

Calibration of the Ultraviolet Cameras on NASA's Carruthers Geocorona Observatory

Author: Alex Zhang

Undergraduate Senior Thesis in Electrical Engineering
University of Illinois Urbana-Champaign
Advisor: Professor Lara Waldrop

May 2023

Abstract

The Carruthers Geocorona Observatory is NASA's first mission dedicated to investigating the fundamental nature of Earth's outermost atmospheric layer known as the exosphere. The Carruthers mission's primary payload is the GeoCoronal Imager, which consists of two co-aligned imagers (channels) for simultaneous, column-volume sensing of ultraviolet emission by exospheric hydrogen atoms. These images will be used to determine the spatial structure and temporal variability of the exosphere on global scales, but a prerequisite of that analysis is isolation of the target emission from instrument and scene backgrounds as well as calibration of the absolute photometric sensitivity of the cameras. This thesis develops the calibration algorithms that will be used for mission data analysis, including: (i) bias and dark current estimation and subtraction, (ii) estimation and subtraction of energetic particle radiation backgrounds, (iii) background photon scene gradient detection and removal, and (iv) flat-fielding. The algorithms are validated using a synthetic image model to measure their accuracy and precision, while laboratory calibration data is used to verify instrument performance against design requirements.

Subject Keywords: Remote sensing, Exosphere, Optical Sensor Calibration, UV Emission

To Sasha for her endless love and support.

Acknowledgements

A special thanks to Professor Lara Waldrop for her advice and direction in this effort. Much gratitude to the Carruthers Algorithms team - Heather Filippini, Carlos Suberviola, Pratik Joshi, Evan Widlowski, and John Wolff - for the design and implementation of the synthetic image model. Thank you to Jason McPhate on the Carruthers Science team for his advice and guidance. Thank you to Robert Irvin for his support on the PATHS cameras.

I would like to express gratitude to all of the ECE Department faculty members for their help and support over the years. I also thank my parents for their love and support.

I would also like to express my sense of gratitude to anyone and everyone who directly or indirectly have lent their hand in this effort.

Table of Contents

1	Introduction	1
1.1	Carruthers Geocorona Observatory Mission Overview	1
1.2	Measurement Concept	2
1.3	Measurement Calibration Overview	7
1.3.1	Validation of Measurement Calibration	7
1.4	Organization of Thesis	10
2	Instrument Background	11
2.1	Carruthers Camera Analysis	11
2.1.1	Engineering Model Camera Overview	11
2.1.2	Bias Frames	11
2.1.3	Dark Current	14
2.2	PATHS Detector Analysis	17
2.2.1	PATHS Overview	17
2.2.2	PATHS Bias Frame	17
2.2.3	PATHS Dark Current	18
3	Particle Radiation	21
3.1	Initial Particle Flux Conditions	21
3.2	MCP-Based Radiation Background	22
3.3	APS-Based Radiation Background	22
3.4	Science Image Corner SNR	23
4	Photon Background	25
4.1	In-Band (IPH) Photon Background	25
4.2	Shot Noise Denoising	26
4.3	IPH WFI Interpolation	28
4.3.1	Problem Statement	28
4.3.2	Assumptions and Notation	28
4.3.3	WFI IPH-Estimator Algorithm	29
4.3.4	Validation	34
4.3.5	Failure of Linearity Assumption	35
4.4	IPH NFI Interpolation	37
4.4.1	Algorithm	37
4.4.2	Validation	38
5	Flat-Field Retrieval	40
5.1	Algorithm	40
5.2	Experimental Validation	40
5.3	Theoretical Validation	46
5.3.1	All-ones Flat-Field Retrieval	47
5.3.2	True Flat-Field Retrieval	49
6	Conclusion	51
	References	52

List of Abbreviations

ADC	Analog-to-Digital Converter
APS	Active Pixel Sensor
C	Celcius
CREME-MC	Cosmic Ray Effects on MicroElectronics Monte Carlo simulation
COSSMo	Carruthers Observatory Student Solar Monitor
DN	Digital Numbers
ECEB	Electrical and Computer Engineering Building
FOV	Field-Of-View
FUV	Far Ultra-Violet
GCI	GeoCoronal Imager
GOES	Geostationary Operational Environmental Satellite
H	Hydrogen
HG	High Gain
hIT	hour Integration Time
ICP	Instrument Control Panel
IOC	Initial On-orbit Checkout
IPH	InterPlanetary Hydrogen
LG	Low Gain
LOS	Line-Of-Sight
MCP	MicroChannel Plate
MMPE	Maximum Magnitude Percent Error
MMMPE	Maximum Magnitude Mean Percent Error
MMPESD	Maximum Magnitude Percent Error Standard Deviation
NFI	Narrow Field Imager
NOAA	National Oceanic and Atmospheric Administration
OOB	Out-Of-Band
PATHS	Photometer Array Tomagraphic Hydrogen Sensing
pfu	proton flux units
Re	Earth-radii
s	seconds
SEP	Solar Energetic Particle(s)
SDL	Space Dynamics Laboratory
SNR	Signal-to-Noise Ratio
Std	Standard Deviation
UV	UltraViolet
WFI	Wide Field Imager

List of Symbols

Symbol	Name	Units
A	Aperture area	[cm ²]
b	Bias	[DN]
B	Noisy bias image(s)	[DN]
c	Mean number of electron counts	[electron counts]
C	Noisy number of electron counts	[electron counts]
d	Mean dark current	$\left[\frac{\text{electron counts}}{\text{second}} \right]$
D	Noisy dark current	$\left[\frac{\text{electron counts}}{\text{second}} \right]$
e	Mean number of photons	[photoelectron events]
E	Noisy number of photons	[photoelectron events]
F	Flat-field factor	
g_{MCP}	mean MCP Gain	$\left[\frac{\text{electron counts}}{\text{photoelectron events}} \right]$
G_{MCP}	MCP Gain	$\left[\frac{\text{electron counts}}{\text{photoelectron events}} \right]$
G_{APS}	APS Gain	$\left[\frac{\text{DN}}{\text{electron counts}} \right]$
H	Ground-truth IPH	[DN]
\hat{H}	Retrieved IPH from estimation algorithms	[DN]
i	Pixel row index	
j	Pixel column index	
k	Frame or trial index	
L	Spectral radiance	$\left[\frac{\text{photons}}{\text{cm}^2 \cdot \text{s} \cdot \text{sr} \cdot \text{nm}} \right]$
m	Mean signal with mean instrument background subtracted out	[DN]
M	Noisy signal with mean instrument background subtracted out	[DN]
\hat{n}	Line-of-sight	
N_{bin}	Number of APS pixels per binned WFI or NFI pixel	
N_{corner}	Number of binned pixels in all corners of an image	
N_{frames}	Number of frames stacked in one image	
N_{mir}	Number of mirrors	
N_{pix}	Number of binned pixels in one row or column of an image	
N_{trials}	Number of trials performed	
P	Percent error	[%]
q_{events}	Mean photoelectron events generated by particle radiation	$\left[\frac{\text{photoelectron events}}{\text{frame}} \right]$
Q_{events}	Noisy photoelectron events generated by particle radiation	$\left[\frac{\text{photoelectron events}}{\text{frame}} \right]$
q_{counts}	Mean electron counts generated by particle radiation	$\left[\frac{\text{electron counts}}{\text{frame}} \right]$
Q_{counts}	Noisy electron counts generated by particle radiation	$\left[\frac{\text{electron counts}}{\text{frame}} \right]$
R	Read noise	[DN]
r_{ij}	Radius of pixel ij in polar coordinates	[binned pixels]
s	Mean signal	[DN]
S	Noisy signal	[DN]
T_f	Frame integration time	$\left[\frac{\text{seconds}}{\text{frame}} \right]$
T_{int}	Integration time	[seconds]
γ	Mean and variance of a Poisson distribution	depends on context
$\Gamma(\lambda)$	Mirror reflectivity	
ϵ	A positive constant	depends on context
$\epsilon(\lambda)$	Channel optical efficiency	$\left[\frac{\text{photoelectron events}}{\text{photon}} \right]$
$\eta(\lambda)$	Photocathode quantum efficiency	$\left[\frac{\text{photoelectron events}}{\text{photon}} \right]$
θ_{ij}	Angle of pixel ij in polar coordinates	[degrees]
λ	Wavelength	[nm]

μ	Mean	depends on context
σ	Standard deviation	depends on context
$\tau(\lambda)$	Filter transmissivity	
ϕ	Particle flux through aluminum shielding	$\left[\frac{\text{particles}}{\text{second} \cdot \text{cm}^2} \right]$
Ω_{ij}	APS pixel solid angle	[sr]

1 Introduction

The Carruthers Geocorona Observatory will explore a frontier in science: the fundamental nature of the boundary between Earth’s neutral atmosphere and interplanetary space, also known as the Earth’s exosphere. The exosphere is defined as the outermost layer of the planetary atmosphere. In Earth’s case, it extends from the exobase (~ 500 km altitude) to $>150,000$ km. Available wide-field images of the exosphere are extremely sparse; only one image from a given platform has ever been acquired. The Carruthers Geocorona Observatory aims to fill this knowledge gap.

1.1 Carruthers Geocorona Observatory Mission Overview

Earth’s exosphere spans many distinct plasma environments where ion-neutral charge exchange (Figure 1) takes place. Charge exchange with exospheric Hydrogen (H) atoms influences the severity and duration of geomagnetic storms, making them extremely important to understand. In addition, planetary atmospheres evaporate through their exospheres. The gravitational escape of H atoms causes permanent loss of surface water on Earth and Mars. However, the role of planetary magnetic fields and trapped plasma populations in mediating the atmospheric mass loss through gravitational escape is poorly understood.

The spatial distribution of exospheric H atoms is a unique tracer of ion-neutral coupling/charge exchange and H escape. However, current analytical theories, which must account for the transition between a fluid regime to a kinetic one, are intractable, while physics-based models are outdated, oversimplified, and/or data-starved. New measurements are needed to understand the fundamental physical processes governing exospheric structure and dynamics, such as diffusion limited supply of H from the lower atmosphere (MLT), the transition from collisional to “collision-less” regime, charge exchange interactions with ambient ions, asymmetric, non-Maxwellian velocity distributions, temperature-driven, global scale, lateral flows, solar radiation pressure deformation into the “geotail”, and H loss via photoionization and solar wind pickup.

The science objectives of the mission are as follows:

1. Determine the drivers of quiet-time exospheric structure on regional and global scales. What is the relative significance of thermal evaporation, charge exchange, and solar radiation pressure in shaping the exosphere during geomagnetically quiet conditions?
2. Determine the nature and origin of transient variability in the exospheric structure. On what spatial and temporal scales does the exosphere respond to impulsive space weather events? What is the relative significance of thermal upwelling and charge exchange in driving this variability?

Exospheric H atoms have a relatively small number of electronic transitions that give rise to emission lines, the most important of which are the Lyman series and the Balmer series (see Figure 2). The H emission is generated via the resonant scattering or fluorescence of solar photons.

Balmer- α emits visible light, which is observable from the ground. However, the emission radiance is not very bright (~ 4 Rayleighs), so it can only be observed at night in clear dark sky conditions. Ground observations typically only support one line of sight, which also gives an unconstrained retrieval problem. An array of cameras for the Photometer Array Tomographic Hydrogen Sensing (PATHS) project are set up on the roof of the Electrical and Computer Engineering Building (ECEB) to detect Balmer- α emissions.




Figure 1: Charge Exchange in the Exosphere

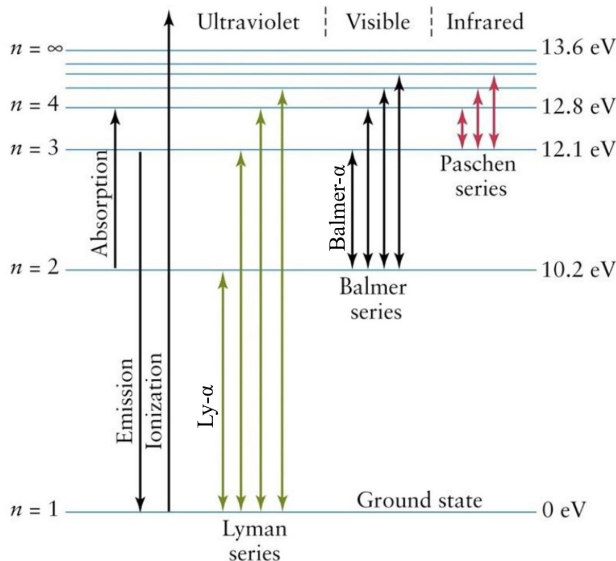


Figure 2: Grotrian diagram for atomic Hydrogen

Calibration algorithms in this thesis will also be tested on PATHS as a part of validating that algorithms function as intended.

On the other hand, the far-ultraviolet (FUV) Ly- α line at 121.6nm is the brightest emission line by far (~ 35 kilo-Rayleighs), but it is not observable from the ground because low-atmospheric molecular oxygen absorbs UV emission. Therefore, the target observable of the Carruthers mission is Ly- α emission radiance, due to its brightness and the fact that it can only be observed by spacecraft.

The Carruthers Geocorona Observatory’s primary payload is the GeoCoronal Imager (GCI), which consists of two co-aligned imagers (channels) for simultaneous sensing of exospheric Ly- α . The Narrow Field Imager (NFI) provides high spatial resolution near the Earth’s limb where the exosphere radiance falls rapidly. The Wide Field Imager (WFI) captures the entire exosphere with the sensitivity to detect the dimmest Ly- α signals at the outermost extent of the exosphere. Both channels employ identical UV-intensified, Active Pixel Sensor (APS)-based cameras and identical filter wheels. A schematic of the GCI is shown in Figure 3.

Ride-share with NASA’s IMAP mission offers the extremely rare opportunity for deployment to an ideal vantage: a halo orbit around the Earth-Sun L1 Lagrange point (~ 1.5 million km from Earth). The GCI will then have dedicated Earth nadir viewing and will provide nearly continuous, high cadence Ly- α imaging of exospheric Ly- α emission. In addition, from this vantage point, the Carruthers Observatory Student Solar Monitor (COSSMo) instrument, a student collaboration experiment attached to the top of the GCI instrument, shown in Figure 3, can monitor the solar UV photon flux which excites the geocoronal Ly- α emission.

The mission will launch in early 2025. After four months of cruising and an additional month for the initial on-orbit checkout (IOC) phase, the 24-month science phase will commence in June 2025. By this point, solar activity is expected to be near the maximum of its 11-year cycle, which will allow the mission to observe several geomagnetic storms and Earth-Sun interactions. The orbital period is around 6 months, with a maximum angle from the Earth-Sun axis of 26° . The trajectory of the mission is shown in Figure 4.

1.2 Measurement Concept

For this section and the remainder of the thesis, anything enclosed in brackets [] are units, unless otherwise stated.

The GCI instrument model is shown in Figure 5, while the optical path is shown in Figure 6. The incident photons pass through a flat MgF2 window, are collimated by curved mirrors, and pass through an interference filter before impinging on a KBr photocathode. Photon detection by the cathode liberates




Figure 3: Carruthers Geocorona Observatory's GCI payload

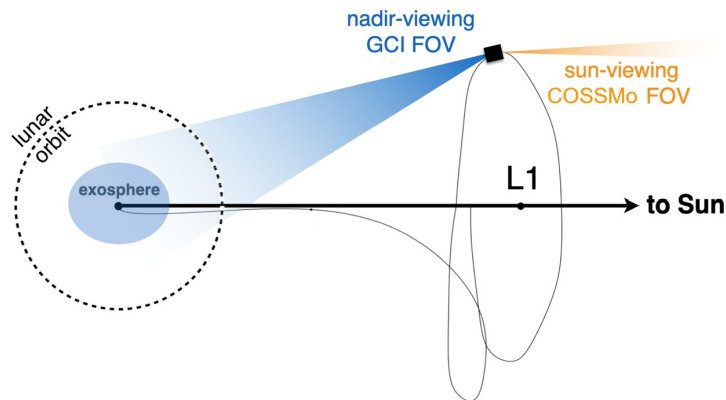


Figure 4: The planned trajectory of the mission. The L1 Lagrange point is defined as the stable orbital position between the Earth and the Sun where their gravitational pulls are balanced.




Figure 5: An overview of the GCI Instrument. The calibration algorithms are given the resulting output of digital numbers (on the very right) and tasked with recovering the original Ly- α signal (on the very left).




Figure 6: Optical Path of photons through the GCI. The black piece that the optical cone feeds into is the MCP + phosphor + taper + detector.

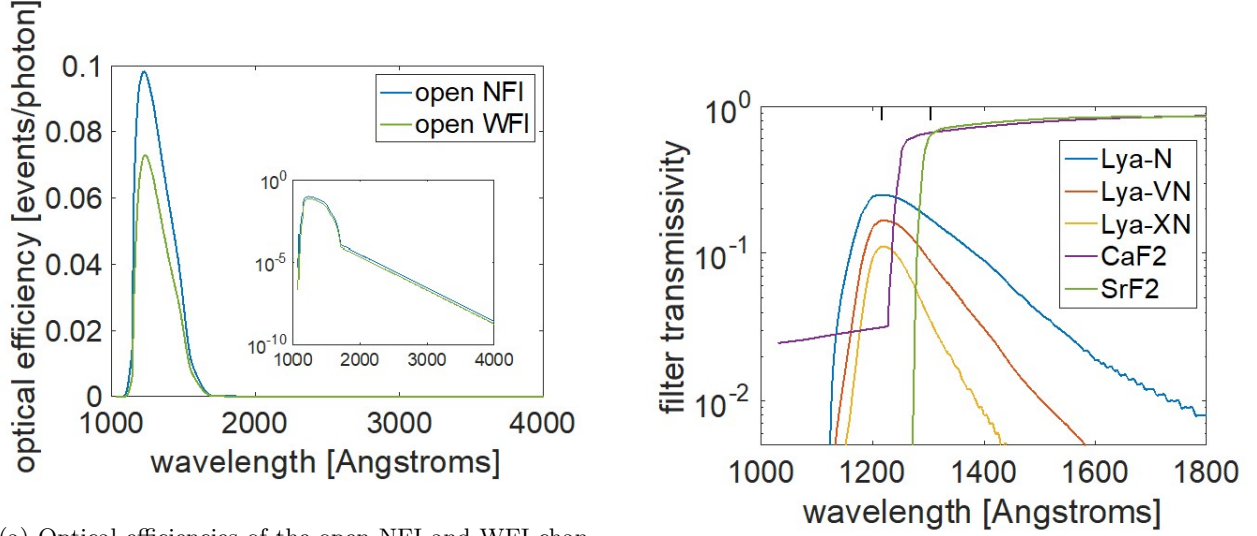
an energetic photoelectron, denoted as a photoelectron “event” in Figure 5. The product of the window transmissivity, mirror reflectivities, filter transmissivity and KBr quantum efficiency comprises the “optical efficiency”, which has units of [detected photons/incident photon], or equivalently, [events/incident photon]. The photoelectron is accelerated through a high voltage drop onto a microchannel plate (MCP), where it produces a cloud of secondary electrons. Note that the total number of secondary electrons released for each event is random. Therefore, the number of MCP electrons created per incident photoelectron event, known as the “MCP gain”, is random and its variance is an important source of noise in the measured signal, which is denoted “gain noise”. The MCP gain is denoted G_{MCP} , while the mean MCP gain is denoted g_{MCP} .

The MCP electrons in turn impinge on a phosphor screen. Phosphor detection of the MCP electrons releases visible wavelength photons, which are fed through a fiberoptic taper onto the APS detector, where they are converted into electron “counts”, which accumulate in pixel wells and are read out by the camera FPGA.

The spectral radiance of diffuse emission at Earth, which is denoted $L(\lambda, \hat{n})$, with units $\left[\frac{\text{photons}}{\text{cm}^2 \cdot \text{s} \cdot \text{sr} \cdot \text{nm}}\right]$, is detected by the GCI instrument. Different wavelengths have differing levels of incident photons, but the number of incident photons also depends on the line-of-sight (LOS) and look direction, denoted \hat{n} .

The WFI and NFI channels have optical efficiencies as a function of wavelength, denoted $\varepsilon(\lambda)$, with units $\left[\frac{\text{photoelectron events}}{\text{photon}}\right]$. Optical efficiency depends on wavelength-dependent mirror reflectivities, denoted $\Gamma(\lambda)$ and photocathode quantum efficiency (QE), denoted $\eta(\lambda)$. The two channels have a different number of mirrors, denoted N_{mir} , as seen in Figure 6; the NFI channel has two, while the WFI channel has four. The “open” WFI and NFI channels achieve wide-band spectral filtering through wavelength-dependent mirror reflectivities and component quantum efficiency. Long-wavelength transmissivity (>170 nm) increases as a function of KBr photocathode activation; deactivation via on-orbit LED illumination will be performed as needed.

Both the WFI and NFI provide further UV spectral targeting options via a 6-position filter wheel. Three filters, labelled N, VN, and XN, support the narrowband Ly- α transmission. The CaF₂ and SrF₂ filters suppress Ly- α transmission (and only transmit longer wavelengths) for out-of-band measurement and subtraction. The blocked filter suppresses all incident photons and is used for calibration purposes. The



(a) Optical efficiencies of the open NFI and WFI channels as a function of wavelength. The two subplots show the same data, but one is on a linear scale and the other is on a logarithmic scale.

(b) Filter transmissivities as a function of wavelength. The blocked filter has no transmissivity at any wavelength.

Figure 7

choice of which narrowband Ly- α filter to use for science operations will be made on-orbit during the IOC phase. The wavelength-dependent filter transmissivity is denoted as $\tau(\lambda)$.

Therefore, the formula for optical efficiency is as follows:

$$\varepsilon(\lambda) = \Gamma(\lambda)^{N_{\text{mir}}} \tau(\lambda) \eta(\lambda) \quad (1)$$

All photons from all wavelengths, over each APS pixel solid angle (Ω_{ij}) and aperture area (A), over the frame integration time (T_f), pass through these optics and filters and are detected as photoelectron events. The indices i and j index pixel row and column respectively and thus run from 0 to 2047, while k indexes the frame number. Therefore, the mean number of detected photons, in units of photoelectron events, per frame k and per APS pixel ij is

$$e_{\text{Earth},ijk} = T_f A \Omega_{ij} \int_0^\infty d\lambda \varepsilon(\lambda) L(\lambda, \hat{n}_{ij}) \quad (2)$$

Since this process counts number of photon arrivals, the true signal is actually “drawn” from a Poisson distribution with mean equal to $e_{\text{Earth},ijk}$. This true signal is denoted as $E_{\text{Earth},ijk}$. For this section and in the remainder of this thesis, the mean is generally denoted with lowercase letters, while the Poisson-distributed quantity (the signal actually detected) is denoted with uppercase letters. Naturally, converting from means to signals introduces shot noise.

However, Earth is not the only source of photons. Background photons arise from the diffuse interplanetary hydrogen (IPH), denoted L_{IPH} . This thesis assumes that these photons are monochromatic at the Ly- α emission line (121.6 nm), but the number of photons is still dependent on the LOS and boresight direction. These photons contribute some photoelectron events per APS pixel ij per frame k :

$$e_{\text{IPH},ijk} = T_f A \Omega_{ij} L_{\text{IPH}}(\hat{n}_{ij}) \varepsilon(121.6\text{nm}) \quad (3)$$

Other major sources of photon background include the Earth limb/disk FUV line emissions and MUV continuum albedo, bright EUV stars, and the lunar disk. These follow similar formula as the above, and are simply denoted as $e_{\text{background},ijk}$.

Finally, non-photon energetic particle radiation, which is part of the instrument background, can also generate photoelectron events at a mean rate denoted $q_{\text{events},ij}$. This value depends on the environment of

the spacecraft and is Poisson distributed. The radiation must be multiplied by the frame integration time in order to get the number of extra photoelectron events added per frame, denoted $E_{\text{radiation},ijk}$.

$$E_{\text{radiation},ijk} = T_f Q_{\text{events},ij} \quad (4)$$

The total number of photoelectron events is then the sum of the individual photoelectron events, which is given as

$$E_{ijk} = E_{\text{Earth},ijk} + E_{\text{IPH},ijk} + E_{\text{background},ijk} + E_{\text{radiation},ijk} \quad (5)$$

Not all pixels respond equally to the same brightness or the same number of incident UV photons. In order to account for this, the photoelectron events are multiplied by a flat-field, denoted F_{ij} , which effectively normalizes the responsivity of all pixels. This flat-field is assumed to be stable across frames, but can vary from pixel to pixel. Once detected as photoelectron events, the signal is gain-amplified and turned into electron-counts through a MCP.

Thermal electron current due to the thermal motion of the electrons within each channel manifests as additional electron counts per second, or “dark current”. The mean number of electron counts as a result of dark current (d_{ij}) accumulate linearly with respect to time, although the dark current itself depends on the temperature of the system. This process is also Poisson distributed. The dark current must be multiplied by the frame integration time in order to get the number of dark electron counts added per frame.

Finally, non-photon energetic particle radiation can also create extra electron counts per second, denoted $q_{\text{counts},ij}$. This value depends on the environment of the spacecraft and is Poisson distributed. The radiation must be multiplied by the frame integration time in order to get the number of extra electron counts added per frame. An energetic particle can either “create” extra electron counts or extra photoelectron events, but not both. Any particle radiation onto the MCP will undergo the intensifier gain and is modeled as extra photoelectron events, while the particle radiation directly onto the APS does not undergo intensifier gain and is modeled as electron counts.

The number of electron counts per frame k and APS pixel ij is then given as

$$C_{ijk} = E_{ijk} F_{ij} G_{\text{MCP}} + T_f D_{ij} + T_f Q_{\text{counts},ij} \quad (6)$$

The APS camera is then read-out by the camera FPGA using an Analog-to-Digital Converter (ADC) as $N_{\text{bin}} \times N_{\text{bin}}$ pixels, where $N_{\text{bin}} = 2$ for the NFI channel and $N_{\text{bin}} = 4$ for the WFI channel. These settings cause the NFI images to have dimensions of 1024×1024 , while WFI images have dimensions of 512×512 . For the remainder of the thesis, the dimensions are denoted N_{pix} , where $N_{\text{pix}} = 1024$ for NFI images and $N_{\text{pix}} = 512$ for WFI images.

The ADC has two gain modes to convert accumulated electrons (analog) into digital numbers (DN): Low Gain (LG) and High Gain (HG). The gain used is denoted G_{APS} . This process is also accompanied by read noise, which is a Gaussian distribution with zero mean. This thesis models read noise as additive noise R_{ij} , which scales with the square root of the number of pixels binned.

Finally, a voltage bias b_{ij} is applied before readout, adding more [DN]. Bias is an artificially induced electronic offset with a fixed pattern, so it does not contribute to noise at all.

Therefore, the signal, in [DN] per frame k and binned pixel ij is

$$S_{ijk} = \sum_{I=N_{\text{bin}}i}^{N_{\text{bin}}(i+1)-1} \sum_{J=N_{\text{bin}}i}^{N_{\text{bin}}(i+1)-1} C_{IJK} G_{\text{APS}} + b_{IJK} + R_{IJ} \quad (7)$$

Frames are stacked onboard before telemetry back to Earth via the Instrument Control Package (ICP). The ICP manages payload functions and is the single point electrical/command/data-interface to the spacecraft. The number of frames stacked, denoted N_{frames} , depends on the required integration time (T_{int} , in units of seconds) and the frame read-out rate. In the baseline operations plan, individual frame integration time, or T_f , is 0.125 seconds. NFI images are integrated over 30 minutes, while WFI images are integrated over 60 minutes.

$$N_{\text{frames}} = \frac{T_{\text{int}}}{T_f} \quad (8)$$

Therefore, the total measured signal, in [DN], per binned pixel ij , is given as

$$S_{ij} = \sum_{k=1}^{N_{\text{frames}}} S_{ijk} \quad (9)$$

1.3 Measurement Calibration Overview

The basic retrieval problem is to determine the global 3-dimensional H-density distribution from a time-series of 2-dimensional arrays of noisy, 1-dimensional measurements of Ly- α radiance from a distant vantage point. The basic measurement problem is to measure the Ly- α emission radiance over a wide field at high cadence and isolate it from photon or instrument backgrounds. This thesis deals mainly with the measurement problem, which is solved through the use of measurement calibration.

Raw images S_{ij} , which are called Level 0 data, need to be properly processed and calibrated before they are usable as science data, since they contain undesirable signals such as dark current, and stars. After distortion correction and registration, the images now become Level 1A data, which is geolocated and annotated. These images are then passed through the following calibration algorithms:

- Instrument Background Subtraction
- Particle Radiation Subtraction
- Out-of-band (OOB) Removal
- IPH Estimation and Removal
- Flat-Fielding
- Absolute Intensity Calibration
- Stellar Calibration (Not covered in this thesis)

The result of these calibration algorithms is denoted as Level 1B data, which are the geolocated, binned, and calibrated science images of Earth's Ly- α emission. The corresponding pipeline is shown in Figure 8. The Level 1B data is then passed through tomography and non-parametric analysis, as well as radiative transfer analysis in order to get Level 2 data, which is the final global 3-dimensional H-density distribution. This part focuses on the retrieval itself (going from two-dimensional images of one-dimensional measurements at a distant vantage point to a three-dimensional global H-density distribution), which will not be discussed in this thesis. The corresponding pipeline is shown in Figure 9.

To help with the calibration, the spacecraft will take images with certain settings at specific cadences in order to provide the necessary constraints on the spacecraft's camera and the environment, which will serve as valuable starting points and baselines for the algorithms. Details about the necessary cadences can be found in later sections.

1.3.1 Validation of Measurement Calibration

A synthetic image model will be used to validate some of the calibration algorithms. The model uses many input parameters, including observing target (boresight direction and spacecraft roll angle), observing conditions (date, solar activity, and particle radiation environment), and operational mode (channel, acquisition time, temperature, beginning-of-life versus end-of-life, and filter). The model then calculates the spectral radiance of the scene arising from the IPH, the Earth and its exosphere, the stars, and the moon. These photon sources are then input to the instrument model, whose output depends on the specified channel and filter.

To simulate shot noise, for each pixel, the model uses a random Poisson-distributed variable with mean equal to the expected value of the signal in the pixel. Finally, the instrument background signal counts are added. The instrument background signal model includes the solar energetic particle radiation, approximated in terms of Poisson-distributed events on the MCP and APS, and dark current, which is specific to the operational mode.

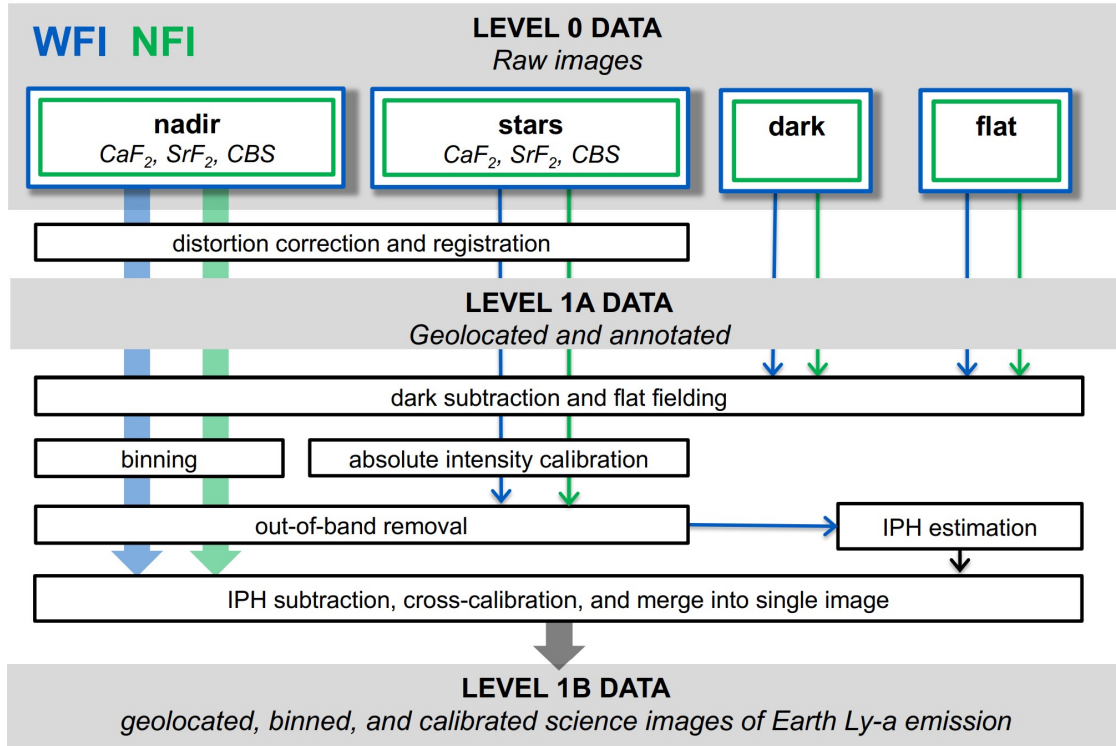


Figure 8: Calibration Pipeline from Level 0 data to Level 1B data. The main focus of this thesis deals with the transition from Level 1A data to level 1B data.

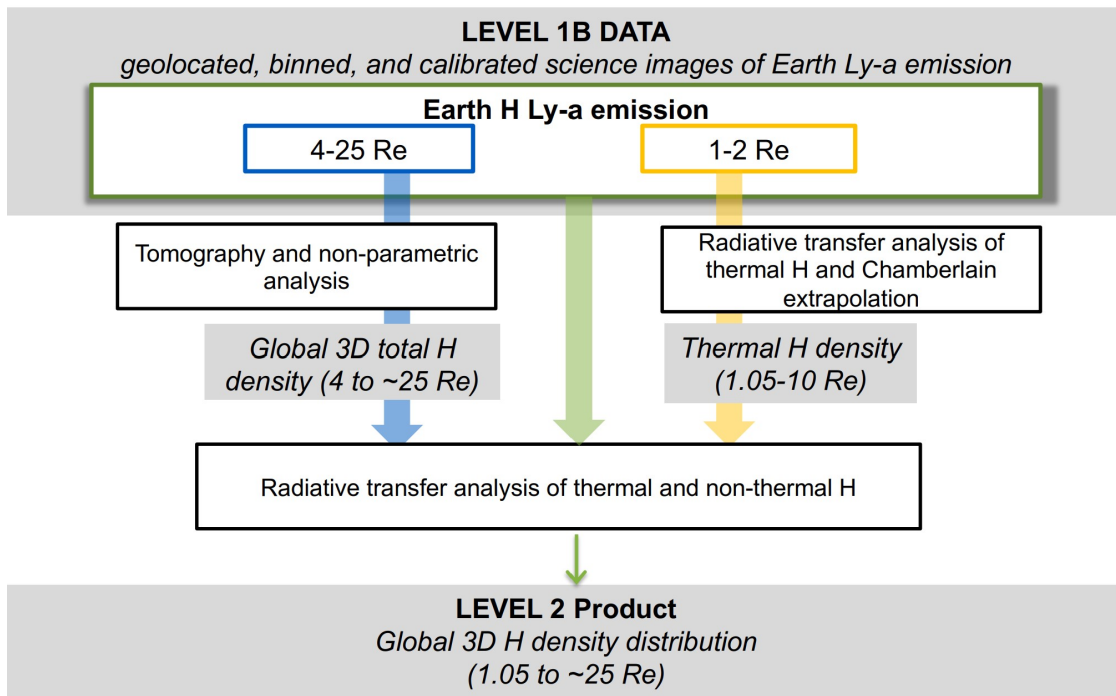


Figure 9: Calibration Pipeline from Level 1B data to Level 2 data. This is included for the sake of completion; the methods shown are not discussed in this thesis.

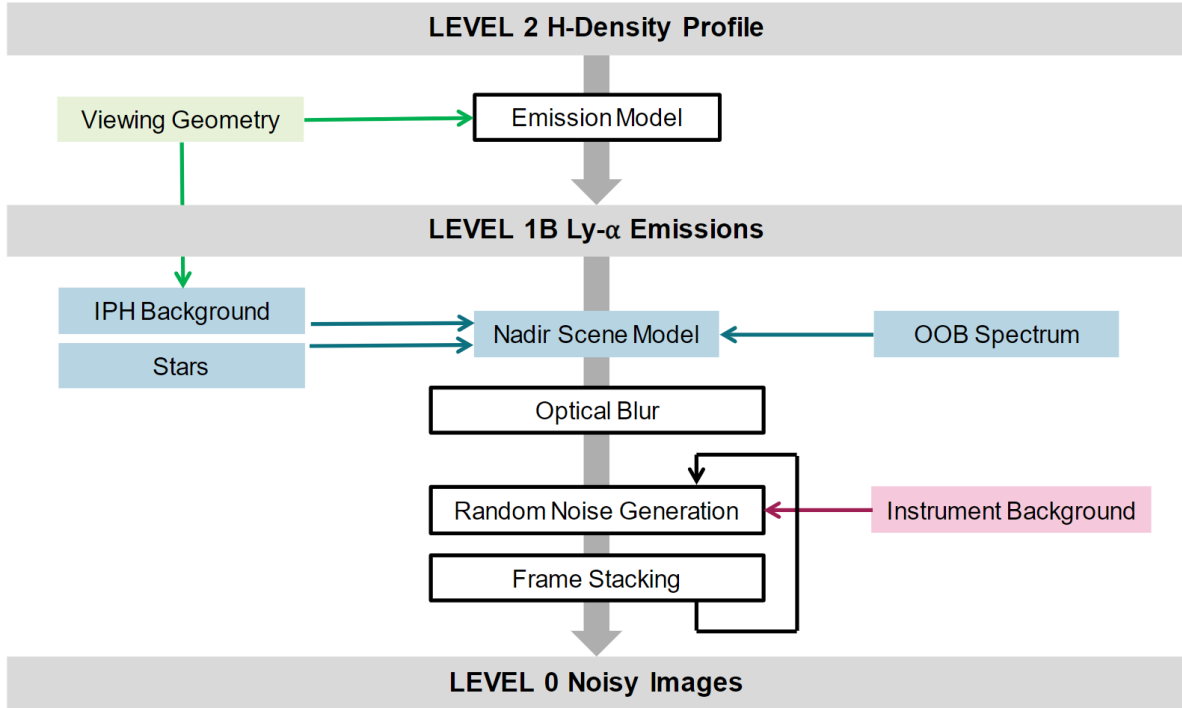


Figure 10: The synthetic pipeline, which goes in reverse order of the retrieval pipeline.

This process is repeated for multiple frames (the number of which depends on the specified integration time), which are then stacked in order to get the final image. Figure 10 shows the full synthetic image model pipeline with some extra details, omitted here for brevity.

In general, validation of the calibration algorithms will follow a common approach, summarized as follows:

1. Validation of a particular parameter retrieval algorithm begins with the specification of an assumed ground-truth value of the parameter that corresponds to a specific environmental/observing condition, which are informed by realistic physical expectations.
2. This assumed parameter value is input into the synthetic image model, which then yields noisy synthetic data. Depending on the validation objective, certain parts of the model may be turned off when creating synthetic data in order to improve speed and efficiency.
3. This synthetic data is then input into the ensemble of retrieval algorithms that are being validated. They will return a retrieved parameter value.
4. The retrieved parameter value and the original ground-truth are compared, and the residual error is computed.

The approach above is repeated over a comprehensive set of test cases, including environmental/exospheric conditions (season, local time, incident solar flux, solar cycle, geomagnetic activity, etc.) and observing/instrument conditions (spacecraft vantage, particle radiation environment, etc.). Each validation test case consists of a Monte Carlo ensemble of independent parameter retrievals, which then give an ensemble of residual errors. Validation success is then judged quantitatively based off the mean or median of the residual errors, which indicates retrieval accuracy, and the standard deviation of the residual errors, which indicates retrieval precision.

Most calibration algorithms have a required accuracy threshold, i.e. a maximum error threshold. The accuracy and precision produced by the validation tests are compared to the required accuracy threshold in order to determine suitability of the proposed calibration algorithm.

1.4 Organization of Thesis

In the introduction we have established the motivation for the development of the calibration algorithms found in this thesis and where they fit in the Carruthers mission. This thesis will devote one chapter to each of instrument background retrieval (including PATHS background retrieval), particle radiation background retrieval, IPH estimation and subtraction, and flat-field retrieval. In each chapter, more detailed background will be introduced, followed by the algorithm itself, followed by the validation of the algorithm and the results.

2 Instrument Background

2.1 Carruthers Camera Analysis

As described above, the instrument background signal consists of the bias and dark current. Dark current is a result of the thermal motion of electrons within each channel. It is expected to accumulate linearly with respect to time, while the slope of this buildup is expected to increase with an increase in the temperature of the system. Dark current is Poisson distributed, thus also adding noise, which is denoted dark noise.

The goal of the bias retrieval algorithm is to retrieve the bias frame from short exposure frames with also contain small amounts of dark signal as well as read noise. Finally, the dark retrieval algorithm will retrieve the dark current slope using data from exposure frames of varying lengths.

2.1.1 Engineering Model Camera Overview

Initial camera data for the Carruthers mission was taken at Space Dynamics Laboratory (SDL) in Utah. For each of 5 temperatures (+35C, +20C, +5C, -10C, -25C), bias, dark, and long-dark frames were taken in both HG and LG modes. In general, there were 10 single images and 1 stacked image of 480 frames for every setting for bias and dark frames. Only one long-dark frame was taken at each temperature and gain setting. All images are 1040×1024 pixels in size. Therefore, the top 8 rows of the top half and the top 8 rows of the bottom half of each image are removed. These were found to be the optically dark and electrically dark rows, and thus should not be part of this analysis.

Bias frames were taken with the shortest integration time possible, $65.6\mu\text{s}$, in order to minimize the effect of dark current. Dark frames were taken with an integration time of 124.97 ms, while long-dark frames were taken with an integration time of 17.193 seconds. The varying integration times allow for a proper recovery of dark current.

In general, it was found that the top and bottom halves of the camera behave very differently, due to the design of the camera.

2.1.2 Bias Frames

All images are 1024×1024 pixels in size after the optically and electrically dark rows are removed. Let i be the row index, j be the column index, and k be the frame index, where indices are 1-indexed. Thus, i and j range from 1 to 1024, while k ranges from 1 to N , where N is the number of images available (usually 10). Let B_{ijk} denote the pixel in the i 'th row, the j 'th column, and the k 'th frame in the raw bias images. Similar notation follows for D_{ijk} for dark and DL_{ijk} for long dark. The stacked image is denoted by replacing the subscript k with the subscript s , like B_{ijs} . There is only one stacked image for every setting, so it is unnecessary to carry around the k index.

According to the camera specifications, top and bottom halves are read out by the FPGA separately, while columns within each half have the same bias voltage applied at readout. Therefore, individual bias images undergo column-averaging, where the same column across all images are averaged to produce one true value. Figure 11 shows a visual representation of the pixels that are averaged for each column.

The standard deviation of each column is also taken. This represents read noise. Let μ_{ij} denote the mean found for pixel ij , while σ_{ij} denotes the standard deviation found for pixel ij . Mathematically, column-averaging uses the following equations:

$$\mu_{ij} = \begin{cases} \frac{1}{512N} \sum_{k=1}^N \sum_{m=1}^{512} B_{mjk} & \text{if } 1 \leq i \leq 512 \text{ (top half)} \\ \frac{1}{512N} \sum_{k=1}^N \sum_{m=513}^{1024} B_{mjk} & \text{if } 513 \leq i \leq 1024 \text{ (bottom half)} \end{cases} \quad (10)$$

$$\sigma_{ij} = \begin{cases} \sqrt{\frac{\sum_{k=1}^N \sum_{m=1}^{512} (B_{mjk} - \mu_{ij})^2}{512N}} & \text{if } 1 \leq i \leq 512 \text{ (top half)} \\ \sqrt{\frac{\sum_{k=1}^N \sum_{m=513}^{1024} (B_{mjk} - \mu_{ij})^2}{512N}} & \text{if } 513 \leq i \leq 1024 \text{ (bottom half)} \end{cases} \quad (11)$$

In order to be able to report a single value for each temperature/gain mode and identify trends, the average of μ_{ij} and σ_{ij} is taken over all ij to return a unique value. The mean of μ_{ij} is denoted as the Mean of Means, while the mean of σ_{ij} is denoted as the Mean of Stds. This process is then repeated for the stacked

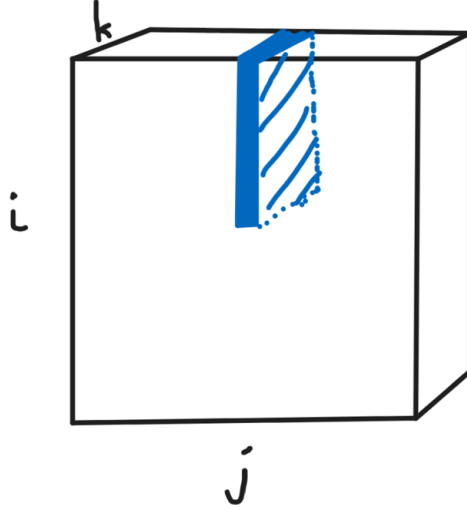


Figure 11: The column-averaging step takes in a 3D datacube of bias images, shown as the black volume, and returns two 2D images μ and σ . The mean and standard deviation of all pixels in the blue volume is taken for the specified column.

Table 1: Bias Mean of Means

Temperature	LG [DN]	HG [DN]
+35C	340.10	252.51
+20C	338.31	242.30
+5C	336.22	237.82
-10C	334.57	235.76
-25C	332.79	231.50

Table 2: Bias Mean of Stds

Temperature	LG [DN]	HG [DN]
+35C	1.14	8.68
+20C	1.13	8.12
+5C	1.12	7.74
-10C	1.11	7.52
-25C	1.10	7.48

image. Tables 1, 2, 3, and 4 report the Mean of Means and Mean of Stds for all temperatures and gain modes.

In order to make sure that the camera is well-behaved, hot pixels, defined as pixels in the column-averaged bias that are at least 3 standard deviations above the mean of the bias, are counted. The top half and bottom halves (labelled top and bottom in the tables) are treated separately. Tables 5, 6, 7, and 8 report the percentage of hot pixels for all temperatures and gain modes.

In general, the mean bias values increase slightly with increase in temperature. Figures 12a and 12b show the percent difference between biases at +35C and -25C. HG bias tends to be more temperature dependent, but this is likely due to larger dark current at high temperatures rather temperature dependence of the actual voltage bias.

In order to minimize dark current, the lowest temperature bias was taken as the actual bias. One bias image was produced for each APS gain mode, denoted b_{ij} . Figures 13a and 13b show the recovered biases.

Table 3: Stacked Bias Mean of Means

Temperature	LG [DN]	HG [DN]
+35C	340.06	252.49
+20C	338.32	242.36
+5C	336.18	237.87
-10C	334.49	235.86
-25C	332.76	231.59

Table 4: Stacked Bias Mean of Stds

Temperature	LG [DN]	HG [DN]
+35C	0.16	4.29
+20C	0.13	2.92
+5C	0.13	2.27
-10C	0.14	2.00
-25C	0.16	1.94

Table 5: LG Bias percentage of pixels that are at least 3 standard deviations above the mean.

Temperature	Top [%]	Bottom [%]
+35C	0.30	0.00
+20C	0.30	0.00
+5C	0.30	0.00
-10C	0.30	0.00
-25C	0.30	0.00

Table 6: LG Stacked Bias percentage of pixels that are at least 3 standard deviations above the mean.


Temperature	Top [%]	Bottom [%]
+35C	0.30	0.00
+20C	0.30	0.00
+5C	0.30	0.00
-10C	0.30	0.00
-25C	0.30	0.00

Table 7: HG Bias percentage of pixels that are at least 3 standard deviations above the mean.

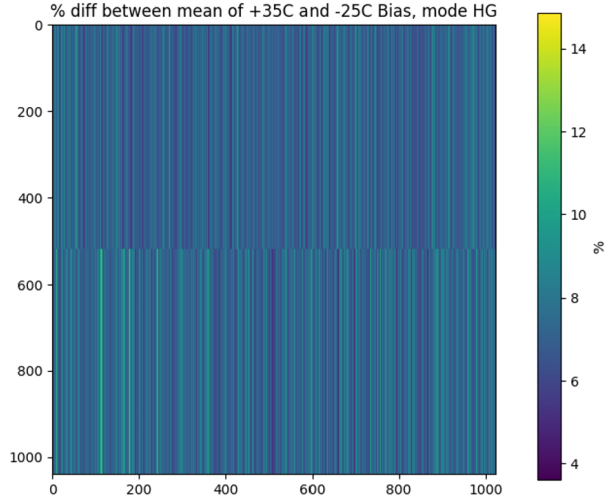
Temperature	Top [%]	Bottom [%]
+35C	0.40	0.10
+20C	0.40	0.10
+5C	0.40	0.10
-10C	0.30	0.10
-25C	0.30	0.10

Table 8: HG Stacked Bias percentage of pixels that are at least 3 standard deviations above the mean.

Temperature	Top [%]	Bottom [%]
+35C	0.40	0.10
+20C	0.40	0.10
+5C	0.40	0.10
-10C	0.40	0.10
-25C	0.30	0.10



(a) LG Bias percent difference between +35C and -25C



(b) HG Bias percent difference between +35C and -25C

Figure 12: Temperature dependence of Bias Frames

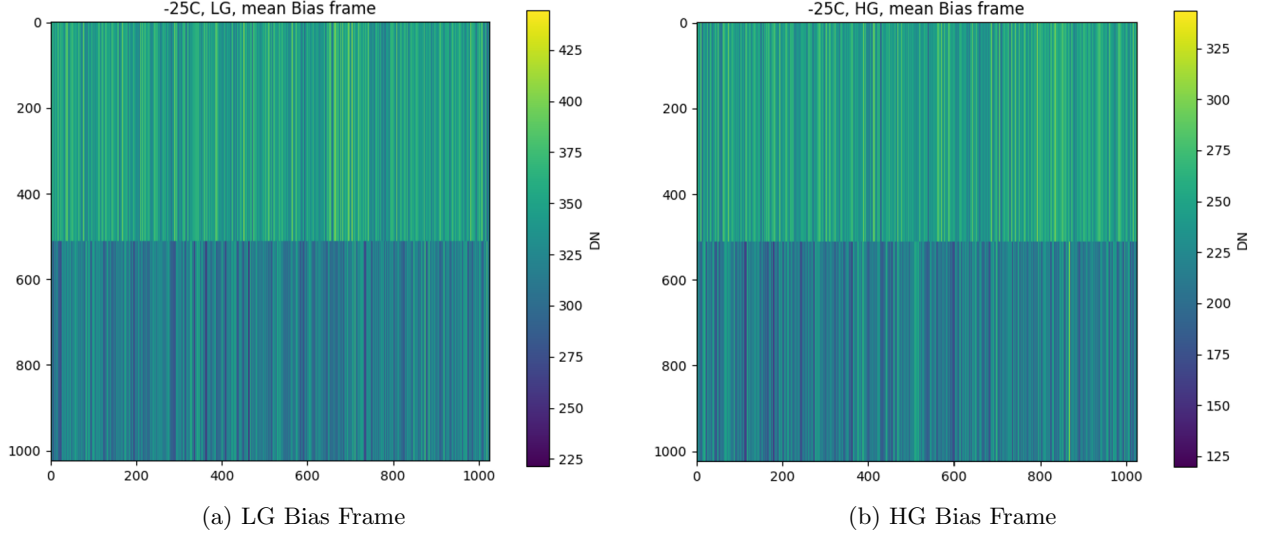


Figure 13: Retrieved Bias Frames

Table 9: Dark Mean of Means

Temperature	LG [DN]	HG [DN]
+35C	0.56	20.34
+20C	-0.02	1.41
+5C	-0.05	-0.95
-10C	-0.22	-1.68

Table 10: Dark Mean of Stds

Temperature	LG [DN]	HG [DN]
+35C	1.08	9.15
+20C	1.04	6.76
+5C	1.03	6.21
-10C	1.01	5.98

These are the biases that are used in the synthetic image model.

2.1.3 Dark Current

In order to calculate dark current, the bias image found in the previous section was subtracted from all individual dark images given (for both dark and long-dark). For the stacked dark image, the bias was multiplied by 480 before being subtracted off, since each frame contributes a bias. The mean and standard deviation of each pixel over all images given is calculated at each temperature and gain mode. Note that if there is only one image, then it is impossible to calculate a standard deviation. In this case, σ_{ij} is set to 0 for all ij . The formulae below use the dark images, but the exact same formulae are used to process the long-dark images.

$$\mu_{ij} = \frac{1}{N} \sum_{k=1}^N (D_{ijk} - b_{ij}) \quad (12)$$

$$\sigma_{ij} = \begin{cases} 0 & \text{if } N = 1 \\ \sqrt{\frac{\sum_{k=1}^N (D_{mjk} - b_{ij} - \mu_{ij})^2}{N}} & \text{else} \end{cases} \quad (13)$$

The mean of both μ_{ij} and σ_{ij} are then taken to represent the mean dark and read noise at the specific temperature and gain modes, denoted Mean of Means and Mean of Stds respectively. Note that $-25C$ is not considered because the data for that temperature was unavailable. Tables 9, 10, and 11 contain the Mean of Means and Mean of Stds. Note that there is no table for the Mean of Stds of the stacked image because there is only one stacked image.

These values should all be positive due to the presence of dark current. However, cold dark images have lower signal values than the bias frame data, yielding negative counts when the bias is subtracted. Tables 12 and 13 show the percentage of pixels that are still non-negative in the resulting dark image. From the

Table 11: Stacked Dark Mean of Means

Temperature	LG [DN]	HG [DN]
+35C	0.63	20.27
+20C	0.02	1.48
+5C	-0.03	-0.94
-10C	-0.15	-1.62

Table 12: Percentage of non-negative pixels in averaged dark image after bias subtraction

Temperature	LG [%]	HG [%]
+35C	79.33	99.97
+20C	48.36	56.25
+5C	45.07	33.33
-10C	32.80	25.85

Table 13: Percentage of non-negative pixels in the stacked dark image after bias subtraction

Temperature	LG [%]	HG [%]
+35C	86.37	99.99
+20C	50.93	60.58
+5C	44.30	28.97
-10C	31.95	21.03

tables, most values of bias-subtracted dark data are negative at cold temperatures. We conclude that the integration time of these frames is insufficient to detect dark signal above read noise level.

The hot-pixel calculations is repeated here as well for completion, contained in Tables 14, 15, 16, and 17.

The long-dark statistics are displayed in Tables 18 and 19. Note that -25C is still not considered because the data for that temperature was unavailable. The data did not have a stacked long-dark frame, so those tables are removed. Finally, only HG mode at $+20\text{C}$ and -10C had multiple images, so these are the only settings that have a Mean of Stds.

In general, the long-dark data appears to not be ideal. Tables 20 and 21 contain the top-half versus bottom-half means. The difference between the two is unexpected and appears to be far too large to be explained by just dark current alone.

Finally, there appears to be a potentially problematic number of hot pixels, as shown in Tables 22 and 23, especially at high temperatures in the bottom half of images read out using the HG mode.

To calculate the dark current, the top and bottom halves are treated separately. In order to reduce margins of error, the stacked dark frames are used in the subsequent calculations. A line was interpolated from the two (time, DN) coordinates, one from the dark frame and one from the long-dark frame, for each half at each temperature and gain. The slope is the dark current at that temperature and gain setting, with units of DN/sec. The y-intercept is the DN count at $t = 0$, which should ideally be 0. Tables 24 and 25 show the results. In these tables, Intercept is abbreviated as Int.

The intercepts seem relatively well-behaved, but negative dark current is unexpected. It is unclear whether this is a result of the data-collection process or a true property of the camera; more data is required to come to a conclusion.

Table 14: LG Dark percentage of pixels that are at least 3 standard deviations above the mean.

Temperature	Top [%]	Bottom [%]
+35C	0.14	0.13
+20C	0.097	0.086
+5C	0.18	0.23
-10C	0.33	0.21

Table 15: LG Stacked Dark percentage of pixels that are at least 3 standard deviations above the mean.

Temperature	Top [%]	Bottom [%]
+35C	0.14	0.13
+20C	0.11	0.096
+5C	0.25	0.17
-10C	0.45	0.42

Table 16: HG Dark percentage of pixels that are at least 3 standard deviations above the mean.

Temperature	Top [%]	Bottom [%]
+35C	0.34	0.32
+20C	0.22	0.17
+5C	0.31	0.26
-10C	0.81	0.79

Table 18: Long Dark Mean of Means

Temperature	LG [DN]	HG [DN]
+35C	70.46	1724.27
+20C	6.61	213.06
+5C	0.98	27.7
-10C	0.14	4.54

Table 20: LG Long Dark Mean of Means by Halves

Temperature	Top [DN]	Bottom [DN]
+35C	-1.12	142.03
+20C	0.36	12.86
+5C	-0.065	2.02
-10C	-0.096	0.38

Table 22: LG Long Dark percentage of pixels that are at least 3 standard deviations above the mean.

Temperature	Top [%]	Bottom [%]
+35C	0.20	0.37
+20C	0.23	0.17
+5C	0.17	0.15
-10C	0.15	0.12

Table 24: LG Dark Current. Int. stands for intercept, and Bot. stands for bottom.

Temperature	Top Slope [DN/sec]	Top Int. [DN]	Bot. Slope [DN/sec]	Bot. Int. [DN]
+35C	-0.092	0.46	8.27	-0.22
+20C	0.035	-0.25	0.74	0.20
+5C	0.0023	-0.10	0.12	0.026
-10C	0.013	-0.33	0.021	0.015

Table 25: HG Dark Current. Int. stands for intercept, and Bot. stands for bottom.

Temperature	Top Slope [DN/sec]	Top Int. [DN]	Bot. Slope [DN/sec]	Bot. Int. [DN]
+35C	-3.85	20.55	203.53	-4.97
+20C	-1.46	1.42	26.25	-1.56
+5C	0.017	-0.87	3.34	-1.44
-10C	0.12	-1.65	0.61	-1.68

Table 17: HG Stacked Dark percentage of pixels that are at least 3 standard deviations above the mean.

Temperature	Top [%]	Bottom [%]
+35C	0.34	0.32
+20C	0.22	0.17
+5C	0.30	0.24
-10C	0.64	0.61

Table 19: Long Dark Mean of Stds

Temperature	LG [DN]	HG [DN]
+35C	-	-
+20C	-	14.19
+5C	-	-
-10C	-	4.56

Table 21: HG Long Dark Mean of Means by Halves

Temperature	Top [DN]	Bottom [DN]
+35C	-45.68	3494.23
+20C	-23.72	449.83
+5C	-0.57	55.98
-10C	0.34	8.74

Table 23: HG Long Dark percentage of pixels that are at least 3 standard deviations above the mean.

Temperature	Top [%]	Bottom [%]
+35C	0.21	1.14
+20C	0.17	1.89
+5C	1.03	1.59
-10C	1.06	0.55

2.2 PATHS Detector Analysis

2.2.1 PATHS Overview

The results in the previous section were abnormal, as the dark current was found to be highly nonlinear and sometimes negative. The same approach for bias and dark current recovery were used on the PATHS cameras in order to make sure the instrument background retrieval algorithms work.

PATHS is a ground-based array of three sets of photometers. The target observable is Balmer- α . PATHS has certain advantages over the Carruthers mission, including ease-of-deployment and cost. However, ground-based observation only supports localized H density retrieval. PATHS uses three sets of photometers, all pointed at the same location in the sky, to alleviate this problem.

PATHS uses three sets of Dual Channel Airglow Photometers, which is composed of the optical telescope, a camera, and a mount/drive system. This photometer system is specially designed for extremely low-level light detection suitable for isolating feeble nighttime airglow emissions from the sky background. The system includes two identical photometers, driven by the same altitude-azimuth drive. One photometer is configured with an “on-band” filter while the other is configured with an “off-band” filter, allowing for easy background subtraction with no other movement of optics or other dispersive elements. This permits robust, long-duration measurement regimens in harsh environments.

The camera layout for PATHS is slightly different than the one used in the GCI. PATHS contains two identical channels, which are labelled Camera 1 and Camera 2. The bias and dark images for PATHS were taken in a lab under dark conditions in the ECEB. Each image for both cameras has dimensions of 2199×2749 pixels.

We collected 100 bias images at temperature 0C with an exposure time of 1ms, then collected 10 dark images for each temperature/exposure time setting combination. The temperatures used were 0C, 10C, and 20C, while the exposure times used were 1s, 10s, 30s, 60s, 120s, 210s, and 300s. We only used one temperature/exposure time setting combination to take the bias frame since it is unnecessary to take more.

2.2.2 PATHS Bias Frame

Let i be the row index, j be the column index, and k be the frame index, where indices are 1-indexed. Thus, i ranges from 1 to 2199, j ranges from 1 to 2749, and k ranges from 1 to N , where N is the number of images available (100 for bias, or 10 for dark images). Let B_{ijk} denote the pixel in the i 'th row, the j 'th column, and the k 'th frame in the raw bias images. Similar notation follows for D_{ijk} for dark.

The top and bottom halves are no longer treated separately, so the column-averaging process changes slightly compared to Equations 10 and 11. Let μ_{ij} denote the mean found for pixel ij , while σ_{ij} denotes the standard deviation found for pixel ij . Mathematically, column-averaging for PATHS uses the following equations:

$$\mu_{ij} = \frac{1}{2199N} \sum_{k=1}^N \sum_{m=1}^{2199} B_{mjk} \quad (14)$$

$$\sigma_{ij} = \sqrt{\frac{\sum_{k=1}^N \sum_{m=1}^{2199} (B_{mjk} - \mu_{ij})^2}{2199N}} \quad (15)$$

As usual, μ_{ij} and σ_{ij} are averaged over all possible ij in order to return one unique value. The mean of μ_{ij} is denoted as the Mean of Means, while the mean of σ_{ij} is denoted as the Mean of Stds. Finally, the hot-pixel calculation is repeated.

For camera 1, the mean of means was 339.05 DN, the mean of stds was 27.35 DN, and 1.24% of pixels were at least three standard deviations above the mean. For camera 2, the mean of means was 337.51 DN, the mean of stds was 26.00 DN, and 1.56% of pixels were at least three standard deviations above the mean. In general, the read noise and number of hot pixels appear to be quite high, but both cameras are consistent. Figures 14a and 14b show the recovered biases for both cameras.




Figure 14

Table 26: PATHS Camera 1 Dark Images Mean of Means.

Temperature	1s [DN]	10s [DN]	30s [DN]	60s [DN]	120s [DN]	210s [DN]	300s [DN]
+0C	-1.47	-0.955	-0.418	0.182	0.579	1.67	2.06
+10C	64.21	62.61	64.66	65.61	68.45	70.92	70.634
+20C	263.16	268.53	261.45	271.13	279.70	288.06	294.36

2.2.3 PATHS Dark Current

In order to calculate dark current, the bias image found in the previous section was subtracted from all dark images. The mean and standard deviation of each individual pixel over all images given is then taken, at each temperature and gain mode. Equations 12 and 13 from Section 2.1.3 can be used here.

Finally, the mean of means, mean of stds, number of non-negative pixels, and number of hot pixels calculations are all repeated. Tables 26, 27, 28, and 29 contain the data for Camera 1 and tables 30, 31, 32, and 33 contain the data for Camera 2. In general, noise (mean of stds) increases as temperature increases, which is expected. Most pixels remain positive after bias subtraction, except for +0C. There are very few hot pixels. Finally, data from both cameras appear to be consistent with each other. These results are all satisfactory.

Each temperature setting has seven (integration time, DN) coordinates, one for each exposure time. A linear fit was performed on these coordinates in order to recover the slope and the y-intercept. The slope is then the dark current at that temperature, with units of DN/sec. The y-intercept is the DN count at $t = 0$, which should ideally be 0. Tables 34 and 35 show the results for both cameras. The intercepts appear to be problematic, although the problem is consistent across both cameras.

Figures 15a and 15b show plots containing the datapoints and the linear fit of the dark current as a function of exposure time. The dark current is highly linear, as expected.

Finally, looking at the percentage of negative pixels that are non-negative after bias subtraction at 0C (Tables 28 and 32), the data remaining is just read noise, with half of the pixels being above 0 and half of the pixels being below 0. This, coupled with a very low dark current slope, imply that dark current can be

Table 27: PATHS Camera 1 Dark Images Mean of Stds.

Temperature	1s [DN]	10s [DN]	30s [DN]	60s [DN]	120s [DN]	210s [DN]	300s [DN]
+0C	23.00	23.02	23.06	23.09	23.10	23.20	23.22
+10C	27.89	27.39	27.50	27.57	27.77	27.85	28.00
+20C	38.68	37.19	37.33	37.52	37.70	37.96	38.18

Table 28: PATHS Camera 1 Dark Images percentage of pixels that are non-negative after bias subtraction.

Temperature	1s [%]	10s [%]	30s [%]	60s [%]	120s [%]	210s [%]	300s [%]
+0C	45.32	46.81	47.90	49.18	49.58	51.77	52.11
+10C	92.61	92.44	92.67	93.07	93.65	94.04	94.35
+20C	99.76	99.79	99.81	99.90	99.91	99.96	99.98

Table 29: PATHS Camera 1 Dark Images percentage of pixels that are at least three standard deviations above the mean.

Temperature	1s [%]	10s [%]	30s [%]	60s [%]	120s [%]	210s [%]	300s [%]
+0C	0.0315	0.00693	0.00766	0.00845	0.00945	0.0103	0.0111
+10C	0.00167	0.00731	0.00873	0.00966	0.0116	0.0440	0.0480
+20C	0.00137	0.00763	0.00954	0.0212	0.0481	0.0532	0.0562

Table 30: PATHS Camera 2 Dark Images Mean of Means.

Temperature	1s [DN]	10s [DN]	30s [DN]	60s [DN]	120s [DN]	210s [DN]	300s [DN]
+0C	-0.630	-0.850	-0.49	0.209	0.735	1.75	2.27
+10C	67.63	67.30	67.98	68.52	71.73	72.90	73.27
+20C	271.38	272.44	264.28	276.10	282.65	289.49	296.85

Table 31: PATHS Camera 2 Dark Images Mean of Stds.

Temperature	1s [DN]	10s [DN]	30s [DN]	60s [DN]	120s [DN]	210s [DN]	300s [DN]
+0C	20.96	20.90	20.92	20.95	20.99	21.07	21.10
+10C	25.87	25.52	25.62	25.61	25.77	25.81	26.08
+20C	35.73	35.37	36.12	35.89	35.90	36.07	36.43

Table 32: PATHS Camera 2 Dark Images percentage of pixels that are non-negative after bias subtraction.

Temperature	1s [%]	10s [%]	30s [%]	60s [%]	120s [%]	210s [%]	300s [%]
+0C	47.92	47.03	47.49	48.99	49.72	51.32	51.92
+10C	91.88	91.75	91.98	92.39	92.83	93.17	93.44
+20C	99.72	99.74	99.77	99.89	99.88	99.92	99.97

Table 33: PATHS Camera 2 Dark Images percentage of pixels that are at least three standard deviations above the mean.

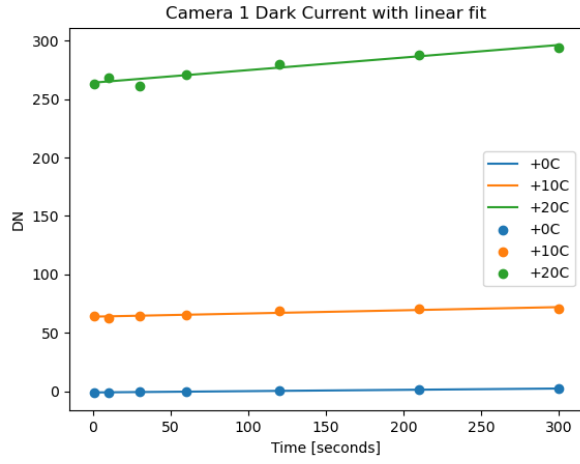
Temperature	1s [%]	10s [%]	30s [%]	60s [%]	120s [%]	210s [%]	300s [%]
+0C	0.00920	0.00412	0.00845	0.00910	0.0107	0.0114	0.0123
+10C	0.00164	0.00769	0.00945	0.0109	0.0127	0.0396	0.0471
+20C	0.00137	0.00817	0.0104	0.0192	0.0466	0.0514	0.0547

Table 34: PATHS Camera 1 Dark Current


Temperature	Slope [DN/sec]	Intercept [DN]
+0C	0.0111	-0.927
+10C	0.0272	63.89
+20C	0.108	263.97

Table 35: PATHS Camera 2 Dark Current

Temperature	Slope [DN/sec]	Intercept [DN]
+0C	0.0106	-0.679
+10C	0.0220	67.61
+20C	0.0949	269.12



(a) PATHS Camera 1 Dark Current datapoints and linear fit



(b) PATHS Camera 2 Dark Current datapoints and linear fit

Figure 15

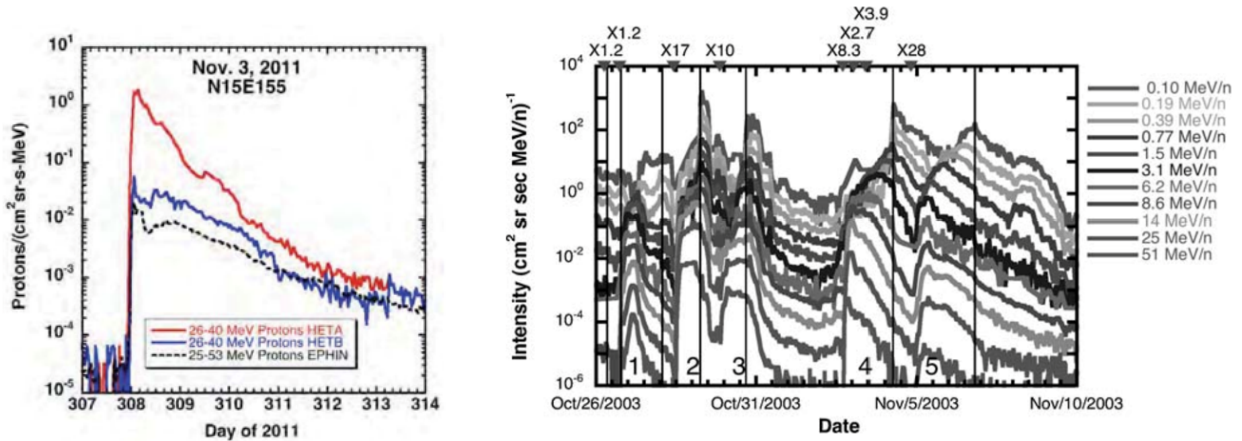
almost ignored for these cameras, which is very convenient.

3 Particle Radiation

The National Oceanic and Atmospheric Administration (NOAA) categorizes solar radiation storms using the NOAA Space Weather Scale, which goes from S1 to S5. The scale is based on measurements of energetic protons taken by the Geostationary Operational Environmental Satellite (GOES) in geosynchronous orbit. The start of a solar radiation storm is defined as the time when the flux of protons at energies ≥ 10 MeV equals or exceeds 10 proton flux units (pfu) ($1 \text{ pfu} = 1 \text{ particle} \cdot \text{cm}^{-2} \cdot \text{s}^{-1} \cdot \text{ster}^{-1}$). The end of a solar radiation storm is defined as the last time when the flux of ≥ 10 MeV protons is measured at or above 10 pfu. This definition allows multiple injections from flares and interplanetary shocks to be considered by a single solar radiation storm. Such storms can persist for time periods ranging from hours to days.

Radiation storms can affect the science data by adding extra photon events in two ways: those created in the MCP that undergo an intensifier gain and those created directly in the APS silicon, which do not undergo an intensifier gain. The use of stochastic modeling requires that these two types of radiation background should be treated differently.

These storms affect the frequency required to collect dark images for calibration purposes, since low cadences can be insufficient during high solar radiation activity, where the high-frequency changes in solar energetic particle (SEP) levels will affect the science data by adding extra background signal. This section analyzes the expected SEP level for the worst-case storm that the Carruthers mission will encounter and discusses a proposed solution to the aforementioned issue.



(a) Fluence spectra from STA, STB, and near-Earth spacecraft show a wide range of spectral indices [3]. (b) Time history of energetic oxygen nuclei measured by the SIS and ULEIS instruments on ACE during the Halloween 2003 series of SEP events [4]. Five large SEP events are evident, numbered 1 to 5. Interplanetary shocks are indicated by vertical solid lines and X-class flares are shown along the top. An intensity of 10^3 is classified as an S3 storm, with decreasing powers of 10 are classified as decreasing levels on the NOAA Space Weather Scale ($10^2 = S2$, etc).

Figure 16: Examples of SEP levels during the November 2011 and Halloween 2003 storms.

3.1 Initial Particle Flux Conditions

A Monte-Carlo simulation for cosmic ray effects on microelectronics (CREME-MC) was used to calculate how many particles would actually land on the MCP and APS and be detected as background signal. The spacecraft is shielded by 5mm of aluminum. Thus, the simulation gave the total integrated particle flux during solar quiet and during an S4 flare event (worst-case) for unshielded microelectronics and microelectronics shielded by 5mm of aluminum.

The simulation yielded the data in Table 36 for the total integrated particle flux ($Z = [1 : 28]$) over energies from 100 keV/nuc to 10^5 MeV/nuc. Note that the “flare” events are assumed to be an S4-equivalent storm; all values can be scaled by a factor of $\frac{1}{10}$ in order to yield S3-equivalent rates.

Table 36: Total integrated particle flux through different levels of shielding in different conditions

Level of Shielding	Condition	Total Integrated Particle Flux
Unshielded	Solar Quiet	4.76 particles/second/cm ²
Unshielded	S4 Flare	9.34×10^4 particles/second/cm ²
5mm Al Shielding	Solar Quiet	4.7 particles/second/cm ²
5mm Al Shielding	S4 Flare	2.23×10^3 particles/second/cm ²

Only nucleons with energies between 1 and 10^5 MeV were considered for the unshielded data, while nucleons with energies between 0.1 and 10^5 MeV were considered for the 5mm of aluminum shielding data. The particle flux through 5mm of aluminum shielding during quiet time will be denoted ϕ_q , while the particle flux through 5mm of aluminum shielding during S4 flare will be denoted ϕ_{S4} .

3.2 MCP-Based Radiation Background

For the MCP-based radiation background, the shielded particle flux that hits the MCP will create an event on the MCP. However, the same particle also creates secondary electrons due to the kinetic energy carried by the incoming particle. For the analysis that follows, a 100% MCP detection efficiency for incident particles and secondaries is assumed; this is a conservative estimate since a lower efficiency implies that cascade isn't triggered for all incident particles, implying fewer events. There is one primary event per particle from the particle flux ϕ calculated above. The secondaries are created at a rate of 20 per incident particle, which is also a conservative estimate; in reality there are usually around 10 secondaries per incident particle.

Once detected as photoelectron events, these particles are gain-amplified through a MCP, such that each photoelectron event generates a cloud of lower energy electrons. These electrons are then incident on a phosphorous screen, which generates visible photons, which are fed through an optical taper onto a detector.

The detector then accumulates the detected signal as electron counts, which accumulate in pixel wells. The scalar multiplier representing the number of electron counts per event is the mean gain and is denoted as G_{MCP} and equals roughly 30000.

Thus, the number of MCP events/sec/cm², denoted MCP_{event} is simply

$$MCP_{event} = (1 + 20)\phi \quad (16)$$

The total MCP-based particle/radiation event rate in quiet time is 105 [events/sec/cm²], while the event rate during an S4 flare is 46×10^3 [event/sec/cm²]. This value can then be converted to [electrons/second/NFI bin] using the following:

$$MCP_{event} \left[\frac{\text{event}}{\text{s cm}^2} \right] \times A \left[\frac{\text{cm}^2}{\text{APS pixel}} \right] \times (1.88)^2 \left[\frac{\text{APS pixel}}{\text{MCP pixel}} \right] \times N_{bin} \left[\frac{\text{MCP pixel}}{\text{bin}} \right] \times G_{MCP} \left[\frac{\text{electrons}}{\text{event}} \right] \quad (17)$$

In the equation above, each APS pixel covers $A = (6.5 \times 10^{-4})^2 [\text{cm}^2]$ area, while the optical taper introduces the 1.88^2 factor. Given the particle flux ϕ as input with units of $\left[\frac{\text{particles}}{\text{s-cm}^2} \right]$ and using $N_{bin} = 4$ for the NFI case, the total factor is $3.763 \left[\frac{\text{electrons}}{\text{particle}} \frac{\text{cm}^2}{\text{NFI bin}} \right]$. The resulting units will be $\left[\frac{\text{electrons}}{\text{s-NFI bin}} \right]$. Calculations for WFI bins are similar, except $N_{bin} = 16$, making the overall factor $15.05 \left[\frac{\text{electrons}}{\text{particle}} \frac{\text{cm}^2}{\text{WFI bin}} \right]$ instead.

Table 37 shows the radiation background for NFI and WFI in quiet time and during an S4 flare event, calculated using Equation 17. Note that the S4 numbers can be scaled by a factor of $\frac{1}{10}$ in order to get the expected dark current during an S3 storm.

3.3 APS-Based Radiation Background

For the APS-based radiation background, the shielded particle flux interacts with the Silicon device, creating an electron-hole pair in the Silicon. This process was simulated in the CREME-MC engine. The

Table 37: Radiation background on the MCP for NFI and WFI under different solar storm conditions

Condition	NFI [electron/second/NFI bin]	WFI [electron/second/WFI bin]
Quiet time	17.69	70.75
S4 Flare	8392	33570

Table 38: Simulation data for SEP interaction with the APS from the CREME-MC engine. All values are in units of electrons/second/ cm².

Condition	Direct SEP interaction with APS	Secondaries	Total
Quiet	30×10^3	30×10^3	60×10^3
S4 Flare	60×10^6	60×10^6	120×10^6

energy deposited into the Si was used to determine the number of electron/hole pairs created, at a rate of 3.6 [eV/e-h pair]. Kruk et al.'s analysis for WFIRST indicated that the secondary rate was roughly equivalent to that from the primaries [2]. Therefore the number of pairs is then multiplied by 2 to get the total number of electrons.

Table 38 contains the data from the simulation. All numbers are in units of [electrons/second/cm²].

A similar calculation to the MCP event case in Equation 17 is used to convert these to [electrons/sec/bin]:

$$\text{APS}_{\text{electrons}} \left[\frac{\text{electron}}{\text{second cm}^2} \right] \times A \left[\frac{\text{cm}^2}{\text{APS pixel}} \right] \times N_{\text{bin}} \left[\frac{\text{APS pixel}}{\text{bin}} \right] \quad (18)$$

Given the number of electrons created as input and using $N_{\text{bin}} = 4$ for the NFI case, the total factor is $3.38 \times 10^{-6} \left[\frac{\text{cm}^2}{\text{NFI bin}} \right]$. Calculations for WFI bins are similar, except $N_{\text{bin}} = 16$, making the overall factor $1.35 \times 10^{-5} \left[\frac{\text{cm}^2}{\text{WFI bin}} \right]$ instead.

Table 39 contains the radiation background for NFI and WFI in quiet time and during S4, calculated using Equation 18. Note that the contributions from the radiation on the APS are much smaller than the contributions from the radiation on the MCP.

3.4 Science Image Corner SNR

Dark images provide a measurement of the scaling factor between the gain amplified radiation signal within the field-of-view (FOV) (where the MCP and APS radiation signals are added together to produce a total radiation signal) and the non-gain amplified signal within the image corners, where only APS radiation is present. The scaling factor may vary in time throughout the storm, due to changes in the energetic electron energy spectrum and consequent changes in the generation of primary and secondary electrons in the APS silicon relative to those generated in the MCP. The temporal stability of the scaling factor will need to be assessed when the spacecraft is in orbit.

The relationship above presents a unique opportunity. During a solar storm, the dark current can be estimated from the measured signal in the image corners in order to determine the particle radiation background. While the science image will only provide the dark current and APS-based radiation background total, the MCP-based radiation background can still be found due to the estimated relationship between the

Table 39: Radiation background on the APS for NFI and WFI under different solar storm conditions

Condition	NFI [electron/second/NFI bin]	WFI [electron/second/WFI bin]
Quiet time	0.1014	0.4056
S4 Flare	202.8	811.2

Table 40: Science image corner SNR for different storm conditions for NFI and WFI.

Condition	NFI SNR			WFI SNR		
	Single pixel	One corner	All corners	Single pixel	One corner	All corners
Quiet	0.004	0.795	1.590	0.010	1.190	2.380
S2	0.071	15.90	31.80	0.201	23.80	47.60
S3	0.711	159.0	318.0	2.011	238.0	476.0
S4	7.111	1590	3180	20.11	2380	4760

APS-based and MCP-based backgrounds. Therefore, with the proper frequency of dark images, the particle radiation background can be suitably calculated from the science image corners.

However, if the dark current vastly overpowers the signal from the APS-based radiation background, then the latter will be undetectable. Therefore, the signal-to-noise ratio (SNR) is used to prove that such a strategy would allow detection of the APS-based radiation background, especially during storm conditions. The SNR for S4 is calculated by using the particle radiation background signal found in the previous section, denoted APS_{S4} . The signal is then scaled appropriately to calculate the SNR for other storm conditions. The mean dark current is assumed to be 1000 [electrons/second/APS pixel], denoted d . This value is a relatively conservative estimate; in reality it is considerably lower, which would boost the SNR ratio higher. The dark current is also scaled appropriately when converting from APS pixels to NFI or WFI bin pixels. The read noise is assumed to be 25 [electrons per APS pixel per frame], denoted R . This value is a conservative estimate; actual values tend to be around 8-9 [electrons per APS pixel per frame].

Each NFI image is taken over an integration time of 30 minutes or 1800 seconds. The number of NFI binned pixels in all image corners is $2 \cdot 10^5$. In the WFI case, each WFI image is taken over an integration time of 60 minutes or 3600 seconds. The number of WFI binned pixels in all image corners is $5.6 \cdot 10^4$. The number of binned pixels in the image corner is denoted N_{corner} . Finally, Equation 8 is used to calculate N_{frames} .

The SNR formula for a single pixel during S4 conditions is then given as:

$$\text{SNR}_{\text{single S4}} = \frac{APS_{S4} \times T_{\text{int}}}{\sqrt{APS_{S4} \times T_{\text{int}} + 19^2 \times d \times N_{\text{bin}} \times T_{\text{int}} + R^2 \times N_{\text{bin}} \times N_{\text{frames}}}} \quad (19)$$

The numerator contains the signal. The denominator contains the signal added to the dark current noise. The calculation includes an extra multiplicative factor of 19, called the dark widener, because while the dark current was assumed to be Poissonian, the distribution actually has a larger tail on the high end compared to the low end. Thus, this extra factor of 19 (conservatively high) is added to correct for this assumption.

The SNR scales with the square root of the number of pixels summed. Thus, the SNR formula for a single corner of the image during S4 conditions is then given as:

$$\text{SNR}_{\text{single corner S4}} = \sqrt{\frac{N_{\text{corner}}}{4}} \text{SNR}_{\text{single S4}} \quad (20)$$

Likewise, the SNR formula for all corners of the image during S4 conditions is given as:

$$\text{SNR}_{\text{all corners S4}} = \sqrt{N_{\text{corner}}} \text{SNR}_{\text{single S4}} \quad (21)$$

Table 40 shows the SNR of the particle radiation background signal in the image corners. The SNR during quiet-time is extremely low, but the particle radiation during quiet-time is also negligible compared to the contribution from dark current, so its effect can simply be ignored at that time.

The science image corner data provides sufficient SNR to measure the non-gain amplified particle radiation background for SEP event levels S2 or greater at the science image cadence (30 minutes for NFI and 60 minutes for WFI). The SNR is high enough to enable detection of particle radiation separately in each corner, which allows for detection of direction-of-arrival effects, which may be useful for other science purposes. The SNR is high enough during S4 events to support detection on a per-pixel basis.

4 Photon Background

Photon background is comprised of photons detected by the camera that are not a part of the target photon emission signal. These photons are split into two categories: the OOB photon background (i.e. photons with a wavelength that do not correspond to Ly- α 's wavelength) and in-band photon background (photons with a wavelength that do correspond to Ly- α 's wavelength, i.e. Ly- α photons that are not generated by exospheric H atoms). The former can be removed through the use of optical filters, while the latter requires independent background scene sensing because it is not spectrally separable from the target signal.

The OOB photon background is dominated by atomic/molecular Oxygen and Nitrogen emissions and continuum albedo from Earth's limb/disk. Only the few innermost science pixels are affected. The line emissions of Hydrogen Ly- α , Oxygen 1304, and Oxygen 1356 are modelled as Gaussian spectral distributions, where integration over wavelength corresponds to line emission radiance. The altitude-dependent emission radiance of Oxygen 1304 and 1356 emission lines is specified using either available observations (e.g. GUVI [1]) or AURIC model [6] output. Exospheric Ly- α is also modeled as a scaled Gaussian spectral line, but the line-integrated emission radiance is specified using radiative transfer modeling, to be discussed in the next section.

The OOB signal is measured daily using designated long-pass optical filters (CaF2 and SrF2), which can then be easily subtracted out, as discussed in Section 1.2. Therefore, this thesis mainly focuses on the task of estimating the in-band photon background.

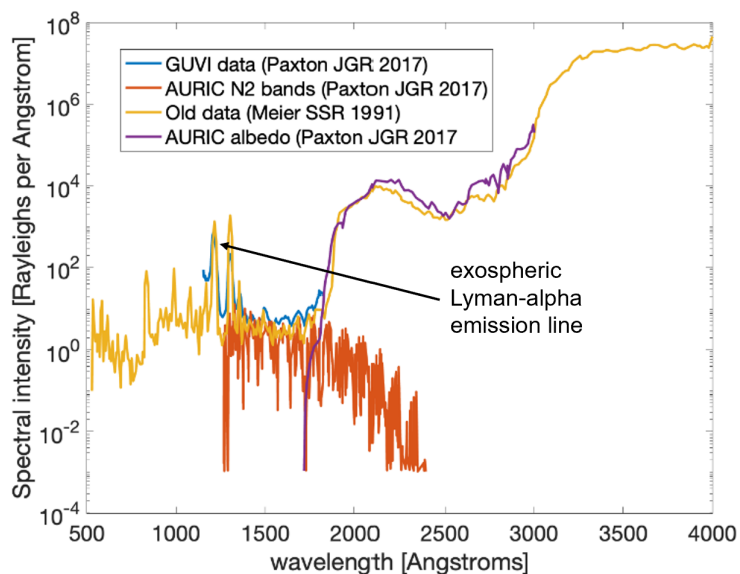


Figure 17: Out-of-band and in-band emission spectrum. The emission radiance from observations and the AURIC model are superposed.

4.1 In-Band (IPH) Photon Background

The in-band photon background is dominated by emission IPH atoms, where the H atoms are of interstellar origin. The IPH background is measured in a scan around the perimeter of the Earth's exosphere during the IOC phase and in the outer annulus of every WFI science image during science phase for interpolation onto exospheric LOS. The inner boundary of the annulus to be used is 30 Earth-Radii (R_E) from nadir: during the mission the Earth and its exosphere will be located within this boundary in the image and obfuscate the IPH signal. The plate scale of the image is calculated and used in the algorithm to determine which pixels in the WFI image can be included in the IPH background retrieval.

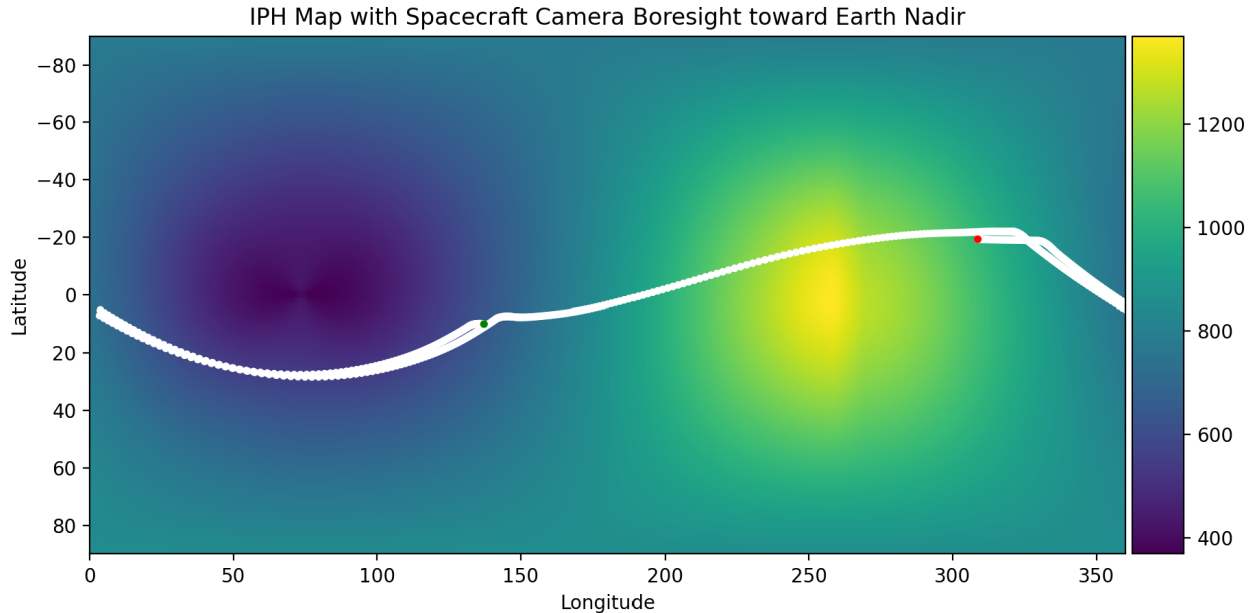


Figure 18: The spacecraft’s camera boresight, directed towards Earth nadir, from July 1st, 2025 to February 18th, 2027. The red dot indicates the camera boresight on July 1st, 2025, while the green dot indicates the camera boresight on February 18th, 2027. The full lifetime of the spacecraft is around two years, but uncertainties in the baseline orbit ephemeris prevent the path from being extrapolated to future dates at this time.

The spacecraft’s camera boresight toward Earth nadir superposed on the IPH map from July 1st, 2025 to February 18th, 2027 is shown in Figure 18. The IPH map itself was produced by first finding the structure, generated using Pryor IPH code [5], then scaled according to observation conditions and observing target parameters. Figure 19 shows the WFI FOV around the camera boresight direction, as well as the relative size of the Earth within the WFI FOV. The Earth’s size in the WFI FOV changes drastically depending on the spacecraft vantage on orbit, which can affect the accuracy and precision of the algorithm, since a closer vantage corresponds to fewer available pixels in the annulus.

4.2 Shot Noise Denoising

The IPH retrieval algorithm is given a noisy Level 0 WFI image, where each pixel is a noisy realization of a discrete Poisson distribution. The initial denoising step in the algorithm involves drawing k independent samples from a Poisson distribution with some large mean and variance γ , then using either the mean or median of these values as the true value to be used during interpolation. In this section we prove that taking the median is the superior choice, which is why the algorithm takes the median of these samples instead of the mean.

Claim: Let ϵ be a relatively small positive constant denoting the amount of deviation from γ . Given a Poisson distribution with some large mean and variance γ , for the same error ϵ off of γ , the mean of k independent samples has a higher probability of being incorrect compared to the median of k independent samples, for some lower bound on k .

Proof: Let N_i be the random variable representing the i ’th sample from the Poisson distribution. Clearly, N_i has mean and variance γ . The sum of all these N_i is defined as N .

$$N = \sum_{i=0}^k N_i \quad (22)$$

However, since the sum of independent Poisson distributions is still a Poisson distribution with mean

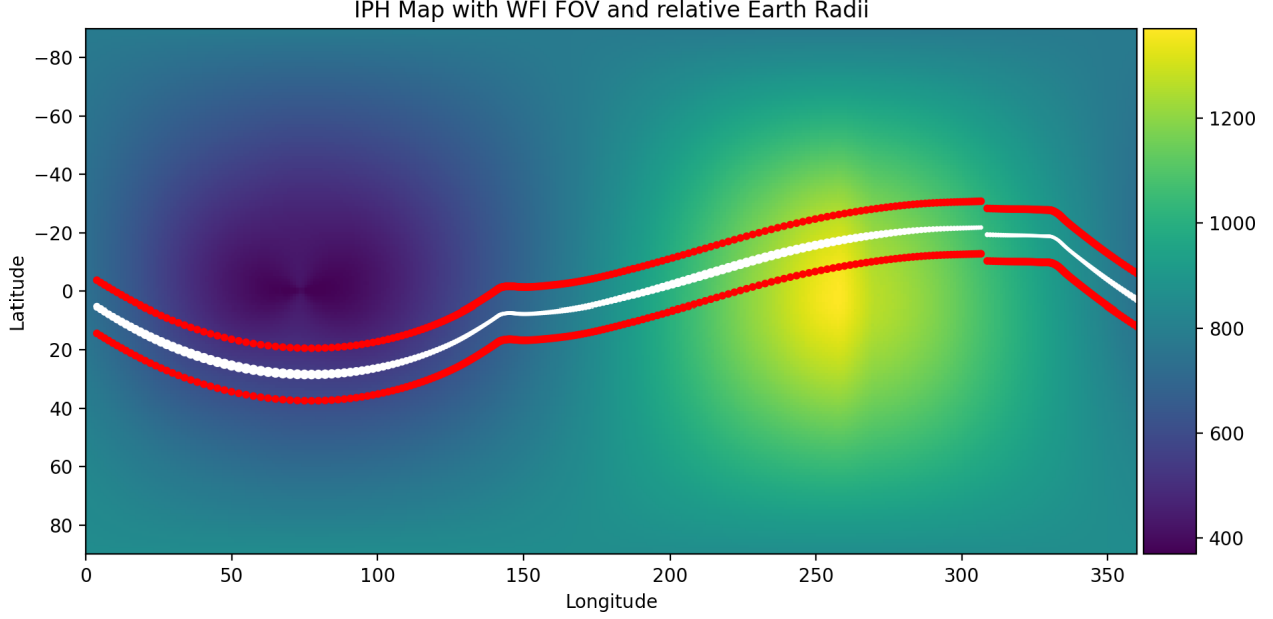


Figure 19: The spacecraft’s camera boresight, directed towards Earth nadir, from July 1st, 2025 to June 30th, 2026. The red lines outline the range of the WFI FOV that the spacecraft will see. The relative size of the white dot represents the relative size of the Earth within the WFI FOV. The size itself is not drawn to scale.

and variance equal to the sum of the individual means and variances, we know that N is itself a Poisson distribution with mean and variance $k\gamma$. We are now interested in the probability of this mean deviating from γ by a constant ϵ or greater. The value of ϵ is restricted to be between 0 and γ to avoid some edge cases (in our practical application, this restriction does not mean much). We’re interested in the mean, i.e. the random variable $\frac{N}{k}$. The equations below represent the probability that the mean is “wrong”.

$$\begin{aligned} \Pr\left(\left|\frac{N}{k} - \gamma\right| > \epsilon\right) &= \Pr\left(\frac{N}{k} - \gamma > -\epsilon\right) - \Pr\left(\frac{N}{k} - \gamma > \epsilon\right) \\ &= \Pr(N > k\gamma - k\epsilon) - \Pr(N > k\gamma + k\epsilon) \end{aligned} \quad (23)$$

Define Y_i as an indicator random variable, where $Y_i = 1$ if N_i is absolutely deviates from γ by more than ϵ , and 0 otherwise.

$$\begin{aligned} \Pr(Y_i = 1) &= E[Y_i] = \Pr(|N_i - \gamma| > \epsilon) \\ &= \Pr(N_i - \gamma > -\epsilon) - \Pr(N_i - \gamma > \epsilon) \\ &= \Pr(N_i > -\epsilon + \gamma) - \Pr(N_i > \epsilon + \gamma) \end{aligned} \quad (24)$$

Given that γ is large, Poisson processes can be approximated as a Gaussian distribution with mean and variance γ . An informal argument for this is the Central Limit Theorem - since Poisson distributions with arbitrarily large γ is the same as summing up γ Poisson distributions with mean and variance 1, by the Central Limit Theorem, these large Poisson distributions approach the Gaussian distribution. Notice that N_i and N can both be approximated as Gaussian distributions, with N being the scaled version of N_i by a factor of k . Notice also that the constants on the right-hand-side of the probabilities in Equations 23 and 24 are scaled versions of each other by a factor of k . Therefore, the probabilities in Equations 23 and 24 are approximately the same and shall be denoted as simply p .

Let $Y = \sum_i Y_i$, and let μ_Y denote the expected value of Y . By linearity of expectation, $\mu_Y = \sum_i E[Y_i] = kp$. For the median to deviate from γ by ϵ or greater, at least $\frac{k}{2}$ indicator variables must be set to 1, i.e. at least half of the trials must absolutely deviate from γ by more than ϵ . Note that the latter condition is weaker but completely encapsulates than the former. Therefore, any probability bound on the latter must also be a probability bound on the former. We then find the probability that $Y \geq \frac{k}{2}$:

$$\begin{aligned} \Pr\left(Y \geq \frac{k}{2}\right) &\leq \Pr(Y \geq (1 + \Delta)\mu_Y) \\ &\leq \exp\left(\frac{-\Delta\mu_Y}{3}\right) \end{aligned} \tag{25}$$

The latter inequality is the Chernoff bound for sums of indicator variables for $\Delta \geq 1$. We can express Δ in terms of p :

$$\begin{aligned} \frac{k}{2} &= (1 + \Delta)kp \\ \Delta &= \frac{1}{2p} - 1 \end{aligned} \tag{26}$$

The Chernoff bound also requires $p \leq \frac{1}{4}$, otherwise it is invalid, which is fine for our applications. What remains is to show the following:

$$\exp\left(-\frac{\left(\frac{1}{2p} - 1\right)kp}{3}\right) = \exp\left(-\frac{k}{6} + \frac{k}{3}p\right) \leq p \tag{27}$$

where the left side of the inequality is the upper bound on the probability that the median is wrong, while the right side is the exact probability that the mean is wrong. After some mathematical manipulations, this inequality simplifies to a lower bound on k that makes the claim true:

$$\boxed{k \geq -\frac{6 \ln(p)}{1 - 2p}} \tag{28}$$

4.3 IPH WFI Interpolation


4.3.1 Problem Statement

Given a noisy Level 0 WFI image that contains only the IPH background, bias, dark, radiation, and accompanying noise, along with information about the spacecraft’s current vantage point (represented by an inputted date), the WFI IPH-Estimator algorithm aims to estimate the IPH background within the entire WFI image. Recall that the algorithm is only allowed to use the outer annulus of the given image (pixel tangent altitudes > 30 Re) to interpolate due to contaminating signals from Earth’s exosphere. Figure 20a shows an example of a noisy image with the annulus, while Figure 20b shows the desired output of the algorithm (the IPH itself), given the noisy Level 0 image as input.

4.3.2 Assumptions and Notation

Below are a few assumptions that the algorithm makes when recovering the IPH:

- The algorithm assumes that it has access to an accurate bias, mean dark current, and mean radiation. These are then subtracted off of the Level 0 image immediately. Sensitivity to bias in instrument background knowledge will be tested as part of future work.
- The algorithm assumes that there is some direction where there is a strong scene gradient (not necessarily linear), while the direction perpendicular to it is roughly constant (or at worst, linear). This assumption is valid for a large majority of the mission’s orbit. A slight modification to this algorithm will be required when the assumption fails (to be discussed later).



(a) Noisy Level 0 image example on July 1st, 2025, over a six-hour integration period. The dot in the middle is Earth itself, while the large red circle, with a radius of $30R_E$, denotes the annulus boundary. Pixels within this boundary cannot be used by the interpolation algorithm.

(b) Ground Truth example on July 1st, 2025, over a six-hour integration period. Used to make the synthetic Noisy Level 0 image.

Figure 20: WFI IPH-Estimator algorithm input (left) and desired output (right)

The image after the bias, dark current, and radiation are all subtracted off is denoted as M , which is indexed by i and j , where $M[i, j]$ denotes the value, in [DN], of the pixel signal in the i 'th row and the j 'th column. The center of the polar coordinate system is placed at the center of the image, where Earth is, at $i = 256, j = 256$. The radius of the pixel in the i 'th row and the j 'th column is denoted as r_{ij} , while similar definitions follow for the angle θ_{ij} . Pixels with $r > 256$ are not considered since they are not within the WFI image's FOV for this particular vantage and associated plate scale. Finally, $30 R_E$ is converted to the number of pixels from Earth, which is denoted as r_E . Any pixels within this radius cannot be used by the algorithm when interpolating.

4.3.3 WFI IPH-Estimator Algorithm

The algorithm first focuses on recovering the main gradient direction, then interpolates across the center of the image FOV.

To recover the gradient direction, 1080 wedges, denoted N_{wedges} , indexed by k , of width $\theta_{\text{wedges}} = 45^\circ$, are used. The sum of the background-subtracted signal in the pixels within the wedge is calculated.

$$W[k] = \sum_{r_E \leq r_{ij} \leq \frac{N_{\text{pix}}}{2}} \sum_{\frac{k}{3} \leq \theta_{ij} \leq \frac{k}{3} + \theta_{\text{wedges}}} M[i, j] \quad (29)$$

The specific values chosen for N_{wedges} and θ_{wedges} is arbitrary, but seems to work well in practice. The direction of the gradient is then identified as the two wedges 180° apart (or equivalently, $\frac{N_{\text{wedges}}}{2}$ indices apart) that have the largest signal difference between them.

$$\alpha = \arg \max_{\ell} \left\{ W[\ell] - W\left[\ell - \frac{N_{\text{wedges}}}{2}\right] \right\} \quad (30)$$

$$W_{\text{max}} = W[\alpha] \quad (31)$$

$$W_{\text{min}} = W\left[\alpha - \frac{N_{\text{wedges}}}{2}\right] \quad (32)$$

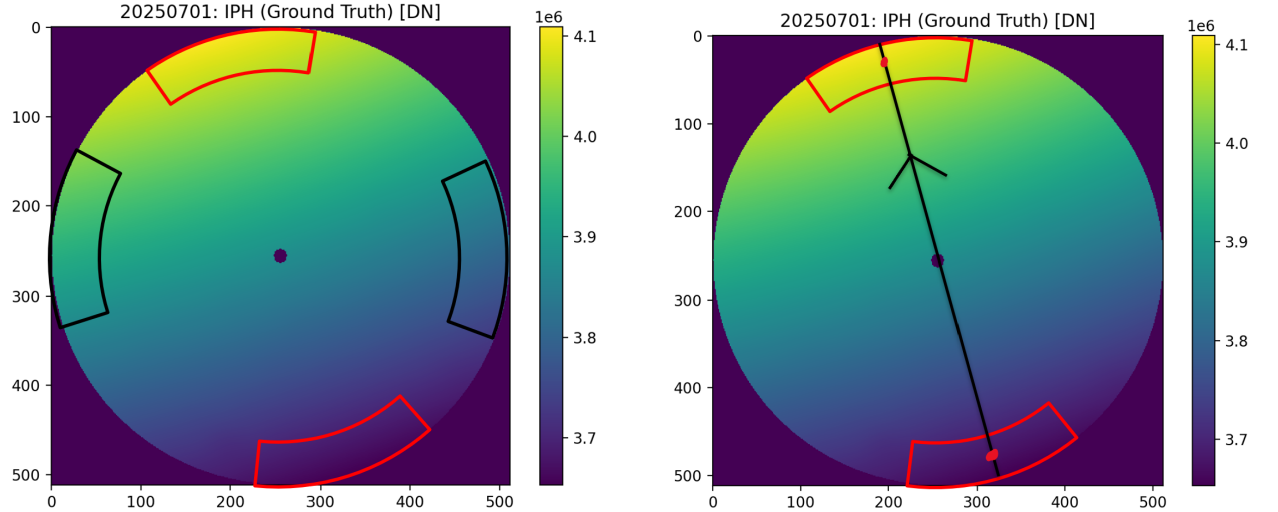
Above, α represents the index at which the pair of wedges 180° apart has the biggest difference. W_{\max} is then the wedge with the largest value in the annulus, while W_{\min} is the wedge with the smallest value in the annulus. Figure 21a shows an example of these wedges and an informal proof-by-picture for why this works. All figures with details about the algorithm are shown using the same ground-truth image as an example; in reality, all of the algorithm’s steps are performed on the noisy images. The use of ground truth in these illustrations is for clarity only.

Once the algorithm identifies the gradient direction, the algorithm then extrapolates the line, in Cartesian coordinates, corresponding to our gradient direction from the centers of the wedges. To do this, the algorithm converts the wedge center coordinates back to Cartesian coordinates, then calculates the slope g_m and intercept g_b of the line that crosses both points. Let $C_{W_{\max},x}$ denote the x coordinate of the center of wedge W_{\max} in Cartesian coordinates. Similar definitions hold for $C_{W_{\max},y}$.

$$g_m = \frac{C_{W_{\max},y} - C_{W_{\min},y}}{C_{W_{\max},x} - C_{W_{\min},x}} \quad (33)$$

$$g_b = C_{W_{\max},y} - g_m \cdot C_{W_{\max},x} \quad (34)$$

Figure 21b shows the recovered gradient direction. The slope and intercept of this line are g_m and g_b respectively.



(a) Example of two pairs of wedges, 180° apart. The difference between the sum of the signal in the pair of red wedges is far larger than the difference between the sum of the signal in the pair of black wedges, so the red wedges will be selected as the gradient direction.

(b) Once the two wedges are chosen, their centers (denoted with the red dots) are used to extrapolate the gradient direction line, shown in black. This line must pass through the center of the image as well, since the original wedges were exactly 180° apart.

Figure 21: Phase 1 of WFI IPH-Estimator algorithm: recovering the gradient direction. Although the wedges and gradient are displayed on the ground-truth, the algorithm still applies to the noisy Level 0 image, since the summation of many noisy pixels tends towards the (scaled) mean of those pixels by the law of large numbers.

The second phase recovers the mean IPH signal itself. To do this, we use $N_{\text{wedges}} = 1080$ wedges, indexed by k , of width $\theta_{\text{wedges}} = 12^\circ$. Once again, the specific values chosen for N_{wedges} and θ_{wedges} is arbitrary, but it seems to work well in practice. The median of pixels within each wedge is calculated. Note that we use median according to the results from Section 4.2, which is valid for the following reasons:

- The mean for the Poisson distribution involved is on the scale of 10^6 , which is large enough that the Gaussian approximation holds well.

- The error needs to be as small as possible, i.e. $p \leq \frac{1}{4}$ is definitely satisfied.
- If $p = 0.0001 = 0.01\%$, then, using Equation 28, we require $k \geq 56$. The number of pixels in each wedge varies and is highly dependent on the size of the annulus. However, even in the worst-case where the annulus starts at a radius of 242 pixels (as seen in Section 4.3.4), we still have $k = \pi \left(\left(\frac{N_{\text{pix}}}{2} \right)^2 - 242^2 \right) \frac{\theta_{\text{wedges}}}{360} = 781$, which meets the restriction. In the more average case where the annulus starts at a radius of 200 pixels, we instead have $k = 2674$. By the claim in Section 4.2, median is superior to mean.

$$I[k] = \text{Median} \left\{ M[i, j] \left| r_E \leq r_{ij} \leq \frac{N_{\text{pix}}}{2}, \frac{k}{3} \leq \theta_{ij} \leq \frac{k}{3} + \theta_{\text{wedges}} \right. \right\} \quad (35)$$

The wedge centers are then calculated for each pair of $I[k]$ that are 180° apart (or equivalently, $\frac{N_{\text{wedges}}}{2}$ indices apart). Each pair of wedges then defines a unique line that crosses through both centers in the pair, with some slope and intercept. The value $I[k]$ is then replaced with the corresponding value at the very edge of the image, linearly interpolated from the values at the pair of wedge centers.

Instead of finding the slope and intercept though, the algorithm instead draws a 1-dimensional number line along the line defined by the wedge centers, with the origin of the line at the FOV origin. The algorithm then linearly interpolates the value along this line based on the Euclidean distance of each wedge center from this origin. The value at the origin is also required, but for this interpolation, since the wedge centers are 180° apart at some constant r value, the average of the values suffices as the value at the origin. Let $D(p_1, p_2)$ denote the Euclidean distance between two points p_1 and p_2 .

$$I[k] = \frac{I[k] + I[k + \frac{N_{\text{wedges}}}{2}]}{2} + \frac{I[k] - I[k + \frac{N_{\text{wedges}}}{2}]}{D(C_{I[k]}, C_{I[k + \frac{N_{\text{wedges}}}{2}]})} \cdot \frac{N_{\text{pix}}}{2} \quad (36)$$

$$I[k + \frac{N_{\text{wedges}}}{2}] = \frac{I[k] + I[k + \frac{N_{\text{wedges}}}{2}]}{2} - \frac{I[k] - I[k + \frac{N_{\text{wedges}}}{2}]}{D(C_{I[k]}, C_{I[k + \frac{N_{\text{wedges}}}{2}]})} \cdot \frac{N_{\text{pix}}}{2} \quad (37)$$

The first term in Equations 36 and 37 represents the signal at the origin. The second term is the signal gained per unit pixel along the line, multiplied by $\frac{N_{\text{pix}}}{2}$ since the value at $r = \frac{N_{\text{pix}}}{2}$ is desired. Therefore, $I[k]$ now holds the value at polar coordinates $r = \frac{N_{\text{pix}}}{2}$, $\theta = \frac{k}{3}$. Figure 22 shows an example of this process.

Next, all points (i, j) within the image, including the ones within the annulus, are projected down to the line defining the gradient direction through the use of an affine transformation. The affine transformation is simply a series of matrix multiplications that first translates the entire system so that the line defining the gradient direction crosses the origin. A linear projection is then performed (which is now possible since the line to be projected upon passes through the origin). The result is then inverse translated back to the original coordinates. Affine translations require an extra 1 to be padded to the end of the vector being transformed.

$$\vec{g} = \begin{bmatrix} C_{W_{\text{max}},x} - C_{W_{\text{min}},x} \\ C_{W_{\text{max}},y} - C_{W_{\text{min}},y} \end{bmatrix} \quad (38)$$

$$P_{ij} = (t_i, t_j) \text{ s.t. } \begin{bmatrix} t_i \\ t_j \\ 1 \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & g_b \\ 0 & 0 & 1 \end{bmatrix} \left[\begin{bmatrix} \frac{\vec{g}\vec{g}^T}{\vec{g}^T\vec{g}} & 0 \\ 0 & 0 & 1 \end{bmatrix} \right] \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & -g_b \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} i \\ j \\ 1 \end{bmatrix} \quad (39)$$

This process will also project the polar coordinates $(r = \frac{N_{\text{pix}}}{2}, \theta = \frac{k}{3})$ to a location somewhere along this line, which is denoted P_k (these points have a value $I[k]$, calculated earlier). Due to the geometry of the image, for nearly all (i, j) there are two values of k , denoted k_1 and k_2 , such that $P_{ij} = P_{k_1} = P_{k_2}$. The points (i, j) that project to the same location along the line as k_1 and k_2 are binned. This bin then contains exactly the points that lie on a line perpendicular to the gradient direction that crosses $(r = \frac{N_{\text{pix}}}{2}, \theta = \frac{k_1}{3})$ and $(r = \frac{N_{\text{pix}}}{2}, \theta = \frac{k_2}{3})$.

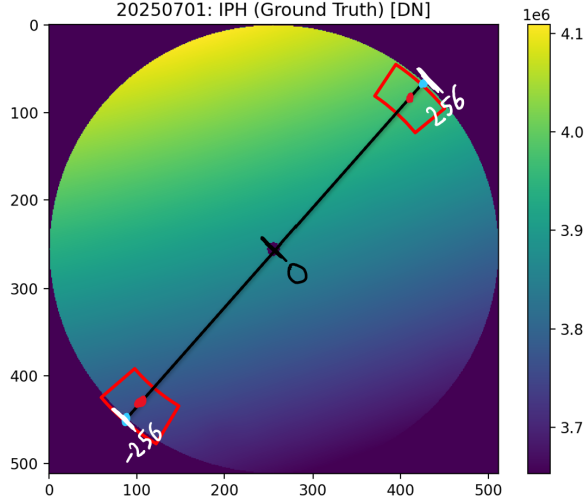


Figure 22: Example of a pair of interpolation wedges with width 12° , 180° apart. The number line is in black, with the origin at the center of the image. The centers of the wedges, marked as red dots, define this number line. The values at the blue dots are then interpolated based on a linear fit.

In practice, k_1 and k_2 are found by calculating the signed distance along the 1-dimensional number-line defined by the gradient direction, with the origin placed at the center of the origin. As a sign convention, positive distances are on the higher side of the gradient, while negative distances are on the lower side. Let $D_{s,g}(P)$ denote this distance for point P . We then have

$$k_1 = \alpha + \frac{1080}{360} \arccos\left(\frac{D_{s,g}(P_{ij})}{\frac{N_{\text{pix}}}{2}}\right)$$

$$k_2 = \alpha - \frac{1080}{360} \arccos\left(\frac{D_{s,g}(P_{ij})}{\frac{N_{\text{pix}}}{2}}\right)$$

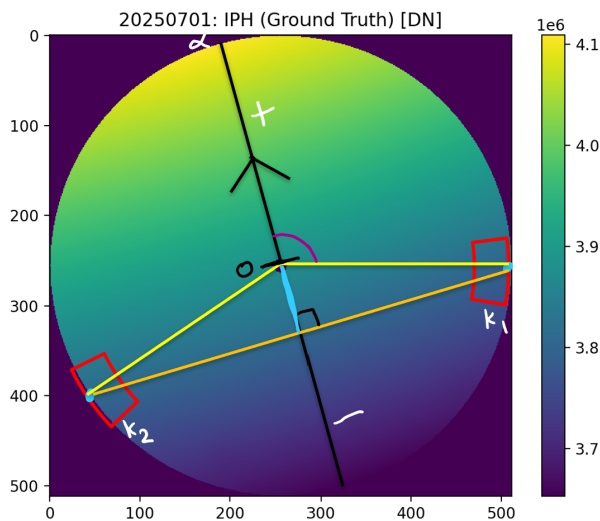
where α is the index of the maximum wedge, as defined earlier. Figure 23a shows an example of the recovery of k_1 and k_2 .

The final step is to then assign values to the points within each bin based on the values $I[k_1]$ and $I[k_2]$. Two options exist: one is to just take the average of $I[k_1]$ and $I[k_2]$, while the second is to do another linear fit. By our assumption, the variation along this direction is expected to be extremely small, so a linear fit serves as a good first-order approximation. Therefore, the algorithm performs the latter to get the recovered IPH estimated value at (i, j) , denoted $\hat{H}[i, j]$. The linear interpolation here is very similar to the linear interpolation performed when recalculating $I[k]$ and $I[k + \frac{N_{\text{wedges}}}{2}]$. Let $D_s(p_1)$ denote the signed distance along the 1-dimensional number-line defined by the center of the k_1 and k_2 edges. If p_1 is closer to k_1 , then the distance is positive, while if p_1 is closer to k_2 , then the distance is negative. The origin of this number line is the point halfway between these centers.


$$D_s((i, j)) = \begin{cases} D((i, j), \frac{k_1+k_2}{2}) & \text{if } D((i, j), k_1) > D((i, j), k_2) \\ -D((i, j), \frac{k_1+k_2}{2}) & \text{if } D((i, j), k_1) < D((i, j), k_2) \\ 0 & \text{else} \end{cases} \quad (40)$$

$$\hat{H}[i, j] = \frac{I[k_1] + I[k_2]}{2} + \frac{I[k_1] - I[k_2]}{D(C_{I[k_1]}, C_{I[k_2]})} \cdot D_s((i, j)) \quad (41)$$

Figure 23b shows the final steps of the algorithm in order to interpolate the values in the middle of the image.



(a) k_1 and k_2 calculation. Any point on the orange line will project to the same spot on the gradient line (black). The blue line is then the signed distance from the origin to these points. Note that the distance is negative in this example because the orange line is closer to the negative side of the gradient. Taking the arccos gives the angle shown in purple. This angle defines the offset from α , top index of the top of the gradient. The offset is scaled by a factor of 3 since there are 1080 wedges and only 360 degrees. This gives k_1 and k_2 .



(b) Once the two wedges are chosen, another number line is setup between them (denoted in orange). Due to the calculations in the previous part, this number line must be perpendicular to the gradient line, shown in black. The values along this line are then linearly interpolated from the values at the blue dots ($I[k_1]$ and $I[k_2]$).

Figure 23: Phase 2 of WFI IPH-Estimator algorithm: Recovering the IPH values. Although the wedges and geometry are displayed on the ground-truth, the process works on the noisy Level 0 image, since using the median of the wedge as the actual value for the blue dot gives an accurate estimate of the true value of the wedge.

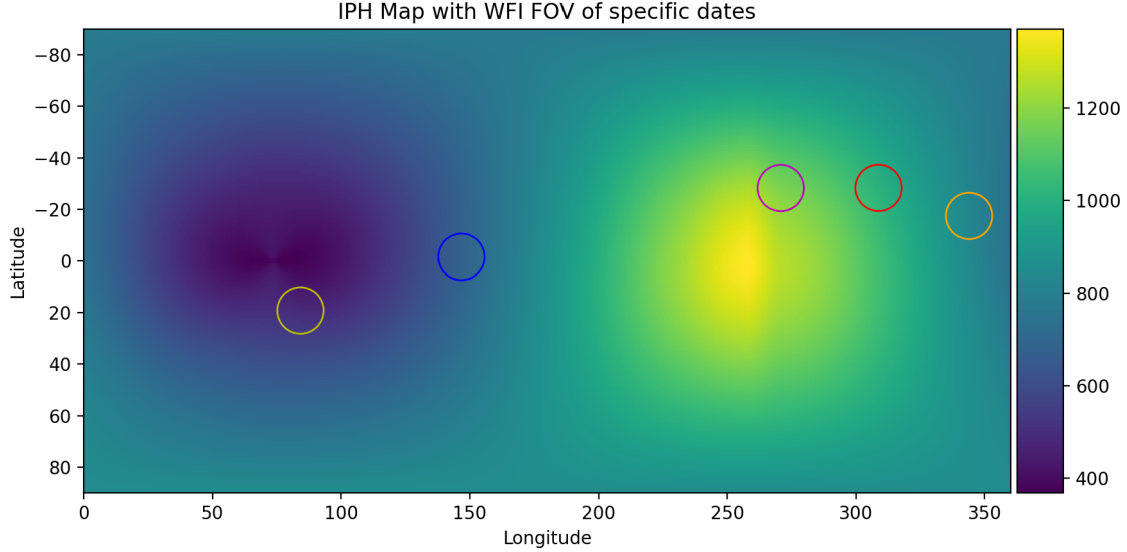


Figure 24: IPH Map with the FOV of the spacecraft at the five different test dates. Pixels within the circle will be captured by the WFI camera. The red is for July 1st, 2025. Orange is for October 1st, 2025. Yellow is for December 1st, 2025. Blue is for March 1st, 2026. Finally, purple is for June 1st, 2026.

The final result tends to have stripes, since different bins (lines perpendicular to the gradient) do not have any explicit relation to each other when being interpolated. However, it is known that these stripes must be a result of noise, since the IPH itself is not expected to have stripes. Therefore, to denoise the image even further, a smoothing function is applied. This step and its effect are discussed in more detail in the next section.

4.3.4 Validation

When validating the WFI IPH-Estimator algorithm, synthetic data is generated using the synthetic image model. The input parameters to the model are such that the spacecraft always points towards earth, but the photons from the Earth, its exosphere, and any stellar radiance are all set to 0. The operational mode is set to use the WFI channel and the narrowband Ly- α (LyaN) filter, with an integration time of 60 minutes. For this section, the flat-field is turned off (it is assumed that the flat-field has already calibrated out).

The algorithm was tested for several possible dates, from the best-case scenario to the worst-case scenario, while varying the number of hours stacked in order to see the effect of smoothing. The dates chosen were July 1st, 2025, October 1st, 2025, December 1st, 2025, March 1st, 2026, and June 1st, 2026. On these dates, the spacecraft's camera boresight is at the locations shown in Figure 24. For each date, h synthetic noisy images were generated using the synthetic image model, corresponding to one per hour for h hours after midnight on each date. The noisy images were then stacked and passed through the WFI IPH-Estimator algorithm to obtain the recovered WFI IPH image, $\hat{H}[i, j]$. The recovered image was then compared to the ground-truth image, $H[i, j]$, which is simply the IPH signal without any noise added on to it. The percent error $P[i, j]$ is then obtained as follows:

$$P[i, j] = 100 \frac{\hat{H}[i, j] - H[i, j]}{H[i, j]} \quad (42)$$

The maximum magnitude percent error (MMPE) was then recorded. In the tables, x -hIT stands for x -hour integration time. Table 41 contains the MMPE for various dates and integration times when smoothing wasn't applied, while Table 42 contains the MMPE for various dates and integration times with smoothing. In general, it appears that the percent error tends to go down when smoothing is applied, but it is hard to

Table 41: One run of the algorithm with various integration times across various dates, without smoothing.

Date	Annulus Start Radius [pixels]	1-hIT MMPE [%]	3-hIT MMPE [%]	6-hIT MMPE [%]	9-hIT MMPE [%]	12-hIT MMPE [%]
July 1st, 2025	198.12	0.737	0.425	0.390	0.343	0.268
October 1st, 2025	195.07	0.998	0.694	0.328	0.318	0.258
December 1st, 2025	241.87	1.638	1.506	1.388	1.424	1.426
March 1st, 2026	186.99	1.024	0.717	0.694	0.575	0.618
June 1st, 2026	226.97	1.151	1.141	1.035	1.108	1.060

Table 42: One run of the algorithm with various integration times across various dates, with smoothing.

Date	Annulus Start Radius [pixels]	1-hIT MMPE [%]	3-hIT MMPE [%]	6-hIT MMPE [%]	9-hIT MMPE [%]	12-hIT MMPE [%]
July 1st, 2025	198.12	0.692	0.469	0.276	0.249	0.215
October 1st, 2025	195.07	0.696	0.509	0.187	0.283	0.233
December 1st, 2025	241.87	1.728	1.183	1.161	1.224	1.218
March 1st, 2026	186.99	0.995	0.616	0.497	0.394	0.356
June 1st, 2026	226.97	1.081	0.983	0.911	0.999	1.090

draw a conclusion from this data since the variance can be extremely high. Therefore, we instead turn to other metrics.

From the data above, an integration time of 6 hours was chosen as a good balance between algorithm accuracy and efficiency while in orbit. The percent error calculation was then repeated for $N_{\text{trials}} = 500$ trials for each date to obtain a 3-dimensional error matrix $P[i, j, k]$, where k indexes the trial number. The mean error per pixel, which denotes our accuracy per pixel, is then calculated as

$$\mu[i, j] = \frac{1}{N_{\text{trials}}} \sum_{k=1}^{N_{\text{trials}}} P[i, j, k] \quad (43)$$

The standard deviation of the error per pixel, which denotes our precision per pixel, is then calculated as

$$\sigma[i, j] = \frac{1}{\sqrt{N_{\text{trials}}}} \sqrt{\sum_{k=1}^{N_{\text{trials}}} (P[i, j, k] - \mu[i, j])^2} \quad (44)$$

This was repeated without smoothing (Table 43) and with smoothing (Table 44) in order to show the effect of smoothing. In order to condense the data, only the maximum magnitude percent error is reported. In the tables, MMMPE stands for Maximum Magnitude Mean Percent Error, while MMPESD stands for Maximum Magnitude Percent Error Standard Deviation. Both have units of percentage. From these statistics, it is clear that smoothing generally has a positive effect on the mean percent error, but also tends to increase the standard deviation. Therefore, smoothing should be implemented.

The requirement for accuracy on the IPH background retrieval is 30%. It is clear from the table of results that the algorithm's performance clearly meets this benchmark. Figure 25a shows an example of the mean percent error over 500 iterations, while Figure 25b shows an example of the percent error standard deviation over 500 iterations.

4.3.5 Failure of Linearity Assumption

From the table above, it appears that the months of July, October, and March have very low percent errors, compared to the months of December and June. Looking at Figure 24, it is noticeable that the

Table 43: Statistics for 500 iterations of the WFI IPH-Estimator algorithm, without smoothing.

Date	MMMPE [%]	MMPESD [%]
July 1st, 2025	0.230	0.103
October 1st, 2025	0.106	0.120
December 1st, 2025	0.861	0.255
March 1st, 2026	0.214	0.140
June 1st, 2026	1.450	0.119

Table 44: Statistics for 500 iterations of the WFI IPH-Estimator algorithm, with smoothing.

Date	MMMPE [%]	MMPESD [%]
July 1st, 2025	0.210	0.139
October 1st, 2025	0.0859	0.152
December 1st, 2025	0.406	0.379
March 1st, 2026	0.193	0.167
June 1st, 2026	1.428	0.141

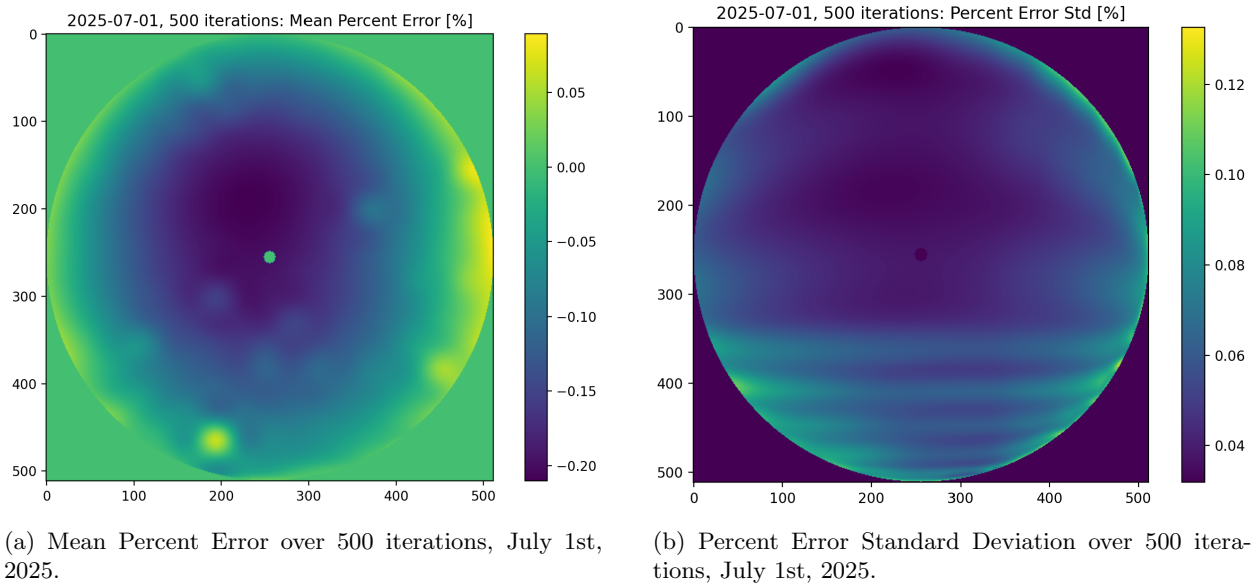


Figure 25: Example accuracy and precision results for July 1st, 2025, with 500 iterations.

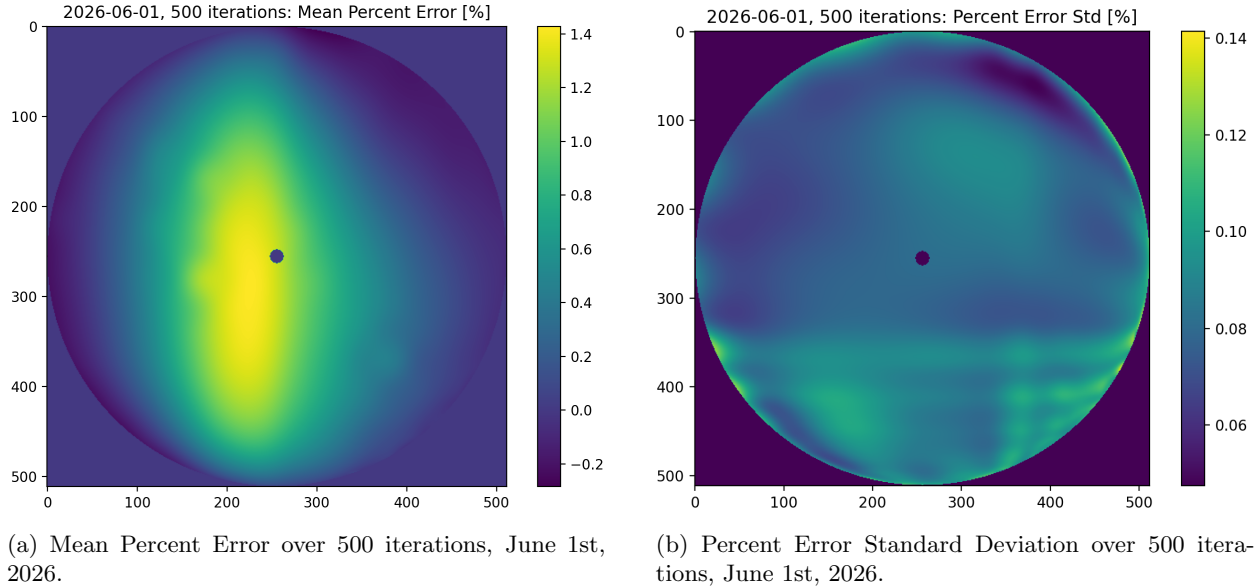


Figure 26: Accuracy and precision results for June 1st, 2026, with 500 iterations. Note the 1.4 mean percent error.

months of December and June are very close to the peaks and valleys of the IPH radiance. These areas are where the linearity assumption begins to break down, thus giving rise to larger percent errors. Figure 26a shows the mean percent error over 500 iterations for June 1st, 2026, while Figure 26b shows the percent error standard deviation over 500 iterations for June 1st, 2026.

A few methods exist to combat this. First, IPH retrieval data from before and after the problematic area could be used to interpolate the IPH data in the area. If the gradient is seen flipping direction over the course of a time period, then it is clear that the spacecraft has passed over a peak or valley of the IPH, thus giving the algorithm enough constraints to be able to interpolate this area. The issue with this strategy is that the IPH background is not stable over long periods of time, since the radiance varies as a function of solar illumination. The algorithm would be able to detect when the spacecraft's boresight direction passed over a peak or valley, but it would still be unable to extract the actual value due to this variability.

Another option is the use of priors - the IPH radiance distribution is quite stable, so we can predict when the spacecraft will reach the problematic areas of the IPH. In these areas, we can then multiply the recovered IPH image by a prior in order to more closely match what we believe should actually be there.

Finally, some of the calibration time can be used to slew the spacecraft towards the problematic area and measure the signal. In this case, Earth and its exosphere will not be in the scene, so the interpolation can then use the entire image to recover the IPH background instead of the outer annulus.


Future work will have to determine which one of these is the optimal approach to solve this problem.

4.4 IPH NFI Interpolation

4.4.1 Algorithm

Given the results of the WFI IPH-Estimator algorithm, the NFI IPH-Estimator algorithm estimates the IPH background within an NFI image that has the same boresight as the WFI image. The NFI's FOV width is too small to be able to use an annulus in the NFI image, which is why a different algorithm is used. Figure 27a shows an example of a noisy image with the annulus, while Figure 27b shows the desired output of the algorithm (the IPH itself), given the noisy Level 0 image as input.

The core algorithm idea remains the same. The implementation details differ slightly because of the FOV change. The WFI has a FOV of 18 degrees and takes images that are 512×512 pixels in size, while the NFI has an FOV of 3.6 degrees and takes images that are 1024×1024 in size. Polar coordinates still work since



(a) Noisy NFI Level 0 image example on July 1st, 2025, over a six-hour integration period.

(b) NFI Ground Truth example on July 1st, 2025, over a six-hour integration period. Note the change in image size and scaling. The hole in the middle of the image is also much bigger, since the size of the Earth is far bigger in the NFI's FOV.

Figure 27: NFI IPH-Estimator algorithm input (left) and desired output (right)

the boresight does not change (i.e. the center of the image still contains Earth). The gradient direction is also unchanged since the NFI and WFI images use the same roll angle.

In fact, all that is required is to scale each NFI pixel up to its WFI counterpart, then pass these through the WFI IPH-Estimator algorithm's translation matrix, arccosine formula, and linear interpolation perpendicular to the gradient direction. Let (i_n, j_n) denote the indices of the pixel in the NFI image, ranging from 0 to 1023. (i_w, j_w) follow similar definitions for the indices of the pixel in the WFI image.

$$(i_w, j_w) = \left((i_n - 512) \cdot \frac{3.6}{18} \cdot \frac{512}{1024} + 256, (j_n - 512) \cdot \frac{3.6}{18} \cdot \frac{512}{1024} + 256 \right) \quad (45)$$

Equation 45 shifts all pixels in the NFI image such that the center is at $(0, 0)$, scales the FOV and each individual pixel, then shifts all pixels such that $(0, 0)$ is now the center of the WFI image $(256, 256)$. These pixel indices are allowed to be non-integers - no such assumption was required for the WFI algorithm.

Once the algorithm finishes, the image is divided by the WFI responsivity and multiplied by the NFI responsivity. These responsivity constants account for the differences between WFI and NFI cameras, including efficiency at Ly- α and etendue. The WFI and NFI cameras are slightly different from each other, resulting in a different efficiency (events per photon) value at Ly- α , while etendue handles the differences in solid angle, FOV, and number of pixels. Finally, the algorithm divides the result by 2 to account for the halved integration time (60 minutes for WFI versus 30 minutes for NFI).

4.4.2 Validation

When validating the NFI IPH-Estimator algorithm, synthetic data is generated using the synthetic image model. The input parameters to the model are such that the spacecraft always points towards Earth, but the photons from the Earth, its exosphere, and any stellar radiance are all set to 0. The operational mode is still set to use the WFI channel and the LyaN filter, with an integration time of 60 minutes. However, the ground-truth image is generated by setting the operational mode to use the NFI channel and the LyaN filter, with an integration time of only 30 minutes. For this section, the flat-field is turned off (it is assumed that the flat-field has already been calibrated out).

The algorithm was tested for several possible dates, from the best-case scenario to the worst-case scenario. The dates chosen were July 1st, 2025, October 1st, 2025, December 1st, 2025, March 1st, 2026, and June

Table 45: Statistics for 500 iterations the NFI IPH-Estimator algorithm, with smoothing.

Date	Annulus Start Radius [pixels]	MMMPE [%]	MMPESD [%]
July 1st, 2025	198.12	0.225	0.0700
October 1st, 2025	195.07	0.0433	0.0825
December 1st, 2025	241.87	0.150	0.228
March 1st, 2026	186.99	0.202	0.0903
June 1st, 2026	226.97	1.437	0.101

1st, 2026 (the same test dates as the WFI IPH-Estimator algorithm). For each date, 6 synthetic noisy images were generated using the synthetic image model, corresponding to one per hour for 6 hours after midnight on each date. The noisy images were then stacked and passed through the NFI IPH-Estimator algorithm to obtain the recovered NFI IPH image, $\hat{H}[i, j]$. The recovered NFI image was then compared to the ground-truth image, $H[i, j]$ in order to obtain the percent error $P[i, j]$.

An integration time of 6 hours was chosen in order to be consistent with the WFI IPH-Estimator algorithm. The percent error calculation was repeated $N_{\text{trials}} = 500$ times, and the MMMPE and MMPESD was calculated for each date. Smoothing was justified by the validation performed for the WFI IPH-Estimator algorithm, so the algorithm uses smoothing.

As with the WFI images, the requirement for accuracy on the IPH background retrieval in the NFI images is 30%. The tables of results show that the algorithm’s performance meets this benchmark. The errors in the NFI IPH recovery generally follow the same patterns as the errors in the WFI IPH recovery, including the abnormally high error rate on June 1st, 2026. This result is expected, since a similar algorithm with the same interpolated values is used for both retrievals.

5 Flat-Field Retrieval

Not all pixels in both GCI cameras respond equally to the same brightness or the same number of incident UltraViolet (UV) photons. In order to account for this, the photoelectron events are multiplied by a flat-field, denoted F_{ij} , which effectively normalizes the responsivity of all pixels.

On the ground, flat-fields are usually calibrated by taking an image (or stacking a few images) of a uniformly lit scene. After background subtraction, the result then corresponds exactly to the flat-field, since any variations in the image must be due to pixel response.

However, there are no uniformly lit scenes that could be used to calibrate flat-field while the mission is in orbit. As a diffuse, narrowband, and spatially-linear scene, the IPH is well suited as a photon source for flat-fielding. Note that the IPH retrieval algorithms in the previous section depend on the flat-field, so those algorithms cannot be used here.

When measuring flat-fields, the spacecraft will first turn off-nadir and point towards a star. Since the star's radiance is known, this immediately gives the responsivity of the center four pixels. The spacecraft then turns again to view the IPH and uses this information to calibrate the remainder of the pixels. Thus, the flat-field will be normalized to the mean of the center four pixels. The current plan is to always use these center pixels to point at a star, but this can be changed on-orbit if it is determined that the center four pixels are becoming burnt out due to heavy usage.

The flat-field will be measured once during IOC phase and monthly during the science phase. Flat-fields are acquired using the open filter position regardless of science filter selection, since the flat-field itself is weakly dependent on filter transmissivity variations. To smear any potential gradients in the IPH scene, the flat will consist of 4 stacked images, each acquired along a common, off-nadir bore sight but at a 90° roll angle with respect to one another. Due to the spatially linear structure of the IPH, as seen in the Section 4, the stacked image will be an approximately uniform scene, subject to instrument background and noise.

5.1 Algorithm

The algorithm itself is quite straightforward. After taking the four images at an off-nadir boresight 90° offset from each other, the algorithm stacks these images, calculates the mean of the four pixels at the center of the image, then divides the whole image by this average. The result is the recovered flat-field.

5.2 Experimental Validation

The flat-field retrieval algorithm is validated on two different flat-fields. The first is the ideal flat-field, where all pixels respond equally to the same brightness, while the second is the flat-field found by testing the camera in the lab. The first flat is mainly used as a benchmark in the retrieval, while we are really interested in how the algorithm performs on the real flat-field. Ground-truth flat-field images are generated separately for the WFI and NFI channels. The WFI flat-fields are shown in Figures 28a and 28b, while the NFI flat-fields are shown in Figures 29a and 29b. Note that the dark spot in the real flat is likely a lab error and not part of our true flat. However, any retrieval algorithm that is able to recover an anomalous spot in the flat-field will certainly also work for a more well-behaved flat.

For the remainder of this section, images generated from the synthetic image model for the purposes of validating the flat-field are denoted synthetic flat images. The synthetic flat images are generated along 4 roll angles in 90° increments around a common boresight direction. For the validation below, the boresight direction is chosen to be latitude 0° longitude 180° . The IPH appears to have a well-behaved linear gradient in the region around this boresight, as shown in Figure 30.

To verify this choice, the mean IPH at this boresight direction is rolled in 90° increments, then stacked, producing the mean IPH signal shown in Figure 31. The variations in the signal are on the order of 0.15%, which is nearly ideal. The NFI FOV is smaller than the WFI, so the variation is expected to be even smaller for the NFI flat-field retrieval.

To generate synthetic flat images, the input parameters to the synthetic image model are such that the spacecraft always points at the specified boresight direction. The operational mode is set to use the WFI or NFI channel, depending on which flat-field is being validated, and use the narrowband open filter, with an

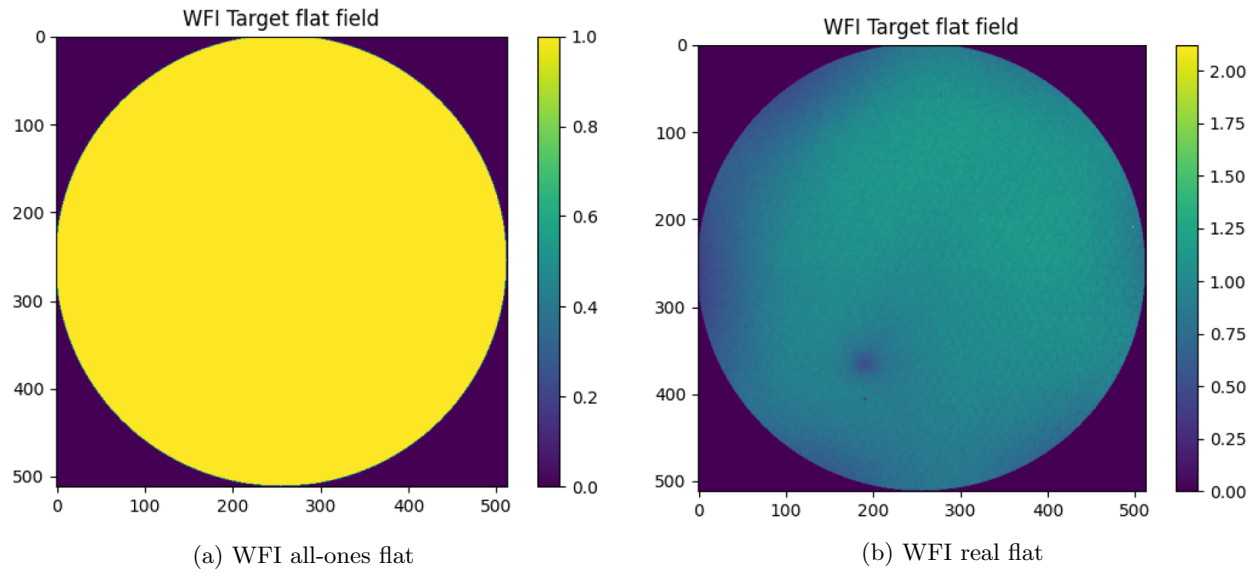


Figure 28: Flat-field ground-truths for the WFI channel




Figure 29: Flat-field ground-truths for the NFI channel




Figure 30: IPH Map. The scale is not displayed, nor is it important - only the relative gradients are needed.




Figure 31: Result from stacking four synthetic rolled WFI images of IPH at boresight latitude 0° and longitude 180°

Table 46: Statistics for mean percent error over 500 iterations of the flat-field retrieval algorithm for all flat-fields.

Flat	Max [%]	Min [%]	Mean [%]	Median [%]	Absolute Median [%]	Std [%]
WFI, all-ones	0.331	-0.369	0.00928	0.0120	0.0855	0.0733
WFI, real	0.528	-0.245	0.183	0.184	0.184	0.0758
NFI, all-ones	0.710	-0.743	-0.0177	-0.0177	0.0991	0.146
NFI, real	1.148	-0.632	0.168	0.168	0.173	0.152

Table 47: Statistics for percent error std over 500 iterations of the flat-field retrieval algorithm for all flat-fields.

Flat	Max [%]	Min [%]	Mean [%]	Std [%]
WFI, all-ones	1.803	1.144	1.583	0.0491
WFI, real	2.880	1.140	1.612	0.136
NFI, all-ones	4.204	2.689	3.630	0.113
NFI, real	8.191	2.577	3.750	0.317

integration time of 90 minutes for each roll angle. The specific flat-field used depends on which flat-field is being validated.

Once the flat images have been generated by the synthetic image model and the retrieved flat has been calculated by the algorithm, the percent error for each trial is calculated by subtracting the target flat from the retrieved flat, dividing by the target flat, then multiplying by 100. The result of this, in the WFI case, is a 512×512 array of percent errors for each WFI binned pixel.

This retrieval is then repeated for $N_{\text{trials}} = 500$ independent iterations, resulting in a datacube of percent errors P_{kij} with dimensions $500 \times 512 \times 512$, indexed k for trial, i for pixel row, and j for pixel column respectively. The mean percent error μ_{ij} and percent error standard deviation σ_{ij} for each pixel ij is then calculated:

$$\mu_{ij} = \frac{1}{N} \sum_{k=1}^N (P_{kij}) \quad (46)$$

$$\sigma_{ij} = \sqrt{\frac{\sum_{k=1}^N (P_{kij} - \mu_{ij})^2}{N}} \quad (47)$$

The NFI percent error datacube has $500 \times 1024 \times 1024$ pixels instead, but the calculations remain the same.

Finally, the maximum, minimum, mean, and standard deviation of μ_{ij} and σ_{ij} over all ij pixels is compiled for each flat-field in Tables 46 and 47. The median of μ_{ij} and the median of $|\mu_{ij}|$ (denoted absolute median) is also calculated in order to get more information on the mean percent error. Corner pixels are not included in these statistics because they are not part of the retrieved flat. For display purposes, these pixels are set to have 0 percent error. Figures 32, 33, 34, and 35 show the mean percent error and percent error standard deviation per pixel.

There are some key observations:

1. The mean percent error for the WFI channel is largely driven by the variations in the underlying IPH image.
2. The mean percent error for the NFI channel is largely driven by the noise instead of variations in the underlying IPH image. This is likely because the NFI's FOV is smaller than the WFI's FOV, resulting in less variation in the background IPH image.

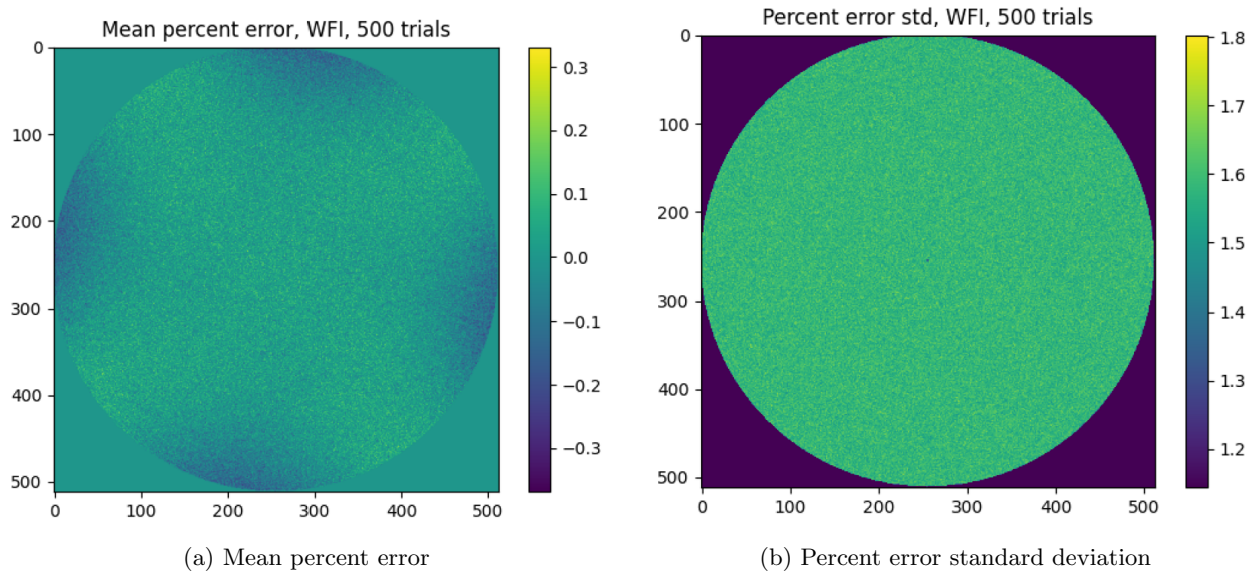


Figure 32: WFI percent error statistics per pixel over 500 iterations when retrieving all-ones flat

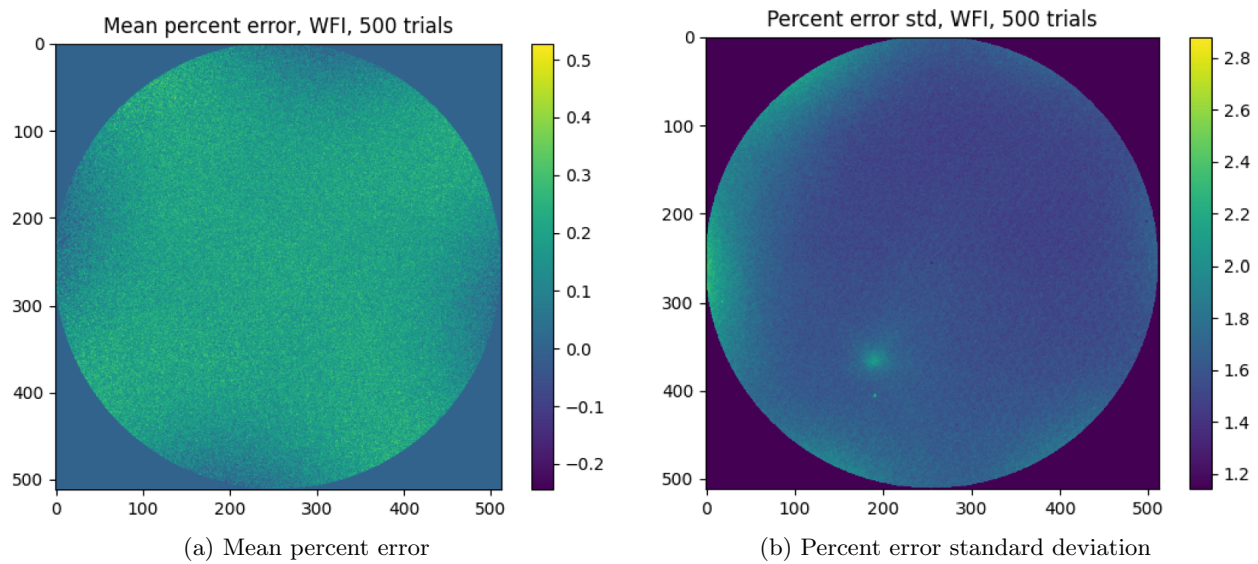


Figure 33: WFI percent error statistics per pixel over 500 iterations when retrieving real flat

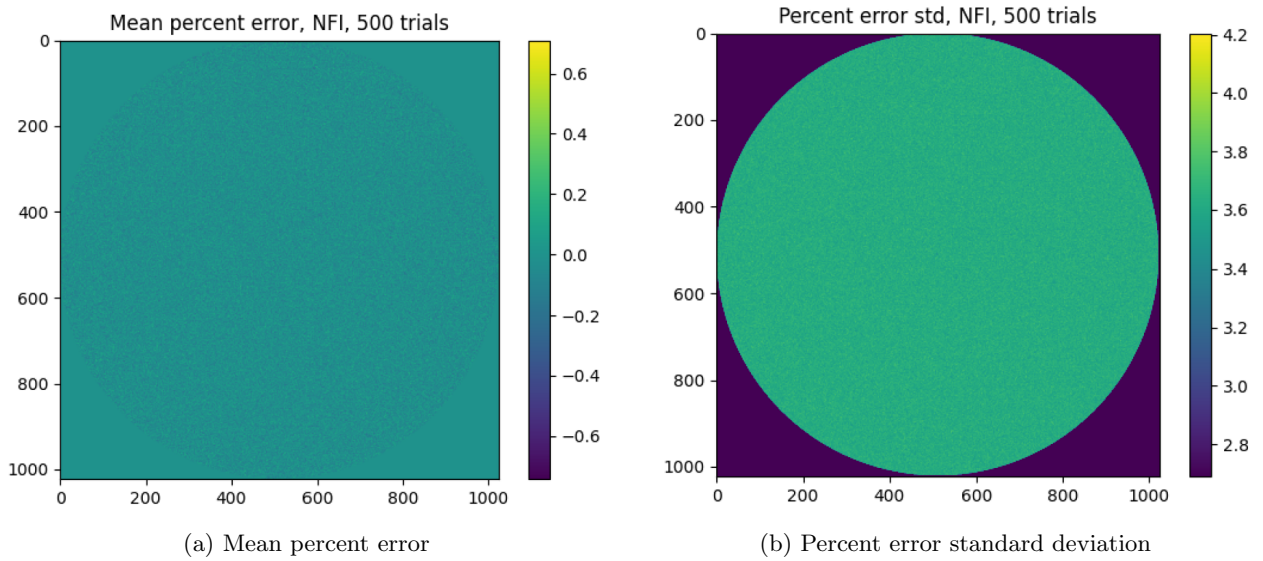


Figure 34: NFI percent error statistics per pixel over 500 iterations when retrieving all-ones flat

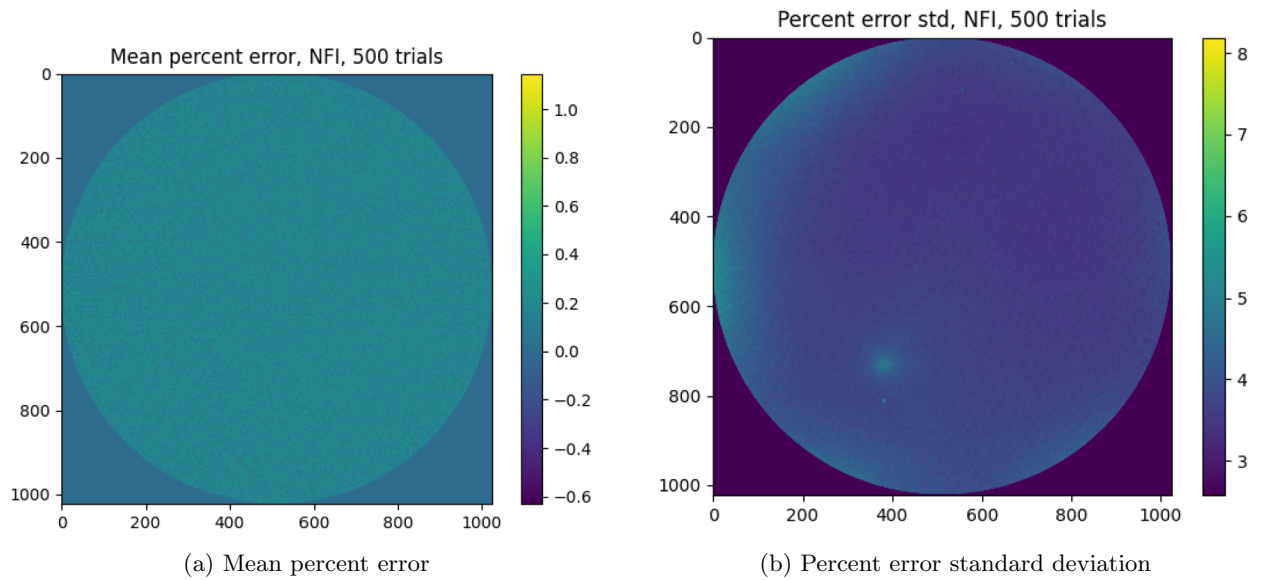


Figure 35: NFI percent error statistics per pixel over 500 iterations when retrieving real flat




Figure 36: Example of percent error when recovering all-ones flat (1 trial)

3. The percent error standard deviations are largely driven by the shape of the flat-field. Pixels with low flat-field value tend to have higher standard deviations, while pixels with high flat-field value tend to have lower standard deviations.
4. The middle four pixels have a much smaller standard deviation compared to other pixels. These are the pixels that the algorithm uses the normalize the stacked image, so having a smaller standard deviation makes sense. This phenomenon is more visible in the WFI images, but still present in the NFI images.

The target retrieval accuracy is 5%. The mean percent errors appear to be well below this threshold. However, the standard deviations appear to be quite high, especially in the NFI case. These statistics are deceiving; an example percent error from a single trial for the all-ones flat case, shown in Figure 36, show that the magnitude of the percent error for some individual pixels do stray above the target retrieval accuracy. The reasons for this behavior are explained in the next section.

5.3 Theoretical Validation

The synthetic image model must be closely examined in order to validate these results theoretically. For each individual pixel, the model adds all of the following, in units of electron-counts, in order to get the final signal for each pixel, in units of DN, in the final image:

- Dark current, which is Poisson distributed. In this section, this is approximated as a Gaussian distribution since the mean is far over 1000 electron counts.
- APS radiation, which is Poisson distributed. In this section, this is approximated as a Gaussian distribution since the mean is far over 1000 electron counts.
- Read noise, which is Gaussian distributed with mean 0.
- Signal, MCP radiation, MCP gain, and flat (denoted “real signal” for the remainder of this section), which is assumed to be Gaussian distributed by the law of large numbers (the 90-minute integration time stacks enough frames for this to hold true).

Since these are all Gaussian distributed and added, each pixel is therefore still a Gaussian distribution, with mean and variance equal to the sum of each distribution’s mean and variance. Finally, each pixel is multiplied by the ADC gain G_{APS} , which also converts the units from electron counts to DN.

For this section, each pixel is the sum of four pixels across four different independent images. However, since each pixel in these four different images is Gaussian distributed, the pixel in the final image is still Gaussian distributed. This yields the final synthetic flat image, in units of DN.

5.3.1 All-ones Flat-Field Retrieval

We first explain the experimental results for the all-ones flat-field retrieval. When processing the image on the all-ones flat, the following steps are taken in order to get the mean percent error over N trials:

1. The mean dark current and mean APS radiation is subtracted out. This affects the mean but not the variance of each pixel.
2. Each pixel is normalized by dividing by the mean of the center four pixels. Since we are working with the all-ones flat, it is safe to assume that the mean of the center four pixels is the mean of each pixel, since all pixels have the same mean. This is an ideal approximation; true errors are likely to be a bit higher than those derived here due to this step. Each pixel now represents the recovered flat-field.
3. Subtract the ground-truth flat. Here, the ground-truth flat is 1 for every pixel. The mean is now 0.
4. Divide by the ground-truth flat. Here, the ground-truth flat is 1 for every pixel, so this does nothing. Each pixel is now the percent error from the recovered flat-field.
5. The above steps (generating a new synthetic flat image, then performing this processing) are repeated N_{trials} times and the mean is taken. This is equivalent to adding N_{trials} of these independent variables together, then dividing by N_{trials} . This does not change the mean, but the variance will decrease by a factor of N_{trials} overall - variance scales linearly (by a factor of N_{trials} here) when adding independent variables, but decreases by a factor of N_{trials}^2 when dividing out N_{trials} . This is our final mean percent error.

All of these operations on a Gaussian distribution result in another Gaussian distribution. Therefore, each pixel in the final mean percent error is Gaussian distributed, with mean 0 and variance σ^2 . Let X denote a Gaussian-distributed random variable with mean 0 and variance σ^2 .

The target retrieval accuracy is 5%. Therefore, in order for the retrieval to be successful, we require that X deviates from 0 by less than 0.05. Mathematically, we have

$$\begin{aligned}
 P(|X| \geq 0.05) & & (48) \\
 &= 2P(X \geq 0.05) \\
 &= 2P\left(\frac{X}{\sigma} \geq \frac{0.05}{\sigma}\right) \\
 &= 2P\left(N(0, 1) \geq \frac{0.05}{\sigma}\right) \\
 &= 2Q\left(\frac{0.05}{\sigma}\right) & (49)
 \end{aligned}$$

The Q function is defined as the tail distribution function of the standard normal distribution.

In order to calculate the actual probability, σ must be known. Therefore, we now pull the actual means and variances of all the signals and processes from the synthetic image model. The WFI case is shown in Table 48, while the NFI case is shown in Table 49. The means and variances for dark and APS radiation are all for a single image, while the mean and variance for the real signal is from the synthetic flat image, but divided by 4 to get the mean and variance for the real signal in a single image. Within the tables, the total signal from the single image is then scaled by multiplying the mean by $4G_{\text{APS}}$ to get the synthetic flat image distribution. The variance is scaled by $4G_{\text{APS}}^2$, since the factor of 4 comes from adding independent variables, while the factor of G_{APS} comes from scaling each variable. The mean dark current and mean APS radiation is then subtracted from each pixel during processing. Note that in order for the subtraction

Table 48: Mean and variance of all signals and processes used to generate the WFI synthetic flat image.

Signal/Process	Mean	Variance	Units
Dark current	8.64e7	8.64e7	counts
APS radiation	4.4e5	4.4e5	counts
Read noise	0	2.1e4	counts
Real signal	4.9e7	1.88e12	counts
Total signal, single image	1.36e8	1.88e12	counts
Synthetic flat image	3.4e7	2.94e10	DN
After mean dark and mean APS subtraction	1.27e7	2.94e10	DN

Table 49: Mean and variance of all signals and processes used to generate the NFI synthetic flat image.

Signal/Process	Mean	Variance	Units
Dark current	2.16e7	2.16e7	counts
APS radiation	1.1e5	1.1e5	counts
Read noise	0	1.05e4	counts
Real signal	9.0e6	3.47e11	counts
Total signal, single image	3.07e7	3.47e11	counts
Synthetic flat image	7.68e6	5.42e9	DN
After mean dark and mean APS subtraction	2.22e6	5.42e9	DN

to be correct, the mean dark current and mean APS radiation are both also multiplied by $4G_{\text{APS}}$ before subtraction.

We can verify these statistics in practice by collecting the values from all of the pixels from the synthetic flat image (before normalization) and plotting them as a histogram. Since this is still the all-ones flat, all pixels have the same distribution. The results are shown in Figure 37. Note that both look like Gaussian distributions as well. In the WFI case, the mean was $1.21e6$ and the variance was $2.93e10$, while in the NFI case, the mean was $2.26e6$ and the variance was $5.41e9$. These are close to the theoretical values derived above.

For the WFI case, after running through the calculations described at the beginning of this section, each pixel in the percent error array for a single trial is a Gaussian distribution with mean 0 and variance $1.820e-4$. When Equation 49 is evaluated for this variance, the result is $2e-4$. There are approximately $2e5$ pixels within the FOV for a WFI image, implying that in practice, some pixels will have more than a 5% error. This agrees with the results from Figure 36a, where approximately 0.04% of the pixels within the FOV had over a 5% error - very close to the theoretical value of 0.02%.

Each pixel in the mean percent error array for $N_{\text{trials}} = 500$ is a Gaussian distribution with mean 0 and variance $3.643e-7$. When Equation 49 is evaluated for this variance, the result is $7.76e-1463$. In practice, no pixel will have more than a 5% error. Moreover, if Equation 49 is recalculated using 0.2% as the threshold for error, the corresponding expression evaluates to 0.000932. This implies that there will be some, but still very few, pixels above 0.2%. Both of these results agree with the experimental mean percent error in Figure 32.

For the NFI case, after running through the calculations described at the beginning of this section, each pixel in the percent error array for a single trial is a Gaussian distribution with mean 0 and variance 0.00110. When Equation 49 is evaluated for this variance, the result is 0.131. This implies that many pixels will have more than a 5% error for a single trial. In practice, the results line up pretty well: in Figure 36b, approximately 13.9% of the pixels within the FOV had over a 5% error.

Each pixel in the mean percent error array for $N_{\text{trials}} = 500$ is a Gaussian distribution with mean 0 and variance $2.199e-6$. When Equation 49 is evaluated for this variance, the result is $8.12e-239$. In practice, no pixel will have more than a 5% error. Moreover, if Equation 49 is recalculated using 0.6% as the threshold




Figure 37: Distribution of synthetic flat image (“incoming pixels”) before normalization

Table 50: Statistics for individual pixel distributions over $N_{\text{trials}} = 500$ iterations of retrieving true flat.

Pixel	Mean	Std	Target	Percent error std
WFI High Pixel	2.128	0.026	2.125	1.22
WFI Low Pixel	0.421	0.0095	0.420	2.26
NFI High Pixel	1.579	0.0482	1.576	3.06
NFI Low Pixel	0.563	0.0263	0.563	4.67

for error, the corresponding expression evaluates to $5.3e-5$. This implies that there will be some pixels with over 0.6% error. Both of these results agree with the experimental mean percent error in Figure 34.

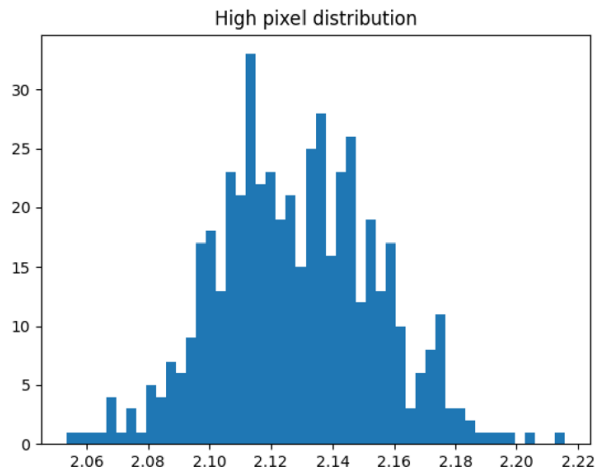
5.3.2 True Flat-Field Retrieval

When retrieving the true flat-field, the normalization step becomes more difficult because we can no longer assume that we are dividing by the mean of the pixel. The normalization step now involves dividing one Gaussian distributed random variable (the current pixel) by another Gaussian distributed random variable (the mean of the center four pixels). This does not necessarily result in another Gaussian random variable.


Therefore, instead of calculating a general formula for all pixels, we simply target individual pixels as examples. The retrieved flat-field value is collected for an individual pixel for all $N_{\text{trials}} = 500$, and plotted on a histogram. The mean and variance are also collected. The individual pixels targeted are chosen to be the extremities of the flat-field, since those are the pixels that correspond to the extremities of the resulting percent error standard deviation. For the remainder of this section, “high pixel” denotes a pixel with a high true flat-field, while “low pixel” denotes a pixel with a low true flat-field.

Figures 38 and 39 show the distributions. While the results look Gaussian, it is not possible to conclude that these individual pixels are Gaussian distributed. Table 50 shows the statistics for each individual pixel. Recall that to calculate the percent error, the target flat-field is subtracted from the retrieved flat-field, then this result is divided by the target flat-field and multiplied by 100. As shown in the table, low pixels generally have higher noise (higher stds) than high pixels, while the NFI generally has higher noise compared to the WFI. These conclusions are seen experimentally in Figures 33 and 35 as well.

From the experimental and theoretical validation, we conclude that the flat-field retrieval algorithm has high accuracy, but low precision. It is almost guaranteed that some individual pixels will have more than 5% error when doing a single flat-field retrieval. One method of combatting this is to stack more frames, which would drive the variance down. More work is needed to determine if this is sufficient.




(a) High pixel




(b) Low pixel

Figure 38: Distributions of retrieved WFI true flat-field for individual pixels



(a) High pixel



(b) Low pixel

Figure 39: Distributions of retrieved NFI true flat-field for individual pixels

6 Conclusion

The calibration algorithms to be used for the Carruthers Geocorona Observatory have been extensively characterized and validated, paving the way for an accurate recovery of the global 3D H-density distribution. The retrieved instrument background (bias and dark current) from the lab data was found to be peculiar, but the algorithm itself was proved to be sufficient when tested on the PATHS cameras. The science image corners were found to be adequate to estimate the energetic particle radiation background. The IPH estimation and retrieval algorithm was found to be sufficient and far surpassed the target retrieval accuracy in most conditions using the synthetic image model. On the other hand, the flat-field retrieval algorithm was validated using the synthetic image model, where it was found that the algorithm results could not meet the target retrieval accuracy threshold.

Below are future work that need to be done before these algorithms are completely finalized:

- The lab data for instrument background calibration needs to be re-acquired in order to get a more accurate measurement of the bias and dark current.
- In Section 3, many approximations were used in the SNR derivations for the Particle Radiation retrieval. A more robust, analytical SNR needs to be performed with true values (instead of estimated values). The main results are not expected to change, but this derivation should still be done for the sake of rigor.
- A rigorous analysis of the effects of smoothing in the IPH Retrieval algorithm need to be performed in order to make sure that none of the variations are artificially introduced by the smoothing itself. These artificial variations have the potential to affect the tomographic analysis and should be avoided.
- The IPH Retrieval algorithm was found to be sufficient for a large majority of the spacecraft's orbit where the IPH linearity assumption holds. However, the algorithm's error shoots up when this assumption breaks down. Possible methods to combat this were discussed in Section 4.3.5; more work needs to be done to decide which method is best.
- The flat-field retrieval algorithm was found to not be sufficient to meet the target retrieval accuracy. The algorithm itself was relatively simplistic - a new algorithm, likely relying on deconvolution, will need to be designed and validated.
- Both the IPH Retrieval algorithm and the flat-field retrieval algorithm assumed perfect knowledge of the instrument background. Their sensitivity to bias in instrument background still needs to be tested.

The development and finalization of these calibration algorithms is an important step towards the successful operation of the Carruthers Geocorona Observatory, which in turn will give crucial insights into the fundamental nature of the exosphere.

References

- [1] A. B. Christensen et al. “Initial observations with the Global Ultraviolet Imager (GUVI) in the NASA TIMED satellite mission”. In: *Journal of Geophysical Research: Space Physics* 108.A12 (2003). DOI: <https://doi.org/10.1029/2003JA009918>. eprint: <https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/2003JA009918>. URL: <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2003JA009918>.
- [2] Jeffrey W. Kruk et al. “Radiation-induced Backgrounds in Astronomical Instruments: Considerations for Geosynchronous Orbit and Implications for the Design of the WFIRST Wide-field Instrument”. In: *Publications of the Astronomical Society of the Pacific* 128.961 (Feb. 2016), p. 035005. DOI: 10.1088/1538-3873/128/961/035005. URL: <https://dx.doi.org/10.1088/1538-3873/128/961/035005>.
- [3] R. A. Mewaldt et al. “Solar energetic particles and their variability from the sun and beyond”. In: *AIP Conference Proceedings* 1539.1 (2013), pp. 116–121. DOI: 10.1063/1.4811002. eprint: <https://aip.scitation.org/doi/pdf/10.1063/1.4811002>. URL: <https://aip.scitation.org/doi/abs/10.1063/1.4811002>.
- [4] R.A. Mewaldt. “Solar Energetic Particle Composition, Energy Spectra, and Space Weather”. In: *Space Science Reviews* 124 (1 2006), pp. 303–316. DOI: <https://doi.org/10.1007/s11214-006-9091-0>. URL: <https://doi.org/10.1007/s11214-006-9091-0>.
- [5] Wayne R. Pryor et al. “Lyman- α Models for LRO LAMP from MESSENGER MASCS and SOHO SWAN Data”. In: *Cross-Calibration of Far UV Spectra of Solar System Objects and the Heliosphere*. Ed. by Eric Quémerais, Martin Snow, and Roger-Maurice Bonnet. New York, NY: Springer New York, 2013, pp. 163–175. ISBN: 978-1-4614-6384-9. DOI: 10.1007/978-1-4614-6384-9_5. URL: https://doi.org/10.1007/978-1-4614-6384-9_5.
- [6] D.J. Strickland et al. “Atmospheric Ultraviolet Radiance Integrated Code (AURIC): theory, software architecture, inputs, and selected results”. In: *Journal of Quantitative Spectroscopy and Radiative Transfer* 62.6 (1999), pp. 689–742. ISSN: 0022-4073. DOI: [https://doi.org/10.1016/S0022-4073\(98\)00098-3](https://doi.org/10.1016/S0022-4073(98)00098-3). URL: <https://www.sciencedirect.com/science/article/pii/S0022407398000983>.