

## Fully-Automated Reduction of Longslit Spectroscopy with the Low Resolution Imaging Spectrometer at Keck Observatory

DANIEL A. PERLEY<sup>1</sup>

<sup>1</sup>*Astrophysics Research Institute, Liverpool John Moores University,  
IC2, Liverpool Science Park, 146 Brownlow Hill, Liverpool L3 5RF, UK*

(Received; Revised; Accepted)

### ABSTRACT

Although conceptually simple, the reduction of longslit optical observations is often difficult and time-consuming in practice, particularly for faint objects and in cases where the instrument is complex and/or reconfigurable. I present and summarize a software package (“LPipe”) for completely automated, end-to-end reduction of both bright and faint sources with the Low-Resolution Imaging Spectrometer (LRIS) at Keck Observatory. It supports all gratings, grisms, and dichroics, and also reduces imaging observations, although it does not include multislit or polarimetric reduction capabilities at present. To demonstrate its capabilities the pipeline is run in fully-automated mode on all available longslit data from the most recent publicly-available semester (2017A); these spectra (and the pipeline itself) are made publicly available to the community.

### 1. INTRODUCTION

Slit spectroscopy provides the cheapest and simplest way of obtaining spectral information at more than one position on the sky. Most classical optical spectrographs support some sort of longslit mode, and many instruments also support the use of customized slit-mask plates (manufactured in advance on- or off-site) to permit the observation of multiple objects that are not spatially collinear.

Most commonly, the desired output of a slit-based observation is a flux-calibrated, one-dimensional spectrum ( $\lambda$  vs  $F_\lambda$ ) of an object or a group of objects.<sup>1</sup> Production of science-quality data requires many reduction steps, including: (1) bias correction and flat-fielding, (2) removal of CCD artifacts and cosmic rays, (3) identification of the source/sources of interest on the slit, modeling the traces their dispersed light makes across the detector, and extraction of these traces, (4) wavelength calibration using arc lamps and/or night sky lines, (5) flux calibration, including correcting for effects arising from Earth’s atmosphere such as telluric absorption. Additionally, if multiple exposures are taken of the same target these must be combined together. If multiple instrument setups are used to cover a wider wavelength region, these must be similarly combined. The impacts of order cross-contamination, scattered/reflected light, and instrument flexure must also be dealt with or at least understood.

As a result, production of even a single science-quality spectrum can take a trained observer many hours (possibly, many days) of effort. This is particularly true when operating a new or unfamiliar instrument. These delays push back the timescale for discovery and publication of interesting results. This problem can be particularly acute when observing transient objects or conducting target-of-opportunity observations: in these cases, rapid characterization of the source (redshift measurement of a GRB afterglow or typing of a young supernova) is often critical to planning further follow-up of the same source. Assistance with these reductions may not be available during the middle of the night when such observations are necessarily obtained.

For these reasons, spectroscopic pipelines are becoming commonplace at a variety of observatories. Many of these pipelines are more accurately described as toolkits, as significant user input and technical knowledge is still required to transform raw data to science-quality spectra. Other pipelines are nearly automated, but only function in a narrow range of circumstances—suitable for survey projects and for telescopes whose configurations are inflexible, but

[d.a.perley@ljamu.ac.uk](mailto:d.a.perley@ljamu.ac.uk)

<sup>1</sup> In other cases, such as in measurement of galaxy rotation curves, a two-dimensional map of position versus spectrum is desired.

of limited applicability to the wider user community for more general-purpose instruments. Also, although it has recently become commonplace to require pipeline or reduction capabilities for newly-developed spectrographs as part of their construction costs, it is difficult for observatories to motivate (fund) the production of similar official pipelines for complex legacy instruments, even when they are still in common use.

One such well-established spectrograph is the Low Resolution Imaging Spectrometer (LRIS) at Keck Observatory, the first science-quality faint-object spectrograph on a 10 m-class optical telescope and still among the most commonly-used and productive instruments in astronomy (Kulkarni 2016). Thanks to several upgrades over its lifetime LRIS is a a complex and flexible instrument: it employs an atmospheric dispersion corrector, two separate channels (red and blue) with separate detectors, dispersers and filter sets, and has multi-object capabilities.

Despite its long history LRIS currently has no widely-available reduction pipeline. A few observatory-maintained IRAF packages<sup>2</sup> exist for basic tasks. A high-quality advanced reduction toolset for LRIS and several other single- and dual-channel spectrographs has been developed by J. X. Prochaska (LowRedux<sup>3</sup>), but it is not completely automatic in operation and not extensively documented. The FLAME pipeline (Belli et al. 2017) supports LRIS and a few other instruments, but requires a large amount of user input, does not perform flux calibration or telluric correction, and does not combine the red and blue arms.

In this paper I summarize a spectroscopic reduction pipeline for LRIS. It is completely automated: given a directory containing a night’s worth of raw FITS data, the pipeline will convert these files to science-quality flux-calibrated spectra of all science targets with no user input whatsoever. The pipeline is interruptible and resumable, logs its progress and actions, and it can interface with other programs. A GUI interface permits fine control over spectroscopic extraction if desired.

An imaging pipeline for LRIS is also available, although we do not describe it in detail here. The pipeline does not yet support LRIS’s multi-object capabilities, although this is a long-term goal.

The pipeline has been widely used, mainly by the supernova groups at UC Berkeley and Caltech, since it was originally developed in 2013. In this paper I provide official, formal documentation of its functionality and capabilities, and demonstrate its capabilities by applying it to a large data set from the public LRIS archive.

The pipeline is written in IDL, a proprietary language that (although formerly common in astronomy) unfortunately makes direct reuse of the code for other instruments limited. The techniques described here may nevertheless be useful for the development of similar high-level pipelines for other facilities.


## 2. SUMMARY OF LRIS

A full technical description of the LRIS instrument can be found in the original paper (Oke et al. 1995, which describes only the original red channel). The blue channel, added in 2000, is described by McCarthy et al. (1998) and Steidel et al. (2004). The atmospheric dispersion corrector (generic to Keck I), added in 2007, is described in Phillips et al. (2006) and spectropolarimetric capabilities—not supported by this pipeline—are described in Goodrich et al. (1995). The upgraded red detector (added in June 2009 and replaced in November 2010) is summarized in Rockosi et al. (2010). We do not attempt to duplicate these discussions, but do provide a brief summary of relevant capabilities for the observer, with emphasis on how LRIS compares to other facilities.

The basic configuration of LRIS—familiar to anyone who has operated the instrument from the observing GUI in the telescope control room—is illustrated in Figure 1. The instrument is located at the Cassegrain focus of Keck I. Light enters through a trapdoor, passes through a slit or slitmask positioned at the focal plane, and continues through collimating optics. A dichroic normally splits the light into red (transmitted) and blue (reflected) components—the separation wavelength chosen by most observers is 5600Å, but several alternative dichroics can be used, and the dichroic can also be removed (for red operations only) or replaced by a mirror (for blue operations only). Dispersive elements—a grating on the red side and a grism on the blue side—split the light into component wavelengths. Several red gratings (providing effective resolutions from 0.4 to 3 Å per resolution element) are available and the grating angle can be changed to control the wavelength range. Several blue grisms are likewise available (resolutions between 1.6 and 9 Å), although the blue wavelength range is fixed for each grism. Dispersed light is sent to the detectors: a pair of red-sensitive deep-depletion 2048×4096 CCDs in the red camera and a pair of blue-optimized 2048×4096 CCDs in the blue camera. Each camera has its own separate shutter (not shown in Figure 1). Calibration lamps can be turned on to illuminate the back of the trapdoor.

<sup>2</sup> <https://www2.keck.hawaii.edu/inst/lris/kecklris.html>

<sup>3</sup> <http://www.ucolick.org/~xavier/LowRedux/index.html>



**Figure 1.** The XLRIS configuration tool seen by observers at the telescope, which demonstrates the basic optical configuration of the instrument and options available to reconfigure the instrument.

Several of these features are worthy of note for performing pipeline reductions. First, the presence of two channels (red and blue) enables very broad wavelength coverage, providing continuous observations from the atmospheric UV cutoff to the sensitivity limit of the detector well beyond  $10000 \text{ \AA}$  in low-resolution mode. This greatly increases the complexity of observing and reductions—the requirements of the two cameras are often competing (when choosing exposure times, dithering methods, calibrations, and so on), and production of a final “combined” spectrum requires putting together two spectra of different resolutions, detector characteristics, and calibration files. Furthermore, the imposition of a dichroic causes strongly wavelength-dependent attenuation in the response curve in both detectors, directly in the middle of the wavelength range of the combined spectrum.

Second, as nearly every element can be moved or changed, the number of observational types that a comprehensive pipeline must support is large. There is no single “standard” setup. Furthermore, since Keck is an entirely classically-scheduled facility, observing and calibration procedures vary even for fundamentally similar observational configurations. Inexperienced users may on occasion commit mistakes in data acquisition that complicate later reduction.

Third, LRIS suffers from significant flexure and has no flexure compensation system. Observations at different positions and orientations can exhibit wavelength offsets in excess of several  $\text{\AA}$ . Corrections to the wavelength solution must be made to avoid significant errors in the wavelength calibration, which in turn would propagate to errors in flux calibration and telluric correction. Spatial flexure is also significant, and can exacerbate problems associated with irregularities in the slit.

Finally, the current LRIS CCDs on both red and blue sides exhibit a spatial gap in their center. The resulting discontinuity requires processing the two chips separately for many reduction steps.

### 3. PIPELINE OPERATIONS

The LRIS pipeline has been developed to handle all the above features, with the exceptions that it does not currently support slitmask observations or polarimetric capabilities. A driving requirement is that the pipeline is able to provide hands-off, intervention-free, end-to-end reductions for standard observing cases (specifically, for observations of single point sources in uncrowded fields with observable continuum flux).

In this section we describe, in order, the steps taken by the pipeline to achieve this aim. In default operations the pipeline automatically steps through every task in order until final spectra are produced at the end. Cautious users desiring more control can also run individual steps or series of steps, or even control reduction steps at the level of

individual targets or individual exposures. We emphasize the spectroscopic pipeline, although a sophisticated imaging reduction pipeline is available as well in the same package; this is briefly summarized in §4.1.

### 3.1. *Data regularization, bias correction, artifact correction, and header setup*

A variety of tasks are performed effectively simultaneously during the first step of the pipeline (in principle they could be separated, but performing them together saves disk write/reads.) The data format is simplified and regularized, the bias level is measured using the CCD overscan region and subtracted, and a variety of useful calculations are saved to the header for later use.

*Data simplification* - Since 2009<sup>4</sup>, LRIS observations (from both red and blue sides) have been stored in multi-extension FITS files with four extensions, one for each amplifier. Each of these also contains an overscan region. These files can be difficult for users to understand without context: orientation, wavelength direction, etc. are not obvious, there is no WCS information written to the header, and (in the case of the red detector) the order of the four files does not match their physical sequence. The first step of the pipeline is to join all four amplifiers together into a simple 2D image that can be stored as a single-extension FITS file.

In LRIS raw, multi-extension output frames, the CCD’s wavelength axis is vertical (y axis) and the spatial axis is horizontal (x axis). This convention is not well-suited to display on the wide-screen monitors that are currently in favor (which have far more pixels along the x direction than the y direction) and it also does not match the orientation of the image display as seen in the telescope control room. To further simplify things for the user, then, the x and y axes are transposed. This also ensures that for a slit PA of zero, “up” on the display is due north and “down” is due south.

*Overscan subtraction* - At the same time, the bias level of the CCD is subtracted using the overscan columns. The present version of the pipeline does not perform true bias subtraction, and any bias frames are ignored, although the capability to combine and subtract bias frames will be added in the future.

*Artifact removal* - A known bad pixel list is loaded and affected pixels are marked with NaN flags.

*Header calculations* - A variety of other parameters are calculated and added to the FITS header. These include calculation of the mid- and end-times of the exposure, the on-sky slit position angle, the average sky background, the sequence ID of the first exposure with the same setup (to help identify blocks of exposures with a common configuration), and the coordinates, elevations, and angular distances of the Sun and Moon. Interpretive comments are added for all header keywords.

Provisional WCS information (using the standard FITS WCS convention, and recognized by DS9 and other display software) is also added to the header. For imaging observations this is a standard WCS based on the RA and DEC keywords. For spectroscopic exposures the (transposed) X-axis unit is the *approximate* wavelength in angstroms while the Y-axis unit is the distance in arcsec from the central chip gap along the slit. These are approximate solutions only for the convenience of the user and do not take into account any calibration data nor any nonlinearity.


The software which saves LRIS images to disk has a curious tendency to drop (not write to disk) certain groups of telescope-related header keywords about 2% of the time. As a result, important values such as the target name, position, and slit orientation are often missing (instrument-specific keywords such as the filter or grating angle are almost never affected). Fortunately because observations are almost always taken in red+blue pairs, the lost information can be recovered by searching for a frame taken simultaneously on the opposite camera (blue or red); the pipeline conducts a check for dropped keywords and replaces them by this procedure. On rare occasions, both members of the pair have critical keywords dropped, or observations were taken in only one camera at a time (the later is fairly common during calibrations). These observations cannot be processed further unless the header keywords are manually readed (the `editheader` task can be used for this if necessary).

Processed files which use the longslit are stored in a separate directory (`spredux2d/`) from processed imaging files (`imredux/`). Output file names follow the convention of other data reduction software by prepending letters to the file name (e.g. in this case, ‘p’ is added to signify that all the preparatory steps above have been run on a file.)

### 3.2. *Flat-field construction*

There are many types of spectroscopic flat fields available with LRIS. The most common two are internal flats (using the halogen lamp) and dome flats (using the flat screen). The halogen lamp is brighter and hotter, and observations

<sup>4</sup> Prior to 2009 both red and blue CCDs stored data as single-extension FITS files



**Figure 2.** Comparison of dome and internal flat-field calibrations and of the 0.7'' and 1.0'' slits. The halogen lamp shows significant vertical banding and an overall gradient that must be removed prior to flat-fielding. The 0.7'' slit shows numerous defects (small obstructions in the slit). Dome flats with the 1.0'' slit provide the cleanest flat-fields.

can be obtained quickly at any position or orientation and before the telescope is released to the observer. However, this convenience is outweighed by the fact that the halogen lamp produces significant variations in both the intensity of illumination and shape of the spectrum along the slit, and it produces substantial scattered red light into the blue camera that dominates the signal over the entire NUV range (Figure 2). Dome flats avoid all of these problems but take much longer to expose (especially in the blue), and contain weak emission lines and weak telluric absorption from air within the dome. The dome lamp is not hot enough to calibrate the NUV range.

Deuterium flats are a relatively new addition to LRIS; a dense set of broad molecular lines produces a pseudo-continuum in the ultraviolet. Sky flats taken during twilight are in principle the best type of calibration, but can only be taken in a short window at the start or end of the night and are not obtained by most users. Both flat types contain strong, narrow line features that must be removed to be useable as flat-fields in practice. Deuterium flats, like halogen flats, also exhibit highly non-uniform illumination along the slit direction.

When run in default mode, the pipeline will only construct the processed flat-fields that it needs for the science configurations that were actually observed that nights (although it can be configured to construct processed flats even for configurations with no science data as well). If multiple flat types are available within a given configuration, the pipeline will produce separate processed flat fields for each flat category and choose between them later.

Flat-field exposures are identified automatically by the pipeline based on their header keywords (e.g., dome flats can be recognized by checking the relative azimuthal offset between the telescope and the dome, which should be 90 degrees exactly), amount of continuum signal, and exposure times. These are sorted into groups of exposures taken with the same configuration, which will be combined together to create a processed flat. Exposures with low S/N are excluded automatically.

For basic flat types (halogen/dome), the wavelength structure of each individual flat exposure is removed by dividing each column by its median value. For flats with strong line structure (sky/deuterium) a response function at the middle of each CCD, then shifted to match each individual column's wavelength offset and divided. In either case, this procedure will have the side-effect of also removing any illumination variations (vignetting) along the wavelength axis of the CCD. LRIS's illumination variations appear to be negligible across most of the CCD but may be significant closer to the edges, in particular at the far ends of the red CCD. This does not matter as long as the wavelength solution remains fixed between observations of the flux standard star and science object, but see the section on flux calibration for more information.

Neither the halogen nor the dome lights present sufficient UV flux to calibrate the blue side shortward of about 4000 Angstroms. The pipeline therefore does not attempt to flat-field data in this region: rather, it simply takes a wavelength average over the high-S/N regions and applies this to every row across the (low-S/N) UV. Note that this procedure can impart a small discontinuity in the effective wavelength response at the position where flat-fielding is no longer possible, although in practice this has never been visible in output spectra. Future pipeline releases may employ a standard pixel flat calculated from twilight frames to avoid this issue and improve flat-fielding generally.



For halogen flats and deuterium flats, the spatial variations along the slit direction are removed by taking horizontal medians and dividing. (This procedure is not necessary for dome or sky flats.)

Air in the dome introduces telluric absorption in the A and B bands for dome flats. The effect is small and restricted to these bands, but is visible as “ripples” in the flats, which are not fully removed by the median division due to the small tilt in the wavelength axis relative to the CCD axis. These are not currently removed.

After the horizontal (and for halogen flats, vertical) response corrections, the flat exposure sequences are median-combined together and saved to disk.

### 3.3. Flat-fielding

In comparison to the construction of the flat-field, the calibration of science data using this field is straightforward, at least in principle: each science exposure is matched to an appropriate flat taken in the same configuration and divided (arc exposures are also flat-fielded although this is not strictly necessary). However, there are some subtleties.

First, it is worth noting the issue of cropping and binning. While most observers will use the same binning settings for the entire night, it is conceivable that some will procure some binned and some unbinned frames either deliberately or by accident, and not acquire flats in the same mode. It is possible to rebin frames in software to deal with this fact and LPipe will do this automatically if necessary. Likewise, while many observers will choose to set the red CCD window settings to only read out the area around the chip to save disk space and transfer time, some frames may accidentally be acquired in a different window setting. LPIPE keeps track of binning and cropping coordinates (both at the telescope and as applied by the pipeline during processing) and will ensure the frames are aligned appropriately, allowing e.g. unwindowed flats to be applied correctly to windowed data. If science observations were taken in full-frame mode, it will also crop the image down vertically to only include the area that is actually exposed by the longslit.

Another important subtlety concerns irregularities/obstructions in the slit, in particular for the 0.7 $\ddot{s}$ lit. This slit has a number of defects in the form of small obstructions along the slit that block a small amount of light at certain positions, an effect that is clearly visible as fine horizontal ‘banding’ in flat-fields taken with this slit (Figure 2). The 1.0 $\ddot{s}$  has a much smaller number of irregularities (the 1.5 $\ddot{s}$ lit has none). In principle, these variations should be removed from science frames by flat-field division. Unfortunately, the physical location of the slitmask moves during the night in response to flexure, and the locations of these bands are observed to move vertically up and down between different exposures. While the slit variations are only a few percent, they are highly undersampled and this will have severe consequences on the quality of 2D sky-subtraction. Dividing a science observation by a (spatially-shifted) flat can also create a positive+negative trace pairs along the slit—which can confuse extraction for observations in twilight when the sky level is strong. Tools have been developed to mitigate these effects, but they are not automated within the current version of the pipeline and users should be aware that for the 0.7 $\ddot{s}$ lit in particular, unsupervised reductions are likely to produce science results that are significantly worse than what is possible in principle with properly flat-fielded observations.

Regions of the flat-field with transmission below 1% of the peak (typically, bad pixels and portions of the CCD away from the longslit) are set to NaN in the final image to avoid immensely magnifying noise.


### 3.4. Splitting

After flat-fielding, all images are still stored as simple (single-extension) FITS files with the two chips merged together. However, it soon becomes justified to split the chips again, for two reasons. First, this saves processing time during the computing-intensive cosmic ray rejection and sky-subtraction steps for the (extremely common) case that only one chip is of interest. Second, this avoids having to continue to recognize the spatial discontinuity at the slit gap position in future processing steps. Therefore at this stage the two chips are split into separate files, prepended by `.r.fits` and `.l.fits`, respectively.

By default only the right chip is split off and the left chip is ignored from this point forward. The user can override this default and retain both chips by specifying this option at the command line when the pipeline is run. However, automated source identification and extraction is currently only available for sources on the right chip since in almost all cases the target of interest is centered on the right chip.

### 3.5. Cosmic ray cleaning

A casual glance at a long exposure on LRIS-red immediately reveals the challenges involved in processing data from this detector: the chip is filled with long charge trails from radiation and cosmic-ray hits, affecting a significant fraction



**Figure 3.** Cosmic ray cleaning demonstration. While the upgraded LRIS-red detector is significantly affected by cosmic rays (left half), these can be efficiently detected and removed by a specialized algorithm (right half, shown after cleaning). The wide horizontal grey band just below the middle of the images is a bad column; the faint horizontal depression in the flux just above the middle is due to a slit defect (§3.3).

(up to 1%) of the pixels per exposure (Figure 3). The trace of the target source is likely to be affected at several points. This is a common feature of deep-depletion CCD arrays.

Old cosmic ray detection algorithms (such as the basic tool in IRAF) work poorly on such features because they affect so many pixels over a large area. Tools developed for space-based imaging may also struggle due to the presence of the bright vertically-structured sky lines which may themselves be moderately undersampled in some binning modes. For these reasons a custom tool was developed, code-named *crayfish* but referred to more prosaically in the pipeline by *pzap* (p for percolation.) This algorithm proceeds in several steps: first, it employs a row/column median subtraction to remove sky lines and source continuum; next, it attempts to model the value of each pixel in the image based on groupings of its neighbors and identify cosmic rays as pixels whose model values disagree both with each other and the real value; finally (and most importantly) it employs a percolation algorithm to try to identify every pixel affected by a long particle trail.

The procedure is not perfect. Cosmic rays with directions of travel along the spatial direction are usually missed because they cannot be distinguished from sky lines using the current algorithm. The percolation procedure sometimes misses a few affected pixels. Most alarmingly, binned observations of galaxies with extremely strong emission lines can cause these lines to be misidentified as cosmic rays in some instances. In general, observers doing narrow emission-line science should employ caution in verifying the pipeline products after this step. Observers with no interest in emission lines (and not employing more than  $2 \times 2$  binning on the red or  $2 \times 1$  binning on the blue side) could safely make the pipeline more zealous and improve the cleaning efficiency.

After cosmic rays are identified they are filled in with a median-model estimate of the pixel values. A mask image indicating the positions of affected pixels is retained for use in flagging compromised regions of the spectrum in later reduction stages.

Cosmic ray rejection is slow; typical processing times are 20 seconds per  $2 \times 1$ -binned image, adding up to 15+ minutes per camera for a typical night. (This time is more than doubled for unbinned data). We are exploring faster and more reliable algorithms.

The blue side is much less-affected by cosmic rays but still contains them; a simplified (and faster) version of the cosmic ray cleaner is applied to blue observations.

Arc files are not cosmic ray cleaned: these are always very short exposures<sup>5</sup> and cosmic rays can be removed by using median-averaging during the extraction.

### 3.6. 2D Sky subtraction

The next step is to 2D sky subtract the images in order to separate signal originating from Earth’s atmosphere (of interest only for calibration purposes and to estimate the noise level that results from them) from the signal originating from astronomical targets (the actual sources of interest).

<sup>5</sup> Except for the Fe arc, which is not used by the pipeline

The basic underlying procedure is to take medians along each wavelength column and subtract them. However, a simple median would not be able to deal with even a small amount of curvature or tilt in the sky lines, so in actuality the medians are calculated in blocks, interpolated, and then fit with a simple regression model. The sky flux is taken out of the science image and written to an extension.

This algorithm works well for modest tilts but fails if the sky lines are tilted with respect to the CCD grid by more than a few degrees, which can be problematic for some gratings. Furthermore, although smooth continuum from sources along the slit are removed by a separate median procedure before fitting the sky lines, absorption and emission lines will subtly influence the sky model and typically result in under/oversubtraction of the sky around them, respectively. Also, very large sources (nearby galaxies) will have any low spatial-frequency components removed as “sky.” Fortunately, this is not generally a problem: the only function of 2D sky subtraction when running the pipeline in default mode is to remove enough sky signal to identify and trace faint sources in later steps, and over- or under-subtraction is unlikely to affect this significantly.

An alternative sky-line subtraction routine based on the method of (Kelson 2003) has also been developed and integrated into the pipeline. Experimentation has shown it to be generally inferior to the spatial median subtraction above for LRIS under most circumstances, but the two algorithms are interchangeable within the codebase and the user can choose which to employ at runtime.

A different type of alternative method for sky-line subtraction employs nodding (dithering the telescope back and forth along the slit, and subtracting). Nodding is not commonly used by LRIS observers because of the long readout times and the competing demands of the blue spectrograph (which is read noise limited for faint-object observations), although it was much more commonly employed in the past in order to mitigate the severe line fringing in the near-infrared observed with the old red detector. Support for nodded pair subtraction does not currently exist in the pipeline but may be added at a later date.

### 3.7. *Shifting and stacking*

It is common, especially on the red side, to acquire repeated exposures for a single object. It is beneficial to stack these prior to source extraction, since this makes it much easier to identify and trace faint sources. Therefore the pipeline looks for a sequence of exposures with the same configuration and at the same position on the sky *except* for small dithers along the slit direction. Using the header positional keywords, all frames after the first are shifted vertically to match the spatial location of the initial exposure (if necessary), and then simply summed, excluding the cosmic-ray affected regions.

In some cases the user may want to extract multiple observations of a source as individual spectra - for example, in order to obtain a time series for a source that is rapidly varying; this option can be specified at runtime. In other cases, a temporary guiding loss or other error may cause repeat exposures in the same setup to not be summed, or the user may want to combine observations of the same source taken at different times in the night. These can be summed explicitly by the user in 2D space or later in 1D space using the appropriate individual commands.

Once exposure summation is complete, the pipeline reaches a transition point: all further steps involve conversion of 2D information to 1D information. In recognition of this, the output files are copied into the `spredux1d/` directory where all further processing will occur. Once this has occurred, the entire `spredux2d/` directory can be safely deleted to save disk space if desired.

### 3.8. *Tracing*

Tracing (determining how the wavelength-dispersed continuum signal from an object tilts and curves as it extends along the CCD in the X direction) is conducted prior to (and separately from) extraction. This aids the identification of very faint sources, and avoids having to repeat the tracing procedure if the aperture extraction details are changed later.

For bright sources (such as standards) tracing is straightforward: the spatial ( $y$ ) centroid is calculated for every wavelength ( $x$ ) column and a fitting function ( $y = F(x)$ ) is applied to calculate a smooth trace curve. This is not always reliable for science observations where the source may be weak at some or all wavelength columns. To improve S/N, the observations are (temporarily) binned up by 32 pixels or more along the wavelength axis, producing a series of spatial profiles along the (entire) slit at different wavelengths. These profiles are then cross-correlated to measure the spatial offset—starting in the middle of the wavelength range where sensitivity and source counts are likely to be maximized, and then progressing outward in both directions, requiring that any variations be small. A polynomial fit



is applied to these to determine the trace function. This function is written as a file ending in `.trace`. If a consistent, slow-varying trace function cannot be found from this method because there are no sources apparent in the image then no trace file is written (this is not a problem: subsequent steps will simply rely on trace files from other, brighter targets.)

In cases where a usable trace is present for only part of the wavelength region occasionally the tracing can get 'lost' in the low-S/N region of the spectrum (this is particularly common when using the low-resolution blue grism, where the trace extends far into the UV where there is usually little signal). This is usually not a problem since the affected region is unlikely to have any usable continuum signal anyway, but it may cause any emission lines to be missed.

### 3.9. *Source and aperture identification*

Next the pipeline must identify what sources to extract. This is by far the most likely step where a user may want to intervene in the automated pipeline operations (generally, to modify apertures or background subtraction behavior, or to extract a different source).

First, the pipeline completely collapses each image along the wavelength axis (the true wavelength axis as determined by the trace solution determined above, not the CCD X-axis) to produce a measurement of the 1D spatial profile along the slit. This is written to disk as an ascii table with a filename ending in `.profile`; the position is referenced to the middle of the CCD). Then, for each observation, the profile file is read in and all significant peaks are identified as candidate sources. The pipeline chooses which source is most 'likely' to be the science target primarily based on proximity to the pointing origin, which is determined by examining the positions of bright star traces observed using the same configuration for which there is no ambiguity about the identify of the target of interest. (The brightness of the source is used as a secondary factor.) The Y position of this source at the reference wavelength (generally the midpoint of the CCD) is saved to two text files: `brightobjectpos.txt` (which contains bright, unambiguous traces that are usually standards) and `objectpos.txt` (which contains the remaining targets). These should not be edited by the user.

This routine is initially carried out independently for blue and red data. In principle this can lead to a different object being automatically extracted on the red vs. the blue side (and these spectra would then be fused together into a chimeric spectrum in future steps), so afterwards the pipeline scans through the lists of extraction coordinates and resolves any such conflicts by calculating the average red-blue  $y$ -offset and identifying any outliers. The aperture and background sizes are also averaged together for consistency.


Because the entire wavelength axis is summed over, this procedure can effectively identify sources as faint as 25th magnitude (in long exposures under dark-sky conditions). Sources with continuum fainter than this, including objects dominated by emission lines, may not be recognized; these will have to be pointed to manually by the user (see below).

In any case, no automated routine can determine with certainty which source(s) a user is interested in or the optimum way to extract them. While extraction of bright stars and QSOs is usually obvious and intervention is rarely needed, for more subtle cases (a faint star in a dense field, a supernova embedded in a bright galaxy) it is usually desirable to specify exactly which source to remove and how to deal with the surrounding background.

As a result, a sophisticated GUI interface for aiding source selection has been developed to supplement the pipeline. This can be started by typing `setapertures` at the IDL command line. This will bring up a window (Figure 4) with a plot of the spatial trace profile (plotted along the monitor's X-axis, but keep in mind that by convention the spatial direction is vertical in 2D frames.) Unless the pipeline failed to identify any candidate sources, several lines will be drawn on this plot: a red line to indicate the trace center, a pair of green lines to indicate the extraction aperture limits chosen by the pipeline, and two pairs of yellow lines to indicate the background aperture. The buttons of these same colors (red, green, yellow) in the left panel can be selected to enable the user to reposition these lines using the mouse and control all aspects of the extraction aperture placement.

The user can request multiple sources to be extracted from a single image by selecting 'New' and placing a new extraction aperture on the diagram (the pipeline automatically chooses only one source). Its width, background, etc. can be controlled as described above. Up to 4 objects can be extracted per longslit; one can toggle between them by selecting the number buttons in top left. It is important to add apertures consistently on the blue and red sides if these are to be connected consistently in later pipeline steps.

If the source has no continuum even in the wavelength-summed profile but does have lines in some parts of the spectrum, apertures can also be selected in 2D space by viewing the entire 2D spectrum. An option to toggle the



**Figure 4.** Screenshot of the aperture selection tool GUI available in the pipeline. Apertures can be resized, moved, added, or subtracted, and the background subtraction algorithm can be changed interactively by the user. This is necessary only for objects with close neighbors and complex regions; in most cases the automated trace identification can be relied on.

1D/profile view and the 2D spectrum view is available. It is also possible to zoom in or out of particular spatial regions.

Tracing options are stored as a simple ASCII text table containing (on each line in order) the file name, the trace center and aperture radius, and the upper and lower background limits (as measured from the trace center). They are saved to `userobjectpos.txt`. These parameters can be specified by the user without use of the GUI by editing the text file `userobjectpos.txt`.

Aperture placement is currently stored in image coordinates, which makes it susceptible to changes in cropping. In the future physical coordinates will be used.

### 3.10. Extraction

The actual extraction of a source is straightforward since all details of the aperture placement have been determined in advance. Extraction is effectively performed by spatial convolution: an extraction kernel of the mathematical form:

$$K(y) = \exp\left(\left(\frac{y - y_{\text{trace}}}{r_{\text{aperture}}}\right)^n\right)$$

is calculated for every value of  $x$ , and the appropriate pixel column is multiplied by this function and the output summed (i.e., the counts at a given wavelength column are calculated as  $C'(x) = \sum_y K(y)P(y)$ ). This kernel is a Sersic function, although our  $n$  is the inverse of the classical Sersic index:  $n = 1$  would correspond to an exponential kernel,  $n = 2$  a Gaussian kernel, and  $n = \infty$  a top-hat kernel (i.e. a classical simple aperture). We typically use  $n = 6$  (a Sersic index of  $1/6$ ) to keep the profile flat in its center while also having soft edges to avoid unphysical discontinuities that result when the profile suddenly jumps from one pixel to another.

To remove the background, the average flux over each background aperture (upper and lower) is calculated and this is interpolated to the source position and subtracted from the source counts above. This is not performed if

background aperture subtraction is turned off (the presumption being that the 2D sky subtraction has removed the background already.)

Unless the use of 2D sky subtraction is turned off, the pipeline repeats this procedure for the sky extension, except with no background aperture. This produces a sky spectrum, which will be saved alongside the object spectrum. (If 2D sky subtraction is off, the sky spectrum is estimated using the background apertures instead.)

This procedure differs from the so-called “optimal” extraction method summarized in [Horne \(1986\)](#). Optimal extraction requires very precise measurement of both the spatial profile and the trace function. Both of these are supplied by the pipeline, but for faint sources or sources with very red or blue colors they may not be reliable. Use of optimal extraction (especially for bright sources) will be explored in the future. This could provide a S/N improvement of 40% or more.

The output is saved to a simple ASCII file with several columns: wavelength pixel, source counts, sky counts, trace centroid, and the fraction of extracted pixels which were bad. A FITS-like header containing all the key information from the original FITS header along with some information about the extraction parameters used is written at the top of the file. This format is used for all 1D spectra written by the pipeline.

### 3.11. *Arc wavelength calibration*

Wavelength calibration of the arc exposures could in principle be performed either before or after source identification and extraction, since neither step is dependent on the other. Currently arc calibrations are run after extraction.

For maximum flexibility (to allow new gratings, CCD changes, etc) the pipeline treats wavelength calibration in a fully general manner: no archival wavelength solution is used and the wavelength-to-pixel solution is treated as a novel problem with no strong priors - other than the approximate initial central wavelength and dispersion, plus the degree of polynomial to fit.

Basic extraction of all valid arcs is first performed (a vertical median is taken over a limited region close to the standard trace position) to produce 1D spectra. A pattern-matching algorithm is used to identify matching groups of bright line triplets in the line list and the arc itself to identify the lines. At that point a linear regression model is used to fit the wavelength solution, which is successively refined by dropping lines that are poor matches (these may be saturated lines, blended lines, or misidentifications from the pattern-match) and re-matching until a good solution is achieved. A high-order polynomial is needed to accurately fit the whole wavelength range for low-resolution setups (currently a 5th order polynomial is used for the blue and a 7th order polynomial for the red.)

A practical challenge is the fact that the LRIS arc lamps have very few lines around 5500 Å (the only bright line between 5085–5769 Å is the Hg line at 5461 Å). This region includes the wavelength limit of the blue 400/3400 and 600/4000 grisms, and will likewise include the limit of the red grating if continuous observations are desired. If the D560 dichroic is used, this is also a region of rapid variation in the transmission on both cameras and small wavelength inaccuracies can impart large flux calibration inaccuracies, which in turn will lead to inaccurate red-blue relative scaling. Users should be alert for flux- and wavelength-calibration artifacts in this region.

Best results are obtained if as many lamps are turned on as possible. Certainly, for evening/morning calibrations, there is little reason not to employ every arc lamp. The original LRIS arc lamps are Cd, Zn, Hg, Ne, Ar. All of these should be turned on for evening and morning calibrations if possible, although strictly only Hg, Cd, and Zn are needed for the blue and only Hg, Ne, and Ar are needed for the red. The Xe and Kr lamps were added in 2015: these are recognized by the pipeline and can be turned on as well. (In contrast, the new Fe lamp is too weak to be of practical use, and De is a continuum lamp.)

Many observers do not turn on the Hg lamp for red calibrations (it is color-coded as a blue lamp in the LRIS observing GUI). The pipeline will still process arcs that employ the Ne and Ar lamps alone. However, use of red arcs without Hg is not recommended because the Hg 5461 line, although heavily attenuated on the red by the D560 dichroic, is important to ensure accurate calibration of the wavelength solution blueward of the last Ne line at 5769 Å.<sup>6</sup>

On the blue side, the lengthy warm-up time for the Cd and Zn lamps often leads to exposures being taken too early before these lines have appeared. These arcs will be inadequate to solve the wavelength solution since the Hg lamp alone does not have enough lines. If multiple identical arcs are taken the pipeline will use only the last one (under the assumption that earlier lamps are more likely to be unwarmed). However, wavelength solutions are always validated

<sup>6</sup> If the grating angle is set to maximize the NIR coverage at the expense of red-blue overlap, this line might not be covered by the CCD anyway and this region cannot be accurately calibrated, even with the Hg lamp on.

internally to check that all expected bright lines are present, so even if “bad” arcs are present they will be ignored in further processing.

Successful solutions are written to disk in the form of a column of coefficients from the polynomial fit. Every available arc exposure is solved individually, except when several identical arcs are taken in series (see above). Decisions about which of these solutions to employ to calibrate which science exposures are taken later. The pipeline is able to translate solutions derived from arcs taken in one binning to science frames taken in a different binning, so it is not necessary for the user to acquire arc calibrations in the same binning as science observations.<sup>7</sup>

### 3.12. Wavelength calibration and flexure correction

Next, the pipeline will apply the wavelength solutions derived from the arcs to the extracted science frames. This requires both matching each science frame to an appropriate arc frame and applying a subsequent flexure correction.

For each configuration, LPipe will look up an appropriate global reference arc exposure (usually one from late in the afternoon) and read in its wavelength solution. Only this arc will be used to estimate the higher-order (dispersion and curvature) terms of the wavelength solution in order to guarantee that all observations taken using the same setting use the same original arc calibration file.

Due to flexure, adjustments to this wavelength solution will be necessary to calibrate observations taken during the night. In most cases the night-sky emission lines in the spectrum can be used: a line list is estimated from the night-sky spectrum and matched to a catalog to determine a linear offset.

If this fails or is impossible due to a short exposure or bright twilight, other methods are employed. If an arc was taken at the precise position of the observation, the zeroth-order term of that arc is substituted for the one from the reference arc. If the object was very bright (e.g. a standard star), the region of the spectrum around the telluric A band (or the B band if A is not covered) is cross-correlated with a reference spectra to estimate the wavelength offset.

There are no bright sky lines at wavelengths shorter than 5577 Å (one weak line is present at 5197 Å but its S/N is always low), and no telluric lines either. As a result, flexure correction is often difficult for LRIS-B. The 5577 Å OI line is always present in the wavelength coverage of the 300/5000 and 400/3400-line grisms, and this line alone is usually enough to measure the flexure correction. This line is normally present in the 600/4000 grism as well, *but* a hardware intervention in 2012 resulted in a shift in wavelength coverage to the blue that meant that the line shifted off the CCD until the D560 dichroic angle was changed in early 2014. Observations with the 1200/3400 grism will never show this line, as will observations with any grism if the D460 or D500 dichroic is used. In these circumstances on-sky arcs provide the *only* way of providing accurate wavelength solution.<sup>8</sup>

In the future it would be desirable to perform wavelength calibration in 2D space. This would require carrying out this step earlier in the process, prior to source extraction, and therefore a significant change to the pipeline workflow. This will be explored in future releases (perhaps in conjunction with multi-slit capabilities.)

### 3.13. Flux calibration

Flux calibration involves calculation of the response function (sometimes called the sensitivity function) using an object whose intrinsic spectrum is known. In its simplest form, the response function is described by:

$$C(\lambda) = F_{\lambda}(\lambda)R(\lambda)$$


Here,  $C$  is the count rate in ADU/s (per extracted pixel) as calculated earlier<sup>9</sup>,  $F_{\lambda}$  is the true intrinsic spectrum of the source, convolved to the resolution of the instrument, in  $\text{erg s}^{-1} \text{m}^{-2} \text{Å}^{-1}$ , and  $R$  is the response function (units:  $\text{ADU erg}^{-1} \text{cm}^2 \text{Å}$ ), which allows us to convert back and forth between these two quantities.

It is simplest to assume that  $R$  is a function of wavelength alone and thus it can be calculated in a simple manner as  $R = C/F$ . In reality, however, it is a function of detector position and wavelength. Pixel-to-pixel variations were removed by flat-fielding, but broader variations associated with the detector or the illumination pattern cannot easily be solved for without a model of the spectrum of the flat-field source. The distinction is unimportant *as long as* the wavelength solution is the same for all spectra, such that pixel and wavelength can be treated as interchangeable. This is approximately true in practice if calibrations are taken in the same setup as science observations, but is not exactly

<sup>7</sup> While most users will acquire all observations in a consistent binning setting, acquiring arcs in 1x1 binning mode (regardless of the science setting) can actually be useful to avoid saturation of the brightest lines for binned observations with low-resolution gratings.

<sup>8</sup> It may be possible to use cross-correlation of the sky continuum against a Solar spectrum to correct blue observations taken in twilight or bright moonlight in the future, but this feature does not currently exist.

<sup>9</sup> The ADUs are corrected for flat-field pixel variations but *not* gain or illumination variations



**Figure 5.** Effective area and total system efficiency (including atmospheric and slit losses, except telluric absorption) of LRIS in a standard wide spectral coverage setup, as measured in July 2016 using the standard star G191-B2B. The 600/4000 grism, 400/8500 grating, D560 dichroic, and 1.0" slit were employed.

true in reality due to the effects of flexure. We assume that the illumination of the CCD is constant within the flexure variation scale.

As a first step, we solve for  $R$  for every standard individually by simply dividing the reference spectrum of an object (interpolated to the wavelengths of the CCD pixel grid) by its observed count spectrum. Reference spectra are taken from the CALSPEC database at StSCI<sup>10</sup>, the ESO libraries<sup>11</sup>, and the ING compilation of standards<sup>12</sup> and come from a variety of individual sources (Oke & Gunn 1983; Oke 1990; Stone 1977; Hamuy et al. 1992, 1994; Bohlin et al. 1995; Bohlin 1996). This is only an initial estimate of  $R$ : it contains noise (especially for short exposures and in the UV), and the limited resolution causes the observed spectrum to be ‘blurred’ slightly on the detector (additionally, many library reference spectra omit the stellar absorption lines entirely!).

These effects must be corrected for. LPipe first looks up what wavelength regions of the spectrum are affected by lines, omits these from both the observed and the intrinsic spectrum, and interpolates over them. Telluric absorption lines are initially interpolated over for the same reason (these will be dealt with with a different procedure momentarily). The response curve is then smoothed by Savinsky-Golay convolution to reduce noise. The telluric absorption is then estimated by comparing the actual counts over the A, B, and water vapor bands to the interpolated counts over these regions: i.e., the pipeline calculates  $a_{\text{tel}}(\lambda) = C_{\text{obs}}(\lambda)/C_{\text{interp}}(\lambda)$ . This is required to be between 0 and 1 (any pixel values  $< 0$  and  $> 1$  are set to 0 and 1, respectively). The non-telluric and telluric components are written to disk along with the wavelength and pixel values in a four-column ASCII file, ending in the file extension .response.

A helpful byproduct of the response calculation is an effective area curve for the observation, which describes the equivalent collecting area for a perfectly efficient instrument located above Earth’s atmosphere that faithfully records every photon incident upon its surface; this provides a directly measure of the real-world performance of the observatory. It can be calculated as  $\mathcal{A} = R G h \nu / \Delta \lambda$ , where  $G$  is the CCD (inverse) gain,  $h \nu$  is the photon energy, and  $\Delta \lambda$  is the wavelength range for one pixel; typically this is observed to be  $\sim 10 \text{ m}^2$  in the blue and  $\sim 15 \text{ m}^2$  in the red for observations in clear weather and average seeing in the regions of optimal sensitivity of each grism and grating. The ratio of this to the physical area of the primary mirror ( $73.3 \text{ m}^2$ ) gives the total throughput of the system as a function of wavelength, approximately 12–20%. This is consistent with the expected losses: a breakdown (at  $6000 \text{ \AA}$  under typical conditions, given values from Oke et al. 1995) is: 0.9 (atmosphere)  $\times$  0.86 (primary)  $\times$  0.88 (secondary)  $\times$  0.8 (slit/aperture loss)  $\times$  0.95 (dichroic)  $\times$  0.4 (rest of spectrograph)  $\times$  0.8 (detector QE).

Telluric absorption, atmospheric extinction, and slit losses all vary as a function of sky location and atmospheric conditions, and thus the response function observed for a standard may not be the same as the response function for

<sup>10</sup> <http://www.stsci.edu/hst/observatory/crds/calspec.html>

<sup>11</sup> <https://www.eso.org/sci/observing/tools/standards/spectra.html>

<sup>12</sup> [http://www.ing.iac.es/Astronomy/observing/manuals/html\\_manuals/tech\\_notes/tn065-100/workflux.html](http://www.ing.iac.es/Astronomy/observing/manuals/html_manuals/tech_notes/tn065-100/workflux.html)



a science object. Slit losses, which may be wavelength dependent (seeing is generally worse in the blue), are difficult to correct for and provide by far the largest uncertainty in both relative and absolute calibration of LRIS data <sup>13</sup>; the pipeline is not able to address this at present. The other two effects (atmospheric extinction and telluric absorption) can be addressed more easily.

The absorption properties of the atmosphere change gradually from night to night (due to variations in water vapor density or the amount of dust and aerosols in the atmosphere) and in principle it would be desirable to solve for atmospheric extinction using the nightly data, e.g. by observing multiple standards at different airmass. Unfortunately, slit loss variations (which also vary with seeing, focus, and pointing) make this not reliable in practice. Instead, differential atmospheric extinction is treated by employing the mean Mauna Kea atmospheric data from [Buton et al. \(2013\)](#), which describes the average opacity at zenith. A polynomial is fit to their data to produce a generic zenith attenuation function, expressed in terms of optical depth as  $\tau_{\text{atm}}(\lambda)$ . This is used to convert the measured response function at the airmass of the flux standard to a predicted response function at the airmass of the science target: i.e.  $R_{\text{cont,sci}} = R_{\text{cont,std}} \times \exp(-\Delta A * \tau(\lambda))$ , where  $\Delta A$  is the difference in airmass  $A$  between science and standard targets ( $\Delta A = A_{\text{std}} - A_{\text{sci}}$ ). Fortunately the opacity of the Mauna Kea atmosphere is quite stable. While variations of 10% are observed in the relative opacity of the blue and the red, these are much smaller than errors associated with slit loss effects.

The telluric component of the absorption is corrected in a similar way, except in this case no model is necessary because the telluric component has been separated out from all other flux-loss effects. Given an airmass difference  $\Delta A$  between the standard and science target, we simply rescale the telluric attenuation component according to the equation  $f_{\text{tel,sci}} = f_{\text{tel,std}} \hat{(1 + \Delta A/A_{\text{std}})}$ .

A good estimate of telluric absorption requires a good estimate of the continuum outside the telluric bands. If the standard is noisy or has weak intrinsic features that are not fully removed this may not be the case. Furthermore, this telluric procedure does not take into account the limited resolution of the instrument. In practice, this limited resolution makes the true airmass correction more complicated than the above equation describes outside the optically-thin limit due to saturation effects. Because resolution is seeing-dependent, the appropriate telluric correction in practice will depend on seeing. These two effects will inevitably lead to some residuals, although these are small in practice as long as standard-star observations are carried out at similar airmass ( $\pm 0.3$ ) to each science target. If science observations are taken at very different elevations from any standard star this will inevitably cause large residuals over the telluric bands.

The pipeline must choose which standards to use to correct which science targets. It will, of course, will only consider standards observed in the same instrumental setup of a given science target. If multiple standard-star observations are obtained (generally a good idea), it must choose between these. The primary consideration in choosing between these is the airmass difference versus the science target. However, many other considerations are also considered, including: whether the standard star is known to have a smooth continuum (vs. many absorption lines that may not be well-characterized in the internal database), whether the standard exposure has a good wavelength calibration (either a long enough exposure time to derive sky lines, or an on-sky arc), whether the standard was observed close in time to the science target, and whether the flexure difference is observed to be small. The pipeline weights these factors together to decide on the “best” standard to use. In some cases two standard response functions may be averaged together.

An important complication that occasionally arises is that of cross-contamination of parts of the spectrum by light at other wavelengths, due to second-order light or reflected light. In most LRIS setups, this is avoided by the use of the dichroic, which cuts off light from orders other than the one desired. However, the dichroic does transmit some blue light and reflect some red light. These effects are currently neglected by the pipeline as no observable impact has been seen in any standard setup, but caution should be exercised when observing very blue or very red targets or in nonstandard setups.

One particularly important case that is commonly encountered in practice involves the use of the 300/5000 blue grism with the D680 dichroic. In this setup the blue spectral region longward of  $\sim 5000\text{\AA}$  is severely contaminated by second-order light and/or reflections. The pipeline will reduce data from this setup (truncating the wavelengths most severely affected by contamination) but the output should be interpreted only with extreme caution since both the

<sup>13</sup> This problem is greatly exacerbated if a the slit is off-target by an amount comparable to the slit size or if the focus is poor.

target and the flux-standard used to calibrate it will suffer from poorly-quantified (but typically large) contamination issues.

A practice employed by some observers to mitigate against light from unwanted orders/reflections (even in setups where it should be minimal) is to observe hot, blue stars to calibrate the blue side and to observe cool, red stars to calibrate the red side. Unfortunately, cool stars usually also have more prominent atmospheric absorption features than the hot, blue white dwarf stars generally preferred as standards. Currently the pipeline exhibits a preference for blue standards with few lines even for the red side and red-specific standards of this type are likely to be ignored in reductions, but the user can specify specific calibration targets if desired.

Regions of very low sensitivity are removed from the output spectrum. Generally this means the UV (the 400/3400 grism nominally extends to 1300Å, but no usable signal ever results shortward of 3000Å), and regions heavily attenuated by the dichroic.

### 3.14. *Red/blue connection*

Up to this point the red and blue cameras have been treated as if they were entirely separate observations conducted with different telescopes (with the single exception of a concordance check to be sure the same source is being automatically extracted). However, most observers want a single file describing the combined spectrum of each source. Therefore the two spectra have to be combined.

As with many pipeline steps, this seems like it should be trivial but can be complex in practice. The red and the blue sides may have been calibrated with different flux standards, with different slit/aperture losses and transmission losses, introducing an absolute (multiplicative) offset between the red and blue science spectra. (Offsets are sometimes seen even when the same standard is used.) This must be removed to avoid imposing an unphysical discontinuity and to avoid turning errors in the absolute calibration (rarely a problem for science observations) into errors in the relative calibration (which are a problem for many applications).

This procedure is only possible for observations in which there is overlap between the red and blue spectra. (If this is not present, the two files are simply merged together with a gap and the final spectrum is written.) The pipeline uses the relative response curves to identify a wavelength region where there is acceptable sensitivity on both cameras.

For bright sources, the pipeline then calculates the flux ratio between the red and blue within this region at each pixel, then determines a statistically-weighted average, which is the overall scaling factor. Flux values on one side are divided by the square root of this average and flux values on the other side multiplied by the square root of this average. (Sky, uncertainty, and response values are also rescaled).

For objects with very faint continuum this procedure is not possible because the S/N is too low. In principle, the sky continuum level or (for the D560 dichroic) the total sky flux under the OI 5577Å emission line could be used, although this requires extremely accurate tracing on both sides right up to the CCD edge, and in practice this technique is not sufficiently reliable for pipeline use. For now, if the S/N is judged to be too low for a direct rescaling (or if there is no wavelength overlap) no scaling is performed: the spectra are joined together at the junction point with no scaling. In such cases the user may wish to carefully examine the spectrum and perform a manual rescaling by eye as needed.

In principle it would be possible to co-add blue and red spectra over the overlap region using a weighting function. This would involve interpolation/resampling, which in general the pipeline avoids in order to keep the noise properties as simple as possible, and the S/N would be improved only over a narrow wavelength region where the transmittance in both directions (blue and red) is similar. Instead, the pipeline calculates using the response curves a junction point where both spectra are similarly attenuated by the dichroic and joins the spectra there. Note that when the 600/4000 blue grism is used or if the red grating angle is set to have minimal overlap, this point may be at the very limit of the blue CCD or the red CCD, respectively.

This is the final step in the pipeline. The result is a science-quality spectrum extending from the blue limit of the blue camera to the red limit of the red camera. Note that the wavelength pixel scale and the resolution will be discontinuous at the junction point because no resampling is performed. It is written to a text file that also includes columns for the sky spectrum, the variance spectrum (uncertainty), and the x pixel and y profile centroids. Each file is written in ASCII format for maximum portability and includes extensive meta-data describing all aspects of the instrument configuration and pipeline steps and settings.

## 4. ADDITIONAL FEATURES

### 4.1. *Imaging pipeline*

An imaging pipeline is also included. Its early steps basically parallel the spectroscopic pipeline (overscan subtraction, flat-fielding, cosmic ray correction) using imaging-optimized routines. From that point the behavior diverges completely as it performs astrometric and photometric registration using catalogs. The ability to solve for and apply airmass solutions via spectrophotometric standards is also included - although this feature is rarely needed now after the release of the Pan-STARRS 3pi catalog, which covers almost any field visible from Keck. (It is still important for *U*-band imaging.) Solving for field distortions is not well-tested. A notable optional feature is the ability to create super-sky correction flats using specific blocks of imaging to remove dust features that have been observed to change positions as a result of changes to the red-side grating angle. Unlike the spectroscopic pipeline (which is entirely written in IDL), the imaging pipeline has some important external dependencies - notably it requires SExtractor, SWarp, and a custom astrometric routine written in Python 2.4. The imaging pipeline will be described in detail in future work.

#### 4.2. Notes about historical LRIS data

LRIS-B has employed two detectors over its lifetime: an engineering-grade SITE 2048x2048 CCD that was used for a brief period after commissioning in 2000, and the science-grade CCD that was installed in June 2002 and remains in use. Reductions of observations from this detector are not supported. The properties of the current instrument are quite stable, aside from occasional motions after instrument observing that have changed the (otherwise fixed) wavelength range of each grism slightly (see §3.12). The pipeline should be able to reduce LRIS-B data from any point in its history after the modern detector was installed.

LRIS-R has been equipped with three red detectors in its lifetime. The most recent detector has been in use since December 2010 and has been stable over this period. The two earlier detectors require some discussion.

The first red detector on LRIS was a standard thin CCD, employed from commissioning in the early 1990s until May 2009. This detector differs from the current detector in many ways, including size (2048×2048), pixel scale (0.210 instead of 0.135), quantum efficiency (much worse in the near-IR), and cosmic ray sensitivity (lower, i.e. better). It is a single CCD with two amplifiers (rather than two CCDs with four amplifiers total) so there is no chip gap and no need to 'split' the data. It suffers from strong fringing in the near-IR, removal of which is extremely challenging. (On-sky halogen flats and/or nodded pair subtraction are required for spectroscopy; fringe frames are required for long-wave imaging.) The pipeline fully supports LRIS R-1 data, but no fringe correction is applied, and the capabilities have not been well-tested. Caution should be exercised in using the pipeline to process LRIS-red data prior to mid-2009.

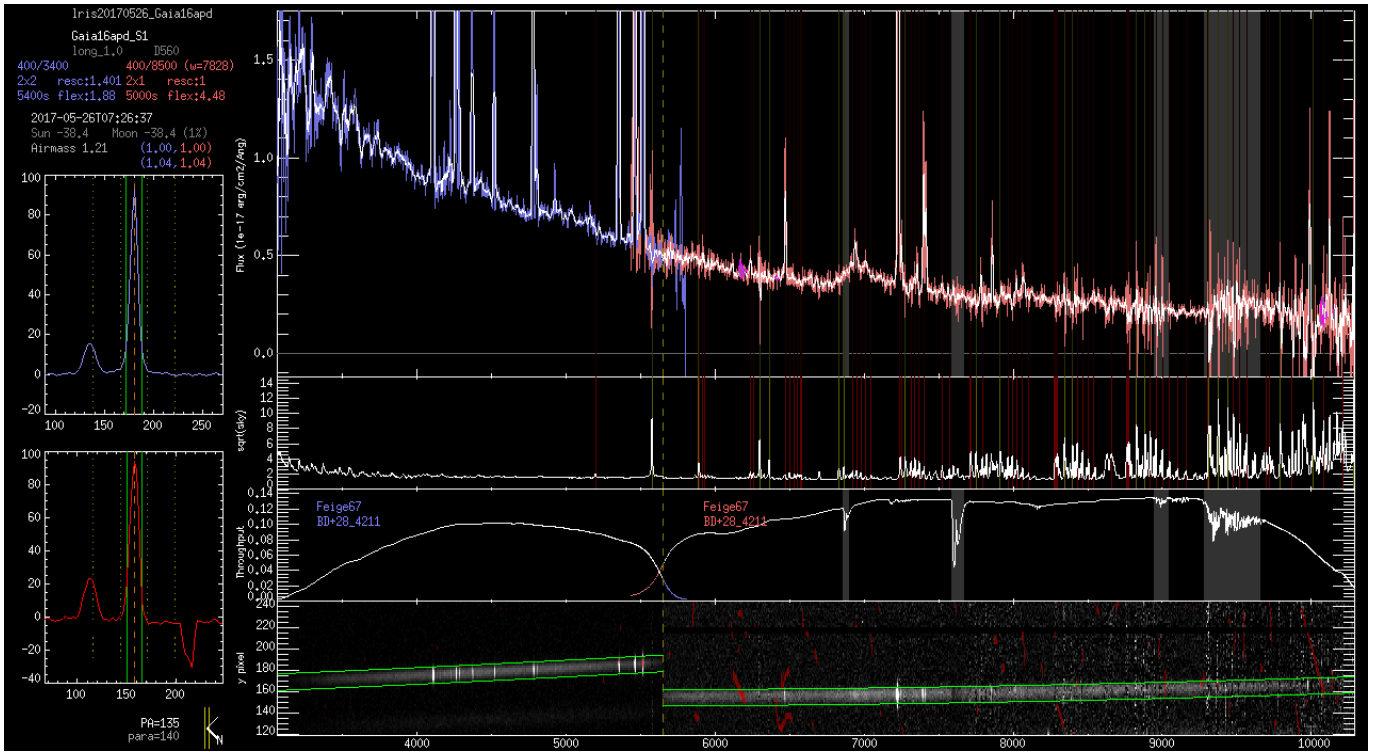
Very early observations (in particular, observations from before the blue channel was added in 2003) lack some header metadata needed by the pipeline and also use different calibration lamps. These observations are unlikely to be successfully processed by the pipeline at all, although I anticipate adding support for LRIS data dating this period eventually.

The second red detector was equivalent to the current one in most respects (size, pixel scale, layout, wavelength-dependent sensitivity, etc.). It exhibited some curious noise and sensitivity patterns, including a significant noise enhancement close to the chip gap that rendered data close to the gap almost unusable. A much bigger note of concern is that after a few months of installation (beginning approximately September 2009) this CCD began exhibiting CTE problems that are clearly visible in the form of one-directional trailing out of cosmic ray hits and sky lines. Initially these affected only some amplifiers but soon appeared throughout the detector, and the issue was solved only by complete replacement of the detector in December 2010. The pipeline will process data from this period through to the end as the CCD format is the same as the current detector, but there is no attempt made to correct for the CTE issues. Results should be interpreted with caution.

#### 4.3. Interruption

The pipeline is fully interruptible. The user can Ctrl+C out of the reductions at any point and effectively resume the pipeline by repeating the same command it was started with. The pipeline does not keep track of where it left off internally, but it can recognize when a step has already been performed based on the pre-existence of the output file it is planning to write—in these cases it will skip over that step. To instruct the pipeline to skip over steps entirely, the user can state which stage in the processing sequence to begin (and, if desired, end) with.

Frequently it is desired to make changes and repeat steps (in particular with regard to aperture placement, see above). By default, if apertures are changed and the pipeline is rerun, nothing will happen because this would require overwriting files generated by the old, automatically-determined apertures. The pipeline can be set to overwrite old files by setting the `"/redo"` option. This by default will redo everything from the beginning and overwrite any pre-



**Figure 6.** Screenshot of the validation tool, showing spectra of an object and the sky, along with details of the extraction and flux calibration (throughput curve). Telluric absorption bands, sky lines, and excised cosmic rays are also marked for reference.

existing files, but can be restricted to specific steps (see above) or specific targets or exposures. The pipeline manual contains instructions and examples for this type of operation.

#### 4.4. Identification of bad data

The pipeline can recognize 'bad' data in a variety of circumstances. Flats with the lamp off or very low counts, observations with the shutter closed, saturated dome or twilight flats, saturated standards, and similar mistakes are identified via either the header keywords or counts in the image and once identified the pipeline will not conduct further processing. The MIRA files used for segment alignment are also recognized automatically. There is no need to formally delete these frames. Indeed, it may be helpful to retain them to be sure the pipeline is able to faithfully record when instrument components were moved.

A frequent, minor observer error at the telescope is to forget to change the object name (written to the FITS headers under the OBJECT keyword) after slewing to a new source. The pipeline does not depend on the OBJECT name in any formal sense: instead it uses TARGNAME which comes from the starlist uploaded at the beginning of the night. (It does look for words like 'sky' and 'flat' to help distinguish twilight sky imaging frames in ambiguous cases but they are not required in any sense.) However, if the telescope position and target are changed with no change in OBJECT name, the OBJECT name of subsequent entries will be set to the TARGNAME.

A known recurrent LRIS hardware error is the failure of the blue shutter. In these cases the trapdoor is usually used as an alternate shutter. Data taken under this scenario does not require any special handling, although any short flux-calibration exposures will be problematic due to the slower shutter speed and these should probably be omitted by the user to ensure they are not used for calibration.

Alternatively, some observers may elect to acquire a 0-s exposure after each image to calculate the signal pattern that accumulated during readout. A tool exists to recognize these exposures, match them with their companion images, and subtract to obtain corrected images, but this is not included in automatic pipeline operations and data taken in this mode may not be scientifically useable if processed using the fully-automated pipeline.

#### 4.5. Logging and diagnostics

**Table 1.** Example processing times

	Run 1	Run 2
Date	2016-09-09	2017-07-15
$N_{\text{cal}}$	18	119
$N_{\text{sci}}$	66	56
binning	1x1	2x1/2x2
<b>Step</b>	<b>Time (s)</b>	
prepare	92	86
makeflat	17	9
flatten	26	9
split	4	2
crclean	595	217
skysubtract	557	175
sum	41	19
trace	20	7
extract	43	18
wavcal	3	3
wavapply	25	26
response	5	2
fluxcal	7	6
connect	5	4
total	1438	583

The pipeline reports what it is doing to the terminal as it proceeds. Almost all of this information is also written to a log file. This includes warning messages (including “reminders” to get data more carefully next time if observations were not conducted in a way considered to be optimal) or error messages if an important file cannot be processed for an unexpected reason such as a wavelength solution failure.

A large number of check plots and summary plots are also written that document the wavelength solution and flexure correction, response curves, flux calibration (including telluric correction), tracing, and so on. DS9 region files containing the output apertures are also written to disk and can be loaded to permit interactive user exploration of the tracing region.

Perhaps most usefully for the end scientific user, the pipeline also includes a validation tool that loads and plots an individual spectrum along with its corresponding sky spectra (for wavelength checking), throughput curves (for flux-calibration checking), and 2D spectra with overplotted trace curves as well as collapsed spatial profiles. Pixels affected by cosmic rays or bad columns are marked. With this tool the user can browse through spectra and identify any problematic results before they are ingested into online databases or used for scientific analysis. The tool is demonstrated in Figure 6

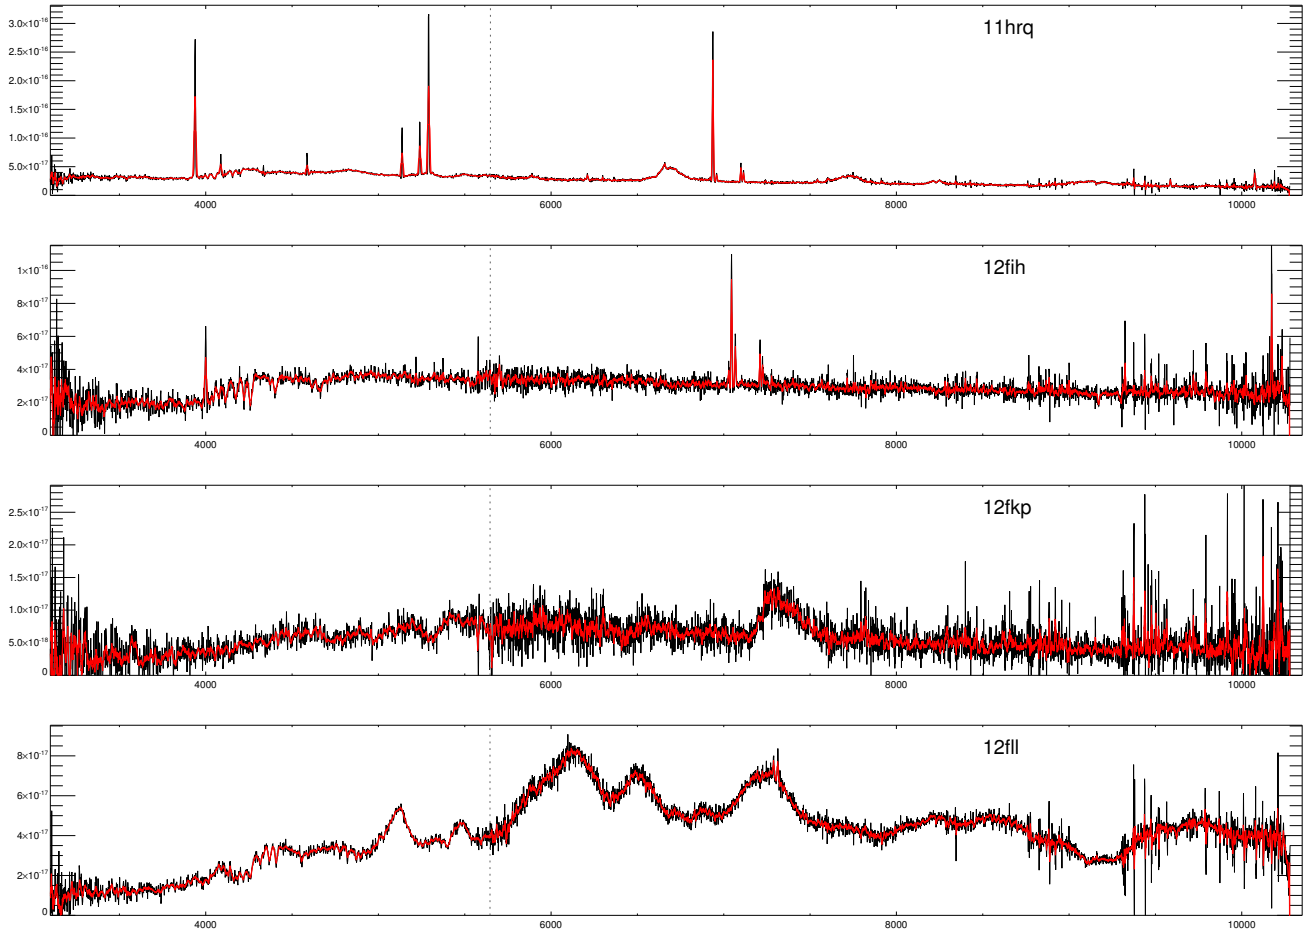
## 5. PERFORMANCE

The processing time is dominated by the cosmic ray and sky-subtraction steps. Some example run times for two recent observing nights are shown in Table 1. The pipeline is much slower when processing unbinned data compared to binned data.

Reducing imaging data is substantially slower (nights primarily or entirely taken in imaging mode can take over an hour to reduce), mostly because of the larger number of files involved (shorter exposures).

Data products (e.g., Figure 7) are science-quality, as long as observations satisfy the basic requirements and the aperture background regions are reasonably “clean” (e.g. free of host-galaxy flux). Use of optimal extraction could improve S/N somewhat, and the flux-calibration accuracy is typically limited by how clean the spectrum of the flux-standard was or how much red-blue overlap is provided by the setup. Given the difficulty of accurately flux-calibrating (and occasionally wavelength-calibrating) over the dichroic overlap region at the CCD edges it is not uncommon to observe small jumps or discontinuities in this area in final spectra, and users should be aware that there is always





**Figure 7.** Examples of data (spectra of supernovae or supernova hosts from an observing run on 2012 July 15) processed by the pipeline.

intrinsic uncertainty in the red-blue scaling for faint objects. Other non-trivial problems are much rarer, although subtle errors can occasionally be introduced at all stages.

Most apparent failures in processing are associated with missing calibration files (or equivalently, bad calibration files since these are generally recognized and not processed). Such issues can of course be solved by acquiring the relevant calibrations: at the telescope (if the pipeline is being employed during real-time reductions of an observing run), or by searching for an appropriate set of calibration files in the public archive and placing them in a pipeline-accessible reduction directory. Future pipeline releases will include a set of archival calibrations for standard observing modes, as well as include an option to continue with no flat-fielding.

## 6. LRIS ARCHIVE REDUCTION

To demonstrate the efficiency of the pipeline in practice and in the interest of making scientific data accessible to the worldwide community, we downloaded all data from the most recent publicly-released semester in the LRIS archive (2017A, which spans from February to July 2017). The data were sorted by observing night, and the pipeline was scripted to run in fully automated mode in each nightly directory.

To allow for the possibility that flats might be obtained on nights different from the science observations, we instructed the pipeline to search the reduction directories from other nights if flat-field files needed for science calibration could not be found in the nightly directory. No other modifications were made, and we did not supervise the reductions while they were in progress and did not inspect or delete any input files, relying entirely on the pipeline to remove irrelevant or problematic calibrations. Results were examined (using the validation tool) only following the completion of all processing.

Some spectra could not be processed due to a lack of any useable calibrations. Beyond this, only a few of the successfully-written output spectra exhibited any notable problems. A small number of red spectra acquired in 2x2 binning mode had some emission lines rejected as cosmic rays, and other minor issues (e.g. small jumps at the dichroic junction) were occasionally evident, but about 90% of the spectra were classified as scientifically useable for most purposes. Two nights suffered from issues with flux calibration: one as a result of an I-band filter being stuck in the red filter wheel during a standard-star observation, and another night during which the blue was severely attenuated for still-unknown reasons.

All validated spectra have been made available online<sup>14</sup>. In the near future we will expand our reductions back at least as far as the commissioning date of the upgraded red detector (2009), and provide some modest additional user supervision to salvage any observations that could not be optimally processed in unsupervised mode due to missing or problematic calibrations. These spectra will be made permanently available as a user-provided data set at the official Keck archive<sup>15</sup>.

Users interested in making use of these spectra for scientific analysis should be mindful of the limitations discussed earlier. In particular, emission-line sources and extended objects (or point sources coincident with extended objects) may be missing some flux, the S/N will be lower than what would be available from an optimum extraction, and there may be issues in extraction of faint targets at the far blue/red ends or in flux calibration across the dichroic junction. Observing runs in which no high-quality standards were observed may also show flux calibration deviations of up to 20% due to standard-star line residuals (especially if reference spectra for the relevant standards are poor), and high-airmass objects will show telluric residuals if standards were not obtained at similar airmass and seeing conditions. Users should also beware that LRIS does not have any straightforward method of robustly tracking slit placement (e.g., slit-viewing images): there is no guarantee that the target specified by the user (and given as the TARGNAME or OBJECT header value) was accurately placed on the slit or extracted by the pipeline.

Users interested in detailed scientific analysis of our public reduced spectra are advised to download the pipeline and re-reduce any relevant observations (potentially using different settings if, e.g., accurate extraction of extended sources is needed).

## 7. CONCLUSIONS

In this paper I have demonstrated a robust, effective, efficient pipeline for fully automated processing of LRIS observations. The pipeline is generic to all configurations of the longslit and requires no user input, except to control fine details of aperture extraction. While highly customized for managing the many complexities of LRIS (flexure, separate red+blue observations, etc.), the design of the pipeline is generic and could easily be extended to other facilities.

The pipeline is currently available at <http://www.astro.caltech.edu/~dperley/programs/lpipe.html> and will also be made available on github and other package repositories in the near future.

This research has made use of the Keck Observatory Archive (KOA), which is operated by the W. M. Keck Observatory and the NASA Exoplanet Science Institute (NExScI), under contract with the National Aeronautics and Space Administration. DAP would like to thank the observatory and the archive for their support of the publication of this paper, as well as the LRIS support astronomers (in particular, Greg Wirth, Marc Kassis, and Luca Ricci) for their advice and assistance during many previous observing runs.

<sup>14</sup> The provisional location is: <http://www.astro.ljmu.ac.uk/~aridperl/lris/archive/>

<sup>15</sup> <https://koa.ipac.caltech.edu/>

## REFERENCES

- Belli, S., Contursi, A., & Davies, R. I. 2017, ArXiv e-prints, arXiv:1710.05924
- Bohlin, R. C. 1996, AJ, 111, 1743
- Bohlin, R. C., Colina, L., & Finley, D. S. 1995, AJ, 110, 1316
- Buton, C., Copin, Y., Aldering, G., et al. 2013, A&A, 549, A8
- Goodrich, R. W., Cohen, M. H., & Putney, A. 1995, PASP, 107, 179
- Hamuy, M., Suntzeff, N. B., Heathcote, S. R., et al. 1994, PASP, 106, 566
- Hamuy, M., Walker, A. R., Suntzeff, N. B., et al. 1992, PASP, 104, 533
- Horne, K. 1986, PASP, 98, 609
- Kelson, D. D. 2003, PASP, 115, 688
- Kulkarni, S. R. 2016, ArXiv e-prints, arXiv:1606.06674
- McCarthy, J. K., Cohen, J. G., Butcher, B., et al. 1998, in Proc. SPIE, Vol. 3355, Optical Astronomical Instrumentation, ed. S. D’Odorico, 81–92
- Oke, J. B. 1990, AJ, 99, 1621
- Oke, J. B., & Gunn, J. E. 1983, ApJ, 266, 713
- Oke, J. B., Cohen, J. G., Carr, M., et al. 1995, PASP, 107, 375
- Phillips, A. C., Miller, J., Cowley, D., & Wallace, V. 2006, in Proc. SPIE, Vol. 6269, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 62691O
- Rockosi, C., Stover, R., Kibrick, R., et al. 2010, in Proc. SPIE, Vol. 7735, Ground-based and Airborne Instrumentation for Astronomy III, 77350R
- Steidel, C. C., Shapley, A. E., Pettini, M., et al. 2004, ApJ, 604, 534
- Stone, R. P. S. 1977, ApJ, 218, 767