AAVSO Guide to Photometric Uncertainty

Supplementary material for the

Carolyn Hurless Online Institute for
Continuing Education in Astronomy (CHOICE)

Table of Contents

Introduction

Every measurement has an uncertainty associated with it. Sometimes the uncertainty is explicitly stated ("The president's approval rating is 50 +/-1%") and sometimes it is implied ("The temperature is 32 degrees", which implies that the real temperature is between 31.5 and 32.4 degrees). In quantitative fields, such as astronomy, it is important to explicitly report your uncertainty. This course will describe the most common methods for determining uncertainty for photometric observations. However, it contains information that will also be useful for anyone interested in time series or any other type of variable star data analysis.

Uncertainty is the result of a battle between two forces: signal and noise. We adopt this motif throughout the document. The issue always comes down to compromise: how much energy are you willing to put into a system to enhance signal and dampen noise? It is a grey area that you, as a scientist, must contend with. As you learn to identify and quantify uncertainty you'll learn more about your system and procedures. It will have a profound effect on future decisions you make regarding your observing program.

So, What's the Point?
Uncertainty tells someone looking at your data how much they should believe it. Consider the following two observations:

Magnitude	Date	Observer
11.5 +/- 0.3	Oct 1.1	BART
11.2 +/- 0.2	Oct 1.2	LISA

The magnitudes reported in these observations are not different from each other. Lisa may want to say the star got brighter in her observation, but she'd be wrong. Bart's observation is really saying "the star is between 11.2 and 11.8 magnitudes" and hers really says "the star is between 11.0 and 11.4 magnitudes". The star could actually be between 11.2 and 11.4 magnitudes in brightness and not change at all between these observations.

Let's consider a different thought experiment. Imagine an astronomer is planning to observe a cataclysmic variable with the Hubble Space Telescope. But HST cannot observe the star if it is brighter than magnitude 10 because it will damage a sensitive instrument. So the astronomer checks ahead of the observing run to make sure it is still in quiescence. They see these observations previously reported.

Magnitude	Date	Observer
10.5	May 1.5	MARG
9.9	May 1.5	HOMR

One of the reports has the star greater than magnitude 10. So based on that info alone, the astronomer will likely cancel the observation to protect their instrument. Now consider the reports when uncertainty is included:

Magnitude		Date	Observer
10.5	+/- 0.2	May 1.5	MARG
9.9	+/- 0.5	May 1.5	HOMR

Since both MARG and HOMR were responsible observers and reported their uncertainties, the astronomer knows the star is likely still fainter than 10. So they can go ahead with their run. The astronomer knows how much to trust the data.

This is a fictional scenario; in real life no observation would be made that close to the criteria. But the scenario is based in reality. In the late 1990s, an astronomer was preparing to observe a cataclysmic variable with HST. They saw a report of it in outburst via another organization's mailing list and was preparing to cancel the observation. They contacted Janet Mattei at the AAVSO for advice. She contacted AAVSO observers for help and received many reports that the star was not in outburst after all. The astronomer went ahead with their observations. What they did, basically, was wait for more data to arrive so they could determine an uncertainty around the reports and then make a decision.

Error vs. Uncertainty

For many, error and uncertainty mean the same thing. So you may hear people ask: "What is the error of your measurement?" It is just a colloquial way to ask for uncertainty. We prefer to use the term uncertainty because it is more precise. Error refers to the difference between the real and the measured value. If it is your first time to measure something you can have an unknown error since you do not know the real value (some theorists prefer to describe this as *no* error!).

Uncertainty is the amount of variation you observe when you repeat a measurement over and over again. Things that can change between measurements cause variation. Note this is not the same as error. Let's say you have an object that is at exactly 10 degrees in temperature. You measure it five times and get readings of 9.8, 9.8, 10.0, 10.1, and 10.2. The average of that group is 9.98 and its uncertainty is 0.18 (we'll explain why later). So the error is 0.02 but the uncertainty is 0.18.

Again, you'll often see the two terms used interchangeably. But in science, and in this course, we'll be focused mainly on determining uncertainty.

Accuracy vs. Precision

Accuracy is the opposite of error – it refers to the level of agreement between the measurement and the real value. Accuracy entails *everything* that can affect a measurement. Precision roughly describes how well an observation can be repeated. That is, if you repeated it over and over again what pattern of measurements would you get?

Precision, on the other hand, only includes changes that can occur *between* measurements.

Let's look at an example regarding making a visual estimate of a variable star. You have a variable star whose brightness is bracketed by two comparison stars of values 10.0 and 10.5 magnitude. Now let's say you make five observations of the star during the night. It is likely that the five observations you make will differ somewhat due to the influence of fatigue, clouds, increasing familiarity with the field, etc. Your precision will be based on the difference between those observations.

At the same time, the accuracy of your observations that night are affected by other things as well. For example, the quality of the comparison stars have a huge effect on accuracy. If the 10.0 comparison star is really a 9.7 comparison star, then your accuracy will be off. It doesn't matter how great your precision is – you're accuracy will suffer. And it will suffer all night long since you are always using that faulty comparison star.

As you may notice, the relationship between accuracy and precision is similar to (but not exactly the same) as the relationship between error and uncertainty.

For an illustration, see Figure 1. It is a real light curve of BZ UMa during one of its legendary short outbursts. The long, gently sloping line is from two observers – let's call them observer A and observer B. The collection of more disorganized observations in the middle are from a third observer, observer C.

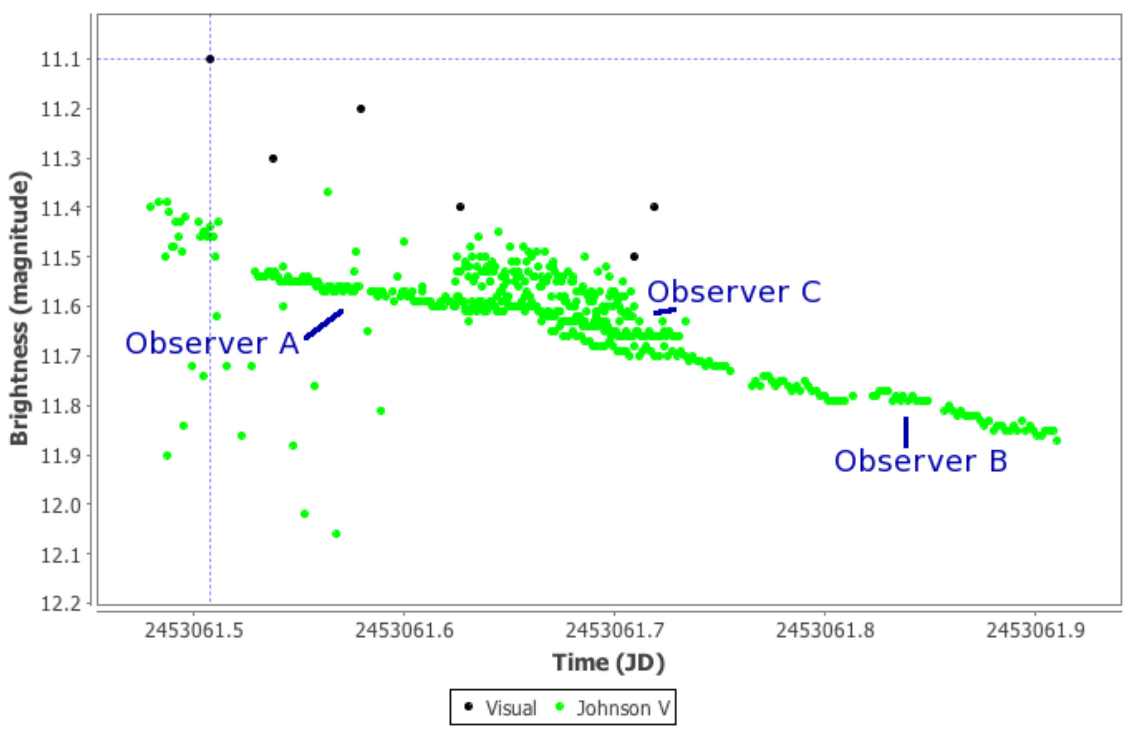

Figure 1: A real light curve of BZ UMa illustrating accuracy and precision differences between observers. (This light curve was generated with the AAVSO's VStar software - available free from aavso.org.)

Observer A collected data for a few hours and then stopped. Observer B began their data collection an hour or so before observer A stopped, so there is some overlap. Notice there is an offset in their reported observations, but the overall slope is similar. This is a difference in *accuracy*. It could be caused by many factors such as using different comparison stars and/or different sequences completely. It could also be caused by not transforming the data or other procedural decisions the observer made. But notice that the difference between the individual observations of Observer A and between the individual observations of Observer B is very similar. This is their *precision*, which is relatively high for both observers in this example.

On the other hand, the precision of Observer C is low compared to the others. There is quite a bit of difference between Observer C's individual observations. One possible cause of this scatter is insufficient exposure time. Observer C also has an overall offset (sometimes referred to as a zeropoint offset) compared to the other two observers. That reflects upon their accuracy. As before, it is possibly due to using different comparison stars than the other observers. However, Observer C's data shows an overall slope very similar to the other two observers. If Observer C binned their data, they would likely increase their precision and have a tighter line with a gentle slope. However, it would still have the offset in accuracy.

This light curve tells us many stories about the observers' night. But what does it tell a researcher? In terms of precision, the researcher will likely average the data from each observer and take the standard deviation to establish their uncertainty. In terms of accuracy, you can't tell that from the light curve alone. Which of the three observers is more accurate? The researcher would look at the data reports submitted by the observer to decide this. One way would be to look at the charts and comparison stars each observer used, and determine which chart/sequence the researcher trusts the most. This is why it is important to report that information to the AAVSO along with your data.

Random vs. Systematic Uncertainty

Random uncertainty (sometimes referred to as stochastic or statistical uncertainty) is the amount of randomness in your measurement. It is always present and cannot be completely eliminated. In variable star astronomy, it is usually dominated by random uncertainty in the amount of light coming into the detector. This random uncertainty follows a Poisson distribution, which is the square root of the number of measured photons associated with the star. Thus, a measurement of 10,000 photons has a random uncertainty of 100 (square root of 10,000) or around 1%. [1]

Systematic uncertainty is related to your specific setup and procedures. This is what most people think of when they think of uncertainty. It includes everything from equipment performance to mistakes in following procedures. For beginners, this is usually the most

[1] Photons also enter the detector from sources other than the star (light leaks, etc.). They will be discussed later on in the section about determining photometric uncertainty.

significant type of uncertainty. As time goes on and observers gain experience, random uncertainty should become more important.

Standard Deviation and Standard Deviation of the Mean

Standard deviation is a type of statistical uncertainty, meaning it is created through statistics – not measurement. Both describe the amount of variance in a system. That is, they let you know how much scatter is in a group of data.

Standard deviation is an estimate of how likely a data point in a group is to deviate from the mean of that group. It is the typical way of reporting uncertainty for means. (For example, error margin in public polling tends to be computed via standard deviation.) It is often reported in parenthesis along with the mean of a set of data. For example: 10.8 (0.9) or 10.8 (SD=0.9) or in the use of error bars on plots (Figure 2).

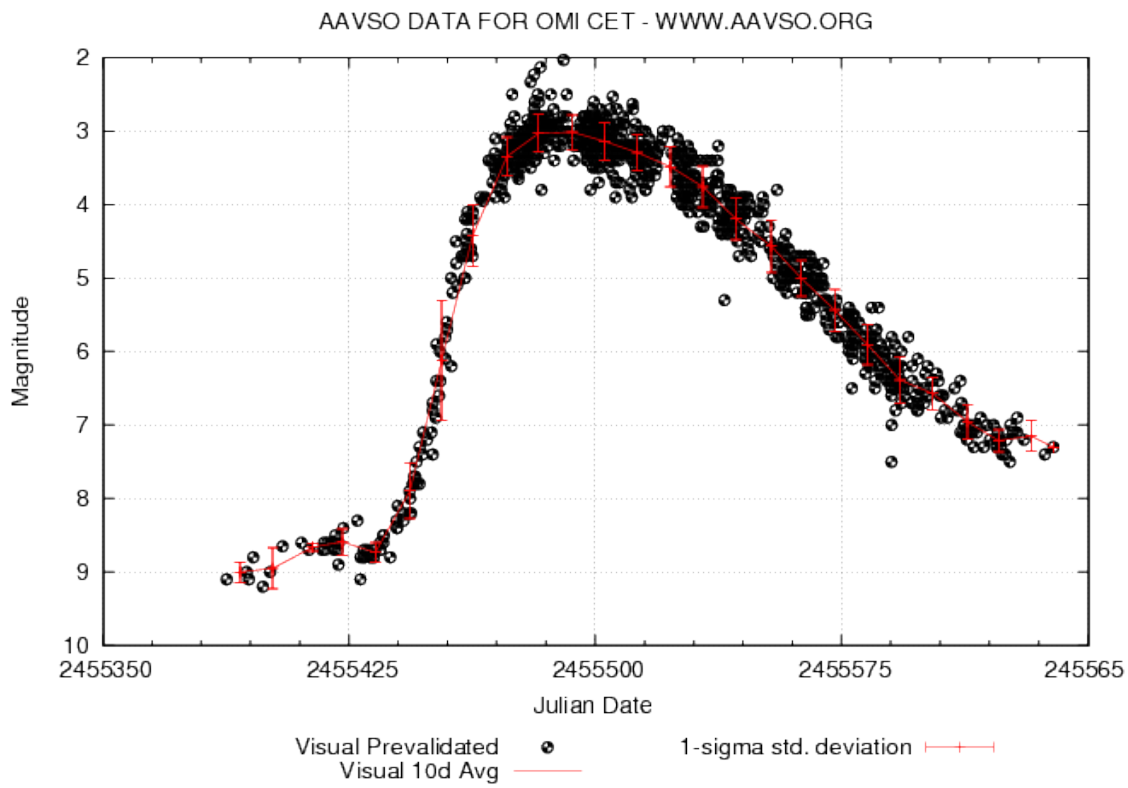

Figure 2. Standard deviation is often shown on graphs as vertical line with brackets at the end. This plot from the AAVSO Light Curve Generator shows 300 days of visual observations of Mira (Omi Cet) with a 10-day mean drawn through the data (red line). The 1-sigma standard deviation is shown at the center of each 10-day section of data. That means that 68% of all data should fall within those brackets. Note that some of the "uncertainty" in the average arises from the fact that the star itself changes in brightness during each 10-day averaged section.

A "standard deviation", as a term, is a specific amount of scatter in the data. That amount is referred to as "sigma" (Figure 3). One sigma of standard deviation means that about 68% of the real data will fall within those values. Two sigma of standard deviation means that about 95% of the real data will fall within those values. Three sigma of standard deviation covers about 99.7%.

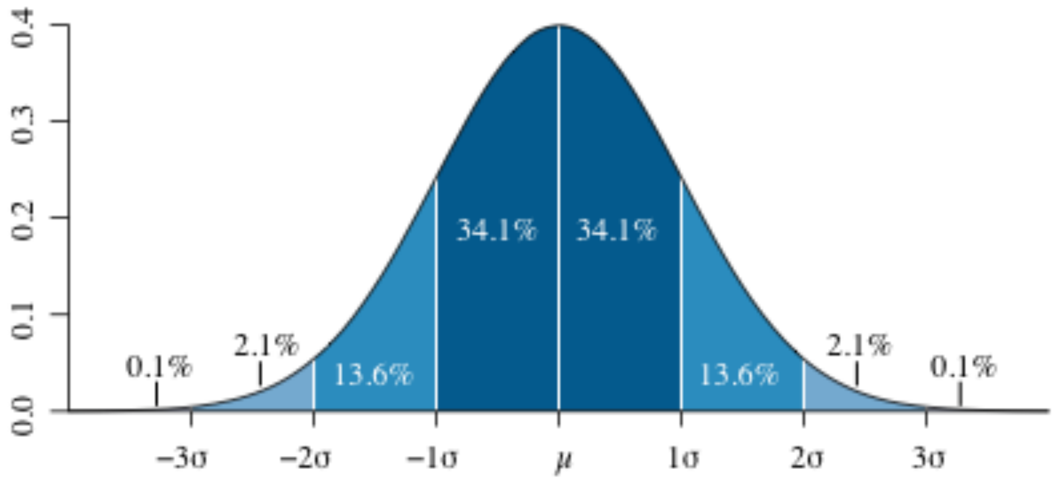

Figure 3. A graph of the standard deviations of a generic population of data. The central peak is the mean of the data. Each white line separates the standard deviations. The dark shaded region is 1-sigma, the moderate blue is 2-sigma and the light blue is 3-sigma. The overall shape of the distribution is known as Gaussian (a.k.a. "normal").

When you see simply "standard deviation" reported with no sigma attached, we recommend treating it as one-sigma. There is no hard standard as to what "standard deviation" by itself means in terms of sigma. But by adopting a conservative measure you are erring on the side of caution.

Standard deviation is relatively easy to calculate manually, if you are so inclined (Figure 4). There are also many online standard deviation calculators. Spreadsheets have built in calculators as well. Excel, for example, will compute it with the STDEV() command.

$$\mathrm{SD} = \sqrt{\frac{\Sigma\ (x_i - \bar{x})^2}{n - 1}}$$

where	$\mathbf{x_i}$	=	the value of each data point
	$\bar{\mathbf{x}}$	=	the average of all the data points
	Σ	=	the Greek letter sigma, meaning "sum of"
	$\mathbf{n}$	=	the total number of data points

Figure 4. The mathematical equation for standard deviation. For more information on basic statistics of datasets consult Chapter 10 of *Variable Star Astronomy* (https://www.aavso.org/education/vsa).

Standard deviation of the mean is related to standard deviation. In fact, it is computed by taking the standard deviation and simply dividing it by the square root of the number of data points in the group. Lets' say you measure the standard deviation of five comparison stars and get a 1-sigma SD of 0.75. You simple would divide 0.75 by the square root of 5 to get a standard deviation of the mean of 0.36.

"Standard deviation" applies to a specific data point within a group of data. As its name suggests, "standard deviation of the mean" applies to the *mean* of that group of data. So when you are comparing means, you should use standard deviation of the mean. But when you are comparing data points, stick with standard deviation. In calculating photometric uncertainty, you'll almost always want to use standard deviation. When analyzing different light curves, you'll want to use standard deviation of the mean most often.

Note: Standard deviation of the mean is sometimes referred to as *standard error* in statistical literature. In astronomy, it is more often referred to as standard deviation of the mean, so we adopt that nomenclature here.

Quadrature

Sometimes you may want to add uncertainties together to come up with a categorical or higher level uncertainty. For example, maybe you want to look at the average brightness of all Mira class variable observations for one night. How do you determine its uncertainty? You can simple take the mean of all the observations and report the standard deviation. However, you lose the uncertainty reported with each observation when you do that. In fact, that assumes that all the measurements had the same level of uncertainty!

So instead of reporting the standard deviation of the mean, you can determine all the uncertainty values associated with the data in *quadrature*. Simply square each of the uncertainty terms, add them and take the square root. For example: To report 3 variable star estimates of 10.0 (0.1), 10.4 (0.2), 10.5 (0.2) with a mean of 10.3 you take the square root of $(0.1)^2+(0.2)^2+(0.2)^2$, which gives an uncertainty of 0.3.

Another conception of how quadrature works is in the discussion of accuracy. Recall accuracy is the combination of two types of uncertainty, precision and systematic. So it can be described with this equation:

$$\text{Accuracy} = \text{sqrt}\{\text{Precision}^2 + \text{Systematics}^2\}$$

Again, you square the sources of uncertainty, add them up and then take the square root.

Significant Figures

A quick and (very) dirty way to determine precision in a measurement is to look at the number of significant digits. A significant digit is any number that reflects the precision of the original instrument. That is, it is a number that was not artificially introduced for another reason besides measurement. In practice, every number is treated as significant if it is not a zero or if it is a zero after a decimal place. For example, both 12,000 and 12,000,000 have two significant digits (the 1 and the 2). However, 12.000 has five significant digits (1, 2, and all the zero's).

So how are significant digits related to uncertainty? First, if there is no uncertainty term associated with a number then the standard practice is to assume the uncertainty is in the range of the location of the last significant digit. Example: 12,000 would be assumed to have an uncertainty of +/-1,000 (the location of the 2). 12.000 would be assumed to have an uncertainty of +/-.001 (the location of the last zero). Another way significant digits are related to uncertainty is in reporting results. The final numbers of a procedure should not have more significant digits than the fewest count of significant digits of any number used anywhere in the analysis. So if you do an analysis using the numbers (12,001 - 12 - 0.012 - 1,222) your final result should only have two significant digits in it (because the twelve only has two significant digits). If you are reporting a true uncertainty with your value, then this rule is irrelevant because the reader does not need to assume anything. However, it is still the accepted practice when writing up publishable results and some editors may complain if you don't follow it.

Caveat: The treatment of significant figures is a debatable topic with no universally adhered to rules. These guidelines are included here to give an idea of how they are generally used in the scientific community. One must always use common sense. There may be times when it makes sense to deviate from these guidelines.

Determining Photometric Uncertainty

Your observing setup and procedures determines the type of uncertainty that you should report to the AAVSO. Figure 5 is a decision tree that describes this process. The reported uncertainty depends on three things: the number of images you have in your bin, the number of comparison stars you use and your photometry software. When possible, we recommend binning your data. However, sometimes more than one observation is not available due to observing restrictions (time, weather, etc.) or all available observations have to be binned in order to get an acceptable signal to noise ratio. Thus, we have created a simplified decision tree to help you decide which technique to used based on characteristics of your observation.

Be warned that this tree is a set of *simplified guidelines*. There are instances where deviation is acceptable. Thus it is very important to not just memorize this tree, but to understand the background concepts described in this document. That will allow you to better decide when you need to deviate from the guidelines.

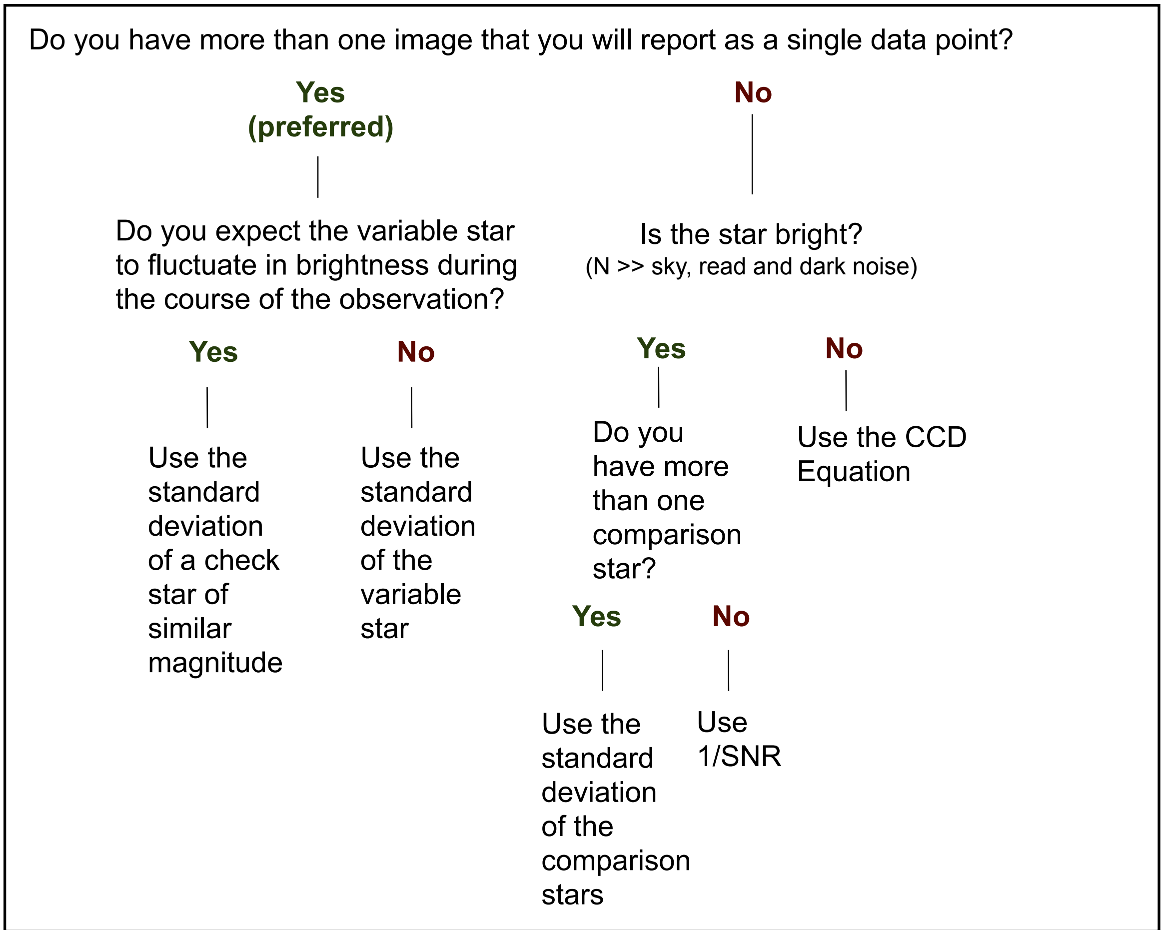

Figure 5. Decision tree for determining how to report photometric uncertainty.

In our descriptions below we rank the preferred methods for determining uncertainty. Try to plan your observations so that you end up near the top of the tree as often as possible as those procedures tend to be more useful. This will increase the quality of your data and make it more useful to the professional community.

The following sections will describe how to determine each of these uncertainties. We will use the VPhot photometry software in our examples (Figure 6). VPhot is a web-based photometry package freely available to members of the AAVSO. It is very easy to use and produces high quality results.

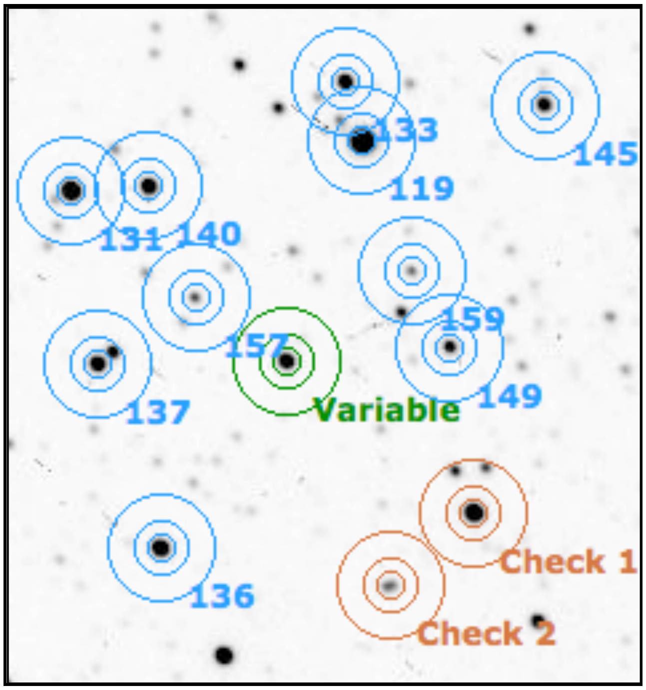

Figure 6. Generic star field in VPhot.

How to Determine if the Variable Star is Fluctuating in Your Data

Part of the Figure 5 decision tree is based on whether you expect the star to be variable within the range of your overall observations. This can be a tough decision. Some star types are known to vary on short time scales (such as the flaring of UGSU's) but the flaring can be random. Should we assume it is flaring right now? Another issue is with the long-term variability. For example, if the star is losing a magnitude per day/week at a consistent rate, then does that slope count as variability?

First, consult the classification of the variable star. Is it of a type that is known to vary on short time scales? If so, then you can assume it will in your data. Another method is to create a scatterplot of uncertainty vs. magnitude with a point for each star in the field. If the variable star is out of line with other stars of similar brightness, then it is likely varying in your data.

Quick Tip: If you are using VPhot with many comparison and check stars, then you can check for variability by visually comparing the uncertainty of the variable star with the uncertainties of the check and comparison stars of similar brightness. This is shown in the "Std." field of the main photometry results table. This is basically the same as the scatterplot method, but with VPhot doing it quantitatively for you. *If the variable's Std. is higher than the comp and check stars of similar brightness, then it is likely varying.*

Technique #1: Standard Deviation of the Variable Star

The first two techniques require more than one observation. It is recommended that you take at least 3-5 exposures of the star and use them to compute both the photometry and its corresponding uncertainty.

For this first technique, simply take the one-sigma standard deviation of the mean estimate of the variable star as measured in all of your frames.

Figure 7 is a screen shot of the results of the time series photometry analysis tool in VPhot. This example is based on an analysis of five images. The "Std" of the "Variable" is the standard deviation of the variable, which is the uncertainty you want to report.

Bins: 1 Sequence: V Stars to plot: ... Ensemble: ... Refresh

	Average	Min	Max	Std	Avg. SNR
Targets					
Variable	13.912	13.901	13.929	0.013	252
Check stars					
Check 1 (0.000)	13.216	13.207	13.222	0.006	356
Check 2 (0.000)	15.452	15.313	15.497	0.079	92
Comparison stars					
119 (11.898)	11.793	11.786	11.802	0.007	751
124 (12.409)	12.286	12.259	12.306	0.018	594
129 (12.858)	12.766	12.752	12.790	0.015	462
131 (13.083)	13.054	13.045	13.070	0.010	400
133 (13.331)	14.142	14.137	14.147	0.004	212
136 (13.630)	13.462	13.448	13.486	0.014	323
137 (13.715)	13.919	13.893	13.934	0.016	252
140 (13.951)	13.864	13.855	13.877	0.009	255
145 (14.468)	14.420	14.377	14.441	0.025	180
149 (14.946)	14.829	14.807	14.843	0.014	138
157 (15.749)	15.575	15.543	15.586	0.018	80
159 (15.892)	15.819	15.785	15.844	0.022	65

Figure 7. Time series results from VPhot.

Technique #2: Standard Deviation of a Check Star

Sometimes the variable star is expected to fluctuate on short time scales. For example, many cataclysmic variables show flares that last on the order of minutes. In these cases, do the same thing as technique #1, except use the check star values instead of the variable star's values. Choose the check star that is closest in magnitude to the variable star and report its standard deviation. The check star used should be within 2 magnitudes of the variable star.

Looking back at the VPhot time series photometry tool screen shot in Figure 7, you can see that we have two check stars. In this case, "Check 1" is the preferred check star to use since it is closer in brightness to the variable. Again, the "Std" field is the number to report as the uncertainty.

Technique #3: Standard Deviation of the Comparison Stars

If you only have one frame but more than one comparison star, you can use the standard deviation associated with the comparison stars. Compute the mean of the estimates of the variable star's brightness when compared to each of the comparison stars, separately. Then report the 1-sigma standard deviation of that mean.

In VPhot, you can do that by simply taking the standard deviation of the mean of all the values in the "Target" field of the single image photometry report (Figure 8).

12 Comparison Stars Toggle Active

Star	IM	SNR	X	Y	Sky	Air	B-V	V-mag	Target estimate	Active
119	-8.389	877	651.498	418.242	337	2.584	-11.898	11.898	14.014	☑
124	-7.880	693	416.946	526.758	331	2.591	-12.409	12.409	14.015	☑
129	-7.399	539	447.129	460.488	334	2.591	-12.858	12.858	13.983	☑
131	-7.133	473	541.935	437.000	328	2.587	-13.083	13.083	13.942	☑
133	-6.125	256	644.883	396.085	332	2.584	-13.331	13.331	13.182	☑
136	-6.697	382	575.806	571.104	333	2.585	-13.630	13.630	14.053	☑
137	-6.317	307	552.234	501.870	328	2.587	-13.715	13.715	13.758	☑
140	-6.328	306	571.345	435.136	330	2.586	-13.951	13.951	14.005	☑
145	-5.810	228	719.720	404.555	332	2.581	-14.468	14.468	14.005	☑
149	-5.376	171	684.012	495.522	333	2.582	-14.946	14.946	14.049	☑
157	-4.595	105	588.652	477.081	329	2.585	-15.749	15.749	14.071	☑
159	-4.357	83	669.732	467.191	332	2.583	-15.892	15.892	13.976	☑

Figure 8. A sample listing of comparison star values from the single image photometry report from VPhot.

VPhot will do this for you automatically – but with a catch. It computes this standard deviation, *then* adds in the statistical uncertainty (computed using the SNR) in quadrature. This gives a slightly more conservative estimate for uncertainty. You may not notice the difference when observing bright stars, but it becomes more apparent with faint stars. In any event, when using VPhot report what appears in the "Err" field of the Target Star Estimate table (Figure 9). If you are not using VPhot and wish to do this manually with the values your own software provides, that is fine.

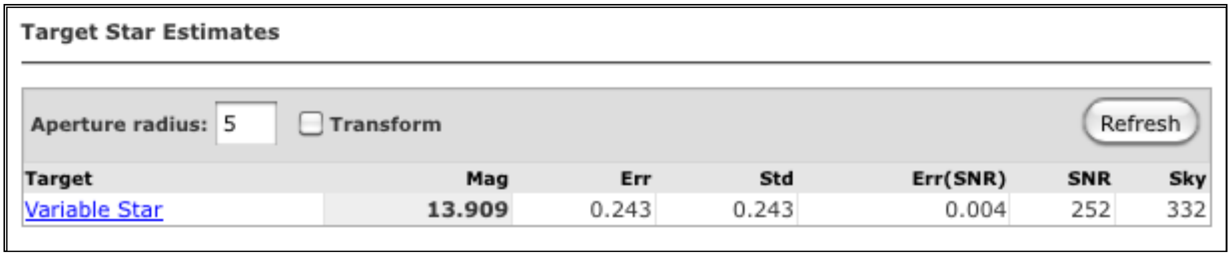

Target Star Estimates

Aperture radius: 5 Transform Refresh

Target	Mag	Err	Std	Err(SNR)	SNR	Sky
Variable Star	13.909	0.243	0.243	0.004	252	332

Figure 9. A sample single image photometry report from VPhot.

Technique #4: CCD Equation

This is perhaps the most complicated procedure for measuring uncertainty since it is done manually (however, one could produce a spreadsheet to automate much of it). First, imagine we don't have to worry about noise sources or the CCD camera or sky background and we just want to know what is the intrinsic uncertainty in our measurement of photons coming from a light source. In the absence of noise, uncertainty is purely Poisson, and are proportional to the square root of the number of photons received in a given time. If you were to receive 10000 photons from your source during a given integration, the Poisson error is 100, or one percent. We could then express the signal-to-noise ratio as

$$N^*/\sqrt{(N^*)}$$

Next, add in the noise sources. Our signal is still N* but what about the noise? If we consider the noise to be the uncertainty in the total number of photons received, then we would need to consider all of the photons from the star (N*) and all of the other photons and other "counts" that fell within our measurement aperture. For the most part, this includes three things: the brightness of the sky itself, the dark current of the CCD (which is effectively like a light source) and the read noise. All three of these affect each pixel within your image, including those within your measurement aperture, so they're added to the term within the denominator.

Let's deal with sky and dark counts first, since they're easiest. Suppose your measurement aperture contains *n* pixels. If each pixel receives Ns photons from the sky, and Nd "photons" of dark current, then you'd have:

$$S/N=N^*/\sqrt{(N^* +n(Ns +Nd))}$$

Next, consider the read noise. Read noise is a type of shot noise and not a Poisson process, and its uncertainty does not go as the square root, but are linear. Since it will go inside the square root, let's add the square of the read noise counts, Nr^2, since squaring cancels out the square root:

$$S/N=N^* /\sqrt{(N^* +n(Ns +Nd +Nr^2))}$$

This is a simple version of what is called the "CCD equation". If you're doing photometry of faint stars, Ns may be of the same order as N*. (If N* and Ns are less than Nd or Nr, then you're probably not exposing long enough!). The standard CCD Equation is (Merline & Howell, 1995; Howell, 2006):

$$\frac{S}{N} = \frac{N_*}{\sqrt{N_* + n_{\rm pix}\left(1 + \frac{n_{\rm pix}}{n_B}\right)\left(N_S + N_D + N_R^2 + G^2\sigma_f^2\right)}}.$$

When your computer records the numbers of photons received, it is actually recording ADUs or "Analog to Digital Units" rather than the raw number of photons you received. Each pixel on your CCD holds photoelectrically generated electrons, and your camera's electronics counts the electrons and converts these counts to ADUs. The gain of the CCD is the number of electrons it takes to create one ADU. If your CCD has a gain of 5, then 5 electrons will yield 1 ADU. If the gain is 100, it will take 100 electrons to get 1 ADU. The CCD equation assumes you're talking about electrons rather than ADU, so in order to use it, you have to include the gain in the calculation. You multiply all of the terms given in ADU by the gain to obtain the number of photons/electrons:

$$S/N = N^*(ADU)\, g \,/\, \sqrt{(\, N^*(ADU)\, g + n\, (Ns(ADU)\, g + Nd(ADU)\, g + Nr^2\, (ADU)\, g)\,)}$$

In this simple case, it simply amounts to multiplying the S/N obtained putting the ADU values into the CCD Equation by the square root of the gain. Now we're practically done. Remember that we are submitting data in magnitudes rather than in the signal- to-noise of the counts, so to find the one-sigma uncertainty in magnitudes, take the inverse of the S/N of each of the variable and comparison star, and add them in quadrature. The result is the one-sigma uncertainty in your magnitude.

For a more detailed explanation of this, as well as a derivation of some second-order terms useful in faint- source photometry, see Chapter 4.4 of Steve Howell's *Handbook of CCD Astronomy*, 2nd edition (2006).

Technique #5: 1/SNR

If your photometry software provides an SNR for the variable star, then simply divide 1 by that SNR and report that figure. This is a very crude uncertainty measurement that only takes into consideration statistical noise. But it is easy to do and better than not reporting uncertainty at all.

Summary

There are many ways to determine and report photometric uncertainty. The important questions to ask are: 1. What type of uncertainty do I want to report? 2. How do I determine it? The flow chart in Figure 5 can help with the former. The information in the rest of this document can help with the latter.

We hope this guide helps you report useful uncertainty values with your photometric data. It is very useful to researchers (both professional and amateur) who use your data. Please feel free to post questions and comments to the online forum for this course.

Further Reading and Acknowledgements

This document was written by Aaron Price with contributions from Arne Henden, Matthew Templeton and Mike Simonsen.

Below are additional resources for further reading on this topic:

- For more discussion of photometry and introductory background information on CCD observing, see the *AAVSO CCD Photometry Guide*: https://www.aavso.org/ccd-photometry-guide
- For a more detailed explanation of photometric uncertainty, as well as a derivation of some second-order terms useful in faint-source photometry, see Chapter 4.4 of Steve Howell's *Handbook of CCD Astronomy*, 2nd edition (2006).
- For a more detailed description of standard deviation, error and uncertainty in time series data see Grant Foster's book *Analyzing Light Curves: A Practical Guide:* http://www.lulu.com/product/paperback/analyzing-light-curves-a-practical- guide/11037112

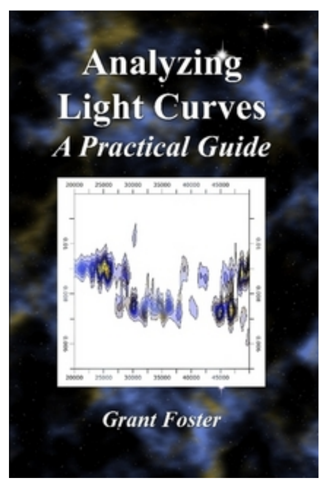

- Arne Henden's preferred reference for uncertainty discussion is *Data Reduction and Error Analysis for the Physical Sciences* by Philip R. Bevington.
- Many examples of applications of uncertainty to specific astronomical research issues are discussed in *Errors, Bias and Uncertainties in Astronomy* edited by Carlos Jaschek and Fionn Murtagh. It was published by Cambridge University Press in 1990.
- The VPhot User's Guide is available to AAVSO members at https://www.aavso.org/download-vphot-user-guide-here.

Appendix A: History of Uncertainty in the AID

The AAVSO International Database has two fields related to uncertainty. One is titled "Uncertainty" and the other "HQ Uncertainty". Prior to 2003, observers were not asked to report uncertainty with their photometric observations. This was largely due to the fact that the AID was historically a visual database that was later extended to accept photometric measurements. This changed in 2003 when the AAVSO asked observers to include an uncertainty term in the "Comments" field[2] of their observation report. The uncertainty estimate was to be prepended by the strings "Err:" or "Error:". Thus, an uncertainty of 0.05 could be reported as "Err: 0.05". Most of the time, this uncertainty was reported using the 1/SNR or a similar method, thus did not include systematic error.

In 2005, the AID was converted into a relational database running MySQL. At this point, a dedicated "uncertainty" field was added. Beginning then, any observation reported with an "Err:" string had that value loaded into the uncertainty field. The field was also reverse populated with Err: strings submitted in prior observations.

In 2006, the AAVSO rolled out a new Extended Format for reporting photometric observations. This new format included a dedicated uncertainty field. Whatever the observer placed in that field of the observation report was loaded into the Uncertainty field in the AID.

During this period, AAVSO observers determined uncertainty using their own techniques. Various online tutorials were written and workshops held at meetings on the topic, but formal instructional guidelines for how to report uncertainty were not provided. As a result, the uncertainty values reported during this time varied greatly from very conservative values that include both systematic and statistical uncertainty to very liberal values that were simply 1/SNR.

In response, a new field was added to the AID called "HQ Uncertainty" in February 2007. This is an observation's uncertainty estimate as determined by AAVSO HQ and not by the observer. It was automatically populated in photometric reports of time series data that were submitted via the AAVSO's online reporting system ("WebObs"). The value is the 1-sigma standard deviation of the variable star's magnitude in the report (hence was only applied to time series data).

The next stage of uncertainty reporting begins with this course. Those who complete the course will have all observations submitted after this date flagged so that the user of the data will know you successfully completed the course. The goal is for the user to know they can apply an extra level of trust to your uncertainty estimates while also standardizing how uncertainty is reported in the AID.

---

[2] CCD Views #314 and #315 available at https://www.aavso.org/ccd-views

Appendix B: Photometry Hints (Apertures, Ensembles and Binning)

This is not a course on photometry so we cannot go into detail in these issues. However, below are two issues that are commonly brought up and have an effect on uncertainty estimates.

- The *annulus size* used for aperture photometry can have an effect both on the photometry measure and the uncertainty estimate. Arne Henden's recommendation is to use 3-5 FWHM diameter when possible. Use a larger FWHM for brighter stars and smaller FWHM for medium and fainter stars. (The default in VPhot is a 1.5 FWHM *radius*, thus the equivalent of 3 FWHM in diameter.)

- What about *the number of stars to be included in ensemble photometry*? Arne Henden made the following suggestion to a student in an earlier cohort of the course:

  "The VPHOT tutorial tells you quite a bit about selection. Your +/- 1.5mag concept is fine; for fainter stars, I tend to use more brighter stars as comps to get better S/N, and for brighter stars, I tend to use fainter comps, not because I want to but because there are fewer stars of comparable brightness to the target from which to choose. If there are plenty of stars to choose from for comparisons, then I typically use 12-20 stars in the ensemble. Kent Honeycutt uses every star in the frame, which is also ok if they are properly weighted. If you get better results with 6 or less stars, then there may be some other reason for the improvement, as the more stars included in the ensemble, the better."

- As you can see from the decision tree, the best way to determine uncertainty is to have multiple observations. *So at what level should you group your data?* You could group all your data into one point but lose time resolution or you could not group at all and sacrifice precision and uncertainty. There is no one answer. It depends on your research goal along with the brightness and expected photometric behavior of the star. However, for general monitoring of a bright star we recommend binning 3 observations. This allows you to determine a good uncertainty measurement while also retaining much of your time resolution. Again, let the science determine your binning size. But when in doubt, bin by 3.

For general photometric advice not related to uncertainty consult the AAVSO Photometry Forum (https://www.aavso.org/forum/5010) or the VPhot Forum (https://www.aavso.org/forum/5015).

Appendix C: About Carolyn J. Hurless

1934-1987

Carolyn J. Hurless was the most active and prolific woman observer in the history of the AAVSO, with 78,876 observations in the International Database. But that only scratches the surface of this remarkable woman's life and career as an AAVSO observer, councilor, officer, mentor and ambassador.

Born in Lima, Ohio, November 24, 1934, Carolyn became interested in astronomy at the age of 13 through her love of science fiction. As a young woman, she was invited to join the Lima Astronomy Club, when President Herbert Speer found her name on the borrowers cards of people who had checked out astronomy books from the public library.

Shortly after that, she decided to make her own 8-inch reflector with the guidance of fellow astronomy club members. When the initial grinding was done, Carolyn found that in her excitement she had hogged out a short focus mirror of f/4, instead of the typical f/8 or f/9 scope most were making at the time. In the end it turned out to be a fine instrument. In fact the short tube length gave her a telescope easily transported and set up for observing. Most of her observations were made with this telescope and she never felt the need to upgrade to something else.

Carolyn learned variable star observing from legendary AAVSO observer and fellow Ohioan, Leslie Peltier. Carolyn would make the trip to Delphos, Ohio to observe faint "inner sanctum" stars with Peltier's 12-inch refractor nearly every week during their lifelong friendship. She was more than happy to pay it forward by mentoring other newcomers and sharing her enthusiasm with other variable star observers around the world.

One way she managed to do that was by publishing the informal monthly newsletter *Variable Views* in which she shared ideas about astronomy, stories of variable stars and amateur astronomers and humorous notes about her own experiences. She started the newsletter at her own expense and published it for 22 consecutive years. Carolyn invited her *Variable Views* readers to summer gatherings each year at Leslie Peltier's home where she was able to inspire young people with her love of the stars and observing.

She managed to reach out and touch people across international boundaries also, in a time when this was not easy to do. She sponsored a Czechoslovakian observer, Jaroslav Kruta, to AAVSO membership. Through persistent correspondence, mainly tape recordings, she taught Jaroslav English, and was able to introduce several other AAVSO members to him by arranging for them to meet when they visited Czechoslovakia.

Besides sharing her enthusiasm for astronomy with the public, she managed to hold down a full-time position as a music teacher, inspiring countless young musicians along the way.

The Carolyn Hurless Online Institute for Continuing Education is proud to carry on in the tradition of this remarkable woman.