

**StarCAT:**  
**A Catalog of Space Telescope Imaging Spectrograph**  
**Ultraviolet Echelle Spectra of Stars**

Thomas R. Ayres

*Center for Astrophysics and Space Astronomy,  
389 UCB, University of Colorado, Boulder, CO 80309;  
Thomas.Ayres@Colorado.edu*

**ABSTRACT**

StarCAT is a catalog of high resolution ultraviolet spectra of objects classified as “stars,” recorded by Space Telescope Imaging Spectrograph (STIS) during its initial seven years of operations (1997–2004). StarCAT is based on 3184 echelle mode observations of 545 distinct targets, with a total exposure duration of 5.2 Ms. For many of the objects, broad ultraviolet coverage has been achieved by splicing together echellegrams taken in two or more FUV (1150–1700 Å) and/or NUV (1600–3100 Å) settings. In cases of multiple pointings on conspicuously variable sources, spectra were separated into independent epochs. Otherwise, different epochs were combined to enhance signal-to-noise (S/N). The processing included a post facto correction to the “calstis” pipeline to compensate for subtle wavelength distortions identified in a recent study of the STIS dispersion relations. An “internal” calibration procedure yielded coherent spectral energy distributions for objects with broadly overlapping wavelength coverage. Although StarCAT represents a milestone in the large scale post-processing of STIS echellegrams, a number of potential improvements are identified, especially related to the ‘final’ calstis pipeline.

*Subject headings:* stars: individual — ultraviolet: spectra

**1. INTRODUCTION**

StarCAT is a Cycle 14 legacy archival project supported by the Guest Investigator program of Hubble Space Telescope (*HST*). The objective of StarCAT was to create an easily accessible collection of high resolution spectral observations of targets broadly identified as

“stars,” obtained by Space Telescope Imaging Spectrograph (STIS) from the time of its installation in 1997, during Servicing Mission 2, to its premature shutdown in August 2004 owing to a power supply failure. This report describes how the StarCAT project was carried out, including selecting the sample, processing the individual echellegrams, coadding multiple exposures, and splicing independent spectral segments to achieve broad ultraviolet coverage (1150–3100 Å). StarCAT is available through an interface maintained by the author (<http://arlsrv.colorado.edu/~ayres/StarCAT/>), which is linked to high level science products (HLSP) stored at the Multimission Archive at Space Telescope (MAST).

## 2. Space Telescope Imaging Spectrograph

The STIS instrument and its in-flight performance are described by Woodgate et al. (1998). A brief overview of the features and operational modes of STIS is provided here, as a context for some of the special characteristics of the pipeline processing and post-processing steps utilized in constructing StarCAT.

The most important predecessor high resolution ultraviolet spectrometers in space were *HST* first generation Goddard High-Resolution Spectrograph (GHRS: Brandt et al. 1994) and the venerable *International Ultraviolet Explorer* (*IUE*: Boggess et al. 1978). STIS was a dramatic advance because it united two formidable technologies tested in the previous instruments: a cross dispersed echelle optical design and high performance digital panoramic detectors. Unlike GHRS, which carried only small format (500 pixel) linear “Digicon” arrays, STIS could achieve broad spectral coverage in a single exposure, even in the highest dispersion modes, thanks to its 2K×2K (in high resolution readout mode) Multi-Anode Microchannel Array (MAMA) cameras, and the compact echelle format (a series of several dozen spectral slices covering narrow, but overlapping, wavelength ranges, optically stacked one on top of another). And, unlike *IUE*, which also could achieve broad spectral coverage with its moderate resolution echelles ( $R \sim 10^4$ ) and then state-of-the-art TV-style vidicon cameras, the STIS MAMAs were capable of high S/N and dynamic range, with almost no background.<sup>1</sup>

In addition to its high resolution echelle modes, STIS could record moderate and low resolution UV spectra utilizing first order dispersers (so-called “G” modes), with higher sensitivity but reduced spectral coverage. It also could reach to longer visible wavelengths with a third camera, a CCD, although only with first order gratings. These lower resolution options will not be considered further, however, because the comparatively narrow spectral

---

<sup>1</sup>At least for the “solar blind” FUV MAMA; the longer wavelength NUV channel suffered elevated backgrounds due to anomalous fluorescence in the camera faceplate.

bands typically captured in each object do not lend themselves to a broad “catalog of atlases” such as envisioned for StarCAT.

Importantly, STIS maintained the wavelength calibration heritage pioneered by *IUE* and continued later with the first generation *HST* spectrometers (GHRS and Faint Object Spectrograph (FOS)); utilizing a hollow cathode discharge source<sup>2</sup> for the purpose. The onboard lamps served three key functions: (1) calibrating the dispersion relations (mapping detector pixel coordinates into wavelengths) of each distinct instrumental setting, on orbit; (2) determining the wavelength zero point of an exposure after movements of the Mode Select Mechanism (MSM, which rotated the appropriate grating into the optical path, to define a specific setting); and (3) compensating for drifts in the zero point, due to short term thermomechanical flexing of the support structure. The latter effect was significant for the *IUE* echelles, but less so for STIS with its more sophisticated design and better control of the thermal environment in *Hubble’s* axial instrument bay. Compensation for the MSM positioning was essential, however: the zero point was not perfectly repeatable from one grating setting to the next (but did not need to be thanks to the “wavecal” procedure).

For echelle spectroscopy, STIS could be operated in two broad wavelength regions: FUV (1150–1700 Å; designated “140” (nm)) and NUV (1600–3100 Å; “230” (nm)), corresponding to the two MAMA cameras in the package. The shortwavelength channel utilized a “solar blind” photocathode material that sharply reduced the sensitivity to longer wavelength ( $\lambda > 1700$  Å) “out-of-band” radiation that might scatter its way into an FUV spectrogram. This was a major improvement over the *IUE* SW spectrometer, and *HST* FOS, especially for observations of NUV bright, but FUV faint, chromospheric emission-line stars (e.g., Ayres et al. 1995; Ayres et al. 1996).

The two broad wavelength bands could be recorded in medium resolution (“M”) or high resolution (“H”), for a total of four “uber” echelle modes: E140M, E140H, E230M, and E230H. The two H modes had similar resolution,  $R \equiv \lambda/\Delta\lambda = 114,000$ , when used with the default narrow spectroscopic aperture (designated 020×009, where the leading value is the slot height in units of 10 marcs, and the trailing value is the width). The FUV M resolution was about half this (45,800), while NUV M was lower still (30,000); again using the default narrow slit for these modes (020×006).

Each of the uber modes utilized sets of discrete MSM rotations to cover its wavelength grasp. The scan positions (or “tilts”) were designated by central wavelength, for example E140M-1425 was the E140M setting with  $\lambda_{\text{cen}} = 1425$  Å. Multiple tilts were required because

---

<sup>2</sup>For STIS, a platinum cathode alloyed with a small amount of chromium, in a neon carrier gas: all three elements contribute to the emission line spectrum of the discharge.

it was not always possible to fit all the echelle stripes on the detector at the same time. The exception was E140M-1425, which could capture the full 1150–1700 Å band in a single exposure (although with a few “gaps” longward of  $\sim 1600$  Å where there was incomplete overlap between adjacent orders owing to a slight mismatch between the trapezoidal echelle format and the square active area of the camera). Likewise, the 1600–3100 Å interval could be covered in just two “prime” tilts of E230M,  $\lambda 1978$  and  $\lambda 2707$ . The H modes were more setting intensive. E140H required three prime scans to cover the FUV band, and E230H needed six for the NUV region.

In addition to the prime tilts, a number of secondary settings were supported (none for E140M, four for E230M, eight for E140H, and 20 for E230H). These provided extra flexibility when, for example, an observer wished to focus on only one spectral feature (or a few) and could dial up a specific tilt to capture it (or them) all at once. A good example is the E230H secondary setting  $\lambda 2812$ , which was optimum for recording the Mg II  $\lambda\lambda 2796, 2802$  resonance lines, together with Mg I  $\lambda 2852$ ; all prominent in chromospheric emission-line stars, as well as interstellar gas (in absorption). This setting was so popular, in fact, that it was routinely calibrated as if it were a prime tilt.

Each distinct mode/scan position could be utilized with a wide range of apertures, depending on the objectives of the program. If high spectral purity was desired, one would choose the default spectroscopic slot; or perhaps the even narrower “Jenkins” slit (010×003) to isolate the extreme sharp core of the point response profile (also called psf). In some instances, the low transmission Jenkins slit was chosen simply to suppress light levels on the MAMA cameras, which were subject to count rate thresholds to minimize damage to the microchannel plates in the intensification stage. The default slot widths correspond to 2 pixels per resolution element in the normal 1K×1K image format (as used in the standard pipeline processing, although the normal data readout is in 2K×2K format), a minimalist approach to the Nyquist condition. The “high res” format could be utilized if better sampling were needed, for example with the Jenkins slit.

If high throughput was desired, then a wider aperture could be selected, say the 020×020 “photometric” slot; at the expense of partially including the broad shallow wings of the psf, thereby lowering the effective resolution somewhat. If very high S/N was important, one could pick the so-called “FP-split” apertures, two sets of five spatially separated slits of the same size (020×006 or 020×020). A sequence of independent exposures taken through the FP slits would place each echellegram on a different set of pixels. Combining the group after the fact on the same wavelength scale then would suppress fixed pattern noise<sup>3</sup>. An additional

---

<sup>3</sup>A consequence of small pixel-to-pixel nonuniformities that survive the flatfielding process, owing to the

set of special purpose apertures included different degrees of neutral density filtering, for observing very bright UV sources that ordinarily would trigger the MAMA safety limits.

Although there were a number of instances where the special apertures were used, by far the majority of the STIS exposures were taken either with the default narrow slit for the particular mode (especially for the sharp absorptions encountered in ISM work), or the photometric aperture (especially for late-type (F–M) stars, whose chromospheric spectra are comparatively faint). The standard apertures were better calibrated radiometrically than their “special” cousins, so in principle one can place more reliance on the pipeline flux scales. Nevertheless, even here one must be careful because slight drifts of the target image in a narrow slit over the course of a several orbit exposure, especially when only a single guide star could be used, might lead to enhanced light loss as the point source approached an aperture edge, thereby invalidating to some degree the default transmission factors applied by the pipeline.

Initial target acquisition typically was accomplished by direct imaging with the CCD, usually through a filter. A “peak-up” then could be performed to more precisely center the object in the specified aperture. This was important for two reasons: a well centered target minimized light loss, particularly with the narrow spectroscopic slits; and a well centered target maximized the accuracy of the radial velocity scale. The peak-up was accomplished by stepping a designated aperture (usually, but not necessarily, the desired science one) in a pattern over the source, and centroiding the resulting intensity map. The peak-up could be performed with the CCD in direct imaging mode, for fainter targets; or using dispersed light, for brighter point sources that otherwise would saturate the CCD.

Finally, the photon counting MAMA cameras could be operated either in ACCUM or T-TAG mode. In the former, a single echellegram was built up in memory by assigning each camera pulse to a corresponding image cell, and accumulating them, compensating for Doppler shifts devolving from *Hubble’s* low-Earth orbit. In T-TAG mode, on the other hand, each pulse individually was logged in a time-tagged “event list,” which sequentially tabulated pixel coordinates and arrival times. T-TAG was valuable for tracking highly variable phenomena, but rarely was used in the echelle modes because it was storage intensive and sacrificed some degree of wavelength fidelity. In StarCAT, for simplicity, all observations were treated as ACCUM exposures.

Figure 1 illustrates schematic examples of echellegrams from the four uber modes, for a range of different types of objects, albeit some type of “star” in each case.

---

difficulty of producing a high S/N flatfield map for the MAMAs on-orbit; see:  
[http://www.stsci.edu/hst/stis/documents/handbooks/currentDHB/ch4\\_stis\\_error2.html#413471](http://www.stsci.edu/hst/stis/documents/handbooks/currentDHB/ch4_stis_error2.html#413471)

### 3. OBSERVATIONS

#### 3.1. Selection of the Sample

An expedient approach was taken to assemble the StarCAT sample. As a first step, it relied on a selection by the Guest Observer (GO) of “STAR” in the target description (so-called “broad category”) of the Phase 2 form. A search was performed via the MAST “HSTonline”<sup>4</sup> portal for all STIS echelle mode observations designated “STAR;A\*,” “STAR;B\*,” and so on up to “STAR;Z\*” as the first priority key word in the broad category line (“\*” is the usual wild card designator). This resulted in about 2,000 separate observations, some of which consisted of multiple subexposures, of about 400 objects. Then, a second search was conducted, for targets specified as “HD\*,” “BD\*,” “WD\*,” “ISM\*,” “CAL\*,” and so forth, as the primary key word, based on an assessment of the *secondary* key words used by GOs when the primary was “STAR;\*.” The second search produced another 1000, or so, exposures of about 150 objects, which were missed in the initial pass. Probably a few additional candidates were overlooked by the two crosscutting searches, especially if the GO had chosen an inappropriate prime broad category; but it was judged a case of diminishing returns to track down these outliers.

Next, the identified STIS datasets<sup>5</sup> were retrieved from the HSTonline archive (circa mid-2008). Each FITS label in the collection was queried for the target name assigned by the GO, and the coordinates of the spacecraft pointing. The header-derived list was collapsed to a table of unique coordinates, which then was passed to SIMBAD to obtain properties of the targets.

At this stage, a unified naming convention was applied to the somewhat diverse target descriptors assigned by individual GOs. The most common were Henry Drapier Catalog (HD) numbers (e.g., HD 34029); variable star names (e.g., AD Leo); entries from the White Dwarf catalog of McCook et al. (latest edition: 2006 (2008 version in VizieR); e.g., WD0839+398 (HHMM+DDd)); BD (Bonner Durchmusterung) numbers; and proper names. The adopted convention was to select in the following order of priority: (1) HD number, (2) variable star name (including supernovae “SN”), or (3) WD number. All other cases, for which a SIMBAD entry unambiguously could be associated with the target, were designated “STARHHMM±DDMM” (e.g., STAR1326–4727: the abbreviated coordinates are  $\alpha_{2000}, \delta_{2000}$ ; epoch 2000). One object was assigned “NOID1911–5957” (tagged as NGC6752-2206 by the GO) because there was no SIMBAD object within 20” of the reported location, whereas

---

<sup>4</sup>See: <http://archive.stsci.edu/hstonline/>

<sup>5</sup>The so-called “x1d” pipeline file, which contains flux tracings of the individual echelle orders.

other targets of that *HST* program were clearly identified in SIMBAD. There apparently is a UV object at the NOID1911–5957 coordinates, nevertheless, based on the successful STIS observation. In a few cases, a GO program targeted several objects closely grouped on the sky, and the STARHHMM±DDMM designation was degenerate with respect to the specified pointings. In such instances, “a,” “b,” etc., were appended to the STAR designations to break the degeneracy. (The coordinates assigned in Table 1, below, have sufficient precision to clearly distinguish the individual objects.)

In a few special cases, the GO had specified pointings on and offset from the main target, and these were treated as separate objects (e.g., “HD39060/POS-CEN” for “BETAPICTORIS,” and “HD39060/POS-NE” for “HD39060-DISK-NE-1”). For one of these — HD39801 ( $\alpha$  Orionis) — a series of spatial scans (sequence of pointings at discrete offsets) had been conducted across the face of the partially resolved stellar disk. These were collected in subdirectory “SCAN.” A second subdirectory — “POS-CEN” — contains the (several) observations taken with the aperture nominally centered on the star. There also was a calibration part of the  $\alpha$  Ori program that obtained an identical spatial scan across a point-like star (STAR0810+7457 in StarCAT nomenclature). Again, the spatial scan frames were placed in a subdirectory called “SCAN.” However, in this case there was an independent series of echelle exposures of STAR0810+7457, from a different GO program. As with  $\alpha$  Ori, these were assigned to a subdirectory “POS-CEN,” although now the single central exposure from the spatial scan was not included, because the material from the other program was judged superior.

Table 1 summarizes the consensus target names and coordinates, with reference to the original GO descriptors and SIMBAD designators. Table 2 provides schematic stellar properties abstracted from SIMBAD.

There are 545 unique “STAR” objects (several of which are multiple pointings on the same target, but offset in some way as mentioned above) for which STIS echelle observations were available, and 3184 unique echellegrams (counting subexposures), about evenly split among the four uber modes. The maximum number of exposures for an object was 270 (HD208816 = VV Cep), although only one other target had more than 100. A third of the stars are represented by only a single exposure, and the majority of the sample has ten, or less (although it should be noted that the entire FUV + NUV interval in principle can be covered by as few as three exposures: one E140M and two E230Ms).

Having identified the sample, and collected the associated STIS exposures, a series of post-processing steps were undertaken to: (1) merge the individual orders of the x1d datasets; (2) coadd the subexposures for observations with multiple “repeats;” (3) coadd exposures in the same mode, setting, and aperture grouped within a single visit; (4) coadd

like-mode exposures taken in different epochs or with different apertures; and, finally, (5) splice together the different wavelength segments available for a given object. The various stages of post-processing were implemented in custom software, written in IDL, and are described next.

### 3.2. Pipeline Datasets and Initial Post-Processing

The `x1d` datasets retrieved from HSTonline had been processed by the most recent incarnation of the “calstis” pipeline, with final versions of the reference (calibration) files adopted in the years following the termination of STIS operations in 2004. Compared with the earlier “CoolCAT<sup>6</sup>” effort (which focussed on late-type (“cool”) stars), the new “final” pipeline processing had incorporated a number of improved reference files, including the echelle ripple correction (relative sensitivity for each order) and the key aperture transmission factors, particularly for the narrower slits where light loss can be substantial. The refined calibrations allow one to connect spectral traces from neighboring wavelength regions (either interorder, within a single echellegram; or between independent settings) more cleanly on a common radiometric scale, especially if the intervals are not directly overlapping (this issue is elaborated later).

The initial post-processing, at the exposure level, consisted of merging the individual echelle orders of the `x1d` file to produce a truly 1D wavelength tracing of flux density ( $f_\lambda$ :  $\text{ergs cm}^{-2} \text{s}^{-1} \text{\AA}^{-1}$ ), photometric error ( $\sigma_\lambda$ : same units), and data quality ( $q_\lambda$ : ranging from 0 (no issues) to various higher values flagging detector blemishes, bad pixels, saturation, and the like). A number of initialization steps preceded the order merging.

#### 3.2.1. Wavelength Distortion Correction

The pipeline processing is deficient in one important, although subtle, respect: the low order polynomials describing the mapping from pixel coordinates in the raw MAMA frames to wavelengths in the extracted spectra (the “dispersion relations”) do not fully capture all the persistent distortions present in the various echelle modes (Ayres 2008). This can be an issue when merging overlapping segments of adjacent orders, because inconsistencies in the assigned wavelengths can lead to imperfect alignment between spectral features in the overlap zones, blurring the resolution.

---

<sup>6</sup>See: <http://arlsrv.colorado.edu/~ayres/CoolCAT/>

The extensive “Deep Lamp Project” (Ayres 2008) — which analyzed long exposures of the STIS wavelength calibration sources — provided a solution to this otherwise vexing problem: a post facto correction that could be applied to the `x1d` pipeline output. As part of the StarCAT effort, empirical distortion models were developed for all 44 supported STIS echelle mode settings, along the lines described in the original Deep Lamp study. Figure 2 schematically depicts the resulting maps grouped in the four uber modes. For some of the little used secondary tilts, the pipeline dispersion relations introduced a more-or-less global displacement, reaching as much as  $2 \text{ km s}^{-1}$  in equivalent velocity units in one example (E230M-2269).

The first step of the `x1d` “initialization” consisted of applying the distortion correction. Again, the main benefit was to achieve accurate alignment between overlapping segments of adjacent echelle orders without the necessity of applying, say, cross-correlation of spectral structure in common (as had to be done, for example, with the independent “FP-split” subexposures disgorged by GHRS). Like the `calstis` dispersion relations, the 2D corrections were expressed in polynomial expansions of two variables, including cross terms; but one that utilized higher order terms where appropriate, and fully orthogonal independent variables (see discussion by Ayres 2008). In practice, the distortions were specified as equivalent velocities,  $v$  (where  $\Delta\lambda = (\lambda / c) \times v$ ), as a function of the grating parameter  $k$  ( $\equiv m \times \lambda$ ) and order number  $m$ . The  $k$  parameter varies exclusively along the order, whereas  $m$  varies exclusively (and trivially) with the orders, thus providing the desirable orthogonality of the independent variables in the bi-polynomial expansions (i.e., containing terms like  $m$ ,  $k$ ,  $m k$ ,  $m k^2$ ,  $m^2 k$ , and so forth). The specific models are summarized in Table 3.

The velocity corrections were applied to the individual orders according to:

$$\lambda_{\text{new}} = \lambda_{\text{old}} - (v[m, k] / c) \times \lambda_{\text{old}}. \quad (1)$$

It would be preferable, of course, to have a more accurate dispersion model implemented in `calstis` itself, to avoid the distortion correction step altogether. But, in the interests of producing StarCAT expeditiously in the face of the development effort required to modify the pipeline, it was decided to proceed with the post facto approach; and later update the catalog when a wavelength-enhanced `calstis` eventually does become available.

### 3.2.2. Resampling to Increase Pixel Density

Following the distortion correction, the parameters of each `x1d` order were linearly interpolated onto a finer wavelength grid with twice the pixel density delivered by the pipeline, to achieve four points per resolution element. The (trivial) double-density resampling was

helpful in the subsequent merging step (and in the various coaddition stages downstream), for which nontrivial interpolations were required to match the wavelengths of one order to those of the next one up: interpolating sub-Nyquist sampled spectra can lead to undesirable smoothing.

### 3.2.3. Recalculation of Photometric Error

At this stage, a slightly modified photometric error was introduced. The pipeline calculates a  $1\sigma$  value according to the square root of the gross number of counts,  $N$ , in the wavelength bin, divided by exposure time and sensitivity. However, a careful analysis of Poisson statistics in the low count regime shows that the  $1\sigma$  error is better represented by  $\sqrt{N + 1}$  rather than  $\sqrt{N}$ .<sup>7</sup> The modified photometric error was calculated according to the  $\sqrt{N + 1}$  model.

Ordinarily, one would utilize for the purpose the GROSS array reported by the pipeline. It was found, however, that occasionally the pipeline GROSS was offset by one entire order, rendering the  $N$  values unsuitable for the local photometric uncertainty calculation. Instead, the appropriate  $N$  was derived directly from the pipeline photometric error (which apparently was calculated properly even when there was the  $\Delta m = 1$  shift of the GROSS counts), by multiplying by the sensitivity ( $s_\lambda$ : counts  $\text{s}^{-1}$  per unit flux density) and exposure time, then squaring the result. The  $1\sigma$  limit in flux density was recovered by adding 1 to the derived  $N$  (which would be the sum of two derived values in regions of order overlap), taking the square root, and dividing by exposure time and sensitivity (sum of individual  $s_\lambda$ 's for overlap zones). In practice, the difference between the modified and pipeline  $\sigma$ 's was minor, except in regions of very low count rate. A benefit of the modified  $\sigma$ , however, is that the photometric uncertainty always is greater than zero, and thus a S/N calculation would never diverge. In the double density interpolation step, the photometric error was boosted by  $\sqrt{2}$  to preserve the average (in quadrature) over the original flux cell.

During verification of the modified  $\sigma_\lambda$ 's it was noticed that occasionally a pair of sharp spikes would appear in the photometric error at the shortwavelength end of E140M spectra

---

<sup>7</sup>This is an average of asymmetric upper and lower bounds in the limiting case of zero background. The asymmetry arises from the positive definite nature of counting in Poisson statistics. The uncertainty, itself, reflects lack of knowledge of the intrinsic source — given a single measurement of a specific number of counts — in the presence of Poisson fluctuations. Even if there are no counts recorded, there still would be a minimum photometric uncertainty, because the empty bin could have been contributed by a downward statistical fluctuation of a nonzero source (see, e.g., discussion by Ayres 2004).

(near 1170 Å, as illustrated later in Figure 8a). Often, the associated flux densities would exhibit a strong negative dip, but sometimes a distinct peak. This is the signature of a bright spot on the detector that most of the time falls in the interorder background, but occasionally inside an order (owing to the slight randomness in the MSM  $y$  positioning). Examination of a random selection of E140M raw images, including many wavecal spectra from the previously mentioned Deep Lamp Project, revealed that the bright spot was highly intermittent: sometimes at a level of a few hundred counts, but more often entirely absent. (Coincidentally, the spot was quite conspicuous in the summed AD Leo E140M exposures of Fig. 1a: it falls in the extreme lower right of the frame, next to the C III 1175 Å blend, and clearly is a detector artifact because a true spectral feature would be replicated in the next order up.)

A check of E140H raw images, taken in three different wavelength settings, also showed the bright spot at the same location on the FUV MAMA, although again its appearance was highly intermittent. In general, the feature was not flagged by the pipeline, although in a few instances it was assigned a data quality flag of 16, indicating an anomalously high dark current. At its worst, the blemish appears as a  $10 \times 3$  patch (in high-res pixels), with up to a few hundred counts. (In the case of AD Leo, the peak was about 40 counts in each 2200–2700 s exposure). Since the feature is not flagged, it can survive the coaddition process. However, it was decided not to attempt a correction at this time, especially since: (1) the issue resides in the pipeline; (2) it affects only a very small number of pixels; and (3) is very intermittent in the first place. One should be wary, nevertheless, of the 1170 Å region of E140M spectra: single or paired spikes in the local  $\sigma_\lambda$  probably signal the presence of the defect. The shortwavelength ends of E140H spectra could be affected similarly. This is another example of a potential area where the pipeline could be improved.

#### 3.2.4. *Data Quality Flag*

The data quality flag plays a vital role in marking valid fluxes for interorder merging as well as in subsequent coaddition stages involving independent exposures. Although the flag should not, strictly speaking, be interpreted as a number (it is a sequence of 16 bits that are either off (0) or on (1) to indicate specific conditions, or combinations of them); roughly speaking, values larger than 500 indicate more serious problems<sup>8</sup>. Of the  $1.3 \times 10^8$  pixels in

---

<sup>8</sup>Specific flagged conditions and the corresponding data quality values are listed in Table 2.8 of the STIS Data Handbook;

See: [http://www.stsci.edu/hst/stis/documents/handbooks/currentDHB/STIS\\_longdhbTOC.html](http://www.stsci.edu/hst/stis/documents/handbooks/currentDHB/STIS_longdhbTOC.html)

the full collection of original `x1d` files, 95% had quality values of 0, 2% were flagged in the range 1–256, and 3% in the range 512–3620 (the latter was the highest value encountered). The most common flag in the middle range was 16, indicating an anomalously high dark rate for the pixel. Visual examination of representative `x1d` files revealed that the spectral points assigned flags in the lower range did not appear to be unusual with respect to their neighbors (except for the one possible conspicuous exception described in the previous section), whereas those in the upper range (i.e.,  $> 500$ ) were more clearly deviant. Thus, the value  $q_\lambda = 500$  was chosen as the cutoff for “bad” fluxes. (The specific value  $q_\lambda = 500$  also was used to mark wavelength gaps that sometimes occur between the `x1d` orders.)

### 3.2.5. *Edge Clipping, Anomalous Gaps, and Flux Roll-Offs*

To avoid any edge effects (physically in the detector or resulting from the pipeline modeling of, say, the scattered light correction), a set number of points were clipped at the beginning and end of each order prior to merging. For the FUV modes, both trim values were set to 40 points uniformly for all  $m$  (the full span of each resampled `x1d` order now was 2048 bins). For the NUV modes, it was found that occasionally the lowest few orders (longest wavelengths) would be affected by unflagged “dropouts” at the shortwavelength edge of the stripe (similar gaps in the higher orders invariably would be flagged correctly). These unmarked dropouts would be treated as valid fluxes, and averaged with the good points at the end of the preceding order, leading to false depressions. While it is easy to spot such defects visually, it was not so easy to design a purely automated procedure to recognize them in an arbitrary echellegram (because they occurred without an obvious pattern). For this reason, all of the NUV modes were treated as if they were affected by the dropouts, and additional points were clipped from the shortwavelength edges of the lowest orders (300 points, total, for the lowest  $m$ ; 240 for the second and third lowest; 160 for the fourth lowest; and the nominal 40 for all the rest). The trim value for the high side of the orders was set to 40 uniformly, same as for FUV. (Note, again, that these are equivalent to “high-res” pixels.)

The expanded dropout exclusion zones were large enough to avoid the worst cases identified in a random examination of NUV exposures, but not so large as to significantly diminish the quality of spectra not affected by unflagged gaps (by eliminating otherwise good fluxes). The strategy resulted in a very substantial decrease in the number of defects in the final spectra, although very occasionally one might still find a small sharp unflagged dropout in one spectrum of an overlapping pair, made obvious, for example, during the splicing step described below, where ordinarily there would be excellent point-by-point agreement between the independent flux traces. Better dropout flagging, particularly in NUV, is another

example of where the pipeline could be enhanced.

The conservative order clipping was dictated by the desire to automate the merging procedure as much as possible, and avoid the more custom hands-on approach that in principle could be applied to a smaller collection of echellegrams. In any event, less than 5% of the data points were removed in the process. In fact, many of these (at the extremes of the orders) already had been flagged by the pipeline (a fair fraction of the 3% bad points noted earlier) and most represent redundant wavelengths, so information loss was minimal. However, a consequence of the trimming is slightly less overlap than would be expected in the low orders, so there are a few more, and somewhat larger, interorder gaps at the longward ends of some of the modes. This is particularly visible in E140M, where small regularly spaced gaps begin to appear at around 1550 Å, and gain in strength toward longer wavelengths.

In the process of tracking down sources of the anomalous gaps in some E230 frames, it was noticed that the E230H-2313 setting, and to a lesser extent neighboring tilt E230H-2363, was affected by a systematic roll-off of the fluxes on the shortwavelength side of each order in the middle range of  $m$ 's for that mode. The roll-off is not flagged as unusual. As a consequence, when the overlap zone was merged, the roll-off part would be averaged with the normal fluxes from the end of the previous order, producing a sharp-edged dip with a gradual curve back to the normal flux level on the longwavelength side. Since this was a *calstis* calibration issue (probably a bad value in one of the echelle ripple parameters) and there was no easy fix, all the E230H-2313 exposures (thankfully, there were only two) were simply excluded from the subsequent processing stages. The E230H-2363 tilt appeared to be mildly affected by this problem, but not so noticeably to justify exclusion of these (more numerous) datasets. Both E230H-2313 and E230H-2363 are little-used secondary settings. The closest prime tilt, E230H-2263, appeared to be perfectly normal.

### 3.2.6. *Merging of Adjacent Orders*

Following the initialization steps described above, the individual echelle orders of the `x1d` file were merged. This was accomplished in ascending wavelength, by linearly interpolating  $f_\lambda$  and  $\sigma_\lambda$  from the end of order  $m$  onto the wavelength scale of the overlapping zone at the beginning of the next order,  $m - 1$ ; then coadding, weighting by the relative sensitivity functions,  $s_\lambda$ . The latter was obtained for each order by dividing the trace of net count rate (counts  $\text{s}^{-1}$ ) by flux density. The sensitivity peaks at the order center, and tapers off to either side. Weighting by (normalized) sensitivity is equivalent to simply adding the net count rates, then dividing by the combined sensitivity, and exposure time. (The sensitivity profile also was utilized in the recalculation of the photometric error, as described earlier.)

The data quality flag was “interpolated” according to a nearest-neighbor function.

In combining the flux-like quantities from the overlapping zones, if both bins were flagged “good” ( $q_\lambda < 500$ , as noted above), then the data quality flag of the merged point was set to the higher of the two  $q_\lambda$ ’s. Since the vast majority of the points have a zero value for the flag, this strategy should effectively capture any potential issues with merged points devolving from a specific flag of one of the pair ostensibly in the good range, but nevertheless  $> 0$ . If one or both points was identified as bad ( $q_\lambda > 500$ ), the point with the lowest  $q_\lambda$  was retained. If wavelength gaps were encountered, additional closely spaced bins with  $q_\lambda = 500$  and  $f_\lambda = \sigma_\lambda = 0$  were interjected at both edges to ensure that the blank interval would be properly treated in later interpolations. (The same hierarchy was followed in subsequent coaddition and/or splicing steps.)

### 3.2.7. Output Files

The rootname of the original pipeline dataset (e.g., “o61s01010”) was retained for the merged file, appended with “\_n” to identify subexposures (multiple repeats within the same observation ‘wrapper’). For example, the merged saveset in the preceding case, a single exposure, would be called “o61s01010\_1”. The entire collection of `x1d` datasets was processed automatically, and the resulting files then were copied to the appropriate target directory (identical to the consensus object name) for subsequent processing: in this example, as “AD-LEO/o61s01010\_1”. The full list of STIS exposures can be found in Table 4, explicitly including the subexposures.

## 3.3. Stage Zero: Coaddition of Subexposures

The previously described merging step was in essence a simple reorganization of the native pipeline `x1d` file. As noted above, however, sometimes an observation consisted of multiple repeats. These were processed separately by the pipeline, and appear as additional data extensions of the `x1d` file. The maximum number was 19 in one exceptional case (VW-HYI), but more typically there were five or fewer, and often just one. The zeroth level coaddition stage was to combine any multiple subexposures; as a prelude to the subsequent stages, which consisted of generally less trivial coadditions of independent exposures.

### 3.3.1. *Velocity Registration*

In the subexposure coadditions, a new issue comes into play. Namely, slight pointing drifts or jitter within a visit can affect the apparent radial velocity of each spectrum of a sequence, no matter how well the occasional routine wavecalcs are able to track the instrumental zero point (which also can change over time, owing to thermal effects). For example, the  $0.2'' \times 0.2''$  aperture (“020×020,” as described earlier), commonly used to achieve the best photometric accuracy, is about seven (low-res) spectral bins wide for E140M-1425, or around  $22 \text{ km s}^{-1}$  in equivalent velocity units. Thus, a slight drift off-center by only a few tenths of the aperture width could produce a spurious velocity shift of several  $\text{km s}^{-1}$  (which incidentally is larger than typical accuracies of bright star radial velocities). Blindly coadding a sequence of exposures containing random shifts of this order could lead to undesirable spectral blurring.

To counter that possibility, one aligns the individual spectra by cross-correlation, if there is enough contrast and S/N in some characteristic spectral feature, or features, within the grasp of the particular mode. Narrow emission lines are good markers in cool star spectra, and sharp interstellar absorptions are similarly valuable in (the more numerous) hot stars. For late-type chromospheric emission-line objects, it is preferable to choose atomic lines like O I  $\lambda 1306$ , rather than highly ionized species such as C IV  $\lambda 1548$ , even though the latter might be brighter, because the higher temperature features are known to display more short term stochastic velocity variability (say, due to microflaring) than their lower temperature cousins. Such considerations do not apply to the same degree, of course, to the interstellar absorptions in hot-star spectra.

In practice, the first subexposure of a sequence was adopted as a template (if a suitable contrasty feature was available), and all the others were cross-correlated against it to establish a set of relative velocity shifts. The initial exposure of a sequence likely will have the best target centering, because it will be closest in time to an acquisition (and peak-up, if any). The remaining observations therefore should be slaved to the first to preserve “velocity memory.”

Because of the excellent internal wavelength precision of the distortion-corrected echellogram, the cross-correlation can be performed on a single spectral feature, as long as S/N is adequate. The rule of thumb for line centroiding, which functionally is analogous to cross-correlation, is that the precision is roughly the line width divided by the peak S/N (e.g., Lenz & Ayres 1992). The latter often is 20, or better, so precisions of a small fraction of the resolution element can be achieved under good circumstances. This is sufficient to prevent spectral blurring. In cases for which suitable narrow features were lacking, or S/N was too low, the subtle effects of spectral blurring would not be recognizable in the first place, so blind coaddition would be acceptable.

The cross-correlation shifts were reported, and applied, as equivalent velocities, because that is how the echelle spectrum responds to a simple spatial displacement of the target away from the center of the observing aperture (i.e., a set of narrow features at different wavelengths in a echellegram would show the same *velocity* offset, rather than, say, a constant *wavelength* shift). Conveniently, this allows the derived shifts to be compared without regard to the specific template wavelength. (The velocity shifts were applied to the individual datasets according to eq. 1.)

### 3.3.2. Coaddition Scheme

The coaddition algorithm was implemented as a semi-automated procedure, with an “operator” in the loop. The operator first was presented with a view of a particular 50 Å interval of the initial spectrum of a set. The specific region was chosen autonomously by the “robot” on the basis of a numerical test that identified areas of combined high contrast and high S/N. The operator then would select a specific feature in the region to serve as the cross-correlation template. Next, the other spectra of the sequence were correlated against the template in a narrow wavelength band (also chosen by the operator), and the resulting usually bell-shaped cross-correlation trace was modeled by a Gaussian, fitting only those values above the  $1/e$  levels (relative to the peak of the profile). Finally, the resulting fits were displayed. If the global cross-correlation was judged a success, the operator would activate the subsequent coaddition step. If the fits appeared to be poor, the operator could choose a new template feature from the same region; move on to a new interval; or if nothing better seemed to be available, default to a blind coaddition (all relative shifts set to zero).

The coaddition proceeded as follows. First, the other subexposures were interpolated onto the wavelength grid of the initial one, after applying the velocity correction. In principle, this implies a somewhat nonuniform treatment, because the first exposure is not subjected to an interpolation, and consequent potential smoothing. The initial double-density resampling, however, should mitigate this bias. Since the wavelength bins of the subexposures were the same, there was no need to further modify the photometric error.

Next, the velocity corrected, resampled spectra were averaged, weighting the flux densities by exposure time, and combining the photometric errors in quadrature,  $\sigma_{\text{tot}} = \sqrt{\sum \omega_i^2 \sigma_i^2}$ , where  $\omega_i \equiv t_i / \sum t_j$  is the normalized exposure time weight. In principle, the weighting factor should be proportional to the total number of net counts in each spectrum. For the subexposures, however, net counts and exposure time are synonymous. (The camera backgrounds are low and do not depend strongly on circumstances of the observations.)

In the averaging step, only the good points (with  $q_\lambda < 500$ ) were included. The  $q_\lambda$  for the average flux then was set to the maximum of the set of the  $q_\lambda < 500$ . If all the flux points in the wavelength bin were flagged as bad, then the parameters of the point with the smallest  $q_\lambda > 500$  were retained. If all the points were flagged with  $q_\lambda = 500$  (indicating a hard gap), then the output  $f_\lambda$  and  $\sigma_\lambda$  were set to zero, and the data quality to 500.

Embedded in Table 4 are summaries of the subexposure coadditions: specific exposure sequences, template wavelengths, cross-correlation bandpass half widths, and derived velocity shifts (zero, by definition, for the first exposure of a group). If a template wavelength and band are not reported, then a blind coaddition was performed. (Incidentally, Table 4 has a similar structure to the ascii scripts that controlled this, and subsequent, stages of coaddition and splicing.)

### 3.3.3. Output Files

The subexposure coaddition created a saveset named according to the root name of the original `calstis x1d` dataset, e.g., “o4o001010.” These merged o-type datasets also were written to FITS<sup>9</sup> files for later public access. The FITS filename follows the MAST HLSP convention and would be, e.g., “h\_o4o001010\_spc.fits” for the example above. The header of FITS extension zero (EXTEN=0) contains information concerning the target, including coordinates, name assigned by the Guest Observer, consensus StarCAT name, and *HST* proposal identifier; the observation, including start time, exposure duration, STIS echelle mode, setting, and aperture; and the processing configuration, including cross-correlation template parameters and derived velocity shifts (when subexposures were involved). Also listed are the wavelength ranges of each echelle order retained in the merging process. From these ranges, one can infer the overlap zones and splice points. If the observation consisted of only a single subexposure, then the StarCAT processing flag would be set to “NULL” and the merged dataset would be stored in the first (and only) extension of the FITS file (EXTEN=1). The parameters are WAVE (wavelength: Å), FLUX (flux density: ergs cm<sup>-2</sup> s<sup>-1</sup> Å<sup>-1</sup>), ERROR (photometric error: same units as flux density), and DQ (data quality flag: unitless). If the observation had two or more subexposures, then the processing flag would be set to “ZEROTH-STAGE” and the *coadded* spectrum would be written to EXTEN=1. Then, the individual subexposures 1 through  $n$  would be written sequentially to EXTEN=2 through  $n + 1$ . Thus, EXTEN=1 of the o-type FITS files always contains the most refined version of the observation, although all the subexposures, if any, are available in the trailing

---

<sup>9</sup>Flexible Image Transport System; see: <http://fits.gsfc.nasa.gov/>

extensions.

### 3.4. Stage One: Coaddition of Like Modes from Same Exposure Sequence

In many cases, one finds a series of independent exposures of a target in the same echelle mode, scan position, and aperture. Sometimes these are similar observations in a single visit, but specified as separate exposure lines instead of as “repeats” (subexposures), perhaps with different durations. Occasionally, there is a sequence of similar echellegrams taken at widely separated times either intentionally, say to assess source variability, or unintentionally simply as a consequence of scheduling. As in the subexposure case, such groups can be combined to boost S/N. Given the possibility of intrinsic source variability, however, the first stage coaddition focussed on like-mode exposures taken in a clearly delineated sequence within a single visit, in all cases not exceeding one day between the initial and final observations. Such groupings, closely analogous to subexposures, were easily recognized by their sequential dataset names.

There were a few exceptions where objects displayed noticeable variability even within the short time frame of a single visit. The Stage 1 coaddition still was carried out, but the variability was noted in the object header in the relevant table (Table 5). In a few instances, there was an issue with one exposure of a sequence. For example, in the first epoch visit for HD94028, the fifth exposure was very short, essentially blank. In such cases, the offending exposure simply was excluded from the visit-level coadd, and will appear in the table isolated from the other datasets of its kind from that visit.

The coaddition procedure was functionally identical to that for Stage 0. Again, the first exposure of a sequence was adopted as the velocity template, and all subsequent exposures of the group were cross-correlated against it to establish the set of relative shifts. The reasoning was the same as in the subexposure case: the initial exposure of a sequence likely will have the best target centering, because it normally will be closest in time to an acquisition (and peak-up, if any). Table 5 summarizes the like-mode, same-visit coadditions in a layout similar to Table 4.

#### 3.4.1. Stage Zero and One Velocity Shifts

Figure 3 illustrates histograms of the derived velocity shifts, separated by uber mode, for Stages 0 and 1. Combining the samples is justified because the distinction between the two types of coadditions is minimal: they both consist of sequences of like-mode exposures

taken close together in time. The Stage 0 processing was done separately mainly to ensure that each individual STIS o-type exposure, say as tabulated in a MAST query result, could be associated with a single refined data product.

The shifts included in the histograms were taken exclusively from the groups having valid cross-correlation solutions (and ignoring the leading exposures for which  $v \equiv 0$  by design). The parameters of the distributions are summarized in Table 6. The  $1\sigma$  standard deviations are slightly more than  $1\text{ km s}^{-1}$  for the M modes, and about half that for the H modes; both at a level where blurring of the cumulative line spread function might be considered minimal (but still worth correcting). The narrowness of the profiles — which mainly must reflect spacecraft drifts and/or jitter — provides additional confidence in the pointing stability of *HST*. This also suggests that velocity blurring in blindly coadded sequences will not be troublesome in general.

Although the vast majority of derived cross-correlation shifts at the visit level (Stage 1) were less than a few  $\text{km s}^{-1}$ , usually much less, there were a few notable exceptions. For HD36285, the target apparently was “dithered” on the  $310\times 005\text{ND}$  slit, which produced very large, although discrete, velocity shifts among the exposures of that visit ( $v = 0$  and  $\pm 808\text{ km s}^{-1}$ ). For HZ-HER, the derived velocity shifts were systematically changing during each visit (owing to rapid orbital motion), with up to tens of  $\text{km s}^{-1}$  increments between exposures, and a maximum shift  $> 100\text{ km s}^{-1}$ . In both cases, the anomalous velocities had an external cause, and do not reflect the usual high precision of the *HST* guiding.

### 3.4.2. Output Files

The like-mode, same-visit coaddition created a saveset named, e.g., “E140H-1271\_020X009\_51773.” The first part is the echelle mode designator and scan setting; the middle part is an aperture code (height times width, in units of 10 marcs, as before); and the final part is the last five digits of the Julian start time, in integer days, of the initial exposure of the sequence (i.e.,  $\text{JD} - 2,400,000$ : “modified Julian Date” or MJD).

There were a few cases in which separate like-mode coadds were from clearly delineated exposure groups, but close enough in time that the integer MJD flag was the same. In these instances, “a” and “b” would be appended to the aperture designator, for example, if there were two such identical file names. There are other examples where “F” or “N” appear at the end of the aperture flag. The former refers to observations taken through the FP-split apertures mentioned earlier, and the latter to one of the ND filtered slits. A few instances of “F”- and “N”-suffix exposures also contained separate sequences close enough in time to

yield identical file names. In such cases, the first coadd would replace suffix “N” with “a”; the second sequence, with “b,” and so forth.

The Stage 1 coadds also were written to FITS, with the same general structure as the o-type files. Again, the EXTEN=0 header contains target information and digests of the Stage 1 processing, including a list of the constituent exposures, the cross-correlation template  $\lambda$  and  $\Delta\lambda$ , and derived velocity shifts. In the exposure list, an o-type spectrum with suffix “\_1” designates a single exposure, while an unadorned o-type indicates a Stage 0 coadd. (This convention is maintained in the Stage 2 and 3 FITS headers described later.) The first (and only) data extension (EXTEN=1) contains the spectral parameters (WAVE, FLUX, ERROR, and DQ). The HLSP FITS file would be called something like, “h\_ad-leo\_e140m-1425\_020x020\_51613\_spc.fits,” where now the StarCAT target name appears at the beginning to minimize potential naming degeneracies.

### 3.5. Stage Two: Coaddition of Like Modes from Different Epochs and/or Apertures

In a number of cases, exposures of a target were taken in the same mode but in different epochs, or using different apertures, say the photometric slot on one occasion, but the Jenkins slit on another. For many of the multiple epoch cases, the motivation was to assess spectral changes of the source (say, an eclipsing binary or a flare star). If the source showed, in fact, a high degree of variability over the time line, there would be little motivation for further coadditions. However, when temporal changes were relatively minor, combining multiple epochs could enhance S/N.

This happens in two ways. First, and trivially, there is the simple accumulation of counts, leading to the normal Poisson  $\sqrt{N}$  improvement. Secondly, and more subtly, observations in different epochs, despite being in the same tilt, will have a slightly different placement of the spectrum on the MAMA camera (from a slight randomness —  $\sim \pm 3$  pixels — in the MSM  $y$  scan positions), leading to a different underlying pattern of pixel sensitivities. The presence of fixed pattern, in fact, normally limits the maximum attainable S/N in a single exposure (said to around 40), but a multiple epoch coaddition can transcend that limitation. (Fixed pattern is suppressed to some extent by the automatic compensation for the changing projected Doppler shift of the spacecraft over the 96 minute *HST* orbit, but the degree depends on the target attitude, time of year, and exposure duration.) For the mixed aperture observations, a coadded spectrum also can boost S/N, although at the expense of having a somewhat more complicated hybrid instrumental profile.

### 3.5.1. Coaddition Scheme

For the mixed epoch and/or aperture sequences, the cross-correlation and coaddition procedures were the same as for the single epoch case, with several important modifications.

First, the derived velocity shifts — as benchmarked against the template spectrum — subsequently were normalized to an average over all the shifts. Each individual observation (which might be a Stage 0 or 1 coadd) can be viewed as an independent realization of the target centering process (which includes acquisition, peak-up, and guiding), which certainly must suffer some degree of random error. Thus, absent a — probably elaborate — scheme to identify which observation of a group is optimum to use as a velocity reference in terms of best target centering, all the contributing spectra should be treated equally. In principle, subtracting the average shift will improve the velocity zero point of the super-coadd by something like the square root of the number of independent contributors.

A second refinement was that because all the exposures were treated on an equal basis, any one of them could serve as the cross-correlation template. The datasets were ordered in descending quality factor (QF: percentage of flux points exceeding  $2.5\sigma$  with respect to the local photometric error (per resolution element), a measure of relative S/N), and the leading spectrum (highest QF value) was selected as the template. This differs from the single epoch case where the first spectrum in a *time sequence* always was designated the template, because in principle its velocity zero point should be the most accurate of the set. Unlike at Stage 0 and 1 — where the constituent exposures typically were similar in duration — now one or more of the spectra might be a multi-repeat coadd, which then would be the preferred template from a S/N standpoint.

The third difference for the mixed epoch/aperture coadds was that occasionally an echelle order or two would be dropped in one set or the other, because of the specific positioning of the MSM in each independent visit. The consequence was that the extreme echelle orders might not fully benefit from the coaddition, which then would be reflected in a higher than expected  $\sigma_\lambda$  for the affected orders.

A fourth difference is that the overall weighting in the coaddition was by  $\sigma_\lambda^{-2}$  instead of  $t_{\text{exp}}$ . This can be shown to be equivalent to weighting by  $(\text{S/N})^2$ , which is nearly the same as weighting by net counts<sup>10</sup>. This was done primarily to account for the mixed aperture cases.

---

<sup>10</sup>One has the freedom, because the weights are normalized, to multiply each  $\sigma_\lambda^{-2}$  by  $f_\lambda^2$ . Regardless of the  $\sigma_\lambda$ , the  $f_\lambda$  for the two (or more) points should be the same. Thus  $\sigma_\lambda^{-2}$  is proportional to  $(f_\lambda/\sigma_\lambda)^2 \equiv (\text{S/N})^2$ . Also, if the background is low, as it usually is for the STIS MAMAs, then  $f_\lambda \sim N$  and  $\sigma_\lambda \sim \sqrt{N}$ , so that  $(\text{S/N})^2 \sim N$ .

While the assigned flux density levels might be the same in two observations of similar exposure time, the underlying count rates might be quite different owing to transmission losses in, say, a narrow slit compared with the broader photometric aperture.

### 3.5.2. *Output Files*

The result of a mixed epoch/aperture coaddition would be called, for example, “E140M-1425\_020X020\_XXXXX-YYYYY” for the sum of a set of E140M-1425\_020X020 exposures taken on different dates encompassing the range of integer MJDs XXXXX–YYYYY. If two different apertures were involved, both would be appended and the name would be, e.g., “E140M-1425\_020X020\_020X060\_XXXXX,” if the observations were in the same epoch, and “. . . XXXXX-YYYYY” otherwise. Any additional suffixes from Stage 0 or 1 (e.g., “a” or “N”) were not carried over to the new filename. Table 7 outlines the coaddition scheme for these types of mixed exposure sequences. Again, the coadditions were restricted to objects displaying minimal variability over their observational time lines. Figure 4 illustrates histograms of relative velocity shifts derived from any cross-correlations performed within this set. Although the samples are smaller, the distributions are similar to those of Stage 0 + 1, although perhaps slightly narrower for the H modes.

The Stage 2 datasets also were written to FITS with the same structure as the Stage 1 files. The resulting MAST HLSP file name would be something like, “h\_ad-leo\_e140m-1425\_020x020\_51613-52427\_spc.fits.”

## 3.6. Stage Three: Splicing of Available Wavelength Segments

The final step in constructing the catalog was to paste together all the separate wavelength segments available for each object into a single coherent spectrum. In many cases, this was trivial because the original GO program had utilized only a single echelle mode. If only a single exposure had been taken, then the final spectrum for that object would consist of the original merged pipeline file, i.e., “xxxxxxxx,” although possibly coadded over multiple subexposures. If several distinct exposures had been taken in a single visit, then the final file would be of the type, e.g., “E140H-1271\_020X009\_51773.” If two or more epochs were available, then the final spectrum would be of the type, e.g., “E140M-1425\_020X020\_XXXXX-YYYYY.” There also would be the various permutations if multiple apertures were involved in the single setting.

### 3.6.1. Splicing Scheme

For the cases in which multiple wavelength regions had been observed, a scheme was devised to knit together the perhaps overlapping segments in wavelength and flux. The procedure was trivial for those sets of observations lacking wavelengths in common, for example an E140M-1425 (1150–1700 Å) plus an E230H-2762 (2623–2900 Å). In such instances, the two spectra (which might themselves be coadds of one type or another) were simply butted together in wavelength, with some points trimmed off of the end of the first spectrum and the beginning of the second, especially if flagged as “bad.” For the other cases involving overlapping intervals, a more elaborate splicing procedure was developed (which will be described shortly).

Table 8 summarizes the specific groupings that were spliced. In one extreme case, the eclipsing binary VV Cep (HD208816: M2 Iab + B6 II), there are 21 separate spliced spectra covering the 21 independent epochs. At the other extreme is WD0501+527 (the calibration target G191-B2B, a DA white dwarf). Here, the extensive collections of M and H modes were separately spliced, owing to the complete coverage at each resolution, and regardless of epoch because the source is expected to be radiometrically very stable. There are seven mixed epoch single-mode coadds and single exposures in the “M” splice (with a maximum monochromatic exposure<sup>11</sup> of 15.6 ks), and a remarkable 36 independent spectra in the “H” splice (with a maximum exposure of 22.6 ks). Also worth mentioning is the single mode multi-epoch coadd of the dM flare star AD Leo “E140M-1425\_020X020\_51613-52427,” which encompasses six independent epochs and a total of 26 separate E140M-1425 exposures, with a maximum monochromatic exposure of 67 ks. Although several “flares” were apparent in the individual pointings, most were relatively small, and it was judged that the super-coadd would have value as the deepest exposure of an active late-type star obtained by STIS. (See, also, Fig. 1a, which is based on a coaddition the first 20 exposures of the sequence, taken within a few days of one another.)

### 3.6.2. Internal Flux Scaffold

As for the procedure itself, it is worthwhile reviewing how the splicing of adjacent spectral bands was handled in the earlier CoolCAT project. There, the individual FUV and NUV segments of a given target were draped over a broadband UV energy distribution abstracted

---

<sup>11</sup>The monochromatic exposure is the total  $t_{\text{exp}}$  at a given wavelength, taking into account all the contributing observations at that  $\lambda$ .

from the extensive collection of *IUE* low resolution ( $R \sim 300$ ) and echelle ( $R \sim 10^4$ ) spectra, covering the period 1978–1995. However, this procedure was viewed as less desirable for StarCAT owing partly to the much larger sample of objects (more than ten times greater) in the face of the labor intensive nature of the *IUE* SED construction. For example, there might be literally hundreds of *IUE* exposures available for popular UV objects, especially the calibration targets (mainly WDs) in common between the two observatories, and collecting and assembling these measurements into a consensus SED would be tedious at best. Furthermore, the *HST* radiometric calibration is superior to that of the earlier observatory, and thus should be preferred. Finally, there is the issue of variability, since the *IUE* spectra were obtained years to decades prior to the STIS pointings, and the SED of the earlier period might not be appropriate to the later *HST* era.

Thus, it was decided to forego the external calibration step, and rely instead on overlapping spectral intervals to serve as an internal radiometric scaffold. Consequently there are a few examples where a target was observed with non-overlapping echelle settings, and the segments do not appear to line up particularly well. In some cases the disjoint appearance might be misleading due to, say, strong interstellar absorption between well separated FUV and NUV segments (e.g., the 2200 Å “bump”). In other cases, it might represent a legitimate throughput issue: a target imperfectly centered in a narrow aperture, for example owing to slight pointing drifts accumulated over several orbits after an initial peak-up, especially when guiding with a single reference star (instead of the normal two). In such cases, it might be possible to identify which of the segments is deviant by finding an anomalous radial velocity (symptomatic of an off-center target), or simply adopting the observation through the widest aperture as likely the most valid.

When one has overlapping spectral intervals, one can measure directly (if enough signal is present in both segments) the relative flux density ratios of the wavelengths in common. Such values (or a default of unity in the absence of any overlap or insufficient signal in either or both adjacent zones) can be propagated through the full sequence to determine a relative flux scaling factor with respect to the first spectrum of the group (which is assigned a default scale factor of 1). Specifically, the scale factor for spectrum 2 is the overlap 2/overlap 1 flux ratio; the scale factor for spectrum 3 is that of 2 multiplied by the overlap 3/overlap 2 ratio; and so forth. At the end, one has a series of scale factors relative to spectrum 1. One next selects the maximum value of the set, which indicates the spectrum that presumably achieved the best throughput. One then divides that value by the others to obtain flux correction factors (“F<sub>scl</sub>,” always  $\geq 1$ ) to pass to the splicing procedure. At worst, the strategy will accurately capture the *shape* of the target SED, but at best the absolute radiometric level as well.

### 3.6.3. Wavelength Registration

As in the previous Stages, wavelength registration was accomplished by cross-correlating contrasty features in common to the overlap regions of adjacent segments. Like the multi-epoch coadds, the velocity shifts for each overlap zone (or zero if there was a gap between neighboring segments) were referenced to the initial segment (shortest wavelength), then the average of the whole set was subtracted; and for the same reason of not wishing to bias the global velocity zero point, treating each measurable overlap region as an independent realization of the target centering process. Of course, many of the overlap zones were contributed by multi-epoch coadds for which the internal velocity zero in principle already had been refined, to the further benefit of the global accuracy.

### 3.6.4. Coaddition of Overlapping Segments

Following alignment in flux and wavelength, the overlapping segments were averaged. The quantities of the shorter wavelength interval (“1”) were linearly interpolated onto the wavelength scale of the longer one (“2”), except for data quality, which was treated in a nearest-neighbor fashion. Prior to interpolation, the photometric error of segment 1 was multiplied by  $\sqrt{\Delta\lambda_1/\Delta\lambda_2}$  to compensate for the, in general, different binning ( $\Delta\lambda$ : Å per bin, variable with  $\lambda$ ) of the two intervals.

As in the Stage 2 procedure, the weighting factor was  $\sigma_\lambda^{-2}$ . This was especially important here, because of the typically much different monochromatic sensitivities of the overlapping modes, and because it often was the case that at the extreme end of a mode, the sensitivity would be varying rapidly with wavelength through the overlap zone. The monochromatic photometric error conveniently captures such differences and variations. The hierarchy for assigning flux values when one or both points were flagged as bad was the same as in the initial order merging step.

Commonly, one of the contributing spectra dominated in the overlap zone, because each exposure originally was set to optimally record some feature typically at the center of its grasp, rather than at the peripheries, so the resulting exposure depths can be quite different. Here, the role of the weaker spectrum would be to set the relative flux scale, and the velocity offset if a suitable marker was present, rather than contributing directly to improving S/N. There were cases, however, where one spectrum dominated, but had flagged gaps in the overlap zone. These gaps then were filled by the weaker spectrum, and the local photometric error consequently assumes a periodic step-like appearance. This effect could be minimized by choosing the longwavelength splice point short enough to exclude as many of the gaps as

possible, although at the cost of eliminating a few tens of Å of higher S/N spectrum between the gaps. Such choices were made on a case by case basis.

### 3.6.5. *Treatment of M and H Overlaps*

The philosophy that guided the specific splicing sequences was to include at each wavelength the highest spectral resolution material available for the particular object, and to minimize, to the extent possible, coaddition of unlike resolutions (e.g., M and H). Within this strategy, M and H modes were separately spliced only if there was complete, or nearly complete, spectral coverage in each set. An alternative strategy might have considered routinely splicing the M and H modes of an object separately. However, the adopted mixed resolution philosophy was dictated by the calibration strategy (of comparing flux ratios in overlapping regions): incorporating all the material available for each target maximized the overlaps for that purpose. Maximizing overlap opportunities also was important for ensuring accurate wavelength continuity throughout the spliced spectrum (by cross-correlating features in the overlap zones). Otherwise, there might be small velocity discontinuities between the different segments contributing to the whole.

For the cases of overlap between M and H modes, the coaddition zone was chosen to be as small as possible, but with sufficient coverage to provide an accurate flux ratio, if possible; and — ideally — a high contrast cross-correlation feature, to ensure accurate velocity alignment as well. Because the coaddition step still was performed, however, narrow features unresolved at M resolution would acquire hybrid lineshapes. Broad features, especially the continuum, would be less affected, and would fully benefit from the coaddition (if the M and H S/N were similar). The coaddition of the M and H overlaps was done — instead of, say, simply introducing a sharp discontinuity — mainly to achieve the higher S/N and ensure a smoother spectral transition between the unlike modes. In any event, all the constituent exposures reside in the StarCAT database, so specific cases where the M/H coaddition was viewed as undesirable could be examined in whichever of the original modes was preferred.

In some instances, one finds a single H mode fully contained within the range of an M mode. The splicing sequence was as follows: (1) retain the M mode values below the shortward edge of the H mode, with a minimal overlap between M and H at the shortward side of the latter as described above; (2) include the full range of the H mode up to its longwavelength edge; and (3) splice back the remainder of the M region, again with a minimal overlap at the longward edge of H. In Table 8, one would find the M spectrum repeated twice, sandwiching the H spectrum, with different slicing parameters reflecting the specific choices of overlap zones at the boundaries of H.

### 3.6.6. *Derived Flux Scale Factors and Velocity Shifts*

Included in Table 8 are the flux scale correction factors for each spectrum; the parameters of the cross-correlation templates, if any; and the derived velocity shifts. Figure 5 illustrates a histogram of the latter (restricted to the cases of successful cross-correlation solutions). Figure 6 is a histogram of the relative flux corrections determined from overlap zones jointly containing sufficient S/N to yield an accurate flux ratio. Because the corrections are inverses of ratios referenced to the largest observed value in a set, they always are  $\geq 1$ . For the histogram, only spectra (specifically, the second of a spliced pair) with a valid cross-correlation solution were considered. If both members of a spliced pair had sufficient S/N to allow a cross-correlation solution, the pair would be assured of having sufficient S/N for a valid flux ratio measurement. Among this sample, only values  $\geq 1.001$  (the next significant figure in Table 8 above  $\equiv 1$ ) were incorporated, to avoid the artificially sharp peak contributed by the reference (normalizing) exposures. The cutoff excludes a few cases where the derived correction was very close to 1, and reported as such owing to roundoff of significant digits in the table, but the main purpose of the figure is illustrative. Most of the scale factors are close to unity, the desirable outcome if the aperture calibrations and overall broad-band radiometry are accurate.

The few cases of large corrections were examined more closely. Observations for which  $F_{scl} > 1.5$  were isolated from the entire list of correction factors, regardless of cross-correlation success. Of the 38 single exposures or coadds identified in this way, almost all were H modes of one type or another (more NUV than FUV). Fully one-third had been taken through the tiny Jenkins slit, for which centering errors can dominate over knowledge of the transmission factors. Another quarter of the exposures had utilized one of the ND filtered apertures, and about a sixth had used the (also nonstandard) 010×020 slot. Only 10% of the observations had been through the 020×020 photometric aperture. The preference of the high  $F_{scl}$  exposures for very small or nonstandard apertures is consistent with the idea that the deviations arise from centering errors on the one hand, and less well calibrated transmission factors on the other.

### 3.6.7. *Output Files*

The final dataset resulting from a successful splicing was named something like, “UVSUM\_1M\_XXXXX-YYYYY”. The “UVSUM” part indicates that multiple wavelength regions were joined; the middle numeral points to a particular grouping of datasets specified for splicing, which might represent a single epoch of a variable source, or multiple visits of a less variable object; the letter appended to the group index encodes whether the constituent

datasets all were medium resolution (“M”), all high resolution (“H”), or mixed (“X”); and the final part summarizes the range of modified Julian start times representing whatever grouping was involved. The MJD designator was composed of the extreme values of the set, considering that one, or more, of the constituent spectra might itself be a mixed epoch coadd. If only a single epoch was represented in the spliced group, then just the one integer MJD would appear.

The Stage 3 UVSUM datasets were written to FITS files of similar structure to the E-type. The EXTEN=0 header contains target information and splicing parameters, while EXTEN=1 stores the spectral data values. The HLSP file name would be something like, “h\_ag-dra\_uvsum\_1m.52755\_spc.fits.”

### 3.6.8. *Final Catalog*

Table 9 lists the set of fully distilled spectra for each target, constituting the final catalog. In the worst cases, this would be just a single exposure. In the best cases, there would be a fully spliced spectrum covering the entire UV range without gaps. In intermediate cases, one would have single mode coadds, possibly multi-epoch; or multiple epochs of partially or fully spliced spectra. Figure 7 illustrates a low resolution tracing of the final spectrum (or spectra) of each object (as indicated in Table 9), with a schematic time line depicting the constituent exposures as a function of mode and epoch. In some cases, like VV Cep, the spectra are highly varied, following the carefully orchestrated multi-Cycle GO program (to catch critical orbital events such as ingress, egress, total secondary eclipse, quadrature, and so forth). In other cases, like G191-B2B, the spectra are time-independent, but combine to produce composites that achieve outstanding S/N throughout the full wavelength grasp of the STIS echelles.

Figure 8 illustrates a few examples of the StarCAT spectra at higher resolution (objects of Fig. 1), showing the remarkable detail that often is present in the coadded and spliced STIS echellegrams. Of course, all of the intermediate data products are accessible through the StarCAT interface, so an investigator can examine the material for a given object at any level of detail desired.

## 3.7. **Uncertainties**

The quality of the StarCAT material in principle is high, thanks to the extraordinary performance level of the STIS spectrometer on the one hand, and that of the *calstis* pipeline

processing on the other. While a thorough assessment has not yet been performed, it is possible to estimate the levels of precision and accuracy of the StarCAT wavelength and flux scales through indirect arguments appealing to parameter measurements obtained at the various stages of post-processing.

### 3.7.1. Wavelength Scales

The StarCAT effort was especially careful with the wavelength scales, implementing an initial distortion correction to the `calstis x1d` files to improve the internal consistency of the relative wavelengths, and to remove small systematic errors in the assigned velocity zero points in some instances (especially for the less heavily used secondary tilts). Taking the Deep Lamp project as a guide, the internal precision of the wavelength scale of an arbitrary *corrected* STIS exposure should be better than the  $600 \text{ m s}^{-1}$  for M modes and  $300 \text{ m s}^{-1}$  for H measured in *uncorrected* wavecals (roughly a tenth of the resolution element in both cases). Such precision could be realized in spectra with  $S/N > 10$ . The cited precision of the relative wavelengths is similar to that quoted in Table 16.2 of the STIS Instrument Handbook, noting that the values there are  $2\sigma$  and per low resolution pixel (with 2 pixels per resolution element).

The accuracy of the velocity zero point of an arbitrary STIS echellogram is more difficult to assess. In principle, it depends mainly on the target centering accuracy; thus, in practice, the time interval from the last peak-up, and the success of the centroiding at that peak-up. One could carry out, say, a careful evaluation of the apparent UV radial velocities of StarCAT objects for which accurate groundbased  $v_r$ 's have been measured. Such a comparison was made for the earlier CoolCAT effort. There, the deviation of sharp chromospheric emission lines from the known stellar (photospheric) radial velocities averaged  $+0.1 \pm 1.4 \text{ km s}^{-1}$  (Ayres 2002; considering only narrow-line dwarfs and giants, excluding dMe dwarfs, from Tables 1 and 2 of that article; thirteen stars in all). Although the analogous StarCAT comparison has yet to be conducted, it would not surprise the author to find an absolute accuracy at the same level (the empirical standard deviation of  $1.4 \text{ km s}^{-1}$ ), or better. This estimate, which refers exclusively to E140M exposures, lies between the 0.5–1 (low-res) pixel absolute accuracy (at  $2\sigma$ ) quoted by the Instrument Handbook.

What can be gleaned straightforwardly from the StarCAT processing are the *relative* velocity shifts between observations in a sequence. These typically show a standard deviation somewhat over  $1 \text{ km s}^{-1}$  for M modes and about half that for H modes. This is a measure of pointing jitter and/or drifts (the latter can be a problem if only a single guide star was available), and would be the velocity blurring expected in a typical sequence of exposures that

were blindly coadded. Such blurring would be considered inconsequential except perhaps in the highest S/N studies of intrinsically narrow spectral features with the H modes. However, in a high S/N spectrum containing conspicuous narrow features, the cross-correlation procedure normally would succeed, and thus would mitigate any relative velocity shifts within that sequence, thereby minimizing any profile blurring. (The “centroiding error” of a high quality cross-correlation measurement — templates with peak S/N > 20 — would be at a level of about 5% of the resolution element, which is better than the internal velocity scale precisions cited earlier.)

In spliced spectra, the relative precision of the velocity scales will be comparable to those of single mode exposures, if there is good overlap between the spliced segments, and if the overlap regions all contain suitable cross-correlation features. Otherwise, particularly when intra-mode gaps are present, one segment of the spliced spectrum might deviate from the velocity scale of the next segment by the absolute accuracy factor, although the relative precision *within* each segment of course still would be high. One could consult the Stage 3 splicing sequence for a particular object to judge whether the spectral continuity was such that the internal velocity scale likely would achieve the single exposure precision over the full range of the splice, or whether some of the segments were isolated from others, and thus the velocity scales might be less well connected.

For the multi-epoch coadds and spliced spectra, where average velocities were subtracted from the group cross-correlation shifts, the accuracy of the absolute velocity in principle would be improved by  $\sqrt{n}$ , where  $n$  is the number of  $v$ 's involved. This would be exclusive, however, of any systematic errors in the assignment of the velocity zero point (including corrections for the spacecraft orbital motion, the telluric component, and so forth). The CoolCAT analysis suggests that any systematic zero-point errors must be small, because the average deviation (over the thirteen stars considered here) was close to zero.

### 3.7.2. Radiometry

As for the StarCAT flux scales, there are two main sources of uncertainty. One is the absolute accuracy of the STIS radiometric calibration, and its repeatability, which includes the knowledge of the aperture transmission factors for each of the modes. Again, for an isolated exposure, the applicability of the default transmission factors will depend sensitively on the quality of the target centering. The absolute radiometric calibration of the STIS echelles — as determined from UV standard stars — is quoted at the 3% level ( $1\sigma$ ) for well-centered targets in the 020×020 aperture, and the repeatability also is 3%, leading to a total error budget of about 4% (in quadrature: Bohlin 1998).

The second major source of uncertainty is the “internal” calibration procedure that joins overlapping spectral segments. As with the velocity scale, if the overlap zones in a group of spectra all have high S/N, so that accurate flux ratios can be derived for each individual pair, then the spliced spectrum should have high relative flux precision throughout the full range, and the absolute accuracy should approach that associated with whatever single exposure or coadd was selected by the splicing procedure as the reference spectrum (for normalizing the F<sub>scl</sub> factors). Because the flux ratios are determined by averaging over many wavelength bins, the precision of the local measurement can be better than 1%, even when the monochromatic S/N is only 10 or so. Thus, the relative flux precision of a well spliced spectrum should be much better than the absolute accuracy (and limited mainly by knowledge of the relative SEDs of the UV standard stars). However, like the velocity case, if the splice included gaps or low-S/N overlaps, the affected segments could not be empirically aligned, and they then might suffer significant relative flux scale errors, especially if one or both sets of underlying observations had a poorly centered target.

Some idea of the radiometric repeatability, in practice, can be gleaned from the F<sub>scl</sub> factors themselves. Considering only the high S/N sample of F<sub>scl</sub>'s (i.e., successful cross-correlation solution), one finds that the average deviation for all the values (including those equal to 1, since the exposure assigned unit F<sub>scl</sub> either was itself the reference spectrum of a group, or had unit flux ratio with the reference spectrum), is 1.06 (with standard deviation 0.14) for the 020×020 aperture (154 cases), and 1.07 (with standard deviation 0.18) when either the 020×009 or 020×006 aperture was used (78 cases). It is somewhat surprising that the narrower slits showed essentially the same average deviation as the photometric aperture. A simplistic interpretation of these values would say that the true repeatability of the STIS radiometry for the echelle modes probably is closer to 7% when supported apertures were used. This would imply an overall flux accuracy of about 8%, including the standard star contribution, for an arbitrary exposure taken with one of the supported apertures. A spliced spectrum, however, could approach the 4% accuracy cited earlier, if it consisted of many constituent spectra, because there would be a better statistical chance that the selected reference spectrum (the one with the highest relative flux ratio) actually achieved ideal throughput (as assumed in the `calstis` processing).

### 3.7.3. Summary

The best StarCAT material likely achieves better than 500 m s<sup>-1</sup> internal velocity precision, and absolute accuracy at the 1 km s<sup>-1</sup> level. The absolute flux accuracy is of order 4%, while the internal relative precision should be several times better (limited mainly

by knowledge of spectral energy distributions of UV standard stars).

#### 4. CONCLUSIONS (AND CAVEATS)

The author sincerely hopes that the community will be able to make productive use of StarCAT. The author is confident that the semi-automated processing algorithms developed for the project fashion optimal spectral traces of each object, at least within the limitations of a variety of underlying assumptions. Furthermore, because each object was treated without any preconceptions that might have been harbored by the original Guest Observer, and at the same time draws on the author’s experience of post-processing many diverse spectra, the StarCAT material should possess a uniformity and lack of bias possibly superior to an individual GO’s attempts to reduce their own more limited set of echellegrams. On the other hand, certain choices had to be made by the operator (namely, the author) who monitored the actions of the robots. Examples include the specific template feature for each level of cross-correlation, the wavelength extremes for determining spectral overlap, and in some cases whether to accept a flux scale factor measurement (which can be biased by high noise levels, say, in one spectrum of a pair) or a cross-correlation solution (which can be affected, for example, by low S/N, or a poor selection of the wavelength limits).

These choices fall into the rubric of “operator bias,” and potentially could be important under certain circumstances. The author has processed the catalog end-to-end several times, making certain global improvements for each pass as reflected in the description, and is convinced that operator bias is minimal. Nevertheless, it is possible that an investigator more intimately familiar with a particular object — especially the highly variable ones — could carry out a superior reduction to the standard fare from StarCAT. The author, of course, would welcome any such comparisons, particularly if they might reveal specific deficiencies of any aspect in the StarCAT post-processing.

At the same time, the StarCAT effort has identified a number of areas where the *calstis* “final” pipeline processing could be enhanced: (1) upgraded dispersion relations to compensate for persistent small-scale wavelength distortions; (2) improved flagging of “dropouts,” especially in the lowest few orders of the NUV modes; (3) proper alignment of the **GROSS** spectrum with the orders; (4) better treatment of intermittent detector bright spots such as the conspicuous one in the lower right corner of the FUV MAMA; and (5) some calibration issues with — albeit rarely-used — E230H secondary tilts.

Finally, on a personal note, the author would like to express appreciation to all of the *HST* investigators who have indirectly, and unwittingly, contributed to StarCAT. It was

fascinating to view the wide diversity of STIS echelle spectra available for the stars, and the often dramatic differences between objects. The fact that such differences exist, and are widespread, is testament to the enormous value of UV spectroscopy to stellar astronomy. It has been a privilege to work with this material.

(Note: the complete Tables 1–5, 7–9, and Fig. 7 are available electronically; only examples of each are provided here.)

This work was supported by grant HST-AR-10638.01-A from the Space Telescope Science Institute, and NASA grant NAG5-13058; and has made use of public databases hosted by SIMBAD and VizieR, both maintained by CDS, Strasbourg, France; and the Multimission Archive at Space Telescope. The author thanks S. Redfield for valuable discussions at various stages of the project; and M. Smith, R. Thompson, and K. Levay for helpful suggestions concerning the final data products.

## REFERENCES

- Ayres, T. R. 2002, The 2002 HST Calibration Workshop : Hubble after the Installation of the ACS and the NICMOS Cooling System, 170
- Ayres, T. R., et al. 1995, ApJS, 96, 223
- Ayres, T. R., Simon, T., Stauffer, J. R., Stern, R. A., Pye, J. P., & Brown, A. 1996, ApJ, 473, 279
- Ayres, T. R. 2004, ApJ, 608, 957
- Ayres, T. R. 2008, ApJS, 177, 626
- Boggess, A., et al. 1978, Nature, 275, 377
- Bohlin, R. 1998, STIS Instrument Science Report 98–18 (Baltimore: Space Telescope Science Institute)
- Brandt, J. C., et al. 1994, PASP, 106, 890
- Lenz, D. D., & Ayres, T. R. 1992, PASP, 104, 1104 McCook et al. (latest edition: 2008
- McCook, G. P., & Sion, E. M. 2006, VizieR Online Data Catalog, III/235B [referred to as McCook+(2008) in VizieR]
- Woodgate, B. E., et al. 1998, PASP, 110, 1183

Table 1. The StarCAT Sample: Star Names

StarCAT Name (1)	$\alpha_{\text{HST}}$ ( $^{\circ}$ ) (2)	$\delta_{\text{HST}}$ (3)	SIMBAD Name (4)	<i>HST/GO</i> Name (5)	Notes (6)
HD108	+1.514	+63.680	HD 108	HD108	
HD166	+1.653	+29.021	HD 166	HD166	
HD256	+1.826	-17.387	HD 256	HD256	
HD1383	+4.574	+61.727	HD 1383	HD1383	
QR-AND	+4.958	+21.948	V* QR And	RXJ0019.8+2156	
HD1909	+5.803	-31.036	HD 1909	HD1909	
HD1999	+6.015	-37.411	HD 1999	HD1999	
HD2454	+7.084	+10.190	HD 2454	HD2454	
HD3175	+8.593	-63.062	HD 3175	HD3175	
HD3369	+9.220	+33.719	HD 3369	HD3369	
HD3827	+10.301	+39.604	HD 3827	HD3827	
HD4128	+10.897	-17.987	HD 4128	HD4128	
HD4174	+11.155	+40.679	HD 4174	HD4174	
STAR0046-7324	+11.676	-73.415	UCAC2 999822	AV15	
WD0044-121	+11.764	-11.872	NGC 246	NGC246-CENTRAL-STAR	

Table 1. —Continued

(1)	(2)	(3)	(4)	(5)	(6)
STAR0047-7306	+11.801	-73.109	UCAC2 999935	AV18	
				AZV18	directories merged
HD4539	+11.872	+9.982	HD 4539	HD004539	
STAR0047-7307	+11.911	-73.130	UCAC2 1000061	AV22	
STAR0048-7325	+12.215	-73.433	2MASSJ00485148-7325585	AV47	
BG-PHE	+12.255	-56.097	V* BG Phe	JL212	
HD4614	+12.277	+57.815	HD 4614	HD4614	
STAR0050-7253	+12.572	-72.892	GSC 09141-07827	AV69	
STAR0050-7252	+12.635	-72.877	UCAC2 1078509	AV75	
STAR0050-7247	+12.682	-72.795	2MASSJ00504380-7247414	AV80	
STAR0050-7242	+12.717	-72.704	AzV 83	AV83	
STAR0051-7244	+12.840	-72.737	AzV 95	AV95	
STAR0051-7248	+12.911	-72.802	2MASSJ00513842-7248060	AV104	
WD0050-332	+13.322	-32.999	WD 0050-33	WD0050-332	
STAR0055-7317	+13.927	-73.292	2MASSJ00554243-7317302	AV170	
STAR0056-7227a	+14.042	-72.459	Cl* NGC 330 ARP 45	NGC330-B22	
STAR0056-7228	+14.087	-72.476	Cl* NGC 330 ELS 4	NGC330-B37	
etc.					

Note. — Objects are listed in increasing right ascension.

Table 2. The StarCAT Sample: Stellar Properties

StarCAT Name (1)	SIMBAD Desc. (2)	Sp. Typ. (3)	$\pi$ ( $''$ ) (4)	$V$ (mag) (5)	$B - V$ (6)
HD108	Em*	O6pe	+0.000	+7.38	+0.11
HD166	BY*	K0Ve	+0.073	+6.13	+0.75
HD256	*	A2IV/V	+0.006	+6.23	+0.10
HD1383	*	B1II...	+0.000	+7.63	+0.26
QR-AND	XB*	Ss	...	+12.73	-0.35
HD1909	a2*	B9IVmn	+0.005	+6.55	-0.07
HD1999	*	B6III	+0.000	+8.30	-0.12
HD2454	PM*	F5Vsr	+0.028	+6.04	+0.43
HD3175	*	B4V	+0.001	+9.33	-0.18
HD3369	SB*	B5V	+0.005	+4.34	-0.11
HD3827	*	B0.7Vn	+0.001	+7.95	-0.19
HD4128	V*	G9II-III	+0.034	+2.04	+1.02
HD4174	Sy*	M2III:e	+0.001	+7.23	+1.61
STAR0046-7324	*	OB	...	+13.17	-0.21
WD0044-121	PN	Op	+0.002	+11.78	-0.35
STAR0047-7306	*	OB	...	+12.48	+0.04
HD4539	*	A	+0.005	+10.29	-0.17
STAR0047-7307	*	OB	...	+12.25	-0.10
STAR0048-7325	*	O8III	...	+13.38	-0.26
BG-PHE	V*	B5	+0.000	+10.18	-0.07
HD4614	SB*	G0V	+0.168	+3.45	+0.58
STAR0050-7253	*	OB:	...	+13.35	-0.22
STAR0050-7252	Em*	OB	...	+12.79	-0.16

Table 2. —Continued

(1)	(2)	(3)	(4)	(5)	(6)
STAR0050-7247	*	O5:n...	...	+13.33	−0.14
STAR0050-7242	*	O7.5Ia+.	...	+13.58	−0.13
STAR0051-7244	*	O7.5III	...	+13.91	−0.30
STAR0051-7248	*	B6	...	+13.13	−0.22
WD0050-332	WD*	sdO	...	+13.36	−0.22
STAR0055-7317	*	O9.5III	...	+14.09	−0.23
STAR0056-7227a	*iC	B2IIe	...	+14.52	−0.29
STAR0056-7228	*iC	B2.5Ib	...	+13.33	−0.07
STAR0056-7227b	*iC	B3Ib	...	+13.03	−0.04
STAR0056-7227c	*iC	B0.5Ve	...	+14.82	−0.32
STAR0058-7210a	*iC	B0IV(Nst	...	+14.98	−0.14
STAR0058-7210b	*iC	OC6Vz	...	+15.01	−0.26
STAR0058-7216	*	OB	...	+12.67	−0.05
STAR0058-7232	*	OB	...	+12.77	−0.11
STAR0058-7210c	*iC	O4V((f+)	...	+14.07	−0.28
STAR0058-7244	*	B1-3II:	...	+14.20	−0.15
STAR0059-7210a	*iC	O3III...	...	+13.50	−0.23
STAR0059-7211a	*iC	O6.5V	...	+15.50	−0.26
STAR0059-7210b	*iC	O5.5:V:	...	+14.18	−0.23
STAR0059-7210c	*iC	O4III...	...	+12.57	−0.20
STAR0059-7210d	*iC	O6.5V	...	+14.53	−0.22
STAR0059-7205	Em*	O7Iab:..	...	+14.50	−0.22
STAR0059-7211b	*iC	O9V	...	+14.91	−0.26
HD5980	EB*	WNp...	...	+11.50	−0.22
STAR0059-7210e	*iC	O7Iaf+	...	+12.31	−0.19
STAR0103-7202	*	O9.7Iab:	...	+12.84	−0.18
HD6268	*	G0	+0.002	+8.10	+0.79
etc.					

Note. —

Table 3. Wavelength Distortion Correction

Polynomial Model	
E140M-1425:	$\tilde{m} \equiv \frac{(m-106)}{10}$ ; $\tilde{k} \equiv \frac{(m\lambda-147950)}{500}$
$v =$	$(0.2168845) + (0.5295246) \times \tilde{m} + (-0.1133152) \times \tilde{m}^2 + (-0.6151966) \times \tilde{m}^3 + (0.0123862) \times \tilde{m}^4$ $+ (0.0813750) \times \tilde{m}^5 + (0.7533000) \times \tilde{k} + (-0.1167010) \times \tilde{k}^2 + (-1.1997739) \times \tilde{k}^3 + (0.1042297) \times \tilde{k}^4$ $+ (0.2603576) \times \tilde{k}^5 + (-0.7276294) \times \tilde{m} \tilde{k} + (0.3840922) \times \tilde{m}^2 \tilde{k} + (0.1872590) \times \tilde{m}^3 \tilde{k} + (-0.1337153) \times \tilde{m}^4 \tilde{k}$ $+ (0.0018007) \times \tilde{m}^5 \tilde{k} + (-0.1091477) \times \tilde{m} \tilde{k}^2 + (-0.2054250) \times \tilde{m}^2 \tilde{k}^2 + (0.1786350) \times \tilde{m}^3 \tilde{k}^2 + (0.0175580) \times \tilde{m}^4 \tilde{k}^2$ $+ (-0.0321286) \times \tilde{m}^5 \tilde{k}^2 + (0.7644269) \times \tilde{m} \tilde{k}^3 + (-0.7003058) \times \tilde{m}^2 \tilde{k}^3 + (-0.3217466) \times \tilde{m}^3 \tilde{k}^3 + (0.2486133) \times \tilde{m}^4 \tilde{k}^3$ $+ (0.0498403) \times \tilde{m}^5 \tilde{k}^3 + (0.1559506) \times \tilde{m} \tilde{k}^4 + (0.1556453) \times \tilde{m}^2 \tilde{k}^4 + (-0.0563305) \times \tilde{m}^3 \tilde{k}^4 + (-0.0372428) \times \tilde{m}^4 \tilde{k}^4$ $+ (-0.0015327) \times \tilde{m}^5 \tilde{k}^4 + (-0.1468092) \times \tilde{m} \tilde{k}^5 + (0.4034192) \times \tilde{m}^2 \tilde{k}^5 + (0.0330954) \times \tilde{m}^3 \tilde{k}^5 + (-0.1312219) \times \tilde{m}^4 \tilde{k}^5$ $+ (-0.0049503) \times \tilde{m}^5 \tilde{k}^5$
E140H-1234:	$\tilde{m} \equiv \frac{(m-339)}{10}$ ; $\tilde{k} \equiv \frac{(m\lambda-421000)}{500}$
$v =$	$(0.0301367) + (0.0382590) \times \tilde{m} + (-0.0383733) \times \tilde{m}^2 + (-0.0418539) \times \tilde{m}^3 + (0.0979899) \times \tilde{k}$ $+ (0.0444978) \times \tilde{k}^2 + (-0.1219705) \times \tilde{k}^3 + (-0.0260690) \times \tilde{m} \tilde{k} + (0.0959600) \times \tilde{m}^2 \tilde{k} + (0.0569775) \times \tilde{m}^3 \tilde{k}$ $+ (0.0352926) \times \tilde{m} \tilde{k}^2 + (-0.0193594) \times \tilde{m}^2 \tilde{k}^2 + (-0.0123493) \times \tilde{m}^3 \tilde{k}^2 + (-0.0278292) \times \tilde{m} \tilde{k}^3 + (0.0192519) \times \tilde{m}^2 \tilde{k}^3$ $+ (0.0079272) \times \tilde{m}^3 \tilde{k}^3$
E230M-2269:	$\tilde{m} \equiv \frac{(m-94)}{10}$ ; $\tilde{k} \equiv \frac{(m\lambda-204100)}{500}$
$v =$	$(-1.1279973) + (0.7839361) \times \tilde{m} + (-0.3184260) \times \tilde{m}^2 + (-0.5095827) \times \tilde{m}^3 + (0.3729259) \times \tilde{k}$ $+ (-0.0702246) \times \tilde{k}^2 + (-0.0920298) \times \tilde{k}^3 + (-0.3107360) \times \tilde{m} \tilde{k} + (-0.0233978) \times \tilde{m}^2 \tilde{k} + (0.1220877) \times \tilde{m}^3 \tilde{k}$ $+ (-0.1275635) \times \tilde{m} \tilde{k}^2 + (0.0957628) \times \tilde{m}^2 \tilde{k}^2 + (0.0868626) \times \tilde{m}^3 \tilde{k}^2 + (0.0405377) \times \tilde{m} \tilde{k}^3 + (0.0244841) \times \tilde{m}^2 \tilde{k}^3$

Table 3—Continued

Polynomial Model	
$+(-0.0092254) \times \tilde{m}^3 \tilde{k}^3$	
E230H-2513: $\tilde{m} \equiv \frac{(m-308)}{10}$ ; $\tilde{k} \equiv \frac{(m\lambda-772300)}{500}$	
$v =$ $(-0.0738475) + (0.0954277) \times \tilde{m} + (-0.0416483) \times \tilde{k} + (0.0130481) \times \tilde{m} \tilde{k}$ etc.	

Note. —

Table 4. Stage Zero Coadditions: Subexposures

Dataset	Mode/Setting	Aperture ( $0.01'' \times 0.01''$ )	Band ( $\text{\AA}$ )	$t_{\text{exp}}$ (s)	Modified JD (d)	Qual. (%)	Xcorr Params. ( $\text{\AA}$ or $\text{km s}^{-1}$ )
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
AD-LEO							
o61s01010_1	E140M-1425	020×020	1140-1729	2200	51613.145	3	
o61s01020_1	E140M-1425	020×020	1140-1729	2700	51613.205	4	
o61s01030_1	E140M-1425	020×020	1140-1729	2700	51613.272	4	
o61s01040_1	E140M-1425	020×020	1140-1729	2700	51613.339	3	
o61s01050_1	E140M-1425	020×020	1140-1729	2700	51613.406	4	
o61s02010_1	E140M-1425	020×020	1140-1729	2200	51614.151	3	
o61s02020_1	E140M-1425	020×020	1140-1729	2700	51614.211	6	
o61s02030_1	E140M-1425	020×020	1140-1729	2700	51614.278	4	
o61s02040_1	E140M-1425	020×020	1140-1729	2700	51614.345	4	
o61s02050_1	E140M-1425	020×020	1140-1729	2700	51614.412	5	
o61s03010_1	E140M-1425	020×020	1140-1729	2200	51615.090	3	
o61s03020_1	E140M-1425	020×020	1140-1729	2700	51615.150	4	
o61s03030_1	E140M-1425	020×020	1140-1729	2700	51615.217	5	
o61s03040_1	E140M-1425	020×020	1140-1729	2700	51615.284	4	
o61s03050_1	E140M-1425	020×020	1140-1729	2700	51615.351	4	
o61s04010_1	E140M-1425	020×020	1140-1729	2200	51616.096	6	
o61s04020_1	E140M-1425	020×020	1140-1729	2700	51616.156	3	
o61s04030_1	E140M-1425	020×020	1140-1729	2700	51616.223	3	
o61s04040_1	E140M-1425	020×020	1140-1729	2700	51616.290	3	
o61s04050_1	E140M-1425	020×020	1140-1729	2700	51616.357	4	

Table 4. —Continued

(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
o6jg01010_1	E140M-1425	020×020	1140-1709	1750	52426.297	3	
o6jg01020_1	E140M-1425	020×020	1140-1709	2880	52426.350	4	
o6jg01030_1	E140M-1425	020×020	1140-1709	2880	52426.416	4	
o6jg01040_1	E140M-1425	020×020	1140-1709	2880	52426.483	4	
o6jg02010_1	E140M-1425	020×020	1140-1709	1750	52427.233	3	
o6jg02020_1	E140M-1425	020×020	1140-1709	2880	52427.285	4	
AG-DRA							
o6ky01010_1	E140M-1425	020×020	1140-1729	2600	52755.815	93	
o6ky01020_1	E140M-1425	020×020	1140-1729	1837	52755.848	91	
o6ky01040_1	E230M-1978	020×020	1607-2365	241	52755.881	54	
o6ky01050_1	E230M-1978	020×020	1607-2365	1872	52755.886	82	
o6ky01030_1	E230M-2707	020×020	2275-3118	374	52755.873	99	
AO-PSC							
o50151010_1	E140M-1425	020×020	1140-1709	2140	51740.494	90	
o50151020_1	E140M-1425	020×020	1140-1709	2540	51740.552	92	
o50151030_1	E140M-1425	020×020	1140-1709	2680	51740.619	91	
o50151040_1	E140M-1425	020×020	1140-1709	2270	51740.686	90	
o50152010_1	E140M-1425	020×020	1140-1709	2140	51741.432	90	
o50152020_1	E140M-1425	020×020	1140-1709	2540	51741.491	91	
o50152030_1	E140M-1425	020×020	1140-1709	2680	51741.558	91	
o50152040_1	E140M-1425	020×020	1140-1709	2270	51741.625	91	
BG-PHE							
o63502010_1	E140M-1425	020×020	1140-1709	1440	52041.327	97	
BW-SCL							
o5b616010_1	E140M-1425	020×020	1140-1729	1977	51433.056	2	
CI-CAM							

Table 4—Continued

(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
o5ci01010_1	E140M-1425	020×020	1140-1729	2242	51623.065	11	
o5ci01020_1	E140M-1425	020×020	1140-1729	1350	51623.120	2	+0.00
o5ci01020_2	E140M-1425	020×020	1140-1729	1540	51623.138	3	+0.00
<b>o5ci01020</b>				2890	51623.120	11	...
o5ci01030_1	E140M-1425	020×020	1140-1729	1350	51623.187	1	+0.00
o5ci01030_2	E140M-1425	020×020	1140-1729	1540	51623.206	2	+0.00
<b>o5ci01030</b>				2890	51623.187	9	...
o5ci01040_1	E230M-2269	020×020	1857-2672	1350	51623.254	35	+0.00
o5ci01040_2	E230M-2269	020×020	1857-2672	1646	51623.271	37	+0.68
<b>o5ci01040</b>				2996	51623.254	41	2328.48 [0.38]
etc.							

Note. — The objects now are listed alphabetically. This is a complete list of the STIS echelle exposures incorporated in StarCAT. Multiple subexposures, if any, are explicitly identified (by `x1d.fits` extension number) and grouped together, followed by the name of the coadded spectrum (and its properties). The “Modified JD” in column (6) is  $\text{JD} - 2,400,000$ . The “Quality” factor in column (7) is the percentage of the valid flux points in the band  $\geq 2.5\sigma$  with respect to the local photometric error (per resolution element): fewer than 1% of the flux points should exceed this significance level by chance (one-sided Gaussian  $\sigma$ ). In column (8), the values for the subexposures are the derived cross-correlation shifts, expressed in equivalent velocities ( $\text{km s}^{-1}$ ). (The first spectrum has  $v \equiv 0$  by definition.) For the coadded spectrum, the first value in column (8), if any, is the wavelength ( $\text{\AA}$ ) of the cross-correlation template feature, followed (parenthetically) by the half-width of the correlation band (also in  $\text{\AA}$ ). If template parameters are not reported, then a blind coaddition was performed (all shifts set to zero).

Table 5. Stage One Coadditions: Like-Mode, Same Visit

Dataset	Mode/Setting	Aperture (0.01"×0.01")	Band (Å)	$t_{\text{exp}}$ (s)	Modified JD (d)	Qual. (%)	Xcorr Params. (Å or km s <sup>-1</sup> )
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
AD-LEO							
o61s01010	E140M-1425	020×020	1140-1729	2200	51613.145	3	+0.00
o61s01020	E140M-1425	020×020	1140-1729	2700	51613.205	4	+0.88
o61s01030	E140M-1425	020×020	1140-1729	2700	51613.272	4	+0.60
o61s01040	E140M-1425	020×020	1140-1729	2700	51613.339	3	+0.50
o61s01050	E140M-1425	020×020	1140-1729	2700	51613.406	4	-0.41
<b>E140M-1425_020X020_51613</b>				13000	51613.145	10	1306.08 [0.12]
o61s02010	E140M-1425	020×020	1140-1729	2200	51614.151	3	+0.00
o61s02020	E140M-1425	020×020	1140-1729	2700	51614.211	6	-0.08
o61s02030	E140M-1425	020×020	1140-1729	2700	51614.278	4	-0.17
o61s02040	E140M-1425	020×020	1140-1729	2700	51614.345	4	-0.47
o61s02050	E140M-1425	020×020	1140-1729	2700	51614.412	5	-0.51
<b>E140M-1425_020X020_51614</b>				13000	51614.151	12	1306.08 [0.12]
o61s03010	E140M-1425	020×020	1140-1729	2200	51615.090	3	+0.00
o61s03020	E140M-1425	020×020	1140-1729	2700	51615.150	4	-0.09
o61s03030	E140M-1425	020×020	1140-1729	2700	51615.217	5	-0.78
o61s03040	E140M-1425	020×020	1140-1729	2700	51615.284	4	+0.34
o61s03050	E140M-1425	020×020	1140-1729	2700	51615.351	4	-0.03
<b>E140M-1425_020X020_51615</b>				13000	51615.090	12	1306.08 [0.12]
o61s04010	E140M-1425	020×020	1140-1729	2200	51616.096	6	+0.00
o61s04020	E140M-1425	020×020	1140-1729	2700	51616.156	3	+0.29
o61s04030	E140M-1425	020×020	1140-1729	2700	51616.223	3	+0.38
o61s04040	E140M-1425	020×020	1140-1729	2700	51616.290	3	+0.55
o61s04050	E140M-1425	020×020	1140-1729	2700	51616.357	4	+0.42
<b>E140M-1425_020X020_51616</b>				13000	51616.096	12	1306.08 [0.12]

Table 5. —Continued

(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
o6jg01010	E140M-1425	020×020	1140-1709	1750	52426.297	3	+0.00
o6jg01020	E140M-1425	020×020	1140-1709	2880	52426.350	4	−1.03
o6jg01030	E140M-1425	020×020	1140-1709	2880	52426.416	4	−0.76
o6jg01040	E140M-1425	020×020	1140-1709	2880	52426.483	4	−0.23
<b>E140M-1425_020X020_52426</b>				10390	52426.297	8	1306.08 [0.12]
o6jg02010	E140M-1425	020×020	1140-1709	1750	52427.233	3	+0.00
o6jg02020	E140M-1425	020×020	1140-1709	2880	52427.285	4	−0.26
<b>E140M-1425_020X020_52427</b>				4630	52427.233	5	1306.08 [0.12]
AG-DRA							
o6ky01010	E140M-1425	020×020	1140-1729	2600	52755.815	93	+0.00
o6ky01020	E140M-1425	020×020	1140-1729	1837	52755.848	91	−1.48
<b>E140M-1425_020X020_52755</b>				4437	52755.815	94	1639.58 [0.60]
o6ky01040	E230M-1978	020×020	1607-2365	241	52755.881	54	+0.00
o6ky01050	E230M-1978	020×020	1607-2365	1872	52755.886	82	−0.56
<b>E230M-1978_020X020_52755</b>				2113	52755.881	83	1891.10 [0.34]
o6ky01030	E230M-2707	020×020	2275-3118	374	52755.873	99	
AO-PSC							
o50151010	E140M-1425	020×020	1140-1709	2140	51740.494	90	+0.00
o50151020	E140M-1425	020×020	1140-1709	2540	51740.552	92	−0.13
o50151030	E140M-1425	020×020	1140-1709	2680	51740.619	91	+1.11
o50151040	E140M-1425	020×020	1140-1709	2270	51740.686	90	+0.73
<b>E140M-1425_020X020_51740</b>				9630	51740.494	95	1334.49 [0.14]
o50152010	E140M-1425	020×020	1140-1709	2140	51741.432	90	+0.00
o50152020	E140M-1425	020×020	1140-1709	2540	51741.491	91	−0.29
o50152030	E140M-1425	020×020	1140-1709	2680	51741.558	91	−0.40
o50152040	E140M-1425	020×020	1140-1709	2270	51741.625	91	+0.76
<b>E140M-1425_020X020_51741</b>				9630	51741.432	95	1334.49 [0.14]
etc.							

Note. — The objects are listed alphabetically. The like-mode datasets are grouped together by visit (if multiple observations were obtained and judged to be suitable for combining), followed

by the name of the coadded spectrum (and its properties). The columns are the same as for Table 4.

Table 6. Average Cross-correlation Shifts and Flux Density Scale Factors

Mode	$\langle v \rangle$ (km s <sup>-1</sup> )	$\sigma_v$	$N$
(1)	(2)	(3)	(4)
Zeroth and First Stage Coadditions			
E140M	+0.06	1.33	339 [28]
E140H	+0.00	0.45	280 [0]
E230M	+0.07	1.24	223 [6]
E230H	+0.03	0.60	242 [7]
Second Stage Coadditions			
E140M	-0.02	1.35	131 [7]
E140H	+0.01	0.26	52 [0]
E230M	-0.17	1.58	67 [1]
E230H	-0.00	0.33	75 [0]
Third Stage Splicing: $v$			
All	-0.06	1.09	372 [4]
Third Stage Splicing: $f_\lambda$ Scale Factors			
...	1.18	0.24	299 [0]

Note. — The Stage 0 and 1 coadditions operate on analogous groupings of exposures, taken close together in time, so the derived velocity shifts were combined. In column (4), the first value is

the number of samples within the bounds of the relevant histogram (e.g., Figs. 3–6:  $-5 \leq V_{\text{shft}} \leq +5 \text{ km s}^{-1}$  for the cross-correlation shifts and  $1.001 \leq F_{\text{scl}} \leq 3$  for the flux density scale factors). The parenthetical value is the number of outliers: these were not included when calculating the averages and standard deviations. The  $f_{\lambda}$  scale factor is dimensionless.

Table 7. Stage Two Coadditions: Like-Mode, Different Epoch and/or Aperture

Dataset	Mode/Setting	Aperture ( $0.01'' \times 0.01''$ )	Band ( $\text{\AA}$ )	$t_{\text{exp}}$ (s)	Modified JD (d)	Qual. (%)	Xcorr Params. ( $\text{\AA}$ or $\text{km s}^{-1}$ )
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
AD-LEO							
E140M-1425_020X020_51614			1140-1729	13000	51614.151	12	+0.37
E140M-1425_020X020_51615			1140-1729	13000	51615.090	12	+0.50
E140M-1425_020X020_51616			1140-1729	13000	51616.096	12	−0.95
E140M-1425_020X020_51613			1140-1729	13000	51613.145	10	+0.34
E140M-1425_020X020_52426			1140-1709	10390	52426.297	8	+0.30
E140M-1425_020X020_52427			1140-1709	4630	52427.233	5	−0.56
<b>E140M-1425_020X020_51613-52427</b>			1140-1729	67024	51613.145	31	1306.08 [0.12]
AG-DRA							
E140M-1425_020X020_52755			1140-1729	4437	52755.815	94	
E230M-1978_020X020_52755			1607-2365	2113	52755.881	83	
o6ky01030	E230M-2707	020×020	2275-3118	374	52755.873	99	
AO-PSC							
E140M-1425_020X020_51741			1140-1709	9630	51741.432	95	+0.49
E140M-1425_020X020_51740			1140-1709	9630	51740.494	95	−0.49
<b>E140M-1425_020X020_51740-51741</b>			1140-1709	19261	51740.494	96	1334.49 [0.14]
BG-PHE							
o63502010	E140M-1425	020×020	1140-1709	1440	52041.327	97	

Table 7. —Continued

(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
BW-SCL							
o5b616010	E140M-1425	020×020	1140-1729	1977	51433.056	2	
CI-CAM							
E140M-1425_020X020_51623			1140-1729	8022	51623.065	47	
o5ci01040	E230M-2269	020×020	1857-2672	2996	51623.254	41	
CY-TAU							
E140M-1425_020X020_51884			1140-1729	5178	51884.739	3	
o5cf03010	E230M-1978	020×020	1607-2365	1260	51884.671	1	
E230M-2707_020X020_51884			2275-3118	1020	51884.690	7	
DG-TAU							
E140M-1425_020X006_51839			1140-1729	12295	51839.408	0	
E230M-2707_600X020_52974			2275-3118	4800	52974.980	3	
E230M-2707_020X020_51960			2274-3117	11114	51960.193	86	
DN-LEO							
o66v04010	E140M-1425	020×020	1140-1729	144	51951.170	90	
E230M-2269_020X006_52653			1841-2673	4721	52653.932	99	
DR-TAU							
E140M-1425_020X020_51785			1140-1729	5794	51785.668	8	
o5cf02010	E230M-1978	020×020	1607-2365	1080	51785.609	13	
E230M-2707_020X020_51949			2274-3117	5207	51949.284	91	−0.72
o5cf02020	E230M-2707	020×020	2276-3119	916	51785.625	85	+0.72
<b>E230M-2707_020X020_51785-51949</b>			2274-3117	6123	51785.625	92	2796.55 [0.48]
etc.							

Note. — The objects are listed alphabetically. The like-mode, different epoch and/or different aperture datasets are grouped together (if multiple observations were obtained and judged to be suitable for combining), followed by the name (and properties) of the coadded spectrum. Note that the datasets in a grouping now are ordered in descending quality factor, so that the template spectrum (first one) will be optimum from a S/N standpoint. The columns are the same as in Table 4.

Table 8. Stage Three Coadditions: Wavelength Splicing

Dataset	Mode/Setting	Aperture ( $0.01'' \times 0.01''$ )	Band ( $\text{\AA}$ )	$t_{\text{exp}}$ (s)	Modified JD (d)	Qual. (%)	Fscl (8)	Vshft ( $\text{km s}^{-1}$ )	Xcorr Params. ( $\text{\AA}$ )
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
AD-LEO									
E140M-1425_020X020_51613-52427			1140.0-1729.0	67024	51613.145	31			
AG-DRA									
E140M-1425_020X020_52755			1150.0-1705.0	4437	52755.815	94	1.025	+1.23	...
E230M-1978_020X020_52755			1610.1-2347.3	2113	52755.881	83	1.025	-0.29	1639.61 [0.37]
o6ky01030	E230M-2707	020 $\times$ 020	2280.0-3118.3	374	52755.873	99	1.000	-0.94	2305.69 [0.42]
<b>UVSUM_1M_52755</b>			1150.0-3118.3	6550	52755				
AO-PSC									
E140M-1425_020X020_51740-51741			1140.0-1709.0	19261	51740.494	96			
BG-PHE									
o63502010	E140M-1425	020 $\times$ 020	1140.0-1709.0	1440	52041.327	97			

Table 8. —Continued

(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
BW-SCL									
o5b616010	E140M-1425	020×020	1140.0-1729.0	1977	51433.056	2			
CI-CAM									
E140M-1425_020X020_51623			1150.0-1708.7	8022	51623.065	47	1.000	+0.00	...
o5ci01040	E230M-2269	020×020	1861.3-2672.6	2996	51623.254	41	1.000	+0.00	...
<b>UVSUM_1M_51623</b>			1150.0-2672.6	8022	51623				
CY-TAU									
E140M-1425_020X020_51884			1150.0-1705.1	5178	51884.739	3	1.000	-0.04	...
o5cf03010	E230M-1978	020×020	1610.1-2342.6	1260	51884.671	1	1.000	-0.04	...
E230M-2707_020X020_51884			2279.9-3118.2	1020	51884.690	7	1.000	+0.08	2326.21 [0.38]
<b>UVSUM_1M_51884</b>			1150.0-3118.2	6438	51884				
DG-TAU									
E140M-1425_020X006_51839			1150.0-1699.9	12295	51839.408	0	1.000	+0.00	...
E230M-2707_020X020_51960			2289.7-3117.7	11114	51960.193	86	1.000	+0.00	...
<b>UVSUM_1M_51839-51960</b>			1150.0-3117.7	12295	51839				
etc.									

Note. — The specific datasets that were spliced are grouped together (if multiple observations were obtained and judged to be suitable for combining), followed by the name (and properties) of the spliced spectrum. The columns are the same as in Table 7, except with the addition of a column (new [8]) for the flux scale factor (“Fsc1”), and now the derived velocity shift and the cross-correlation template parameters are listed on each line (columns [9] and [10]). The second exposure of a sequence is assigned template parameters referring to the overlap with the first spectrum; the third exposure (if any) will have the parameters for the spectrum 2 and 3 overlap; and so forth. The velocity shifts (“Vshft”) initially were benchmarked to the first exposure, then the average for the whole group was subtracted from each entry. If Vshfts and template parameters are blank, then the adjacent spectra either were simply concatenated, if there was no overlap between them; or were “blindly spliced” in cases of overlapping wavelengths, but insufficient S/N for a reliable Vshft determination or simply lack of a suitably constrasty template feature. Similarly, an Fsc1 of 1.000 in a list of generally  $> 1$  values would typically be the spectrum with the highest apparent throughput, against which the other wavelength regions were normalized. Unit Fsc1 also will occur as the default if the segments of an exposure group were not overlapping, including cases of insufficient S/N in one, or both, components to reliably measure the flux density ratio.

Table 9. StarCAT Final Catalog

Object Name	Dataset	Mode/Setting	Aperture ( $0.01'' \times 0.01''$ )	Band ( $\text{\AA}$ )	$t_{\text{exp}}$ (s)	Modified JD (d)	Qual. (%)
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
AD-LEO	E140M-1425_020X020_51613-52427			1140-1729	67024	51613.145	31
AG-DRA	UVSUM_1M_52755			1150-3118	6550	52755	96
AO-PSC	E140M-1425_020X020_51740-51741			1140-1709	19261	51740.494	96
BG-PHE	o63502010	E140M-1425	020 $\times$ 020	1140-1709	1440	52041.327	97
BW-SCL	o5b616010	E140M-1425	020 $\times$ 020	1140-1729	1977	51433.056	2
CI-CAM	UVSUM_1M_51623			1150-2672	8022	51623	45
CY-TAU	UVSUM_1M_51884			1150-3118	6438	51884	3
DG-TAU	UVSUM_1M_51839-51960			1150-3117	12295	51839	29
DN-LEO	UVSUM_1M_51951-52653			1150-2673	4721	51951	95
DR-TAU	UVSUM_1M_51785-51949			1150-3117	7203	51785	31
DS-TAU	UVSUM_1M_51780			1150-3119	6986	51780	33
EF-PEG	E140M-1425_020X020_51713			1140-1709	6883	51713.080	0
EK-TRA	o5b612010	E140M-1425	020 $\times$ 020	1140-1709	4302	51384.772	39
EV-LAC	E140M-1425_020X020_52172			1140-1729	10920	52172.698	4
EX-HYA	E140M-1425_020X020_51682-51694			1140-1709	15201	51682.894	97
FQ-AQR	E230M-2269_020X006_52527			1840-2672	7620	52527.226	99
FU-ORI	E230M-2707_020X020_51962			2274-3117	5217	51962.936	72
HD100340	o63557010	E140M-1425	020 $\times$ 020	1140-1729	1440	51999.496	97

Table 9. —Continued

(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
HD100546	UVSUM_1H_51747			1150-2887	3134	51747	95
	UVSUM_2H_51852			1494-2887	1970	51852	99
HD101131	o6lz48010	E140H-1489	020×020	1390-1586	300	52740.936	99
HD101190	o6lz49010	E140H-1489	020×020	1390-1586	300	52577.357	99
HD101436	o6lz51010	E140H-1489	020×020	1390-1586	600	52786.196	99
HD102065	UVSUM_1H_50900-50901			1150-2888	8879	50900	97
HD102634	o4ao06010	E230M-2707	020×006	2275-3118	1440	51128.343	100
HD103095	UVSUM_1X_52839			1150-3158	11097	52839	53
HD103779	UVSUM_1H_51284-51959			1160-1901	1466	51284	97
HD104705	UVSUM_1H_51171			1160-1586	6220	51171	97
HD106343	UVSUM_1H_51284			1163-1901	1486	51284	96
HD106516	UVSUM_1M_51163-51202			1150-3118	1728	51163	34
HD106943	UVSUM_1H_52786			1160-1551	2220	52786	95
HD10700	E140M-1425_020X020_51757			1140-1709	13450	51757.124	18
HD107113	UVSUM_1M_50736			1607-3118	3481	50736	88
HD107213	UVSUM_1M_51378			1607-3116	3539	51378	76
HD107969	E140M-1425_020X020_52281			1140-1729	5197	52281.514	97
HD108	UVSUM_1H_51677			1163-1901	3988	51677	95
etc.							

Note. — Refer to Tables 4, 5, 7, and 8 for the specific groupings of exposures in each final dataset. In a few cases for which complete spectral coverage was available in both M and H modes, these were separately spliced. In most situations, however, the single final UVSUM spectrum is either entirely M, or entirely H, and only occasionally a mixture of the two (“X-type”).

Fig. 1.— Visualization of STIS echellegrams in the four “uber modes.” In each frame, the free spectral range is divided into a series of slices (echelle orders), stacked one on top of another; usually with some wavelength overlap between the end (right side) of one order,  $m$ , and the beginning (left side) of the next,  $m - 1$ . The ordinate scale to the left marks the individual echelle orders; the scale to the right indicates the central wavelength of the particular order. Within each order, wavelengths increase from left to right. The schematic echellegram was assembled from multiple observations (up to 20, in the case of AD Leo) from the so-called x2d file, which provides a spatially resolved image of each order, wavelength corrected and with compensation for scattered light, on an absolute flux scale. The intensity stretch is linear and in negative (darker shading indicates higher fluxes). (a) E140M mode, setting 1425 Å, of the red dwarf flare star AD Leo. The dark streaks toward the bottom are H I  $\lambda$ 1215 Ly $\alpha$  (repeated in adjacent orders), strongly in emission from the  $\sim 10^4$  K stellar chromosphere. The stretch saturates the broad H I feature, but allows one to see the many other emissions present in this region. The slight dimples below the redward peaks of Ly $\alpha$  is Si III  $\lambda$ 1206 in the next order below. A faint continuum, which arises mainly in the chromosphere, also is visible. The smudge in the middle left of order 109 is the famous coronal forbidden line [Fe XXI]  $\lambda$ 1354, which forms at very high temperatures  $\sim 10^7$  K in the stellar outer atmosphere. The bright double feature visible on the left side of order 95 is C IV  $\lambda\lambda$ 1548,50, which arises at intermediate temperatures  $\sim 10^5$  K: the shortwavelength component also is seen at the end (right side) of the preceding order. (b) E140H mode, setting 1271 Å, of the peculiar Of star HD 108. The spectrum is dominated by continuum emission throughout the region, interrupted by numerous absorptions. The extended blank area between 1202 Å and 1220 Å is H I Ly $\alpha$ , now strongly in absorption, probably mainly interstellar. A prominent P-Cygni wind profile of the N V doublet is visible near 1238 Å: the darker streaks represent the blended blueshifted emission peaks (repeated in adjacent orders), and the lighter streaks just below are the absorption trough. The double narrow dips just shortward of the N V peak, repeated in orders 340 and 341, are interstellar Mg II. (c) E230M mode, setting 2707 Å, of YZ CMi, another red dwarf flare star. The bright pair of emissions at the center of order 73 is the Mg II  $\lambda\lambda$ 2796,2802 doublet, a prominent radiative cooling channel for stellar chromospheres. A faint continuum, of photospheric origin, also is visible. Note that the higher, more crowded orders now are at top (reverse of previous E140 frames), and wavelengths decrease from bottom to top, although they still increase from left to right within an order. (d) E230H mode, setting 2513 Å, of the white dwarf G 191-B2B (WD0501+527 in StarCAT). Here, the spectrum is a smooth and mostly featureless photospheric continuum, only occasionally interrupted by a narrow interstellar absorption line. The bland appearance well suits the use of this object as a radiometric calibrator.

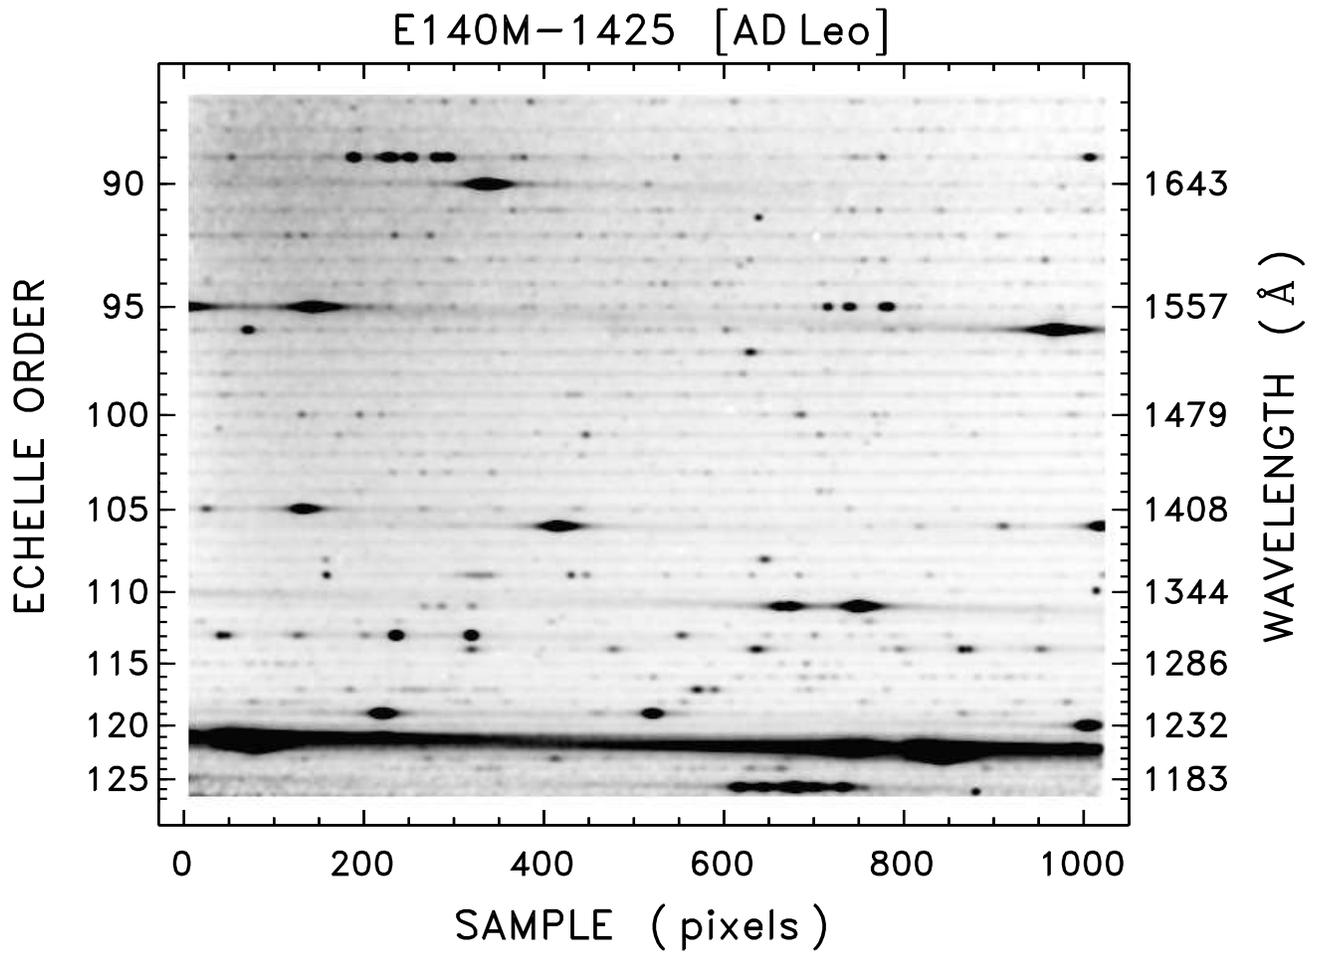


Fig. 1a.—

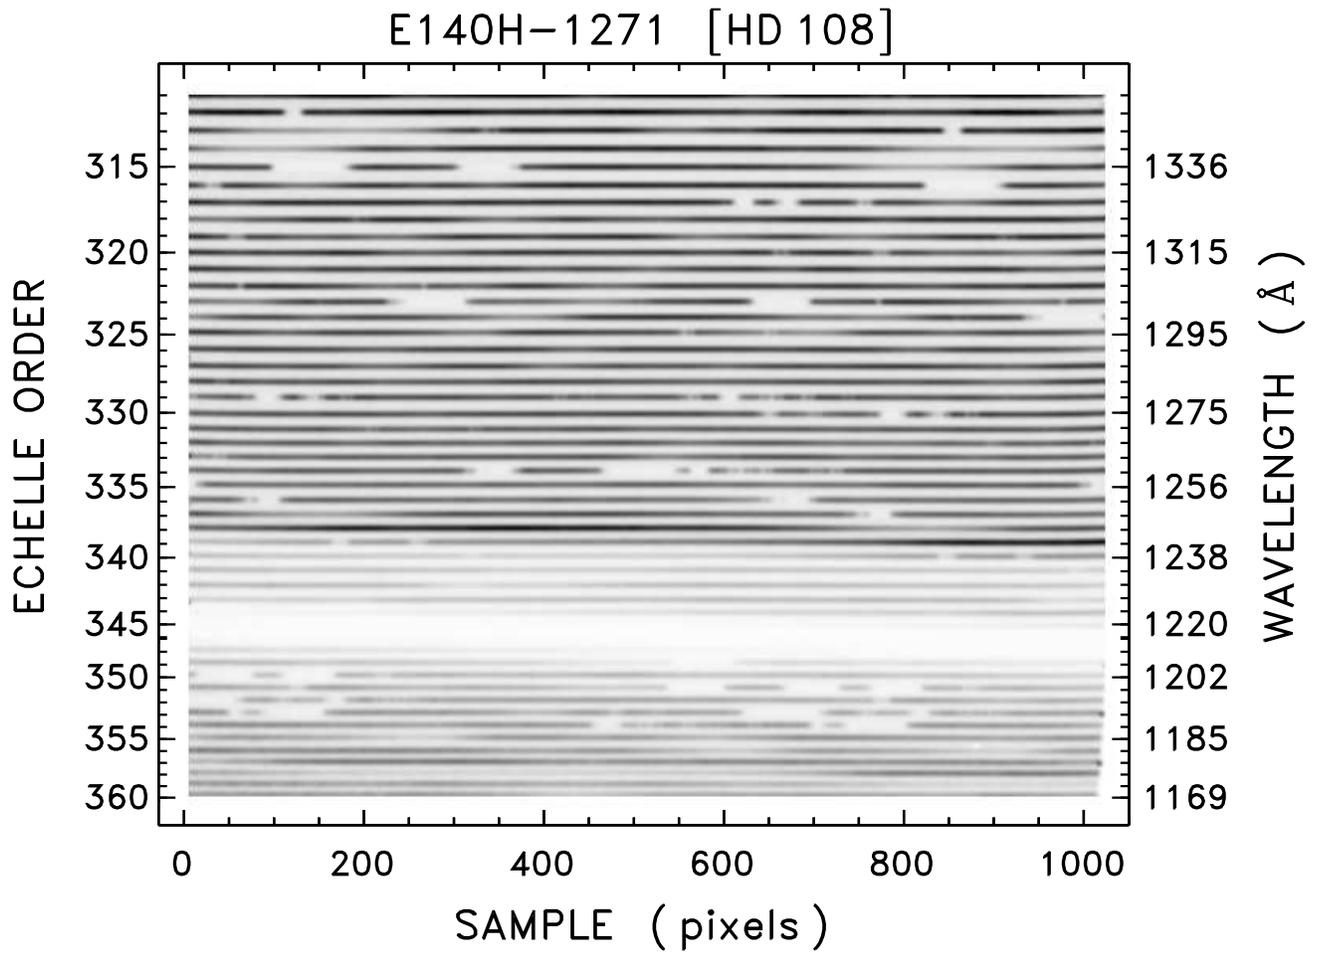


Fig. 1b.—

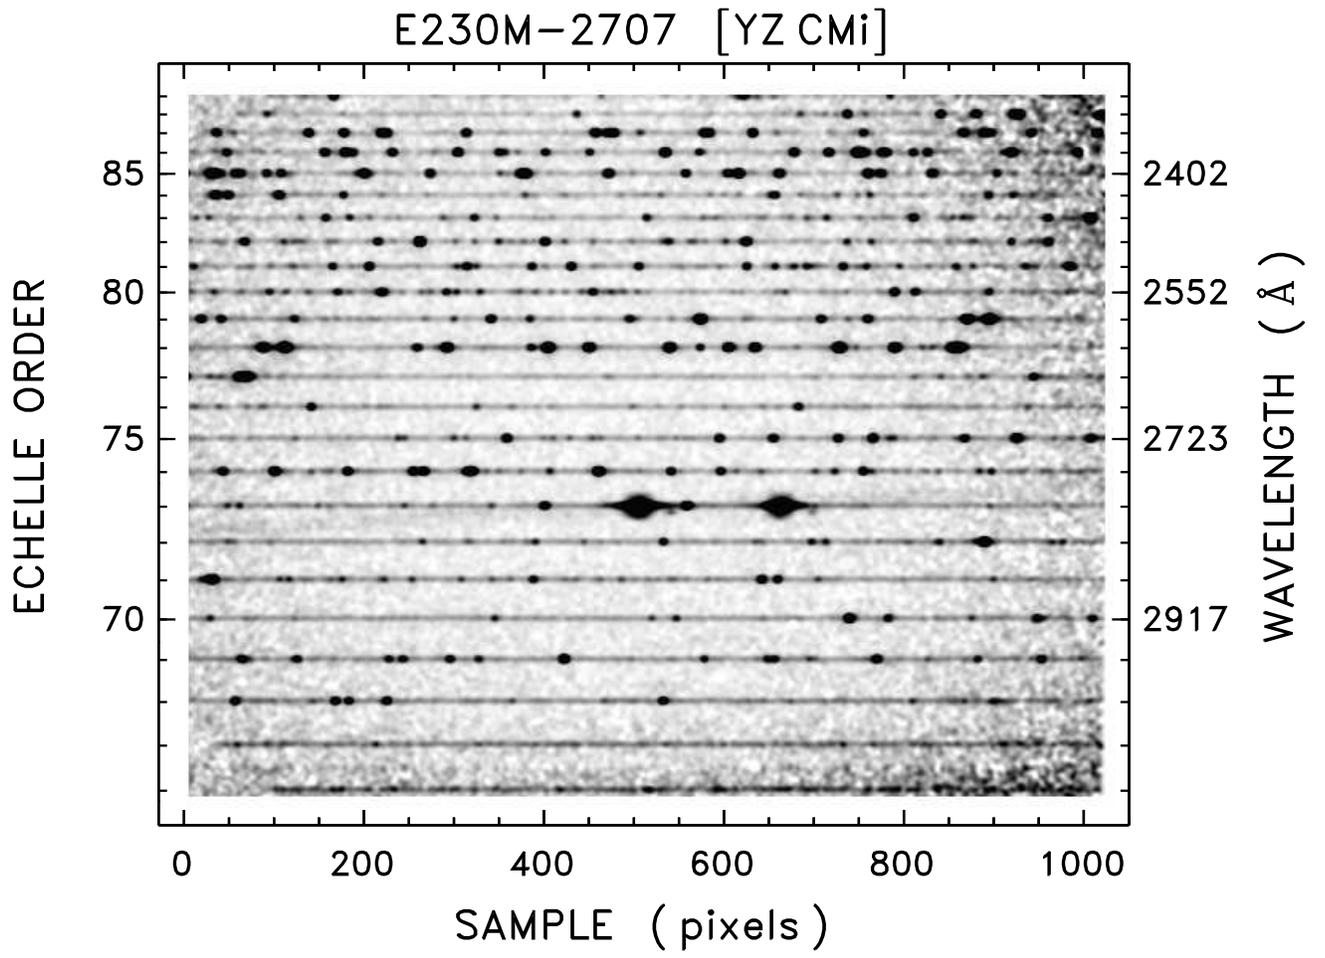


Fig. 1c.—

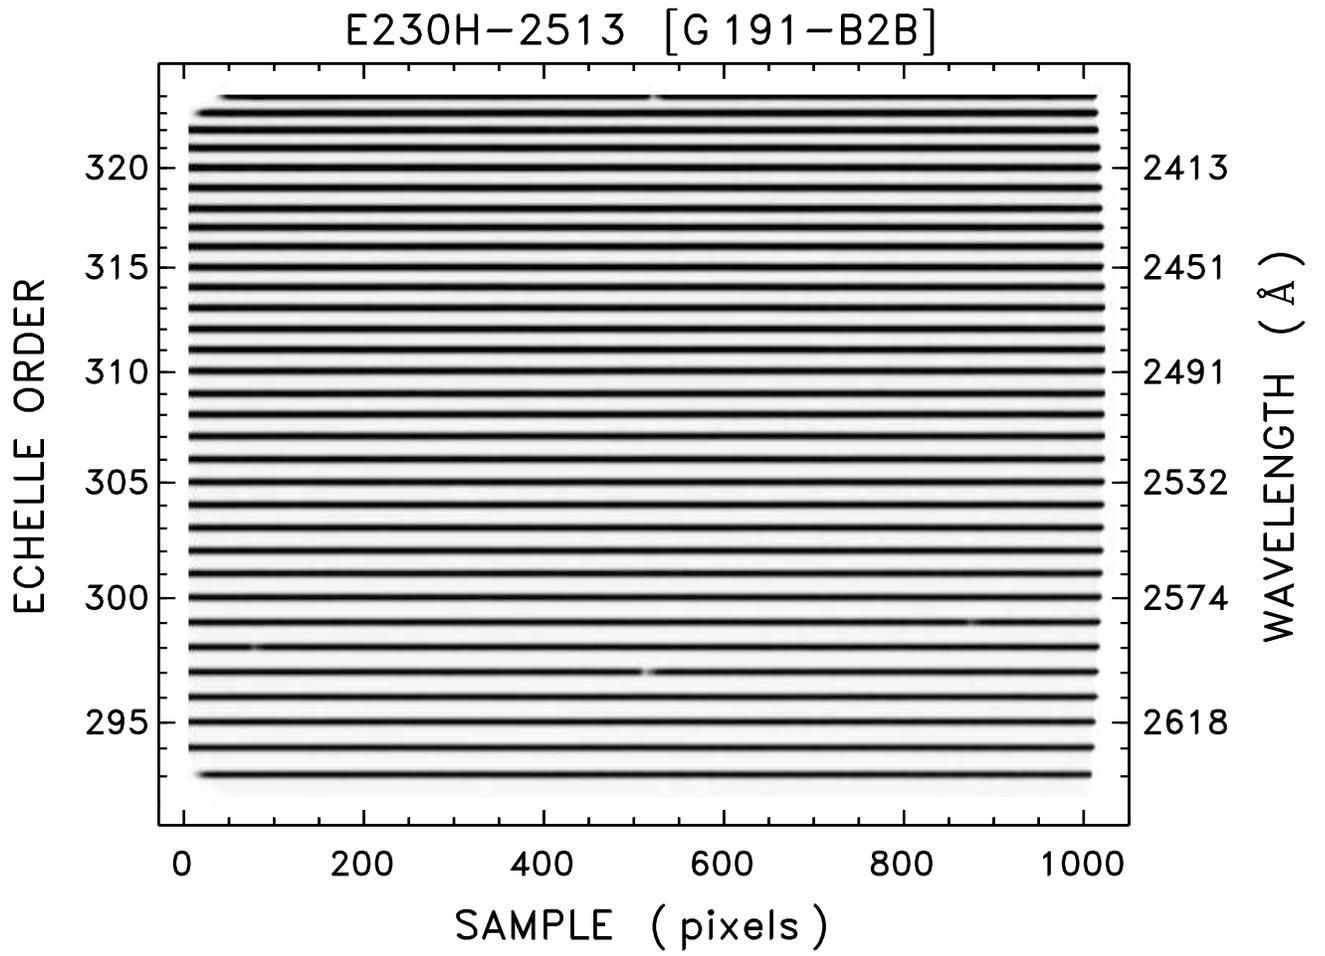


Fig. 1d.—

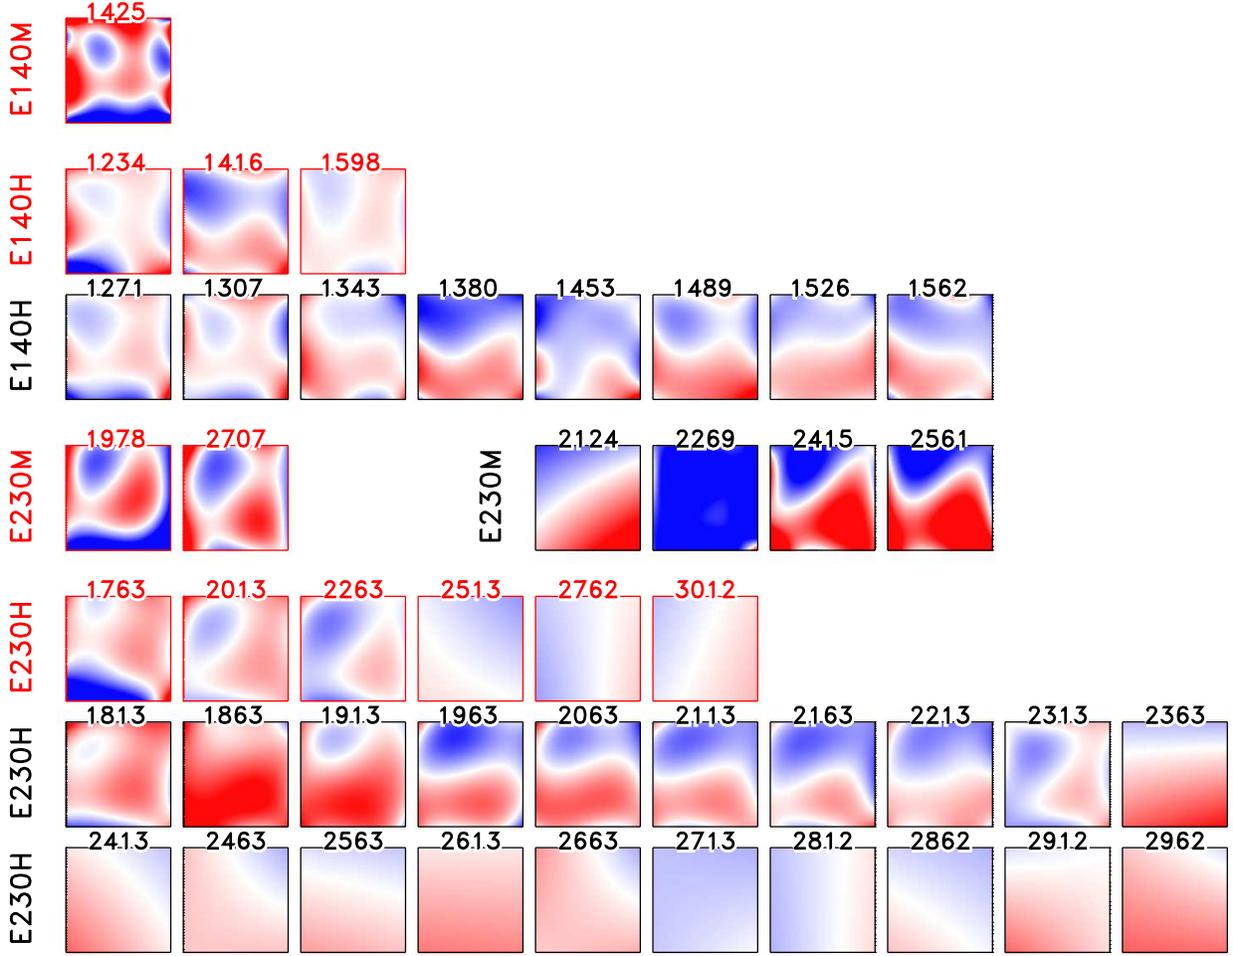


Fig. 2.— Schematic wavelength scale distortion maps for the 44 supported ‘tilts’ of the 4 main echelle modes of STIS. The format is analogous to that of Figs. 1 (orders on ordinate, wavelengths on abscissa; low orders at top for E140, but at bottom for E230). The maps were constructed from the polynomial models of Table 3, treating  $m$  as a continuous, rather than discrete, variable. The ordinate also assumes a linear variation of  $m$  between the  $m_{\min}$  to  $m_{\max}$  for the particular mode/tilt, rather than the nonlinear spacing in the real echellegram. Red shading indicates a redshift; blue, a blueshift; and the darkest colors represent an amplitude of about  $2 \text{ km s}^{-1}$ . The frames bordered in red are the prime settings; the others (black borders) are secondary tilts. In general, the prime settings are better calibrated than the secondary ones; and the high resolution modes show lower amplitude, smoother corrections than the medium resolution counterparts. This is because the distortions arise mainly from the relatively low order of the polynomial dispersion relations implemented in the pipeline, and become more conspicuous in the modes where the polynomial model has to cover more spectral “territory.”

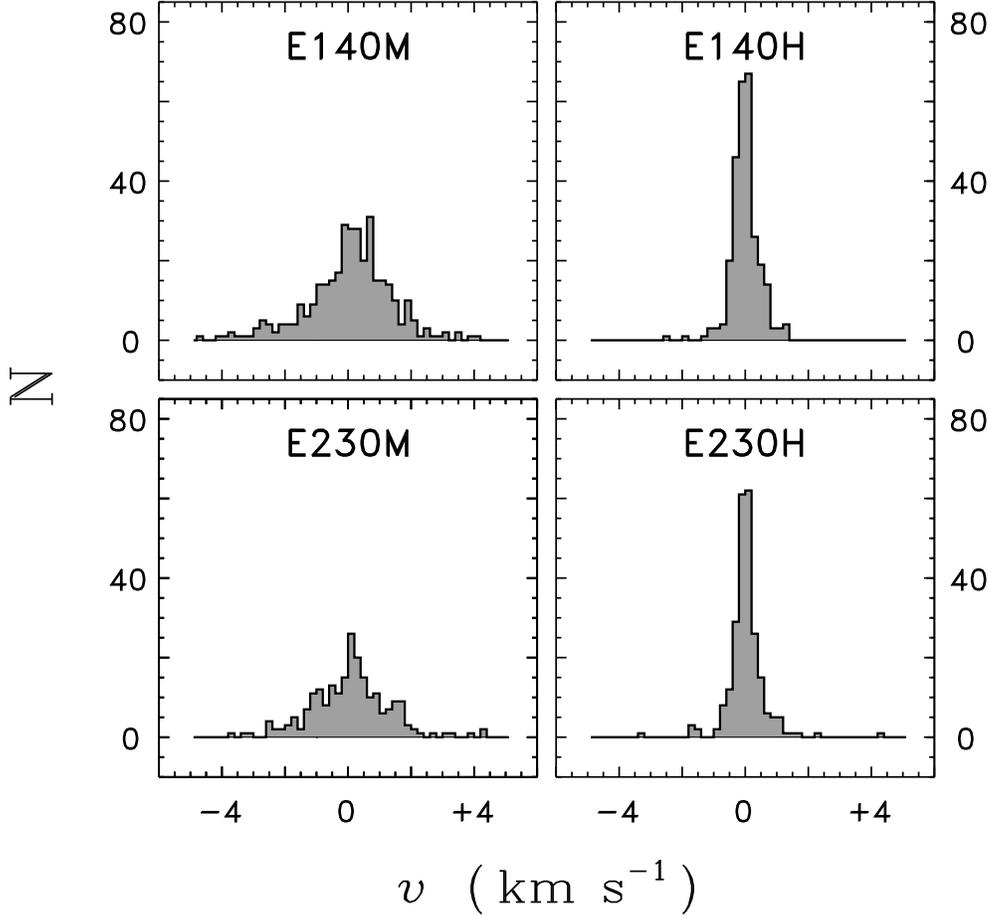


Fig. 3.— Histograms of velocity shifts measured in the Stage 0 and 1 coadditions, which combined series of like-mode spectrograms taken close together in time. The velocity shifts were determined by cross correlation, relative to the first exposure of each well defined sequence. Only the second and higher values are displayed ( $v \equiv 0$  for the first exposure, by definition), and only in cases of a successful cross-correlation solution. Binning is in steps of  $0.2 \text{ km s}^{-1}$ . In the two high resolution modes, one detector pixel (1K $\times$ 1K format) corresponds to about  $1.3 \text{ km s}^{-1}$ , whereas in medium resolution 1 pixel is about  $3 \text{ km s}^{-1}$  (E140M) or  $5 \text{ km s}^{-1}$  (E230M). Thus, the different distributions would appear more similar if cast in pixels rather than velocity. (Consistent with the idea that the shifts devolve mainly from small, random spatial offsets, since the camera pixel scale is synonymous with the plate scale.)

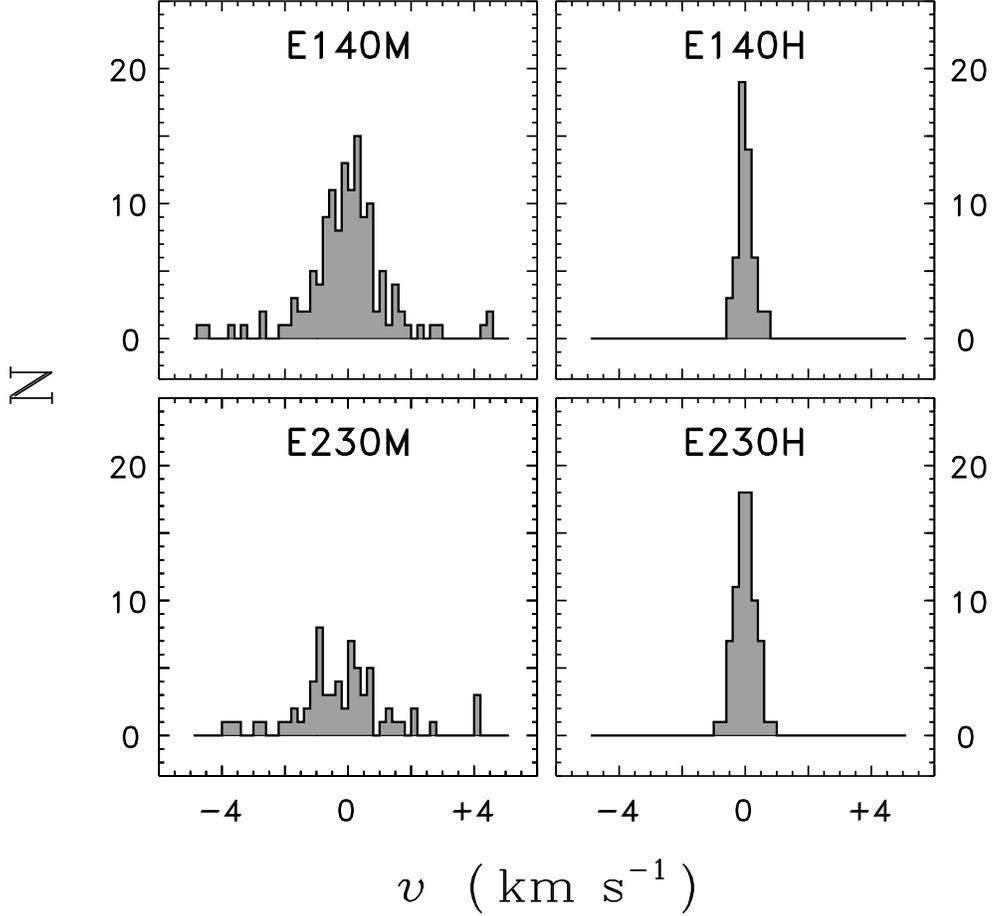


Fig. 4.— Histograms of velocity shifts derived in the Stage 2 coaddition process, which combined series of like-mode spectrograms taken in different epochs (or, in some cases, with different apertures). The velocity shifts were determined by cross correlation, relative to the highest S/N exposure of each sequence; but then the average was subtracted (unlike the Stage 0 and 1 procedure) to avoid bias, and the values for all the datasets are displayed (excluding the cases for which a cross-correlation solution was not obtained). Binning again is in steps of  $0.2 \text{ km s}^{-1}$ . Although there are fewer samples in this comparison, again the H modes display narrower distributions than the M modes, and the widths are similar to the Stage 0 + 1 counterparts.

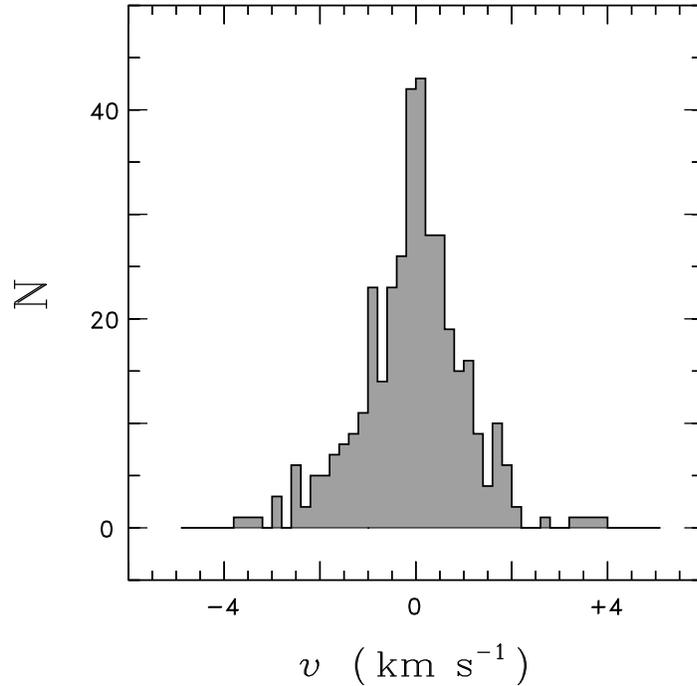


Fig. 5.— Histogram of velocity shifts derived in the Stage 3 splicing procedure, which joined spectra of the same object taken in different wavelength bands, perhaps separated by epoch in the event of conspicuous variability. The distributions were restricted to observations with a successful cross-correlation solution. Binning is in steps of  $0.2 \text{ km s}^{-1}$ . The composite distribution (which represents cross-correlations between pairs of M modes, pairs of H modes, and M-H pairs) has a hybrid appearance befitting the mixed nature of the constituent datasets: broader than the like-mode H profiles, but narrower than the M examples.

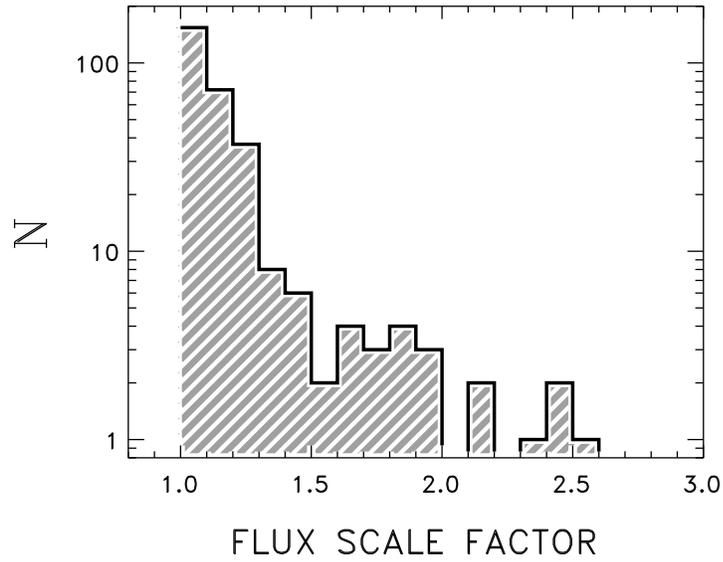


Fig. 6.— Histogram of flux scale factors derived in the stage 3 splicing, restricted to the best quality spectra. Note the logarithmic ordinate. Only the values  $\geq 1.001$  were included, to avoid the reference spectrum of each set (highest initial flux ratio: see text) whose scale factor was 1.000 by design. Binning is in steps of 0.1. The distribution falls rapidly toward higher scale factors. The modest secondary peak beyond  $F_{\text{scf}} \sim 1.5$  is contributed mainly by the tiny Jenkins slit and other nonstandard apertures.

Fig. 7.— Schematic global views of StarCAT final spectra for several representative objects. In each panel, the right hand frame provides a time line of the observations, showing the mode/setting as a vertical line at the elapsed time of the pointing relative to the initial epoch: orange and red marks are for E140M and E140H tilts, respectively; blue and green are for E230M and E230H, respectively. The ordinate of the panel runs from 1000 Å at the bottom to 3300 Å at the top. The gray zone delimits the FUV modes. The left hand panel illustrates the final spectrum (or spectra if more than one, say due to separation into individual epochs for variable objects) on a logarithmic flux scale, together with the  $1\sigma$  photometric error (per resampled bin: the noise level per resolution element would be smaller by a factor typically of 2). Both  $f_\lambda$  and  $\sigma_\lambda$  were smoothed, for display purposes, by two passes of a 51-bin running mean, following removal of gaps by eliminating the affected wavelengths. The first spectrum is coded thick blue and its error is thin red; the second spectrum (if any) is thick green, and its error is thin orange. If there were more than two spectra in a set, only the first two error curves are shown, to avoid confusion, and the third spectrum is displayed in thin black. For additional spectra, the color coding repeats. Frames (a)–(d) correspond to the objects of Fig. 1. (a) Red dwarf flare star AD Leo. (b) Peculiar Of star HD108. (c) Red dwarf flare star YZCMi. (d) White dwarf G 191-B2B (WD0501+527 in StarCAT; separated into M and H traces; the third, lower, curve is for a series of exposures through the Jenkins slit). (e) Herbig Be (B9 Vne) star KR Mus (HD100546 in StarCAT). (f) Symbiotic star RW Hya (M III:pe; HD117970 in StarCAT). (g) Red bright giant  $\alpha$  TrA (K2 II-III; HD150798 in StarCAT). (h) Algol eclipsing binary 31 Cyg (K2 III+; HD192577 in StarCAT). (i) Algol eclipsing binary VV Cep (M2 Iab+; HD208816 in StarCAT).

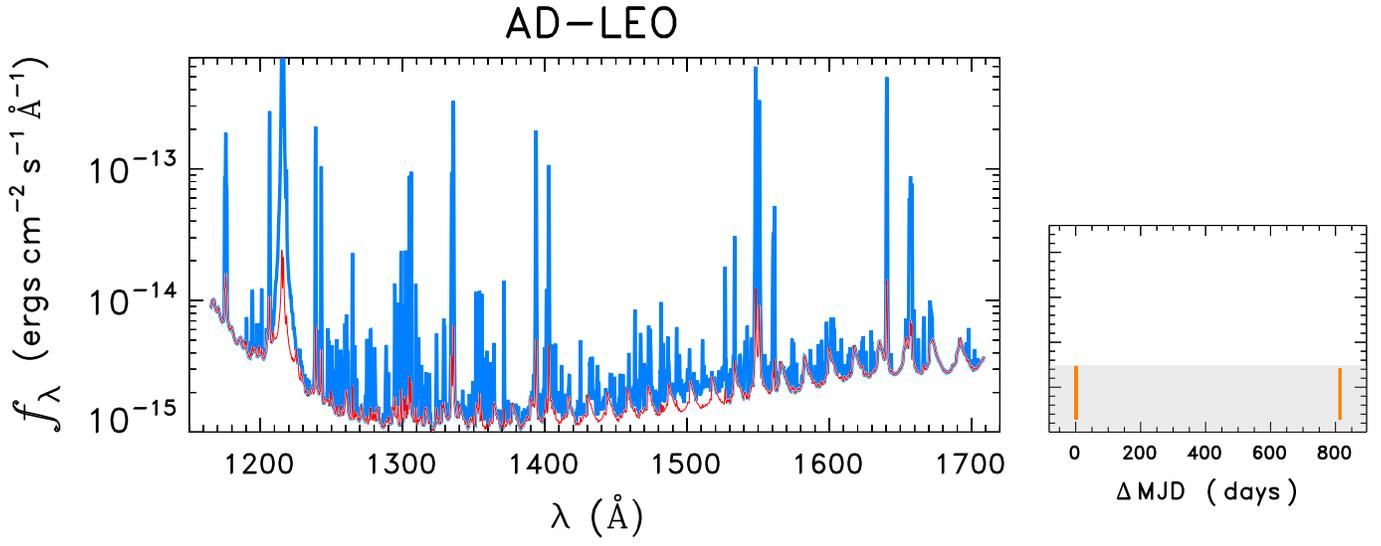


Fig. 7a.—

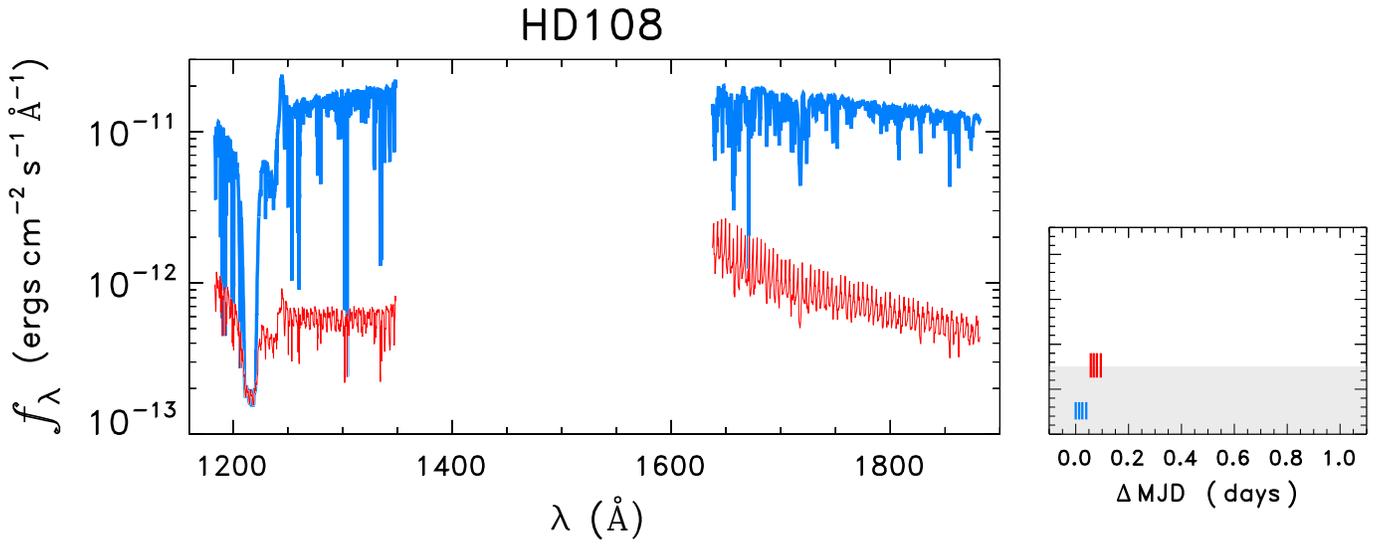


Fig. 7b.—

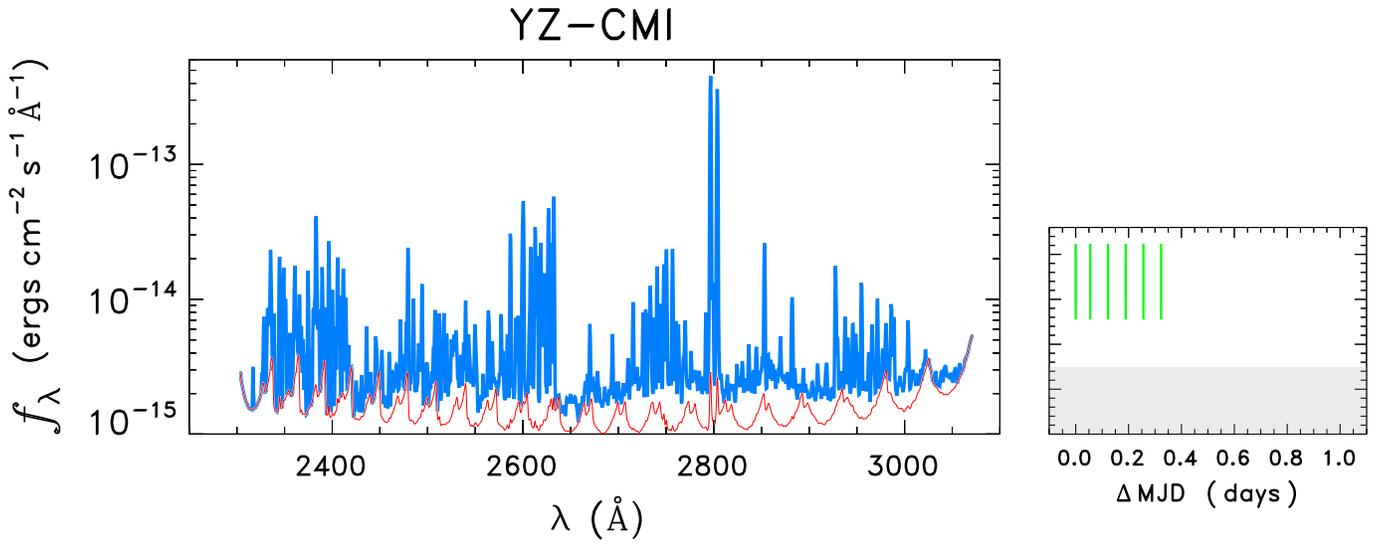


Fig. 7c.—

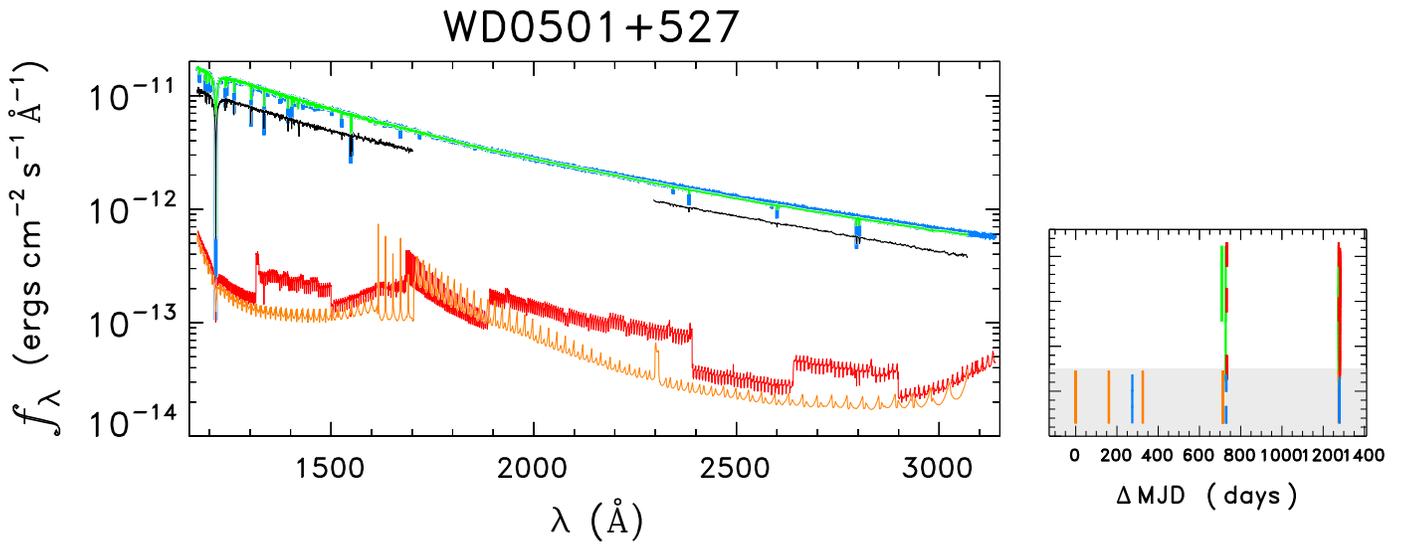


Fig. 7d.—

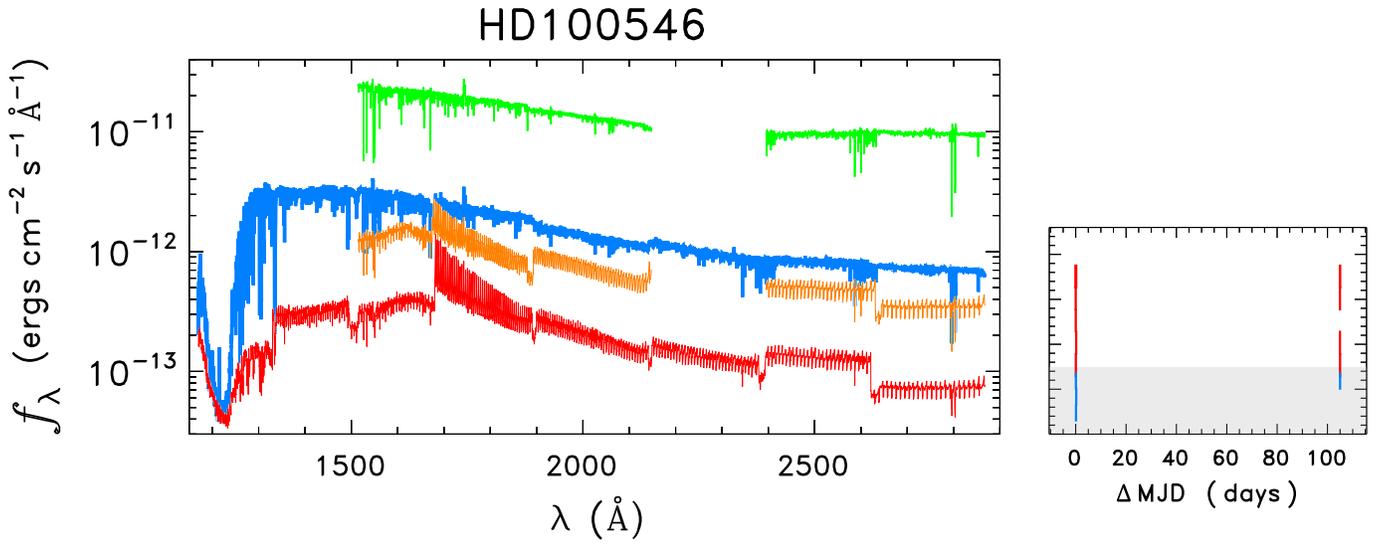


Fig. 7e.—

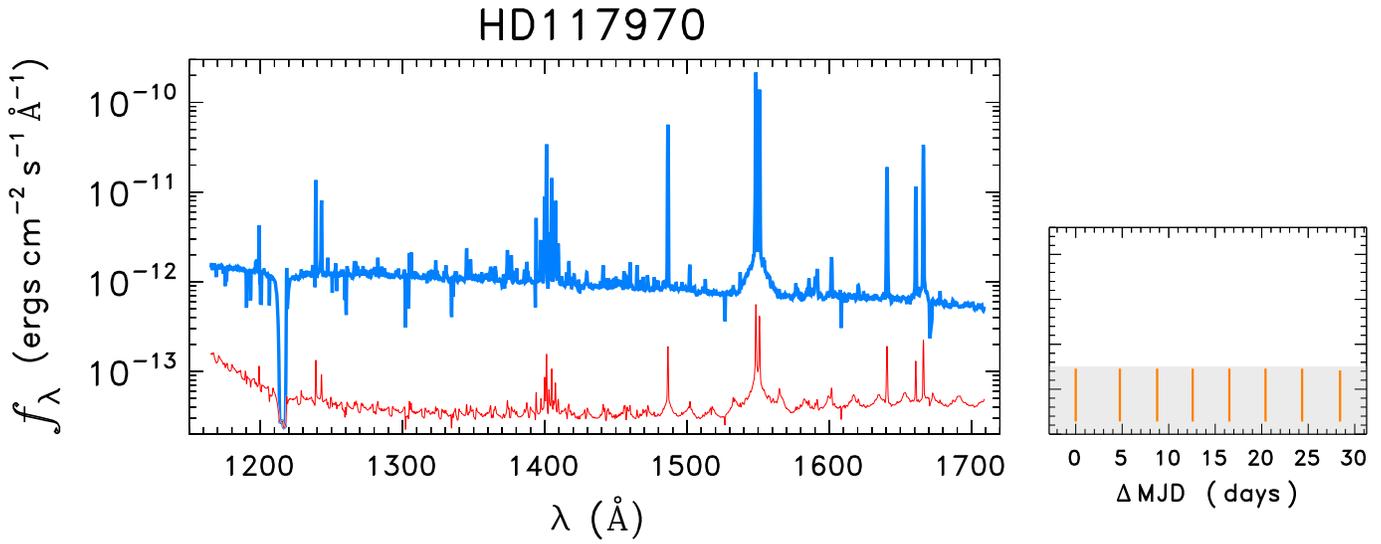


Fig. 7f.—

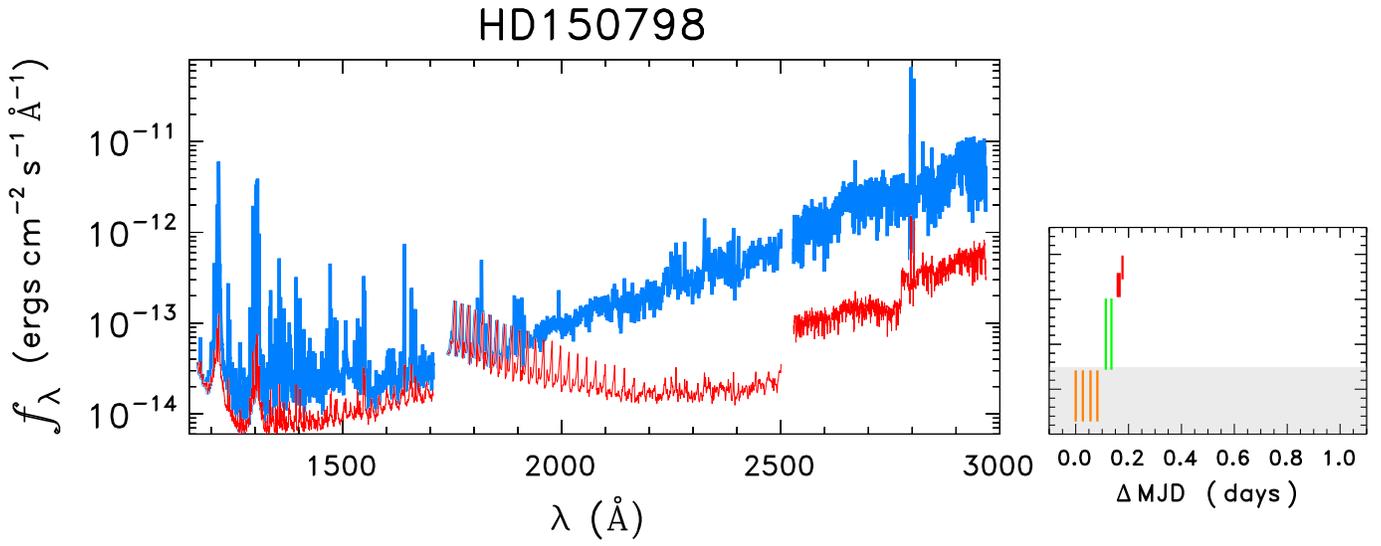


Fig. 7g.—

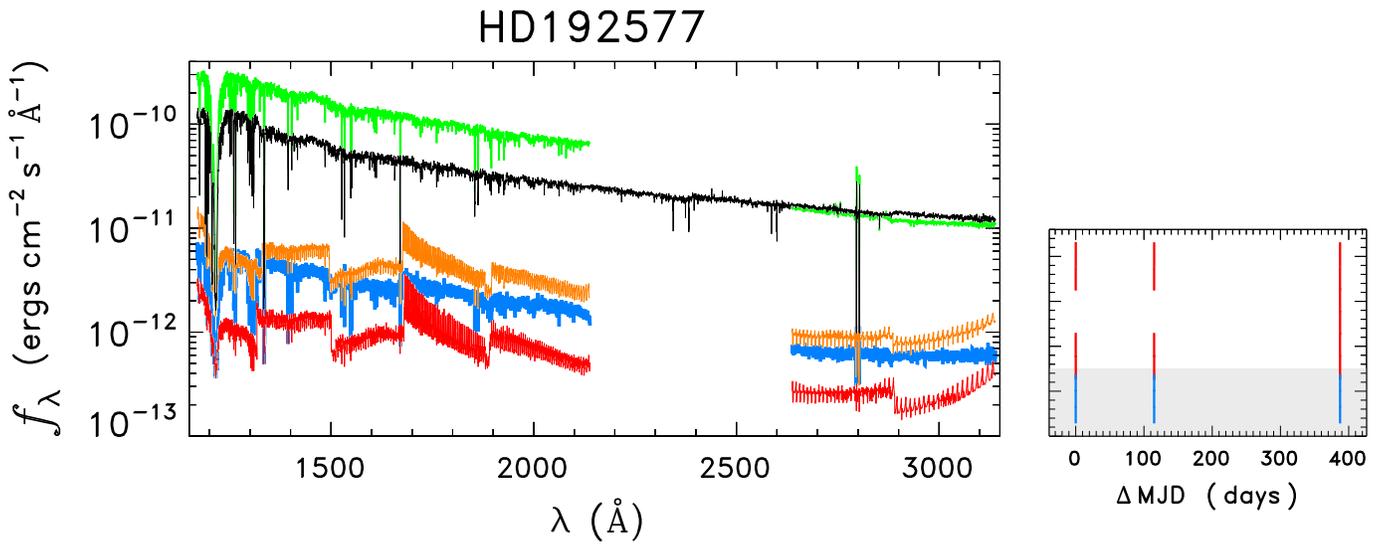


Fig. 7h.—

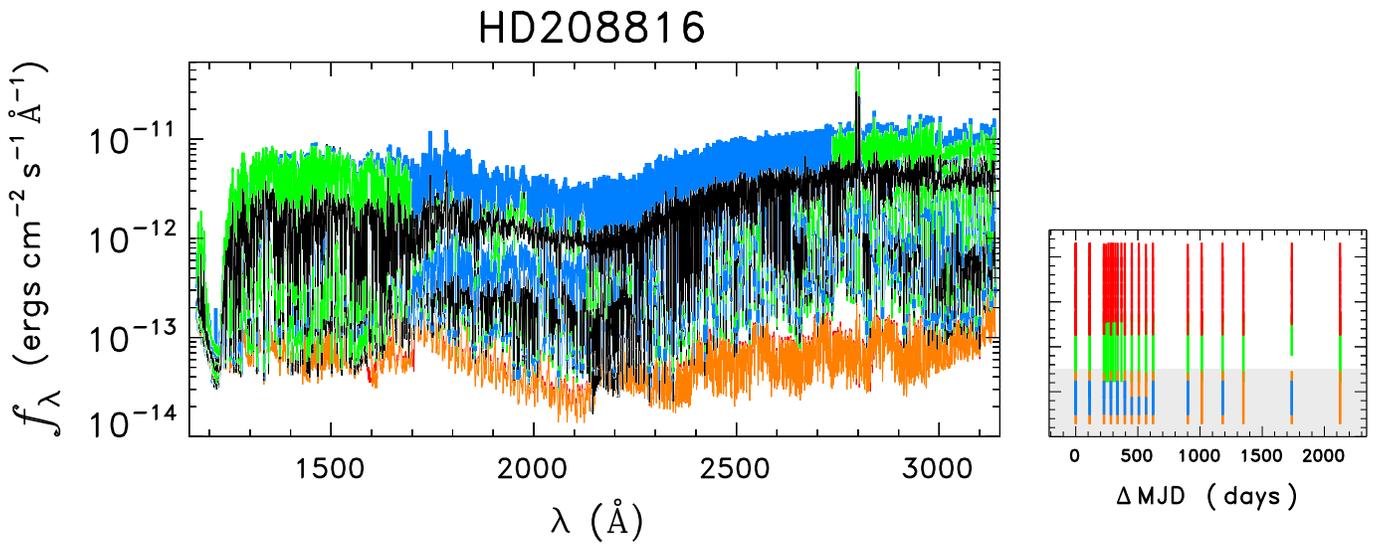


Fig. 7i.—

Fig. 8.— Higher resolution views of the final catalog spectra for representative objects (those illustrated in Figs. 1a–1d). The spectrum is the thin black trace and the photometric error is in red. Points flagged as “bad” (only a few of which normally survive the coaddition/splicing process) are marked with green X’s. The  $f_\lambda$  and  $\sigma_\lambda$  curves were smoothed by two passes of a running mean corresponding roughly to the local resolution element. Again, the  $\sigma_\lambda$  per resol would be a factor of  $\sim 2$  smaller. (a) Red dwarf flare star AD Leo. Note the pair of sharp spikes in  $\sigma_\lambda$  at 1168 Å and 1178 Å, caused by an unflagged intermittent detector bright spot (see, also, Fig. 1a). In the constituent echellegrams, the spot falls mainly in the interorder zone between  $m = 126$  and 127, leading to negative dips in the background subtracted spectra of these orders. (b) Peculiar Of star HD 108. (c) YZ CMi, another red dwarf flare star. (d) White dwarf G 191-B2B (WD0501+527 in StarCAT). Here, the spliced H-mode spectrum is shown.

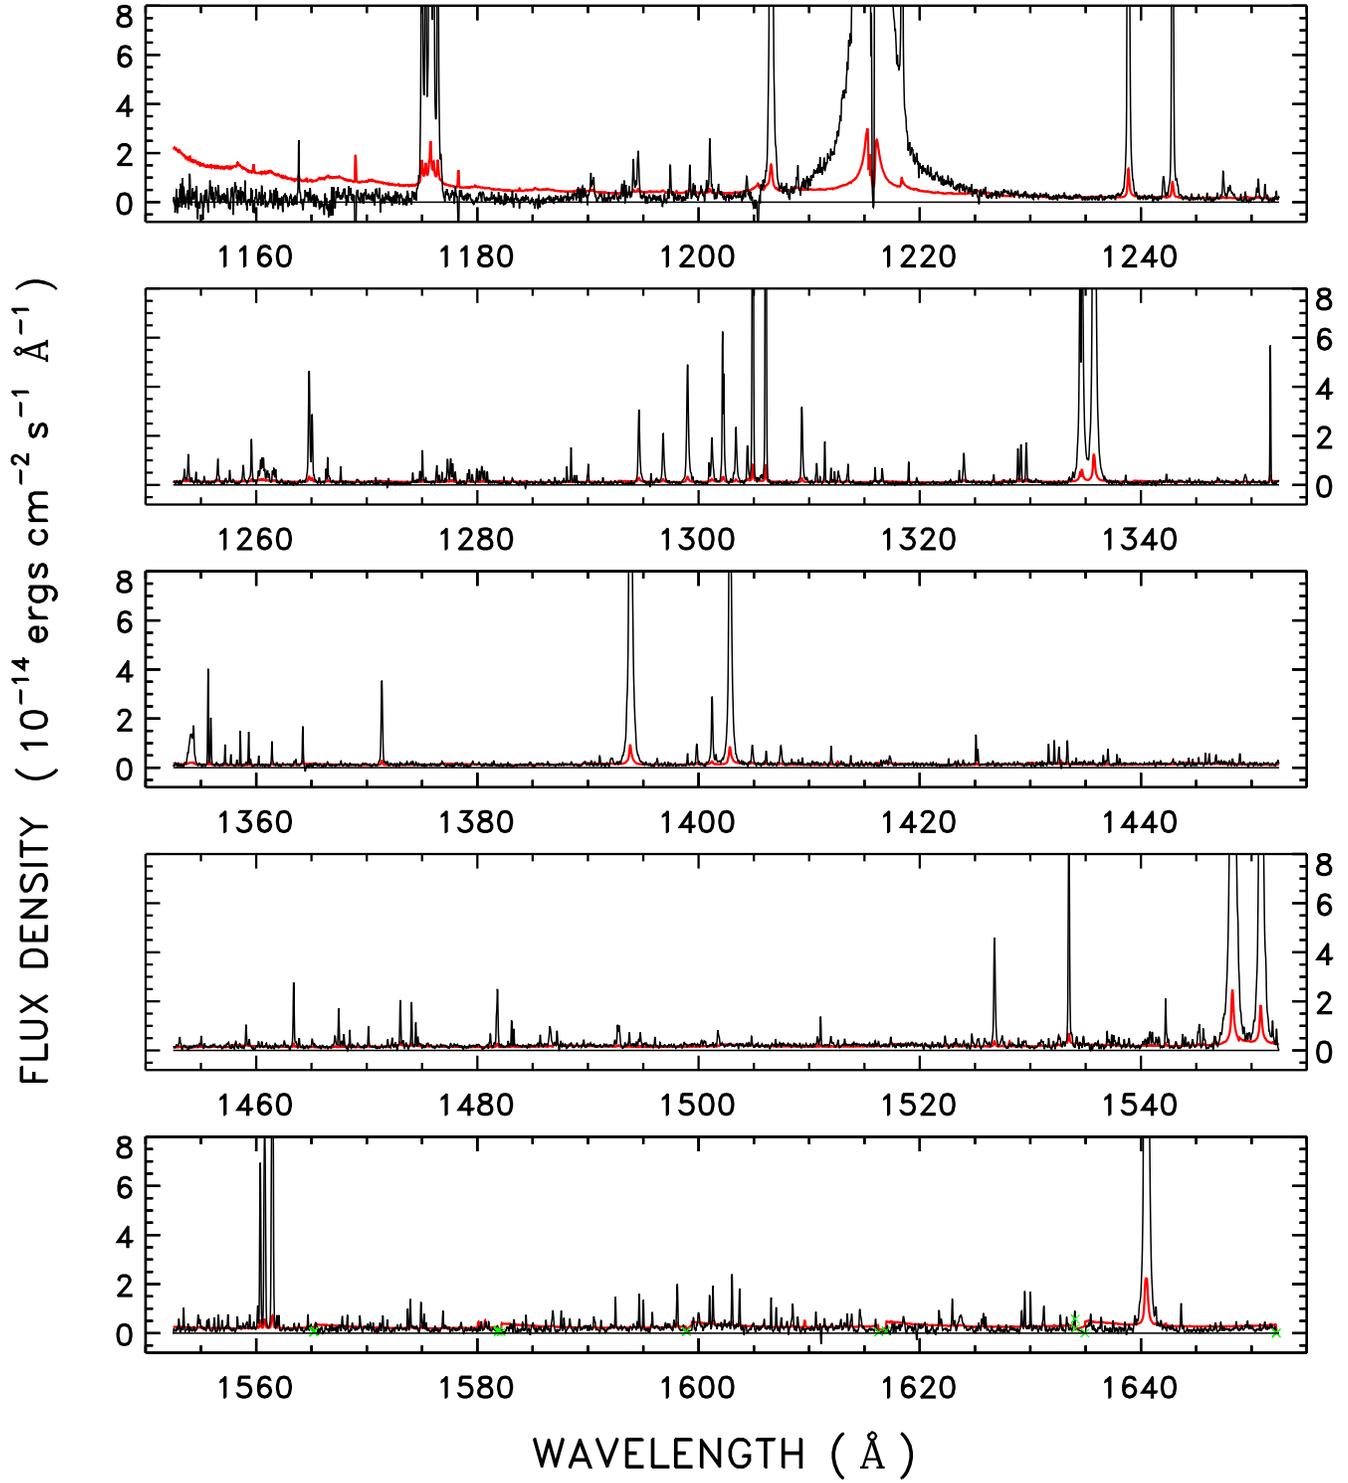


Fig. 8a.—

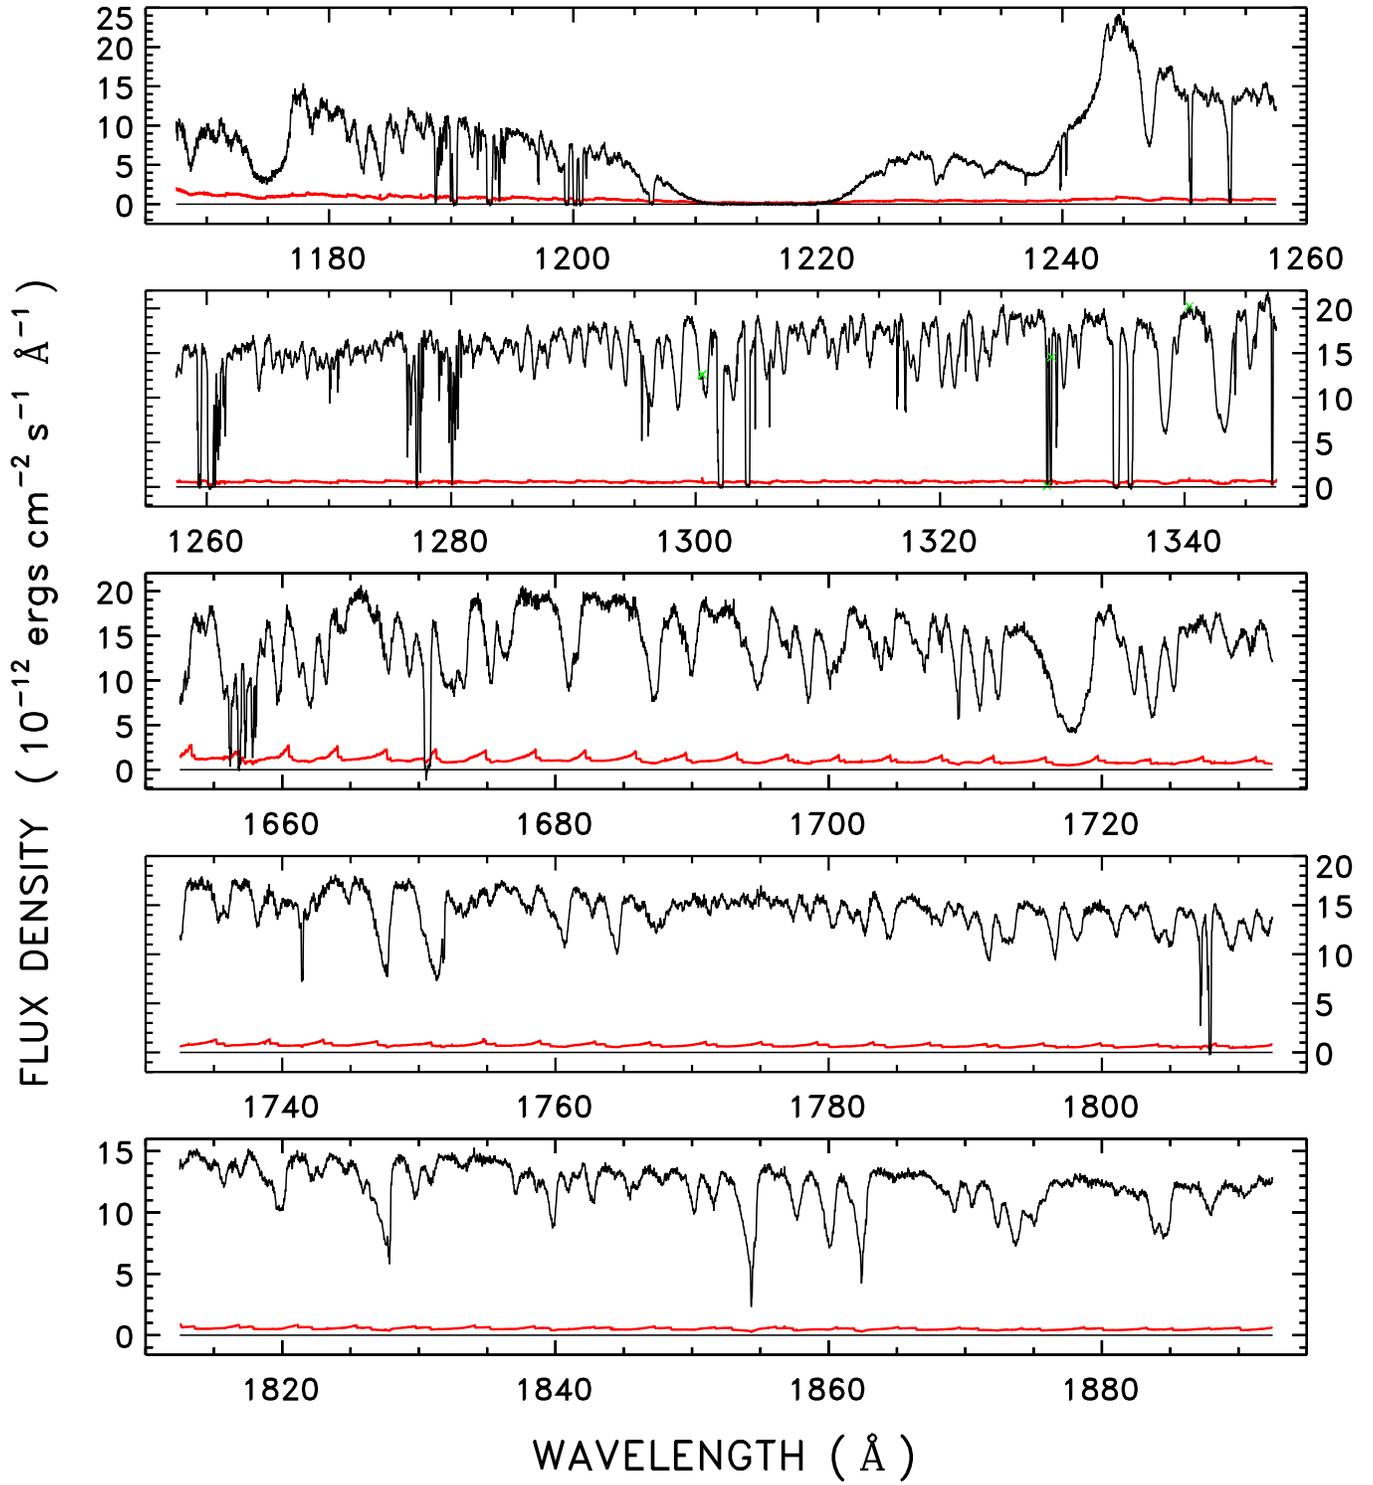


Fig. 8b.—

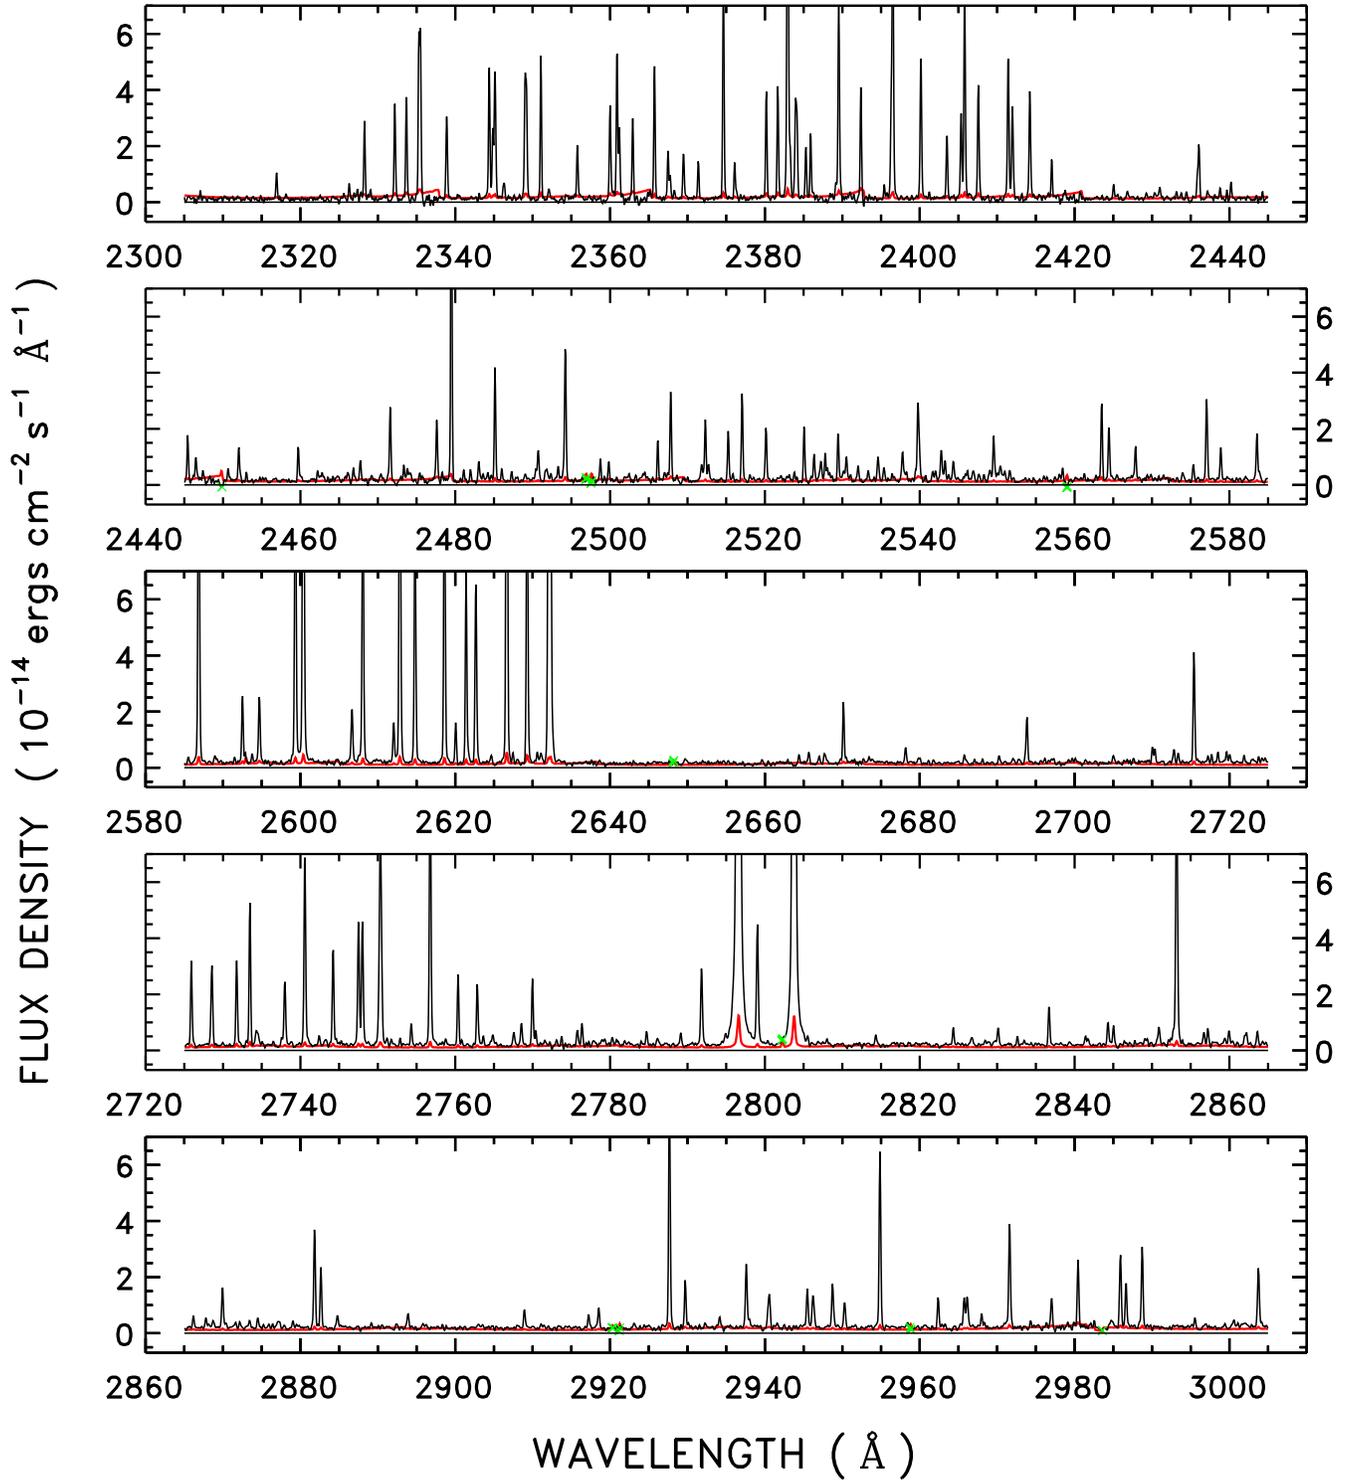


Fig. 8c.—

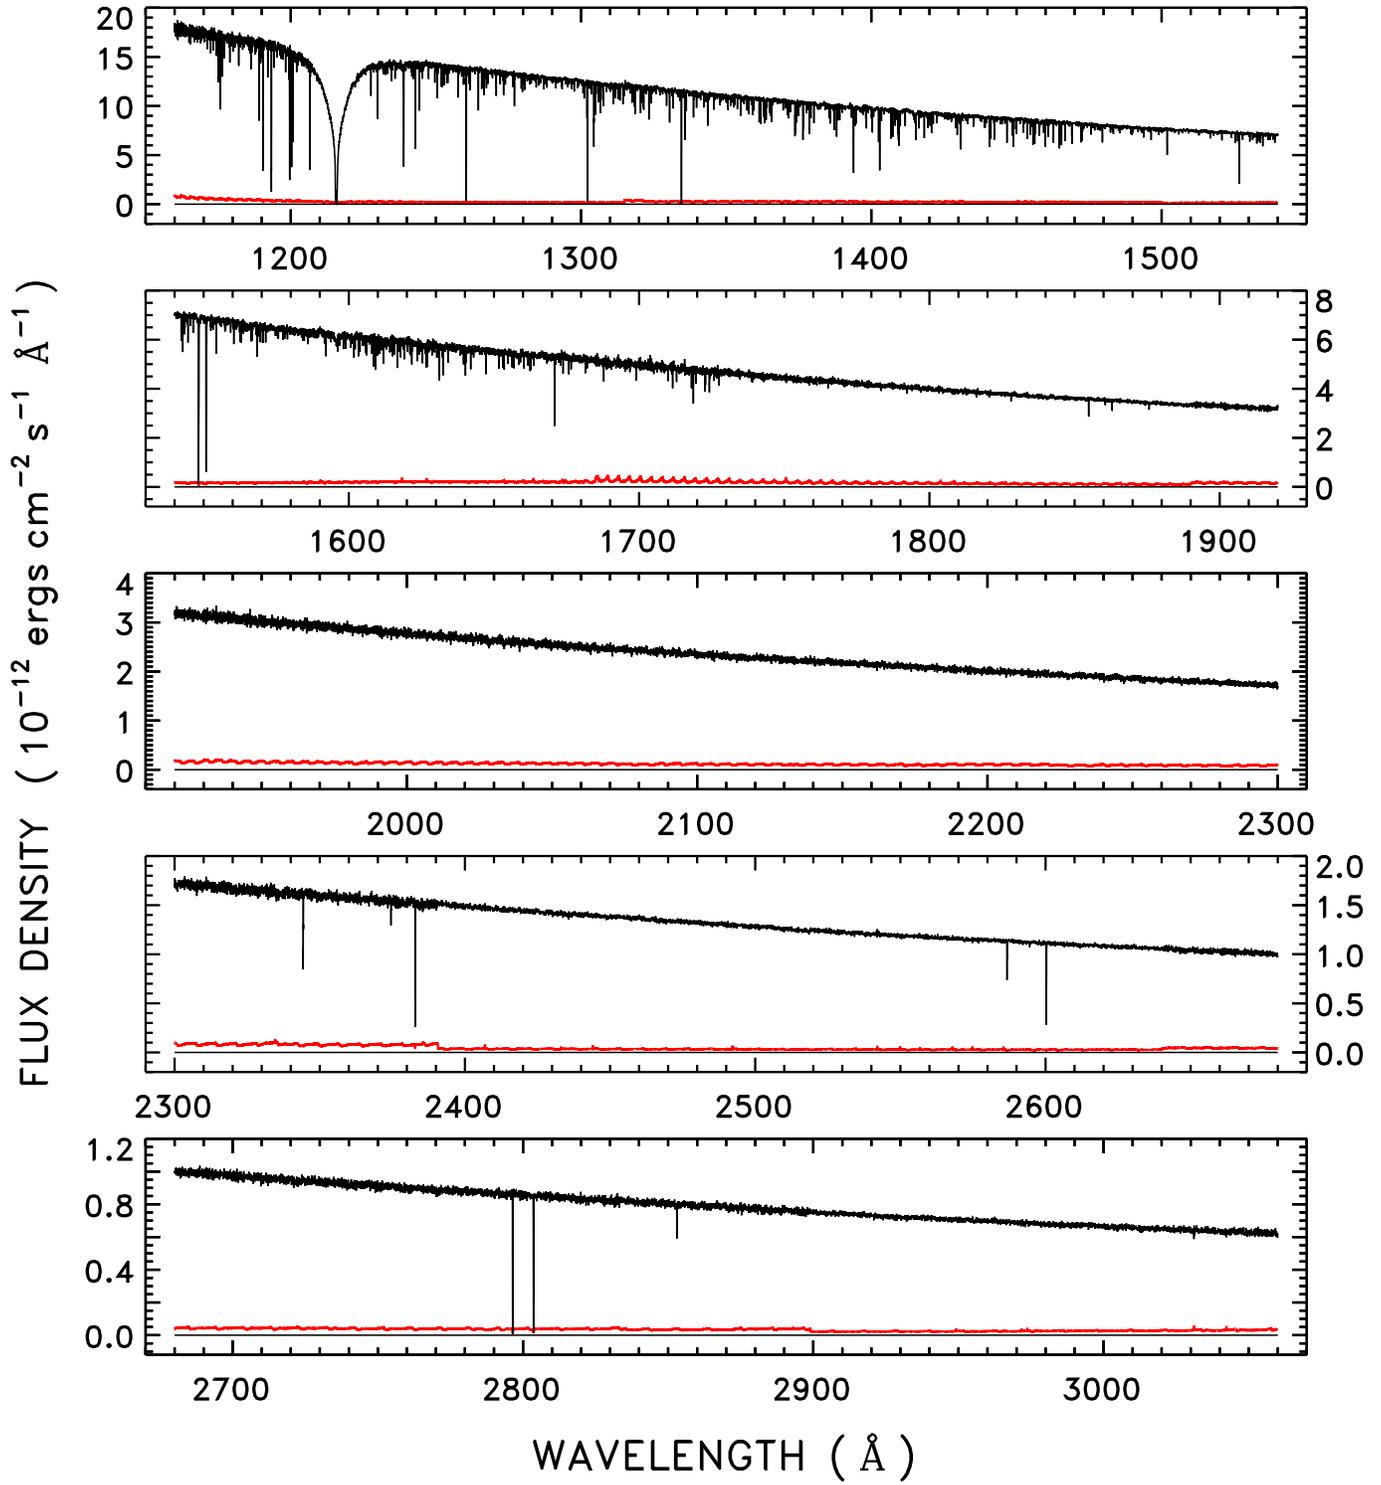


Fig. 8d.—