

Data reduction pipelines for the Keck Observatory Archive

H. D. Tran^{*a}, R. Cohen^a, A. Colson^a, J. A. Mader^a, M. Swain^b, A. C. Laity^b, M. Kong^b, C.R. Gelino^b, G. B. Berriman^b

^aW. M. Keck Observatory, 65-1120 Mamalahoa Hwy., Kamuela, HI, USA 96743; ^bNASA Exoplanet Science Institute, Mail Code 100-22, 770 South Wilson Ave., Pasadena, CA, USA 91125

ABSTRACT

The Keck Observatory Archive (KOA) currently serves ~ 42 TB of data spanning over 20 years from all ten past and current facility instruments at Keck. Although most of the available data are in the raw form, for four instruments (HIRES, NIRC2, OSIRIS, LWS), quick-look, browse products generated by automated pipelines are also offered to facilitate assessment of the scientific content and quality of the data. KOA underwrote the update of the MAKEE package to support reduction of the CCD upgrade to HIRES, developed scripts for reduction of NIRC2 data and automated the existing OSIRIS and LWS data reduction packages. We describe in some detail the recently completed automated pipeline for NIRSPEC, which will be used to create browse products in KOA and made available for quicklook of the data by the observers at the telescope. We review the currently available data reduction tools for Keck data, and present our plans and anticipated priorities for the development of automated pipelines and release of reduced data products for the rest of the current and future instruments. We also anticipate that Keck's newest instrument, NIRES, which will be delivered with a fully automated pipeline, will be the first to have both raw and level-1 data ingested at commissioning.

Keywords: Data archive, data reduction pipelines, DRP, KOA, W. M. Keck Observatory, NExSci

1. INTRODUCTION

The Keck Observatory Archive (KOA)^{1,2,†} began operations in August 2004, initially serving raw, level-0 (uncalibrated) data from the High-Resolution Echelle Spectrometer (HIRES). It now serves data from all ten active and retired WMKO facility instruments. Next to more instruments, the single item most often requested by KOA users is more reduced, possibly publishable-quality data in the archive. Currently, for all but four of the instruments KOA serves only level-0 data. We offer for HIRES, 1-D extracted, wavelength calibrated spectra, for the Near-Infrared Camera (NIRC2), distortion corrected images with detector signature removed, for the integral field unit OH-Suppressing Infra-Red Imaging Spectrograph (OSIRIS), sky-subtracted reduced data cubes of two-dimensional array as a function of wavelength, and for the Long Wavelength Spectrometer (LWS), reduced co-added 2-D images from 6-D arrays. Produced by fully automated data reduction pipelines (DRPs), these calibrated/reduced or level-1 products are intended for browsing purposes only, and not meant to be used for scientific publication. However, even as “quick-look” browse products, they can give valuable insight into the quality and scientific content of the data. The accompanying DRPs then offer users a tool for them to perform more careful data reduction ideally tailored to their observations and to support quick-look reductions at the telescope.

Producing level-1 data requires a DRP that is fully automated so that so that the archive operations can remain seamless, without manual human interaction. KOA has demonstrated the effectiveness of automating existing interactive data reduction packages or tools, often developed by the instrument team or the observers themselves. This approach was used to develop level-1 pipelines for HIRES and OSIRIS. For HIRES, KOA underwrote an adaptation of the *MAKEE* pipeline to process data obtained with the 3-CCD camera upgrade. A similar approach was used for OSIRIS, NIRC2, and LWS. KOA is currently completing the development of a Python-based pipeline for the Near-Infrared Echelle Spectrograph (NIRSPEC).

We plan to use the same model to make level-1 products available for all the operating modes of all Keck instruments.

^{*}htran@keck.hawaii.edu

[†]<http://koa.ipac.caltech.edu/>

Many data reduction tools have been made available for nearly all of the instruments (see e.g., <http://www2.keck.hawaii.edu/koa/public/drp.html>). We are developing ways and assessing efforts to adapt and automate these reduction packages with a view to developing a delivery schedule. New instruments (such as NIRES) are required by Keck to be delivered with a DRP, and KOA will cooperate with developers to ensure that an automated pipeline will be archive-ready from the first day of operation.

In this paper we provide an overview of the pipelines currently available to produce level-1 data for KOA. We will go into some detail on the development of the NIRSPEC DRP which KOA has been developing for KOA, and will provide a roadmap for reaching our goal of having a DRP and level-1 data for all of the current and future instruments in KOA.

2. CURRENT PIPELINES AT KOA

Currently, there are four instruments with fully automated DRPs producing quick-look level-1 data in KOA. We have just completed a DRP for a fifth instrument, NIRSPEC, and will discuss it in more detail here. We anticipate that KOA will begin producing and ingesting level-1 data for NIRPEC sometime this summer.

2.1 HIRES

The HIRES pipeline is the modification of a more general purpose DRP called *MAKEE* (<http://www.astro.caltech.edu/~tb/makee/>). *MAKEE* was upgraded to support a new 3-CCD detector, and is in wide use by the HIRES community, producing over 100 citations.

Using the *MAKEE* pipeline, KOA produces the extracted one-dimensional spectra of counts vs. wavelength for each order of the 2-D raw frame. These extracted spectra have been bias-subtracted, flat-fielded, sky-subtracted, and wavelength-calibrated. We also provide diagnostic products for assessing the quality of the automated spectral extractions, such as signal-to-noise ratio estimates, sky spectra, flux profile along slit, and centroids of the profiles, all as a function of wavelength. KOA also provides a qualitative “pass” or “fail” grade for each of the extractions, along with an interactive plotting tool to facilitate and enhance the browsing experience of these extracted spectra from the results page of the user interface.

2.2 NIRC2

KOA developed a simple DRP that provides the basic calibration of NIRC2 imaging data. The script determines the appropriate dark and flat calibration files from information provided in the data headers, but does not identify a suitable sky frame for background subtraction or co-add multiple exposures. Nonetheless, by correcting for bad pixels and removing the detector signature, the pipeline creates an improved product that facilitates the assessment of the data. If available, a master flat and/or dark is used. If not, the median of at least one and a maximum of 10 flats and/or darks found by the calibration association will be used. In addition, KOA determines and adds world coordinate system (WCS) information, which did not exist in the original header, and removes optical distortion for full-frame images.

To help users get a better assessment of the data, KOA recently released an interactive visualizer for NIRC2 calibrated images that supports functions such as panning, zooming, and adjusting the image intensity stretch. This greatly improves the browsing experience over the previously available static jpeg images.

2.3 OSIRIS

OSIRIS is an integral field spectrograph behind AO that obtains a grid of spectra over a small field of view. Although raw data are stored in a 2-D FITS file, it is very difficult to assess the quality or science content of the data. A DRP exists (<http://www2.keck.hawaii.edu/inst/osiris/tools/>) that processes the raw data into a data cube that provides spatial and wavelength information. The DRP consists of a set of scripts that can be executed either via GUI or command line. The online mode provides quick real-time reductions which are used primarily for basic data visualization and quality assessment done at the telescope while observing. The offline version includes an expanded reduction options, does more iterations for a better construction of the data cubes, and is used to produce publication-quality products. KOA has adopted this pipeline to produce quicklook level-1 data for KOA. It needs a parameter file providing the list of calibration frames, including “rectification matrices”, and processing options, and in most cases we are able to determine these calibration files for a particular file directly from information in the FITS header. An important step in fully automating the OSIRIS DRP for KOA is the identification of a suitable sky frame for background subtraction. This process has been described in Holt et al. (2014), which can be consulted for more details.

OSIRIS has just undergone an upgrade of its SPEC detector. The DRP is currently in the process of being updated to support data from the new detector. Also, new rectification matrices (“recmats”), which provide the required precise mapping of the flux between each pixel in the raw image and the reduced cube, need to be created. Normally, these recmats are very stable, and there’s no need to update them unless there’s a significant configuration change or service to the instrument (i.e., the dewar gets opened). There is also a current issue with the DRP, where the flux from the raw frame is being mis-assigned to the wrong lenslets, resulting in the presence of light “echos” in wavelength space of the reduced data cube, most noticeably in high-contrast or bright objects. This issue is actively being investigated.

As for NIRC2 calibrated data, KOA provides interactive viewer to help users browse the reduced OSIRIS data cubes, with functions to pan, zoom, select wavelength slices to display, and make depth plots of object flux counts as a function of wavelength for any selected spatial regions.

2.4 LWS

Raw LWS imaging data are 6-dimensional data files in which entire observing “chop-nod” sequences (image repeats, small dithers, sky frames) are stored. A quicklook representation of the raw data is not possible, so we process the data to produce a single, co-added 2-dimensional image. The basic code to create this product already existed (<http://www2.keck.hawaii.edu/inst/lws/lws-idl.html>). KOA modified it to return a FITS file rather than a floating point array.

3. THE NIRSPEC DRP

3.1 NIRSPEC Instrument

NIRSPEC (McLean et al. 1999) is an all-reflective, near-infrared, high-resolution spectrograph designed to operate over the wavelength range 0.95 to 5.4 microns. A cross-dispersed echelle grating is used to obtain this wavelength coverage at high spectral resolution. A low-dispersion mode is also available using the cross-disperser, giving a spectral resolution of ~ 2,200. NIRSPEC is a vacuum-cryogenic instrument operating at a temperature of about 60K. Although not optimized for adaptive optics (AO), NIRSPEC can be located behind the AO system. NIRSPEC has two detectors, a spectrographic indium antimonide (InSb) detector with 1024x1024 pixels and a slit-viewing mercury-cadmium-telluride (HgCdTe) infrared array with 256x256 pixels and a cutoff wavelength of 2.5 micron.

Approximately 2000 nights of data have been gathered by NIRSPEC since commissioning in January 2000, producing roughly 165k science frames with the SPEC detector. High-resolution mode makes up approximately $\frac{3}{4}$ of this data.

3.2 Available Reduction Packages for NIRSPEC

There are a few tools available to reduce NIRSPEC data:

- Developed by the instrument team at delivery, the REDSPEC package is comprised of a set of IDL procedures. It performs the rectification of array data in which both spatial and spectral corrections are applied, and provides a basic spectral extraction which yields a text file containing wavelength and intensity. The wavelength calibration step requires highly interactive identification and verification of arc lines.
- The WMKONSPEC package is a suit of simple IRAF tasks related to the reduction of low-dispersion spectra.
- The Kelsonware package is written in Python and is one of the more complicated packages to use since it is more a general collection of routines that can be used for a large number of different instruments.

All three of these tools rely heavily on user interaction during the data reduction process. Although it is desirable to re-use packages that are already available, the archiving of Keck data is a fully automated process and thus this interaction is not desirable. Therefore, the KOA team has developed a new automated NIRSPEC DRP.

3.3 The NIRSPEC DRP (NSDRP)

Written in Python, the NIRSPEC DRP (NSDRP) is designed to be fully automated to run without any user interaction. In addition, it is intended to produce only quick-look, browse extracted spectra for the purpose of assessing the scientific quality and content of the raw data. It may not be able to reduce all types of NIRSPEC data. Specifically, it currently can only reduce data obtained in high-dispersion mode, and is optimized for and works best on frames having:

- single point source in slit
- reasonably bright targets with detectable continuum
- NIRSPEC-1 through NIRSPEC-7 filters
- well-separated orders without overlapping
- sufficient exposure times (≥ 30 s) with detectable sky lines

Future updates to the pipeline may add new capabilities and improvements that can handle other cases. The NSDRP can be run in batch or “KOA mode”, which processes all files for a whole night at once, or “user” or command-line mode for a single frame, which is useful for observers at the telescope. In KOA mode, the DRP is designed to process all object frames taken on any given night of observing. KOA utilizes a calibration association algorithm to find the dark and flat frames required for the reduction process. The calibration association is performed during the level-0 (raw) data processing and a new keyword (IMAGETYP) is added to each FITS header specifying the type of observation. The DRP uses this keyword along with the instrument configuration keywords to determine which calibrations to use for the object frame. In this mode, the NSDRP should be run only on files downloaded from KOA since it needs to know the correct image type (on-sky object or calibrations) of the data frames. Only NIRSPEC files ingested into the archive have gone through the data evaluation and preparation process^{5, 6} and are guaranteed to contain this information. The latest version of the pipeline can be obtained from Github at https://github.com/2ichard/nirspec_drp.git.

3.3.1 The Design

The central tasks of NIRSPEC DRP are to locate and extract individual spectral orders from object frames and flats, rectify orders in the spatial and spectral dimensions and combine OH emission sky lines identified in each order to find a global (to the frame) wavelength calibration solution.

Flat field frames taken with the same filter, slit and echelle and cross disperser angles as the object frame are required for data reduction. Flats are used not only to account for intensity variations in pixel-to-pixel sensitivity but also to locate the spectral orders on the detector.

Figure 1 shows a raw object frame of a point source taken with the NIRSPEC-3 (J) filter and a corresponding flat. These images are histogram equalized. Note the many sky lines that uniformly illuminate the slit in the object frame. If dark frames with the same exposure time as the object frame are available they are used in dark frame subtraction to minimize fixed-pattern noise.

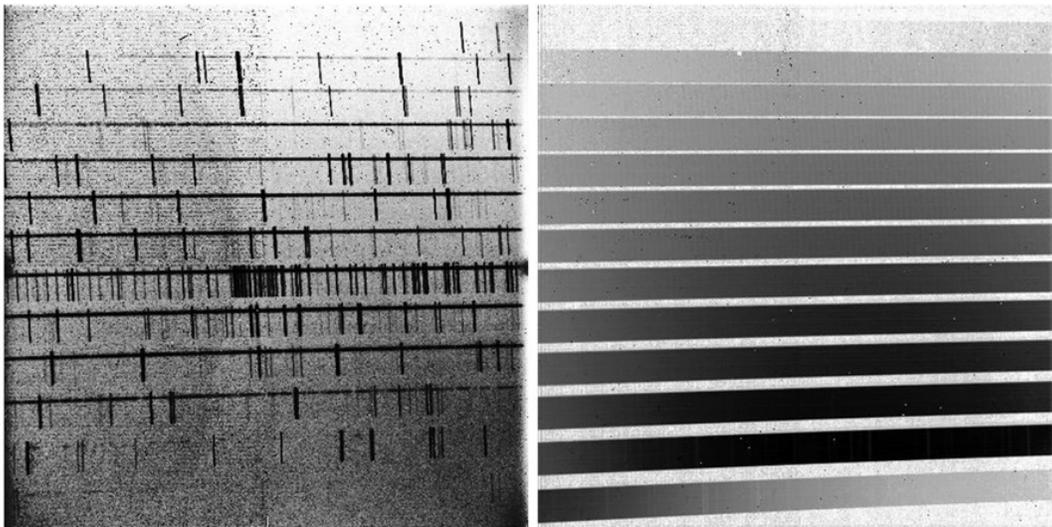


Figure 1. Object frame (*left*) and flat field in J.

Extended objects are not well accommodated by the DRP though the program will produce output as long as there is a discernable peak in the spatial intensity profile. For successful wavelength calibration, exposures long enough to yield discernable OH sky lines are required. In the absence of sky lines sufficient in number and intensity, wavelength calibration is not possible and wavelength scales are approximated based on the grating equation. Arc lamp frames for wavelength calibration are not supported by the DRP at this time but this has been identified as an area for future development.

The NIRSPEC DRP can reduce data collected with NIRSPEC filters 1 through 7 and a wide range of echelle and cross disperser angles and all of the high resolution slits. It makes extensive use of common python libraries including numpy, astropy, scipy and matplotlib. The software is written mainly in the procedural style though the main application-specific data structures are implemented as Python classes.

The KOA data preparation and assessment process (DQA) invokes NSDRP automatically. DQA supplies the DRP with input and output directory paths as command line arguments. In addition to level 1 data products, two types of log files are produced, one to record per-night summary data and another to record details of data reduction on a per object frame basis. Figure 2 presents an overview the DRP.

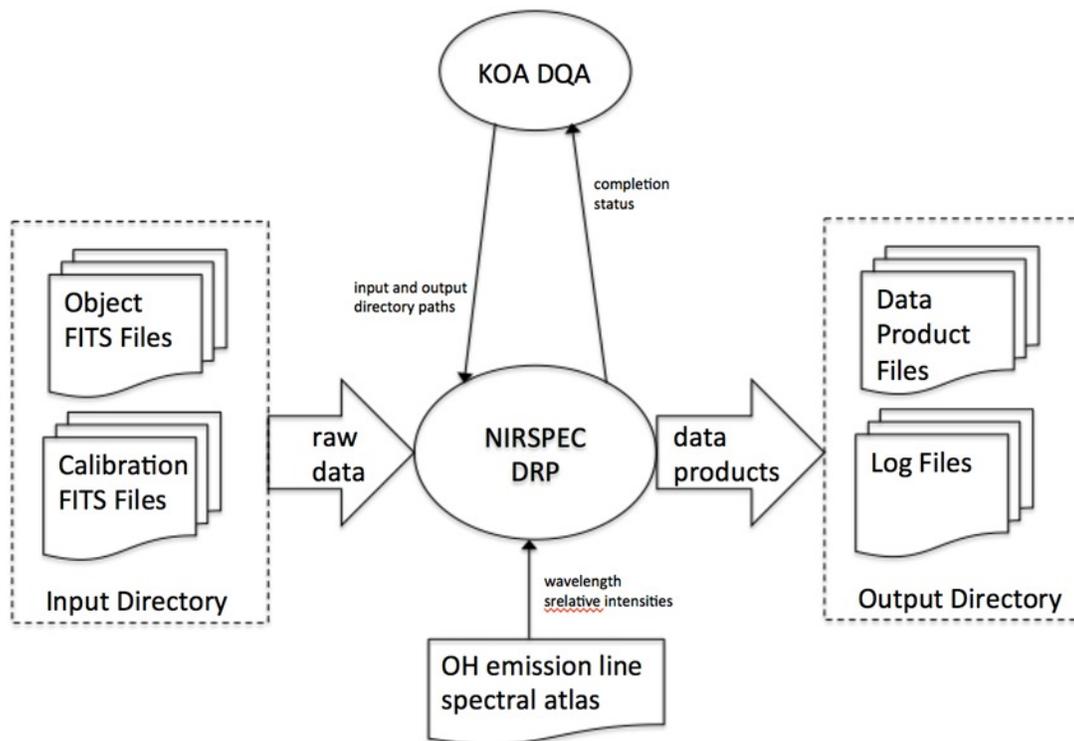


Figure 2. NSDRP System Context Diagram

Written entirely in the Python programming language, NSDRP has a modular design and uses four central application-specific data structures implemented as Python classes. This design choice simplifies construction and management of multiple simultaneous instances of these classes. The four central classes are:

1. RawDataSet – Represents an associated set of raw image files—object frames, darks and flats.
2. ReducedDataSet – Represents the results of an object frame reduction. Has fields representing reduction results pertaining to the frame as a whole as well as a list of Order objects that represent reduction results for individual orders.

3. Order – Represents reduction results for an individual order. Has fields representing reduction parameters pertaining to the order as a whole and a list of Line objects representing identified sky lines.
4. Line – Represents an observed sky line that has been associated with an OH emission line of known wavelength.

In addition to object files, flats and optionally dark frames, the only other input data required by the DRP is a file containing an atlas of OH emission lines in the near infrared⁷.

The NSDRP processes the data using the following steps:

- **File association**

The DRP first scans the source directory and identifies object frames that can be reduced. For each object frame that can be reduced, `create_raw_data_set.create()` associates object frames with corresponding calibration frames and constructs a series of `RawDataSet` objects. Each raw data set consists of an object frame, one or more flat field frames and zero or more darks.

- **Remove instrument signatures**

An outer loop iterates through each raw data set and calls `reduce_frame.reduce_frame()` to reduce each object frame. A `ReducedDataSet` object is instantiated by `reduce_frame()` and all reduction results for the frame are stored in this object. If there are multiple darks or flats they are median combined. Combined darks are subtracted from object and combined flat. Cosmic ray artifacts are removed using a Laplacian edge detection algorithm.

- **Locate orders**

The grating equation is used to determine which orders are expected to be on the detector. For each order expected to be on the detector, `reduce_order.reduce_order()` is called to instantiate an `Order` object. The orders are located using the grating equation of the form

$$\lambda(c, \theta, \phi, M) = \frac{k_1 \sin(\theta) + k_2 \left(\frac{w}{2} - c\right) \cos(\phi)}{M}$$

where θ is the echelle angle, ϕ is the cross disperser angle, and M is the order number. k_1 and k_2 are constants the values of which have been determined empirically and w is the width of the detector in pixels along the dispersion axis. Corrections are applied to the expected wavelengths by shifting the wavelength scale by an order-dependent constant value. Corrections are also applied to calculated order positions by adding slit- and filter-dependent constants.

- **Extract and rectify orders**

Order processing involves cutting out a rectangular region containing the order from the full frame. Using spline interpolation, rectification in the spatial dimension is accomplished by tracing order edges in flat frames, and in the spectral dimension by tracing sky lines along the slit. The flat image used to find the lower and upper bounds of the order by tracing the peaks of the difference image of slightly shifted flat images. These order traces are used to extract the object spectrum.

- **Determine wavelength solutions from sky lines**

For each order, wavelength calibration is done using known OH emission line wavelengths and relative intensities. First, a synthetic sky spectrum is generated from a list of the known OH lines. This synthetic spectrum is then cross-correlated with the observed sky spectrum to determine the wavelength shift. Figure 3 illustrates this process for one order as an example.

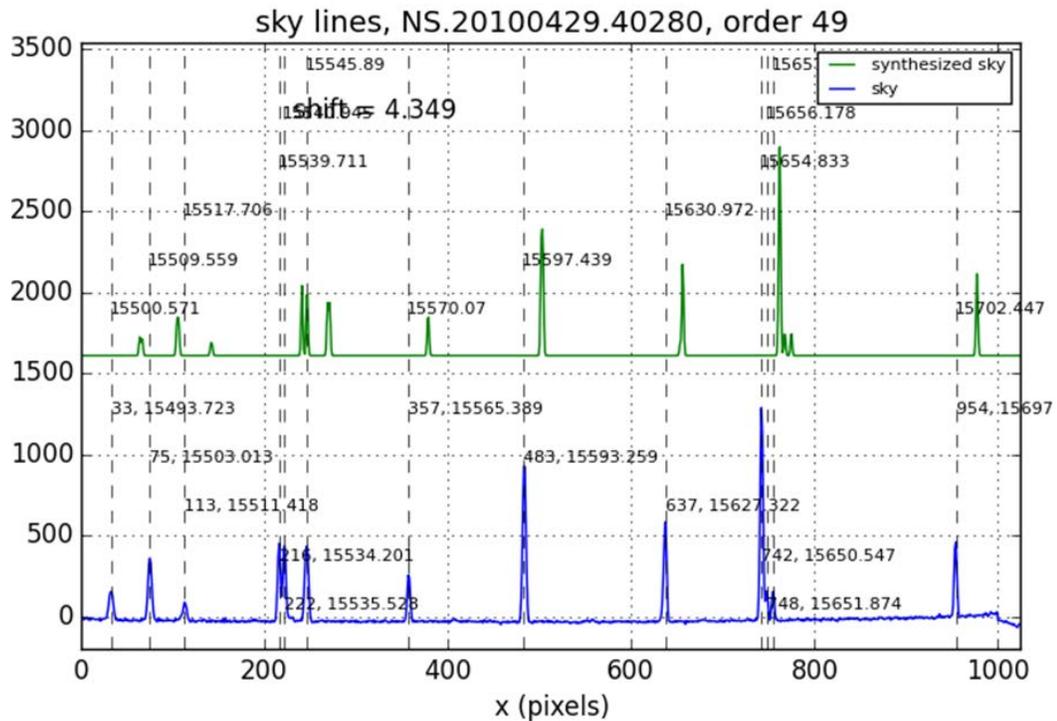


Figure 3. Sky line identification by cross-correlation synthetic OH lines with observed sky spectrum

When all orders have been reduced and sky lines identified, `reduce_frame.find_global_wavelength_soln()` is called to find coefficients to the wavelength equation. Identified lines from all the orders are used in a 2-D fit of wavelength as a function of order, column location of order to known wavelengths. A two-dimensional polynomial fit is used for wavelength calibration:

$$\lambda(c, M) = r_0 + r_1c + r_2c^2 + r_3M^{-1} + r_4cM^{-1} + r_5c^2M^{-1}$$

where c is the column pixel, M is the spectral order and the coefficients $r_0..r_5$ are determined by the method of least squares. This wavelength solution is then applied to all spectral orders.

If too few sky lines are identified in the orders or the global fit is poor, the wavelength scale falls back to the theoretical solution from the grating equation.

- **Extract 1-D spectra**

The rectified order is collapsed along the spectral axis to compute the mean spatial profile along the slit. Order reduction continues with spectral extraction and sky line identification. Spectra are non-optimally extracted using fixed size extraction windows, while the background is determined by fixed windows on either sides of the object profile. Spectra of background-subtracted object, sky and estimated noise are produced.

- **Write products**

At the end of the reduction process, NSDRP writes out the following products for each order of each frame of NIRSPEC spectroscopic data: FITS and ascii tables containing column numbers, wavelength, flux counts, error, sky background, trace and profile fits. It also produces FITS files of both the rectified 2-D images as well as 1-D extracted spectra of the flux, sky background, noise, trace, and profile for each order. Assorted jpeg images of these products are also provided for preview purposes along with log files. Shown below are some examples of the output products.

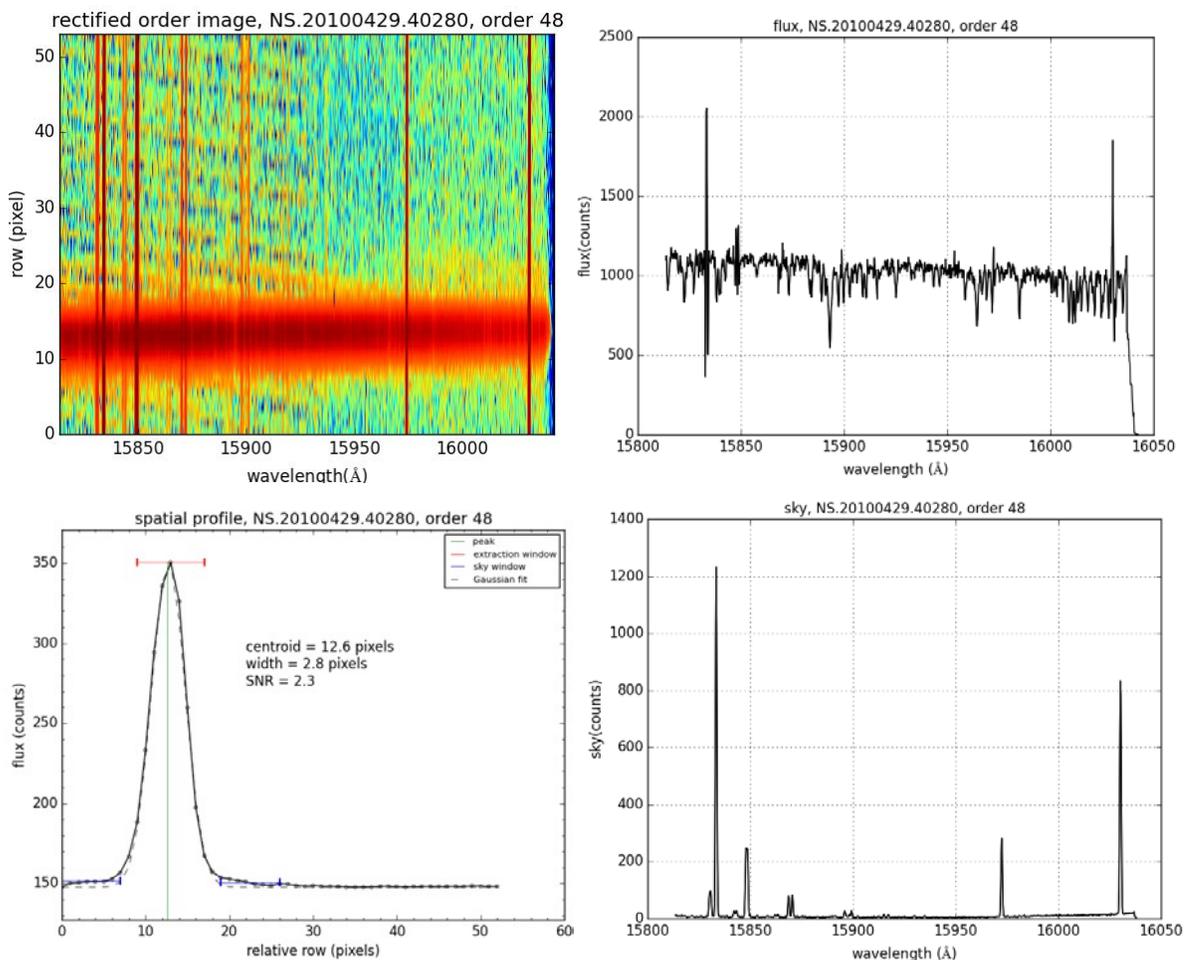


Figure 4. Examples of reduced data products from NSDRP: *Top left*: 2-D rectified, wavelength calibrated order image; *Top right*: extracted object flux spectrum; *Bottom left*: profile showing the extracted and background windows used; *Bottom right*: extracted sky spectrum.

3.3.2 Future Improvements

Several areas for future improvement have been identified. NSDRP in KOA mode does not currently support pair subtraction, co-adding, telluric correction or any form of photometric calibration. Since there is not enough information in the FITS headers to make the necessary file associations, some mechanism for associating related files will be needed in order to implement these features. Pair subtraction for background removal is being developed for the command-line mode for quicklook assessment of data at the telescope, and we expect to release it at the next update.

The wavelength calibration could be further improved by using arc lamps in addition to the sky lines. This is especially helpful for very short exposures (≤ 30 sec) where sky lines are often either very weak or absent. Another strategy is to combine the wavelength solutions in successful, presumably longer-exposure cases, and apply it those that failed the wavelength calibration but are otherwise having identical instrumental setup (i.e., same filter, slit, echelle and cross-disperser angles, etc.). We are currently in the process of implementing this option for the next release of the DRP.

NSDRP uses conventional extraction but the noise calculations already performed could be used to implement optimal extraction to maximize signal-to-noise ratio while preserving spectrophotometric accuracy.

Other areas for improvement are dynamic adjustment of extraction window widths and separations, improvements to order edge detection algorithm to handle step inversions on overlapping orders, and optimization of empirically determined grating equation coefficients.

4. FUTURE DRPS FOR KOA

4.1 Multi-Object Spectrometer For Infra-Red Exploration (MOSFIRE)

The current version of the MOSFIRE DRP can be found at <https://keck-datareductionpipelines.github.io/MosfireDRP/>. It was written in python by N. Konidaris and C. Steidel at Caltech. Recent improvements and bug fixes, including automating the wavelength calibration and removal of dependency on IRAF, have made it suitable and nearly ready for use by KOA for systematic production of level-1 data. Currently, the DRP produces a 2-D rectified, wavelength calibrated, and background subtracted spectral image as the final product, but a 1-D spectral extraction module using both simple and optimal background subtraction algorithms is being developed. The module will be universal and usable by any other pipeline capable of producing a 2-D rectified and wavelength calibrated array with any number of sources. This is a step towards developing a common infrastructure for spectroscopic pipelines at WMKO. Once implemented, we anticipate that the MOSFIRE DRP will be the next pipeline to be incorporated into KOA.

4.2 ESI, LRIS, DEIMOS, NIRC

The *MAKEE* DRP written for HIRES has the capability to also reduce Echellette Spectrograph and Imager (ESI) data, and we are planning to implement it as the pipeline for ESI in ways similar to what has been done for HIRES. Incorporating the pipeline and ingesting the level-1 data into KOA should be fairly straightforward in this case. For the 1st generation near-infrared camera NIRC, a simple DRP to calibrate the imaging data similar to that used for NIRC2 could also be used.

For the Low-Resolution Imaging Spectrometer (LRIS) and DEep Imaging Multi-Object Spectrograph (DEIMOS), no automated DRPs exist, although there are reduction packages available, such as Low-Redux⁸ by J. Prochaska and DEEP2⁹. Producing fully automated DRPs suitable for use by KOA for these optical multi-object spectrometers is challenging, and requires significant resources and effort. A recent study by B. Holden (internal study commissioned by the Keck Science Steering Committee) estimates that it would take ~ 16 months of effort to produce a fully functional LRIS DRP for use both by KOA and the general user community. It would be written in python by converting the Low-Redux IDL package with some reuse of codes from the MOSFIRE DRP.

4.3 Future instruments: NIRES and KCWI

Two upcoming instruments at Keck will be delivered with DRPs. For the Near-Infrared Echellette Spectrometer (NIRES), a pipeline called *TSX* (<http://www.astro.caltech.edu/~tb/tsx/>) has been written by T. Barlow (the developer of the HIRES pipeline *MAKEE*) based on the data for Triplespec at Palomar. Since NIRES is a close copy of Triplespec (one of four such clones¹⁰ available in the world; the other two are at Apache Point Observatory and Cerro Tololo Inter-American Observatory) with a single configuration and no moving parts, and the data format is essentially the same, very minimal modifications in the pipeline are expected to support NIRES data. *TSX* runs fully automatically and in a very similar way to the *MAKEE* pipeline, and we expect that it would be fairly straightforward to adopt and automate it as a DRP to produce NIRES level-1 data for KOA.

The Keck Cosmic Web Imager (KCWI) currently under development at Caltech will be delivered with a DRP written in IDL called *KDERP* (<https://github.com/kcwiddev/kderp>). The pipeline will be able to support the “nod & shuffle” observational technique (for best sky subtraction) and also have a quick-look mode that will help observers assess the quality of their data in real time. Like *TSX* for NIRES, the *KDERP* pipeline is also being developed based on data and expertise gained from the closely related CWI instrument at Palomar.

5. CONCLUDING REMARKS

While there is little doubt that having level-1 data in data archives such as KOA can be very useful to astronomers in helping them assess the quality and scientific content of the data, as a benefit to the Observatory, these level-1 products will enable the development of a rigorous instrument performance analysis program, which monitors and analyzes these data and leads to the improvement of calibration and maintenance of the instruments. Using this feedback process, the latest and best calibrations can be used for “on-the-fly” data processing, helping observers to get the best-quality data. Higher-level science data products, suitable for scientific analysis and publication, can also be accepted and served by KOA in the form of contributed datasets from observers. We have made such datasets available on KOA for the

KODIAQ¹¹ project, and will continue to encourage and host additional contributions, as has been done for example, at the Mikulski Archive for Space Telescopes (MAST)¹².

Acknowledgements: The Keck Observatory Archive (KOA) is a collaboration between the W. M. Keck Observatory (WMKO) and the NASA Exoplanet Science Institute (NExSci). Funding for KOA is provided by the National Aeronautics and Space Administration (NASA). WMKO is operated as a scientific partnership among the California Institute of Technology, the University of California, and NASA. The Observatory was made possible by the generous financial support of the W. M. Keck Foundation. NExSci is sponsored by NASA's Origins Theme and Exoplanet Exploration Program, and operated by the California Institute of Technology in coordination with the Jet Propulsion Laboratory. This material is based in part upon work supported by NASA under Grant and Cooperative agreement No. NNX13AH26A.

REFERENCES

- [1] Berriman, G. Bruce, Gelino, C. R., Laity, A. C., Kong, M., Swan, M., Holt, J., Goodrich, R. Mader, J, and Tran, H. D., "The Operation and Architecture of the Keck Observatory Archive," *Astronomical Data Analysis Software and Systems XXIII*, 485, 123 (2014).
- [2] Berriman, G. B., et al., "The Design and Operation of the Keck Observatory Archive," *Proc. SPIE 9152*, 91520A (2014).
- [3] Holt, J., Tran, H. D., Goodrich, R., Berriman, G. B., Gelino, C. R., and the KOA Team, "Automating OSIRIS Data Reduction for the Keck Observatory Archive," *Astronomical Data Analysis Software and Systems XXIII*, 485, 127 (2014).
- [4] McLean, Ian S., Becklin, Eric E., Bendiksen, Oddvar et al., "Design and development of NIRSPEC: a near-infrared echelle spectrograph for the Keck II telescope," *Proc. SPIE 3354*, 566-578 (1998).
- [5] Tran, Hien D., Mader, Jeff A., Goodrich, Robert W., Tsubota, Myrna, "Technical Solutions in Preparing Data for The Keck Observatory Archive (KOA)," *Proc. SPIE 8451*, 8451E-29 (2012).
- [6] Tran, H. D., Holt, J., Goodrich, R. W., et al., "Metadata and data management for the Keck Observatory Archive," *Proc. SPIE 9152*, 91522I (2014).
- [7] Rousselot, P., Lidman, C., Cuby, J.-G., Moreels, G., and Monnet, G., "Night-sky spectral atlas of OH emission lines in the near-infrared," *Astronomy and Astrophysics*, 354, 1134-1150 (2000).
- [8] The Low-Redux page: <http://www.ucolick.org/~xavier/LowRedux/>.
- [9] The DEEP2 DEIMOS Data Pipeline: <http://deep.ps.uci.edu/spec2d/>.
- [10] Wilson, J. C., Henderson, C. P., Herter, T. L., et al., "Mass producing an efficient NIR spectrograph," *Proc. SPIE 5492*, 1295-1305 (2004).
- [11] The KODIAQ data repository at KOA: <https://koa.ipac.caltech.edu/applications/KODIAQ/>
- [12] High Level Science Products at MAST: <https://archive.stsci.edu/hlsp/index.html>