Instrument Simulator and Data Reduction Pipeline for the iLocater Spectrograph

Eric B. Bechter[1] , Andrew J. Bechter[1], Justin R. Crepp[1], Jonathan Crass[1], and David King[2]

[1] University of Notre Dame 225 Nieuwland Science Hall Notre Dame, IN 46656, USA; ebechter@nd.edu

[2] University of Cambridge Madingley Road Cambridge, CB3 OHA, UK

Abstract

iLocater is a near-infrared (NIR) radial velocity (RV) spectrograph that is being developed for the Large Binocular Telescope in Arizona. Unlike seeing-limited designs, iLocater uses adaptive optics to inject starlight directly into a single-mode fiber. This feature offers high spectral resolution while simultaneously maintaining a compact optical design. Although this approach shows promise to generate extremely precise RV measurements, it differs from conventional Doppler spectrographs, and therefore carries additional risk. To aid with the design of the instrument, we have developed a comprehensive simulator and data reduction pipeline. In this paper, we describe the simulation code and quantify its performance in the context of understanding terms in a RV error budget. We find that the program has an intrinsic precision of $\sigma < 5\,\mathrm{cm\,s^{-1}}$, thereby justifying its use in a number of instrument trade studies. The code is written in MATLAB and available for download on GitHub.

Key words: instrumentation: spectrographs – methods: data analysis – methods: numerical – techniques: radial velocities – techniques: spectroscopic

Online material: color figures

1. Introduction

A new generation of Doppler radial velocity (RV) spectrographs aim to use adaptive optics (AO) to inject starlight directly into single-mode fibers (SMF) (Crepp et al. 2016). Unlike multi-mode fibers (MMF) or conventional slit-fed designs, SMFs propagate only the fundamental spatial mode, resulting in a stable Gaussian beam that is completely decoupled from input imaging conditions and fiber stresses (Bland-Hawthorn & Horton 2006; Bland-Hawthorn et al. 2010). The small mode field of SMFs, typically only several microns in diameter in the near-infrared (by definition comparable to the wavelength of light), further allows for high spectral resolution observations while offering a compact opto-mechanical design (Crepp 2014; Schwab et al. 2014; Jovanovic et al. 2016). Realizing the advantages of improved wavelength sampling and instrument stability would allow diffraction-limited spectrographs to reach instrumental noise floors well below $1\,\mathrm{m\,s^{-1}}$ (Bechter et al. 2018). Moreover, ultra-high resolution measurements ($R \geqslant 150{,}000$) present an opportunity to study asymmetries in the profiles of stellar absorption lines, a first step toward disentangling the effects of stellar variability (astrophysical "jitter") from the signal of orbiting planets (Davis et al. 2017).

The RV semi-amplitude induced by an Earth-mass planet orbiting in the habitable zone of a Sun-like star corresponds to a measurement precision of only 10^{-4} pixels. However, the majority of data reduction pipelines currently available have only been tested at the level of $\sim 10^{-3}$ pixels. Due to the fact that many subtle effects are expected to be encountered at the sub-meter-per-second level, it is important to start developing data reduction pipelines early-on in the design process for new RV instruments (Fischer et al. 2016). In the case of diffraction-limited systems, for which there is no archival repository of representative frames, numerical simulations can serve as an important alternative that captures the physics relevant to SMF spectrographs and its impact on RV extraction methods. Investing time into thorough models also provides the framework necessary for trade-studies that offer feedback during the design process.

Modern optical design tools, e.g., Zemax OpticStudio or Synopsys CODE V, can use the optical model to generate spectral traces and orders as they would be measured by a megapixel array (Terrien et al. 2014; Gibson & Wishnow 2016). The optical footprint generated from these tools allows spectra to be mapped to a physical detector plane. However, additional details are needed to simulate the entire light path from stellar source through the atmosphere to the telescope and instrument. Specifically, this includes simulated spectral sources, such as stellar spectra or calibration sources, atmospheric emission and absorption spectra, wavelength dependent instrument efficiency curves, instrument point-spread functions (PSFs), and the addition of photon noise, detector noise, and other sources of uncertainty. Self-consistently combining all of

these tools into a powerful simulation program that, when paired with a data reduction pipeline, can be used to investigate a number of important topics before the instrument sees first light:

1. Inform instrument optical design (e.g., spectral order and intra-order spacing requirements);
2. Quantify RV error budget terms;
3. Simulate entire RV surveys and analyze on-sky performance;
4. Investigate secondary science cases following commissioning;
5. Optimize calibration system parameters; and
6. Verify algorithms used in the data reduction process.

Such calculations are particularly valuable when connecting scientific goals of a program to hardware requirements through the instrument design review process and laboratory testing.

Motivated by the benefits of detailed software modeling, we have developed an end-to-end simulator as well as an initial data reduction pipeline for the iLocater spectrograph: a diffraction limited PRV instrument that will be installed at the Large Binocular Telescope (LBT) in Arizona (Crepp et al. 2016). In this paper, we present the iLocater instrument simulator (Section 2), including calibration of spectral models, Doppler broadening and RV shifts (planets, barycentric motion, noise), atmospheric simulations, instrument throughput calculations, instrument response function simulations, and detector simulations. In Section 3, we describe the structure of each part of the modular data pipeline: image processing, spectral extraction, wavelength calibration, pre-processing, and RV extraction. We demonstrate the simulator's intrinsic RV precision in Section 4. Finally, we present a number of useful applications from simulations in Section 5. The simulation code described has been designed to adaptable such that minor code modifications can introduce additional astronomical and calibration sources or even different spectrograph optical designs. In order to make these tools accessible for the astrophysics community, we have made the source code publicly available for download on GitHub.[3]

2. Instrument Simulator

The software simulator uses an object-oriented, modular, class structure for its code architecture. The simulator has been developed initially for the iLocater system and is comparable to simulations built for instruments such as HPF and HARPS-N (Artigau et al. 2012; Terrien et al. 2014). The majority of code is written in MATLAB, although it relies on optical design exports from Zemax, and detector noise generated from a Python package. The iLocater instrument is organized into three main parts: the acquisition camera, spectrograph, and

calibration system (Crepp et al. 2016). The acquisition camera couples AO-corrected light from the telescope into a single-mode fiber that feeds the spectrograph. The calibration system provides the option of injecting a Fabry–Pérot etalon spectrum, Uranium-Neon (U-Ne) emission lamp, or a white light spectrum, depending on the type of frame needed. Each of these optical systems are simulated in this code as well as the wavelength dependant throughput, AO system, atmosphere, and stellar spectra.

The simulator class design uses the physical divisions of the iLocater instrument to define classes whenever possible (Figure 1). There are six primary super-classes: (1) the Spectrum class instantiates the type of spectrum or spectra needed according to the mode selected by the user; (2) the Instrument class allows changes to the optical model to be made as well as modular selection of the instrument; (3) the Throughput class models each optical path of the instrument as a wavelength dependent efficiency curve; (4) the Simulation class combines results from the first two classes, imports Zemax optical design details, and handles converting spectra from 1D to 2D and binning them on the detector in a flux preserving manner; (5) the Atmosphere class creates atmospheric spectra according to specific input conditions; and (6) the Fiber Coupling class, which computes coupling losses due to fiber misalignments and other coupling issues. The following sections describe the flow of the simulation code which also shown visually in Figure 2.

2.1. Input Options and Variables

The simulator is a flexible and complex code package containing many user input options and settings that need to be initialized. These include parallel settings to enable parallel processing; simulation variables that allow upsampling for more precise calculations; spectral options for stellar, atmospheric and calibration sources; spectrograph design and polarization settings; the number of spectral orders and sub-orders to be simulated; fiber coupling parameters; and observing conditions. Table 1 lists input options and variables currently available to the user.

2.2. Preparing the Spectra

An instance of each spectral class is created according to the specific settings defined by the user inputs. Each of these classes inherits a common set of properties and methods from the parent Spectrum class as well as having their own specific properties and methods (Figure 1).

2.2.1. Stellar Spectra

The simulator uses a catalog of BT-Settl FGKM models (Allard et al. 2012). The requested spectrum is retrieved from the catalog and scaled to the specified apparent magnitude
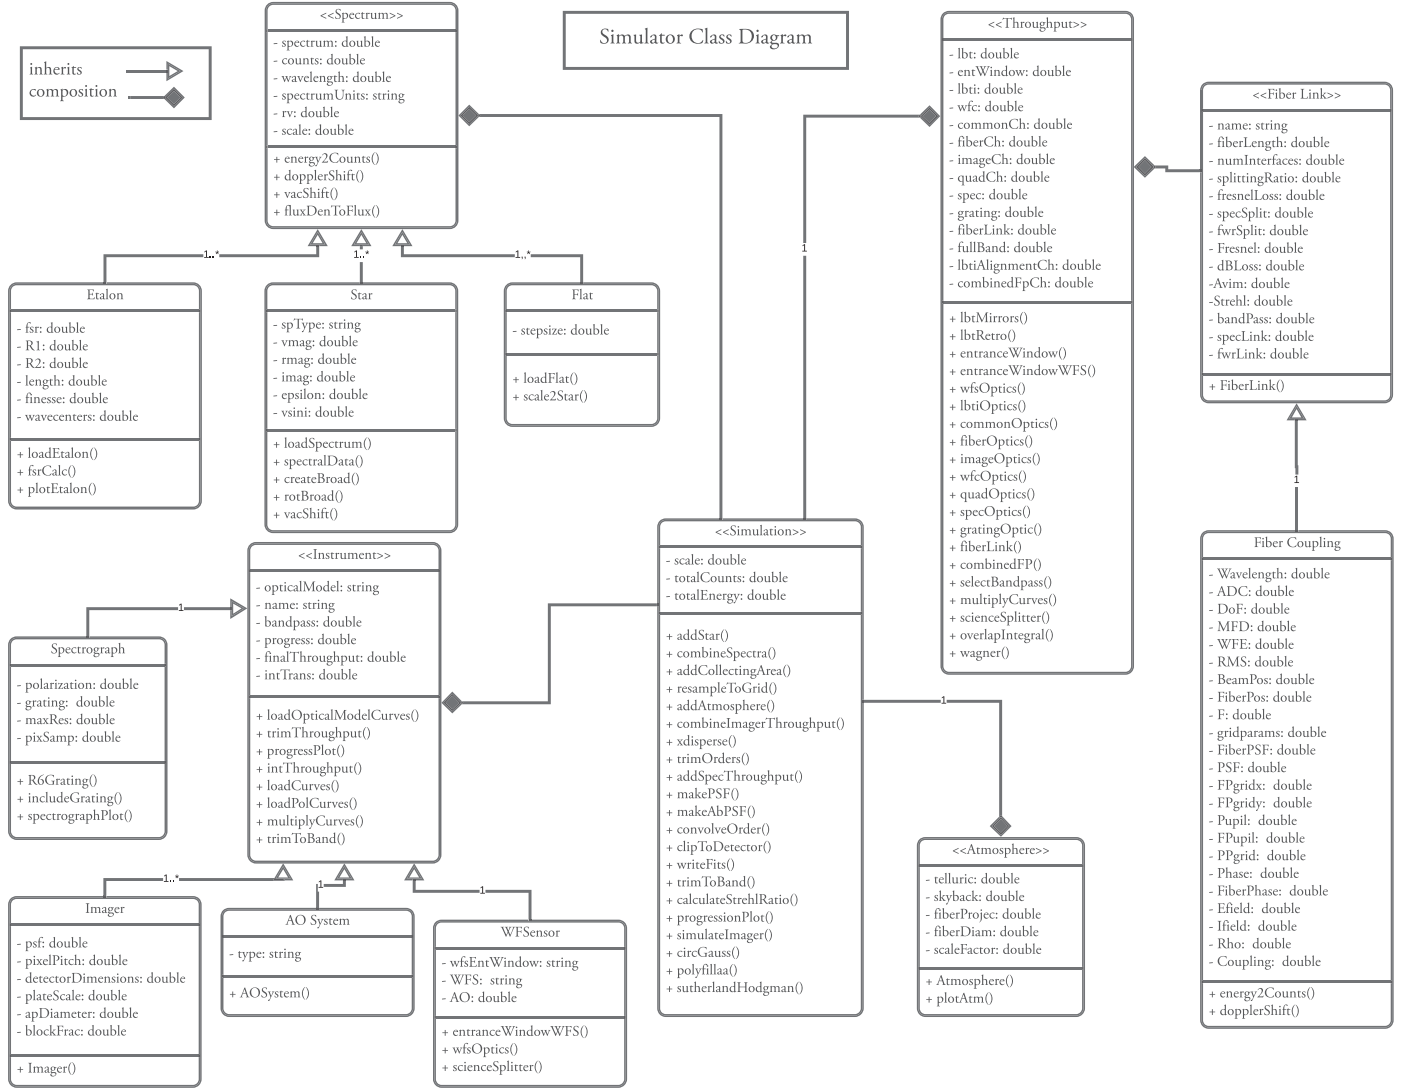
---

3 https://github.com/ebechter/InstrumentSimulator

Figure 1. UML class diagram for the Simulator. This chart shows the simulator's object-oriented nature for the six main super-classes and their subclasses. Each super-class is named within ≪≫ and subclasses are shown by the inherit open-tipped arrow. Composition lines note object existence dependencies and association line contains an indication of the one (1) or many (1..*) allowed instances of a class. Each box indicating a class show their properties first, including the data type and a list of their methods below.

using the stellar color and effective temperature sequence table provided in Pecaut & Mamajek (2013). All spectra in the catalog are generated at vacuum wavelengths and a resolution of 0.02 Å. Figure 3 shows a model M2V spectrum, indicating iLocater's wavefront sensor, imaging, and spectroscopy bands.

Simulated stellar spectra first need to be Doppler broadened as it is largely responsible for the resulting profile of each line. Rotational line broadening manifests from stellar rotation, quantified by the projected equatorial velocity, $v \sin i$. Each surface element of the star with a different velocity will result in a different Doppler shift. To calculate the magnitude of this effect, the line-of-sight components of these surface velocities are integrated over the disk of the star, resulting in a cumulative

Doppler shift distribution. To implement this rotational broadening effect in the simulator, the Doppler shift distribution is calculated according to Equation (18.14) in D. F. Gray's "The Observation and Analysis of Stellar Photospheres" (Gray 2005). This formulation also takes into account the relative flux levels changing from the stellar edge to the center via the linear-limb darkening law. Additionally, stellar spectra can be Doppler shifted according to a user input radial velocity. The relativistic Doppler equation is used to compute the wavelength shift:

$$\lambda_{\text{shift}} = \lambda * \sqrt{\frac{1 + \beta}{1 - \beta}}, \qquad (1)$$
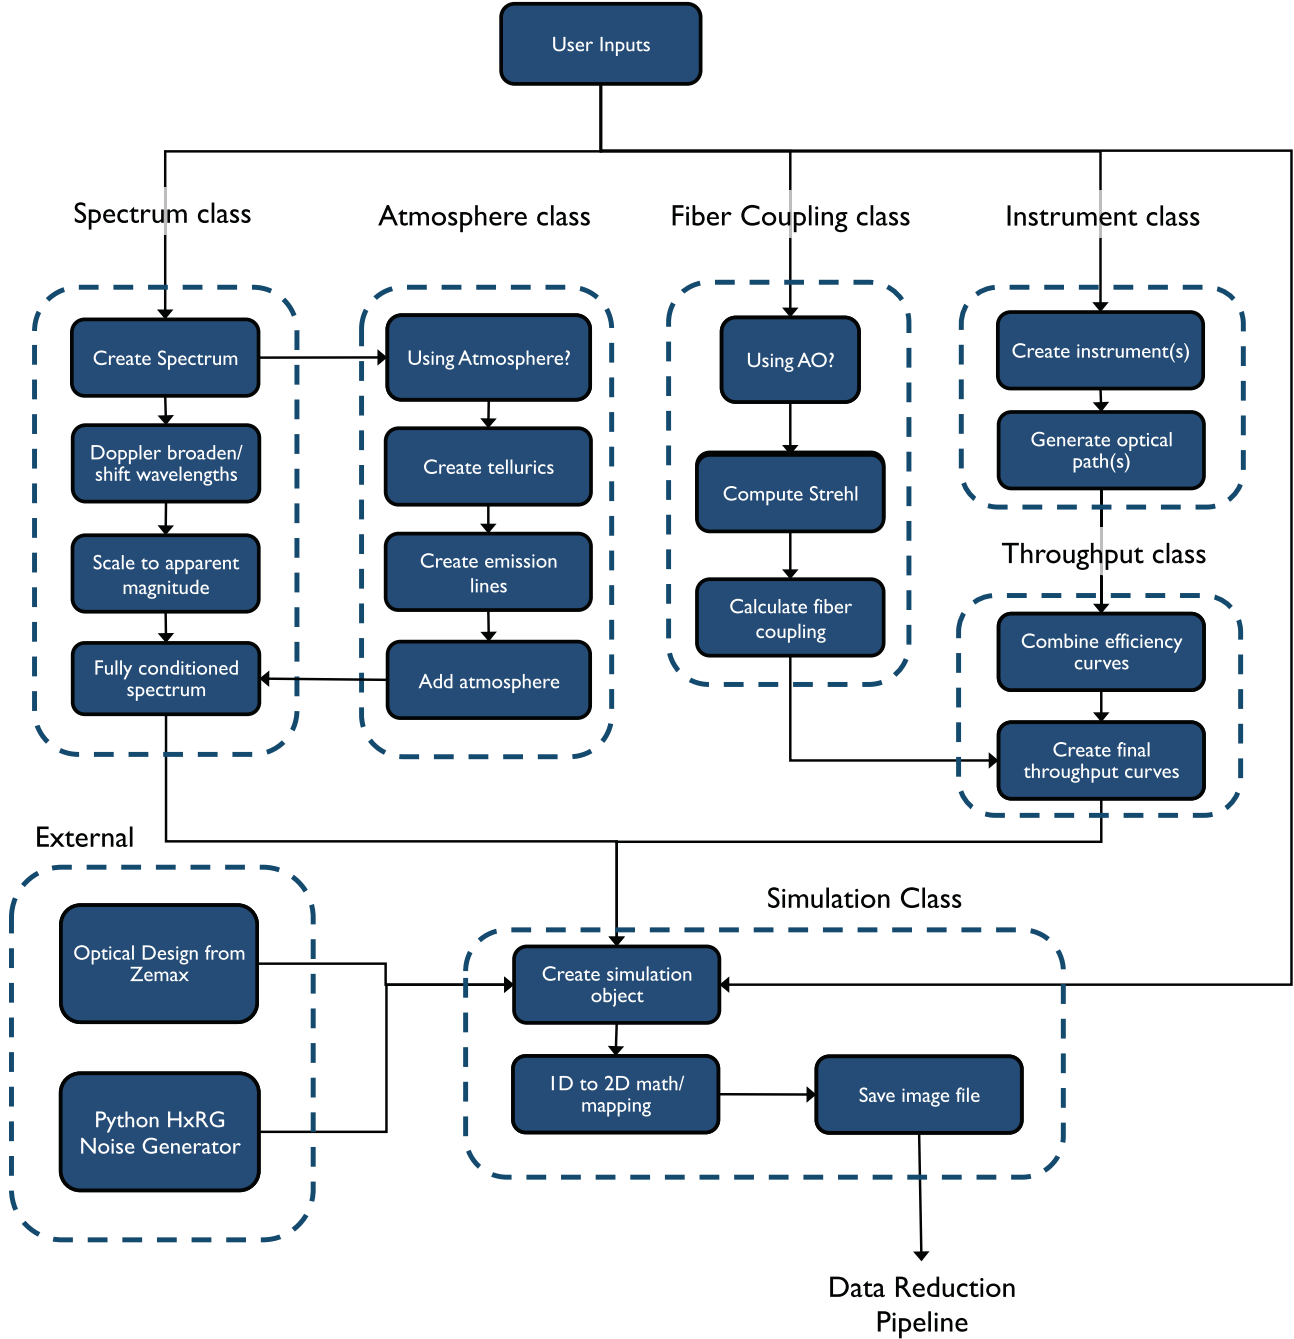
Figure 2. Flow diagram of simulation code. Each class is responsible for generating the specific subclass of Spectrum or Instrument that will be used in the simulation class. Throughput curves are generated and condensed into a single efficiency curve which is combined with the final spectrum product in the simulation object. External noise generators and optical formatting files are also used by the simulation object.
(A color version of this figure is available in the online journal.)

where λ_{shift} is the Doppler shifted wavelength, λ is the rest-frame wavelength, and $\beta = \Delta v_{\mathrm{rad}}/c$.

2.2.2. Flat-field Spectra

There are two options for generating flat field spectra. The first source is a discretely sampled uniformly flat spectrum that is scaled to a specified source power. This source is primarily used for debugging purposes and assessing throughput efficiences. The second, more realistic, choice of flat-field source is a flattened white light spectrum generated from the same source that injects light into the Fabry–Pérot calibration unit.

Table 1
List of Input Variables and Initialization Parameters

Name	Default	Unit	Description
Parallel Settings:			
parflag	0	1/0	Flag indicating the use of parallel processing
numworkers	⋯	⋯	Set number of parallel processing cores
Simulation Inputs:			
scale	1	⋯	Up-sampling factor used to minimize numerical errors during simulation
dname	⋯	⋯	Assign directory where files are saved
fname	⋯	⋯	Set file names for simulated frames
Stellar Inputs:			
type	M0V	⋯	Stellar spectral type model from catalog
vmag	11	⋯	Apparent visual magnitude of star
vsini	2	km s^{-1}	Stellar rotational velocity
rv	0	m s^{-1}	Inject a stellar radial velocity
epsilon	1	⋯	Limb darkening parameter used in rotational broadening
Etalon Inputs:			
FSR	10	⋯	Free spectral range of Fabry–Pérot
R1	0.93	⋯	Reflectively of first mirror in Fabry–Pérot
R2	0.93	⋯	Reflectively of second mirror in Fabry–Pérot
finesse	⋯	⋯	Specify finesse as an alternative to R1 and R2
rv	0	m s^{-1}	Inject a radial velocity
Spectrograph Inputs:			
nOrders	36	⋯	Number of spectral orders generated on detector
tracenum	[1, 2, 3]	⋯	Specify which of the three traces to be used
footprint	12.22	⋯	Optical layout version to be used in simulations
waveSolution	0	1/0	New wavelength solution (1) or load previous (0)
throughput	⋯	⋯	Cell array of strings indicating path of light
polarization	[0, 0.5]	⋯	The degree of polarization and P-fraction of light
Fiber Coupling Inputs:			
fiberpos	$[\frac{4}{7}, \frac{4}{7}, 13]$	μm	Global fiber position offset in X, Y, and Z
rndfiber	⋯	1/0	Gaussian random fiber position
fibermas	42	mas	Fiber diameter in milliarcseconds
Observation Conditions:			
zdeg	45	degrees	Zenith angle
AO	SOUL	⋯	AO system. LBT's FLAO or SOUL AO system
tellurics	1	1/0	Include telluric spectrum (1) or ignore (0)
skyback	1	1/0	Include sky-background spectrum (1) or ignore (0)

Note. For a more complete description of these, see the getting started section in the readme file included with the Github distribution.

2.3. Fabry–Pérot Etalon

Fabry–Pérot etalons are a powerful and cost effective tool used for precise wavelength calibration in high resolution spectrographs. The etalon cavity produces a comb-like function in a way that can be tuned to match the the spectrograph wavelength band and resolution while offering numerous calibration peaks. In the simulator, the spectrum generated by a Fabry–Pérot cavity is computed using the transmittance function,

$$T = \frac{1}{F \sin^2(\delta/2)}, \qquad (2)$$

where F is the coefficient of finesse and delta is the phase difference between successive reflected pairs. Section 5.3 outlines the optimization procedure used to select the specific parameters of iLocater's Fabry–Pérot.

2.4. Atmospheric Effects

Accounting for the spectral contaminating effects of Earth's atmosphere (namely, tellurics and OH emission lines) is essential when designing an instrument to work in the NIR as telluric absorption features in this wavelength region can negatively impact RV measurements by several meters per second (Bean et al. 2010). Telluric absorption features are comprised of a number of molecular species (e.g., H_2O, O_2,
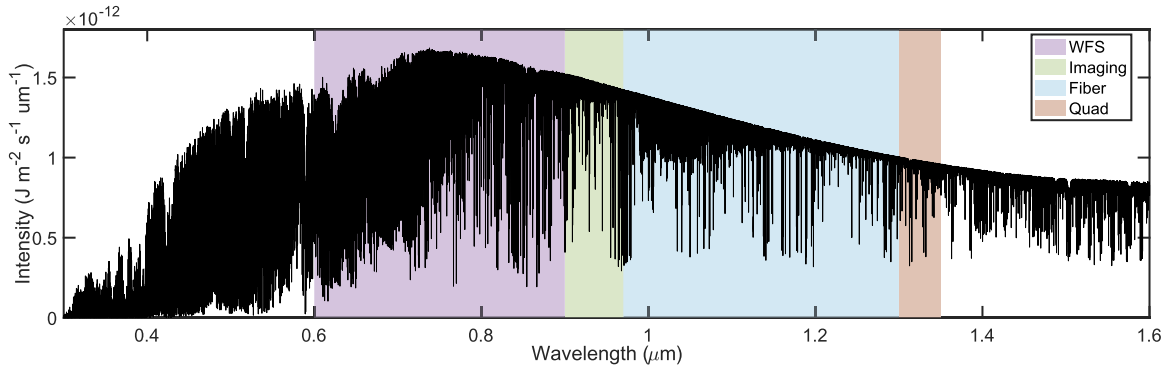
Figure 3. BT-Settl synthetic spectrum: $T_{eff} = 3600$ K (~M2V) scaled to an apparent magnitude of $v = 12$. The shaded regions show the wavelengths of light used in each iLocater subsystem. The purple shaded band represents the wavefront sensing channel (0.60–0.90 μm), the green band is used in the acquisition camera imaging channel (0.90–0.97 μm), the blue band is directed to the the single-mode fiber spectroscopic channel (0.97–1.27 μm), and the red band is used in the quad cell (1.3–1.5 μm).
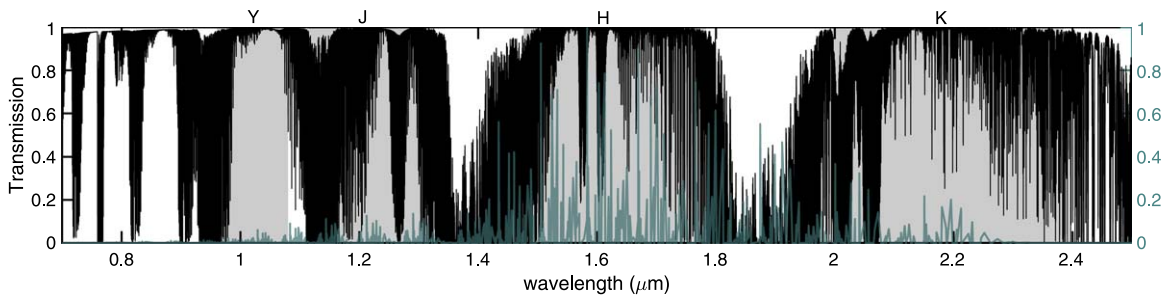(A color version of this figure is available in the online journal.)

Figure 4. Transmission spectrum of the atmosphere in the NIR. The Y-, J-, H- and K-bands are shaded. Relative OH emission line strengths are shown in teal. iLocater has been designed to work in the Y- and J-band regions due to the relative lack of telluric and emission lines compared to other NIR regions.
(A color version of this figure is available in the online journal.)

CH_4, CO_2) which contaminate stellar light as it travels through the atmosphere. Although few of these species have small seasonal variations (CO_2 and CH_4), most will fluctuate on a much more frequent timescale (on the order of 10 minutes), presenting a significant obstacle to the data reduction process (Vacca et al. 2003).

To simulate the absorption effects of Earth's atmosphere, a telluric spectrum is generated using the online TAPAS code which can then combined with stellar models (Bertaux et al. 2014). Given an R.A., decl., location (latitude, longitude, altitude), and a date, TAPAS computes the atmospheric transmittance from the top of the atmosphere down to the observatory, based on the HITRAN (high-resolution transmission molecular absorption database) molecular database and the LBLTRM (Line-By-Line Radiative Transfer Model) radiative transfer code. The program also provides separate transmittances associated with H_2O, O_2, O_3, CO_2, CH_4, N_2O and Rayleigh scattering.

Atmospheric emission lines at the NIR (1–2.5 μm) are mainly comprised of numerous narrow OH emission lines.

These lines fluctuate on a timescale of 5–15 minutes with relative flux variations between 5% and 10%. To incorporate sky background into the simulator, high resolution, simulated sky background spectra are downloaded from the Gemini website.[4] These models incorporate high resolution sky emission spectra, zodiacal continuum emission approximated by a $T = 5800$ K blackbody, and thermal emission from the atmosphere treated as a blackbody with $T = 273$ K, available at a variety of airmasses and water vapor column depths. Figure 4 is an example synthetic telluric spectrum generated at the LBT site (Mt. Graham, Arizona) with standard observing conditions and a normalized sky background spectrum overplotted.

Simulating ground-based observations also includes a wavelength shift through the atmosphere. This is accounted for by

---

4   http://www.gemini.edu/sciops/telescopes-and-sites/observing-condition-constraints/ir-background-spectra

using the IAU standard conversion,

$$\frac{\lambda_{\text{vac}} - \lambda_{\text{air}}}{\lambda_{\text{air}}} = 6.4328 \times 10^{-5} + \frac{2.94981 \times 10^{-2}}{146 - \sigma^2}$$
$$+ \frac{2.5540 \times 10^{-4}}{41 - \sigma^2} \quad (3)$$

where $\sigma = 10^4/\lambda$, with λ in angstroms.

2.5. Flux Simulation

We estimate incident flux, $F(\lambda)$, based on the number of photons per second delivered to each of the instrument's optical channels using

$$F(\lambda) = \int s_*(\lambda) \, a(\lambda) \, T(\lambda) \, d\lambda, \quad (4)$$

where $s_*(\lambda)$ is the stellar spectrum, $a(\lambda)$ represents the atmospheric effects, and $T(\lambda)$ is the instrument throughput.

Steps taken to estimate flux are

1. Rotationally broaden stellar models s_* using a projected rotational velocity, $v \sin i$.
2. Scale flux to desired apparent magnitude (using Vmag as reference).
3. Construct atmospheric spectrum, $a(\lambda)$, using synthetic telluric models for atmospheric absorption and sky-background for atmospheric emission.
4. Multiply spectrum by the throughput curve of the channel, $T = (T_{\text{spec}}, T_{\text{WFS}}, T_{\text{image}})$.

Flux values ($\text{J s}^{-1} \text{m}^{-2} \, \mu\text{m}$) are converted to detector counts first by multiplying by the collecting area of LBT, calculated using a primary mirror diameter of $D = 8.22$ m with a secondary obstruction of $d = 0.91$ m (11% diameter). To convert from Joules to photon counts, each discrete spectral bin is divided by the average energy of a photon in that bin. The product is multiplied by the width of each bin (in microns) and multiplied by the integration time.

2.6. Adaptive Optics

Calculating expected Strehl ratio is important when simulating AO-fed instrument performance as it is necessary for fiber coupling estimation. The baseline AO system performance for a given star is calculated using a look-up table that converts the simulated number of incident photons per second at the wavefront sensor pupil plane to a Strehl ratio at the center of the V, R, Z, Y, J, H, and K bands. The table was provided by the LBT AO group, which calculates Strehl ratio for a given zenith seeing (Pinna et al. 2016). To approximate Strehl degradation with increasing airmass, we incorporate the zenith distance, z, analytically:

$$S(\lambda, \theta, z) = \exp(-|\log_e(S_0(\lambda, \theta))| \sec z), \quad (5)$$

where $S_0(\lambda)$ represents the Strehl ratio at $\sec z = 1$. The result is a changing chromatic throughput dependent on airmass.

2.7. Modeling Instrument Throughput

Careful modeling of instrument throughput is essential for precision instruments, especially high resolution spectrographs, as they are typically photon starved. Accurate throughput budget models are used to inform many design choices, including estimating observational efficiency through exposure times, determining an instrument bandpass, and limiting stellar magnitudes for science cases.

Instruments are generated as individual objects, as seen in Figure 1. This allows for each instrument to be simulated individually or as a combination of subsystems. The result of throughput modeling is to generate the wavelength dependent instrument throughput $T(\lambda)$ for each object and the specified combination of all instrument objects.

To create an instrument object, the user can define an arbitrary number of sequential optical elements. Each element can be defined with a surface coating name, number of surfaces, number of passes, angle of incidence, etc. Efficiency curves for each surface are drawn from a database of standard coatings and throughput is numerically calculated by sequentially multiplying the efficiency curves $T(\lambda)$ at each surface. If polarization is specified, then S and P curves of each optic are taken into account and the total throughput and degree of polarization are computed.

For iLocater, the user can specify a number predefined optical paths including those which represent the optical channels of the telescope, acquisition camera, spectrograph and calibration source. Each optical path is instantiated as an object and object parameters are populated both from initially defined parameters and the interaction between other classes during the simulation process. New optical paths can be defined and used interchangeably with existing ones to model instrument upgrades or as a stand-alone object for new instruments.

The object-oriented nature of the simulation code allows the user to quickly identify the subcomponents that introduce highest throughput losses. Design changes that alter optical coatings or the number optics can be assessed by quantifying the system performance both at the local subsystem and the global scale of the entire instrument.

2.8. Creating the Simulation Object

The Simulation class first imports the necessary optical model values from Zemax. To accomplish this, we use a Zemax macro to export focal plane coordinates (x, y) and wavelength (λ) data for several points per spectral order. Each order has three fiber traces that are physically offset to avoid cross-contamination. This data will be used to generate a wavelength solution and a physical mapping of each

wavelength to the detector. The simulation code also has the ability to import PSF models at each desired wavelength. The purpose of a PSF model is to determine the instrument response function and incorporate the effects of optical diffraction. Zemax's physical optics propagation (POP) analysis supplies the appropriate PSF information.

Once the optical model information has been generated in Zemax and read into MATLAB, the wavelength solution is created by fitting either a low order polynomial to each spectral order or using a 2D model of Chebyshev polynomials that fits all spectral orders at the same time. The specific model is also used in the wavelength calibration module in the data reduction pipeline and is described in detail in Section 3.4.

Next, final throughput curves are combined with fully conditioned spectra from the appropriate subclass of spectrum used to generate them. While the spectra are still 1D at this point, they are split into respective spectral orders in preparation for individual PSF convolution. Depending on the inputs provided and the mode specified, the simulator can use a single representative PSF for a spectral order or, using a custom convolution algorithm, spatially varying PSFs across every order. The simulator also has the ability to inject optical aberrations in simulated PSFs, allowing for users to simulate realistic PSFs without the need for POP analysis in Zemax. After convolution, each spectral order has two spatial axes (x, y) and one wavelength axis. This rectangular array is then warped and clipped onto the detector sampling such that it matches the physical curvature and location of each order. Flux-preserving clipping is accomplished by assigning vertices to each discrete sample in the rectilinear array, representing each sample as a polygon. Each polygon's vertices are mapped to the detector coordinates using the previously acquired Zemax optical information and distributed into each pixel accordingly using a Sutherland-Hodgeman clipping algorithm.

Once the spectral signal has been computed across the detector, the user has the option of adding photon noise as well as detector noise. Photon-noise is governed by Poisson statistics and added to the array by a random draw from a Poisson distribution for each pixel. Detector noise is added by making use of a NIR detector system noise generator written for the *James Webb Space Telescope* Near Infrared Spectrograph (Rauscher 2015). Figure 5 shows a single output frame from the simulator with iLocater's spectral format imprinted on an H4RG detector. The array is saved as a fits file with a custom header that contains many of the specific parameter settings of the particular simulation run.

3. Data Reduction Pipeline

The iLocater data reduction pipeline is comprised of five modular sub-pipelines: image processing, spectral extraction, wavelength calibration, post-processing, and RV computation. An overview of iLocater's data reduction pipeline is shown in

Figure 6, showing the flow of data required to recover a single RV measurement. All data reduction software algorithms have been developed in MATLAB. The following section provides a description of each sub-pipeline.

3.1. Image Processing

The image processing sub-pipeline converts raw science frames or cubes from the H4RG into cleaned, processed, 2D images, that are ready for spectral extraction. Figure 5 shows an example raw image with iLocater's spectral format imprinted. The standard steps for our spectroscopic data reduction are

1. Master calibration frame creation,
2. Pixel nonlinearity correction,
3. Dark subtraction,
4. Bad pixel correction,
5. Flat division.

3.1.1. Master Calibrations

The calibration frames used in image processing are flat fields and dark frames. Each master calibration frame is created using standard median combination. Flat fields are generated using the same white light source that powers the wavelength calibration unit combined with a flattening filter to minimize chromatic intensity modulations and neutral density filters to scale the output power according to the desired integration time.

3.1.2. Pixel Nonlinearity

HxRG's nonlinear response to light is a well documented effect in the literature, the most commonly mentioned effect being reciprocity failure (Biesiadzinski et al. 2011; Plazas et al. 2017). iLocater's H4RG nonlinear response will be thoroughly tested and characterized in the laboratory, prior to spectrograph integration and on-sky commissioning. The image processing sub-pipeline will apply the prescribed linearizing factor for each pixel to restore count values in raw image frames.

3.1.3. Dark Subtraction

Dark frames are a way in which dark current, detector bias and other background sources can be subtracted from science and calibration images. Additionally, hot/dead or poorly responding pixels on the detector can be identified at this stage. Dark frames will be created by integrating for a time equal to the exposure time of science frames while blocking all light to the spectrograph or alternatively existing dark frames can be rescaled to match the exposure time of the science frame. As they will contain the same level of bias as other frames, separate bias frames are not necessarily needed. After linearity correction, master dark frames are subtracted from each frame.
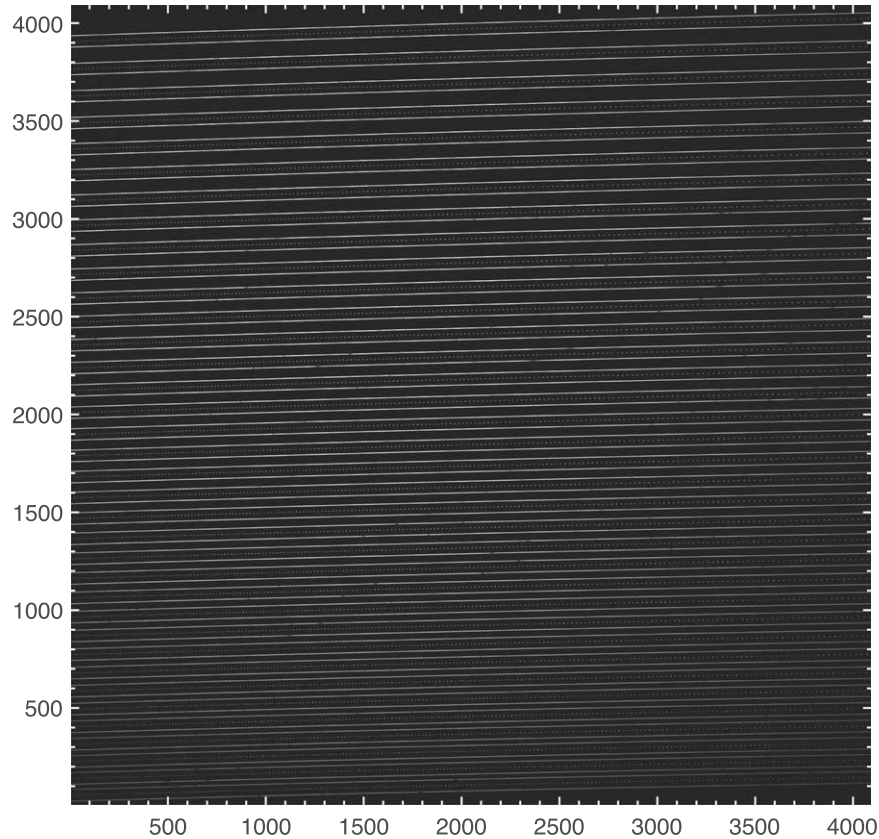
Figure 5. Example of a simulated detector frame. The top and bottom traces in each spectral order contain the spectrum of an M0V star while the central trace displays simulated 10 GHz etalon spectrum for simultaneous calibration purposes. The 36 spectral orders span orders $m = 117$–152, top to bottom with the dispersion direction going left to right. There are no physical noise sources or detector effects shown in this frame, however the entire throughput of the telescope and instrument can be seen in the Gaussian shaped intensity profile across each order and along the cross-dispersion (vertical) direction.

3.2. Bad Pixel Correction

Bad pixels identified in dark frames or during detector characterization in the laboratory are first saved in a bad pixel map, a 2D binary grid indicating the locations of bad pixels. Bad pixels are corrected by replacing each mapped bad pixel with a linearly interpolated value using a specified window surrounding the bad pixel.

3.2.1. Flat Division

Fiber-fed spectrographs, including iLocater, typically have no way to uniformly illuminate the entire detector once it is mounted in the spectrograph focal plane. Because of this they cannot correct full frame pixel-to-pixel variation, standard in most image processing pipelines. Fiber flats will be created by illuminating the entrance fibers with a white light source which will be used primarily to help remove the grating's blaze function (Section 3.5.3). However, full frame flat-fields will be taken in the lab during detector characterization and testing.

The fidelity of the lab flat-fields will be quantified throughout instrument commissioning.

3.3. Spectral Extraction

Spectral extraction is a particularly important step in the data reduction pipeline. In general, it refers to the operation of systematically recombining the signal around a central order's trace in the cross-dispersion direction, resulting in a 1D spectrum for each echelle order.

3.3.1. Order Identification

Before a spectral order can be recovered, it must first be identified using a continuum source (super-continuum source or tungsten) to illuminate the fiber orders. Numerous vertical slices, each a few pixels wide, are taken across the detector, median combined and smoothed using a Gaussian filter. Each of these is fit using a sum of Guassians model with the number of Gaussian functions equal to the number of orders present in the slice. The resulting centroids are paired with the corresponding horizontal
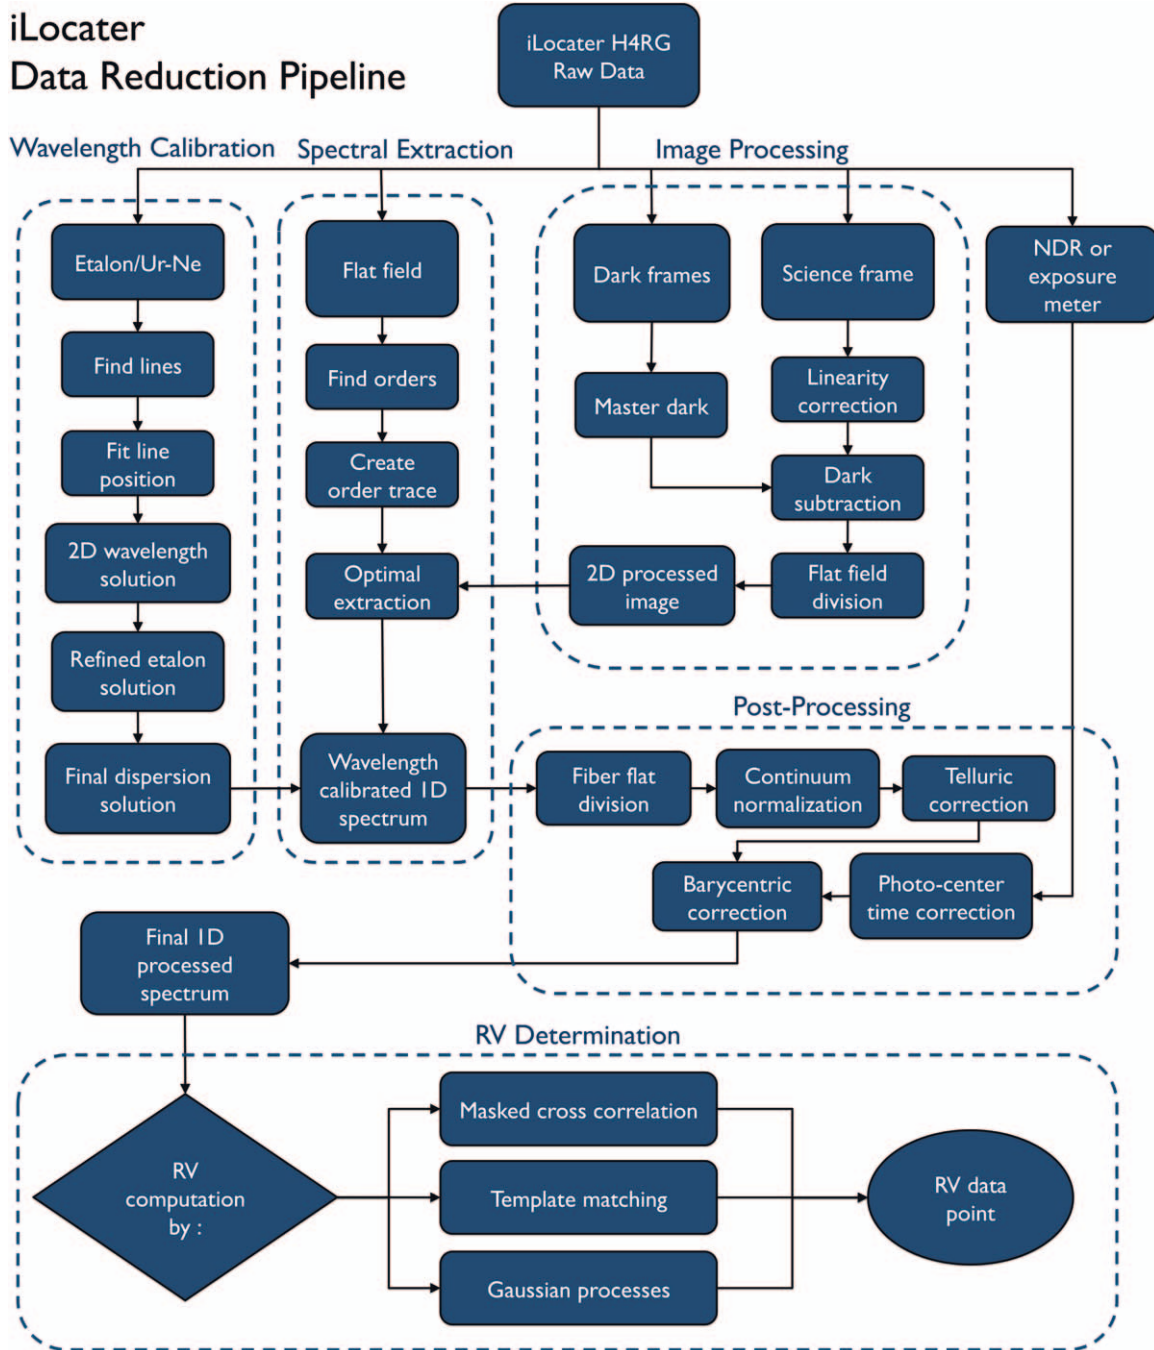
Figure 6. Data reduction pipeline organizational chart showing each sub-pipeline, their processes, and the flow of data required to recover a single RV measurement. (A color version of this figure is available in the online journal.)

pixel coordinate and fit using a low order polynomial across the dispersion direction. The polynomial coefficients are saved and used for subsequent science and wavelength calibration order extractions. Figure 7 (left panel) shows a small portion of traced orders from an image generated using the spectrograph simulator and the order identification algorithm.

3.3.2. Order Extraction

The data reduction pipeline currently has two options for performing spectral extraction: (1) a simple, sum-over-columns algorithm uses a vertical pixel window to sum all of the signal contained within, and (2) a more popular algorithm used by RV
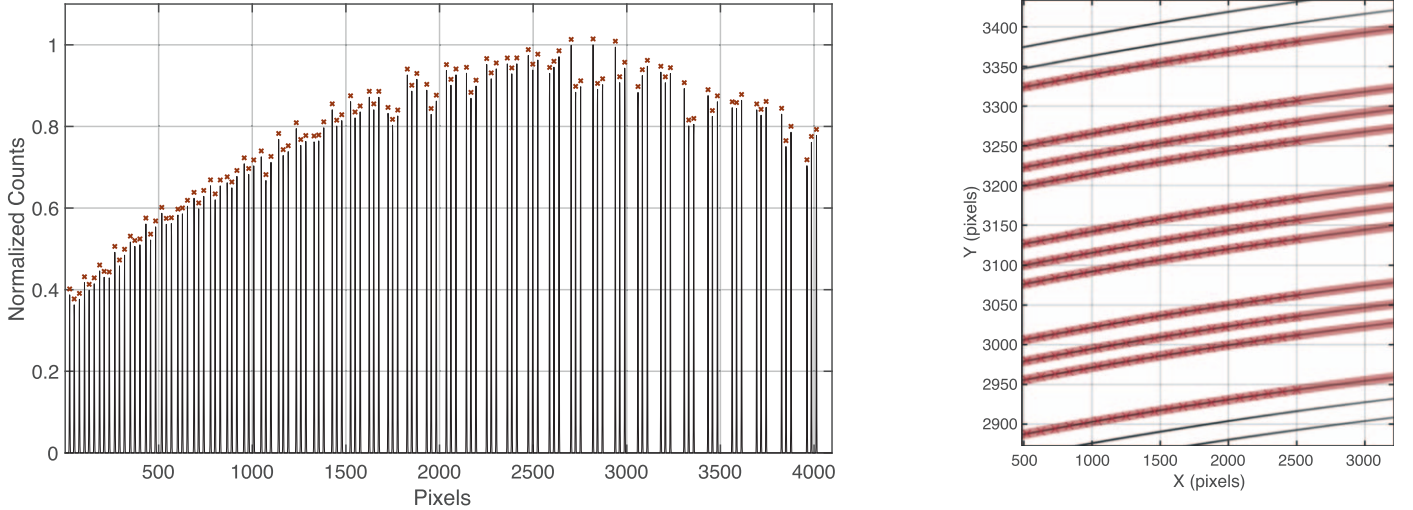
Figure 7. Left: Gaussian-smoothed, vertical slice of 10 median-combined columns showing all of iLocater's 36 spectral orders duplicated three times for each entrance fiber. Right: sub-frame image showing the trace algorithm and order identification algorithm working on a simulated flat field image. Color scale has been inverted to more clearly show the orders (black) and empty detector space (white).

(A color version of this figure is available in the online journal.)

instruments, optimal extraction (Horne 1986; Marsh 1989; Piskunov & Valenti 2002), which scales 1D cross-sectional profiles to the imaged spectrum where the scaling factor is based on a flux estimate. Most algorithms will also attempt to model and reconstruct the spatial profile/slit function with a polynomial/Gaussian.

There are also a few new algorithms being developed by other RV groups that could be added to iLocater's spectral extraction pipeline in the future. First, a "flat-relative" optimal extraction algorithm developed by M. Zechmeister (Zechmeister et al. 2014) offers improved speed and efficiency compared to classical optimal extraction. For stabilized spectrographs like iLocater, order profiles and positions are, for the most part, object- and time-independent, which means the spatial profile does not necessarily need to be modeled empirically. A high S/N master flat is used as an extraction mask where the extracted spectrum is scaled relative to the cross-section of the flat. Essentially, the extracted spectrum is measured relative to the flat spectrum and because the flat field contains the same spectral signature, e.g., spatial profile, pixel-pixel variations, and blaze function, these are all automatically incorporated into the extraction routine depending on how static they are are in practice.

Bolton & Schlegel (2010) outlines the "perfect" extraction of one-dimensional spectra from two-dimensional digital images of optical fiber spectrographs, based on accurate 2D forward modeling of the raw pixel data. This new technique promises statistically independent extracted samples of the 1D spectrum as well as no degradation of the 2D spectrograph resolution. However, this method requires very large matrix inversions. Another, more practical approach is

offered in Kos et al. (2018), where photonic combs are used to precisely map aberrations. Forward-modeling convolves this map with template spectra and attempts to reproduce the observed image. Results show this reconstruction method simplifies a number of reduction steps and reliably extracts spectra with 2–3 times nominal resolution. New methods such as these represent future spectral extraction improvements that are possible for stabilized, fiber-fed instruments like iLocater.

3.4. Wavelength Calibration

The wavelength calibration sub-pipeline handles tracking long-term wavelength drifts of the instrument and calculating the nightly dispersion solution. The calibration sources available to iLocater are U-Ne emission lamps and a laser-locked Fabry–Pérot etalon.

3.4.1. Dispersion Solution

The dispersion solution will provide science traces with a wavelength solution. To achieve this, iLocater will use a combination of the U-Ne emission lamp and a stabilized Fabry–Pérot etalon. The Fabry–Pérot module provides outstanding line density and stability, while the U-Ne emission lamp aids in practical identification of the etalon lines through an, absolute wavelength solution that is performed initially and periodically, as necessary. The more numerous and intrinsically stable etalon lines will be used to improve on the dispersion solution, providing a very precise wavelength calibrated spectrum, suitable for precision RV measurements.

The wavelength calibration procedure has been adapted for iLocater, but in general follows the procedure provided in Section 2.4 of Brahm et al. (2017). First, the U-Ne calibration source is fed through all three of iLocater's spectrograph entrance fibers. This illuminates every possible trace with U-Ne light, resulting in a "full U-Ne" exposure. Then, the etalon is used to illuminate all three fibers in the same way. The spectral extraction for both of these frames follows the same procedure described Section 3.3.2. After image processing and spectral extraction, U-Ne lines are fit using Gaussian functions and centers recorded. For most of the reduction and extraction process, spectral orders are referred to by their relative order numbers (e.g., 1–36). However, to enable a 2D wavelength solution $\lambda(x, m)$, the actual spectral orders, m, are needed. This requires that the offset, m_0, is found that satisfies

$$m_0 + i = m_i, \qquad (6)$$

where i represents the spectral orders numbered on the detector (e.g., $i = 1$–36, $m_0 = 116$, $m_i = 117$–152). This is done by minimizing the slope in

$$y(i) = (m_0 + i)\lambda_{\mu_i}, \qquad (7)$$

where $\lambda\mu_i$ is the mean wavelength of the ith order (Brahm et al. 2017). A 2D wavelength is then calculated using an expansion of the grating equation using Chebyshev polynomials (Baranne et al. 1996). The 2D wavelength solution takes the form:

$$\lambda_{2D}(x, m) = \frac{1}{m} \sum_{i=1}^{n_m} \sum_{j=1}^{n_x} c_i(m) c_j(x) \chi_{ij} \qquad (8)$$

where x is pixel location, m indicates spectral order number, c_i indicates the Chebyshev polynomial at order i, n_m and n_x are the degrees of the polynomial, and χ_{ij} is coefficient matrix of the wavelength solution.

Next, the extracted etalon spectral orders are fit using a multi-Gaussian model, order by order, to determine each etalon peak in pixel space. The specific model used is a sum of Gaussians:

$$\sum_{i=1}^{N_{\text{Gauss}}} a_i \cdot e^{\frac{-(x-b_i)^2}{2 \cdot c_i^2}} + d_{\text{off}} + x \cdot e_{\text{slope}} \qquad (9)$$

where N_{Gauss} indicates the number of Gaussian functions chosen to fit an order, a_i is the amplitude free parameter, b_i is the centering parameter, c_i is the width parameter, d_{off} is a single global offset parameter, and e_{slope} is a global linear slope parameter designed to account any residual continuum during the fitting process. Figure 8 shows a single order of an extracted etalon spectrum with an overlapping best-fit model in red. As the etalon FSR results in an extremely regular and close spacing of emission peaks, differentiating individual peaks can be difficult. Using the U-Ne wavelength solution, an existing line list of etalon peak wavelength centers, or a combination of both, a wavelength is assigned to each etalon peak. Now the etalon can be used to derive a 2D wavelength solution in the

same form as Equation (8), but with a precision far exceeding the limitations of U-Ne (Figure 9). U-Ne frames used to create an initial wavelength calibration are valid as long as the instrument remains stable enough that etalon lines will not spatially drift so much they would be mistaken for an adjacent peak. As iLocater is intended to be very stable over long periods of time, U-Ne calibration will rarely be used in comparison to the etalon.

3.4.2. Instrument Drift

iLocater has three simultaneous traces for each spectral order, fed by two fibers, one connected to each telescope, and one simultaneous reference fiber that is set between the two science fibers. During science exposures, and between bracketed calibration frames, this reference fiber is used to correct for instrument drifts on timescales comparable to the integration time. Drifts are calculated using Equation (8) but holding the χ_{ij} coefficients constant and introducing a new free parameter, δv_{drift}, the velocity drift, to Equation (8):

$$\lambda_{2D}(x, m) = \frac{1}{m}\left(1 + \frac{\delta v_{\text{inst}}}{c}\right) \sum_{i=1}^{n_m} \sum_{j=1}^{n_x} c_i(m) c_j(x) \chi_{ij}. \qquad (10)$$

3.5. Post-processing

The post-processing sub-pipeline will produce the finalized spectral data product. The remaining steps included in post-processing are blaze function removal, continuum normalization, telluric corrections, and barycentric correction.

3.5.1. Telluric Correction

Telluric removal in the infrared requires careful consideration and is extremely difficult to achieve a correction residual of less than $1\,\text{m}\,\text{s}^{-1}$ (Bean et al. 2010). Currently, most groups attempt to model telluric lines with synthetic models using a comprehensive line list and radiative transfer with accurate atmospheric models. iLocater's telluric removal strategy will follow a similar process to TAPAS (Bertaux et al. 2014), TelFit (Gullikson et al. 2014), Molecfit (Smette et al. 2015), and TERRASPEC (Bender et al. 2012), using the line-by-line radiative transfer model (Clough et al. 2005); and the High Resolution Transmission (HITRAN) line database (Rothman et al. 2013) to generate and subtract a representative synthetic telluric spectrum. This method is typically more accurate than empirical correction (Gullikson et al. 2014; Smette et al. 2015), achieving line removal precision of 2%–5%. Poorly fitted lines will be masked out following the techniques in Bean et al. (2010), Seifahrt et al. (2010), Blake et al. (2010).

None of the current techniques have yet reached the necessary precision for sub meter-per-second measurements in the NIR. The reasons for this are primarily missing lines in
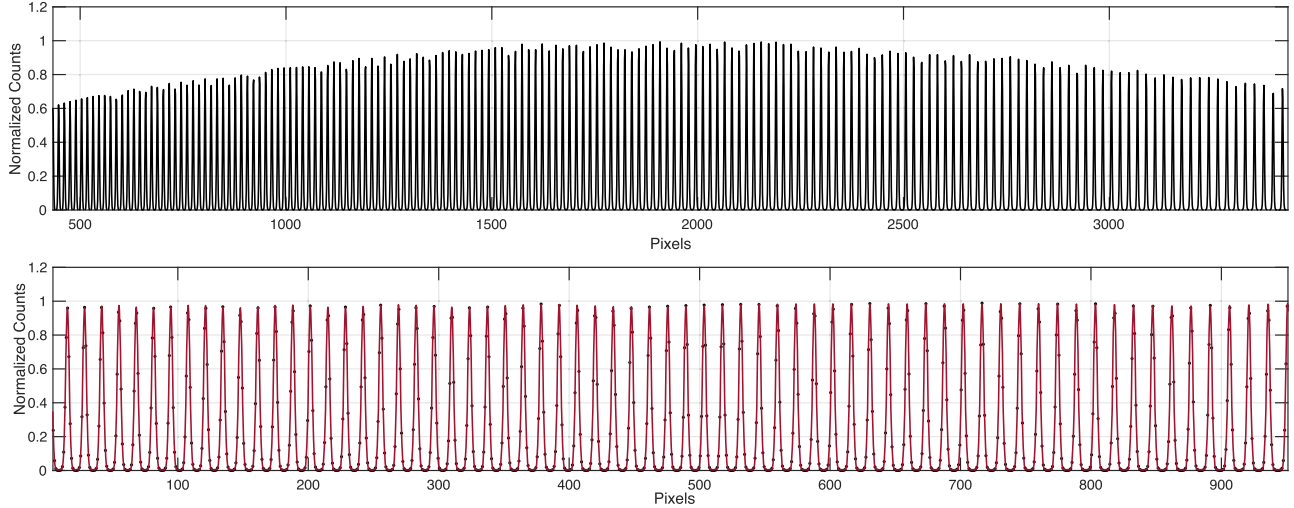
Figure 8. Top: extracted, noise-free simulation of a 10 GHz etalon, spectral order 152. Bottom: blaze corrected, normalized, sub-frame of the same spectral order demonstrating an automated, simultaneous 233 Gaussian model fit to derive pixel centers of each etalon line. The x-axis shows pixel numbers along the dispersion direction.

(A color version of this figure is available in the online journal.)

the HITRAN database, uncertainties or errors in attributes such as the line position, strength, or shape, limitations in current modeling codes for deriving correct line profiles (i.e., velocity dependence, line mixing effects), insufficient knowledge of real time atmospheric conditions (e.g., water column density variations), and wind-induced line shifts (Fischer et al. 2016).

3.5.2. Sky-background Subtraction

iLocater is using a 5.8 μm diameter single-mode fiber that has a very small on-sky angle of 1.4×10^{-3} arcsec^2, compared to a typical large diameter multi-mode fiber angle of 1 arcsec^2. Because of this, there is significant natural suppression of sky-background emission, making dedicated sky measurements and removal procedures largely unnecessary. However, in the case that iLocater requires simultaneous sky measurements, one of the science fibers could be used to sample sky background while the other collects starlight using the LBT's differential pointing capabilities.

3.5.3. Blaze Removal

The efficiency function of each order in an echelle spectrograph is normally dominated by a characteristic slope known as the blaze function which must be corrected prior to RV computation. A measurement of the blaze function can be made by injecting a flat field source into the fibers. The effect of the blaze function can be mitigated by normalizing and dividing the science spectrum by the fiber flats (Figure 10). However, this will not completely correct the spectrum because of the difference in continuum between the observed star and flat field source. Varying polarization states in time resulting

from intrinsic single-mode fiber properties will also cause a slightly different blaze measurement each time it is measured, emphasizing the need for further continuum normalization.

3.5.4. Continuum Normalization

The remaining residual continuum left in the spectrum after de-blazing must be removed as it can result in spurious Doppler measurements. The continuum is sampled using a moving box that rejects absorption lines. The bottom panel of Figure 10 shows very minor residual slopes being removed. Continuum normalization should be applied very cautiously as in some cases the correction induces its own residual tilt if the continuum is improperly sampled.

3.6. Barycentric Correction

When using the RV method to search for planets orbiting nearby stars, the dominant signal is due to the Earth's motion about the solar system barycenter. Barycentric correction consists of computing the observatory velocity with respect to the solar system barycenter, projected in the direction of the observed star. The two principal velocities to be computed are the movement of the Earth around the barycenter and the Earth's rotation at the geographical coordinates of the observatory. We implement the Wright & Eastman (2014) correction which uses full, relativistic, calculations precise to 1 cm s^{-1}.

3.6.1. Exposure Meter

Barycentric corrections can only be calculated for a single instant in time, however, iLocater's exposure times are
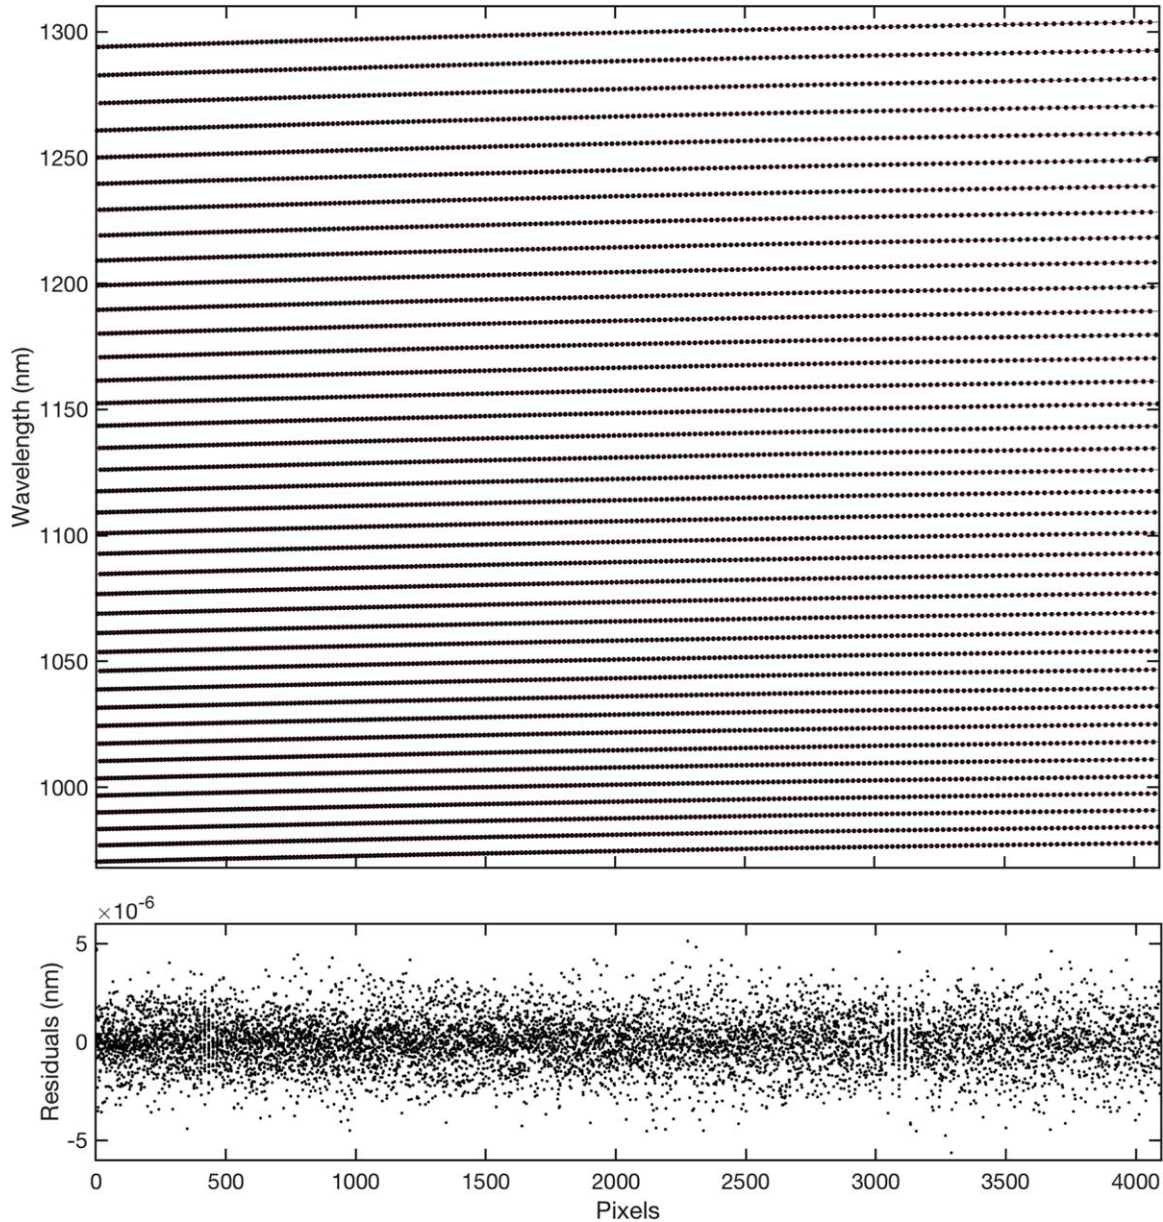
Figure 9. 2D wavelength solution derived from a simulated etalon frame. Points represent etalon peaks in wavelength and *x*-pixel direction while the black line denotes the 2D wavelength solution. Residuals are shown below and have an rms of 1×10^{-6} nm.
(A color version of this figure is available in the online journal.)

expected to last 30 minutes or longer. Additionally, the barycentric correction does not scale linearly with time so a flux-weighted average needs to be calculated throughout the observation (approximately every 1 minute to achieve $1\,\mathrm{cm\,s}^{-1}$ correction error). To calculate this flux-weighted average, iLocater will use readouts from a femtowatt photoreciever, located just before the entrance fiber in the spectrograph. For a ground-based observation, atmospheric extinction will introduce a wavelength dependence in the transmittance of photons to the instrument, possibly requiring a wavelength dependent

barycentric correction. iLocater's H4RG detector is capable of non-destructive readouts which could potentially be used as a chromatic exposure meter by calculating a different barycentric correction for each spectral order.

3.7. Radial Velocity Determination

3.7.1. Building a Binary Mask

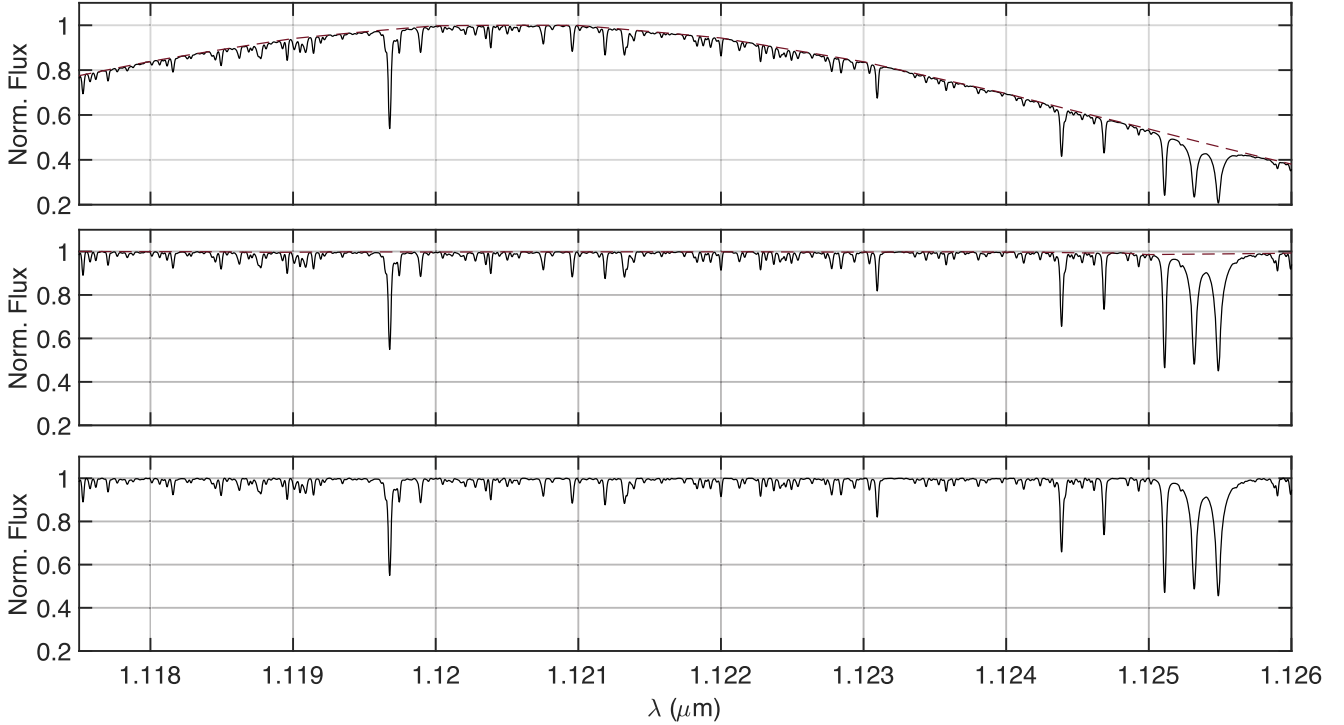A binary mask used in cross-correlation serves the purpose of only considering the RV information contained in certain

Figure 10. Top panel: extracted spectrum of an MOV star for a particular echelle order. Center panel: same order divided by the blaze function derived from a normalized flat field image. Bottom panel: continuum normalized spectrum by fitting the deblazed continuum with a polynomial.
(A color version of this figure is available in the online journal.)

selected stellar absorption lines. This helps avoid blended lines, lines with minimal depth compared to the continuum, and lines with other unsuitable characteristics. In general, the mask should include as many "clean" lines as possible to maximize the RV signal. Additionally, lines should be weighted by their depth relative to the normalized continuum.

We have written a MATLAB function to scan across a synthetic stellar spectrum, given a spectral type, rotational velocity, and bandpass, and builds a line list for masked cross-correlation. First, the synthetic spectrum is normalized, its sign inverted and brought to a range of [0 1]. The spectrum is divided into 0.1 nm pieces and a peak-finding function computes all observable peaks. A small interval is chosen around the first and last peak in the piece and is fit with a multi-Gaussian function where the number of Gaussians is equal to the number of peaks found. Each Gaussian has three parameters describing its amplitude, center, and width and two parameters for a constant offset and linear slope to remove any residual local trends in the inverted continuum. An example of this process is shown in Figure 11.

The line list is formed by appending the Gaussian center parameters as well as the ratio of the amplitude to the offset value. In order to avoid blended lines, which are present in many spectral chunks, the software identifies pairs of lines whose separation is less than a critical value. In this situation, only the line with largest amplitude is left in the mask list. To

further clean the line list from unsuitable lines, only lines with a value of the FWHM roughly in agreement with that expected for the rotational velocity of the star are kept. Figure 12 shows a portion of a simulated order (order 152) of a normalized M0V spectrum with the final line list overlaid in red.

3.7.2. Cross Correlation Function

The RV information of an observed star is contained within the wavelength positions of its spectral absorption lines, specifically in the line "wings," i.e., where the derivative is maximized. An efficient way to measure the variation of spectral absorption lines in time is using the cross-correlation function (CCF), with a binary mask taking values equal to 1 in the regions where a typical stellar spectra contains narrow absorption lines and equal to 0 elsewhere. For the RV extraction sub-pipeline, we have written our own version of the standard masked CCF technique Baranne et al. (1996), Pepe et al. (2002). This technique is currently used by the HARPS team and has performed very well in precision RV applications (Bonfils et al. 2013). The CCF is given by

$$CCF(v) = \frac{\int_{\lambda_i}^{\lambda_f} w(\lambda')m(\lambda')s(\lambda)d\lambda}{\sqrt{\int_{\lambda_i}^{\lambda_f} m(\lambda')^2 \, d\lambda}}, \qquad (11)$$

where λ_i and λ_f are the initial and final wavelengths of an order, $s(\lambda)$ is the spectrum, $w(\lambda')$ is a list of weights stored in the
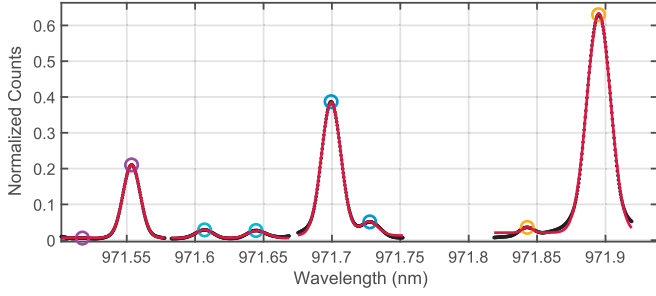
Figure 11. Small subsection of the mask building process showing an inverted synthetic stellar spectrum along with all identified peaks (circles) and multi-Gaussian fits (red line).
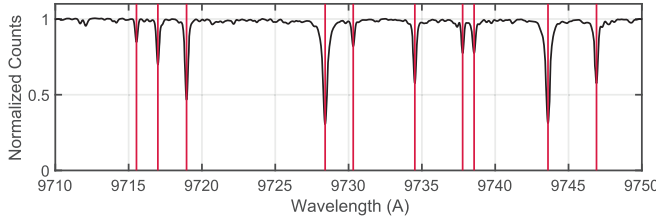(A color version of this figure is available in the online journal.)

Figure 12. Masked stellar lines for a portion of order 152 overlaid on an M0V synthetic spectrum.
(A color version of this figure is available in the online journal.)

binary mask, $m(\lambda')$, and λ' is the Doppler shifted wavelength (Brahm et al. 2017).

4. Verification of Simulator and Pipeline Performance

4.1. Numerical Simulation Errors

Measuring Doppler shifts on the order of $10\,\mathrm{cm\,s^{-1}}$ requires a very precise software pipeline. This also requires that the simulation code is not introducing algorithmic or numerical errors at a level that would overwhelm this measurement. To quantify the simulator's numerical errors in RV values, we simulated data frames of an M0 star. These frames contain no physical noise, no throughput modifications, or any modifying effects in order to isolate computational errors from simulated instrument systematic errors. 20 total frames were simulated, varying only injected RV, between $-30\,\mathrm{km\,s^{-1}}$ and $30\,\mathrm{km\,s^{-1}}$, chosen to reflect the typical magnitude of RV fluctuations under barycentric motion. The data reduction pipeline was then used to extract each order, apply an ideal wavelength solution, and compute the RV using masked cross correlation, which also serves as a test of the pipeline's ability to recover RV's under noiseless conditions. The residual error measured between the pipeline's recovered RV and the injected RV into the simulator are shown in Figure 13. The residual scatter shows structural errors do not exceed $10\,\mathrm{cm\,s^{-1}}$, with an rms = $0.03\,\mathrm{cm\,s^{-1}}$.
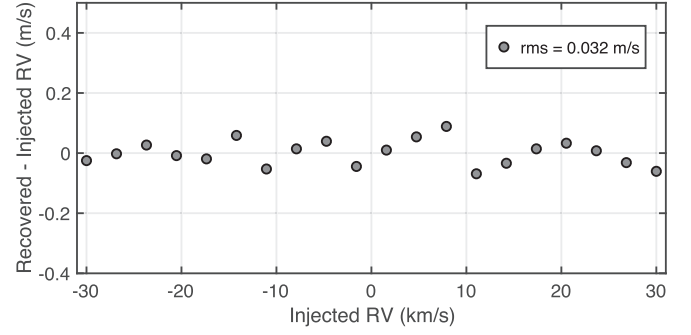

Figure 13. RV residuals from simulated noiseless data frames where a synthetic stellar spectrum (M0V) was shifted by precise Doppler velocities between -30 and $+30\,\mathrm{km\,s^{-1}}$. Velocities were then recovered using masked cross correlation (Section 3.7.2). These residuals represent the inherent algorithmic noise injected through data reduction processing.

From this, we conclude numerical simulation errors are sufficiently suppressed for the purposes of probing design choices and developing robust extraction/analysis software. Note that this test was performed under a simulation scale factor of 1, described in Table 1. Numerical errors can be further reduced by increasing the scale factor but at the cost of computation time, where the time to complete a simulated frame approximately increases with the square of the scale factor.

4.2. Masked CCF Performance

Another test of the simulator and pipeline performance is to verify the quality of stellar masks derived in Section 3.7.1 using photon noise to investigate various signal-to-noise (S/N) levels. This experiment is conducted by varying the average S/N, per collapsed pixel, in simulated frames. Here, collapsed pixel is referring to the conversion of 2D extracted spectral orders to 1D collapsed orders by recombining signal in the cross-dispersion direction. Masked CCF performance is measured against the Doppler information content contained in iLocater's spectral orders, measured using the formalism in Butler et al. (1996):

$$\sigma_{\mathrm{ph}} = \frac{1}{\sqrt{\sum\left(\frac{dI\,/\,dV}{\epsilon_I}\right)^2}}, \qquad (12)$$

where dI/dV represents the slope of the measured stellar intensity as a function of wavelength (expressed in velocity units) and $\epsilon_I = \sqrt{N_{\mathrm{ph}}}/N_{\mathrm{ph}}$ is the fractional Poisson error. Figure 14 shows the achievable RV precision computed at each S/N as well as the recovered RV from masked cross-correlation, measured by the rms of 50 repetitions at each S/N. This shows the simulated photon-noise routines, as well as the derived masks are close to the expected performance of a masked CCF routine.
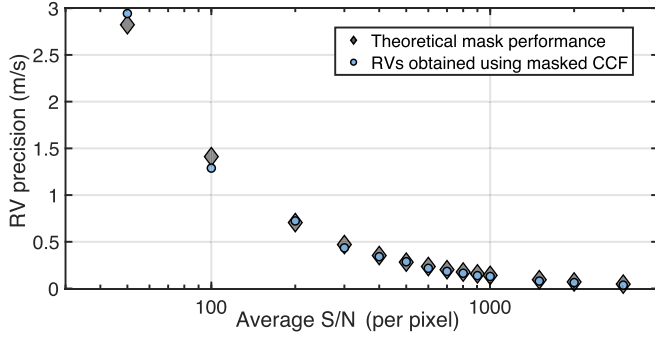
Figure 14. rms scatter in RV residuals for a set of simulation runs with varying levels of photon noise, as a function of collapsed S/N per pixel. The implementation of a masked CCF RV measurement routine approximates the expected ideal performance for the mask used. This set of results focus on mask performance and do not include sources of error other than photon noise.
(A color version of this figure is available in the online journal.)

Figure 15. Achievable RV precision for each of iLocater's 36 spectral orders. The target chosen was an M0V star with an average S/N between 190 and 195 across the spectral orders. Combining the RV information in every order gives a single measurement RV precision of $1.02 \, \mathrm{m \, s^{-1}}$.
(A color version of this figure is available in the online journal.)

5. Applications

5.1. Assessing Instrument Performance

One of the first and most important uses of the software tools described within is to determine achievable on-sky photon-noise limited signal-to-noise. This is accomplished by using the simulator to generate different spectral types at a range of apparent magnitudes, integrating the signal on the detector over the desired integration time, and then assessing the achievable RV precision using the formalism established in Section 4.2. Photon noise limited RV precision, σ_{ph}, is calculated using Equation (12). This could be used is to assess the photon-noise limited RV precision contained in each of iLocater's spectral orders for a specific spectral type and average S/N per order. Figure 15 shows each of iLocater's individual spectral orders containing an M0V star with approximately constant S/N. This plot shows there is a clear trend in achievable Doppler precision, decreasing with longer wavelengths as well as the necessary average S/N in order to achieve $\mathrm{m \, s^{-1}}$ RV precision.

5.2. Spectral Order Cross-contamination

Choosing the correct cross-disperser in echelle spectrograph design is an important step to avoid adjacent orders contaminating each other while maximizing information content by optimizing the number of orders that are captured by the detector. We use the simulator to visualize the optical design and inspect the inter-order and intra-order separation. For iLocater, we set an upper limit of less than 0.1% cross-contamination between orders. Therefore, we define the "edge" of an order to be the point at which the ratio of intensity relative to maximum is 0.1%. As the vertical profile of each order follows a Gaussian profile, this quantity is easily computed, given a known FWHM. Adopting a pessimistic value of 5 pixels sampled in a FWHM in the cross-dispersion direction,
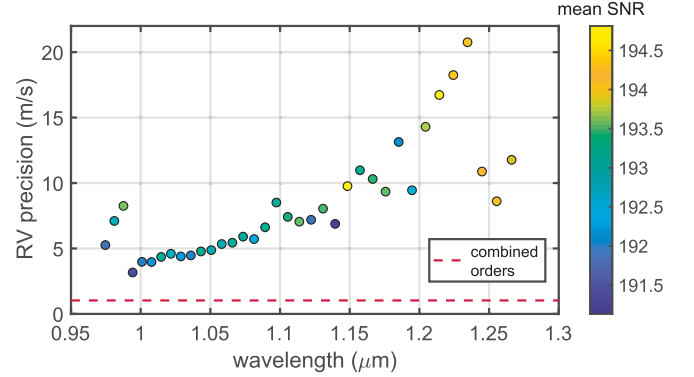
this gives a requirement of 16 pixels between adjacent order centers. To test this, we generate a noiseless full frame detector image with a flat-field spectrum in each fiber trace. Vertical cross cuts are taken with a focus on spectral orders with the shortest wavelengths as they will experience the smallest separations. Figure 16 shows a cross-cut of all 36 spectral orders and 3 traces within each order. The right side shows a plot of the first three spectral orders, where the closest separation exists between the third trace of order one and the first trace of order two. Fitting these with Gaussian profile and computing the separation gives a value of 27.3 pixels, well within the contamination requirement.

5.3. Optimizing Fabry–Pérot Finesse and Free Spectral Range

We derive optimal finesse ($\mathcal{F}$) and free spectral range (FSR) values for our Fabry–Pérot etalon and verify these calculations using visuals created with the simulator. First, we define three straightforward requirements for the etalon:

1. The etalon needs to span the primary science bandpass of iLocater (0.97–1.27 μm) to provide satisfactory wavelength calibration information in every spectral order.
2. To maximize information content in the calibration source and simplify the peak-fitting process, the etalon modes should not be resolved by the spectrograph.
3. In order to optimize the number of peaks in each order but avoid overlapping, we require that the minimum line separation (peak-to-peak) be at least 3 PSF FHWM of the spectrograph PSF.

From these requirements, we can derive the spectral characteristics of an optimal Fabry–Pérot etalon cavity, specifically the finesse, which sets the etalon peak width, and free spectral range which defines peak-to-peak separation.
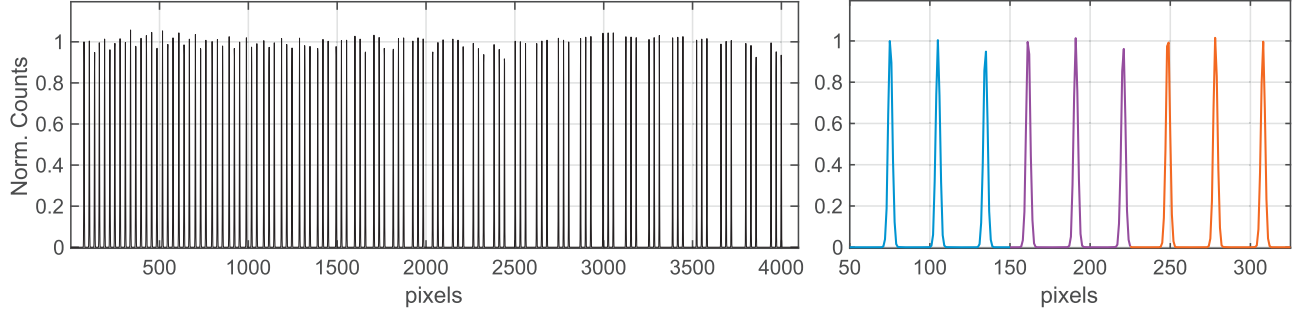
17

Figure 16. Left: slice of uniformly illuminated fiber traces on detector in the cross dispersion direction showing 36 spectral orders with 3 traces in each order. Note, there is no modulation in counts as throughput effects are not included here. Right: windowed plot focusing on the first three spectral orders, going left to right. Each spectral order is identifiable by a different color.

(A color version of this figure is available in the online journal.)

5.3.1. Finesse

Following the formalism presented in Section 3.1 of Cersullo et al. (2017) and using iLocater's bandpass and an under-sampling factor of 5 instead of 3, we obtain the following condition for the finesse:

$$\mathcal{F} = 15 \cdot \frac{\lambda_R}{\lambda_B}. \tag{13}$$

Using iLocater's wavelength bounds ($\lambda_B = 0.97\,\mu$m and $\lambda_R = 1.27\,\mu$m) we specify a minimum finesse of 20. Due to small batch-to-batch reflectivity ripple in the custom-manufactured etalon mirror surfaces, the finesse of the final system is being targeted at ∼40, with an expected margin of 20–70. This range fulfills our requirement.

5.3.2. Free Spectral Range

Condition (3) requires there to be a minimum separation equivalent to 3 PSF FHWMs at the bluest wavelength. For this part of the calculation we will adopt a more conservative definition of PSF width, moving from FWHM to $1/e^2$ (≈13.5%). Using the following relation to compute Gaussian width at intensity, I, for a known σ,

$$\text{Gauss width}(I) = 2 \cdot \sigma \cdot \sqrt{2 \cdot \log\left(I^{-1}\right)}, \tag{14}$$

with $I = 1/e^2$ gives an equivalent separation of 5.1 FWHMs, slightly modifying the value given in condition (3). Inserting this value in Equation (7) of Cersullo et al. (2017) gives

$$\text{FSR}_B \geqslant 5.1 \cdot \frac{\lambda_b}{R}, \tag{15}$$

where λ_B is the bluest spectral bandpass wavelength and R is the spectrograph resolution. Computing Equation (15) with $R = 150{,}000$ and $\lambda_B = 970$ nm results in a FSR = 10.5 GHz.

To visualize the results derived above, we generate a synthetic etalon spectrum with an $\mathcal{F} = 40$ and FSR = 10 GHz and pass it through our spectrograph simulation code.

Figure 17 shows two 100×100 pixel windows taken from the full simulated detector frame. The left frame is centered on the shortest wavelengths of the bluest spectral order, showing that even at the minimum etalon peak separation, each peak is well separated from adjacent peaks. The right frame centers on the longest wavelengths of the reddest order, showing the maximum separation of each order.

5.4. H4RG Detector Noise

iLocater is implementing an H4RG-10 NIR detector. Although it is relatively untested in Doppler spectrographs compared to the H2RG, it offers the necessary array and pixel size to accommodate iLocater's spectral resolution, bandpass and pixel sampling needs. An independent study has been undertaken to focus on the suitability of H4RG detectors used in precision RV work by translating their noise characteristics into RV errors (E. Bechter et al. 2019a, in preparation). This work makes use of the spectrograph simulator, HxRG noise generator, and data reduction pipeline to characterize detector noise and unique NIR detector effects and translate them directly into RV errors in addition to investigating potential mitigation strategies.

5.5. RV Error Budget

Modern Doppler spectrographs are pursuing extraordinary precision in single RV measurements, seeking to reduce instrument systematic errors and maximize resolving power to disentangle stellar activity signals from observed spectra (Davis et al. 2017). Many instruments are investigating new and relatively untested designs to push the current precision boundaries. To verify their fidelity, quantified RV error budgets are becoming standard practice during the spectrograph design phase, in which each design choice is assessed by its impact on RV precision (Halverson et al. 2016). Using the simulation tools outlined in this paper, a comprehensive RV error budget that includes: photon noise, instrument systematic terms,
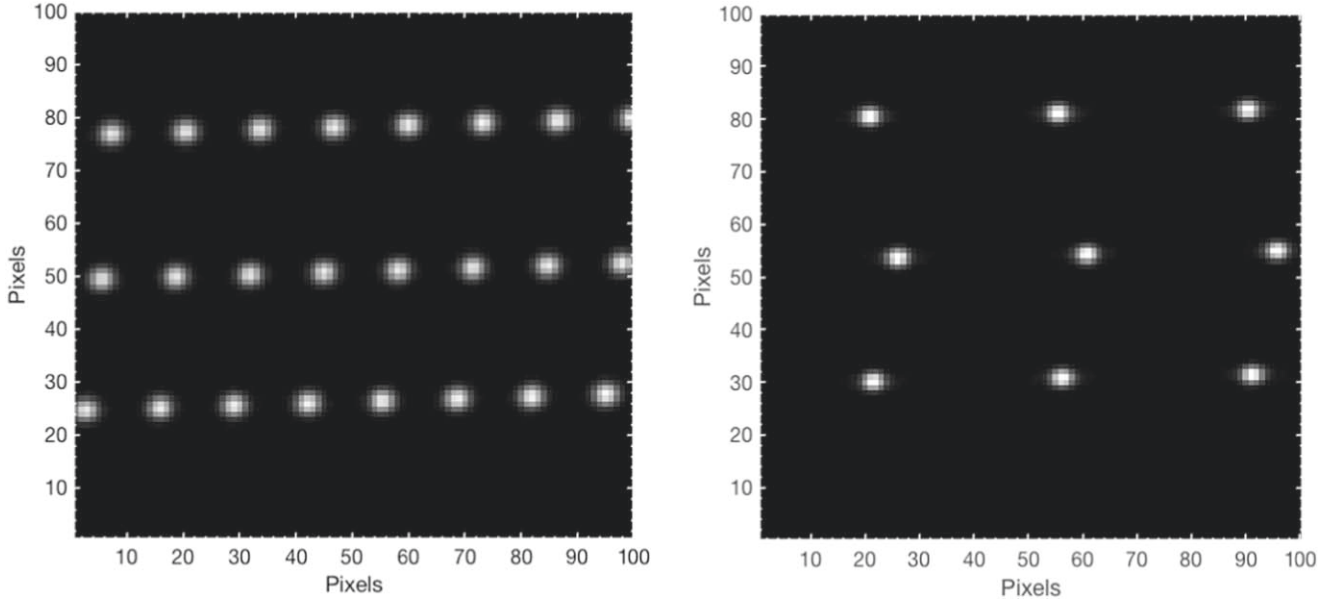
Figure 17. Left. Simulated detector frame centered at the shortest wavelengths of iLocater's bandpass. This region corresponds to the closest etalon spacing and visually confirms sufficient separation according to requirement (2). Right. Simulated detector frame centered at the longest wavelengths. Etalon peaks still appear unresolved, adopting the shape of the instrument PSF.

software and barycentric removal residuals, and time-varying atmospheric contaminants has already been assembled for iLocater, the details of which can be found in Bechter et al. (2018).

5.6. Optical Aberrations

Optical aberrations in spectrographs originate from many possible sources including: manufactured optical surfaces, misalignments in the optical system, thermal effects, etc. If the aberrations vary over time, they can impart asymmetries on the spectral signal through the instrument PSF, leading to a shift in the center of the light distribution on the detector focal plane. Typically, spectrograph optical designs are optimized using spot sizes or wavefront error budgets. We use the simulator and pipeline in a sensitivity study specifically designed to investigate the impact of optical aberrations on RV measurements (E. Bechter et al. 2019b, in preparation).

6. Summary

Astronomers are building planet-finding spectrometers that aim to measure Doppler shifts at the level of $1\,\mathrm{m\,s}^{-1}$ and below. The ultimate goal of studying the masses and orbits of terrestrial worlds in the habitable zone requires RV precisions that are an order of magnitude better than the current state of the art. This level of performance corresponds to routinely measuring translational line shifts of only several atomic radii on the detector. A number of subtle effects can impact the retrieval of these extraordinarily small RV variations.

Data reduction pipelines play a key role in being able to reliably extract the motion of stellar absorption lines in RV time-series data. Software represents not only an error term itself, but also permits the evaluation of many other terms in the error budget through numerical simulations. Developing a data reduction pipeline early-on in the instrument design and development process is thus essential as it provides insight into the effects that limit precision.

We have developed a comprehensive simulation code and RV data reduction pipeline for the iLocater spectrograph, an AO-fed SMF Doppler instrument being constructed for the LBT in Arizona. This paper provides a detailed description of the code structure and overview of the how the various classes interact. This infrastructure has been used to inform design decisions for the spectrograph and quantify RV error budget terms.

Using conventional methods for cross-correlating stellar spectra, we find that iLocater's software pipeline typically results in residual RV variations of several centimeters per second. This level of precision is well below the effects introduced from astrophysical jitter or that expected from photon noise and instrument stability. Among other things, these tools have been used to study the impact of signal-to-noise (e.g., throughput budget) and to optimize and verify wavelength calibration line spacings, order spacing, masked-CCF performance, and fiber contamination effects. In the future, the pipeline will be used to further study the effects of optical aberrations, barycentric correction, telluric contamination, the use of HgCdTe devices, and characterizing absorption line asymmetries at high spectral resolution.

ORCID iDs

Eric B. Bechter ⓘ https://orcid.org/0000-0001-8725-8730

References

Allard, F., Homeier, D., & Freytag, B. 2012, RSPTA, 370, 2765
Artigau, É., Bouchy, F., Delfosse, X., et al. 2012, Proc. SPIE, 8451, 84513I
Baranne, A., Queloz, D., Mayor, M., et al. 1996, A&As, 119, 373
Bean, J. L., Seifahrt, A., Hartman, H., et al. 2010, APJ, 713, 410
Bechter, A. J., Bechter, E. B., Crepp, J. R., King, D., & Crass, J. 2018, Proc. SPIE, 10702, 107026T
Bender, C. F., Mahadevan, S., Deshpande, R., et al. 2012, ApJL, 751, L31
Bertaux, J. L., Lallement, R., Ferron, S., Boonne, C., & Bodichon, R. 2014, A&A, 564, A46
Biesiadzinski, T., Lorenzon, W., Newman, R., et al. 2011, PASP, 123, 958
Blake, C. H., Charbonneau, D., & White, R. J. 2010, APJ, 723, 684
Bland-Hawthorn, J., & Horton, A. 2006, Proc. SPIE, 6269, 62690N
Bland-Hawthorn, J., Lawrence, J., Robertson, G., et al. 2010, Proc. SPIE, 7735, 77350N
Bolton, A. S., & Schlegel, D. J. 2010, PASP, 122, 248
Bonfils, X., Curto, G. L., Correia, A. C. M., et al. 2013, A&A, 556, A110
Brahm, R., Jordán, A., & Espinoza, N. 2017, PASP, 129, 034002
Butler, R. P., Marcy, G. W., Williams, E., et al. 1996, PASP, 108, 500
Cersullo, F., Wildi, F., Chazelas, B., & Pepe, F. 2017, A&A, 601, A102
Clough, S. A., Shephard, M. W., Mlawer, E. J., et al. 2005, JQSRT, 91, 233
Crepp, J. R. 2014, Sci, 346, 809
Crepp, J. R., Crass, J., King, D., et al. 2016, Proc. SPIE, 9908, 990819
Davis, A. B., Cisewski, J., Dumusque, X., Fischer, D. A., & Ford, E. B. 2017, ApJ, 846, 59
Fischer, D. A., Anglada-Escude, G., Arriagada, P., et al. 2016, PASP, 128, 066001
Gibson, S. R., & Wishnow, E. H. 2016, Proc. SPIE, 9911, 99112D
Gray, D. F. 2005, The Observation and Analysis of Stellar Photospheres (Cambridge: Cambridge Univ. Press)
Gullikson, K., Dodson-Robinson, S., & Kraus, A. 2014, AJ, 148, 53
Halverson, S., Terrien, R., Mahadevan, S., et al. 2016, Proc. SPIE, 9908, 99086P
Horne, K. 1986, PASP, 98, 609
Jovanovic, N., Schwab, C., Cvetojevic, N., Guyon, O., & Martinache, F. 2016, PASP, 128, 121001
Kos, J., Bland-Hawthorn, J., Betters, C. H., et al. 2018, MNRAS, 480, 5475
Marsh, T. R. 1989, PASP, 101, 1032
MATLAB 2017, version 9.3.0.713579 (R2017b) (Natick, MA: The MathWorks Inc.)
Pecaut, M. J., & Mamajek, E. E. 2013, ApJS, 208, 9
Pepe, F., Mayor, M., Galland, F., et al. 2002, A&A, 388, 632
Pinna, E., Esposito, S., Hinz, P., et al. 2016, Proc. SPIE, 9909, 99093V
Piskunov, N. E., & Valenti, J. A. 2002, A&A, 385, 1095
Plazas, A. A., Shapiro, C., Smith, R., Rhodes, J., & Huff, E. 2017, JInst, 12, C04009
Rauscher, B. J. 2015, PASP, 127, 1144
Rothman, L. S., Gordon, I. E., Babikov, Y., et al. 2013, JQSRT, 130, 4
Schwab, C., Leon-Saval, S. G., Betters, C. H., Bland-Hawthorn, J., & Mahadevan, S. 2014, in IAU Symp. 293, Formation, Detection, and Characterization of Extrasolar Habitable Planets, ed. N. Haghighipour (Cambridge: Cambridge Univ. Press), 403
Seifahrt, A., Käufl, H. U., Zängl, G., et al. 2010, A&A, 524, A11
Smette, A., Sana, H., Noll, S., et al. 2015, A&A, 576, A77
Terrien, R. C., Bender, C. F., Mahadevan, S., et al. 2014, Proc. SPIE, 9152, 915226
Vacca, W., Cushing, M., & Rayner, J. 2003, PASP, 115, 389
Wright, J. T., & Eastman, J. D. 2014, PASP, 126, 838
Zechmeister, M., Anglada-Escudé, G., & Reiners, A. 2014, A&A, 561, A59