SOLAR-1/CCOR-2 Data Release
Provisional Data Quality
Read-Me for Data Users

Date: June 2, 2026

Version 1.0

Contents

List of Acronyms

API - Application Programming Interface
CCOR-2 – Compact Coronagraph-2
CME – Coronal Mass Ejection
DN – Data Number
ECEF – Earth-Centered Earth-Fixed
FITS – Flexible Image Transport System
FOV – Field of View
GPA – Ground Processing Algorithm
GOES – Geostationary Operational Environmental Satellite
HDU – Header Data Unit
HEE – Heliocentric Earth Ecliptic
HGA – High Gain Antenna
MSB – Mean Solar Brightness
NCEI – National Centers for Environmental Information
NOAA – National Oceanic and Atmospheric Administration
OSPO – Office of Satellite and Product Operations
PQF – Pixel Quality Flag
PVR – Product Validation Review
SOLAR-1 – Space weather Observations at L1 to Advance Readiness - 1
SWFO – Space Weather Follow-On
SUVI – Solar Ultraviolet Imager
SWPC – Space Weather Prediction Center
UTC – Coordinated Universal Time
WCS – World Coordinate System

List of Tables

List of Figures

- Figure 1: Schematic diagram to visualize the structure of the header data units (HDUs) inside a file that follows the FITS format.
- Figure 2: Two difference images computed between `level 1A` products without/with the coalignment applied in the left and right panel respectively. This particular example was chosen to straddle the daily HGA gimbal stepping period when we are usually guaranteed to see some image motion. The panel on the right shows the resulting improvement in the straylight artifacts achieved by shifting the images to coalign the centers of the occulters.
- Figure 3: Two difference images computed from `level 0B` data by subtracting sequential images, for a case where the lens motions are small (left) and large (right).
- Figure 4: Example `level 0B` images with corrupted or missing data. These images have `IMGBLK_Q==F` either `BADBLK_N` or `MISBLK_N` set to a non-zero value.
- Figure 5: Example of retrospective `level 1A` (left) and `level 2` (right) images. The `level 1A` image intensity is dominated by the stray light, while the `level 2` images have a monthly minimum (`ccor2-mm`, see Section 7.2) subtracted from the image which removes the bulk of the average stray light and F-corona. However, due to lens movements other types of stray light like the diffraction peaks and inner curved artifacts are not removed. Furthermore, bad pixels due to dust or bright pixels can become transformed into larger artifacts during the image co-alignment process and appear as static bright regions in the `level 2` images.
- Figure 6: Example level 0B images with a daily minimum background subtracted to highlight examples of transient features in the CCOR-2 data.
- Figure 7: Schematic overview of the SWPC and NCEI GPA processing and product levels. Inside each product shape we also indicate the units for each product level.

1 Introduction

The SWFO CCOR-2 Provisional-maturity Product Validation Review (PVR) was held on June 2, 2026. Based on the results of this review, CCOR-2 data were determined to be ready to support operational forecasts. The PVR panel recommended that the CCOR-2 data be promoted to Provisional Validation Maturity.

The NCEI will host both operational and retrospective data streams (FITS files) for all product levels here:

- https://www.ncei.noaa.gov/cloud-access/space-weather-portal/

For a more detailed explanation on using an application programing interface (API) for accessing data from the NESDIS cloud please read the information here.

- Please note that NCEI provides access to archived data files. This process can have latency as large as 72 hrs, but the daily data typically are available within a few hours of the end of a UTC day.

To access images for real-time space weather forecasts, the SWPC provides an operational data stream with `level 1A` (FITS) and `level 3` (JPEG) files in a rolling 60-day window:

- https://services.swpc.noaa.gov/products/ccor2/fits/
- https://services.swpc.noaa.gov/products/ccor2/jpeg/

2 Product Overview

CCOR-2 provides white-light images of the solar corona to monitor coronal mass ejections (CMEs) and other solar features. It is located on board the SOLAR-1 mission which is in orbit around the Lagrange 1 point. Operating in the visible wavelength range, the instrument employs an occulter to block direct sunlight and reveal the faint corona. For comprehensive technical details, refer to Thernisien et al. (Space Sci. Rev. 2026, accepted)[1].

Instrument Specifications:

- Bandpass full width at half maximum: 469 - 755 nm.
- Image Size: 2048 x 1920 pixels (`levels 0A, 0B, 1A, 2`), 1024 x 960 pixels (`level 3`). Due to instrumental vignetting, the corner regions and the central pylon are obscured in each image.
- Plate Scale: Varies from ~24.3 arcsec/pixel to ~23.7 arcsec/pixel (inner -> outer FOV).
- Field of View: Approximately 4.4 to 25 solar radii (1.2 - 6.5 deg in elongation). Diagonally the FOV extends out to ~27 solar radii (7.3 deg in elongation).
- Spatial Resolution: ~2 pixels (48 arcsec) across most of the image with a median point spread function across most of the image around ~1.5 pixels (32 arcsec).

Data Format:

---

[1] https://arxiv.org/abs/2508.13467

- All CCOR-2 products are stored in the hardware-independent Flexible Image Transport System (FITS) format[2]. Each FITS file is organized into a series of Header Data Units (HDUs), consisting of a primary unit followed by one or more extensions. Each HDU contains a text header (metadata) followed by the corresponding binary data (Figure 1). Different products will have different numbers of HDUs (see Section 7).
- Product metadata includes necessary navigation information to orient the Sun, adhering to World Coordinate System (WCS) conventions (e.g., Thompson et al., 2006).
- To assist with data interaction, we plan to release Python Jupyter notebooks. In the interim, users can refer to the NOAA Solar Ultraviolet Imager (SUVI) data examples for guidance on handling FITS files.

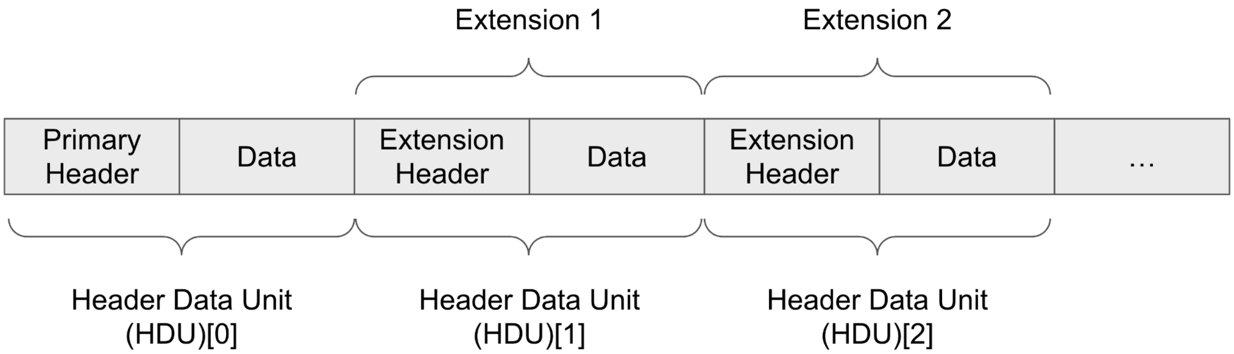

Figure 1: Schematic diagram to visualize the structure of the header data units (HDUs) inside a file that follows the FITS format.

3 Provisional Validation

Provisional validation of SWFO CCOR-2 data indicates:

- Validation activities of the data are ongoing and we encourage participation and feedback from the broader research community.
- Impactful anomalies have been identified in the data and are under analysis; solutions are in development.
- Incremental product improvements may still occur.
- Product performance has been validated through comparisons with independent coronal observations.
- Analysis establishes performance relative to the official performance baseline.
- Documentation includes some recommended remediation strategies for known anomalies but work is still ongoing.
- All testing procedures are fully documented.
- The product is ready for operational use and comprehensive calibration/validation activities.

4 Issues

---

[2] https://fits.gsfc.nasa.gov/fits_resources.html

Users of SOLAR-1/CCOR-2 data are responsible for inspecting the data and understanding the following known caveats, which are under active investigation. CCOR-2 is a space-based coronagraph and many of the caveats listed below represent standard considerations that a data user might encounter for imaging instruments, space-borne observatories, or coronagraphs generally.

4.1 Image to image motions

The most impactful artifact in the CCOR-2 dataset that users need to be aware of stems from a mechanical issue with the instrument that was not detected until the door opened on-orbit. It is believed that one or more of the lenses that focus the image for CCOR-2 are slightly loose within their mount. There are two main sources of mechanical impulses exciting the lenses:

- **HGA Gimbal Adjustments**: The High Gain Antenna (HGA) gimbals perform daily adjustments between 03:05 and 03:10 UTC, timed to avoid CCOR-2 exposure windows. However, this regularly results in image displacement when differencing data captured at 03:00 and 03:15 UTC. See Figure 2 for an example of the impact on image differencing.
- **Reaction Wheel Resonance**: Fluctuating reaction wheel speeds can induce lens resonance, causing motion even when the HGA gimbals are stationary. These disturbances are often most pronounced leading up to the weekly momentum adjustments (scheduled Mondays around 03:04 UTC). These maneuvers are occasionally synchronized with HGA stepping to minimize cumulative displacement; characterization and mitigation efforts remain ongoing.

This lens motion impacts the data in two main ways:

- First, the image moves slightly around the detector in a measurable but unpredictable way. The amplitude of this motion is typically less than 5 pixels, but larger excursions of up to 10 pixels have been observed. This effect is mitigated for `level 1A+` products in both the operational and retrospective data (see Figures 2 and 3), but users of `level 0B` data will see more noticeable artifacts.
    - During processing from `level 0B` to `level 1A`, the image shift and coalignment algorithm determines the occulter center by performing a grid search to maximize angle-integrated intensity in polar coordinates. This search is constrained to a ±7 pixel range along both the X and Y axes. Users of `level 1A+` data should be cautious when metadata keywords `SHIFT_X` or `SHIFT_Y` equal ±7 in the HDU[1] header; this may indicate that the applied correction has reached the limits of the search grid. A more robust center-determination algorithm is currently in development.
    - The instrument's vignetting function includes a static outer region (~50 pixels wide) where transmission drops to zero. Because this mask is fixed to the detector, rather than the moving image, correcting for lens motion creates a residual annular artifact at the edge of the field of view (FOV). This effect, visible as a dark/bright band in the right panel of Figure 2, renders the affected pixels photometrically unreliable. These pixels are marked under `FLAG6` in the HDU[2] retrospective headers (see Section 7.2.2.3).

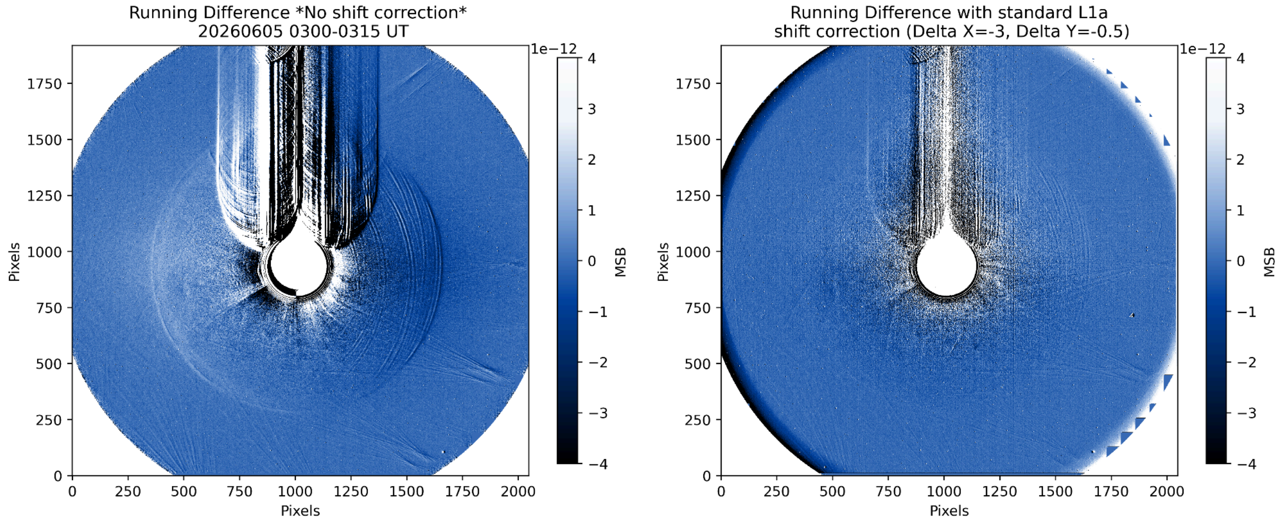

Figure 2: Two difference images computed between `level 1A` products without/with the coalignment applied in the left and right panel respectively. This particular example was chosen to straddle the daily HGA gimbal stepping period when we are usually guaranteed to see some image motion. The panel on the right shows the resulting improvement in the straylight artifacts achieved by shifting the images to coalign the centers of the occulters.

- Second, instrumental straylight patterns evolve over time. Because lens movement introduces optical distortions into the straylight, this effect is not currently correctable. As with previous coronagraphs, scattered or refracted sunlight affects the image in the following ways:
    - **Primary Straylight**: In products where the background has not been subtracted, the strongest straylight dominates the image signal. This is most intense near the central occulter and along the pylon, often manifesting as diffraction ring artifacts.
    - **Circular Ring**: A large, prominent circular ring (radius ~600 pixels) appears in the middle of the images, roughly corresponding to the edge of the vignetting region.
    - **Curved Arcs**: Large, curved arcs (several hundred pixels long) with a distinct bright-dark-bright intensity structure may be visible within the vignetted region (see Figure 3 right plot). These likely stem from imperfections in the pylon or occulter; because the lenses shift, these light paths map to different detector pixels between frames, preventing consistent subtraction. The origin of these features remains under investigation.
    - **Diffraction Spikes**: The "spiky" artifacts at the edge of the field of view are caused by residual diffraction from the entrance aperture's sawtooth pattern interacting with lens assembly components (such as spacers and retainers). When the lenses remain stationary these patterns also remain static, so they subtract effectively in background-corrected images. However, lens motion causes these patterns to shift and change orientation through a combination of refraction/reflection/scattering at lens interfaces. This creates residual artifacts in running difference and background subtracted images. In the example shown in

Figure 3 below, the plot on the left shows a difference-image made from two consecutive images where the lens motion is small so many of the stray light features get subtracted. The plot on the right shows an example where this is not the case and the straylight changes from image to image, causing artifacts.

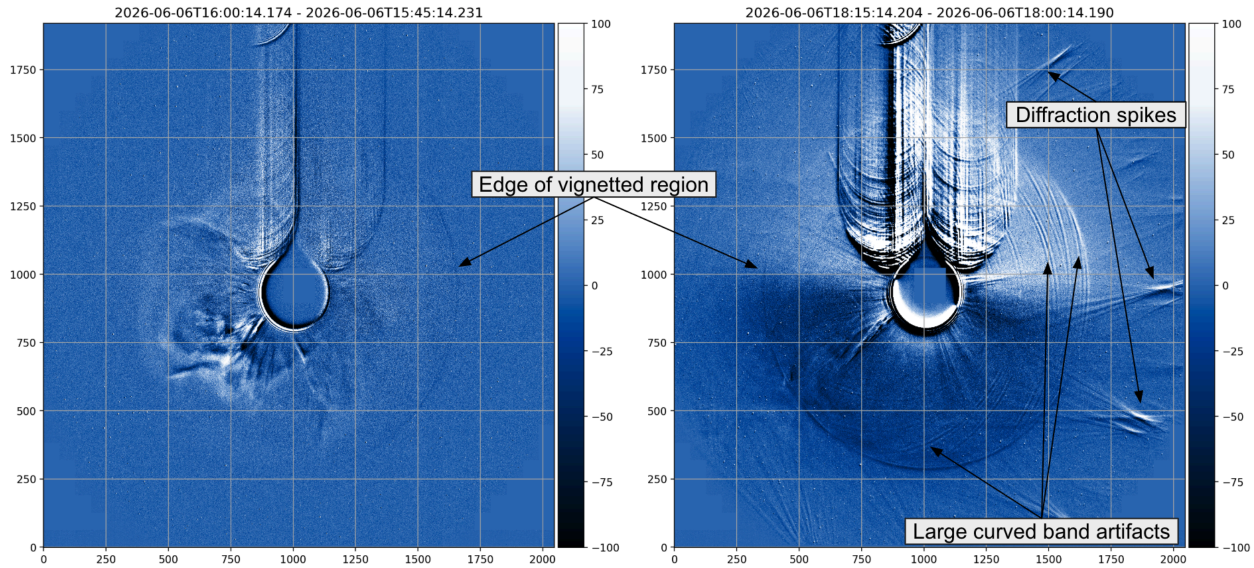

Figure 3: Two difference images computed from `level 0B` data by subtracting sequential images, for a case where the lens motions are small (left) and large (right).

4.2 Data validity start date

CCOR-2 data are provisionally valid beginning with observations on June 2, 2026. Observations prior to this date may have incorrect or incomplete navigation metadata. The National Centers for Environmental Information (NCEI) will reprocess and release early mission data using the Provisional Maturity algorithms and updated look-up tables.

4.3 Reading products using IDL

Because the files use internal compression the typical IDL FITS reader programs (readfits.pro, mreadfits.pro) may not work. A number of workarounds have been found so far:

- Use mreadfits_tiecomp.pro (or read_sdo.pro) with the keyword /USE_SHARED enabled.
- Use a Python-to-IDL bridge. One can then use the Python astropy package to read in the compressed FITS data from IDL.
- Use a simple Python script, like in the example below, to open the compressed file and save it to an uncompressed FITS file before reading it in using IDL.

```python
from astropy.io import fits
```

```
in_filename="/example_directory/example_image_in.fits" #replace this with the
path/filename of the image you want to open in IDL
out_filename="/example_directory/example_image_oout.fits" #replace this with the
path/filename of the uncompressed image
ext_num=1 #which HDU do you want? 1 is the standard image data for L0b, L1a, etc

fits_data=fits.getdata(in_filename,ext_num)
fits_header=fits.getheader(in_filename,ext_num)
fits.writeto(out_filename,fits_data,fits_header)
```

4.4 Image data gaps/corruption

CCOR-2 operates at a nominal 15-minute cadence. The NCEI-hosted products will be the most complete record that exists. Under standard operating procedure, the NCEI archive will continue to fill in missing data up to 9 days after the data was collected.

- The image is segmented into 64 x 64 pixel blocks before it is downlinked to the ground station. To conserve bandwidth when processing/downlinking data blocks in the corners of the image and at the center of the image array (roughly behind the occulter) are not downlinked. These pixels are marked using `FLAG7` (see section 7.2.2.3) in the pixel quality flag (PQF) arrays in retrospective products, and do not count towards the values recorded in `BADBLK_N` and `MISBLK_N`.
- Data is occasionally lost or badly corrupted (see Figure 4 for some examples). When this happens the data quality flag `IMGBLK_Q==F` and the keywords `BADBLK_N` and `MISBLK_N` have non-zero values (see Section 8 for full definitions).
- As of June 2, 2026, there is a known issue with the `level 0` decompression algorithm such that images with corrupted data packets in the telemetered stream will create a failure that corrupts the full image. As such, images with any amount of corruption should not be trusted currently. A patch is in progress for the decompression algorithm and for labeling corrupt pixel blocks in the retrospective data.

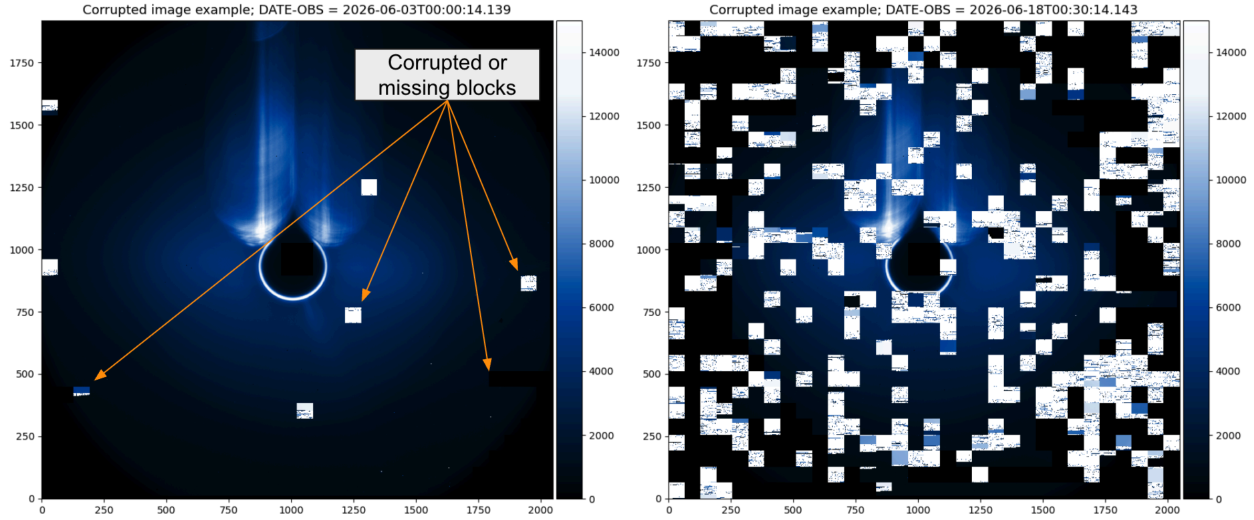

Figure 4: Example `level 0B` images with corrupted or missing data. These images have `IMGBLK_Q==F` either `BADBLK_N` or `MISBLK_N` set to a non-zero value.

4.5 Image dimensions

- 2048 x 1920 (full resolution) for all data products between `level 0A` and `level 2`.
- 1024 x 960 for `level 3` products.

4.6 Different filename conventions

Due to differences in data processing systems, operational and retrospective products follow distinct filenaming conventions. Detailed explanations of these conventions can be found in Sections 7.1.1 and 7.2.1.

4.7 Backgrounds

The image signals are dominated by instrumental stray light created inside the instrument by the pylon and occulter assembly, as well as sunlight scattered by solar system dust (F-corona). In order to bring out the K-corona signal in individual `level 2` images, backgrounds need to be subtracted from the images. Using established methodology for other coronagraphs, time-averaged backgrounds are computed using a pixel-wise minimum over median images computed for each UTC day. By time-averaging the images the backgrounds are meant to remove the K-corona signal which varies on the order of hours to weeks.

- Operational products use a 7-day lookback period to compute the background since they rely on the smallest possible cadence.
- Retrospective products use a 29-day window centered on the day for which the background is generated. Retrospective daily-median and monthly-minimum (`ccor2-dm` and `ccor2-mm`) products are available for download.

Due to the need to co-align images to correct for the lens movements, time-averaging induces some new artifacts in the backgrounds, which then feed into the `level 2` products. Figure 5 shows an example `level 2` file made using a retrospective background as well as some of the artifacts users may encounter.

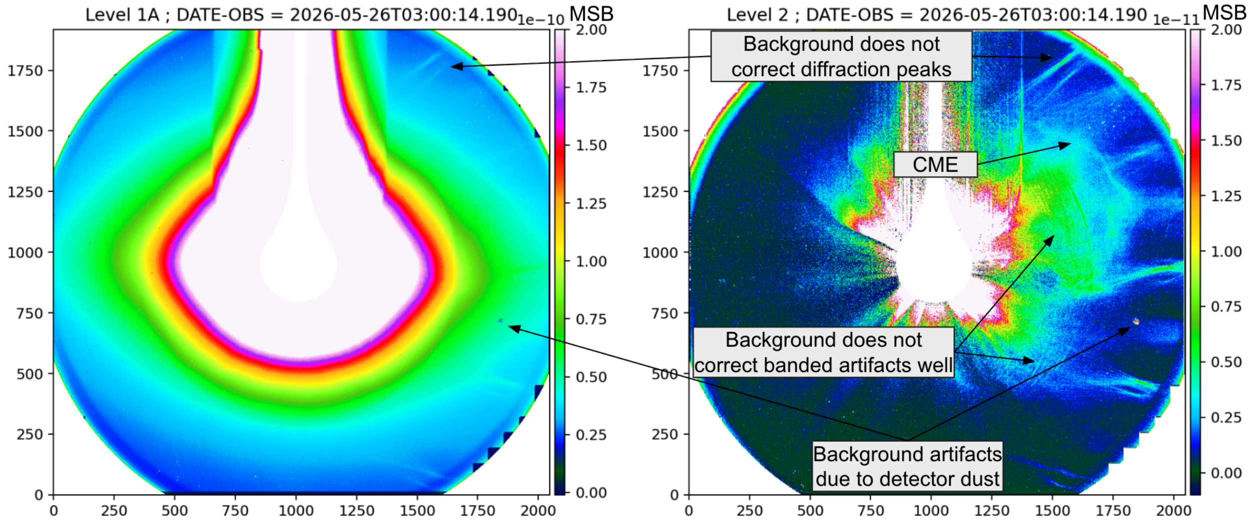

Figure 5: Example of retrospective `level 1A` (left) and `level 2` (right) images. The `level 1A` image intensity is dominated by the stray light and F-corona, while the `level 2` image has a monthly minimum (`ccor2-mm`, see Section 7.2) subtracted from data which removes the bulk of the stray light and F-corona. However, due to lens movements, some stray light residuals from the diffraction peaks and inner curved artifacts are not removed. Furthermore, bad pixels (e.g. due to dust) can become transformed into larger artifacts during the image co-alignment process and appear as stationary bright pixel clusters in the `level 2` images and animations.

4.8 Transient features

Some images or sequence of images may contain visual artifacts due to comets, planets, and debris coming off the spacecraft itself. Figure 6 shows some examples of types of transients observed so far:

- CCOR-2 images contain on average hundreds of stars transiting the FOV in any single image.
- Comets can transit the FOV for both short and long periods of time. Cometary appearance is similar to what is observed from the ground, with a bright cometary body (often saturated) with an elongated tail.
- Debris or dust falling off the spacecraft may traverse the telescope entrance and cause extra scattered light and streaks in the images.
- Optical “ghosts” appear when debris scatters light causing multiple reflections to appear in the image. These can appear as groups of bright sources moving across the image.

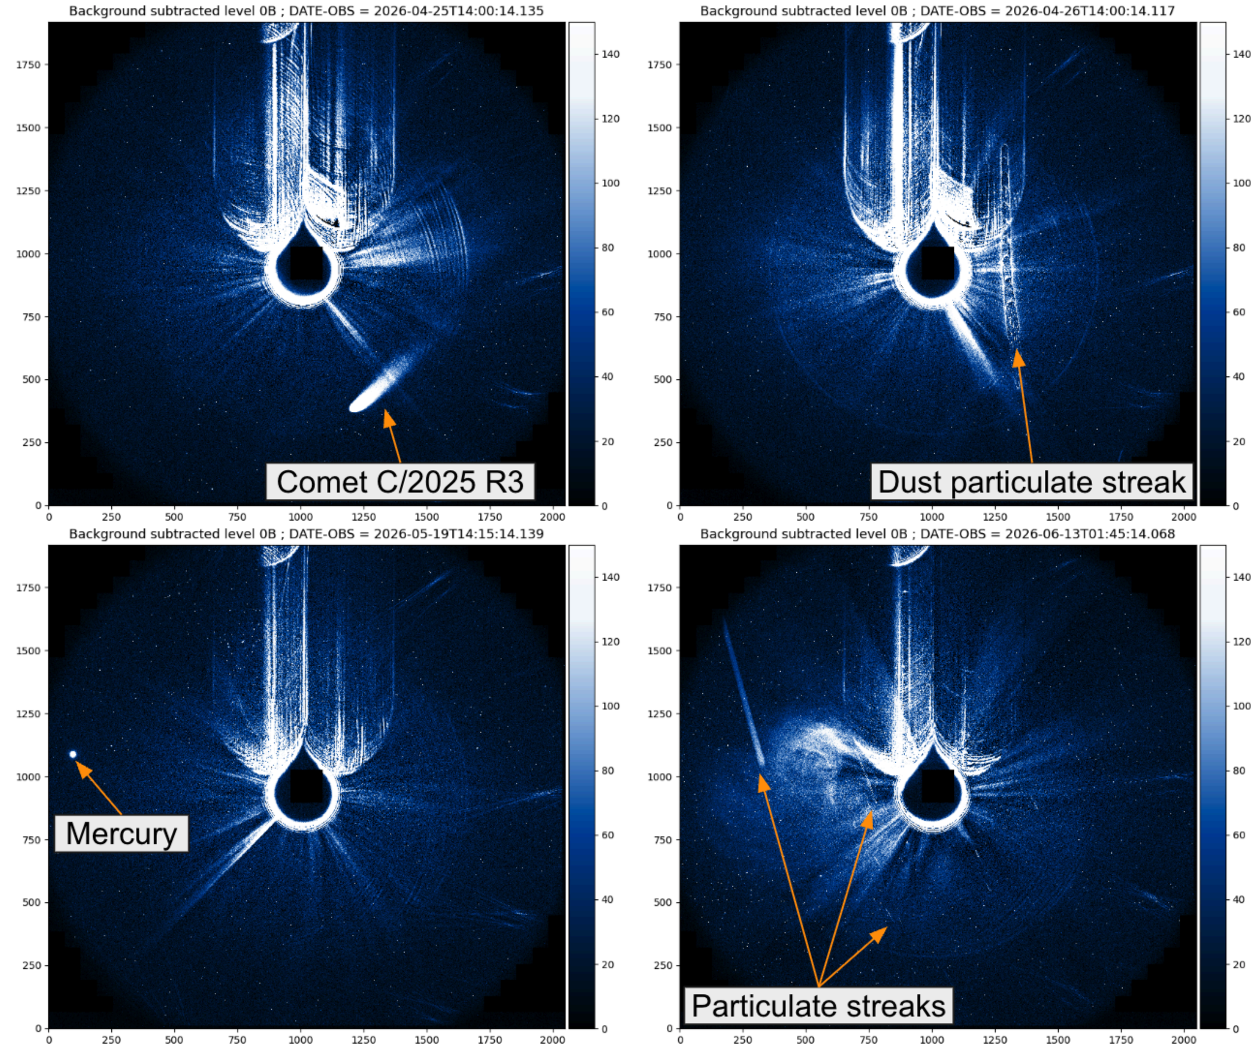

Figure 6: Example level 0B images with a daily minimum background subtracted to highlight examples of transient features in the CCOR-2 data.

4.9 Optical distortions of sources

Because the telescope has a varying point spread function (PSF) across the image point sources change shape as they traverse different regions of the image plane.

- For unsaturated sources like faint stars the effect is small enough to occur below resolution and noise thresholds.
- Saturated sources like the planet Mercury, other satellites or very bright stars show a change in their shape as they traverse the image from circular near the edge to more elongated in the inner region.

The lens motions do not appear to significantly distort the parts of the image coming from distant sources like stars and space weather phenomena. However, the local straylight artifacts like the occulter diffraction ring and diffraction spikes around the outer edge of the image do change shape due to optical effects. The occulter ring is sometimes enlarged by approximately 1-2%.

4.10 Attitude Information

CCOR-2 attitude information is stored in standardized WCS metadata within the first extension header (HDU[1]).

- **Reference Frames:** The default WCS uses a Helioprojective frame (Sun-centered with Solar North as the rotation reference), while WCSA and WCSB metadata use a Celestial frame.
- **Pointing Accuracy:**
    - `level 0B`: Pointing is typically accurate to within ±5 pixels, largely due to the lens-induced motions described in Section 4.1. Rare, temporary deviations of up to 10 pixels may occur due to unusual lens configurations within the lens barrel.
    - `level 1A+`: Post-coalignment processing significantly improves pointing accuracy for both operational and retrospective `level 1A+` products, reducing errors to a standard deviation of approximately 1 pixel.
    - Future improvements: NCEI plans to implement star-fitting optimization to achieve sub-pixel pointing accuracy (~0.3 pixels) for retrospective `level 1A+` products once validation is complete.

4.11 Metadata accuracy and completeness

Some metadata entries in CCOR-2 files may be incomplete or preliminary. Updates will be provided to reprocessed retrospective data.

- For strings, the default fill value is ‘FILL’, for numeric data it is -9999.
- Some FITS header comments may be truncated as we continue to adjust the headers.

4.12 Spikes and bad pixels

CCOR-2 uses onboard scrubbing to remove cosmic-ray hits, and images on the ground will be generally free of spikes due to cosmic ray hits. Defective pixels are currently not corrected in CCOR-2 images and users may encounter negative pixel values or NaNs in the image arrays. Each retrospective CCOR-2 product contains a corresponding pixel quality flag (PQF) mask added as an extra header data unit (see Section 7.2.2.3). Users should consult these masks before detailed analyses.

4.13 Platform location

The metadata contains information about the SOLAR-1 platform location specified in the Heliocentric Earth Ecliptic (HEE) system within the header metadata: `HEEX_OBS`, `HEEY_OBS`, `HEEZ_OBS`.

- The products also include placeholder values (still under analysis) for the Earth-Centered Earth-Fixed (ECEF) coordinate system within the following metadata: `OBSGEO-X`, `OBSGEO-Y`, and `OBSGEO-Z`.

- Users requiring heliocentric coordinates should consult references like Hapgood (1992) for conversion details.

4.14 Additional information from SWPC

We encourage users to also read any complementary documentation that is available now and will likely be continuously updated by SWPC:

- https://www.spaceweather.gov/products/ccor-1-coronagraph-experimental

5 Standard Use Cases

Each of the CCOR-2 data products has its own merits. Figure 7 provides a reference for the calibrations steps applied to different data products.

If you want to:

- Work with detector counts instead of calibrated intensity units
- Visually inspect artifacts such as those described in Section 4

Then you should work with `level 0B` data, but be aware that:

- The position of the image on the detector changes frame to frame (with a range of approximately 5 pixels for both axes). Before you do any time-dependent analysis of the data, the data should be co-aligned using star fitting or diffraction ring fitting (as is done in the `level 1A+` products).
- The vignetting correction should only be applied in the co-aligned frame. The occulter center is assumed to be at X=1010.0 pix, Y=935.0 pix.

If you want to:

- Analyze running difference imagery in physical units for identifying coronal mass ejection or temporal transients

Then you should work with `level 1A` data, but be aware that:

- The data has been bilinearly interpolated from a regular rectilinear grid onto a regular rectilinear grid to temporally co-align images
- The data has not been flatfielded before being interpolated so detector artifacts (low response and high response pixels) will introduce spurious time-dependent signals. These transients are only a few contiguous pixels in size. The apparent motion of transients should match the Cartesian path described by subsequent coalignment shifts (`SHIFT_X` and `SHIFT_Y`)
- Pixel Quality Flags in HDU[2] should be consulted for visualization and analysis. There is no PQF for hot/cold detector pixels currently.

If you want to:

- Analyze the evolution of quiescent K-corona structures or visualize transients in the context of the surrounding K-corona

Then you should work with `level 2` data, but be aware that:

- CCOR-2 background generation relies on 29 days of images to produce a very crude model of the F-corona+instrumental straylight+some residual fraction of the K-corona. Given that CCOR-2 instrumental straylight varies over minutes to hours timescales, `level 2` images will include a component of under-subtracted straylight as well the time-evolving K-corona. It would not be scientifically accurate to construe CCOR-2 `level 2` data as simply “the K-corona”.

6 Ground Processing Algorithm (GPA)

The schematic in Figure 7 summarizes the processing actions and products generated by two parallel but separate data processing flows running on the SWPC and NCEI systems.

The operational data processing flow is run by SWPC and combines telemetry and data CCSDS packet information into images and metadata that can be stored in FITS files. The most important driver for operational products is achieving low latency so forecasters can make space weather predictions. The current latency on these products is around 20 minutes and SWPC is making publicly available real-time `level 1A` data over the past 30 days at:

- https://services.swpc.noaa.gov/products/ccor2/

NCEI will be responsible for storing and disseminating operational products older than this time window. Operational data through NCEI’s archive will be available with a delay that can be as long as a few days due to how data is flowed internally between SWPC and NCEI.

Retrospective data products are produced by NCEI. The retrospective algorithm takes as input the operational `level 0B` products made by SWPC, delivered to NCEI after the end of every UTC day. SWPC and NCEI products at equivalent levels are generated by similar but distinct algorithms with key differences found in how backgrounds are produced. NCEI data may be reprocessed to take advantage of instrument trending work throughout the mission lifetime, which should not be expected of SWPC operational products.

NCEI plans to make available a Python library that contains the GPA code along with auxiliary routines that can be used to perform data analysis and visualization.

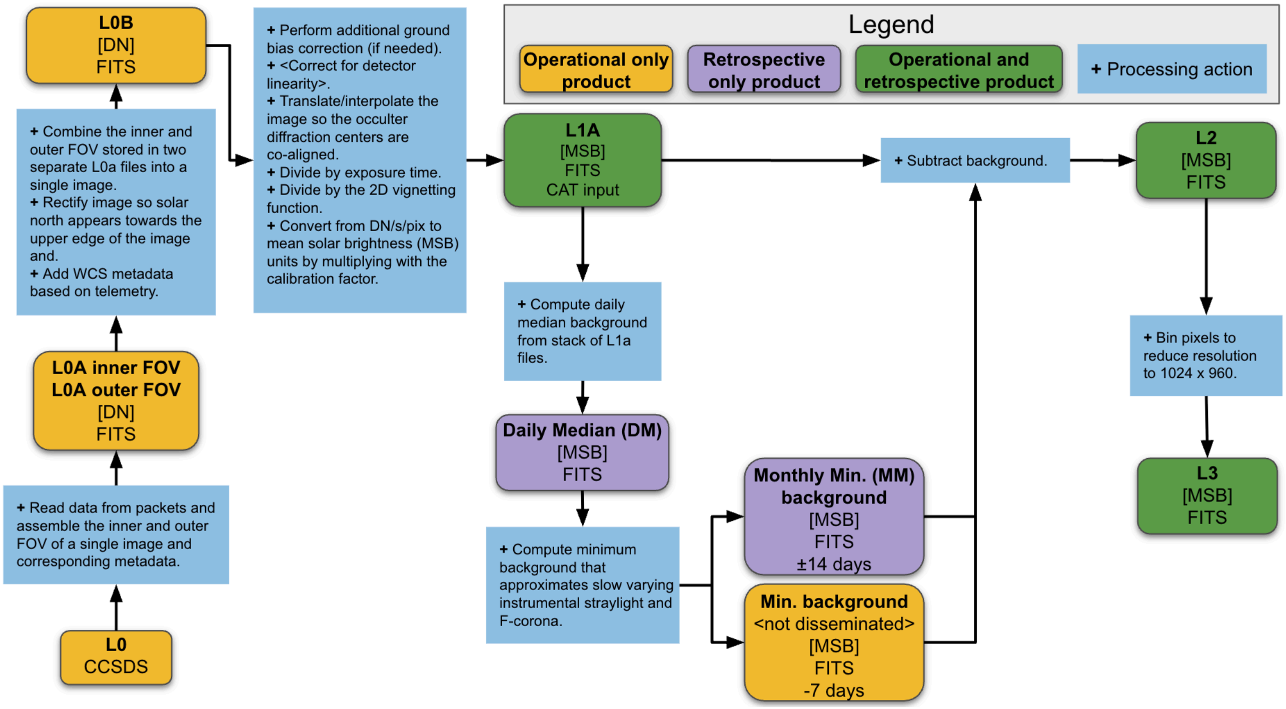

Figure 7: Schematic overview of the processing steps and products generated at each step of the operational (SWPC, low latency) and retrospective (NCEI, scientifically authoritative) GPAs. Inside each product shape we also indicate the units and file type for each product level.

7 Product descriptions

As indicated in Figure 7, operational (SWPC) and retrospective (NCEI) products are closely related, with the main differences being latency, header structure and as needed reprocessing to include calibration updates.

The latency impacts the time window used to compute the background image which represents an estimate of stray light and F-corona that gets subtracted when processing to level 2 and 3. The SWPC operational backgrounds are created using data looking backwards up to 7 days from the present date such that operational `level 2` and `level 3` files can be made available with low latency (<30 min).

The retrospective backgrounds utilize a longer 29-day time window that is centered on the day when a file is processed. As such there is a 14-day delay in producing `ccor2-l2` and `ccor2-l3` retrospective products. In principle by averaging over a full solar rotation the algorithm can achieve good mixing of transient features in the K-corona and computed improved backgrounds with fewer artifacts. In practice the need to co-align the images has been shown to induce artifacts due to the temporal filtering that is used to create the backgrounds (also see Section 4.1 and Figure 5).

Another difference between operational and retrospective products is that operational products are typically not reprocessed once created. Thus, while the operational products are much more suitable for space weather forecasting purposes due to their low latency (<30 min), the retrospective products are more suitable for detailed scientific analysis since they incorporate new information about instrument calibration. NCEI plans to make all products and processing levels available to allow community members to use the data most suitable for their specific analysis needs.

7.1 Operational products

7.1.1 Filenaming convention

```
[instrument]_[prod_level]_[date_obs]_[version]_[sucode].[extension]
```

Example operational CCOR-2 filenames:

```
CCOR2_0A_20260609T054514_V00_T1.fits
CCOR2_0A_20260609T054514_V00_N2.fits
CCOR2_0B_20260609T054514_V00_NC.fits
CCOR2_1A_20260609T061514_V00_0C.fits
```

Table 1: Detailed filename field descriptions for the operational products.

Field	Description	Allowed (or example) values
`instrument`	The instrument from where the data came from. The GOES-19 spacecraft contains the CCOR-1 instrument.	`CCOR2`
`prod_level`	The operational data product levels names.	`0A, 0B, 1A, 2, 3` See Table 2 below for a full description of each product level.
`date_obs`	Start datetime (UTC) indicating the start of the first exposure (`DATE-OBS` metadata). Format is: `[yyyy][mm][dd]T[hh][mm][ss]`	Example: `20260609T054514`
`version`	[legacy] GPA software version number. Format is: `V[version]`. This labelling method is not being used. The operational GPA version is instead recorded in the primary header of the files (HDU[0]) and as part of the HISTORY processing stream for the files.	Examples: `V00, V01`

`sucode`	Combination code that indicates the image type. It is used to clasify images taken at different cadence or as part of calibration maneuvers.	`T1` : Present in the 0A products. Indicates the first (inner) portion of the image. `N2` : Present in the 0A products. Indicates the second (outer) portion of the image. `NC` : Present in 0B and higher products. Indicates a nominal 15 minute cadence image. `0C, EC` : Present in 0B and higher products. Indicates a full resolution image taken at an irregular cadence, typically as part of calibration maneuvers or long term instrument trending. `0B` : Present in 0A products. Indicates a bias image frame which is only 64x2048 pixels in size. Other examples that may appear in `0A` products: `R2, L1, 01, 00`.
`extension`	The data file extension.	`fits`

7.1.2 Product structure

All operational products use the astropy.io.fits module to internally compress the images using GZIP_1. Using internal compression enforces that the written files contain at least two header data units (HDUs) with the primary containing an empty data portion and the compressed image being located in the first extension. All the operational products contain exactly two HDUs.

7.1.2.1 Primary HDU

The header only contains information about the SWPC operational GPA version and a `HISTORY` keyword describing the structure of the file.
The binary data portion of the primary HDU is empty.

7.1.2.2 First extension HDU

This HDU contains the operational metadata and image data. Table 2 provides a summary description of the data. Currently only a small subset of metadata items in the headers contain comments. A fuller description of what the header metadata entries mean can be found in a separate document called CCOR-2 Operational Product Metadata Definitions. This document continues to be updated so may not always contain a complete list of the keywords found in the products. Incremental processing applied to each image is tracked through the HISTORY keyword.

Table 2: Description of the operational data stored in the first extension HDU.

Product level	Image dimensions (pixels)	Units	Processing applied
`0A`	2048 x 1920	DN	Constructed from lowest level CCSDS packets. Contains the inner or outer portion of the image.
`0B`	2048 x 1920	DN	Contains the combined inner and outer portions of the image. Rotated/reflected so that Solar N points towards the top edge of the image.
`1A`	2048 x 1920	MSB	Time normalized. Flat field corrected by division through the vignetting function. Converted from DN to MSB units through multiplication with the photometric calibration factor stored under the keyword `CALFAC`. Detector nonlinearity correction is currently NOT applied pending more testing. Bias subtraction reconstruction is currently NOT applied pending more testing.
`2`	2048 x 1920	MSB	A background image that approximates the instrumental straylight + F-corona background is subtracted from the 1A products. The background is computed as a minimum over daily median images computed from `level 1A` files over the past 7 days.
`3`	1024 x 960	MSB	Image resolution is downsampled by a factor of 2 along each image axis. This product level is more suitable for making animations that are smaller in size.

7.2 Retrospective products

7.2.1 Filenaming convention

```
[env]_[dsn]_[satellite]_[startDate]_[endDate]_[processDate]_[access].
[extension]
```

Example retrospective CCOR filenames:

```
sci_ccor2-l1a_solar1_s20260530T000014Z_e20260530T000043Z_p20260531T07
0730Z_pub.fits
sci_ccor2-l2_solar1_s20260530T000014Z_e20260530T000043Z_p20260531T070
730Z_pub.fits
sci_ccor2-mm_solar1_s20260530T000000Z_e20260530T235959Z_p20260531T143
300Z_pub.fits
```

Table 3: Detailed filename field descriptions for the operational products.

Field	Description	Allowed (or example) values
`env`	The ground system environment where the product was created and/or the type of data.	`sci`
`dsn`	The product Data Short Name.	`ccor2-l1a, ccor2-l2, ccor2-l3, ccor2-dm, ccor2-mm` See Section 6.2.2 for a full description of each product level.
`satellite`	The platform/satellite from which the product came.	`solar1`
`startDate`	The coverage start datetime of the data. For `ccor2-l1a, ccor2-l2` and `ccor2-l3` products this time refers to the start of the first exposure time (`DATE-BEG` metadata). For `ccor2-dm, ccor2-mm` products this time refers to the start of the UTC day for which the files are valid. Format is: `s[yyyy][mm][dd]T[hh][mm][ss]Z`	Example: `s20260530T000014Z`

`endDate`	The coverage end datetime of the data. For `ccor2-l1a, ccor2-l2` and `ccor2-l3` products this time refers to the end time of exposures used in the image average (`DATE-END` header keyword). For `ccor2-dm, ccor2-mm` products this time refers to the end of the UTC day for which the files are valid. Format is: `e[yyyy][mm][dd]T[hh][mm][ss]Z`	Example: `e20260530T000043Z`
`processDate`	The product file creation datetime. Format is: `p[yyyy][mm][dd]T[hh][mm][ss]Z`	Example: `p20260531T070730Z`
`access`	Public access level. Either “pub” for public release or “emb” for embargoed (no public release).	`pub, emb`
`extension`	The data file extension.	`fits`

7.2.2 Product structure

All retrospective products use the astropy.io.fits module to internally compress the images using RICE_ONE. Using internal compression enforces that the written files contain at least two header data units (HDUs) with the primary containing an empty data portion and the compressed image being located in the first extension. As detailed below, retrospective products contain extra HDUs to store the pixel quality flag (PQF) mask and lists of files used to generate the aggregate daily median and monthly minimum data.

7.2.2.1 Primary HDU

The header contains a set of keywords summarizing the information in the extension HDUs. The `HISTORY` keyword is used to transmit details about the type of compression used to make the retrospective product.
The binary data portion of the primary HDU is empty.

7.2.2.2 First extension HDU

This HDU contains the main science and telemetry metadata in the header and the image data itself (Table 4). The metadata entries contain comments to summarize the meaning of each keyword but the users should also consult the document on CCOR-2 Operational Product Metadata Definitions. To improve readability some of the engineering keywords that are included in the operational products will not be included in the retrospective products. More detailed documentation for the retrospective product metadata will be made available in the future.

Incremental processing applied to each image is also tracked through the `HISTORY` keyword. Retrospective versions of products corresponding to the operational `level 0A` and `level 0B` are currently not being produced but may become available in the future.

Table 4: Overview of the image data stored in the retrospective products stored in the first extension HDU.

Data Short Name	**Image dimensions (pixels)**	**Units**	**Processing applied**
`ccor2-l1a`	2048 x 1920	MSB	The input data is the operational `level 0B` product which is then: ● Time normalized. ● Flat field corrected by division through the vignetting function. ● Converted from DN to MSB units through multiplication with the photometric calibration factor stored under the metadata keyword CALFAC. Detector nonlinearity correction is currently NOT applied pending more testing. Bias subtraction reconstruction is currently NOT applied pending more testing.
`ccor2-dm`	2048 x 1920	MSB	Daily medians are aggregate images that are computed as a median over `ccor2-l1a` images that do not contain any bad or missing blocks (BADBLK_N = 0 and MISBLK_N = 0) are deemed valid (ISVALID = True).
`ccor2-mm`	2048 x 1920	MSB	Monthly minimums are aggregate images that are computed as a pixel-wise minimum over `ccor2-dm` images. The files are picked to span a 29 day window centered on the day for which the `ccor2-mm` file is valid. Only images with the same yaw flip orientation are used.
`ccor2-l2`	2048 x 1920	MSB	A `ccor2-mm` background image that approximates the instrumental straylight + F-corona background is subtracted from the `ccor2-l1a` products.
`ccor2-l3`	1024 x 960	MSB	Image resolution is degraded by a factor of 2 along each image axis. This product level is more suitable for making animations that are smaller in size.

7.2.2.3 Second extension HDU

The data portion of the HDU stores a pixel quality flag (PQF) mask which is an array of the same dimensions as the image stored in the first extension HDU (e.g. 2048 x 1920 or 1024 x 960 pix). The values stored for each pixel are integer values that are sums of powers of 2 which encode different flags. This type of scheme allows a pixel to have several flags “raised on it” simultaneously.

The flag types are recorded in the header portion of the HDU and summarized Table 5 for each product. Many of the flags are set based on the `level 0B` product pixel values which are measured in DN.

For the binned `ccor2-l3` products, a bitwise OR is performed on the pixels that are being binned such that all the flags from the pixels that get binned are preserved in the lower resolution pixel.

Table 5: Definitions for the flags used by the retrospective products to compute the PQF data array stored in the second extension HDU.

Flag name	**Flag value**	**Flag meaning**	**Products that include this flag.**
`FLAG0`	0	`Level 0B` pixel value in linear range (1000 ≤ I(DN) ≤ 11580)	`ccor2-l1a, ccor2-l2, ccor2-l3, ccor2-dm, ccor2-mm`
`FLAG1`	1	Vignetting function value < 0.1	`ccor2-l1a, ccor2-l2, ccor2-l3, ccor2-dm, ccor2-mm`
`FLAG2`	2	Vignetting function value < 0.01	`ccor2-l1a, ccor2-l2, ccor2-l3, ccor2-dm, ccor2-mm`
`FLAG3`	4	`Level 0B` pixel value below linear range (I(DN) < 1000)	`ccor2-l1a, ccor2-l2, ccor2-l3`
`FLAG4`	8	`Level 0B` pixel value in correctable nonlinear range (11580 < I(DN) < 15300)	`ccor2-l1a, ccor2-l2, ccor2-l3`
`FLAG5`	16	`Level 0B` pixel value is saturated (I(DN) > 15300)	`ccor2-l1a, ccor2-l2, ccor2-l3`
`FLAG6`	32	Pixels with unreliable photometry.	`ccor2-l1a, ccor2-l2, ccor2-l3, ccor2-dm, ccor2-mm,`

`FLAG7`	64	Dead pixels or no data being downlinked intentionally or due to failure (I(DN) = 0)	`ccor2-l1a, ccor2-l2, ccor2-l3`
`FLAG8`	128	Pixel value is equal to 0/NaN/Inf in the background. This flag is fundamentally a way to flag pixels that get affected by the image aggregation when the `ccor2-dm, ccor2-mm` products are created.	`ccor2-dm, ccor2-mm, ccor2-l2, ccor2-l3`

Extracting Flags with Bitwise Operators

To extract pixels from the PQF data pertaining to single or multiple flags, bitwise comparative operators must be used. Because each flag corresponds to a unique power of two, their bits do not overlap, allowing multiple flags to be stored in a single integer value (see Table 6). We plan to release more detailed usage examples as well as include routines in the Python libraries to decode the flags and make masks. The examples below are meant to illustrate a simple way the users could use the PQF mask.

Table 6: Example integer and binary representation for flag combinations.

Flag(s)	Integer Value	8-bit Binary
`FLAG1`	1	`00000001`
`FLAG2`	2	`00000010`
`FLAG1 + FLAG2`	3	`00000011`
`FLAG6`	32	`00100000`
`FLAG1 + FLAG5 + FLAG6`	49	`00110001`

You can use bitwise operators to interact with these flags. The OR (|) operator combines flags into a single mask, and the AND (&) operator checks for the presence of specific flags within the PQF data.

```python
# Assume pqf_data is the 2D array from the second extension HDU

# 1. Combine flags to create a mask for multiple conditions
```

```
# Example: Create a mask for pixels that are either Vignetting level 1 (1) OR level 2 (2)
# Binary: 00000001 | 00000010 = 00000011 (Integer: 3)
mask_vignetting = 1 | 2

# 2. Extract pixels where any of the flags in the mask are set
# This returns True for pixels having FLAG1, FLAG2, or both
pixels_vignetting = (pqf_data & mask_vignetting) > 0

# 3. Check for specific conditions (e.g., FLAG6 for unreliable photometry)
# Bitwise AND results in > 0 if the flag is present
pixels_unreliable = (pqf_data & 32) > 0

# 4. Check for pixels that are BOTH saturated (FLAG5=16) AND have unreliable photometry (FLAG6=32)
# Integer value: 16 | 32 = 48 (Binary: 00110000)
mask_saturated_unreliable = 16 | 32
pixels_both = (pqf_data & mask_saturated_unreliable) == mask_saturated_unreliable
```

7.2.2.4 Third extension HDU

This HDU is only included in the daily median and monthly minimum (`ccor2-dm, ccor2-mm`) products. The data portion stores a list of filenames used to generate the daily median and monthly minimum aggregate products. For the daily median products this list may be as long as 96 entries but may have fewer entries due to missing files and exclusion of images that contain bad or missing data blocks.

8 Data quality flags

In addition to the PQF mask (see Section 7.2.2.3) both operational and retrospective products have a number of data quality flags that are used to automatically identify images with anomalies in them based on the information available in the telemetry. A summary of the flags is included in Table 7.

Table 7: Summary of the data quality flag metadata included in the operational and retrospective products.

Keyword	Value type	FITS comment
`ISVIABLE`	`bool`	`T = image is nominal -- no degraded inputs`
`ISNORMAL`	`bool`	`T = image is full resolution (2048x1920) and taken with the regular inner/outer mask`

DCMPRS_Q	bool	T = decompression algorithm was successfully applied
IMGBLK_Q	bool	T = no bad or missing blocks were found in the image data
BADBLK_N	int	Number of bad blocks (valid only for NORMAL images) reported by the decompression code
MISBLK_N	int	Number of missing blocks (valid only for NORMAL images) reported by the decompression code
TMTIME_Q	bool	T = Temperature measurement time (APID 982) is within 60s of image acquisition time
TMTIME_N	int	[s] Time difference between temperature measurement time (APID 928) and image acquisition
ADCS_Q	bool	T = ACDS data time (APID 110) is within 60s of image acquisition
ADCS_N	int	[s] Time difference between the ADCS data time (APID 110) and image acquisition
ADCSSQ	int	ADCS data (APID 110) sequence counter
EPTIME_Q	bool	T = Ephemeris (APID 100) time is within 60s of image acquisition time
EPTIME_N	int	[s] Time difference between ephemeris data record (APID 100) and image acquisition
EPTIMESQ	int	Ephemeris data (APID 100) sequence counter
EPVALID	bool	T = Ephemeris data valid
ATTVALID	bool	T = Attitude data valid
SUNPNT_Q	bool	T = Instrument is in nominal sun-pointing mode

9 Contact for Further Information

For general inquiries, contact:
swx.coronagraph@noaa.gov
SPSD.UserServices@noaa.gov (OSPO User Services)

Users are also encouraged to contact the NCEI GOES-R CCOR team for further assistance, questions, or to report any issues with CCOR files. Additional data access information is available on the NCEI website:
https://www.ncei.noaa.gov/products/space-weather/swfo