image?

- Same for satellite tracks.
- Have I tested for residual bulk images (RBI)? If this is known to affect my camera, do I have a plan in place to minimize the effect of RBI on photometry?

Field Uniformity:

- Can I detect coma or other radial aberrations in my images? This can have a profound effect on photometric accuracy. If so, what am I doing about it?
- How confident am I that field illumination (after application of flat fields) is constant across the field of the camera's sensor? Have I measured residual flatness?

Calibration Procedure:

- Am I making a sufficient number of bias, dark, and flat field exposures?
- Am I combining bias, dark, and flat field exposures in a way that protects against defects (e.g., cosmic ray hits) in a single image from influencing the combined master calibration image?
- Do I have configuration control over my calibration frames? Do I know where each one came from? Do I know how old each one is?
- Do I refresh calibration masters often enough? How do I know?
- How do I test my master flat images? How much confidence do I have in my master flats?

Filters:

- Am I using filters that adhere to a standard photometric color system?
- Am I using a different master flat for each filter?
- Do I monitor the filters for slow degradation over time?
- How accurately are filters positioned in the optical path? (Do filter-specific dust motes fall on the same pixels even after the filter has been repositioned?)

Color Transformations:

- Do I know my color transformations?
- How old are my color transformations?
- Am I submitting transformed measurements to the AAVSO?

Aperture selection/placement:

- What aperture and sky annulus radii am I using? Do I know why I'm using the values that I'm using?
- How/when do I check for contaminating faint stars in my target (or comparison star's) annulus?

Comparison Stars:

- How am I choosing comparison stars?
- How well do comparison star colors align with target stars' colors? How well do comparison star colors align with each other in any given field?
- How do I ensure that I pick up changes to the AAVSO charts or sequences?
- How well do comparison star brightness match my target stars' brightness?

Check Stars:

- Am I using a check star? How am I using the check star?
- If using a single comparison star, the other sequence stars are all candidate check stars. In addition to the reported check star, the other sequence stars can be used as additional, unreported check stars. Am I taking advantage of this?
- Am I getting check star photometry from the AAVSO sequences?

Reporting:

- Do I report data with a cadence suitable for the science I'm supporting?
- Am I taking advantage of all the reporting fields available to me in the AAVSO reporting format?
- Do my reported observation errors make sense? Do I know where they come from? How have I validated them?
- (For precise-time observing programs...) Do my time tags reflect the middle of the exposure time? Is my time source synchronized to a UTC standard? Are my time tags sufficiently precise and accurate for the science I'm supporting? If minutes or seconds matter, am I using heliocentric time?

Use of Standard Fields:

- Am I using standard fields as a part of my observing program?
- Am I using standard fields for both finding color transformation coefficients and for measuring my photometric accuracy?
- How often do I capture and analyze standard field images?
- How am I using my standard field observations to influence my long-term improvement activity?

Comparison with Other Observers:

- How am I checking every observation for consistency (and obvious errors) prior to submission?
- How closely do my measurements align with other observers? When I see inconsistency, what do I do about it?

Appendix F: Getting Started, Step By Step

In this appendix, we walk through the first steps associated with getting out at night with your system and starting to perform photometry. (Depending on your starting point and your goals, your journey through this process may differ somewhat from this *Guide*.) Section numbers have been included in square brackets to provide references to more in-depth discussions elsewhere within this *Guide*.

Although this appendix is intended to be useful for any photometry system, it uses a specific camera and telescope in its examples: a QHY183M monochrome CMOS camera with two 1.25" photometric filters (Johnson B and V) paired to an 8" (20cm) f/10 Schmidt-Cassegrain telescope on an equatorial mount. You need to substitute the characteristics of your own system.

As an important note, this particular camera uses a 12-bit analog-to-digital converter. Per [3.3.4] each pixel can have an ADU value of up to 4,095. However, the camera designers chose to multiply each pixel value by 16 to give an ADU range of up to 65,535. This will become relevant when checking for saturation, below.

This appendix is also specific in places about the software used in its examples: AstroImageJ and VPhot. This software pair is somewhat unique by not being tied to any specific operating system; they can be used with Windows, MacOS, and Linux. Again, there are many other choices available to you. We simply chose this pair for our example here.

- **Pick an initial target:** SS Cyg makes a great first target for northern hemisphere observers. This star is generally in one of two states: quiescent (around magnitude 12) and outburst (around magnitude 8.5).
- **Calculate image scale and field of view:** The QHY183M sensor measures 13.3mm by 8.87mm with 2.4 micron pixels. Use the formula from [3.2.1]:

FOV = (57.3 × width/focal length) by (57.3 × height/focal length)

(FOV in degrees, focal length in mm, height & width of the chip in mm)

$$FOV = 57.3\left(\frac{13.3\,mm}{2000\,mm}\right) \text{ by } 57.3\left(\frac{8.87\,mm}{2000\,mm}\right) = 0.38° \text{ by } 0.25° = 22\,arcmin \text{ by } 15\,arcmin$$

The image scale uses the formula from [3.2.2]:

Image scale = (CCD pixel size/focal length) × 206.265

(image scale in arcsec/pixel, CCD pixel size in microns, focal length in millimeters)

$$\text{Image Scale} = \left(\frac{2.4}{2000}\right) 206.265 = 0.25\,\text{arcsec/pixel}$$

With typical seeing of 3 arcsec FWHM, this means that FWHM is 12 pixels (which is quite oversampled). As discussed in 3.2.3.2, we can accept this oversampling or improve it by either adding a focal reducer or binning our images. In this example, we choose to accept the oversampling.

- **Select an observing cadence:** For a star with this typical lightcurve, a cadence of one data point (report) per day is appropriate (one report per color: a data point in B and a data point in V).
- **Fetch a chart and identify comparison and check stars:** Using the Variable Star Plotter on the AAVSO website, fetch a chart and overlay the field of view, as seen in Figure F-1. Looking for candidate comparison stars, many seem suitable (98, 109, 86, 119, 114, 123, etc), but we tentatively choose the 98 comparison star for those times when SS Cyg is in outburst and choose the 128 comparison star for those times when SS Cyg is quiescent. We choose the 98 comp based on it being relatively bright (to give good SNR) and having few nearby contaminating companions. We reject the 86 sequence star because of the potential of close companions and because its B-V color index is 1.3, making it somewhat redder than the others [2.2]. We choose the 109 star as a check star; we only choose a single check star so that we can look for shifts in the plot of the check star as we shift back and forth between the outburst comp star and the quiescent comp star. The target star, both comparison stars, and the check star are all well inside the field of view.

Planning Checklist:
- Star high enough in the sky?
- Star far enough from the moon?
- Comp star a reasonable color match to the variable?
- Same comp and check as last time?
- Exposure long enough to give reasonable counts on var, comp, and check?
- Exposure short enough to avoid saturation?
- New bias, dark, or flat images needed

- **Make test exposures and create an exposure plan:** An exposure strategy consisting of exposure time, gain, offset, readout mode, and number of exposures is selected. Based on manufacturer curves, a gain setting of 0 provides the highest dynamic range. This camera does not offer a choice of readout modes. An offset setting of 100 is selected based on examination of the background histogram for a test exposure at a gain setting of 0. The test exposure of 60 seconds was about the longest the telescope mount would permit without guiding (no guidescope is currently present), and the test exposure showed at least a factor of 2 margin to saturation-related nonlinearity, so the planned exposure time will be 60 seconds. Images need not be stacked, since the test image suggested that adequate SNR would be available for all stars of interest in the frame. The number of exposures is somewhat arbitrarily set at 10 per filter (a minimum of 3 is needed) to provide a good indication of measurement precision.

- **On a cloudy night, create master calibration files:** We need a raw Master Dark for 30 seconds, flats for B and V, and a Master Bias (to dark-correct the flats). A total of 20 dark frame exposures are made by installing the telescope tube end-cover and making 20 exposures for 30 seconds each. The Master Bias is made from 50 bias exposures (end-cover in place, zero exposure time). A light panel and T-shirt are used for the B and V flats: 20 exposures each at around 2.5 seconds for the V and 4 seconds for the B. Those three image sequences (bias, dark, and flat) are provided to AstroImageJ, which creates the raw Master Dark, Master Bias, and two Master flats.
- **Set up for an observing session:** On the observing night, the telescope is set up, and a polar alignment is completed (if needed). The camera is turned on and the camera cooler power is slowly raised until the camera has cooled to a temperature requiring about 80% power (steady state). The camera is focused.

Observing Night Checklist:

- Sky free of haze, smoke, thin clouds, aircraft contrails?
- Equipment cooled down?
- Dewing prevented?
- Laptop has adequate charge (or is plugged in)?
- Cables firmly attached and out of the way?

- **Find the field and make a test exposure:** The telescope is pointed and a 30-second test exposure is made. The field is identified, and pointing is adjusted so that the SS Cyg target star is approximately centered in the field. Several checks are made:
    - Verify that the target, comparison, and check stars all appear in the image (as expected based on the FOV overlay on the chart).
    - Check the brightest of the relevant stars (SS Cyg, comparison, and check star) for saturation by plotting the stars' PSF profiles using the software you are using for in-the-field analysis. Note the highest ADU count. (User experience with this camera suggests that saturation-related nonlinearity sets in with an ADU level of about 45,000.)

Ready for Imaging Checklist:

- Are var/comp/check reasonably well-centered in field?
- Is the centering exposure free of obvious defects?
- Camera in temperature lock?
- Exposure length correct?
- Proper binning and gain?
- Sensible file names and directories to save images?

- **Make your science exposures:** Ten 60-second exposures are made through each filter. If any are contaminated by satellite trails or poor tracking or "wind shakes," additional 60-second exposures are added to the sequence.

Science Image Checklist:

- Check the images as they pop up.
- Listen for expected filter wheel movements.
- Keep checking sky for clouds, etc.

- **Shutdown the observing session:** The telescope is shutdown, the power level of the cooler is slowly reduced over a period of 15 minutes, and then the camera is shut down.
- **Calibrate your images with darks and flats:** AstroImageJ is used to subtract the raw Master Dark frame and scale/divide the result using the flat frames. Your 20 science images are calibrated into 20 new, calibrated science images. [4.2]
- **Assess calibrated images:** Each of these calibrated images is now assessed, checking (again) for
  - saturation-related nonlinearity[InfoBox 3-B],
  - passing satellite trails close to the target, comparison, or check stars[4.8.7],
  - tracking/wind problems[page 70],
  - cosmic ray defects within a few pixels of the target, comparison, or check stars[4.8.6].

  Any images with problems are discarded.
- **Load into photometry tool:** The 20 calibrated science images being used for photometry are loaded into VPhot and plate-solved, identifying the stars in the images.
- **Choose photometry apertures:** Use VPhot to measure FWHM of a sample of stars. We find we have an FWHM of about 11.5 pixels. We use 2xFWHM (23 pixels) as the radius of the inner aperture. Use 3x the inner aperture radius (69 pixels) as the inner radius of the background annulus.
- **Perform photometry:** We check the appearance of SS Cyg to determine whether it is in outburst or quiescent. It is clearly fainter than the 98 sequence star, so we proceed for the quiescent case (using the 128 comp star). Within VPhot, we designate the 128 star as the comparison star. Generate a VPhot photometry report for each image and click on "**`Keep this`**." (Note that the VPhot Users Guide assumes that you are using a comparison star ensemble. We aren't; we have a single comparison star. Where the VPhot Users Guide talks about multiple comparison stars being selected,

we have only one.)

- Generate an untransformed report: In VPhot, go to the Analysis Log, select all 20 images, and click on **`Create AAVSO Report`**. If everything looks as expected, click on **`Create Report File`**. This will create an untransformed report on your computer with one report entry for each of your 20 images.
- **Perform color transformations:** In the VPhot Images display, click on **`TransformApplier`**. Use the **`Browse`** button to select the untransformed report that you just downloaded. Click on the **`Aggregate observations`** checkbox (which will combine your 20 images into 2 transformed data points). Click on **`Create Transformed AAVSO Report`**. Then click on **`Download Output`** to load the final, transformed report onto your machine. This report will have a single reported value for B and a separate one for V. The reported uncertainties will be derived from the standard deviations of your measurements (in each set of 10 images), your photometry SNR data, and the uncertainties in your transformation coefficients.
- **Assess against AAVSO LCG:** Look at the AAVSO Light Curve Generator for SS Cyg. Find your date/time of measurement and compare your two values (one per color) against the curve. The uncertainties in your report establish an expectation of how well your data could align with other photometrists' data. If yours doesn't match that well, go back and double-check everything.

> **Analysis Checklist:**
> - Appropriate dark-and-bias and flat field corrections applied?
> - Appropriate aperture and annulus sizes?
> - Transform applied?
> - Do the results make sense?
> - Metadata fully populated in the AAVSO report?

- **Submit measurements to AAVSO:** Submit your two measurements (one per color, reported as B and V)
- **Fill out Appendix E report card:** Do a self-assessment using the criteria described in Appendix E.

Appendix G: Observer Resources

Books

Berry, Richard, and James Burnell. *The Handbook of Astronomical Image Processing* (second edition). Willmann-Bell, Inc. 2005. ISBN 978-0943396828.

Budding, Edwin, and Osman Demircan. *Introduction to Astronomical Photometry* (second edition). Cambridge University Press, 2007. ISBN 978-0521885263.

Buchheim, Robert K. *The Sky is Your Laboratory*, Springer Science+Business Media, 2007. ISBN 978-0387718224.

Hall, Douglas S., and Russell M. Genet. *Photoelectric Photometry of Variable Stars — A Practical Guide for the Smaller Observatory* (second edition). Willman-Bell, Inc., 1988. ISBN 978-0943396194.

Henden, Arne A., and Ronald H. Kaitchuck. *Astronomical Photometry, A Text and Handbook for the Advanced Amateur and Professional Astronomer*. Willman-Bell, Inc., 1990. ISBN 978-0943396255.

Howell, Stephen B., *Handbook of CCD Astronomy* (second edition). Cambridge University Press, 2006. ISBN 978-0521617628.

Warner, Brian D., *A Practical Guide to Lightcurve Photometry and Analysis*. Springer Science+Business Media, 2006. ISBN 978-0387293653.

Appendix H: Glossary

absolute photometry – The measurement of the brightness of a star using measuring equipment that has been previously calibrated using other stars of known brightness. An alternative to differential photometry.

accretors – Stars gaining material (typically, hydrogen) as a result of gravitational attraction; the process of accretion frequently causes the inbound material to settle in a semi-stable disk around the receiving star.

amplitude – The change in brightness of a variable star; typically measured as the difference between (ratio of) the star's brightest and dimmest apparent brightness.

apparent brightness – The brightness of a star as seen by an observer on the Earth. Apparent brightness is reduced from absolute brightness by interstellar absorption and by distance from the Earth.

binary stars (binaries) – Star pairs that formed together and that orbit each other

CCD – Charge Coupled Device: a semiconductor technology that enables electric charge to be moved from one cell (pixel) to an adjacent one; a technology used in digital camera sensors

charge – a measure of the quantity of electrons that have accumulated in a material

check star – A star (or multiple stars) of known brightness used for quality control purposes while conducting *differential photometry*; the check star's brightness is measured as if its brightness was not already known.

CMOS – Complementary Metal–Oxide–Semiconductor: a semiconductor technology for making transistors that is an alternative to CCD in the manufacture of digital camera sensors

coma – an optical aberration that causes stars to become misshapen and stretched; coma typically becomes worse as stars are moved from the center of an image to the edges

comparison star – A star (or multiple stars) of known brightness used as references in *differential photometry*

differential photometry – The measurement of the brightness of a star by comparing its flux to the flux of a star of known (previously-measured) brightness; the two stars are typically imaged simultaneously (requiring them to be close in the sky).

dynamic range – The ratio of the brightest thing faithfully rendered by a camera to the dimmest thing captured. The brightest thing faithfully rendered is usually determined by the onset of saturation nonlinearity. The dimmest thing captured is usually determined by the level of background noise. Thus, a sensor's dynamic range is typically the ratio of full well depth (or onset of ADC saturation, which ever occurs at the lower exposure level) to the read noise, with both numbers measured in electrons.

eruptive variable (eruptives) – Variable stars that exhibit rapid increases in light output as a result of an instability in the star's atmosphere or in a shell (or disk) of material surrounding the star

exoplanet – a planet orbiting a star other than our sun.

filter – an optical element put into a light path to selectively pass specific wavelengths of light and block other specific wavelengths

focal reducer – an optical element added to a telescope to reduce its apparent (effective) focal length

flux – a measure of the rate of arrival of photons from a star

Fork Mount – a telescope mounting system that supports the telescope with a large fork-shaped structure; the telescope moves *inside* the tines of the fork.

gain – the ratio of the number provided by a camera for each of its pixels and the number of electrons that were released due to photoelectric effect; in cameras with an adjustable gain, the *gain setting* is a user-selectable value that determines the gain. The *gain* and the *gain setting* rarely use the same scale.

German Equatorial Mount – a telescope mounting system that provides two telescope motions, one that is in an east-west direction in the sky and one that is north-south

luminosity – an absolute measure of the total amount of power radiated by a star (as distinct from its *apparent brightness*)

magnitude – (Within this document,) a nonlinear measurement scale for reporting the brightness of stars. *Flux* (a linear measure) is converted to *magnitude* (a nonlinear measure), as described in section 5.1

off-axis – The center of an image is commonly called "on-axis" to indicate that it sits at the intersection of the imaging surface and the longitudinal optical axis that connects the centers of all the optical elements of the system; *off-axis* refers to locations that are *not* in the center of the image (e.g., edges, corners)

photoelectric effect – a process in which electrons are released from a material as a result of the absorption of an inbound photon of light

photometrist – someone who measures the light from a star

photometry – the process of measuring light from a star

pulsators – Variable stars that vary in size, shape, or temperature over time

residual error – The error that remains after a measurement has been corrected for all known error sources.

Residual error is usually defined as the difference between a measurement and a community-accepted standard value.

saturation – the reduction in a pixel's sensitivity to light due to the amount of light already received.

standard color system – an astronomical standard that defines measurement color bands by specifying *response curves* that describe how much light of what wavelengths are to be included in a measurement of that color. As an example, astronomers use the Johnson standard color system to define colors named "ultraviolet," "blue," and "visual". Other standard color systems provide alternative definitions of these three colors and add definitions of dozens of additional colors.

star field – a specific region of the sky, including the stars found in that region.

vignetting – a reduction in the brightness of an image as one moves away from the center of the image, typically caused by some physical object that blocks some of the light rays being focused to create the image.