



STScI | SPACE TELESCOPE
SCIENCE INSTITUTE

When there is a discrepancy between the information in this technical report and information in JDox, assume JDox is correct.

JWST TECHNICAL REPORT

Title: A Revised Implementation of the ATOCA Algorithm for Spectral Extraction of NIRISS/SOSS Data	Doc: JWST-STScI-009313, SM-12 Date: February 5, 2026 Rev: -
Authors: Timothy D. Brandt Edward M. Molter Aarynn L. Carter :	Release Date: 24 March 2026

Abstract

This Technical Report presents a revised implementation of the ATOCA algorithm for spectral extraction of NIRISS data taken in the Single Object Slitless Spectroscopy (SOSS) mode. ATOCA involves forward-modeling the dispersed image on the detector in order to extract spectral orders while minimizing cross-order contamination. Our implementation represents a factor of $\approx 15\text{--}30$ time savings over the implementation in JWST builds 1.20.X and earlier, imputes missing data to avoid biases due to cosmic ray hits or other data loss, and constructs a representative integration to optimize the parameters of the ATOCA reduction. This revised implementation ultimately has very little impact on the extracted spectra aside from its treatment of missing data. It has been implemented in Version 1.21.0 of the JWST pipeline.

1 Introduction

The Single Object Slitless Spectroscopy [SOSS, Albert et al., 2023] mode of the Near-Infrared Imager and Slitless Spectrograph [NIRISS, Doyon et al., 2012, 2023] is a time-series observation mode primarily designed for the spectroscopic characterization of transiting exoplanet systems. Utilizing the GR700XD grism, the SOSS mode cross-disperses the light of a given target across multiple spectral orders from $0.6\text{--}2.8\ \mu\text{m}$ at a resolving power, $R \approx 700$. Crucially, these spectral orders are not spatially isolated on the detector, causing notable order contamination. This is most apparent at the longest wavelengths of Order 1 and Order 2, where the spectral orders overlap, but is also present to some degree at all wavelengths due to the broad cross-dispersion wings. With

individual pixels experiencing illumination from different spectral orders at different wavelengths, there is no single wavelength solution, and an accurate spectroscopic extraction cannot be obtained using standard techniques.

The ATOCA algorithm [Darveau-Bernier et al., 2022] enables extraction of the overlapping spectral orders. The basic procedure is to forward-model the two-dimensional image on the detector using a detailed characterization of the instrument and optics. The forward model can then be derived for all spectral orders except for the desired one and subtracted from the image. This step provides a relatively clean image of a single spectral order. That order can then be extracted using one-dimensional approaches, e.g., simple summation cross-dispersion (box extraction). By applying this approach to each of the three spectral orders in turn, all three may be extracted.

ATOCA forward-models the detector image with a spectrum of much finer sampling than the detector resolution. If it were using this approach to infer the spectrum, it would result in strong wavelength-to-wavelength correlations. The use of an oversampled spectrum solely to model and remove contaminating orders is not a problem in this respect. However, it introduces so many free parameters that the problem is no longer sufficiently overconstrained and ATOCA’s least squares approach can become numerically unstable. ATOCA solves this problem by Tikhonov regularization, with an appropriate regularization factor to be determined from the data. ATOCA uses a threshold on the derivative of χ^2 of the two-dimensional model with respect to the regularization factor.

As implemented in the JWST pipeline [Bushouse et al., 2025] version 1.20.0, the ATOCA algorithm is computationally expensive, often requiring an hour or more to process 100-200 integrations. This Technical Report summarizes small changes to the algorithm that result in a large performance improvement. After these changes, the extracted spectra are very similar but not identical to those derived using the legacy pipeline. We demonstrate the reasons for the differences, arguing that the new results reduce systematic errors due to the handling of missing data and also reduce biases from noise correlations with signal.

We organize the Report as follows. Section 2 briefly summarizes the ATOCA algorithm. Section 3 presents the core of our revised approach: treating the per-pixel uncertainties as being constant across integrations. Section 4 discusses the regularization that the problem requires, and Section 5 discusses the oversampled wavelength grid. Section 6 quantifies and discusses the differences between the legacy and revised implementations and Section 7 summarizes the performance improvements. We conclude with Section 8.

2 The Least Squares Approach of ATOCA

The ATOCA algorithm begins by constructing a forward model of the two-dimensional image on the detector. This model image is a linear function of the spectrum at each wavelength. We denote this spectrum as a vector \mathbf{f} and the linear relationship between spectrum and detector image by the matrix \mathbf{A} ; the dimensions of \mathbf{A} are $n_{\text{pixels}} \times n_{\text{wavelengths}}$. We further denote the flattened detector image by the vector \mathbf{b} and the flattened uncertainties by the vector $\boldsymbol{\sigma}$. The χ^2 value for the model is then

$$\chi^2 = \left| \frac{\mathbf{A}\mathbf{f} - \mathbf{b}}{\boldsymbol{\sigma}} \right|^2 \quad (1)$$

where the absolute value symbols represent the L2 norm and the fraction represents element-wise division.

Equation 1 is not regularized. Minimizing it numerically may become unstable with a sufficiently oversampled wavelength array due to a poor condition number of \mathbf{A} . To overcome this problem, ATOCA uses Tikhonov regularization with an $n_{\text{wavelengths}} \times n_{\text{wavelengths}}$ regularization matrix $\mathbf{\Gamma}$ and a factor α ; the χ^2 to be minimized is then

$$\chi^2 = \left| \frac{\mathbf{A}\mathbf{f} - \mathbf{b}}{\boldsymbol{\sigma}} \right|^2 + \alpha |\mathbf{\Gamma}\mathbf{f}|^2. \quad (2)$$

Following the standard approach for the solution of a linear system by the normal equations [Press et al., 1992], we next define \mathbf{b}' as the element-wise ratio of \mathbf{b} and $\boldsymbol{\sigma}$ and \mathbf{A}' as the row-wise ratio of \mathbf{A} and $\boldsymbol{\sigma}$. Equation (2) can then be written

$$\chi^2 = |\mathbf{A}'\mathbf{f} - \mathbf{b}'|^2 + \alpha |\mathbf{\Gamma}\mathbf{f}|^2. \quad (3)$$

To minimize Equation (2), we can differentiate and set the derivative to zero. This yields

$$\left(\mathbf{A}'^T \mathbf{A}' + \alpha \mathbf{\Gamma}^T \mathbf{\Gamma} \right) \mathbf{f} = \mathbf{A}'^T \mathbf{b}'. \quad (4)$$

Equation (4) can be solved for \mathbf{f} with standard linear algebra, using either an inversion or a decomposition of the matrix $\mathbf{A}'^T \mathbf{A}' + \alpha \mathbf{\Gamma}^T \mathbf{\Gamma}$. In ATOCA, the matrix \mathbf{A} represents a detector model and is computed once; it is the same for all integrations. Similarly, the matrix $\mathbf{\Gamma}$ encodes the regularization scheme and only needs to be computed once. Unfortunately, the matrix \mathbf{A}' includes the uncertainties of every pixel. As a result, it is computed separately for each integration. The subsequent matrix multiplications and inversion or decomposition must also be done separately for each integration, at significant computational expense.

With the machinery summarized in this section, the best-fit \mathbf{f} may be computed integration by integration, and from this, the best-fit scene $\mathbf{A}\mathbf{f}$. To extract an order, a separate model matrix \mathbf{A}_i may be constructed excluding the desired order i . From this a decontaminated image of the order may be derived as $\mathbf{b} - \mathbf{A}_i\mathbf{f}$.

3 Fixed Uncertainties

The most expensive parts of the ATOCA calculation are the construction of the model matrices \mathbf{A} and \mathbf{A}_i for each order i , their multiplication, and the solution of Equation (4) for \mathbf{f} by either matrix inversion or matrix decomposition. The model matrices only need to be computed once as they are the same for all integrations. However, the matrix in Equation (4),

$$\mathbf{M} = \mathbf{A}'^T \mathbf{A}' + \alpha \mathbf{\Gamma}^T \mathbf{\Gamma}, \quad (5)$$

differs for all integrations because the pixel uncertainties $\boldsymbol{\sigma}$ change from integration to integration, and $\mathbf{A}' = \mathbf{A}/\boldsymbol{\sigma}$. In this section we develop an approach that holds the uncertainties constant across integrations. There are a number of motivations for doing so:

1. NIRISS SOSS time series show only modest variations in signal from one integration to the next; the uncertainties change little.
2. By adopting fixed uncertainties, the matrix \mathbf{M} is the same for all integrations and may be inverted or decomposed once.

-
3. Photon noise is typically derived from the best-fit count rate. This introduces a covariance that biases the best-fit model image low, which can slightly underestimate contamination.

This section will focus on the first two points; we will defer discussion of the third point to Section 6.

For the constant, integration-independent uncertainty, we adopt the median of σ across integrations excluding invalid data (represented by NaNs). We then compute the matrices \mathbf{A}' , \mathbf{M} , and \mathbf{M}^{-1} using this uncertainty. The matrix \mathbf{b}' will still be computed integration-by-integration, but will always use the same uncertainty regardless of the pixel values \mathbf{b} . The solution for the oversampled spectrum \mathbf{f} is then given by the product of a matrix and a vector, which is several orders of magnitude less expensive computationally than the inversion or decomposition of the large matrix \mathbf{M} .

Our use of a fixed uncertainty prevents us from masking pixels in individual integrations. While this may seem to be a drawback, masking pixels in individual integrations is not possible in aperture-based spectral extraction either (this is the current default extraction method of the JWST pipeline). As of JWST Pipeline Version 1.20.2, invalid pixels (and their uncertainties) are set to zero for the purposes of spectral extraction. We instead impute these pixels' values. We first replace invalid data with the median across all integrations. We then overwrite replaced pixel values with the median of those pixels across a maximum of nine neighboring integrations. This preserves the temporal localization of the imputation. To flag replaced data, we inflate their uncertainties by a factor of ten relative to the median uncertainty across integrations. The extracted spectrum thus preserves a record of the use of imputed data. We make no attempt to impute data that are missing across *all* integrations due to, e.g., a static bad pixel. Future work could use the ATOCA-derived model of the detector image to do so.

4 The Optimal Tikhonov Factor

Our updated ATOCA implementation runs much faster than the previous implementation due to the fact that large matrices only need to be computed, inverted, or decomposed once across all integrations. However, the χ^2 optimization of Equation (2) still requires a regularization factor α . The previous implementation of ATOCA computed the best Tikhonov factor α from the first integration alone. This process proceeded in two stages. First, χ^2 was computed according to Equation (2) using a grid of 10 trial Tikhonov factors across ten orders of magnitude. Several different optimizations were then applied, and the smallest Tikhonov factor among them was adopted. Second, 20 trial Tikhonov factors were used across four orders of magnitude centered on the initial best-fit value. In this stage, a threshold was applied to $d\chi^2/d\log\alpha$ to determine the appropriate factor.

We revisit this algorithm for several reasons. We would like to avoid the use of different criteria in the first iteration and subsequent iterations, and if possible, we would like to improve the computational efficiency. It is also possible that the first integration may not be representative of subsequent integrations, and we would like our approach to be robust to this case.

We proceed by constructing a representative integration for use in determining the Tikhonov factor. For this purpose we use the mean of all integrations, ignoring invalid data in our averaging. We choose the mean rather than the median so that noise will average down even in the presence of astrophysical variability; the median would not have this property. Given many integrations, this approach results in an image with much less noise than any individual integration. While we

could scale the noise down appropriately to find a good Tikhonov factor, this would change the typical noise relative to underlying systematics that do not average down. In practice it proves to be an unreliable way to generate good Tikhonov factors. Instead, we use a random number generator to add Gaussian noise to our typical (mean) integration. We scale this Gaussian noise so that its standard deviation matches the median uncertainty that we use across all integrations, adjusted for the fact that our mean integration retains a fraction $1/\sqrt{n_{\text{integrations}}}$ of the original noise. To make the pipeline repeatable in its output, we initialize the random number generator to a fixed seed.

We also change how the best Tikhonov factor is calculated. We retain the initial coarse grid of trial factors across ten orders of magnitude in α , but then construct a modified Akima spline [Akima, 1970] through these points as a function of $\log \alpha$. We next compute the derivative of the spline and generate a large number of points. We choose the first point at which the derivative exceeds the adopted threshold as our initial guess for the optimal Tikhonov factor. We next add two additional trial factors, one on either side of the coarse grid closest to this initial best fit factor, and recompute the spline and its derivative. If the spline derivative never crossed the threshold within the initial domain, this refinement step extends the domain. We repeat this refinement procedure three times to improve the resolution on the Tikhonov factor by a typical factor of eight for an additional cost of six trial computations rather than the previous twenty. Because the cost of this step is dominated by the cost of evaluating χ^2 for each Tikhonov factor, our approach results in a factor of nearly 2 in runtime savings. It also retains the same criterion—a threshold on $d\chi^2/d\log \alpha$ —for inferring the best regularization factor throughout the process.

5 The Adaptive Wavelength Grid

The final change we make is to the oversampled wavelength grid. The ATOCA algorithm constructs an adaptive wavelength grid intended to enable a nearly perfect model of the detector. This wavelength grid is limited to a maximum of 20,000 points, but there is no limit to its resolution. This permits the adaptive algorithm to construct a wavelength grid oversampling the pixels by many orders of magnitude. This is undesirable given the use of sparse matrices to represent the mapping between spectrum and detector image, and generally unproductive given the regularization of the problem. A very highly oversampled wavelength grid will produce a locally dense matrix in the linear mapping between wavelength and pixel.

We limit this behavior by imposing a minimum separation of $10^{-5} \mu\text{m}$ between elements of the wavelength grid. This is still very oversampled, at ≈ 0.01 pixels, but it limits the denseness of the resulting matrices and improves performance. It avoids a handful of pixels where the legacy implementation would compute wavelength resolutions as fine as $\approx 10^{-7} \mu\text{m}$, corresponding to $\sim 10^{-4}$ pixels.

6 Quantitative Differences in the Outputs

The previous sections detailed a number of small changes to our implementation of the ATOCA algorithm. These will result in differences in the extracted spectra. We divide the changes that could cause differences in the end results into four categories:

1. Differences due to the imputation of missing pixel data;

-
2. Differences due to the use of a single set of uncertainties for all integrations;
 3. Differences due to a slightly different Tikhonov factor; and
 4. Differences due to limiting the wavelength grid’s oversampling.

In this section we will quantify and study each of them. To assess the impact of our changes, we focus on the dataset `jw02734002001_04101_00001-seg001`, the first segment of the same dataset used in the NIRISS SOSS demonstration notebook¹. This example dataset is from Early Release Observation (ERO) Program 2734 (PI: K. Pontoppidan). The program consists of time series observations (TSO) of confirmed exoplanets HAT-P-18b [Fu et al., 2022, Fournier-Tondreau et al., 2024] and WASP-96b [Radica et al., 2023, Wang et al., 2025], intended to demonstrate the power and precision of the JWST TSO modes. We will begin by holding both the wavelength grid and the Tikhonov factor fixed, and then proceed to investigate the impact on the Tikhonov factor itself and downstream effects on the extracted spectra. Finally, we show the impact of our change to the model wavelength grid.

6.1 The Extracted Spectra and Imputed Data

We first fix both the wavelength grid and Tikhonov factors and investigate quantitative changes in the extracted spectra. We compute the mean and median of the difference between the extracted spectra across the 100 integrations in Segment 1 of our sample data. The mean difference will show the effects of lost pixels in individual integrations, which will systematically bias the spectra low when they are set to zero. The median difference will show systematics from the use of a fixed set of uncertainties.

Figure 1 shows the mean and median differences in the extracted spectra across integrations. The differences are multiplied by a factor of 1000 to be visible on the same axes as the spectra themselves. The mean differences are systematically negative in Orders 1 and 3, reflecting the impact of pixels with missing data that were set to zero. The median difference in both cases is many orders of magnitude smaller than the measured spectra, and many orders of magnitude smaller than the noise.

Order 2 in Figure 1 presents a more complex story. The blue half of the Order 2 spectrum shows the same behavior as Orders 1 and 3, with the mean difference being systematically negative due to the loss of missing data. The red half, where there is significant overlap between Orders 1 and 2, is the only place where the median difference is significantly nonzero and where the mean has significant positive fluctuations. In the following subsection we discuss the reason for this.

6.2 Systematic Differences in Order 2 due to Fixed Uncertainties

The red section of Order 2 overlaps part of Order 1 on the NIRISS detector. The extracted spectrum of Order 2 is very sensitive to the removal of contamination from Order 1. Figure 1 shows that the extracted spectrum in the new reduction is smaller than the spectrum in the old reduction in this spectral region. This may be understood as a consequence of the modeled contamination being larger in the new reduction. In this subsection we discuss the reason for the change.

¹<https://github.com/spacetelescope/jwst-pipeline-notebooks/>

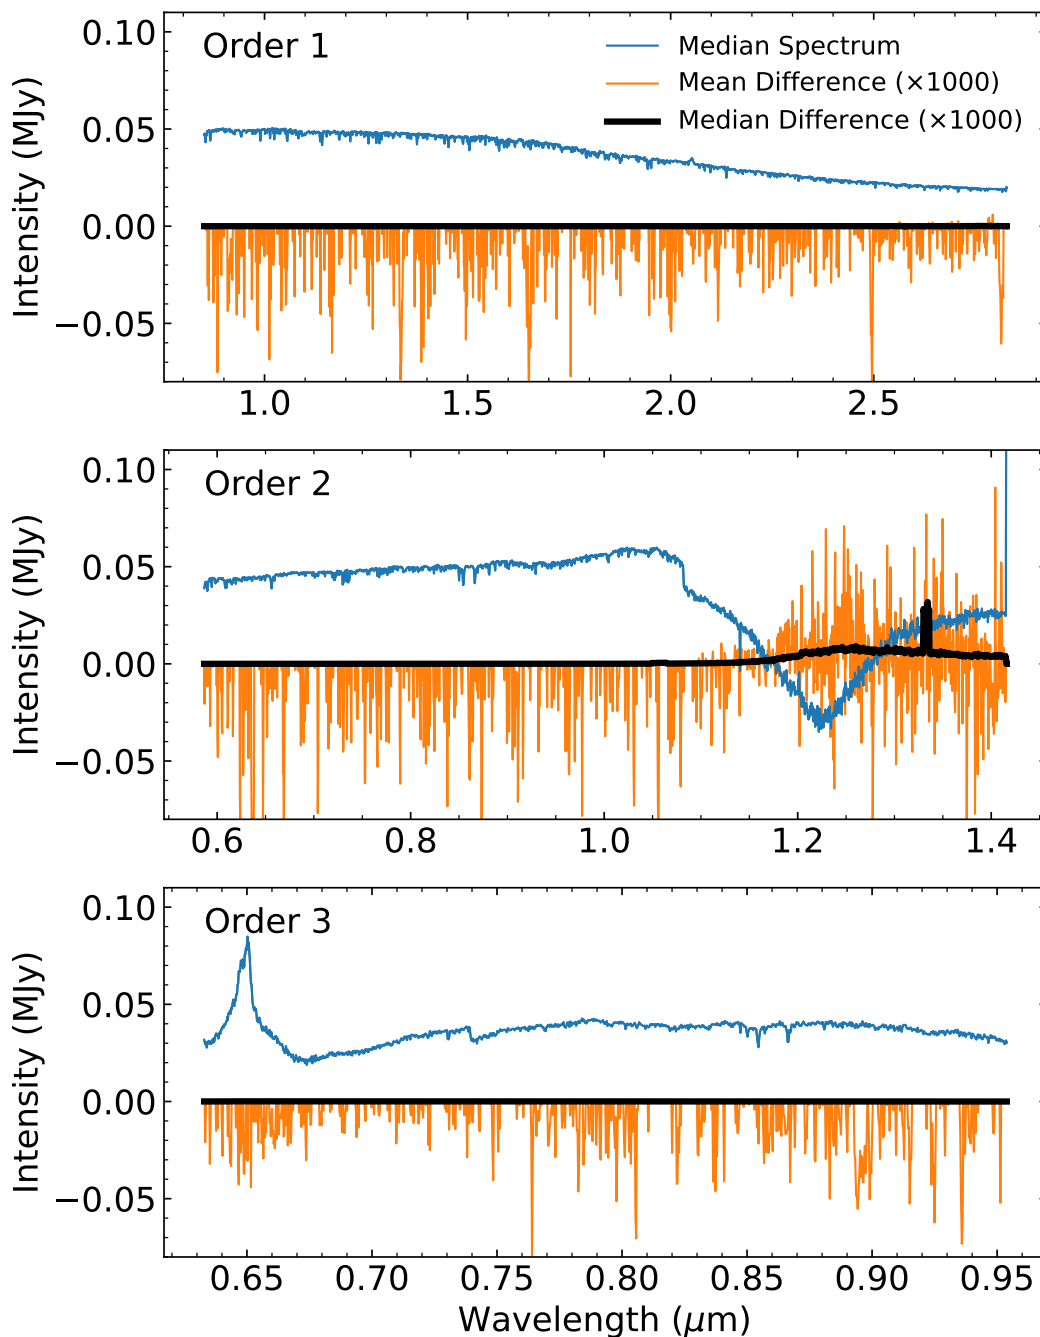


Figure 1: Mean and median across integrations of the difference between spectra extracted with the previous and updated ATOCA implementations. Negative values mean that the previous implementation gives a smaller value. The measured spectrum is shown in blue; the differences are multiplied by 1000 to be visible on the same axis limits. The mean difference shows the effects of imputed data for lost pixels and rounding errors that affect small numbers of pixels more or less randomly. The median difference shows systematic effects. These are negligible apart from Order 2, which we discuss further in Section 6.2 .

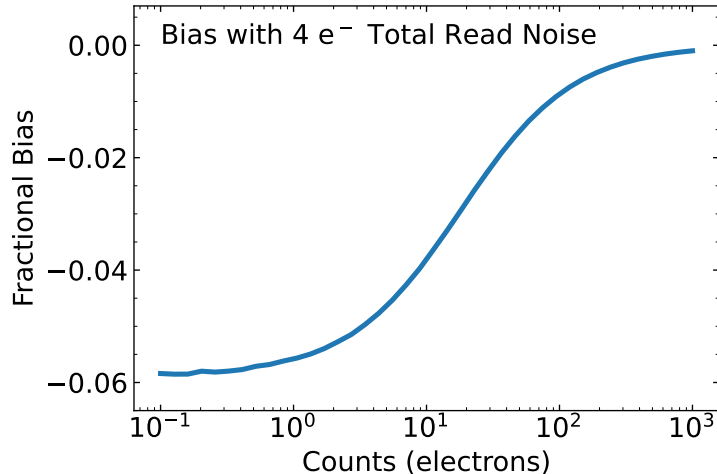


Figure 2: Bias in the estimator of Equation (6). This bias results in an underestimate of the signal when using a χ^2 fit to data for which photon noise was estimated from the data themselves.

When assigning the uncertainties associated with a pixel, the read noise may be independently measured but the photon noise depends on the counts that pixel received. The expected number of counts will not be independently known (otherwise there would be no need to measure it with JWST). Instead, the photon noise is estimated from the data themselves. If, by chance, a pixel happens to receive more photons than expected, it will have a larger photon noise uncertainty assigned to it. The result of this is a covariance between the realization of the noise and the assigned value of the uncertainty.

The covariance between the realized noise and the inferred uncertainty introduces a bias in a maximum likelihood estimate of the flux, like χ^2 . To see this, we set the read noise variance to be σ_r^2 , assign $\sigma_i^2 = n_i$ to be the number of measured photoelectrons in integration i , and assume that the mean number of photoelectrons $\langle n_i \rangle = N$ is unbiased. A maximum likelihood estimator gives

$$\tilde{n} = \left(\sum_i \frac{n_i}{n_i + \sigma_r^2} \right) \left(\sum_i \frac{1}{n_i + \sigma_r^2} \right)^{-1}. \quad (6)$$

In the limit of $\sigma_r^2 \rightarrow 0$, a single measurement of $n_i = 0$ results in $\tilde{n} = 0$.

Figure 2 shows the empirical bias in the estimator given by Equation (6) for a read noise of 4 electrons. This read noise is the effective noise from fitting all reads up-the-ramp, and can be smaller than the noise associated with an individual read. The bias is always negative: the recovered signal is always, on average, an underestimate of the true signal. This results in subtracting less contamination from the red half of Order 2.

By using the median uncertainty across integrations, the correlation between the estimated photon noise and the signal is much reduced. This results in less bias in the revised implementation of ATOCA, a larger inferred contamination of the red half of Order 2, and a slightly smaller extracted intensity. The difference is small but the results from the new extraction are to be preferred.

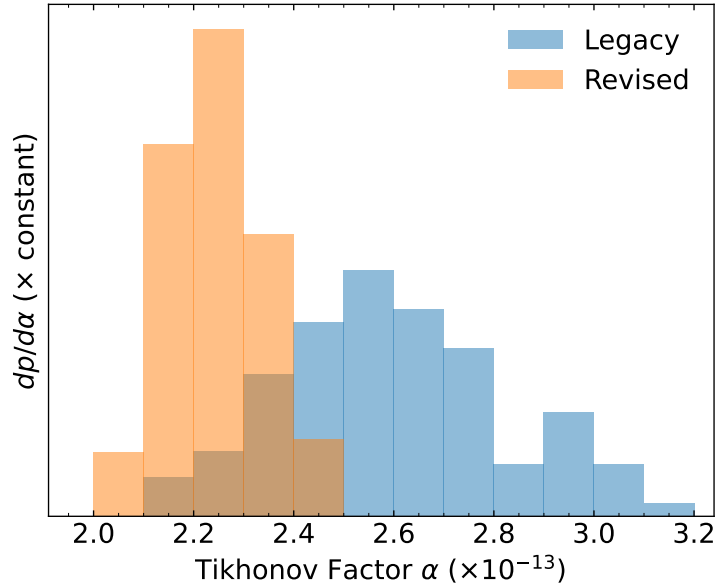


Figure 3: Distribution of best-fit Tikhonov factors for Orders 1+2 under the legacy ATOCA implementation (blue) and the revised one (orange). The previous implementation shows a larger spread depending on which of 100 integrations is used as the reference. The revised version shows a smaller spread and a small offset, a factor of ≈ 4 smaller than the refined grid resolution in the legacy implementation.

6.3 The Tikhonov Factor

As described in the previous section, we now compute the Tikhonov factor using a reference integration, taken to be the average of all integrations with Gaussian noise added to ensure that the median uncertainty is representative of its error. To assess the impact of this change, we compute the distribution of the best-fit Tikhonov factor across all 100 integrations of Segment 1 in the sample dataset. We then compute the same distribution for our reference integration with 100 realizations of Gaussian noise. In all cases, we fix the wavelength grid to match that used in the previous version of the ATOCA implementation.

Figure 3 shows the results for the Tikhonov factor for Order 1, the only Tikhonov factor meaningfully used in the ATOCA algorithm. The distributions of the legacy and revised implementation are shown in shaded blue and orange histograms, respectively. The two distributions overlap, but they do show slight differences. The range of factors with the updated approach is smaller because the differences arise solely from different realizations of Gaussian noise. Systematics and mis-estimated errors are absent by construction. The difference in the centers of the two distributions is small, $\approx 15\%$. With 20 points spanning four orders of magnitude, the Tikhonov factor grid spacing in the legacy version of ATOCA was a factor of $20^{4/19} \approx 1.62$, or ≈ 4 times the difference in distribution centers. Such a difference may be due to the change from a finite difference calculation of derivatives to a spline fit, or it could be due to slight differences in the properties of real and idealized noise.

The grid spacing in the legacy ATOCA implementation is larger than the width of the distribution of best-fit factors shown in Figure 3 even when performed integration-by-integration. The best-fit factor also uses a somewhat arbitrary threshold on $d\chi^2/d\log\alpha$ to determine the best

Tikhonov factor. For these reasons, it is very unlikely that the results would improve from separate tunings of the Tikhonov factor integration-by-integration.

Finally, we verify the insensitivity of our results to differences in the Tikhonov factor like those shown in Figure 3. The first integration in our sequence has a best-fit Tikhonov factor of 2.97×10^{-13} for Orders 1+2, larger than 90% of all integrations. With our default random number generator seed, we obtain a best-fit Tikhonov factor of 2.23×10^{-13} for our revised ATOCA implementation. The difference between these two values is \approx twice the difference between the centers of the distributions. We then reprocess the entire sequence of integrations with these Tikhonov factors and analyze the results as in Figure 1.

Figure 4 shows our results. As in Figure 1, the mean difference mainly shows the impact of imputed data apart from the red half of Order 2. The median difference largely tells the same story as in Figure 1, but with larger excursions from zero at a few wavelengths in the red half of Order 2. These remain at a level $\approx 0.1\%$ of the peak value of this spectrum, and $\lesssim 1\%$ of the uncertainties. The typical impact will be smaller than this due to the lower average discrepancy between the two distributions shown in Figure 3, and it will be stochastic.

6.4 The Wavelength Grid

Finally, we check the effects of changing the minimum wavelength spacing on the results. We show the difference between extracted spectra with and without applying a limit on the wavelength resolution as discussed in Section 5. Figure 5 shows these differences in spectra extracted from the same file as that used for Figures 1 and 4. The differences are barely visible even when inflated by a factor of 1000, and are negligible for the purposes of science.

7 Implementation and Performance

Many of the changes summarized in this Technical Report are designed to offer performance improvements over the legacy ATOCA implementation. In this section we review the performance improvements on JWST Programs 1091 and 2734 in the `extract_1d` step of the `calwebb_spec2` pipeline. Program 1091, PI Andre Martel, was a calibration program using the A-type star BD+60 1753.

The primary performance improvement comes from the re-use of precomputed matrices. There is overhead associated with these computations, about a minute on a 2023 Macbook Pro. However, the cost of processing an individual integration on the same machine drops from about 20 seconds in JWST build 1.20.0 to about 500 ms with the new implementation. The computational cost of the initial Tikhonov factor optimization also drops by a factor of 2-20 depending on the details of the wavelength grid as discussed in the previous section. The net result for the first segment of the Program 2734 dataset is a decrease in runtime from about 30 minutes to slightly more than 2 minutes.

A few other pipeline changes modestly improve runtime between pipeline versions 1.20.0 and 1.21.0. The model spectral trace positions and their corresponding wavelength values based on reference information [i.e., PASTASSOSS Baines et al., 2023a,b] and the spectral resolution kernel $\kappa(\bar{\lambda}, \lambda)$ are both identical integration-by-integration. The legacy software repeats these calculations every integration; they are now pre-computed alongside the other pre-computed matrices. The reconstruction of decontaminated spectra from trial Tikhonov factors, which is not necessary

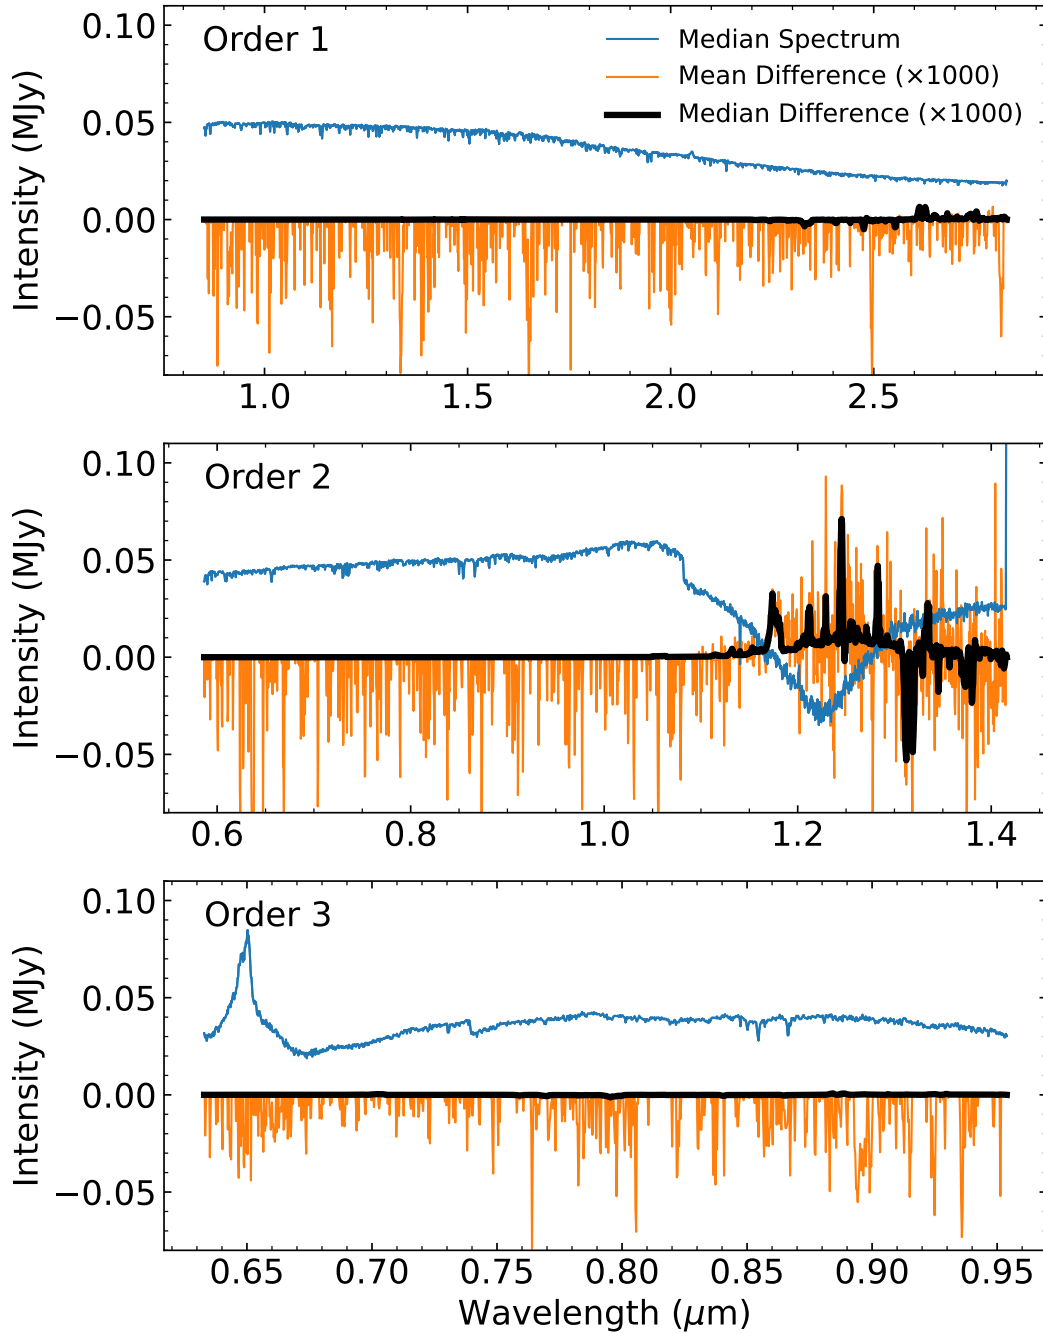


Figure 4: Similar to Figure 1, but without fixing the Tikhonov factors to match between the two ATOCA reductions. The Tikhonov factor for Orders 1+2 derived from the legacy approach is 2.97×10^{-13} , about twice as discrepant from the 2.23×10^{-13} of our revised approach as the factor would be from an average integration (c.f. Figure 3). Despite this, median differences are very small, with the largest deviations appearing at a few wavelengths in the part of Order 2 that overlaps Order 1.




Figure 5: Differences in the extracted spectra, similar to Figure 1, but with and without imposing a resolution limit on the wavelength grid as discussed in Section 5. The mean and median differences are nearly identical; only the median is shown. The differences between the extracted spectra are negligible.

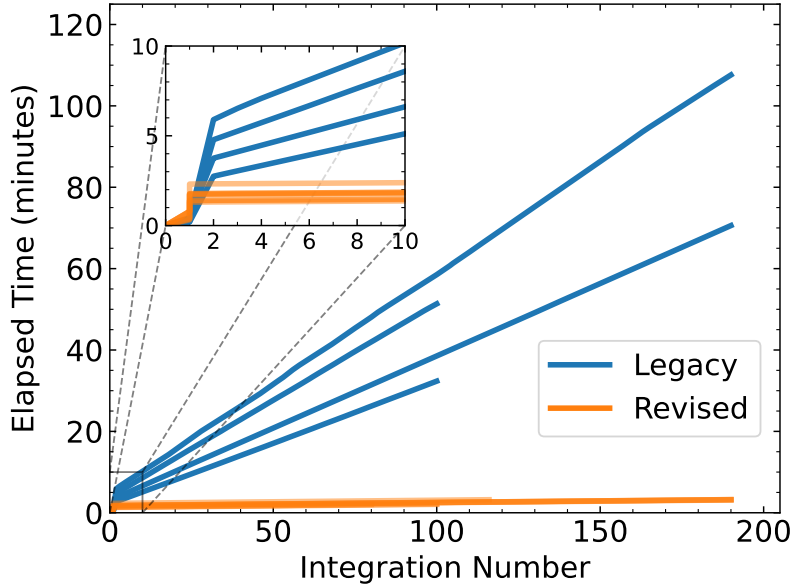


Figure 6: Performance of the legacy (blue) vs. revised (orange) ATOCA implementations on the JWST pipeline step `calwebb_spec2` for segments from JWST programs 1091 and 2734. The revised implementation has a ≈ 1 – 2 minute initial cost but then takes just ≈ 0.5 seconds per integration. Speedups for datasets with many integrations can be as high as a factor of ≈ 30 .

to determine the best factor, is now only carried out if intermediate diagnostic products are requested. While none of these changes affect the numerical results of the step, they contribute to the performance improvements discussed here.

Figure 6 shows the performance of our new implementation on several segments from JWST programs 1091 and 2734 with up to 190 integrations per file. The performance is measured for spectral extraction, `calwebb_spec2`. The revised implementation has a significantly smaller overhead from Tikhonov optimization. Combined with matrix construction and multiplication and matrix inversion, ATOCA overheads now total ≈ 1 – 2 minutes on the 2023 Macbook Pro used for testing. After the initial Tikhonov optimization and matrix calculations, the subsequent cost per integration is ≈ 0.5 seconds.

The legacy algorithm has significantly variable overheads for the Tikhonov optimization and a range of time costs per integration. This is due to the use of sparse matrix solvers throughout whose performance depends on the details of the matrices, and in particular, their local denseness due to occasionally extreme oversampling in the wavelength grid. This apparently stochastic variation is the cause of the differing slopes of the blue lines in Figure 6.

Our revised ATOCA implementation has a cost that is relatively insensitive to dataset or to the number of integrations. For the five segments of Program 1091 and the three segments of Program 2734, with 68 to 190 integrations, run times of `calwebb_spec2` range from 2.1 minutes to 3.3 minutes. This represents a runtime improvement factor ranging from 15 to 30 over the legacy ATOCA implementation in JWST pipeline version 1.20.0.

8 Conclusions

This Technical Report has summarized a new implementation of the ATOCA algorithm to extract one-dimensional spectra from NIRISS/SOSS data. The new implementation offers substantial performance improvements while changing the data little. Some of the changes, e.g. from a slightly different regularization factor, are stochastic and unlikely to affect data quality. Others, such as the imputation of missing data rather than assuming those data to be zero, are likely to improve data quality. The revised approach also largely avoids a slight bias in the modeled spectrum due to covariance between the realization of photon noise and the inferred value of photon noise. All quantitative changes in the extracted spectra, other than from the effects of invalid data, are $\lesssim 1\%$ of the uncertainty.

The changes summarized here have been implemented in the JWST pipeline as of Version 1.21.0. Future work could further develop data imputation for static bad pixels, e.g., using the model derived by the ATOCA algorithm.

References

- Loïc Albert, David Lafrenière, Doyon René, Étienne Artigau, Kevin Volk, Paul Goudfrooij, André R. Martel, Michael Radica, Jason Rowe, Néstor Espinoza, Arpita Roy, Joseph C. Filippazzo, Antoine Darveau-Bernier, Geert Jan Talens, Anand Sivaramakrishnan, Chris J. Willott, Alexander W. Fullerton, Stephanie LaMassa, John B. Hutchings, Neil Rowlands, M. Begoña Vila, Julia Zhou, David Aldridge, Michael Maszkiewicz, Mathilde Beaulieu, Neil J. Cook, Caroline Piaulet, Pierre-Alexis Roy, Pierrot Lamontagne, Kim Morel, William Frost, Salma Salhi, Louis-Philippe Coulombe, Björn Benneke, Ryan J. MacDonald, Doug Johnstone, Jake D. Turner, Marylou Fournier-Tondreau, Romain Allart, and Lisa Kaltenegger. The Near Infrared Imager and Slitless Spectrograph for the James Webb Space Telescope. III. Single Object Slitless Spectroscopy. *PASP*, 135(1049):075001, July 2023. doi: 10.1088/1538-3873/acd7a3.
- René Doyon, John B. Hutchings, Mathilde Beaulieu, Loïc Albert, David Lafrenière, Chris Willott, Driss Touahri, Neil Rowlands, Micheal Maszkiewicz, Alex W. Fullerton, Kevin Volk, André R. Martel, Pierre Chayer, Anand Sivaramakrishnan, Roberto Abraham, Laura Ferrarese, Ray Jayawardhana, Doug Johnstone, Michael Meyer, Judith L. Pipher, and Marcin Sawicki. The JWST Fine Guidance Sensor (FGS) and Near-Infrared Imager and Slitless Spectrograph (NIRISS). In Mark C. Clampin, Giovanni G. Fazio, Howard A. MacEwen, and Jacobus M. Oschmann, Jr., editors, *Space Telescopes and Instrumentation 2012: Optical, Infrared, and Millimeter Wave*, volume 8442 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 84422R, September 2012. doi: 10.1117/12.926578.
- René Doyon, Chris J. Willott, John B. Hutchings, Anand Sivaramakrishnan, Loïc Albert, David Lafrenière, Neil Rowlands, M. Begoña Vila, André R. Martel, Stephanie LaMassa, David Aldridge, Étienne Artigau, Peter Cameron, Pierre Chayer, Neil J. Cook, Rachel A. Cooper, Antoine Darveau-Bernier, Jean Dupuis, Colin Earnshaw, Néstor Espinoza, Joseph C. Filippazzo, Alexander W. Fullerton, Daniel Gaudreau, Roman Gawlik, Paul Goudfrooij, Craig Haley, Jens Kammerer, David Kendall, Scott D. Lambros, Luminita Ilinca Ignat, Michael Maszkiewicz, Ashley McColgan, Takahiro Morishita, Nathalie N.-Q. Ouellette, Camilla Pacifici, Natasha Philippi, Michael Radica, Swara Ravindranath, Jason Rowe, Arpita Roy, Niladri Roy, Karl

-
- Saad, Sangmo Tony Sohn, Geert Jan Talens, Driss Touahri, Deepashri Thatte, Joanna M. Taylor, Thomas Vandal, Kevin Volk, Michel Wander, Gerald Warner, Sheng-Hai Zheng, Julia Zhou, Roberto Abraham, Mathilde Beaulieu, Björn Benneke, Laura Ferrarese, Ray Jayawardhana, Doug Johnstone, Lisa Kaltenegger, Michael R. Meyer, Judy L. Pipher, Julien Rameau, Marcia Rieke, Salma Salhi, and Marcin Sawicki. The Near Infrared Imager and Slitless Spectrograph for the James Webb Space Telescope. I. Instrument Overview and In-flight Performance. *PASP*, 135(1051):098001, September 2023. doi: 10.1088/1538-3873/acd41b.
- Antoine Darveau-Bernier, Loïc Albert, Geert Jan Talens, David Lafrenière, Michael Radica, René Doyon, Neil J. Cook, Jason F. Rowe, Romain Allart, Étienne Artigau, Björn Benneke, Nicolas Cowan, Lisa Dang, Néstor Espinoza, Doug Johnstone, Lisa Kaltenegger, Olivia Lim, Tyler Pauly, Stefan Pelletier, Caroline Piaulet, Arpita Roy, Pierre-Alexis Roy, Jared Splinter, Jake Taylor, and Jake D. Turner. ATOCA: an Algorithm to Treat Order Contamination. Application to the NIRISS SOSS Mode. *PASP*, 134(1039):094502, September 2022. doi: 10.1088/1538-3873/ac8a77.
- Howard Bushouse, Jonathan Eisenhamer, Nadia Dencheva, James Davies, Perry Greenfield, Jane Morrison, Phil Hodge, Bernie Simon, David Grumm, Michael Droettboom, Edward Slavich, Megan Sosey, Tyler Pauly, Todd Miller, Robert Jedrzejewski, Warren Hack, David Davis, Steven Crawford, David Law, Karl Gordon, Michael Regan, Mihai Cara, Ken MacDonald, Larry Bradley, Clare Shanahan, William Jamieson, Mairan Teodoro, Thomas Williams, Maria Pena-Guerrero, Brett Graham, Edward Molter, Timothy Brandt, Christian Hayes, Rachel Cooper, Melanie Clarke, and Joseph Filippazzo. JWST Calibration Pipeline, July 2025.
- William H. Press, Saul A. Teukolsky, William T. Vetterling, and Brian P. Flannery. *Numerical recipes in FORTRAN. The art of scientific computing*. 1992.
- Hiroshi Akima. A new method of interpolation and smooth curve fitting based on local procedures. *Journal of the ACM*, 4(17):589–602, 1970.
- Guangwei Fu, Néstor Espinoza, David K. Sing, Joshua D. Lothringer, Leonardo A. Dos Santos, Zafar Rustamkulov, Drake Deming, Eliza M.-R. Kempton, Thaddeus D. Komacek, Heather A. Knutson, Loïc Albert, Klaus Pontoppidan, Kevin Volk, and Joseph Filippazzo. Water and an Escaping Helium Tail Detected in the Hazy and Methane-depleted Atmosphere of HAT-P-18b from JWST NIRISS/SOSS. *ApJL*, 940(2):L35, December 2022. doi: 10.3847/2041-8213/ac9977.
- Marylou Fournier-Tondreau, Ryan J. MacDonald, Michael Radica, David Lafrenière, Luis Welbanks, Caroline Piaulet, Louis-Philippe Coulombe, Romain Allart, Kim Morel, Étienne Artigau, Loïc Albert, Olivia Lim, René Doyon, Björn Benneke, Jason F. Rowe, Antoine Darveau-Bernier, Nicolas B. Cowan, Nikole K. Lewis, Neil J. Cook, Laura Flagg, Frédéric Genest, Stefan Pelletier, Doug Johnstone, Lisa Dang, Lisa Kaltenegger, Jake Taylor, and Jake D. Turner. Near-infrared transmission spectroscopy of HAT-P-18 b with NIRISS: Disentangling planetary and stellar features in the era of JWST. *MNRAS*, 528(2):3354–3377, February 2024. doi: 10.1093/mnras/stad3813.
- Michael Radica, Luis Welbanks, Néstor Espinoza, Jake Taylor, Louis-Philippe Coulombe, Adina D. Feinstein, Jayesh Goyal, Nicholas Scarsdale, Loïc Albert, Priyanka Baghel, Jacob L. Bean, Jasmina Blečić, David Lafrenière, Ryan J. MacDonald, Maria Zamyatina, Romain Allart1, Étienne Artigau, Natasha E. Batalha, Neil James Cook, Nicolas B. Cowan, Lisa Dang,

René Doyon, Marylou Fournier-Tondreau, Doug Johnstone, Michael R. Line, Sarah E. Moran, Sagnick Mukherjee, Stefan Pelletier, Pierre-Alexis Roy, Geert Jan Talens, Joseph Filippazzo, Klaus Pontoppidan, and Kevin Volk. Awesome SOSS: transmission spectroscopy of WASP-96b with NIRISS/SOSS. *MNRAS*, 524(1):835–856, September 2023. doi: 10.1093/mnras/stad1762.

Le-Chris Wang, Zafar Rustamkulov, David K. Sing, Joshua Lothringer, Patrick McCreery, Daniel Thorngren, and Munazza K. Alam. A Comprehensive Analysis of the Panchromatic Transmission Spectrum of the Hot-Saturn WASP-96 b: Nondetection of Haze, Possible Sodium Limb Asymmetry, Stellar Characterization, and Formation History. *arXiv e-prints*, art. arXiv:2511.16771, November 2025. doi: 10.48550/arXiv.2511.16771.

Tyler Baines, Néstor Espinoza, Joseph Filippazzo, and Kevin Volk. Characterization of the visit-to-visit Stability of the GR700XD Spectral Traces for NIRISS/SOSS Observations. Technical Report JWST-STScI-008448, STScI, 2023a.

Tyler Baines, Néstor Espinoza, Joseph Filippazzo, and Kevin Volk. Characterization of the visit-to-visit Stability of the GR700XD Wavelength Calibration for NIRISS/SOSS Observations. Technical Report JWST-STScI-008571, STScI, 2023b.