PyWiFeS: A Rapid Data Reduction Pipeline for the Wide Field Spectrograph (WiFeS)

Michael J. Childress[1,2] • Frédéric P. A. Vogt[1] •
Jon Nielsen[1] • Robert G. Sharp[1]

Abstract We present PyWiFeS, a new Python-based data reduction pipeline for the Wide Field Spectrograph (WiFeS). PyWiFeS consists of a series of core data processing routines built on standard scientific Python packages commonly used in astronomical applications. Included in PyWiFeS is an implementation of a new global optical model of the spectrograph which provides wavelengths solutions accurate to ∼0.05 Å (RMS) across the entire detector. The core PyWiFeS package is designed to be scriptable to enable batch processing of large quantities of data, and we present a default format for handling of observation metadata and scripting of data reduction.

1 WiFeS and PyWiFeS: An Introduction

The Wide Field Spectrograph (WiFeS – Dopita et al. 2007, 2010) is an image-slicing integral field spectrograph built at the Research School of Astronomy and Astrophysics (RSAA) of the Australian National University (ANU). WiFeS is continuously mounted on the ANU 2.3m telescope at Siding Spring Observatory (SSO) in New South Wales, Australia. The medium wavelength resolution and wide field of view of WiFeS make it ideal for a multitude of studies such as galaxy kinematics, radial velocity measurements of stars hosting planets (Bakos et al. 2013; Bayliss et al. 2013),

Michael J. Childress

Frédéric P. A. Vogt

Jon Nielsen

Robert G. Sharp

[1] Research School of Astronomy and Astrophysics, Australian National University, Canberra, ACT 2611, Australia.

[2] ARC Centre of Excellence for All-Sky Astrophysics (CAAS-TRO).

spatially-resolved emission line measurements in galaxies undergoing star-formation (Green et al. 2010) or gas-rich mergers (Vogt et al. 2013), and the study of narrow emission lines from supernovae interacting with a circum-stellar medium (Fraser et al. 2013; Inserra et al. 2013).

The first data reduction pipeline for WiFeS was developed for use with the NOAO IRAF software and was modeled on the data reduction package for the Near-IR Integral Field Spectrograph (NIFS – McGregor et al. 2003). A need was identified for a rapid data reduction pipeline which could yield fully processed data in nearly real time, thereby adding the capability of real-time classification of astrophysical transients. The PyWiFeS data reduction package was developed to meet this new criterion.

PyWiFeS is designed to make use of existing Python libraries commonly used in astronomical research. The only core requirements are the NumPy, SciPy, and PyFITS packages, with one non-required but strongly recommended dependency on Matplotlib. PyWiFeS is written as a series of data processing routines (functions) that operate on data stored in Flexible Image Transport System (FITS) format and produce output which is also stored in FITS files, without the need for complex new data classes. It is thus intended to be a flexible package which can be run directly via function calls in the Python interpreter, or scripted to perform batch data reduction.

Much effort was expended to make data formats described by abstract variables as much as possible, so that the pipeline is not tailored specifically to WiFeS except where absolutely necessary. The code is open source and could potentially be adapted for reduction of data from future image-slicing integral field spectrographs, which will play an important role in the next generation of ground-based (e.g., GMTIFS; McGregor

et al. 2012) and space-based (e.g., WFIRST; Spergel et al. 2013) telescopes.

The focus of this paper is to present the core structure and data processing routines of PyWiFeS. This software will be continuously maintained by the RSAA and made available to the public via the PyWiFeS Wiki [1]. The primary Python routines comprising PyWiFeS are briefly described in Section 2, while details on function call syntax and examples are available from the PyWiFeS Wiki. The most complex process in WiFeS data reduction is derivation of the wavelength solution, which we describe thoroughly in Section 3. Finally, the default data reduction script built from the PyWiFeS routines and the preferred PyWiFeS metadata format are described in Section 4.

2 PyWiFeS Data Reduction Routines

The WiFeS instrument has a $25''\times38''$ field of view (FOV), which the WiFeS image slicer splits into 25 $1''$-wide "slitlets". Light from these slitlets passes through a beamsplitter (dichroic) and is sent to the blue and red arms of the spectrograph where it passes through a volume-phased holographic (VPH) grating. This disperses the light into the equivalent of longslit spectra for each slitlet, which are collected on 4k×4k CCDs. An example of raw WiFeS data from an observation of the supernova SN 2012ec in the nearby galaxy NGC 1084 is shown in Figure 1. The profile of each slitlet on the CCD is significantly separated from that of its neighboring slitlet in order to enable nod-and-shuffle observations for optimum sky subtraction (see Dopita et al. 2010, for details).

PyWiFeS image processing routines come in a variety of styles: some operate on raw WiFeS CCD frames, some operate iteratively on data in each slitlet, and some operate on the global three-dimensional data (commonly referred to as a "data cube"). We outline the operations employed by the major PyWiFeS processing routines in the sections that follow. Many of the operations require knowledge of certain aspects of the detector characteristics, or default line lists for wavelength solutions, etc. All of this information is stored in a Python dictionary packaged in Python pickle format in the WiFeS metadata file, which we refer to as needed below. Some routines also require knowledge of the instrument wavelength solution, which is derived using the detailed routine described in Section 3, but for this Section we treat it as having already been derived.

2.1 Image Pre-processing

The first step in CCD data reduction is to extract the science pixels from the raw data and convert that data from ADUs to real photon counts (i.e., electrons). This is done by measuring and subtracting the overscan level from the overscan regions of the data, then multiplying the overscan-subtracted ADU values by the gain of the readout electronics. This is done seamlessly for all epochs of WiFeS data with the `subtract_overscan` routine, using detector characteristics stored in the WiFeS metadata file.

The second step of image pre-processing is to remove any cosmetic defects in the detector. The first generation of detectors (Fairchild CCDs) in WiFeS were free of any cosmetic defects. In 2013 a second generation of detectors (E2V detectors) were installed in March and May of 2013 for the red and blue detectors, respectively. These detectors provided much improved performance in terms of long-term stability and read noise, but each have a few columns (1 for red, 2 for blue, out of over 4000) of the CCD which are defective in most rows. Correction of these bad pixels is implemented by simply interpolating across the bad columns using the PyWiFeS routine `repair_blue_bad_pixels` (and its red counterpart). These bad pixels are also flagged in the data quality extensions when the data is separated into the multi-extension FITS format. Fortunately these bad columns are typically far from any important galaxy emission lines in the red ($\lambda \sim 5550$Å for R3000 and $\lambda \sim 5590$Å for R7000) and despite potential impact in the blue ($\lambda \sim 4940$Å for B3000 and $\lambda \sim 5040$Å for B7000), these columns comprise only a few hundred $\mathrm{km\,s^{-1}}$ gap in velocity coverage.

Any residual two-dimensional structure in the CCD bias level is removed using bias frames. Typically, this step is performed by median combining several bias frames into a *superbias.* This method is limited by the stability of the detector bias level, which was found to vary both spatially and temporally (on a time scale of minutes to hours at the level of a few e^-) in the first generation of detectors in WiFeS. We illustrate these variations in Fig. 2, by comparing 11 blue biases observed over a 13 hours interval corresponding to one observing night (including afternoon calibrations). As illustrated in the top panel, the shape of WiFeS biases is highly non-linear along the x dimension. The mean bias level is subject to variations of 1 e^- within the first hour, and again at the end of the night. The largest shape variations of the mean bias level occur for x positions around 3500 and beyond.

These spatial and temporal variations limit the precision of a standard superbias subtraction method. A
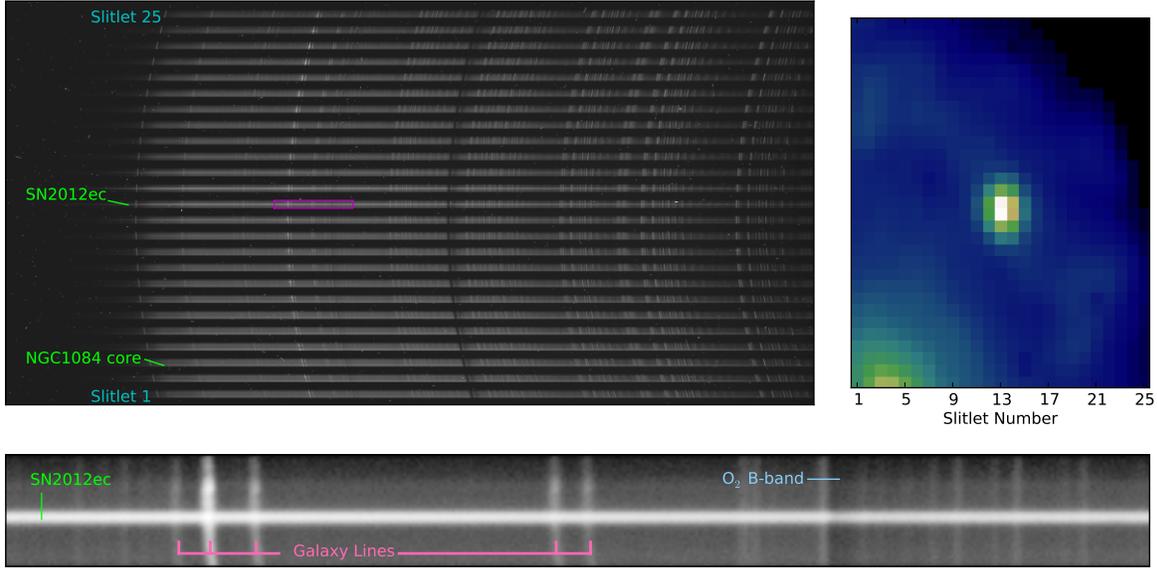
Fig. 1 Raw data (top left) from a WiFeS observation of SN 2012ec (Maund et al. 2013) in the nearby galaxy NGC 1084. The curved wavelength solution of WiFeS can be seen from both the night sky lines and the atmospheric O_2 A-band (dark band near the center). In the collapse image from the final process data cube (top right) SN 2012ec is visible near the center of slitlet 13, while the core of NGC 1084 is centered near the bottom of slitlet 3. These objects can also be seen in the full raw image. A zoomed-in view (bottom panel) of a region of the raw data (highlighted in purple in the raw data image) shows the SN 2012ec trace, as well as emission lines from NGC 1084 (specifically Hα straddled by N II $\lambda\lambda$6548/6584 along with the S II $\lambda\lambda$6717/6731 doublet), and the atmospheric O_2 B-band.

solution to this problem, first implemented on WiFeS data (in a 4-amplifiers read-out mode) by Rich et al. (2010), is to fit a multi-dimensional surface to the bias taken closest to any given exposure, and use this *locally reconstructed* bias instead of a global superbias. This bias fitting method was later on officially added to the original IRAF data reduction pipeline for WiFeS, and is implemented in PyWiFeS with a slightly different algorithm.

A typical blue, raw *superbias* (y-binning of 2 and 1-amplifier read-out mode) is shown if Fig. 3 (top frame). The underlying structure, although somewhat masked by noise, is clearly visible. Red biases display a similar behavior, and are not illustrated here. Typically, a large number of bias frames are co-added (using the `imcombine` routine in PyWiFeS) to create a raw superbias, from which the bias structure is fitted in a two step process:

1. The raw bias is collapsed and median averaged along the y dimension for all x position (see the top panel of Figure 2). This step was introduced first, to correct for the largest spatial variations which occur along the x dimension.

2. The spatial variations along the y dimension are obtained by performing a 2D smoothing (using a symmetric gaussian kernel of 50 pixels) of the raw superbias minus the 1-dimensional correction obtained in step 1, and subtracted on a row-by-row basis.

In Fig. 3, we show in the middle panel the reconstructed superbias based on the raw superbias in the top panel. In the bottom panel, we show the residual after subtracting the reconstructed bias from the raw bias. As expected, the large scale spatial variations have been removed and the residual is flat at a sub noise-level.

In the default reduction script (see Section 4.2), this master *reconstructed superbias* is used as the default to perform the bias-subtraction step on any given exposure. It is usually constructed from a series of biases acquired during the afternoon. Individual bias frames acquired by the observer throughout the night can be associated with particular science frames (see Section 4.1), and will be used by the reduction script to create *local reconstructed biases* which are then subtracted from the requested science frame. Although these local reconstructed biases rely on individual bias frames, they are in fact noise free, and represent the
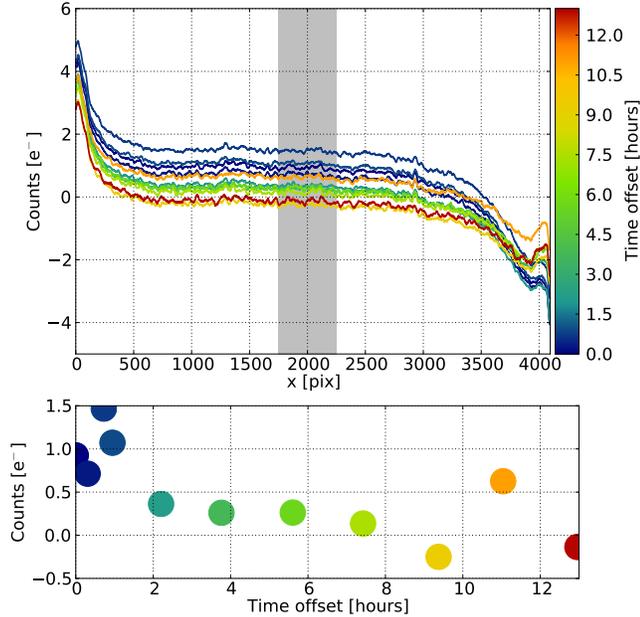
Fig. 2 Comparison of 11 blue biases (y-binning of 2 and 1-amplifier read-out mode) acquired over 13 hours on August 14th, 2012. Top: median bias level of each frame for all x positions. Each curve was smoothed using a Savitzky-Golay routine for clarity (box radius: 51, order: 3, see Savitzky and Golay 1964; Press et al. 2007). Bottom: average of each curve in the top panel, calculated over the grey interval [1750:2250] along x, plotted as a function of the time offset after the first bias was acquired.

most accurate way to perform the bias subtractions step with the WiFeS instrument with first-generation detectors, and best account for the spatial temporal variability of their bias levels.

It is unclear at this stage how the bias levels of the recently installed second-generations detector behave. Once they will have been characterised, PyWiFeS will be updated (if required) to ensure the most appropriate bias subtraction method is being used.

2.1.1 Separating the Slitlets

After basic image pre-processing, reduction algorithms for WiFeS data generally operate on individual slitlets. This requires data for each slitlet to be isolated from the full CCD frame, which is accomplished by first measuring the regions of the CCD occupied by each slitlet. These "slitlet profiles" are measured from a flat lamp image by calculating the normalised flux in the flat exposure as a function of the vertical direction on the CCD. The center of the slitlet is measured as the midway point between the two edges of the slitlet, which we define as the highest and lowest rows for which the flat lamp flux exceeds 10% of the maximum flux. The
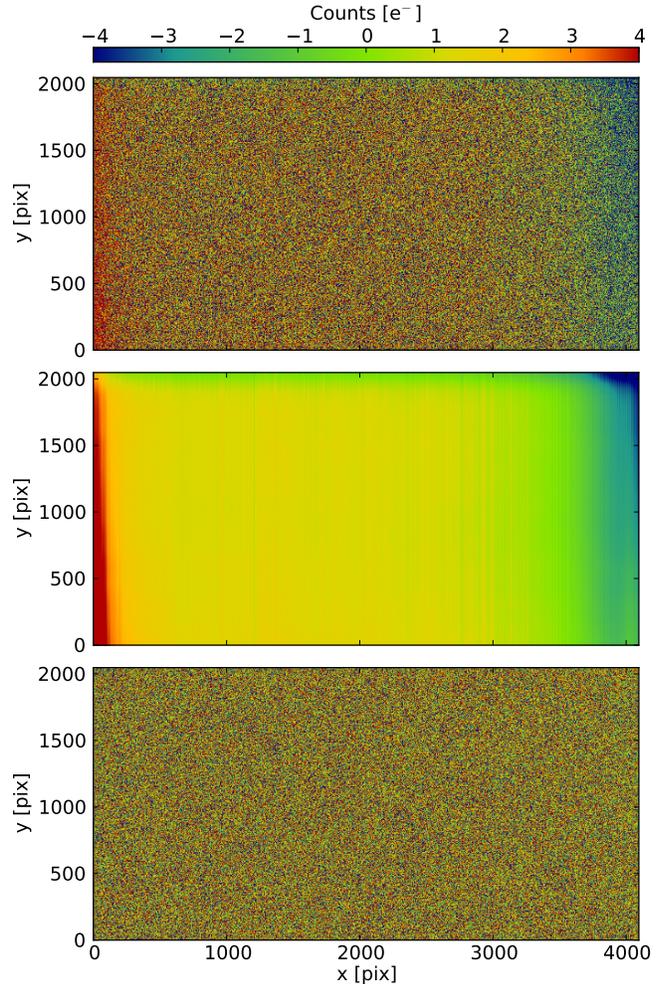
Fig. 3 Top: blue, raw superbias, obtained from combining 5 bias frames previously overscan-subtracted individually. Middle : reconstructed, noise-free, superbias. Bottom: raw superbias minus the reconstructed superbias.

slitlet is then defined as the 86 rows of the CCD centered on that middle row. This principle is illustrated in Figure 4.

Data from the full WiFeS CCD is split into individual slitlets and saved in a multi-extension FITS (MEF) file with the `wifes_slitlet_mef` routine. The first extension of this file contains the original FITS header, as well as some keywords added by the preprocessing routines. The next 25 extensions contain the pre-processed data from each slitlet. The middle 25 extensions are the variance extensions, which when created account for the Poisson noise from photon counting as well as the detector read noise. A final set of 25 extensions serve as "data quality" extensions, which are currently only used to flag pixels interpolated over in the cosmic ray rejection step or bad pixels of the second generation detectors repaired in preprocessing. These data quality
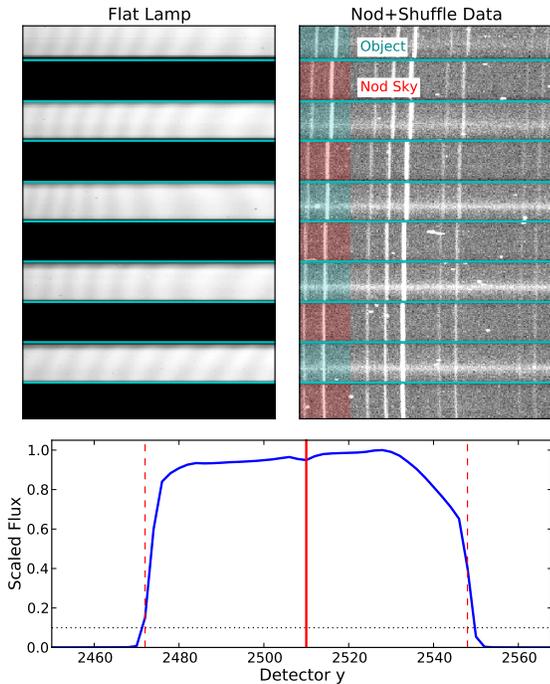
Fig. 4 Top Left: Flat lamp observation used to measure the slitlet boundaries (cyan lines). Top Right: Nod-and-shuffle (N+S) observation, with object slitlets shaded cyan and nodded sky slitlets shaded red. The WiFeS cameras were carefully designed to remove distortion in the y direction, allowing for separation of the slitlets by simply selecting the relevant detector rows. Bottom: Fit of the slitlet boundaries (vertical dotted red lines) and center (vertical solid red line) from the flux vs. row profile of the flat lamp observation, using a threshold of 10% of the maximum flux (horizontal dotted black line).

extensions are built into the data format to allow for complex data quality flag tracking in future incarnations of PyWiFeS.

For observations in nod-and-shuffle (N+S) mode, the pixels between slitlets are used to store photon counts from nodded sky field observations (see top right panel of Figure 4). A variant of the `wifes_slitlet_mef` routine designed for N+S observations, `wifes_slitlet_mef_ns` saves data from the nodded sky slitlets into a second MEF file specified by the user. Automatic identification of N+S observations and calls to the appropriate MEF saving function are incorporated in the default data reduction script (see Section 4).

2.2 Cosmic Ray Rejection

Cosmic ray (CR) rejection for WiFeS data is accomplished in PyWiFeS by means of a custom implementation of the Laplacian kernel technique devised by van Dokkum (2001). The PyWiFeS CR rejection routine performs additional operations on the data to account for the slanted (sometimes curved) wavelength solution of the instrument, as well as the multiplicity of operating on twenty five slitlets of data.

A schematic representation of the CR rejection steps is shown in Figure 5. The first step in identifying CRs is to subtract a smooth model of the sky background. This is complicated for WiFeS data by the slanted wavelength solution, so the sky model is derived in three steps: (1) resample the data to a rectilinear wavelength grid, (2) smooth with a box-median kernel along the cross-dispersion (y) direction, then (3) transform the smoothed sky spectrum back to the original pixel wavelength sampling. This sky model is then subtracted from the original data, and that subtracted data is convolved with a Laplacian kernel as devised by van Dokkum (2001). A threshold based on noise statistics of the data and detector read noise is then applied to the convolved data and those pixels above the threshold are identified as being contaminated by CRs. A slightly lower threshold is then applied to neighboring pixels to those identified as CR-contaminated. The data in all CR-contaminated pixels is then replaced with by interpolating from nearby CR-free pixels, and the entire process is repeated for the requested number of iterations (the default is three). The CR rejection procedure is performed for all 25 WiFeS slitlets, and can take advantage of parallel processing capabilities by simply passing the keyword argument `multithread=True` to the CR rejection function.

2.3 Flat-Fielding

2.3.1 Constructing the Response Function

The throughput of the WiFeS instrument varies pixel to pixel. At each point, the throughput is a product of losses along the optical path of the instrument and the quantum efficiency of the detector pixel, all of which can vary with position and wavelength. The smooth wavelength-dependent throughput of the instrument (and the atmosphere) is corrected by observing spectrophotometric standard stars, as discussed in detail in Section 2.5. The spatial variations in the overall throughput are corrected by means of *flat-fielding*, where observations of a spectral source that is both spatially and spectrally smooth is used to directly measure and correct these spatial variations.

For WiFeS, the smooth spectrum lamp source is part of the instrument's internal calibration unit. Our tests have indicated that the spatial illumination from the internal calibration unit differs slightly from the on-sky optical path. For this reason, a complete flat-fielding
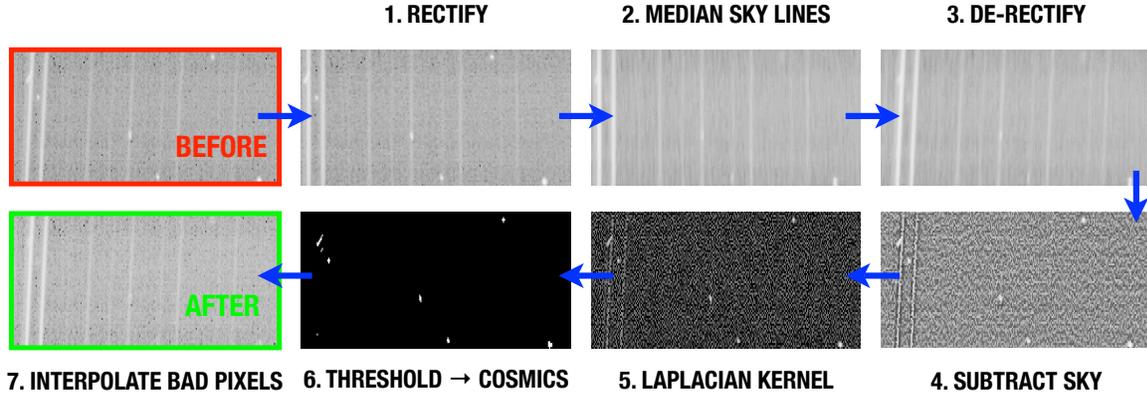
Fig. 5 Schematic diagram describing the cosmic ray rejection procedure employed for WiFeS data in PyWiFeS. The core algorithm is based on van Dokkum (2001), but requires iteration over all slitlets (accomplished via parallel processing) and significant modification to account for the slanted (even curved near detector edges) wavelength solution for WiFeS.

solution requires a complex combination of smooth spectrum source data from the internal lamp unit (*lamp flat* observations) and spatially flat on-sky data obtained by observing the ambient glow of the twilight sky (*sky flat* observations). These two types of calibrations provide respectively the spectral flat-field correction and the spatial flat-field (illumination) correction, which are derived with techniques outlines below and illustrated schematically in Figure 6.

The spectral flat-field correction is derived from lamp flat observations in the following steps:

1. The raw spectrum of the flat lamp is first derived by calculating a median across all rows of the center slitlet. (Note that the dispersion in the center slitlet is acceptably close to vertical to achieve a reliable measure of the lamp spectral shape.)
2. The smooth shape of the lamp spectrum is derived by fitting a low order polynomial to the lamp spectrum in logarithmic flux versus wavelength space.
3. Each row of each slitlet is divided by this smooth spectrum shape, after scaling the smooth spectrum to match the observed flux in the middle half of that row (in practice, we use the middle 2000 columns). The ratio of the observed flux to the smooth lamp spectrum shape gives the spectral response for that row.
4. Division by the (scaled) smooth spectrum shape is repeated for all rows of all slitlets to derive the final spectral flat-field correction.

The spatial flat-field correction (often called the "illumination" correction) is derived from sky flat observations, taking advantage of the fact that each row of each slitlet corresponds to roughly a single spatial pixel ("spaxel") on the sky. The illumination correction is derived in the following steps:

1. The sky flat is first divided by the spectral response function to correct CCD features or variation in the dichroic throughput. It is expected that the dichroic throughput will vary across the instrument FOV due to the differing angles of incidence and the resulting change in the (angle-dependent) dichroic throughput. This must be corrected before the illumination correction (driven by the other instrument and telescope optics) can be derived.
2. The baseline spatial flat spectrum is measured from the middle row of the middle slitlet.
3. Each row of each slitlet is divided by the baseline spatial flat spectrum (after it has been resampled to the wavelength sampling of that row).
4. The illumination correction for each row is calculated as the median value of this ratio (observed spectrum divided by the baseline spatial flat spectrum) in the columns corresponding to a fixed wavelength range, which was chosen to be the wavelength range covered by the middle half (i.e., 2000 columns) of the middle row of the middle slitlet. (Our inspection of the spectrum ratios for each row showed these ratios to be very flat in this central wavelength range, providing a reliable measurement of the throughput for each spaxel.)
5. The preceding two steps were repeated for all rows of all slitlets, yielding a spatial flat-field (illumination) correction for all spaxels in the instrument FOV.

The final flat-field response function is the product of the spectral flat-field correction and the spatial flat-field correction. This master response function effectively corrects for three effects: (1) variations in quantum efficiency of particular pixels on the detector (i.e., "cold" or "hot" pixels), (2) moderate spectral "wiggles" in the overall instrument throughput caused by the dichroic beamsplitter (including spatial variation
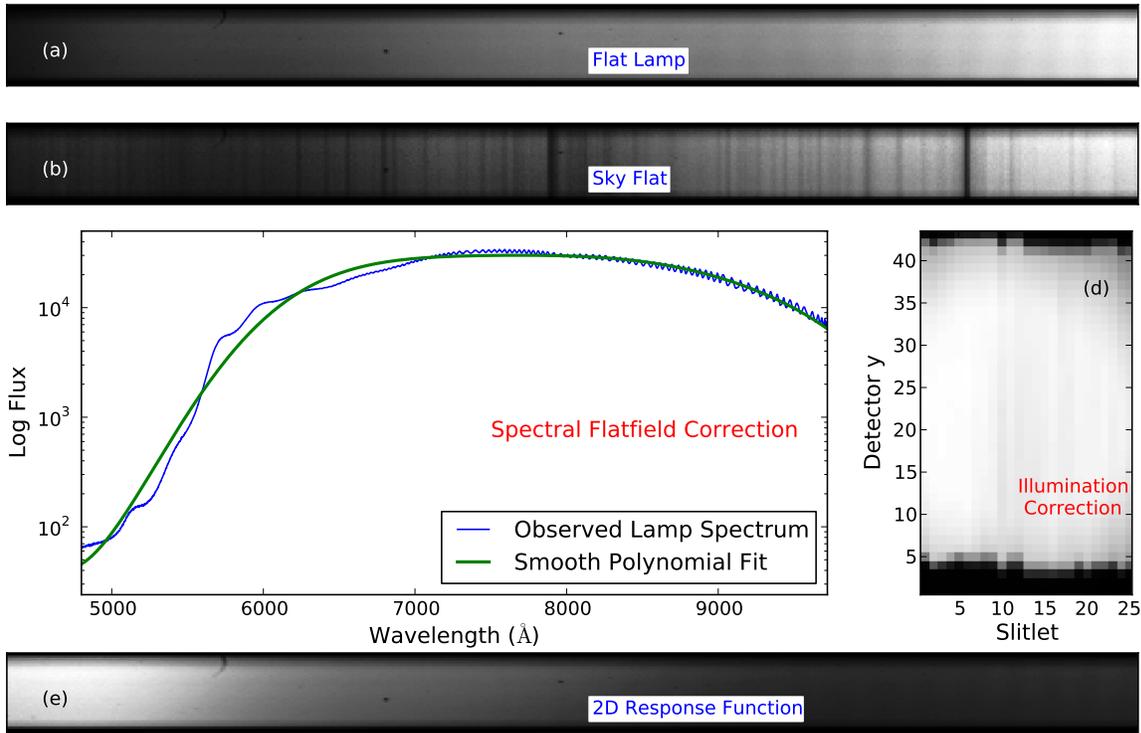
Fig. 6 (a) Flat lamp data, with "cold pixels" and fringing effects clearly visible. (b) Sky flat (twilight) data. (c) Fit of smooth function (green curve) to the flat lamp spectrum (blue curve), which corrects for cold pixels, fringing, and "dichroic wiggles". (d) Illumination correction determined from sky flat data (before correction of the vertical offsets between slitlets). (e) Final two-dimensional master response function used to flatfield data.

of the wiggles), and (3) non-uniformity in the spatial illumination across the instrument FOV.

2.3.2 Accounting for Scattered Light

After the development of the PyWiFeS flatfielding algorithm, examination of the data revealed very subtle irregularities in the spectral flat-field solution which were determined to be caused by scattered light within the instrument. A representative example of a blue master lamp flat is shown in the top panel of Figure 7 (red flat-fields are similar). The color scale has been set to reveal the horseshoe-shaped internal reflection to the right of the frame, and also reveals the diffuse light between each of the 25 slitlets visible towards the center of the frame as a red glow. Over most of the spectral range (the x dimension), the diffuse glow only amounts to 1-2% of the slitlet fluxes at any given position. To the right of the frame (corresponding to shorter wavelengths), the lamp becomes very faint. The horseshoe reflection largely dominates the signal in this region and strongly affects the flat-fielding of the data below 4500Å if left uncorrected.

To prevent contamination of the flatfield response function by this scattered light, we reconstruct the global internal reflection structures (both the horseshoe and the diffuse glow) from the regions of the detector between slitlets (the "interslit" regions), then subtract this correction from the raw master lamp and master sky flats (see Figure 7, middle and bottom panels). The detailed algorithm is a 4 step process, illustrated in Figure 8 for one slitlet, and is as follows :

- Step 1: for all 25 slitlets, extract the slitlet region and the two surrounding interslit regions; panel (a).
- Step 2: from the slitlet definitions obtained previously (see Section 2.1.1), cut out the slitlet, and smooth the interslit regions with a symmetric gaussian kernel of 10 pixels; panel (b).
- Step 3: extract the interslit count values on a finite grid with a resolution $\Delta y = 3$ and $\Delta x = 10$, and use them as input for a Bivariate Spline fitting routine (`RectBivariateSpline` in the `scipy.interp` module) to reconstruct the reflections across the slit; panel (c).
- Step 4: subtract the reconstructed contamination from the data; panel (d)

Fig. 8 Illustration of the 4 steps implemented in PyWiFeS to remove the internal reflection and the diffuse glow in WiFeS flats (slitlet number 10). (a) Raw frame (1 slitlet and 2 interslit regions). (b) Smoothed interslit regions only. (c) Smoothed interslit regions, extraction grid points and reconstructed contamination over the slitlet region. (d) Final, corrected slitlet and interlsit regions. (e) Correction intensity in percent of the original data. In each panel, the dashed lines delineate the slitlet/interlsit limits.

Fig. 7 Top: Raw master lamp flat (B3000, y-binning of 2 and 1 amplifier read-out). The internal horseshoe-shaped internal reflection is clearly visible to the right of the image. Middle: reconstructed internal reflection and diffuse glow from the interslit regions. Bottom: cleaned master lamp flat.

In panel (e) of Figure 8, we show, in percent of the original frame intensity, the amount of correction applied to the slitlet. The correction has only very little effect (∼1-2%) over most of the slitlet, where the lamp flux is strong, but is very efficient when the lamp becomes less efficient, and the internal horseshoe reflection dominates the signal.

The Bivariate Spline interpolation routine depends very little on the grid resolution (chosen in Step 3) because the grid points are distributed on either sides of a large gap. As a result, only the grid points closest from the slitlet strongly influence the reconstructed pattern across the slitlet. While it is impossible to exactly separate light traversing the primary optical path from scattered or reflected light, this routine allows removal of both the internal horseshoe reflection and the diffuse

interslit glow to within the noise level of the data as illustrated in Figure 7 (bottom panel).

2.3.3 Correcting for Anomalous "Sag" in Data

A careful inspection of reduced WiFeS data revealed that even after careful bias correction (see Section 2.1), a residual gradient in the background of every science frame was present along the x (spectral) direction. This "sag" is easily seen in the interslit regions of the bias subtracted frames. In the final reduced frames, this residual signal results in a drop in the spectrum intensity in regions where the flat lamp is faint, presumably due to impact of the sag on the flatfield solution. This effect is especially visible with the B3000 grating below 4200 Å.

The *desag* step was implemented to provide an optional way to correct this effect in the data. In practice, it is similar to the flat correction step described above, but is applied to all science frames in the night. Currently we categorise this step as optional in data reduc-

tion for two primary reason: (1) in its current form the desag step only works for "point-and-stare" observations (and not N+S); and (2) the origin of the residual intensity gradient in the observation background is still unclear. PyWiFeS users who think they require this step to reduce their data are welcome to employ it, but are also strongly encouraged to consult with the developers. Preliminary tests on the new WiFeS data taken after the detector upgrades suggest that the desag step is no longer required, and should only be considered for WiFeS users reducing observations taken prior to January 2013.

2.4 Data Cube Generation

The format of data generated by integral field spectroscopy is flux values as a function of wavelength at each spatial location on the sky. This sampling of the flux is commonly referred to as a "data cube" as it samples flux in three dimensions. This data format requires the three coordinates (x, y, λ) be derived for each pixel on the detector. While the x coordinate is simply the slitlet number, determining a consistent y coordinate for all slitlets requires a measurement of the spatial zeropoint in each slitlet. This is accomplished using a "wire" frame observation, which we describe in detail in Section 2.4.1.

The (x, y) coordinate pairs calculated within the instrument field of view do not correspond to the same physical coordinates on the sky (α, δ) at all wavelengths. This is due to the influence of atmospheric differential refraction (ADR – Filippenko 1982), where the wavelength-dependent index of refraction of the atmosphere deflects light of different wavelengths by subtly different values. Calculation of this effect is a solved astrophysical problem, and we describe the implementation in PyWiFeS in Section 2.4.2.

For most applications, the irregular spatial and wavelength sampling realised by the actual WiFeS detector pixels is too complex to easily extract physical quantities over the full instrument FOV. Instead the observed fluxes are resampled onto a regular (rectilinear) grid of (x, y, λ) for the final data cube, and this step for PyWiFeS is outlined below in Section 2.4.3.

2.4.1 Spatial Zeropoint Derivation

For each column of each slitlet of WiFeS data, the true vertical coordinate y is determined by measuring the zeropoint with a "wire frame" observation. An occulting wire coronograph that is flat in the (true) y direction is placed across the instrument FOV, and the flat lamp is shone onto the detector so that the wire's

shadow appears on the detector in each slitlet. The position of the wire is fitted from this shadow using the PyWiFeS `derive_wire_solution` routine, whose procedure is outlined graphically in Figure 9.

The wire shadow in each column is isolated by subtracting the observed flux from a smooth continuous model of the flat lamp flux, which is fitted as a low order (default is 4th order) polynomial in detector y over two regions where the lamp flux is high. The wire position is then measured as the centroid of the absorption spectrum in the region bounded by the high lamp flux regions. Once the wire centroid has been measured for each column of the detector, the wire profile is fitted as a simple linear function of detector column, with the fit restricted to columns where the lamp flux is not diminished by the dichroic. The wire solution for all columns of all slitlets is then saved in a FITS file used as input in the data cube generation routine.

2.4.2 Correction for Atmospheric Differential Refraction

Atmospheric differential refraction (ADR – Filippenko 1982) causes an excess deflection of light of bluer wavelengths. This causes the apparent position of an object to shift as a function of wavelength, and must be corrected for in integral field spectroscopy data. In PyWiFeS, this is corrected at the data cube generation step by applying a Python implementation of the ADR equations of Filippenko (1982). In Figure 10, we show a high airmass ($secz = 1.85$) WiFeS observation of a standard star (LTT2415). The trace of the star shows excellent agreement (as expected) with the predicted deflection. These deflection curves are calculated directly from the observation details stored in the WiFeS data FITS header, and applied in the data cube generation step.

2.4.3 Coordinate Rectification

Once spectral and spatial coordinates have been determined for each pixel of each slitlet, the data must then be resampled onto a rectilinear (x, y, λ) coordinate system. This is accomplished by performing a two-dimensional resampling of the data for each slitlet. The spatial zeropoint (wire profile) and vertical ADR deflection are used to assign a true y value to each pixel, and the wavelength solution defines the λ value for each pixel. The final (y, λ) grid for each slitlet is defined by the central $38''$ of the slitlet, and a wavelength array which by default spans the common wavelength coverage of all slitlets with a pixel scale equal to the mean pixel scale of all the data. Data are resampled
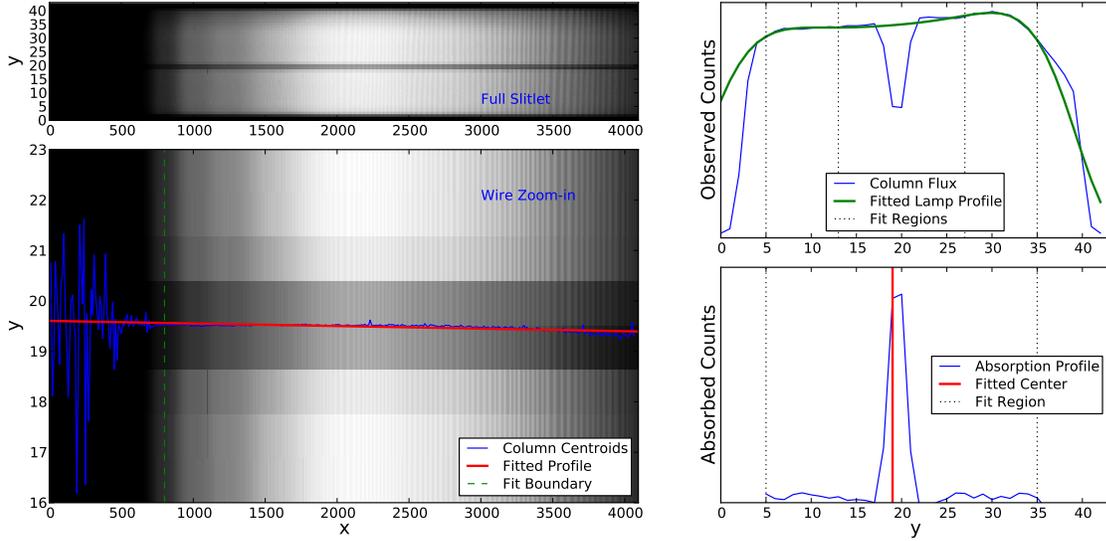
Fig. 9 Top left: Full slitlet from a wire frame observation. Top right: Column slice from a single slitlet (blue curve). The unabsorbed lamp profile is fitted as a smooth polynomial (green curve) from the regions where the lamp flux is high (dotted vertical lines). Bottom right: Wire absorption profile (blue), with fitted wire centroid (vertical red line). Bottom left: Zoom in of the wire absorption region, with fitted centroids for each column (blue curve), and fitted linear wire profile (red line) from bounded fit region (vertical dashed green line).
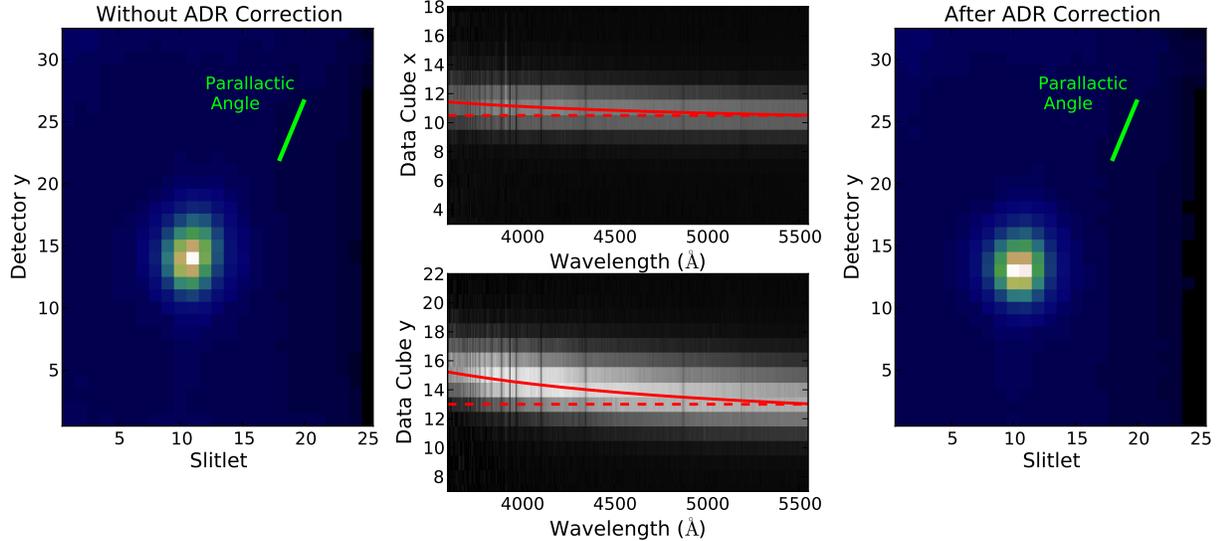
Fig. 10 Left: Processed WiFeS data cube of a high airmass ($secz = 1.85$) observation of the standard star LTT2415, with no correction applied for ADR. Elongation along the parallactic angle (shown in green) is clearly evident. Middle panels: Collapsed stack of flux vs. x and y spaxel coordinates for the uncorrected data cube. The star trace clearly follows the predictions (solid red curves) of the equations of Filippenko (1982), compared to zero deflection (dashed lines). Right: The same WiFeS data cube but processed with ADR correction. The previous elongation along the parallactic angle has been remedied.

from the observed (y, λ) coordinates for all pixels to the desired (y, λ) grid using the `griddata` routine in the `scipy.interpolate` Python package.

After this first resampling, correction of ADR deflection *across* slitlets requires a second resampling of the data (currently there is no efficient three-dimensional

resampling routine in Python, understandably due to the challenging computational requirements). This second resampling is accomplished via a simple one-dimensional resampling along all x values for a fixed (y, λ) position. By definition, flux beyond one edge of the detector will be required to properly resample the true flux along that edge. Since there is no way to produce this absent information, the data cube generation routine by default fills this unsampled edge with a copy of the observed edge data, meaning this edge will also have slightly erroneous flux. For this reason, it is still valuable for observers to align their observations along with parallactic angle to avoid this edge resampling problem.

2.5 Flux Calibration

Once WiFeS data have been corrected for pixel-wise variations of detector efficiency and spatial variations of the instrument throughput, the wavelength dependence of overall throughput of the instrument must still be corrected. This is achieved through observation of astronomical sources with known spectral energy distributions, referred to as spectrophotometric standard stars (see Bessell 2005, for a thorough review). This process of measuring the absolute sensitivity of the instrument and correcting observational data to this absolute scale is known as flux calibration, and is achieved in PyWiFeS using routines defined in the `wifes_calib` sub-module. Flux calibration is typically accomplished in three steps:

1. Standard star spectra are extracted from data cubes using the `extract_wifes_stdstar` routine.
2. The instrument sensitivity function is derived by comparing observed standard star spectra to their reference spectra using the routine `derive_wifes_calibration`.
3. The flux calibration solution is applied to uncalibrated WiFeS data cubes using the `calibrate_wifes_cube` routine.

Standard star spectra are extracted using a simple aperture extraction technique employed by the `extract_wifes_stdstar` routine. This is illustrated schematically in the left panel of Figure 11, where the sky background spectrum is estimated as the mean spectrum of spaxels outside some fiducial radius. After this sky spectrum is subtracted from each spaxel, the flux within some object aperture is summed to obtain the final object spectrum. The center of the extraction aperture is determined by fitting the flux centroid in the (x, y) plane, which is generally a robust measure of the star's location. The extraction and sky radii are set by default to $5''$ and $10''$, respectively, though these

values (and the extraction center) can be modified by passing the desired values as keyword arguments to the extraction routine.

After the standard star spectra have been extracted, the instrument sensitivity function can be derived by comparing these to their reference spectra. In PyWiFeS this is accomplished in a (mostly) automated fashion using the `derive_wifes_calibration` routine. This function takes a list of standard star cubes as input (or optionally a list of extracted spectra) and extracts all spectra according to the algorithm described above.

The `derive_wifes_calibration` routine then attempts to identify the name of the standard star from the 'OBJECT' header field of the data cube file (alternatively the user can pass a list of object names to this routine). These names are compared to a list defined in the `wifes_calib.py` file, comprising a Python dictionary containing the reference spectrum file name for each standard star. Future planned upgrades to PyWiFeS include the ability to cross-reference standard star lists by the object coordinates. Once the reference star name has been identified, its reference spectrum is loaded and the ratio of observed flux to reference flux is calculated for each star in the input list. This ratio is then corrected for the smooth atmospheric extinction using the default Siding Spring Observatory extinction curve measured by Bessell (1999).

The final flux calibration solution is fitted from the 'counts-to-flux' ratio values (i.e., the sensitivity curves) obtained for all stars in the list passed to `derive_wifes_calibration`. Points falling in regions of strong telluric absorption are not included in the fit. The final sensitivity curve can be calculated with a single star or with a list of several stars. The sensitivity curves from all stars can be normalised to one another using the 'norm_stars' keyword, which normalises all stars based on the middle 200 pixels, scaled to the maximum sensitivity curve. This allows a higher signal-to-noise measurement of the instrument sensitivity while not being hindered by grey extinction variations throughout the night.

Once the flux calibration solution has been derived, it can be applied to uncalibrated WiFeS data cubes using the `calibrate_wifes_cube` routine. This routine interpolates the derives flux calibration solution to the wavelength values of the input data cube, and multiplies the observed flux (in photon counts) to convert it to the true flux (in F_λ). This routine also applies the extinction correction, again based on the Bessell (1999) canonical Siding Spring extinction curve.

We note that the current flux calibration routine in PyWiFeS does not allow for a fit of residual extinction (i.e., deviation from the input extinction curve),
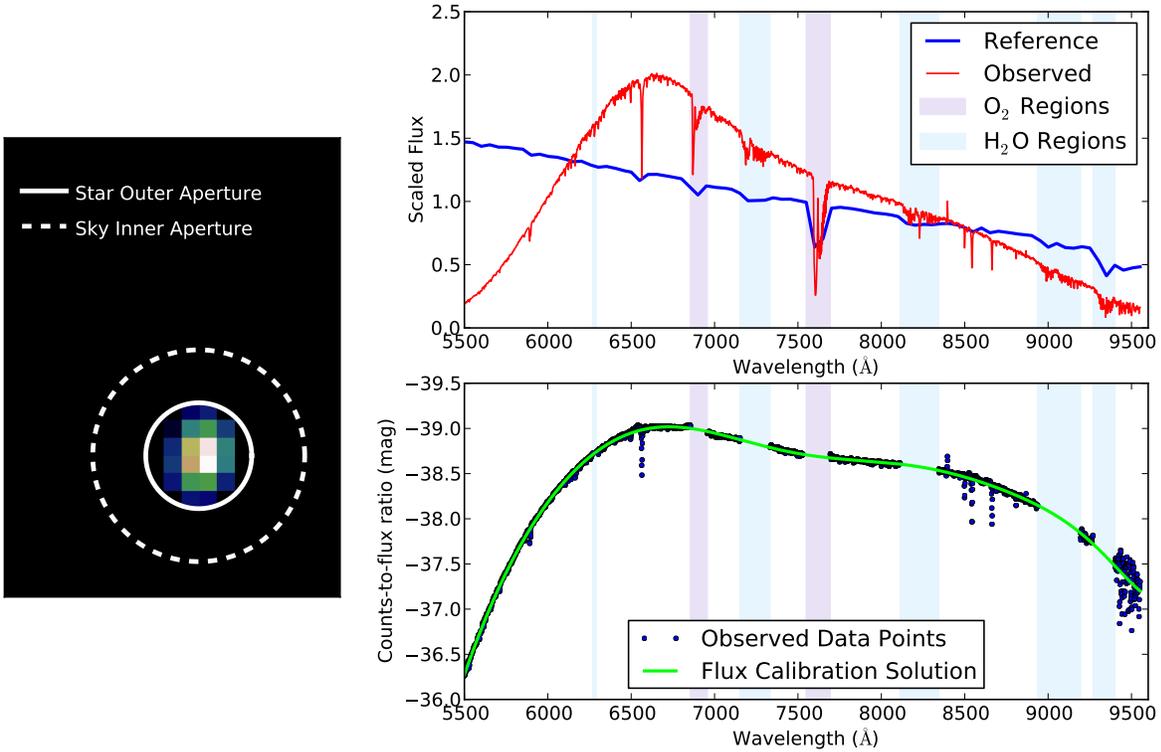
Fig. 11 Left: Collapsed data cube image for a standard star observation (LTT7379 – from Oke 1990), showing the inner boundary of the sky region (dashed line) and the star extraction aperture (solid line). Right Top: Observed flux (red curve – in photon counts, scaled by its mean flux) compared to the reference star flux (blue curve – in F_λ, again scaled by its mean flux). Right Bottom: Flux calibration solution (solid green curve) derived from flux-to-counts ratio (in magnitudes) observed in data (blue points). In the right two panels, telluric regions due to O_2 and H_2O are shown as shaded purple and blue regions, respectively.

though that capability is planned to be included in future upgrades to PyWiFeS. However, the star spectrum extraction routine is capable of saving the extracted spectrum in a format readable by the IRAF `onedspec` package, and the flux calibration routine can accept IRAF-style flux calibration solutions and extinction curves. Thus, users who require stringent flux calibration and extinction corrections can work with existing IRAF routines by employing the appropriate formatting keywords when calling the PyWiFeS extraction and flux calibration functions.

2.6 Correction of Telluric Absorption

The broad smooth absorption of light by the Earth's atmosphere is corrected using the atmospheric extinction curve in the flux calibration step in PyWiFeS. However, narrow structured absorption features caused by molecular species in the atmosphere (primarily molecular oxygen and water) persist in all object spectra after this step. These telluric features are most efficiently corrected by observing sources with smooth spectra at similar airmass and time of the night as a science target (e.g., using the smooth star division technique of Bessell 1999). Alternatively, measuring telluric absorption in a number of smooth spectrum objects over a range of airmass can provide an acceptable estimate of the average telluric absorption behavior for a given night.

In PyWiFeS, telluric features are measured using the routine `derive_wifes_telluric` and corrected using the routine `apply_wifes_telluric`. As with the flux calibration routine, the `derive_wifes_telluric` routine takes a list of star cubes (or extracted spectra) as input. Regions of the object spectra not affected by telluric absorption are fitted with a low order polynomial, as demonstrated in Figure 12. The ratio of observed flux to the predicted smooth flux defines the telluric absorption for each spectrum.

Because telluric absorption strength depends on airmass, the absorption from each each spectrum in the list is corrected to airmass $secz = 1.0$ and the mean absorption across all objects defines the final telluric absorption correction functions. The airmass dependence is a product of the saturation level of the absorp-
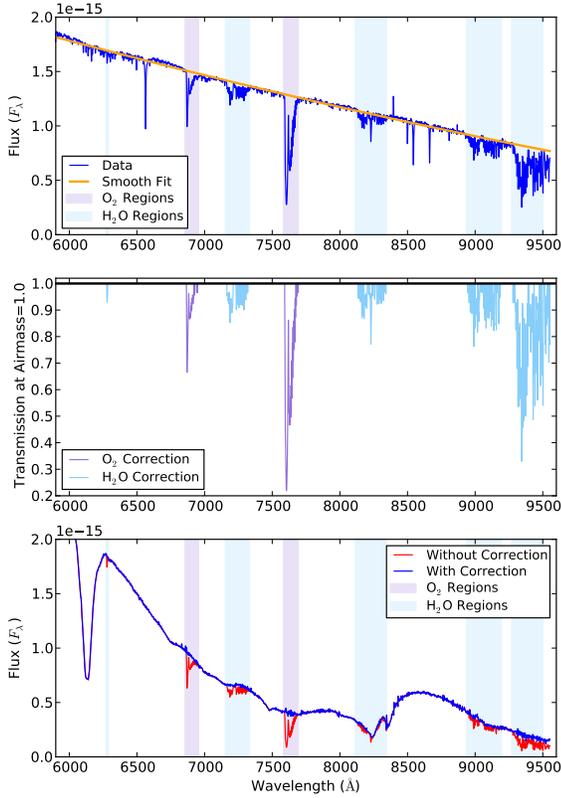
Fig. 12 Top: Flux-calibrated standard star spectrum, along with low-order polynomial fit of the smooth continuum fitted excluding regions of strong telluric absorption. Middle: Telluric correction functions for O_2 and H_2O, scaled to airmass $secz = 1.0$ (see text). Bottom: Example of a smooth spectrum source (the type Ia supernova SN 2012fr Childress et al. 2013) before (red) and after (blue) correction of telluric absorption features.

tion, which often differs between O_2 (which is more saturated) and H_2O (see, e.g., Buton et al. 2013). The logarithmic power law index for these two species can be set by the user as keyword arguments to the telluric fitting routine, and are stored along with the respective correction curves in a Python pickle (.pkl) file. From an observing night in 2013 we measured these indices to be 0.40 and 0.72 for O_2 and H_2O, respectively, which are both intermediate between the optically thick (saturated) value of 0 and the optically thin limit of 1. A more detailed study of the airmass dependence of telluric absorption features is planned for future work, but can be easily modified in practice using the keywords built into the routine.

Finally, once the telluric correction functions for O_2 and H_2O have been determined, they can be applied directly to data (with the appropriate airmass dependence) by calling the `apply_wifes_telluric` routine.

3 WiFeS Wavelength Solution

The greatest challenge in the reduction of WiFeS data is to derive the wavelength solution for the entire instrument from an arc lamp observation. In typical longslit spectroscopy reductions, the wavelength solution is derived from arc lamp data via an interactive procedure where the user finds emission lines by eye and identifies the associated reference wavelength for those lines. The fitted positions of the arc lamp emission lines and their true wavelengths are then used to extrapolate the wavelength value associated with each pixel across the entire detector. Performing line identification manually for multiplexed spectroscopic data such as that obtained with IFUs is not only time consuming, but also particularly prone to human error due to its repetitive nature.

For WiFeS data, our goal was to develop robust tools for accomplishing the same steps but without requiring the user to identify emission line positions and reference wavelengths by hand. Furthermore, we sought a physical model for the WiFeS wavelength solution which incorporates the correct analytical parametrization of how the wavelength solution varies across the detector within each slitlet, and how the wavelength solutions of all slitlets are related by the optics of the instrument.

In this Section we outline the algorithms employed in PyWiFeS to find emission lines in arc lamp data and identify their true reference wavelengths (Section 3.1), and the implementation of the WiFeS optical model in PyWiFeS (Section 3.2). The details of the instrument optical model are beyond the intended focus of this paper, so we defer such details to future work. Instead we focus on software implementation of the model in its application to real data.

3.1 Finding and Identifying Emission Lines

Discovering and fitting emission lines automatically from a large volume of data without user interaction requires a robust set of algorithms. For each row of each slitlet, we first identify all pixels which potentially have emission line flux in them with the following technique. We measure the mean and RMS of the background flux level (i.e., from pixels without emission line flux) via an iterative outlier rejection technique. Similarly, we measure the mean and RMS *derivative* of flux along the dispersion axis. Pixels must then pass a series of cuts to be identified as a potential emission line center: the flux value must be below the saturation level and above a certain flux threshold, the flux in the adjacent pixels must also exceed the flux threshold, and the derivative of adjacent pixels must exceed a threshold set from the

flux derivative statistics (by default thresholds are the mean plus 3σ for flux and flux derivative as determined from the background statistics). Any adjacent sets of pixels which pass these cuts are grouped together and the potential line center is identified as the middle pixel value of the group.

Once potential emission lines have been identified, their *exact* centers must be measured. PyWiFeS provides two algorithms which can be chosen by setting the 'find_method' keyword of the `derive_wifes_wave_solution` routine. The first, called 'loggauss', is fast but less accurate, and is intended for rapid wavelength solution fits. It fits the logarithm of the emission line profile as a quadratic function (equivalent to a Gaussian in linear flux) but naturally must excise the pixels with zero flux from the fit. The second technique, called 'mpfit', is somewhat slower but extremely accurate, and should be considered the standard for science-grade data reduction. This method performs a Gaussian fit to the emission line flux profile using the `mpfit` Python package (a Python port of the IDL implementation of the non-linear least squares fitting program MINPACK Markwardt 2009). Though this procedure takes up to thirty minutes on a standard laptop, it is designed to naturally take advantage of parallel processing capabilities and generally takes about two minutes on the Linux clusters at RSAA. A visual comparison of these two methods is presented in Figure 13, which illustrates the sub-pixel errors in line centres found using the 'loggauss' method as compared to the more robust 'mpfit' method.

Once a series of emission lines have been identified in the data, the true wavelength of the atomic transition producing that line must be identified. Classically, this process is done interactively where an observer must visually inspect the arc lamp spectrum and compare it to a reference spectrum where lines have been identified. After identifying a few lines in the center of the longslit, a low order polynomial fit to the dispersion relation (i.e., wavelength versus pixel number) is used to automatically identify the remaining lines, and these lines are cross-identified in all remaining rows of the data. This visual inspection and line identification procedure was deemed to be cumbersome, especially given the excellent wavelength stability of the WiFeS instrument (thermal drift of less than a few angstroms).

In practice, emission lines in WiFeS arc lamp spectra are associated with a reference wavelength in a multi-step process. First a baseline guess for the wavelength at each found emission line position (in x, y on the detector) is calculated from a known reference wavelength solution. A coherent shift of the wavelength guess for each row is applied by performing a cross-correlation
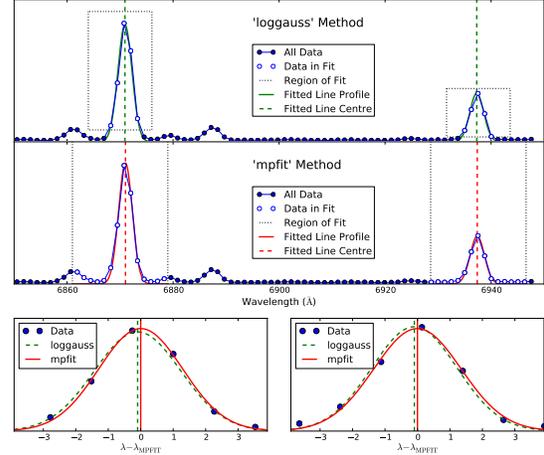
Fig. 13 Top: Fits to several arc lamp emission lines using the 'loggauss' method. Data points near a line peak with flux values exceeding 10% of the peak flux (outlined by black dotted boxes) are fitted as a Gaussian in logarithmic flux space. Middle: Fits to the same lines using the 'mpfit' method, which utilizes *all* nearby data points including those with near-zero flux. Bottom Panels: Comparison of the line profiles fitted with the two methods ('loggauss' – dashed green; 'mpfit' – solid red), alone with their fitted centres (vertical lines). Data points are shown as blue circles.

of the spectrum for that row with a reference spectrum (this is optional but strongly recommended). Next, a comparison of the found line wavelengths to reference line wavelengths is performed (for each detector row), so that each found line is assigned a reference if and only if that reference wavelength is the closest wavelength in the reference list *and* the found wavelength is the closest found line to that reference line (i.e., a nearest neighbors requirement from both the found and reference lists). A baseline cut is applied at this stage so that associated line pairs with highly discrepant wavelength values (default is 5Å) are rejected as false associations.

Finally, once lines have been reliably measured and associated with a reference wavelength, the fitting of the wavelength solution proceeds. The technique employed in PyWiFeS is built on a global model of the instrument optics, which we describe in the following Section.

3.2 Implementing the WiFeS Optical Model

The manner in which light of different wavelengths is dispersed within the WiFeS instrument can be described analytically using the geometric optics of VPH gratings (Barden et al. 2000). Light from all spatial positions in the WiFeS FOV passes through the same VPH grating, but light from each spaxel has a different

angle of incidence and physical point of entry due to the slit geometry from the image slicer (Fig 14a). The slicer separates the slitlets by reflecting each spatial slice with a differential in the vertical plane spreading the light out into a pseudo-slit. Each slitlet is also offset in the horizontal direction by widths of the neighboring slits producing the staircase slit geometry. Slitlets are further offset in the vertical direction by their own length to allow the necessary storage space on the CCD for nod-and-shuffle observations.
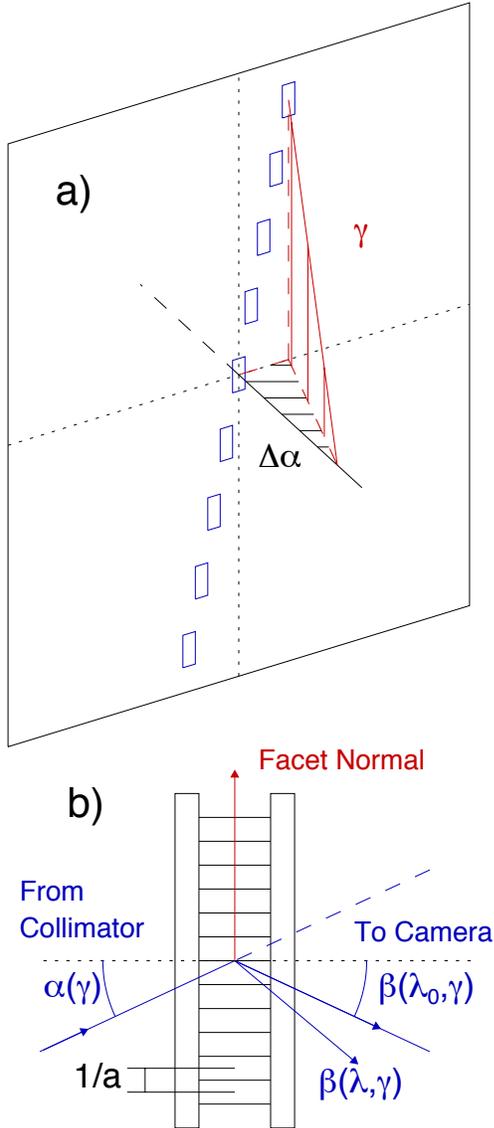
Fig. 14 Top: Schematic diagram of grating geometry, defining pertinent parameters of the grating equation. Bottom: Geometric distribution of slitlet images as passed through VPH gratings (after being split on the image slicer and reflected by other optics).

The light from each slitlet is dispersed following the grating equation:

$$m\lambda = a \cdot [\sin(\alpha) + \sin(\beta)] \cdot \cos(\gamma) \qquad (1)$$

with a grating line density a lines per mm. The off-axis angle, γ, introduces additional optical path difference for off-axis slits, resulting in significant spectral curvature towards the edge of the CCD. The staircase slit geometry also results in a slit-number dependence for angle of incidence α introducing a shift to the diffraction pattern with slit-number in addition to the classic spectral curvature.

The WiFeS optical model is described by a total of 43 parameters. The 18 primary parameters include values such as the grating line density and blaze wavelength, the primary angle of incidence of the optical beam, the centre of the beam on the detector (in x and y), radial distortion terms, focal length of the camera, tilt of the detector (in x and y), as well as other geometric terms. In addition to these, each of the 25 slitlets has some very small deviation (which we label $\Delta\alpha_i$) from uniform angular separation α between adjacent slitlets, caused by limits in the manufacturing precision of the image slicer. Though incredibly small ($\lesssim 10^{-5}$ radians), these offsets produce a coherent shift in wavelength of order 0.05Å for each slitlet.

The 43 WiFeS optical model parameters are fitted by minimising the residual offset between the optical model wavelengths and reference wavelengths for the successfully found emission lines. In practice this is done in a series of steps which each fit for a specific group of optical model parameters (e.g., the 25 $\Delta\alpha_i$ values are fitted together as the final group of parameters). Residuals from this optical model method are typically of order 0.05Å for the R=7000 gratings, and 0.1Å for the R=3000 gratings. Figure 15 shows the final diagnostic plots for the optical model wavelength solution fit for a R7000 grating Ne-Ar arc lamp exposure. Structure in the residuals is largely removed, with a final residual dispersion of 0.053Å.

In principle, many of the parameters in the WiFeS optical model should be constant over time (the $\Delta\alpha_i$ values, the beam centre and detector tilt, detector focal length, etc.). Analysis of large volumes of WiFeS data is currently underway to determine the stability of these parameters, and future versions of PyWiFeS will likely reflect knowledge of the static optical model parameters.

3.3 Stability of the Wavelength Solution

WiFeS is located on the Nasmyth focus of the 2.3m ANU telescope. This design ensures that the instrument is subject to a constant gravity vector, and makes
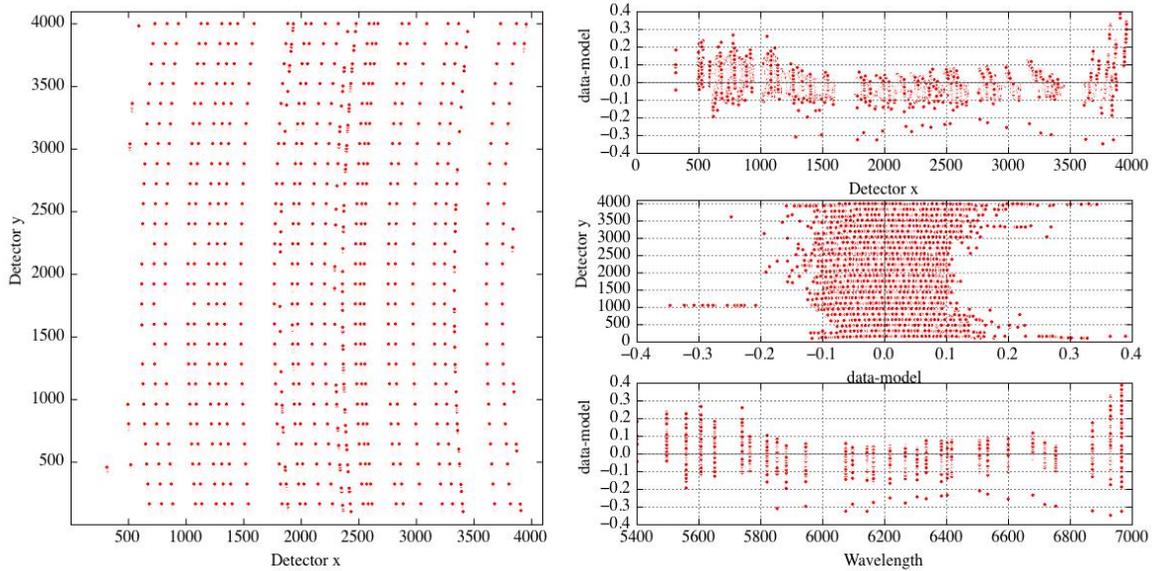
Fig. 15 Output from the PyWiFeS optical model wavelength solution for WiFeS, using a Ne-Ar arc lamp exposure with the R7000 grating. Left: All 19539 found arc lines in the detector plane, illustrating the curved shape of the WiFeS wavelength solution. Right Residuals of fitted emission line centres from prediction for the best fit optical model prediction, as a function of detector x (top), detector y (middle), and wavelength (bottom).

it a very stable spectrograph overall. WiFeS can however be subject to temperatures changes of the order of ∼5-15 degrees during an observing night. When subject to a change in temperature, the WiFeS gratings will expand or contract, increasing or decreasing their resolution. To account for and correct this effect, observers are required to acquire arc frames during the night to monitor the changes of the wavelength solution. In Figure 16, we illustrate the effect of temperature drift on the wavelength solution.

In the top two panels, we show the spectra of a Ne-Ar arc lamp observed with the B3000 and R7000 gratings. We reduce the arc lamp exposure similarly to a science frame using PyWiFeS, and plot the 25 spectra extracted along the middle slice (along the y direction) of the final data cube. The arc frame has been acquired at the end of a winter observing night on 16th August 2012. The wavelength solution is derived from a similar arc exposure acquired early in the preceding afternoon. The reference wavelengths for the different arc lines are marked with vertical dashed lines. As expected, using an arc far-away from a given observation results in a bad wavelength calibration of the data set, with errors of the order of 1-2Å(typically 40-60 km/s). We note that the error results in a blueshift of the spectra for the blue frame, and a redshift of the spectra for the red frame. In the bottom plot, we show the same Ne-Ar spectra, but using the frame itself to derive the wavelength solution. As expected, the arc lines are in this case in perfect agreement with their reference

wavelengths. We urge WiFeS observers to keep this point in mind when planning their observing run, and to take arc frames at regular time interval to monitor the temperature changes of the grating. Even if one's science goals do not require an absolute wavelength calibration, regular arc frames are still required to obtain a wavelength solution consistent between the red and blue frames. Our experience suggests that an interval of ∼1 hour between arcs is a reasonable choice, but we leave it to every observer to decide which cadence is most appropriate to their science goals.

4 Data Reduction Scripting and Metadata

PyWiFeS was originally designed to be a generic toolkit for WiFeS data reduction which could be scripted to match the user's desired reduction procedure. Toward this end, we designed a metadata structure and reduction script format which has become the default data reduction script included in the PyWiFeS package. The general principle underlying our preferred data reduction scheme is that the user provides the context for the data they wish to process in a generic way which can be stored in the metadata structure. That abstract structure can then be used as input to the data reduction script, which interprets the metadata structure in a completely generic fashion. This allows new data sets to be interchangeable to a degree that reduction
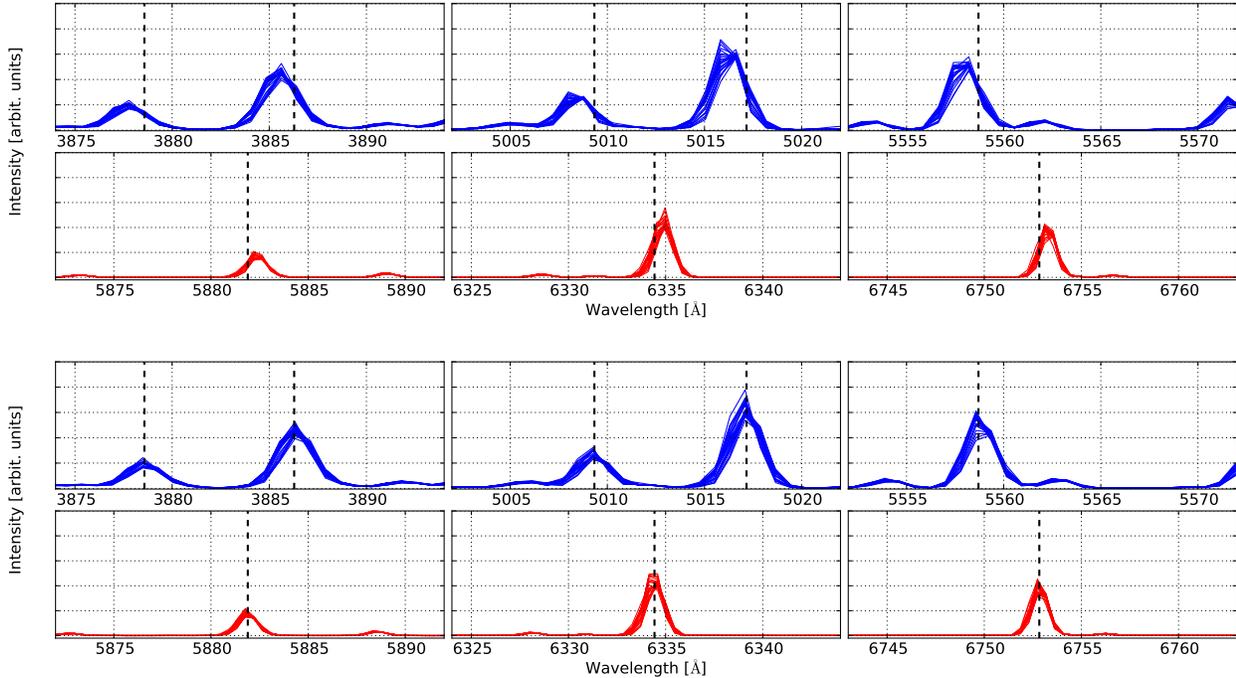
Fig. 16 Top: B3000 and R7000 Ne-Ar arc lamp spectra observed at the end of the night on 16th August 2012, and reduced using an Ne-Ar arc lamp frame acquired in the afternoon ($\sim$8 hours earlier). The vertical dashed lines denote the reference position of the arc lines. Each panel spans 20 Å. Bottom: idem, but with the wavelength solution derived from the late-night Ne-Ar arc itself. The shifts visible in the top plots highlight the importance of taking arc frames at regular interval during an observing run to account for the temperature variations inside WiFeS.

procedures can be repeated on different data sets without requiring the user to manually specify names of data files at each step of the reduction. In this Section we first describe the default metadata structure, then briefly outline the operation of the default reduction script. A detailed walkthrough of how to generate the metadata and run the reduction script is provided on the PyWiFeS Wiki.

4.1 Default PyWiFeS Metadata Structure

High-level information characterizing important properties of data are typically referred to as *metadata*. For spectroscopic observations such as those collected with WiFeS, examples of metadata include the observation type, the target being observed, or whether multiple exposures should be associated with one another. This information is then used to decide which data reduction procedures should be used, and is perhaps most commonly handled manually. Because spectroscopic data reduction often follows a repeated pattern, we desired an abstract metadata structure which could be assigned to a particular set of data then handled generically by a standard data reduction script.

In practice, the PyWiFeS metadata structure for a given set of data is stored in a Python dictionary which is saved in a Python pickle file. This dictionary is defined in the scripts `save_blue_metadata.py` and `save_red_metadata.py` (the metadata structures must be defined separately for the red and blue cameras), which we will call the "metadata saving" scripts. These are included in the PyWiFeS distribution and each script simply defines the metadata dictionary and saves it to a pickle file.

The metadata dictionary has several keywords which store information about the various types of calibration images, as well as science observation groups, and observations of standard stars. Generic instrumental calibration files – such as flat field observations, arc lamp exposures, and bias frames – which can be applied to calibrating all science and standard star data are stored in simple lists. These files are considered to be the "master" calibration files for a given data reduction run, and are applied in the calibration of all files unless specific calibration files are specified for a given observation.

Science and standard star observations are each stored as a list of Python dictionaries. Each dictionary describes an observation group, including all science frames which are ultimately co-added, as well as associated calibration frames ("local" calibrations) taken explicitly for use in calibrating that particular science

target (e.g., a bias frame and arc frame taken immediately after observing a science target). Each keyword of the dictionary is a list, with the 'sci' keyword containing a list of all observations to be co-added as the final science frame, the 'sky' keyword containing a list (typically of length 1) of sky frames to be subtracted during data reduction, 'arc' containing the "local" arc frame, and so on. Standard star dictionaries also have the special keyword 'type' whose value is a list consisting of one or both of the strings 'flux' and 'telluric', used to distinguished if that star is used in the flux calibration and/or telluric correction procedures.

For example, if the only science observation is a pair of science frames followed by a sky frame, then a bias and an arc frame, the pertinent part of red metadata saving script would appear as this:

```
sci_obs = [
    {'sci'  : ['r0001', 'r0002'],
     'sky'  : ['r0003'],
     'bias' : ['r0004'],
     'arc'  : ['r0005'],
     },
]
```

This illustrates the naming convention employed in the metadata structure, where the observation root name is used to identify the data so that 'r0001' refers to the observation whose raw data is in the file "r0001.fits". At the end of the metadata script, the list 'sci_obs' is saved as the 'sci' keyword of the metadata dictionary. A summary schematic diagram of the PyWiFeS metadata structure is shown in Figure 17.
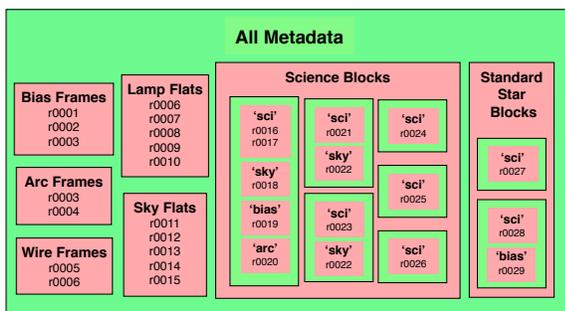
Fig. 17 Schematic diagram of the default PyWiFeS metadata structure. Python dictionaries are green, and Python lists are red. Names or descriptions of the lists or dictionaries are in bold face, while place face text denotes example file names. Master calibration files are stored in Python lists within the metadata dictionary, which the science and standard star observations are lists of dictionaries. Science observation block dictionaries must at minimum have the 'sci' keyword defined, but may also have calibration and sky frames associated to them as appropriately named lists in the observation block dictionaries.

The metadata structure is saved to a Python pickle file by running the `save_red_metadata.py` script from the command line, and this pickle file is passed as input to the data reduction script. The metadata saving script is a critical part of the data reduction scheme, and should be carefully prepared by the user before each reduction.

For convenience, we include with PyWiFeS a Python script which inspects all FITS files in a directory and attempts to automatically generate the metadata saving scripts (for blue and red). We caution that this script is very simplistic, and identifies calibrations by the 'IMAGETYP' header keyword set by the WiFeS controller. All calibration files are by default assigned to the "master" calibration lists, and any calibrations which the user took for specific science frames should be assigned by the user to those frames as in the example provided above. The PyWiFeS metadata structure is designed to be as abstract as possible to enable ease of data reduction, but requires careful construction of the metadata by the user in order to ensure data reduction is performed correctly.

4.2 The PyWiFeS Data Reduction Script

The PyWiFeS data reduction scripts `reduce_blue_data.py` and `reduce_red_data.py` take a PyWiFeS metadata file as input and perform the data reduction procedures outlined by the user. The "standard" version of these scripts are included in the PyWiFeS distribution, but can be edited to add or modify data reduction steps as desired. In this Section we describe the default structure of the reduction script.

The data reduction procedure performed by the reduction script is defined in a Python list `proc_steps`, which is a list of the data reduction steps. Each step is defined by a Python dictionary which has four keywords:

- `step` – This is the main label for each reduction step and identifies which function (described below) should be called. For example, the first step "overscan_sub" calls the function `run_overscan_sub`, which in turn calls the appropriate PyWiFeS functions to operate on the data.
- `run` – A simple Python Boolean (`True` or `False`) which sets whether that step in the reduction is run. This keyword is useful when individual steps of the reduction need to be re-run, e.g., if the wavelength solution does not converge to the user's satisfaction. It is important to note that a step which has been set to 'run:False' will be interpreted by the reduction script as having been successfully executed, and the following step will look for the output of the

"switched-off" step. If users wish to completely remove a step from the reduction chain, it should either be commented out or deleted.

- **suffix** – This defines how the output of the reduction step will be labeled. For example, an image whose root name is 'r0001' (as defined in the metadata file) and is processed through the first reduction step with **'suffix':'01'** would have output "r0001.p01.fits" where the 'p' indicates the file has been through the processing step labeled by '01'. The reduction script knows to look for the suffix of the previous step to define the input image to that step. Following the same example, 'r0001' passed to the next step with **'suffix':'02'** would result in that step looking for "r0001.p01.fits", processing it, and saving the output in "r0001.p02.fits".

- **args** – This keyword is set as a dictionary of arguments which should be passed to the main function being called in that data reduction step. For example, in the "cube_gen" step which generates final data cubes, one can set the minimum wavelength of the cubes by setting the **args** keyword as **'args': \{'wmin_set':5400.0\}**. This will pass the argument on so that every time **generate_wifes_cube** is called in this step, it will be called with the keyword argument **wmin_set=5400.0**. This functionality allows the user to adjust the output of the PyWiFeS functions easily from the processing step list.

Each reduction step calls a "master reduction function" (such as **run_overscan_sub** in the example above) which is defined in the reduction script (which can be modified by the user). These functions *must* be defined to take as input the metadata structure (the Python dictionary described in Section 4.1), the suffix from the previous reduction step (**prev_suffix**), the suffix to be used in output from the current step (**curr_suffix**), and finally the additional arguments to be passed the pertinent PyWiFeS functions. A loop at the end of the reduction script automatically parses the processing step list (**proc_steps**) and calls the master reduction functions with the correct suffixes.

Each master reduction function parses the metadata in a specific way and calls the pertinent functions from PyWiFeS to perform that reduction step. For example, the **run_overscan_sub** function calls **pywifes.subtract_overscan** on all raw frames, **run_obs_coadd** calls **pywifes.imarith_mef** on all groups of science and standard star frames, and **run_extract_stars** calls **pywifes.extract_wifes_stdstar** on all standard star data cubes.

Some steps in the processing list do not perform operations on all data, but instead prepare master calibration files such as the flatfield response function or

the wavelength solution. The master reduction function which generates these master calibration files also checks to see if any science observation groups require the use of "local" calibration files, and derives the appropriate calibrations from those files as well. The names of the master files are defined early in the reduction script (and again can be changed as desired by the user) and are global variables which are referenced when calibrations must be applied in the processing steps, but use of local calibrations is also handled automatically. For example, in the data cube generation step, science frames are transformed into data cubes using the master wire and master wavelength solution by default, but using the local wavelength solution for observation groups where a "local arc" was defined.

The full list of the default processing steps is:

Step 00 Overscan subtraction (must happen before *any* other step).

Step 01 Bad pixel repair.

(CAL) Creation of "superbias", and bias model from superbias and any local biases.

Step 02 Bias frame subtraction.

(CAL) Creation of "superflats", both lamp and sky.

(CAL) Measurement of the slitlet profiles.

(CAL) Desag (optional)

(CAL) Flat frame cleanup.

(CAL) MEF flat creation, both lamp and sky.

Step 03 MEF creation for all frames.

(CAL) Derivation of the wavelength solution.

(CAL) Derivation of the wire (spatial) solution.

(CAL) Derivation of the flatfield response function (requires wavelength solution).

Step 04 Cosmic ray rejection.

Step 05 Subtraction of sky frames.

Step 06 Co-adding of frames in observation groups.

Step 07 Application of flatfield response function.

Step 08 Generation of data cubes.

(CAL) Extraction of uncalibrated standard star spectra.

(CAL) Derivation of flux calibration solution from standard star spectra.

Step 09 Flux calibration of all data cubes.

(CAL) Extraction of calibrated, but not telluric-corrected, standard star spectra.

(CAL) Derivation of telluric correction from standard star spectra.

Step 10 Correction of telluric absorption.

Step 11 Reformatting into 3D data cube format.

Steps which perform processing operations on all science frames (and other frames requiring that reduction step) are labeled by number in the above list, while steps labeled as (CAL) denote processing steps which prepare or operate on master calibration files.

We also note that nod-and-shuffle ("N+S") frames are handled automatically in the PyWiFeS reduction scripts. Special treatment of N+S data needs to occur when the MEF files are first created, when subtraction of sky frames occurs, and any intermediate steps between. In the default scheme described above, MEF files are created at the 'p03' step, and N+S frames are automatically processed with the N+S version of the slitlet separation function. For example, if the raw WiFeS image was "r0001.fits", the N+S version of this step would save the object pixels as the MEF file "r0001.p03.fits" and the nodded sky pixels are saved as the MEF file "r0001.s03.fits". By default, these files are both processed in the cosmic ray step to produce "r0001.p04.fits" and "r0001.s04.fits" files, then in the sky subtraction step the nodded sky pixels in "r0001.s04.fits" are subtracted from the object pixels in "r0001.p04.fits" to produce the sky-subtracted "r0001.p05.fits". All of these steps are executed automatically without the user needing to identify N+S frame, as these are identified from the 'WIFESOBS' FITS header keyword.

Finally, we note that all PyWiFeS data processing functions are designed to work seamlessly with WiFeS data taken with detector binning, or in half-frame readout mode. However, the data reduction script and metadata creation script are not currently designed to handle multiple observing modes at the same time (though that functionality is planned for future upgrades). For example, attempting to subtract a full-frame bias from a half-frame image will cause the reduction to fail. We caution that these failures modes are not Python exceptions built specifically into PyWiFeS, but rather are the results of normal Python errors being raised (e.g., due to incompatibility of data size). This means that it is possible to perform reduction steps which are wrong but do not raise Python errors, such as subtracting a full-frame bias with binning of 2 from a half-frame image with binning of 1, or flux calibrating R7000 data with a calibration solution derived for the R3000 grating. Again, users should be careful in compiling their metadata so that only compatible data are being sent to the data reduction script.

5 Conclusion

We presented PyWiFeS, a new rapid data reduction pipeline for the WiFeS instrument. PyWiFeS consists of a series of data processing functions written in Python which utilise existing functionality from the standard Python modules NumPy, SciPy, and PyFITS. PyWiFeS does not make any function calls outside of Python,

and thus requires no special system-dependent software installation nor does it require other astronomical data processing environments such as IRAF.

Data processing functions in PyWiFeS are designed to operate on FITS format data as both input and output. This means data processing naturally occurs as a modular process, and the order of data reduction steps can be easily adjusted by constructing Python scripts. More importantly, the scriptable nature of PyWiFeS enables repeat processing to accommodate upgrades to the reduction code or to process new data in the same way as previous data sets. PyWiFeS data reduction functions were also carefully constructed to define instrument characteristics with abstract variables wherever possible so that the code could potentially be adapted for use in reducing data from other image-slicing integral field spectrographs.

Beyond the functional data reduction procedures defined in the core PyWiFeS modules, the reduction pipeline also includes a default metadata format and reduction script favored by the developers. The main goal of this structure is to decouple the assignment of calibration or processing roles to specific files from the data reduction procedures performed once those roles are assigned. Specifically, if an observer uses arc, flat, and bias frames in the same data reduction roles for every night of observing, the observer should not have to repeat or edit calls to data reduction procedures simply to reflect the different file names from different nights. Our default PyWiFeS metadata structure accomplishes this goal by storing the observation metadata in an abstractly labeled Python dictionary. For example, all bias frames for a night are stored in a Python list within the metadata dictionary, and the data reduction script operates on bias frames when needed by iterating through members of that list (without requiring knowledge of the exact file name). Metadata for science observations are stored as a list of Python dictionaries containing lists of the science frames themselves and any associated calibration frames, so that calibration data needing to be applied to a specific frame can be easily tracked.

While this paper outlines the official release version of PyWiFeS, future improvements to the pipeline will be implemented as needed and released from the RSAA PyWiFeS Wiki. The current version of the pipeline already accounts for all standard integral field data reduction procedures, as well as several sophisticated algorithms designed to account for realistic behavior of the WiFeS instrument. These include the scattered light flatfield correction procedures, which result in excellent spatial flatfielding, and the global wavelength solution based on the instrument optics. Future data reduction

improvements based on better understanding of the instrument behavior will be incorporated when identified, and improvements in data reduction scripting will be released with future developments as well.

References

Bakos, G.Á., Csubry, Z., Penev, K., Bayliss, D., Jordán, A., Afonso, C., Hartman, J.D., Henning, T., Kovács, G., Noyes, R.W., Béky, B., Suc, V., Csák, B., Rabus, M., Lázár, J., Papp, I., Sári, P., Conroy, P., Zhou, G., Sackett, P.D., Schmidt, B., Mancini, L., Sasselov, D.D., Ueltzhoeffer, K.: Publ. Astron. Soc. Pac. **125**, 154 (2013). 1206.1391. doi:10.1086/669529

Barden, S.C., Arns, J.A., Colburn, W.S., Williams, J.B.: Publ. Astron. Soc. Pac. **112**, 809 (2000). doi:10.1086/316576

Bayliss, D., Zhou, G., Penev, K., Bakos, G., Hartman, J., Jordán, A., Mancini, L., Mohler, M., Suc, V., Rabus, M., Béky, B., Csubry, Z., Buchhave, L., Henning, T., Nikolov, N., Csák, B., Brahm, R., Espinoza, N., Noyes, R., Schmidt, B., Conroy, P., Wright, D., Tinney, C., Addison, B., Sackett, P., Sasselov, D., Lázár, J., Papp, I., Sári, P.: ArXiv e-prints (2013). 1306.0624

Bessell, M.S.: Publ. Astron. Soc. Pac. **111**, 1426 (1999). doi:10.1086/316454

Bessell, M.S.: Annu. Rev. Astron. Astrophys. **43**, 293 (2005). doi:10.1146/annurev.astro.41.082801.100251

Buton, C., Copin, Y., Aldering, G., Antilogus, P., Aragon, C., Bailey, S., Baltay, C., Bongard, S., Canto, A., Cellier-Holzem, F., Childress, M., Chotard, N., Fakhouri, H.K., Gangler, E., Guy, J., Hsiao, E.Y., Kerschhaggl, M., Kowalski, M., Loken, S., Nugent, P., Paech, K., Pain, R., Pécontal, E., Pereira, R., Perlmutter, S., Rabinowitz, D., Rigault, M., Runge, K., Scalzo, R., Smadja, G., Tao, C., Thomas, R.C., Weaver, B.A., Wu, C., Nearby SuperNova Factory: Astron. Astrophys. **549**, 8 (2013). 1210.2619. doi:10.1051/0004-6361/201219834

Childress, M.J., Scalzo, R.A., Sim, S.A., Tucker, B.E., Yuan, F., Schmidt, B.P., Cenko, S.B., Silverman, J.M., Contreras, C., Hsiao, E.Y., Phillips, M., Morrell, N., Jha, S.W., McCully, C., Filippenko, A.V., Anderson, J.P., Benetti, S., Bufano, F., de Jaeger, T., Forster, F., Gal-Yam, A., Le Guillou, L., Maguire, K., Maund, J., Mazzali, P.A., Pignata, G., Smartt, S., Spyromilio, J., Sullivan, M., Taddia, F., Valenti, S., Bayliss, D.D.R., Bessell, M., Blanc, G.A., Carson, D.J., Clubb, K.I., de Burgh-Day, C., Desjardins, T.D., Fang, J.J., Fox, O.D., Gates, E.L., Ho, I.-T., Keller, S., Kelly, P.L., Lidman, C., Loaring, N.S., Mould, J.R., Owers, M., Ozbilgen, S., Pei, L., Pickering, T., Pracy, M.B., Rich, J.A., Schaefer, B.E., Scott, N., Stritzinger, M., Vogt, F.P.A., Zhou, G.: Astrophys. J. **770**, 29 (2013). 1302.2926. doi:10.1088/0004-637X/770/1/29

Dopita, M., Hart, J., McGregor, P., Oates, P., Bloxham, G., Jones, D.: Astrophys. Space Sci. **310**, 255 (2007). 0705.0287. doi:10.1007/s10509-007-9510-z

Dopita, M., Rhee, J., Farage, C., McGregor, P., Bloxham, G., Green, A., Roberts, B., Neilson, J., Wilson, G., Young, P., Firth, P., Busarello, G., Merluzzi, P.: Astrophys. Space Sci. **327**, 245 (2010). 1002.4472. doi:10.1007/s10509-010-0335-9

Filippenko, A.V.: Publ. Astron. Soc. Pac. **94**, 715 (1982). doi:10.1086/131052

Fraser, M., Inserra, C., Jerkstrand, A., Kotak, R., Pignata, G., Benetti, S., Botticella, M.-T., Bufano, F., Childress, M., Mattila, S., Pastorello, A., Smartt, S.J., Turatto, M., Yuan, F., Anderson, J.P., Bayliss, D.D.R., Bauer, F.E., Chen, T.-W., Förster Burón, F., Gal-Yam, A., Haislip, J.B., Knapic, C., Le Guillou, L., Marchi, S., Mazzali, P., Molinaro, M., Moore, J.P., Reichart, D., Smareglia, R., Smith, K.W., Sternberg, A., Sullivan, M., Takáts, K., Tucker, B.E., Valenti, S., Yaron, O., Young, D.R., Zhou, G.: Mon. Not. R. Astron. Soc. **433**, 1312 (2013). 1303.3453. doi:10.1093/mnras/stt813

Green, A.W., Glazebrook, K., McGregor, P.J., Abraham, R.G., Poole, G.B., Damjanov, I., McCarthy, P.J., Colless, M., Sharp, R.G.: Nature **467**, 684 (2010). 1010.1262. doi:10.1038/nature09452

Inserra, C., Smartt, S.J., Scalzo, R., Fraser, M., Pastorello, A., Childress, M., Pignata, G., Jerkstrand, A., Kotak, R., Benetti, S., Della Valle, M., Gal-Yam, A., Mazzali, P., Smith, K., Sullivan, M., Valenti, S., Yaron, O., Young, D.: ArXiv e-prints (2013). 1307.1791

Markwardt, C.B.: In: Bohlender, D.A., Durand, D., Dowler, P. (eds.) Astronomical Data Analysis Software and Systems XVIII. Astronomical Society of the Pacific Conference Series, vol. 411, p. 251 (2009). 0902.2850

Maund, J.R., Fraser, M., Smartt, S.J., Botticella, M.T., Barbarino, C., Childress, M., Gal-Yam, A., Inserra, C., Pignata, G., Reichart, D., Schmidt, B., Sollerman, J., Taddia, F., Tomasella, L., Valenti, S., Yaron, O.: Mon. Not. R. Astron. Soc. **431**, 102 (2013). 1302.0170. doi:10.1093/mnrasl/slt017

McGregor, P.J., Hart, J., Conroy, P.G., Pfitzner, M.L., Bloxham, G.J., Jones, D.J., Downing, M.D., Dawson, M., Young, P., Jarnyk, M., Van Harmelen, J.: In: Iye, M., Moorwood, A.F.M. (eds.) Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, vol. 4841, p. 1581 (2003). doi:10.1117/12.459448

McGregor, P.J., Bloxham, G.J., Boz, R., Davies, J., Doolan, M., Ellis, M., Hart, J., Jones, D.J., Luvaul, L., Nielsen, J., Parcell, S., Sharp, R., Stevanovic, D., Young, P.J.: In: Ground-based and Airborne Instrumentation for Astronomy IV. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, vol. 8446, September 2012. doi:10.1117/12.925259

Oke, J.B.: Astron. J. **99**, 1621 (1990). doi:10.1086/115444

Press, W.H., Teukolsky, S.A., Vetterling, W.T., Flannery, B.P.: Numerical Recipes 3rd Edition: The Art of Scientific Computing, 3rd edn. Cambridge University Press, ??? (2007)

Rich, J.A., Dopita, M.A., Kewley, L.J., Rupke, D.S.N.: Astrophys. J. **721**, 505 (2010). 1007.3495. doi:10.1088/0004-637X/721/1/505

Savitzky, A., Golay, M.J.E.: Analytical Chemistry **36**, 1627 (1964)

Spergel, D., Gehrels, N., Breckinridge, J., Donahue, M., Dressler, A., Gaudi, B.S., Greene, T., Guyon, O., Hirata, C., Kalirai, J., Kasdin, N.J., Moos, W., Perlmutter, S., Postman, M., Rauscher, B., Rhodes, J., Wang, Y., Weinberg, D., Centrella, J., Traub, W., Baltay, C., Colbert, J., Bennett, D., Kiessling, A., Macintosh, B., Merten,

J., Mortonson, M., Penny, M., Rozo, E., Savransky, D., Stapelfeldt, K., Zu, Y., Baker, C., Cheng, E., Content, D., Dooley, J., Foote, M., Goullioud, R., Grady, K., Jackson, C., Kruk, J., Levine, M., Melton, M., Peddie, C., Ruffa, J., Shaklan, S.: ArXiv e-prints (2013). 1305.5422

van Dokkum, P.G.: Publ. Astron. Soc. Pac. **113**, 1420 (2001). arXiv:astro-ph/0108003. doi:10.1086/323894

Vogt, F.P.A., Dopita, M.A., Kewley, L.J.: Astrophys. J. **768**, 151 (2013). 1303.0290. doi:10.1088/0004-637X/768/2/151 ▪