

The AAVSO Spectroscopy Manual

Version 1.1 - October 2018



AAVSO

49 Bay State Road
Cambridge, MA 02138
Phone: +1 617 354-0484
Email: aavso@aavso.org

Copyright 2018 AAVSO
ISBN 978-1-939538-42-0

Table of contents

	<i>page</i>
Introduction	3
Acknowledgements	4
Section 1 - Equipment	5
1.1 Principles of Spectrographs	
1.2 Other Equipment Considerations	
Section 2 - Calibration Data and Procedures	7
2.1 Introduction	
<i>A few thoughts about IRAF versus ISIS</i>	
2.2 Calibration Data Dependent Only on the CCD	
2.3 Calibration Data Dependent on Spectrograph Configuration	
2.4 Additional Reference Data for Specific Science Objectives	
Section 3 - Data Processing	14
Section 4 - Common Data Quality Problems	17
4.1 Raw Data Issues	
4.2 Data Processing Issues	
Resources	22

Introduction

The intent of this manual is to provide a thorough introduction to astronomical spectroscopy, to give guidelines for the steps required to calibrate and processes spectroscopic data, and to give an overview of the equipment required. This manual is not specific to any particular telescope/spectrograph system or image processing software, and as such is a guide, not a “recipe book.”

Section 1 gives an overview of the design and principles of spectrographs, and briefly details the equipment required to produce a fully calibrated set of spectroscopic data; Section 2 describes the calibration data required and how it is applied; Section 3 takes a comprehensive look at processing spectroscopic data; and Section 4 briefly discusses potential data quality issues to be aware of.

Acknowledgements

This Manual was written by Dr. Ryan Maderak, Benedictine College with contributions from Joe Daglen, Ken Hudson, and Dr. Stella Kafka, AAVSO Director.

Section 1 - Equipment

1.1 Principles of Spectrographs




Figure 1.1 – Diagram of a Spectrograph (courtesy of Francois Cochard).

Components of a Spectrograph

Slit: acts as the light source, restricts the angular size of the object/source (both to isolate the source and to achieve higher spectral resolution), and causes the initial dispersion through diffraction. In a fiber-fed spectrograph, the fiber acts as a slit.

Collimator lens: transforms the divergent light from the slit into a collimated (parallel) beam; poor collimation will result in vignetting (loss of light from edges of the beam) which will reduce signal or even reduce wavelength coverage.

Dispersion element:

Prism: disperses light due to the wavelength dependence of the index of refraction; blue light is dispersed more than red. Prisms have higher throughput than gratings but with much lower dispersion.

Diffraction Grating: glass plate with a large number of finely-ruled grooves (typically hundreds of lines per mm), each of which acts as a diffraction source; interference between light from each groove increases resolution by making each wavelength “sharper”. In contrast to prisms, gratings disperse red light more than blue. Gratings also have multiple diffraction orders that can be used. Note that while in the basic design described here a transmission grating is assumed, there are also reflection gratings.

Objective lens: focuses light from the dispersion element onto the detector.

Dispersion versus Resolution

As alluded to above, there is a distinction between dispersion and resolution. Dispersion describes the extent to which the light is spread out, i.e. how much wavelengths are separated. But resolution describes the smallest spectral feature that can be clearly distinguished; for example, the resolution is equal to the width of the thinnest spectral line (more specifically, it is the full width at half maximum, FWHM, of the line). Lines closer together than that are blended and not distinguishable. As noted above, resolution is set in part by the size of the slit, and in part by the dispersing element. Slit spectrographs often have slits of multiple sizes that can be selected; alternatively, some have a Vernier which allows adjustment of the slit size.

1.2 Other Equipment Considerations

In addition to a spectrograph, you will of course need a telescope OTA, mount, CCD, computer, control and imaging software, etc. Spectrographs do not usually require additional control software. There are no special requirements for the CCD imager, so long as the pixel size is well-matched to the resolution (2 – 4 pixels per line FWHM), and the format is large enough for the wavelength range of the spectra; it is also helpful if the sensitivity is roughly uniform over optical wavelengths.

Guiding

This is perhaps the most challenging aspect of acquiring a spectrum, as the star must remain on the slit. This could in principle be accomplished with a standard off-axis guide camera, but instead you should guide on an image of the slit itself, and in fact need a camera in this position to center the star on the slit initially.

Flat fielding

The principles of flat fielding for spectroscopy are discussed in Section 2. Spectrographs often have an internal flat field lamp (typically tungsten) which can be used for this purpose. The source only needs to be spatially uniform over the slit of the spectrograph; spatial uniformity is not relevant for fiber-fed spectrographs, unless a second fiber is used to image the background sky spectrum, in which case it must be uniform over both.

Comparison Spectra

The principles of wavelength calibration spectra are discussed in Section 2. Spectrographs often have an internal comparison spectrum lamp (typically Neon and/or Argon). The emission lines are identified and their wavelengths entered, and then a polynomial fit is applied to find the wavelength solution (this assigns wavelength units to individual pixels of the CCD). Neon and Argon are usually used since those gases have a rich emission line spectrum in the optical. In general, you should use a lamp with as many spectral lines as possible to get a good wavelength solution.

Section 2 - Calibration Data and Procedures

2.1 Introduction

The description of data processing that follows is based on the “nuts and bolts” approach that you would apply using a package like IRAF (developed by NOAO – the National Optical Astronomical Observatory), and comparisons between the “IRAF approach” and the increasingly popular ISIS software (written by Christian Buil) will be made.

A few thoughts about IRAF versus ISIS

IRAF (<http://iraf.noao.edu>) has been the principle workhorse of optical astronomy for decades. However, IRAF is no longer being regularly updated by NOAO, and development is now driven primarily by users, specifically through <http://iraf.net> (with NOAO’s approval and encouragement). This shift is in part a result of the development of new image processing routines developed in languages such as IDL, which are perhaps more “elegant.” IRAF remains a valuable tool due to those decades of development and a still very substantial user base. Many professional astronomers prefer a package like IRAF because of the direct control it provides over every aspect of image processing – in fact the array of “knobs and switches” in IRAF is probably daunting to a new user. However, if one needs to understand how the fine details of processing impact finished spectra at the 1% level, the very mathematical structure of IRAF is an extremely useful tool.

Nonetheless, for non-professional spectroscopists, IRAF is probably not the best option. The many steps in IRAF required to produce a finished spectrum use a number of different packages; locating and learning how to use these may be frustrating to a new user. In addition, you will need to locate a good user-written manual for IRAF. In contrast, ISIS, for example, has implemented many of the steps one would perform in IRAF in an elegant way that does not lack rigor (albeit without the level of precise control possible with IRAF). Undergraduates at small academic observatories may also benefit from the more teaching-friendly ISIS. For most new users, software like ISIS may very well result in a better up-front understanding of what the processing steps applied to spectra actually do, in addition to a less-steep learning curve.

The ISIS website provides thorough manuals for processing data for each of Shelyak’s spectrograph models, and are a recommended resource. http://www.astrosurf.com/buil/isis/isis_en.htm

Spectroscopy requires the same basic calibration data as for photometry, but with a few important additions (for example, wavelength calibration data). Also, some of the calibration steps, while similar in purpose to the analogous ones for photometry, must be handled differently (for example, flat fielding). We describe the calibration procedures here, and they are presented in order, but note that there are additional steps, and we will put all of this in the sequential context of the overall processing procedure in Section 3.

Calibration image sets of each type (with the exception of wavelength calibrations) should be taken at least once during every observing session.

2.2 Calibration Data Dependent Only on the CCD

For these calibration types, only a single set is absolutely necessary (for each CCD binning configuration) per session, even for multiple spectrograph configurations.

Biases (aka Zeros)

A snapshot of the CCD background level (zero exposure time), to be subtracted from all images. Best practice is to acquire a substantial number (at least 10, if not 20) and then take the median of them (pixel by pixel) to produce a master bias. If observations span many hours, consider taking a set at both the beginning and end of the night.

Bias correction can be handled with *ccdproc* or *imarith* in IRAF. ISIS can be used to both produce a master bias image and apply the correction (see their manual for details); alternatively, you could first produce the master bias image using another image utility, and then simply supply that file to ISIS.

Darks

Integrated exposures with the shutter closed. Due to thermal noise, extra electrons accumulate in each pixel and introduce additional electronic noise to the data; this is called dark current. It is temperature dependent, and can be minimized by cooling the CCD to 20 – 30°C below ambient; most CCDs have built in thermoelectric cooling that can achieve this. In addition, there may be “hot pixels” that accumulate charge at a higher rate.

There are two approaches to correcting this. (1) If you have a limited set of exposure times, you can simply take a substantial number of darks at each exposure time (at least 10), make a “master” dark for each (by median combining them), and then subtract the master dark from the corresponding images. This also eliminates the need for a separate bias correction. (2) If you have many different exposure lengths, then assuming that the dark current is scalable (linear) with exposure length, take several long (30 minutes or more) dark exposures, bias subtract them, scale to the appropriate length, and subtract from the bias corrected images.

As with bias subtraction, dark subtraction can be done in IRAF using *ccdproc* or *imarith*, but note that *ccdproc* can do dark scaling automatically, if that option is enabled. As with master bias subtraction, ISIS can be used to both create the master dark (if you did not do this with another image utility) and apply the correction; ISIS can also scale darks.

2.3 Calibration Data Dependent on Spectrograph Configuration

A calibration set for each of these types will be necessary for each combination of grating/wavelength region/dispersion used.

Flats

Well-exposed spectra of a continuous (thermal) source. As with biases, the best practice is to acquire a substantial number (at least 10), and average them to produce a master.

As is the case for photometry, flats are used to correct for pixel-to-pixel variations in sensitivity (i.e. spatial variations in the instrument response). This is done by fitting a curve to the flat field spectrum in the dispersion direction, and then dividing the flat by the fit to produce a “normalized” flat. All other spectra can then be divided by this normalized flat. In this way, only the pixel-to-pixel variation is removed, rather than other components of the instrument response.




Figure 2.1 – Plot along the dispersion direction for a raw flat-field spectrum, showing the pixel-to-pixel variation and the instrumental response.

Flats should not be used to correct for the total instrument response (which depends on the CCD sensitivity as a function of wavelength, the grating, etc.), because that assumes a uniform flux across all wavelengths. That correction, called *flux calibration*, requires a different procedure to

be described later. It is not absolutely necessary to do a flux calibration, except for a few specific kinds of science objects.

Once again, ISIS can both produce a master flat and apply the flat field correction; you do not need to produce a normalized flat, ISIS takes care of this for you. In IRAF, the steps described above can be done using *fit1d* and *imarith*.

Comparison Spectra

A spectrum of a source with known emission line wavelengths, which is used to find the wavelength scale (dispersion solution) for your spectra. Only one well exposed comparison spectrum is needed, but be careful not to over expose and saturate strong lines, or else the line centers cannot be accurately determined.

The solar spectrum could be used if necessary. Simply take a spectrum of the daytime sky, which gives you the solar spectrum via scattered sunlight. Refer to a solar spectrum atlas for wavelengths of lines in the solar spectrum. Be careful to avoid using atmospheric absorption lines (aka telluric lines).




Figure 2.2 – Plot along the dispersion direction for a comparison spectrum, showing emission lines.

Identifying lines in the comparison spectrum is similar in ISIS and IRAF (though the interface is different – in IRAF you use the routine *identify*): you first identify a line (or a few lines) manually by entering the wavelength, and then other lines can be identified automatically from a supplied list (IRAF has comparison line lists included, you simply have to specify which type of lamp

was used). In ISIS you first supply an estimate of the dispersion, whereas in IRAF this is not necessary. In each case, you must specify the order of the polynomial fit. The lowest order that fits the data well is best; even in the case where you have a large number (20 or more) reliably identified comparison lines, a polynomial order greater than 5 – 6 is probably not appropriate. In the case where you have only a few lines, the order of the fit should be at least one degree less than the number of points (for example, for 3 lines, at most order 2 should be used).

For identifying lines, it may be helpful to find a comparison line atlas (see for example <http://iraf.noao.edu/specatlas/>). But keep in mind that relative strengths of the emission lines depends in part on the specific lamp used, so pay attention to both the strength and the spacing of the lines.

ISIS will apply the wavelength solution to your object spectra more-or-less automatically. In IRAF, after *identify* you must use the routines *refspec* and *dispcor*.

Sky Background Spectra

Spectra of the sky background should be acquired near each of your object spectra. The atmosphere produces various emission lines, and the sky might also contribute a scattered white light background, both of which should be subtracted out. For slit spectrographs, the portions of the slit not on a star, image the background sky; this is ideal as the sky background from the slit can be subtracted without additional correction. In the case of fiber fed spectrographs, you must either take a second spectrum of the same exposure length away from the star, or use a second fiber to acquire a simultaneous sky (but note, if another fiber is used, you will need to apply a *throughput correction* to account for the different response of the two fibers).




Figure 2.3 – Raw multi-object fiber-fed image of an open star cluster, showing both object spectra and sky background spectra. Note the bright sky emission line “dots” on both the faint sky background spectra and the brighter object spectra.

For slit spectra, ISIS and the IRAF routine *doslit* both perform sky subtraction by sampling then subtracting the sky background from user-defined intervals outside of the object spectrum. For a fiber fed spectrograph, you would simply treat the fiber(s) with the background sky exposure as if it were an additional object spectrum, then subtract it.

Note: for short exposures of bright objects, sky subtraction may not be necessary. The contribution is significant for long exposures. Sky subtraction is important for long exposures and/or faint objects, because the sky signal will be a substantial fraction of the object signal. Also, the sky background is dependent on the observing conditions.




Figure 2.4 – Sample sky background spectrum, showing sky emission lines.

Flux Calibration Standards

For some stars, high quality reference spectra exist that have been carefully corrected for instrumental response. Dividing your spectrum by the reference spectrum (for the same star) will allow determination of your instrumental response, which can then be used to produce spectra on an absolute rather than relative flux scale. These calibration spectra must be taken under photometric (completely cloud free) conditions.

Flux calibration, while beneficial, is not absolutely necessary, except for certain specific science objectives (for example, if you intend to fit a Plank function, a.k.a. blackbody curve, to your spectrum).

ISIS has a built in function for flux calibration, and even has the reference spectra for some flux standards pre-loaded. For IRAF, you would need obtain a reference spectrum, and use the routine *sarith*.

2.4 Additional Reference Data for Specific Science Objectives

Radial Velocity Standards

Spectra of stars with well-known radial velocities, which can be used to check your radial velocity scale, and if necessary, correct it. Recommended if you intend to produce radial velocities on an absolute scale.

Daytime Solar Sky Spectrum

The spectrum of the Sun as acquired from sunlight scattered by the atmosphere. Recommended if you intended to perform a spectroscopic chemical abundance analysis – solar abundances are often used as zero points for the abundance scale.




Figure 2.5 – Sample daytime solar sky spectrum, in the H α region.

Section 3 - Data Processing

We will now give a sequential list of calibration and processing steps. Those steps already described in Section 2 will simply be listed; steps being introduced here will be described.

Note: it is not advisable to co-add (stack) raw images, even if they were taken one after the other. Produce processed, wavelength calibrated spectra first, then co-add.

A. Bias subtraction

B. Dark subtraction

– Note, as alluded to above, a separate bias subtraction is not necessary if you are not using scaled darks (which have had the bias subtracted).

C. Flat fielding

D. Aperture Tracing and Extraction

A raw spectrum spans at least several pixels in the direction perpendicular to the dispersion. This is called the aperture. These pixels must be defined and then summed into a single one-dimensional spectrum. Defining the pixels to sum over is called aperture tracing, and summation itself is called extraction.

ISIS traces the aperture by having the user select the row or column closest to the center of the aperture, and then measuring the slant angle between the actual aperture and that row/column. This can be thought of as a linear trace. The aperture is then extracted more-or-less automatically.




Figure 3.1 – Plot perpendicular to the dispersion direction, showing the spectral aperture that must be summed over.

E. Cosmic Ray Subtraction

Cosmic rays are familiar to anyone who has taken CCD data. They appear as bright dots or streaks in the image. ISIS can remove them as part of the aperture extraction process. There are a number of different routines in IRAF (including external user-written ones) that can be used.

It can be argued that in the case of a spectrum, it might be better not to remove cosmic rays, because the removal is likely to leave affected spectral features distorted, and it might be better to have the cosmic ray there to let us know that the spectral feature has been affected.




Figure 3.2 – Spectrum showing a cosmic ray spike.

F. Wavelength solution (using comparison spectra)

G. Sky subtraction

H. Heliocentric Doppler Correction

The observed radial velocities of objects depend on the Earth's orbital motion. Therefore observed radial velocities vary not only over the course of the year but also over the course of the night, and depend on the coordinates of the object. If we want to produce a radial velocity curve for an object, then we must correct for this, and place the radial velocity in the Sun's reference frame.

In IRAF this can be done using the routine *dopcor*. ISIS does not have a straight forward way of performing this correction (if you could calculate it independently you could simply shift the spectrum, but this calculation is non-trivial).

I. Continuum Normalization and/or Flux Calibration

The continuum is the “baseline” level of the spectrum, visible where there are no spectral lines. Continuum normalization involves fitting a polynomial to the continuum and then dividing by the fit to normalize the continuum level to 1. One of the primary purposes of continuum fitting is if you want to measure the **equivalent width** of spectral line, which is a measure of the total emission or absorption.

IRAF and ISIS handle continuum fitting in the same way, but with different interfaces. In IRAF, use the routine *continuum*. You should use the lowest polynomial order possible, especially in regions with strong lines.




Figure 3.3 – Comparison of a wavelength-calibrated spectrum before and after continuum fitting.

Section 4 - Common Data Quality Problems

4.1 Raw Data Issues

The following issues occur during data acquisition, and can in principle be identified in real-time.

Saturation (and Non-Linearity)

As with photometry, saturation means irrecoverable loss in data accuracy. Indicators that your spectrum is compromised due to saturation or non-linearity include (but are not limited to):

- Emission lines with flat tops (they should be roughly Gaussian).
- Flattening of the spectral response (in a plot along the dispersion direction); this can usually be checked by comparison to rows away from the center of the aperture.
- Flattening of the aperture profile in the direction perpendicular to the dispersion (this should be roughly Gaussian).
- A perfectly uniform continuum; even a very high signal-to-noise spectrum should show some random variation (though you may have to zoom in to see it).




Figure 4.1 – Plot in the dispersion direction for a raw flat field spectrum showing flattening of the spectral response due to saturation.

Under-Sampling

If your spectrum contains no substantial stellar spectral features, then you have probably (a) under-exposed; or (b) not centered, drifted off of, or missed your target.

Under-exposure: at the bare minimum, your target is underexposed if the weakest spectral lines are not statistically distinguishable from the noise. There are rigorous criteria for this (for

example, if the *equivalent width* of the line is less than about 3 times that of a typical noise feature), but practical indicators are that there are no spectral features that are clearly stronger than the largest noise features and/or that there are no spectral features that include the expected number of pixels for a “resolution element” (at least 2 – 4 pixels, as noted in Section 1).




Figure 4.2 – An underexposed spectrum, i.e. low signal-to-noise ratio.

Off-center or Missed Target: This will result in sky emission lines being disproportionately strong compared to the stellar continuum. If telluric absorption lines (due to our atmosphere) are visible in the spectrum, this is an indication that some star/object light has been received (some telluric lines, such as the A and B bands due to O₂, are strong enough to be seen even a low continuum exposure).




Figure 4.3 – Spectrum in which the stellar continuum is disproportionately dominated by sky background, indicating poor centering or a missed target.

Wrong Object

If you have not observed a target before, it may not be readily obvious that you have selected the wrong star. A helpful indicator when centering the star is whether it is as bright as stars of similar magnitude; this same indicator of course applies after you have acquired the spectrum, if the continuum counts are much lower than expected (but see the above discussion of under-exposure) for a star of that magnitude. A more precise, but more difficult indicator is whether the star is of the correct spectral type. This can be gauged by comparing the spectrum to an appropriate template; broadly speaking, the absorption line strength of most elements increase with decreasing temperature, i.e. later spectral types.

For stars showing emission lines, such as spectral types Be or WR, the emission lines themselves provide an excellent indicator that you have the correct object. We also call attention to resources such as the BeSS database (for Be stars specifically), where users can easily compare their spectra to others for the same star.




Figure 4.4 – Sample stellar spectra. Note how the number and strength of absorption lines increase from O stars (hottest) to M stars (coolest).

4.2 Data Processing Issues

Poor Normalization

This occurs when the continuum has not been fit appropriately, and results in substantial variations in the continuum level. Normalization issues most likely result during flat-fielding or continuum fitting. This can be evaluated by examining a region of continuum with few strong spectral features. Another possible manifestation of poor normalization is that if the wings of very strong lines have been over-fit, then part of the line profile will be normalized out, which can be more difficult to recognize.




Figure 4.5 – Spectrum showing poor continuum normalization.

Poor Wavelength Solution

This can occur if there are too few comparison lines used, or if an inappropriate fit order is used. The result is a wavelength scale that is skewed or not very uniform over some parts of the spectrum (note that this is likely to affect only parts of the spectrum, portions without many comparison lines being most susceptible). This may not be readily obvious in the finished spectra. Good indicators are that (a) lines in one part of the spectrum give radial velocities that are systematically different from other parts, or (b) the FWHM of the lines being systematically different from other parts of the spectrum (though this would be an extreme case).

Poor Sky Subtraction

Perhaps the easiest processing issue to recognize, as the effects on the spectrum can be quite dramatic. This is more likely to be an issue for fiber-fed spectrographs, in which the object and sky cannot be imaged simultaneously with the same throughput. There are three cases:

1. Under-subtraction, which will leave residual emission lines
2. Over-subtraction, which will result in artificial “absorption lines”
3. Wavelength shift between the object spectrum and the sky spectrum, which will result in artificial absorption line next to a residual emission line, but neither having the correct FWHM for a spectral feature.

#1 and #2 are most likely due to poor normalization between the object and sky spectra, but could also result, for example, if using a fiber-fed spectrograph with a sky fiber that is too close to another star. #3 is simply a wavelength solution issue.

Note that poor sky subtraction can also affect the continuum level.




Figure 4.6 – Comparison of a spectrum before and after imperfect sky subtraction.

Resources

AAVSO - <https://www.aavso.org>

BeSS - <http://basebe.obspm.fr/basebe/>

Comparison line atlas - <http://iraf.noao.edu/specatlas/>

IRAF - <http://iraf.noao.edu>

ISIS - http://www.astrosurf.com/buil/isis/isis_en.htm