

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

Understanding Galactic AI-26 with the Compton Spectrometer and Imager

Permalink

<https://escholarship.org/uc/item/9qp8f40t>

Author

Beechert, Jacqueline Nicole

Publication Date

2023

Peer reviewed|Thesis/dissertation

Understanding Galactic ^{26}Al with the Compton Spectrometer and Imager

By

Jacqueline Nicole Beechert

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Physics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Dr. John Tomsick, Co-chair
Professor Stuart Bale, Co-chair
Professor Daniel Kasen
Professor Marjorie Shapiro

Spring 2023

Understanding Galactic ^{26}Al with the Compton Spectrometer and Imager

Copyright 2023
by
Jacqueline Nicole Beechert

Abstract

Understanding Galactic ^{26}Al with the Compton Spectrometer and Imager

by

Jacqueline Nicole Beechert

Doctor of Philosophy in Physics

University of California, Berkeley

Dr. John Tomsick, Co-chair

Professor Stuart Bale, Co-chair

Aluminum-26 (^{26}Al) is a radioactive isotope produced in massive star processes. High-resolution spectroscopy of its 1.809 MeV gamma-ray (γ -ray) decay signature constrains the dynamics of its emission as it is ejected from its progenitor sites and incorporated into the interstellar medium of the Milky Way Galaxy. Imaging reveals dominant emission in the Inner Galaxy and emission in localized regions of massive star activity. Taken together, spectroscopy and imaging of ^{26}Al shed light on the chemical evolution of the Galaxy over millions of years.

The Compton Spectrometer and Imager (COSI) is a compact Compton telescope designed to measure astrophysical γ -rays of energy 0.2–5 MeV. Its high-purity germanium detectors track incident photons as they Compton scatter throughout the detector volume. In 2016, COSI flew on a NASA ultra-long duration balloon for 46 days. This dissertation details the first analysis of ^{26}Al in the flight data and reports a measurement of 3.7σ significance above background and an Inner Galaxy flux of $(8.6 \pm 2.5) \times 10^{-4} \text{ ph cm}^{-2} \text{ s}^{-1}$. All scientific achievements from the flight are predicated on calibrations performed before launch. These calibration procedures were repeated in advance of an intended 2020 balloon flight. Analyzing calibration data validates instrument performance and informs studies of detector effects, including charge sharing and charge trapping, that complicate the measurement process.

The next generation of COSI as a NASA Small Explorer satellite is slated for launch in 2027. It is anticipated to strengthen the spectroscopic measurement of ^{26}Al in the balloon flight and yield the most detailed images of ^{26}Al to date. Extensive testing of the imaging algorithm in COSI's new analysis toolkit is presented as a first step towards producing these images. The desired culmination of these efforts is an enhanced understanding of Galactic ^{26}Al by way of thorough calibrations, novel insight into undesirable detector effects, and advanced analyses of high-resolution data from the COSI balloon and satellite instruments.

Contents

Contents	i
List of Figures	vi
List of Tables	xiii
Acknowledgments	xiv
1 Introduction	1
2 Aluminum-26 Science	3
2.1 History and observations of Galactic Aluminum-26	4
2.2 Dominant progenitors	7
2.2.1 Multi-wavelength analyses	8
2.2.2 Theoretical estimates	10
2.2.3 Complementary observations	12
2.2.3.1 Iron-60	12
2.2.3.2 Positrons	14
2.2.3.3 Geologic searches	14
2.2.4 Summary	16
2.3 Kinematics	16
2.3.1 Injection along Galactic rotation	17
2.3.2 Motion within localized regions	19
2.3.3 Summary	20
2.4 Summary of current and future observations	22
3 Soft Gamma-ray Instrumentation and COSI	25
3.1 Gamma-ray telescopes	25
3.1.1 Coded mask	26
3.1.2 Classic Compton Telescope	27
3.1.3 Compact Compton Telescope	27
3.1.4 Laue lens	28
3.2 Exploring the MeV gap with germanium CCTs	29

3.2.1	The MeV gap	29
3.2.2	High-purity germanium detectors	32
3.3	Principles of gamma-ray detection in Compton telescopes	34
3.4	The Compton Spectrometer and Imager	37
3.4.1	Detectors	37
3.4.2	Cryostat	37
3.4.3	Cryocooler and thermal control	39
	3.4.3.1 The Liquid Cooling System	41
	3.4.3.2 Thermal environment during the COSI 2016 flight	42
3.4.4	Shields	43
3.4.5	Electronics	45
3.4.6	Aspect system	46
3.4.7	NASA Balloon Program	46
3.4.8	Flight heritage	47
3.5	The COSI 2016 balloon flight	48
3.5.1	Flight overview	48
3.5.2	Scientific achievements	49
	3.5.2.1 Positron-electron annihilation	50
	3.5.2.2 GRB detection and polarization	52
	3.5.2.3 Detection of the Crab Nebula	53
3.5.3	Attempted 2020 mission	53
3.6	Analysis methods	54
3.6.1	MEGAlib	54
3.6.2	cosipy and COSI Data Challenges	55
4	Calibrations of the Compton Spectrometer and Imager	56
4.1	Calibration pipeline	57
4.2	Data collection	58
	4.2.1 Low-energy configuration	58
	4.2.2 High-energy configuration	59
4.3	Energy calibration	60
	4.3.1 Linearity investigations	62
	4.3.2 Temperature correction	65
4.4	Cross-talk correction	67
4.5	Strip pairing algorithm	68
4.6	Depth calibration	69
	4.6.1 “Classic” depth calibration	70
	4.6.2 Comparison to 2016	71
	4.6.3 Future improvements	71
4.7	Detector Effects Engine	73
4.8	Instrument performance	73
	4.8.1 Energy resolution	74

4.8.2	Angular resolution	75
4.8.3	Effective area	77
4.8.4	Polarization response	80
4.9	Summary	81
5	Measurement of Galactic Aluminum-26 in the COSI 2016 Flight	83
5.1	Data selection	84
5.2	Residual method	87
5.3	Maximum likelihood method	90
5.3.1	Sky model	91
5.3.2	Background model	92
5.3.3	Propagating background uncertainties in a joint fit	95
5.3.4	Results	95
5.3.4.1	Signal region	95
5.3.4.2	Line parameters	96
5.3.5	Method validation	98
5.3.5.1	Different template maps	99
5.3.5.2	Signal region altitude	99
5.3.5.3	Background region altitude	100
5.3.5.4	Separate 10-, 9-detector portions	100
5.3.5.5	Rigidity	102
5.3.5.6	Broader energy range	104
5.3.5.7	Systematic uncertainties in flight data analysis	104
5.3.6	Flux implications	106
5.3.6.1	Mass of Aluminum-26	106
5.3.6.2	Positron production rate	107
5.3.6.3	Core-collapse supernova rate	107
5.4	Validating the method with simulations	108
5.4.1	Simulated data sets	108
5.4.2	Complete flight simulation	109
5.4.3	Simulations with different template maps	111
5.4.4	Increasing the signal	113
5.4.5	Background-only simulations	113
5.5	Discussion	114
5.5.1	Comparison to previous measurements	114
5.5.2	Systematic uncertainties	116
5.5.3	Galactic diffuse continuum emission	117
5.6	Summary	119
5.7	Future work	120
6	Imaging with the Compton Spectrometer and Imager	121
6.1	The Compton Data Space	121

6.2	Instrument response	123
6.3	Model fitting	124
6.4	Richardson-Lucy deconvolution	126
6.4.1	Overview	126
6.4.2	Richardson-Lucy Algorithm	127
6.5	Testing the Richardson-Lucy algorithm	128
6.5.1	“Fixed” orientation	129
6.5.1.1	Point sources	129
6.5.1.2	Extended sources	130
6.5.1.3	Few photons	131
6.5.2	“Simple” orientation	135
6.5.2.1	Point sources	135
6.5.2.2	Extended sources	137
6.5.3	“Fine” orientation	137
6.5.3.1	Point sources	137
6.5.3.2	Extended sources	140
6.5.4	“Step 6 degrees” orientation	140
6.5.4.1	Point sources	140
6.5.4.2	Extended sources	144
6.5.5	Scanning the Galactic Plane	144
6.5.5.1	Point sources	144
6.5.5.2	Extended sources	146
6.5.6	Summary	147
6.6	Data Challenge 1	148
6.6.1	Point sources	149
6.6.1.1	No background	150
6.6.1.2	Including background	151
6.6.2	Positron-electron annihilation	153
6.6.2.1	No background	153
6.6.2.2	Including background	156
6.6.3	Aluminum-26	159
6.6.3.1	No background	161
6.6.3.2	Including background	161
6.7	Summary	166
7	Conclusion	168
	Bibliography	170
A	Temperature correction	182
B	Charge sharing in GRIPS detectors	188

B.1	Overview of charge sharing	188
B.2	The GRIPS detector	189
B.3	Measurements	189
B.4	Results	191
B.5	Conclusions	197
C	Charge trapping in COSI detectors	200
C.1	Background	200
C.2	Manifestation in COSI	201
C.3	Future work	205
D	Derivation of the Compton equation	211
E	Optimization of the Compton scattering angle in Aluminum-26 analysis	213
E.1	“Step down”	214
E.2	“Step up”	216
E.3	Simultaneous optimization	218
F	Activation and atmospheric background simulations	220
F.1	Activation background	220
F.2	Atmospheric background	222
	F.2.1 Atmospheric quantities for the Ling model	223
	F.2.2 Generating an orientation file for the Ling model	224
F.3	Event selections	225

List of Figures

2.1	Energy spectrum of ^{26}Al along the Galactic Plane from HEAO-3.	5
2.2	COMPTEL 1.8 MeV images of Galactic ^{26}Al	6
2.3	SPI 1.8 MeV image and spectrum.	6
2.4	A multi-wavelength comparison of COMPTEL's 1.8 MeV image to other all-sky maps across the electromagnetic spectrum.	9
2.5	Probability from a population synthesis model that the present-day ^{26}Al in Upper Scorpius originated from Wolf-Rayet (WR) stars, supernovae (SN), or a mixture of the two.	11
2.6	Spectral intensities of the ^{60}Fe and ^{26}Al γ -ray signatures from SPI.	13
2.7	Longitude-resolved measurements of Galactic ^{26}Al by SPI.	18
2.8	SPI model of how ^{26}Al flows into superbubbles from spiral arms.	19
2.9	The Scorpius-Centaurus region as a potential target of ^{26}Al kinematic observations.	20
2.10	Velocity flows in the Large Magellanic Cloud.	21
2.11	Simulated observations of ^{26}Al assuming the sensitivities and resolutions of (a) COMPTEL and (b) the satellite version of COSI.	21
2.12	Maps of ^{26}Al from Population SYnthesis COde (PSYCO) simulations compared to COMPTEL and SPI observations.	24
3.1	Schematic view of the operating principle of SPI on INTEGRAL, a coded mask imager.	26
3.2	Instrument sensitivities and attenuation of photons in the MeV gap.	30
3.3	Dominant interaction mechanisms as a function of detector atomic number and photon energy.	33
3.4	Illustration of a typical γ -ray event and source localization in a CCT.	36
3.5	COSI's cross-strip germanium detectors (GeDs) and full detector stack.	38
3.6	The COSI cryostat.	38
3.7	The COSI cryocooler and coldfinger.	40
3.8	Sunpower CryoTel CT cryocooler lift vs. cold tip temperature.	40
3.9	The COSI liquid cooling system as seen (a) on the cryocooler itself and (b) from the top of the gondola.	41
3.10	A schematic of fluid flow in the LCS.	42
3.11	A 24-hour portion of the 2016 flight demonstrating the dramatic change in exterior temperature and the relatively stable temperature inside the electronics bay.	43

3.12	Cryocooler collar temperature during the 2016 flight as compared to the external temperature seen by the radiator plate.	44
3.13	Proper detector temperature, cold tip temperature, and cryocooler power during the 2016 flight.	44
3.14	Balloon altitude and flight path of the COSI 2016 balloon flight.	49
3.15	Exposure map of the 2016 balloon flight.	50
3.16	Image of 511 keV emission from positron-electron annihilation in the Milky Way Galaxy, as observed by SPI.	51
3.17	COSI observations of positron-electron annihilation during the 2016 balloon flight.	52
4.1	COSI calibration structures in 2016 and 2020.	60
4.2	The Melinator window showing the 2020 energy calibration of strip 8 on the DC side of detector 11.	61
4.3	Single-strip energy resolution (FWHM) as a function of energy in 2020 and 2016 energy calibrations.	63
4.4	Deviation from linearity in the 2020 energy calibration of three COSI strip electrodes.	63
4.5	Residuals of MCMC linear, second-order, third-order, and fourth-order polynomial fits to 2020 energy vs. ADC calibration data.	64
4.6	Results of MCMC polynomial fits to 2020 energy vs. ADC calibration data.	65
4.7	Temperature correction: deviation of fitted photopeak energy from the true photopeak energy.	66
4.8	Temperature correction in 2020 ^{22}Na (511.0 keV) calibration data.	67
4.9	The cross-talk effect in 2016 ^{137}Cs calibration data.	68
4.10	Illustration of the strip pairing algorithm.	69
4.11	Comparison of depth calibration in 2020 and 2016 calibration data.	72
4.12	Flowchart of calibration and event reconstruction in MEGALib.	74
4.13	Illustrations of the angular resolution of a Compton telescope.	76
4.14	The effective area as a function of energy in 2020 and 2016 calibration data.	79
4.15	The measured ASAD from partially polarized calibration data collected at SSL in 2019.	81
5.1	The signal and background regions of the COSI balloon ^{26}Al analysis displayed over the SPI 1.8 MeV image.	85
5.2	The COSI 2016 flight path displayed over the signal and background regions of the ^{26}Al analysis.	85
5.3	The full COSI 2016 flight spectrum of events which pass the signal and background region event selections (top) and the background-subtracted spectrum (bottom).	87
5.4	Residual counts between signal and background region measurements indicate an excess contribution of signal region counts of energy 1803–1817 keV.	89
5.5	COSI 2016 flight spectra in the signal and background regions of the ^{26}Al analysis.	91

5.6	The spectral sky model defined by COSI's response to the DIRBE $240\ \mu\text{m}$ map over 50 2016 flights.	92
5.7	Empirical fit to COSI flight data in the background region, with minimal event selections, which provides a smooth description of the background template shape.	94
5.8	Posterior distributions of the sky amplitude α and background amplitude β in the COSI 2016 signal region.	94
5.9	Background-subtracted spectrum and fitted sky models from the COSI 2016 flight.	97
5.10	Zenith response of COSI to 2 MeV photons at a flight altitude of 33 km, indicating strongest sensitivity to photons originating from within $\leq 35\text{--}40^\circ$	97
5.11	Summed (Sky+BG) and individual sky and background models plotted over the flight signal region spectrum.	98
5.12	Significance above background of the ^{26}Al measurement as a function of minimum flight altitude in the signal region.	100
5.13	COSI flight spectra in the signal and background regions for separate (a) 10- and (b) 9-detector portions of the flight.	101
5.14	Posterior distributions of the sky amplitude α and background amplitude β in the COSI 2016 signal region for separate (a) 10- and (b) 9-detector portions of the flight.	101
5.15	Empirical fit to COSI flight data for separate (a) 10- and (b) 9-detector portions of the flight in the background region, with minimal event selections, which provides a smooth description of the background template shape.	102
5.16	COSI 2016 flight data in the signal and background regions from the 10- and 9-detector portions of the flight, binned by Earth latitude.	103
5.17	Background-subtracted COSI 2016 spectrum of the signal and background regions after weighting the flight data by latitude, i.e. geomagnetic rigidity.	104
5.18	(a) COSI flight spectra in the signal and background regions between 1650–1950 keV. (b) Posterior distributions of the sky amplitude α and background amplitude β in this expanded energy range.	105
5.19	Empirical fit to COSI flight data in the background region, with minimal event selections, using an expanded energy range of 1650–1950 keV.	105
5.20	Combined simulations of one 2016 flight over the DIRBE $240\ \mu\text{m}$ template image, instrumental activation background, and atmospheric background in the signal and background regions.	109
5.21	Power law plus three Gaussian empirical fit to the instrumental activation and atmospheric background simulations with minimal event selections.	110
5.22	Summed (Sky+BG) and individual sky and background models plotted over the signal region spectrum in the complete flight simulation.	111
5.23	Background-subtracted spectrum and 1σ sky model contours from the complete flight simulation.	112
5.24	Significance vs. estimated Inner Galaxy flux for simulated data sets containing n DIRBE, SPI, and COMPTEL simulations of the flight combined with activation and atmospheric background simulations.	114

5.25	Distribution of the test statistic (TS) from 1000 simulated data sets.	115
5.26	The Galactic diffuse continuum emission in COSI 2016 balloon flight measurements.	118
6.1	The connection between the Compton scattering process in a detector and the Compton Data Space (CDS).	123
6.2	Visualization of the measurement process, which sorts detected photons into the three-dimensional Compton Data Space (CDS).	125
6.3	A simplified illustration of the “image space:” the two-dimensional sky is spanned by Galactic longitude ℓ and Galactic latitude b and divided into pixels of specified size.	126
6.4	Images in (a) <code>mimrec</code> and (b) <code>cosipy</code> of a simulated point source at $(\ell, b) = (3^\circ, 21^\circ)$. The instrument stares directly at the source with fixed orientation for 10^4 s.	130
6.5	Images in (a) <code>mimrec</code> and (b) <code>cosipy</code> of five simulated point sources: $(\ell, b) = (3^\circ, 21^\circ), (21^\circ, -3^\circ), (-3^\circ, -3^\circ), (21^\circ, 3^\circ),$ and $(21^\circ, 33^\circ)$	131
6.6	True images of the two-dimensional Gaussian sources, i.e. plotted directly from their functional forms (not reconstructed).	132
6.7	Images in <code>mimrec</code> of extended emission when COSI stares with fixed orientation at two-dimensional Gaussian sources.	133
6.8	Successive iterations of the <code>cosipy</code> RL algorithm when COSI stares with fixed orientation at a composite “X”-shaped source.	134
6.9	Convergence of the background amplitude, acceleration parameter, and likelihood of the “X” shape in the <code>cosipy</code> RL algorithm when COSI stares with fixed orientation.	135
6.10	The “X” and “plus” as imaged by <code>cosipy</code> when COSI stares with fixed orientation at the constituent two-dimensional Gaussian sources.	136
6.11	Images in (a) <code>mimrec</code> and (b) <code>cosipy</code> of the first ten photons of a simulated point source at $(\ell = -3^\circ, b = -3^\circ)$	136
6.12	Images in (a) <code>mimrec</code> and (b) <code>cosipy</code> of five simulated point sources: $(\ell, b) = (3^\circ, 21^\circ), (21^\circ, -3^\circ), (-3^\circ, -3^\circ), (21^\circ, 3^\circ),$ and $(21^\circ, 33^\circ)$. The instrument orientation shifts 60° after 5000 s for a total observation time of 10^4 s (“simple” orientation).	138
6.13	Images in (a, b) <code>mimrec</code> and (c, d) <code>cosipy</code> of “X”- and “plus”-shaped sources. The instrument orientation shifts 60° after 5000 s for a total observation time of 10^4 s (“simple” orientation).	139
6.14	Images in (a) <code>mimrec</code> and (b) <code>cosipy</code> of five simulated point sources: $(\ell, b) = (3^\circ, 21^\circ), (21^\circ, -3^\circ), (-3^\circ, -3^\circ), (21^\circ, 3^\circ),$ and $(21^\circ, 33^\circ)$. The instrument orientation shifts 10° off-axis every ~ 2000 s (“fine” orientation).	140
6.15	Images in (a, b) <code>mimrec</code> and (c, d) <code>cosipy</code> of “X”- and “plus”-shaped sources. The instrument orientation shifts 10° off-axis every ~ 2000 s (“fine” orientation).	141
6.16	Images in (a) <code>mimrec</code> and (b) <code>cosipy</code> of five simulated point sources: $(\ell, b) = (3^\circ, 21^\circ), (21^\circ, -3^\circ), (-3^\circ, -3^\circ), (21^\circ, 3^\circ),$ and $(21^\circ, 33^\circ)$. The instrument orientation shifts 6° off-axis every 1000 s (“step 6 degrees” orientation).	142

6.17	COSI's normalized response, weighted by the number of simulated counts in each source, compared to the true (raw) number of simulated counts in each source.	143
6.18	Images in (a, b) <code>mimrec</code> and (c, d) <code>cosipy</code> of "X"- and "plus"-shaped sources. The instrument orientation shifts 6° off-axis every 1000 s ("step 6 degrees" orientation).	145
6.19	Appearance of point sources in COSI's field of view as the instrument moves along a custom orientation file that scans the Galactic Plane.	146
6.20	Images in (a) <code>mimrec</code> and (b) <code>cosipy</code> of five simulated point sources: (-99°, -3°), (-51°, -3°), (-3°, -3°), (51°, -3°), (99°, -3°). The instrument orientation scans the Galactic Plane.	147
6.21	Images in (a, b) <code>mimrec</code> and (c, d) <code>cosipy</code> of "X"- and "plus"-shaped sources. The instrument orientation scans the Galactic Plane.	148
6.22	The combined (a) lightcurve and (b) energy spectrum in <code>cosipy</code> of the Crab Nebula, Cygnus X-1, Centaurus A, and Vela, each at 10× their true flux with no background.	150
6.23	Images in (a) <code>mimrec</code> and (b) <code>cosipy</code> of the Crab Nebula, Cygnus X-1, Centaurus A, and Vela, each at 10× their true flux with no background.	151
6.24	Convergence of the background amplitude, acceleration parameter, and likelihood of the Crab Nebula, Cygnus X-1, Centaurus A, and Vela without background in the <code>cosipy</code> RL algorithm.	152
6.25	The combined (a) lightcurve and (b) energy spectrum in <code>cosipy</code> of the Crab Nebula, Cygnus X-1, Centaurus A, and Vela, each at 10× their true flux with atmospheric background.	153
6.26	Images in (a) <code>mimrec</code> and (b) <code>cosipy</code> of the Crab Nebula, Cygnus X-1, Centaurus A, and Vela, each at 10× their true flux with atmospheric background.	154
6.27	Convergence of the background amplitude, acceleration parameter, and likelihood of the Crab Nebula, Cygnus X-1, Centaurus A, and Vela with atmospheric background in the <code>cosipy</code> RL algorithm.	155
6.28	Images in (a) <code>mimrec</code> and (b) <code>cosipy</code> of 511 keV emission, simulated at 10× its true flux with no background.	157
6.29	Convergence of the background amplitude, acceleration parameter, and likelihood of 511 keV emission without background in the <code>cosipy</code> RL algorithm.	158
6.30	Images in (a) <code>mimrec</code> and (b) <code>cosipy</code> of 511 keV emission, simulated at 10× its true flux with atmospheric background.	159
6.31	Convergence of the background amplitude, acceleration parameter, and likelihood of 511 keV emission with atmospheric background in the <code>cosipy</code> RL algorithm.	160
6.32	Images in (a) <code>mimrec</code> and (b) <code>cosipy</code> of the 1809 keV signature of ²⁶ Al, simulated at 10× its true flux with no background.	162
6.33	Convergence of the background amplitude, acceleration parameter, and likelihood of ²⁶ Al emission without background in the <code>cosipy</code> RL algorithm.	163
6.34	Images in (a) <code>mimrec</code> and (b) <code>cosipy</code> of the 1809 keV signature of ²⁶ Al, simulated at 10× its true flux with atmospheric background.	164

6.35	Convergence of the background amplitude, acceleration parameter, and likelihood of ^{26}Al emission with atmospheric background in the <code>cosipy</code> RL algorithm. . . .	165
6.36	An ensemble of simulations of 511 keV emission as observed by COSI-SP illustrates the expected variation in source morphology with statistical differences between realizations.	167
A.1	Temperature correction in 2020 data: 81.0 keV, 122.0 keV, and 136.0 keV.	183
A.2	Temperature correction in 2020 data: 276.4 keV, 302.9 keV, and 356.0 keV.	184
A.3	Temperature correction in 2020 data: 383.0 keV, 661.7 keV, and 898.0 keV.	185
A.4	Temperature correction in 2020 data: 1173.2 keV, 1274.5 keV, and 1332.5 keV.	186
A.5	Temperature correction in 2020 data: 1836.0 keV.	187
B.1	Schematic of charge sharing in a detector with strip electrodes.	188
B.2	Charge sharing measurements. (a) Spare GRIPS GeD with 149 strips per side of 0.5 mm strip pitch. (b) Experimental setup.	190
B.3	Sample spectra collected on November 10, 2021 from channel 2 on the DC side of the spare GRIPS GeD.	191
B.4	Percentage of triggered strips as a function of energy and (a) DC- and (b) AC-side placement.	192
B.5	Summed charge (^{241}Am and ^{57}Co) on middle channels 1 and 2 plotted against the charge on channel 2. Only events which trigger channels 1 and 2 (zero energy on channels 0 and 3) are considered.	194
B.6	Summed charge (^{57}Co and ^{137}Cs) on middle channels 1 and 2 plotted against the charge on channel 2. Only events which trigger channels 1 and 2 (zero energy on channels 0 and 3) are considered.	195
B.7	Left column: Summed charge on channels 0 and 2 for ^{241}Am placed on the DC and AC side of the detector. Events deposit zero energy above threshold on channels 1 and 3. Right column: Summed charge on channels 0 and 3, where events deposit zero energy above threshold on channels 1 and 2.	196
B.8	Percentage of total event energy measured on each channel for ^{241}Am , ^{57}Co , and ^{137}Cs placed on the DC (left) and AC (right) sides of the detector.	198
C.1	Depth distribution of all 652–672 keV single-site events from COSI 2016 (left) and 2020 (right) ^{137}Cs calibration data.	202
C.2	Spectra in the nine most populated depth bins of detector 0 in COSI (a) 2016 and (b) 2020 ^{137}Cs calibration data.	203
C.3	The centroid of the Gaussian fit to COSI ^{137}Cs spectra as a function of interaction depth in detector 0.	204
C.4	Evidence from Hull et al. (2014) of electron trapping in a segmented planar HPGe detector.	205
C.5	Fitted centroid vs. depth in COSI detectors 1, 2, and 3.	207
C.6	Fitted centroid vs. depth in COSI detectors 4, 5, and 6.	208

C.7	Fitted centroid vs. depth in COSI detectors 7, 8, and 9.	209
C.8	Fitted centroid vs. depth in COSI detectors 10 and 11.	210
E.1	Compton scattering angles of (a) simulated ^{26}Al and atmospheric background events with incident energies 1803–1817 keV, and (b) real flight data (1803–1817 keV; 10 detectors).	215
E.2	“Step down” and “step up” procedures for optimizing the choice of Compton scattering angle.	216
E.3	The percentage of signal events removed as a function of maximum Compton scattering angle.	217
E.4	Simultaneous optimization procedure. The color scale indicates the percentage of events (left: atmospheric background, right: ^{26}Al signal) which have ϕ between the minimum (x -axis) and maximum (y -axis) limits.	219
E.5	Estimated significance as a function of cuts in ϕ defined by the minimum and maximum values indicated on the axes.	219
F.1	Simulated spectra of delayed emission from instrumental activation due to protons, neutrons, and α -particles.	221
F.2	Strength of the source function as a function of energy in three models (Peterson et al. 1972; Ling 1975; Peterson et al. 1973).	223
F.3	Representative exposure map of the 2016 flight, scaled down to one day, generated by dividing the sky into pixels and plotting the percentage of total pointings in each.	225
F.4	The change in azimuthal (x) and zenith (z) orientations during the COSI 2016 flight for one arbitrary pixel of size $\sim 20.3 \text{ deg}^2$	226
F.5	The fraction of Compton events and single-site events present in each of the simulated background data sets.	227
F.6	Individual event selections applied to (a) α -particles and (b) neutrons in the activation simulations.	228
F.7	Individual event selections applied to (a) protons in the activation simulations and (b) atmospheric photons in the Ling model simulation.	229
F.8	Percentage of simulated activation and atmospheric Compton events which pass or fail the event selections listed on the x -axis.	230
F.9	Application of event selections in succession to the 12-detector α -particle simulation and 10-detector atmospheric simulation.	231

List of Tables

2.1	Flux of ^{26}Al in massive star regions as measured by SPI.	23
4.1	The seven radioactive isotopes used to calibrate COSI, each listed with their γ -ray line energies and respective branching ratios.	58
4.2	Single-strip energy resolution (FWHM) in 2020 and 2016 calibration data.	62
4.3	Fully-reconstructed energy resolution (FWHM) in 2020 and 2016 calibration data.	75
4.4	Angular resolution (FWHM) in 2020 and 2016 calibration data.	77
4.5	Effective area in 2020 and 2016 calibration data.	79
5.1	The longitudinal (ℓ) and latitudinal (b) definitions of the signal and background regions in the COSI 2016 ^{26}Al analysis.	86
5.2	Event selections on flight data in the ^{26}Al signal and background regions.	86
5.3	Fit parameters of a power law plus three Gaussian fit to the flight data in the background region with minimal event selections.	93
5.4	Aluminum-26 line parameters from the COSI 2016 flight.	99
5.5	Line rates and uncertainties after the rigidity-weighted background subtraction.	103
5.6	Summary of flight data results from various tests of method validation.	106
5.7	Fit parameters of a power law plus three Gaussian fit to the instrumental activation and atmospheric background simulations in the background region with minimal event selections.	109
5.8	Summary of results from complete flight simulations of various tracer maps.	112

Acknowledgments

Completing this dissertation would not have been possible without the generous support of my colleagues, friends, and family.

Thank you to my advisor, Dr. John Tomsick, for his dependable guidance that has shaped my growth as a scientist and his trust in me as a valued member of the COSI team. Thank you to Dr. Andreas Zoglauer for his constant efforts to inform my analyses with novel tools and explanations. Thank you to Professor Steven Boggs for his enthusiastic collaboration and interest in my work. To Dr. Thomas Siegert, I extend sincere gratitude for years of patient pedagogy and conversation about space, statistics, and so viele süße Apfelkuchen.

Thank you to Dr. Hadar Lazar, my COSI co-captain, whose creativity and camaraderie are a daily source of motivation. I am endlessly grateful that our time together has yielded not only a fruitful working relationship, but also a lifelong friendship. Thank you to Nicholas Pellegrini for his sharp hardware and software skills and positive presence in the lab. Thank you to Brent Mochizuki, Steve McBride, Dr. Carolyn Kierans, Dr. Alexander Lowell, Dr. Jarred Roberts, and Dr. Clio Sleator for their reliable assistance, kind encouragement, and countless campaign stories.

Thank you to everyone who has reviewed this document and helped me synthesize years of research into a readable composition. Thank you to the members of my dissertation committee, Professor Stuart Bale, Professor Daniel Kasen, and Professor Marjorie Shapiro, for their instructive comments and experienced perspective.

My path to this PhD did not begin in Berkeley. Thank you to my professors and undergraduate research group at the University of Michigan. Thank you to Mr. Ralph Mazzio, my high school physics teacher in New York, whose captivating curriculum inspired me to pursue physics and remains among the most illuminating of my career.

Lastly, there are no words which can adequately express my gratitude for my friends and family. Their enduring love and support has selflessly enabled me to chase what has at times felt like an ever-distant goal. As I look towards the future, I am dually heartened by long-awaited reunions on the East Coast and the knowledge that the relationships which enriched my years on the West Coast will never be far from reach.

Chapter 1

Introduction

Gamma-rays (γ -rays) are a window into the highest-energy components of the universe, carrying with them a relatively unobstructed view of objects spanning the most disruptive black holes, pulsars, γ -ray bursts (GRBs), and supernovae, to the slow-burning processes of stellar nucleosynthesis that, over millions of years, enrich the Milky Way Galaxy with the very elements that comprise our day-to-day lives.

Despite the penetrating nature of γ -rays that permits detection of highly energetic astrophysical sources from vast distances, deeper exploration of these objects proves a formidable task. The precise emission mechanism of γ -rays from GRBs, for example, is not understood. Attempts to study the magnetic environment surrounding GRBs and compact objects through measurements of γ -ray polarization remain uncertain and few in number. Even the details of well-measured signals of interest, namely the bright 511 keV signal emanating from the Galactic Center and the 1.809 MeV signal of ^{26}Al from the Galactic Plane, continue to elude observers after decades of research.

It is clear that the mysterious nature of astrophysical γ -rays matches their vast scientific potential. Robust instrumentation and analysis techniques designed specifically to accommodate the difficulties posed by the field are a prerequisite to success. Historical and ongoing development of these tools pave the way to a more transparent γ -ray sky.

This dissertation details a recent probe of the soft γ -ray regime: the Compton Spectrometer and Imager (COSI; Chapter 3). Developed as a balloon-borne instrument, COSI follows a lineage of previous Compton telescopes launched and recovered through the NASA Balloon Program Office. COSI took flight on NASA's superpressure balloon from Wanaka, New Zealand in 2016 and stayed afloat for 46 days. It was a technological and scientific success, boasting detection and imaging of the 511 keV positron-electron annihilation signature (Kierans 2018; Kierans et al. 2020; Siegert et al. 2020), the Crab Nebula (Sleator 2019), and GRB 160530A (Lowell 2017). This document expands this list of accomplishments to include stellar nucleosynthesis via the first measurement of Galactic ^{26}Al in the COSI 2016 balloon flight (Chapter 5). This also marks the first measurement of ^{26}Al with a compact Compton telescope.

These scientific results, however, are attainable only through precise calibrations of the

instrument. The calibration procedures honed by the COSI team leading up to the 2016 flight and executed prior to an attempted 2020 mission are explained thoroughly (Chapter 4). The depth of these explanations is twofold: (1) Knowledge of the instrument’s subsystems and the purpose of each calibration is necessary to understand the capabilities, and equally important, the limitations of the instrument. In particular, this dissertation explores the deleterious detector effects of charge sharing (Appendix B) and charge trapping (Appendix C) to inform future mitigation efforts. (2) The quality of scientific analysis is predicated on the quality of instrument calibrations.

These calibration procedures will be integral to the next iteration of COSI as a NASA Small Explorer satellite mission, slated for launch in 2027. Drawn heavily from the design of the balloon mission, the COSI satellite will improve upon the balloon results with two years of continuous observation time, during which it will survey the entire sky each day, improved effective area, finer angular resolution, and less atmospheric background than the balloon mission by nature of its higher altitude in low-Earth orbit. Expected achievements of the satellite are defined by extensive simulations. An additional component of this dissertation focuses on the adaptation of a Richardson-Lucy deconvolution algorithm to test COSI’s imaging capabilities in the balloon configuration (Chapter 6). The algorithm was deployed in a public data challenge that introduced COSI analysis procedures to the astrophysics community. It will be used in future COSI satellite data challenges and analyses of real satellite data.

Thus, as an instrument uniquely equipped to detect photons of energy 0.2–5 MeV, the COSI balloon has yielded valuable technological and scientific insight and will help to fill in the historically under-explored “MeV” gap of γ -ray astrophysics (0.1–100 MeV) as a satellite mission. The calibration data, associated detector analyses, measurement of Galactic ^{26}Al , and imaging tests comprising this dissertation together lay the groundwork for a promising next chapter in COSI history.

Chapter 2

Aluminum-26 Science

The stars populating our Galaxy differ greatly in age, size, metallicity, color, luminosity, and temperature, among other characteristics. Driving this wide array of features is the origin story of each star that fundamentally governs the journey through its life cycle, its contributions to the surrounding interstellar medium (ISM), and its ultimate fate. This life cycle is of interest to astronomers because stars seed the universe with heavy elements, transforming its primordial abundance of hydrogen and helium into the materials which comprise our universe today. We can study this transformation by observing nuclear de-excitations of radioactive isotopes distributed across the sky. In particular, we can use γ -ray telescopes to detect the MeV γ -rays emitted by radioactive isotopes and place constraints on the models of stellar and explosive nucleosynthesis which create them.

MeV γ -ray telescopes with excellent energy resolution and a wide field of view can establish a timeline of the synthesis and expulsion of isotopes into the ISM. In other words, observations of a variety of isotopes can paint a spatial and temporal picture of where certain radioactive isotopes are synthesized, by which objects, and when. For example, Cobalt-56 (^{56}Co) emits 0.847 MeV and 1.238 MeV γ -rays with a short half-life time of 77 days, enabling short-term follow-up of Type Ia supernovae. The overall brightness and temporal appearance of ^{56}Co lines reveal the initial quantity of ^{56}Ni and the explosion structure (e.g., Siebert and Diehl 2015). Titanium-44 (^{44}Ti) also informs supernova history on a short time scale of $\sim 10^2$ years. It emits 0.068, 0.078, and 1.157 MeV γ -rays with a half-life time of ~ 59 years, tracing young supernova remnants in the Milky Way (Cassiopeia A: Iyudin et al. (1994); Iyudin et al. (1997); Grefenstette et al. (2014). Tycho: Troja et al. (2014)). Lines from ^{44}Ti have also been measured from SN 1987A in the Large Magellanic Cloud (Boggs et al. 2015).

Longer-lived isotopes can complement the short-term measurements with an integrated view of nucleosynthesis in the Galaxy over millions of years. Iron-60 (^{60}Fe) emits 1.173 and 1.332 MeV γ -rays with a half-life time of 2.6 Myr and is thought to be released only in core-collapse supernova (CCSN) events. Conversely, Aluminum-26 (^{26}Al) is thought to be released during the lifetimes of massive stars, in terminal CCSN explosions, and to a lesser extent in asymptotic giant branch (AGB) stars and novae. It decays to an excited state of ^{26}Mg with a half-life time of 0.7 Myr, after which it immediately deexcites to the ground

state of ^{26}Mg and emits a 1808.72 keV γ -ray ([IAEA Nuclear Data Section](#)). The 0.7 Myr half-life ensures that ^{26}Al lives long enough to survive ejection from its progenitor sites and decay as it is distributed throughout the ISM. Thus, observations of ^{26}Al via its 1.809 MeV signature can provide unique insight into stellar nucleosynthesis, broader stellar feedback, and the chemical evolution of the Milky Way Galaxy over millions of years.

This chapter will review the history of ^{26}Al observations in the context of two broad goals: understanding (1) its dominant progenitors and (2) the kinematics of its emission. The subsequent chapters of this dissertation will detail the ways in which MeV γ -ray telescopes, with an emphasis on COSI, are well-suited to studying this isotope.

2.1 History and observations of Galactic Aluminum-26

The emission of ^{26}Al in the Milky Way Galaxy was discovered by the High Energy Astronomy Observatory (HEAO)-3 in the 1980s. HEAO-3 was launched by NASA into low-Earth orbit in 1979 and observed the Galaxy in the 0.05–10 MeV bandpass for two years. Comprised of four coaxial high-purity germanium detectors, HEAO-3 had a $\sim 42^\circ$ field of view, effective area of about 8 cm^2 , and average full-width half-maximum (FWHM) energy resolution of approximately 3.3 keV at 1809 keV (0.18%). In 1984, the HEAO-3 team published the first evidence of ^{26}Al emission at $1808.49 \pm 0.41\text{ keV}$ (FWHM $\leq 3.0\text{ keV}$) along the Galactic Plane with 4.8σ significance (Mahoney et al. 1984). The intensity was reported as $(4.8 \pm 1.0) \times 10^{-4}\text{ ph cm}^{-2}\text{ s}^{-1}\text{ radian}^{-1}$, which corresponds to a Galactic ^{26}Al mass of approximately $3M_\odot$. The analysis was limited to spectroscopy (Figure 2.1); as HEAO-3’s cesium iodide (CsI(Na)) shielding was largely transparent to γ -rays of energy greater than 1 MeV, the angular response of the instrument suffered at 1.8 MeV and was not sensitive enough to obtain a detailed map of the emission.

The Imaging Compton Telescope (COMPTEL) on board NASA’s Compton Gamma-Ray Observatory (CGRO) significantly advanced knowledge of ^{26}Al over the course of its operation from 1991 to 2000. COMPTEL was a “classic” Compton telescope (see Section 3.1.2) with an angular resolution of $\sim 4^\circ$ (FWHM), energy resolution of 8.8% (FWHM) at 1.27 MeV, and a $\sim 1\pi\text{ sr}$ field of view in the energy bandpass 0.7–30 MeV (Schönfelder et al. 1993). COMPTEL produced the first all-sky images of ^{26}Al with maximum entropy (ME) and multi-resolution regularized expectation maximization (MREM) algorithms (Figure 2.2; Oberlack et al. 1996; Oberlack 1997; Plüschke et al. 2001). The ME method tends to emphasize localized emission and lumpy structure at late stages of the algorithm’s progression and can underestimate the flux of the emission at earlier stages. Hence, the ME map in Figure 2.2a is from an intermediate stage. On the contrary, the MREM algorithm seeks the smoothest attainable image and suppresses low-significance features and potential artifacts via a user-controlled threshold. The MREM map is shown in Figure 2.2b. Comparing the ME and MREM images provides a balance between a distribution potentially contami-




Figure 2.1: Energy spectrum of ^{26}Al along the Galactic Plane from HEAO-3. The solid line shows a fit of linear background plus a Gaussian photopeak to the data. Image from Mahoney et al. (1984).

nated by statistical fluctuations and artifacts and a distribution that is overly smoothed and hides real features of structured emission. Notably, the ME and MREM images both reveal concentrated ^{26}Al emission in the Inner Galaxy ($|\ell| \leq 30^\circ$, $|b| \leq 10^\circ$) and localized sites of emission in regions of massive star activity along the Galactic Plane, including Cygnus, Carina, and Vela.

Aluminum-26 observations since the early 2000s are largely characterized by the success of the INTERNATIONAL Gamma-Ray Astrophysics Laboratory (INTEGRAL) satellite, launched by the European Space Agency in 2002. Of particular relevance to this document is the SPectrometer on INTEGRAL (SPI), a coded mask instrument (see Section 3.1.1) which over more than two decades of observation time has contributed immensely to spectroscopic and imaging studies of ^{26}Al . SPI’s high-purity germanium detectors have excellent energy resolution of $\sim 2\text{ keV}$ (FWHM) at 1 MeV (0.2%) and detect photons across an energy bandpass of 20 keV–8 MeV. The field of view is approximately 16° and the angular resolution is about 2.5° . SPI produced an image of ^{26}Al (Figure 2.3a) that largely corroborates the main conclusions of the COMPTEL image: emission is concentrated in the Inner Galaxy with localized emission in regions of massive star activity (Bouchet et al. 2015). The full-sky spectrum of ^{26}Al from SPI (Siegert 2017) is shown in Figure 2.3b. Further spectroscopic achievements of SPI, including models of ^{26}Al motion into low-density supershells and flows of ^{26}Al in specific regions like Orion-Eridanus and Upper Scorpius, are discussed later in this chapter.

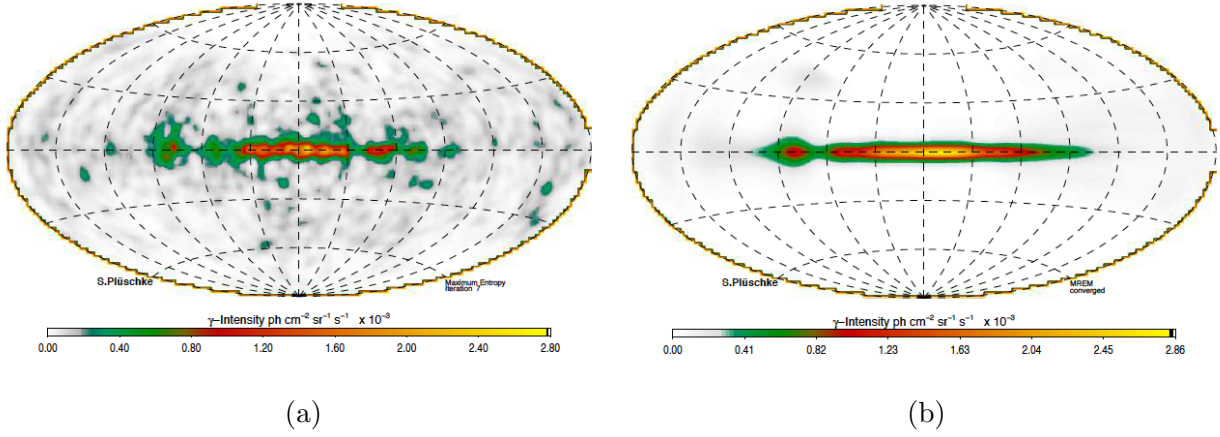


Figure 2.2: COMPTEL 1.8 MeV images of Galactic ^{26}Al . (a) The maximum entropy (ME) image shows concentrated emission in the Inner Galaxy ($|\ell| \leq 30^\circ$, $|b| \leq 10^\circ$) and hints of emission in the Cygnus ($\ell \sim 80^\circ, b \sim 0^\circ$), Carina ($\ell = 286.5^\circ, b = 0.5^\circ$), and Vela ($\ell = 264^\circ, b = -3^\circ$) regions. The resolution is $\sim 4^\circ$. (b) The MREM image provides a smoother view which also indicates concentrated emission in the Inner Galaxy and along the Galactic Plane ($b \sim 0^\circ$). Both images from Plüschke et al. (2001).

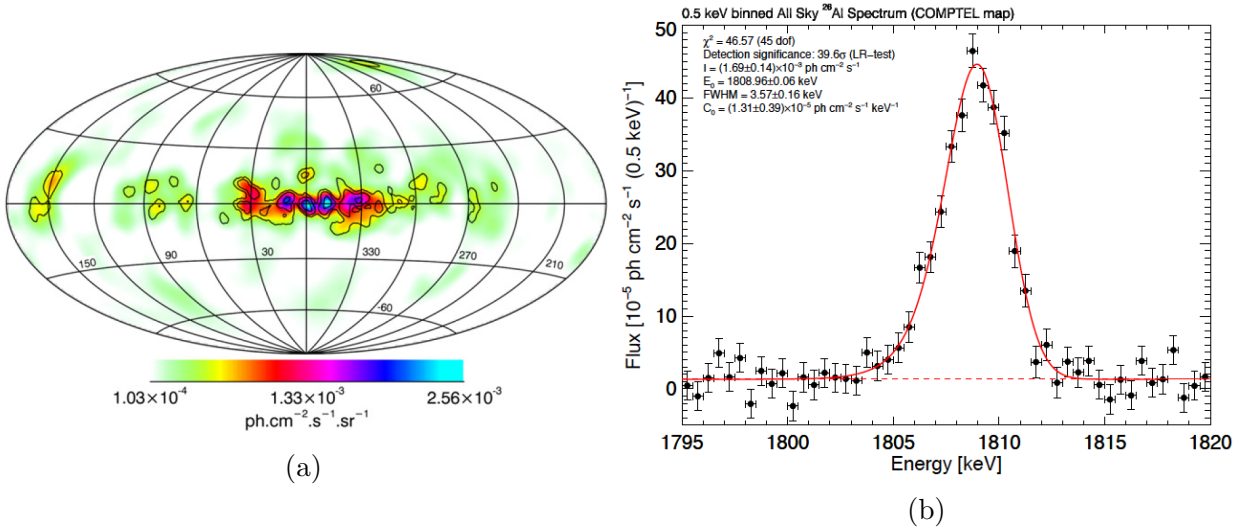


Figure 2.3: SPI 1.8 MeV image and spectrum. (a) The image was produced from 2×10^8 seconds of observation time between 2003 and 2013. Emission is concentrated in the Inner Galaxy and extends to localized regions along the Galactic Plane, including Cygnus, Vela, Carina, and Scorpius-Centaurus ($328^\circ < \ell < 355^\circ, 8^\circ < b < 30^\circ$). Image from Bouchet et al. (2015). (b) Spectrum of ^{26}Al emission from SPI. Image from Siegert (2017).

2.2 Dominant progenitors

Today, Wolf-Rayet (W-R) stars and core-collapse supernovae (CCSNe) are considered the dominant progenitors of ^{26}Al with possible smaller contributions from novae and asymptotic giant branch (AGB) stars. A brief summary of these objects is provided below as background to the following discussion about observational constraints:

Wolf-Rayet stars

Wolf-Rayet (W-R) stars, a subset of O-type stars, are promising candidates for ^{26}Al synthesis given their large masses ($\approx 20\text{--}40 M_{\odot}$, with the most massive stars $\geq 60 M_{\odot}$), high temperatures ($\gtrsim 10^4 \text{ K}$), and high metallicity required for synthesis of heavy elements (Abbott and Conti 1987). Fast rotation ($\sim 190 \text{ km s}^{-1}$; Huang and Gies 2006) and large thermal gradients inside of the stars lead to enhanced mass loss rates ($10^{-5}\text{--}10^{-3} M_{\odot} \text{ yr}^{-1}$; Abbott 1982; Vink and Koter 2005; Dale 2015), convection, and substantial mixing of stellar material which generate strong winds that dredge up ^{26}Al from the stellar interior and eject it into the ISM during the star’s lifetime. Initial wind velocities are $\sim 10^3 \text{ km s}^{-1}$ (Howarth and Prinja 1989) and terminal velocities are $\sim 10^2 \text{ km s}^{-1}$ (Dale 2015). W-R stars also liberate large quantities of photons that ionize atoms in the ISM.

Core-collapse supernovae

Core-collapse supernovae (CCSNe) are powerful explosions that mark the end of a massive star’s life. When nuclear fusion in the core of a massive star ceases after iron production, the core cannot combat the force of its self-gravity and it collapses in on itself (Hoyle and Fowler 1960; Janka 2012). The rebound of supersonic in-falling material at the core generates shocks that propel products of nucleosynthesis into the ISM. Assuming that a CCSN progenitor star of mass $M \sim 20 M_{\odot}$ has a lifetime $\tau \sim 9 \text{ Myr}$ (Schaller et al. 1992), CCSN progenitors of mass $M > 20 M_{\odot}$ will have $\tau < 9 \text{ Myr}$; those with $M < 20 M_{\odot}$ will have $\tau > 9 \text{ Myr}$.

Novae

Novae are explosive events precipitated by the sudden ignition of hydrogen fusion on the surface of a white dwarf in a binary system. The ignition occurs when the distance between the white dwarf and its stellar companion is small enough for accreted matter, mostly hydrogen, to fall onto the surface of the white dwarf (Gallagher and Starrfield 1978). The white dwarf heats this new atmosphere of hydrogen until it thermally ignites, sparking nuclear fusion and synthesis of intermediate-mass elements. A review of classical novae is provided in Chomiuk et al. (2021).

Asymptotic giant branch stars

Low- to intermediate-mass stars ($\sim 0.5\text{--}8 M_{\odot}$) undergo an asymptotic giant branch (AGB) phase after helium burning. The helium shell in AGB stars can suddenly ignite in “helium shell flashes,” which create strong convective currents and “thermal pulses” that dredge up products of nucleosynthesis from the stellar interior. Extensive mass loss in the form of winds can disperse this material into the ISM (Höfner and Olofsson 2018).

2.2.1 Multi-wavelength analyses

As sites of either stellar or explosive nucleosynthesis, all of these objects are candidate progenitors of ^{26}Al . Historically, however, the estimated contributions of these sources to the total observed yield of ^{26}Al have varied. Studies of Galactic ^{26}Al from the HEAO-3 mission concluded that CCSNe are not major sources of ^{26}Al in the Galaxy. The observed mass fraction of $^{26}\text{Al}/^{27}\text{Al}$ in the ISM was $\sim 10^{-5}$, which though comparable to that obtained in studies of meteoritic material, was greater than an order of magnitude beyond predictions from supernova models at the time. This implied that supernovae could not be the dominant progenitors of Galactic ^{26}Al . Rather, a likely origin from millions of novae and thousands of supernovae was proposed (Mahoney et al. 1984).

The subsequent COMPTEL image opened the door to several analyses which instead favored W-R stars and CCSNe as the dominant sources of ^{26}Al and largely excluded novae from consideration. First, the image clearly shows ^{26}Al emission along the Galactic Plane, implying an association of ^{26}Al with the star-dense spiral arms of the Galaxy. To safeguard against over-interpreting specific features of the ME and MREM images, both of which are subject to imaging artifacts, extensive comparisons of the COMPTEL 1.8 MeV image to 31 others were performed to correlate the emission with well-constrained astrophysical tracers (Knödlseeder et al. 1999). This multi-wavelength approach, spanning wavelengths of ~ 10 m to energies greater than 100 MeV, focused the list of candidate ^{26}Al sources using complementary physical interpretations of the observed morphology. The chosen astrophysical maps for comparison were pre-treated to remove contamination from known instrumental background and extragalactic sources. They were also convolved with the COMPTEL point spread function to establish a common angular resolution of 3.8° (FWHM). A maximum likelihood approach then determined the best fit of these “tracer” maps (the source model) and instrumental background (the background model) to the COMPTEL 1.8 MeV map. The detection significance of each tested tracer map is shown in Figure 2.4.

The analysis suggests that the best tracers of ^{26}Al emission are those of young stellar populations and those concentrated at low Galactic latitudes, i.e. near the Galactic Plane. Maps that trace old stellar populations do not correlate strongly with ^{26}Al emission. This association with young stellar populations thus favors ^{26}Al from massive stars that live on short timescales. Considering both the maximum likelihood results and residuals between the longitudinal profiles of the 1.8 MeV and tracer maps, the study finds that the 53 GHz free-free map tracing ionized gas most strongly correlates with ^{26}Al emission. It is the only map which shows similar intensity in Cygnus and describes the 1.8 MeV emission at negative longitudes. The Diffuse Infrared Background Experiment (DIRBE) $240\ \mu\text{m}$ map, which traces thermal dust at approximately 12 K, was the only other map that exhibited small residuals similar to those of the 53 GHz free-free map, albeit with greater residuals near Carina and Vela.

Overall, the correlation of COMPTEL’s 1.8 MeV map with star-forming regions, spiral arms, and young stellar populations points to active synthesis of ^{26}Al in massive stars. Additionally, the association of ^{26}Al with the 53 GHz free-free map suggests that the sources

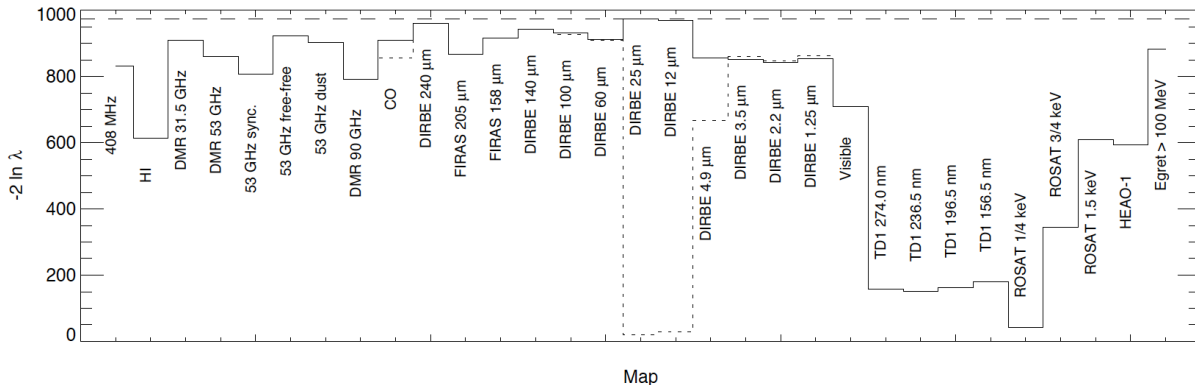


Figure 2.4: A multi-wavelength comparison of COMPTEL’s 1.8 MeV image to other all-sky maps spanning the electromagnetic spectrum. The detection significance is defined as $\sqrt{-2\ln\lambda}$ in units of σ for maximum likelihood ratio λ . Maps that trace young stellar populations at low Galactic latitudes follow the 1.8 MeV emission most closely. Image from Knödlseeder et al. (1999).

of ^{26}Al must sustain a population of ionizing stars. These are stars of mass $M \geq 20 M_{\odot}$ and short lifetimes of $\sim 3\text{--}9$ Myr (Schaller et al. 1992) that can ionize atoms in the ISM. Free electrons interact with these ions to produce free-free radiation. Correlation with the ionized ISM strengthened COMPTEL’s argument in favor of W-R stars as dominant sources of ^{26}Al . Not only are W-R stars known sources of ionizing photons, but they live short enough lives to coexist with the population of ionizing stars in the Milky Way. CCSNe also have short enough lifetimes to feasibly sustain ionization. Novae, on the other hand, as tracers of old stellar populations, live too long to do so and are consequently excluded as significant sources of ^{26}Al . The same timescale argument disfavors AGB stars, which typically have masses of $< 9 - 10 M_{\odot}$, as dominant sources (Knödlseeder 1999). Note that while AGB stars may not be dominant contributors, theoretical models and meteoritic abundances of ^{26}Al do support some AGB origin (Mowlavi and Meynet 2000).

Estimating the mass of Galactic ^{26}Al from the continuum luminosity of ionizing photons yields $3.1 \pm 0.9 M_{\odot}$ (Knödlseeder 1999). Theoretical nucleosynthesis predictions indicate that contributions from both W-R stars and CCSNe are required to account for this mass, though uncertainties in stellar nucleosynthesis models (especially with respect to stellar metallicity and rotation) make it difficult to constrain W-R yields. Yields from CCSNe are also uncertain because not all stars which end as CCSNe are sites of substantial ^{26}Al emission (e.g., CCSN progenitors of mass $< 8 - 10 M_{\odot}$).

Analysis of the SPI 1.8 MeV image also suggests a massive star and CCSN origin of ^{26}Al . The image itself exhibits concentrated emission in regions of massive star activity. Testing template maps with a model-fitting approach corroborated the COMPTEL multi-wavelength analysis: maps tracing dusty regions of young stellar environments and free-free emission are

highly correlated with 1.8 MeV emission (Bouchet et al. 2015). CCSN ejecta also generate substantial dust and gather large dust grains, supporting the likelihood of their contributions to the observed abundance of ^{26}Al .

2.2.2 Theoretical estimates

Theoretical estimates of ^{26}Al production are derived from simulations of Galactic chemical evolution and hydrodynamical models. Kaur and Sahijpal (2019) model the widespread distribution of ^{26}Al (and other short-lived radioisotopes (SLRs) including ^{60}Fe , ^{36}Cl , ^{41}Ca , and ^{53}Mn) with simulations of Galaxy-wide evolution. They model the enrichment of finer regions ($0.1\text{--}1\text{ kpc}^2$) by simulating the evolution of sequential stellar generations. The models suggest that W-R winds from a massive star in a neighboring stellar cluster ($\sim 20\text{ pc}$ away) likely seeded the presolar molecular cloud of our early Solar System with ^{26}Al . Note that ^{26}Al and the other isotopes listed above are called “short-lived” in this context because these studies often concern behavior during the time of the early Solar System on $\sim\text{Myr}$ timescales, compared to the present $\sim\text{Gyr}$ age of the Solar System.

It is also possible that a CCSN was the responsible actor behind this injection of ^{26}Al into the solar neighborhood, as explored in SN-triggered solar system formation (Cameron and Truran 1977) and 3-D hydrodynamical models of molecular cloud cores (Boss and Keiser 2014; Boss 2017). Modeling the yields of ^{26}Al from such a supernova, however, is complicated by sensitivities to progenitor mass, initial metallicity, and the production and destruction channels of ^{26}Al which proceed both before and during the explosion. Lawson et al. (2022) find that it is difficult to assess CCSN properties from ^{26}Al because the ^{26}Al produced during the lifetime of the star before the supernova, which comprises the majority of the total yield, lives long enough to survive the CCSN explosion and mix with the ^{26}Al produced during the explosion itself.

Divergent from the hypothesis of external injection by a strong CCSN event or W-R wind is an alternative method of enrichment that argues that the early Solar System exhibits a notable abundance of heavy elements simply because star-forming regions are correlated (Fujimoto et al. 2018). Molecular clouds preferentially accrete gas from nearby regions that contain products of nucleosynthesis, making it likely that sites of star formation are contaminated with SLRs like ^{26}Al . Thus, it is a physical association with enriched regions and stellar feedback, rather than a chance encounter with a nearby event, that may be responsible for SLR abundances in young stellar systems.

Studying the surrounding environment of prestellar cores in the star-forming complex Ophiuchus may also help classify the dominant progenitors of ^{26}Al . Specifically, its proximity to the Upper Scorpius region, in which COMPTTEL measured ^{26}Al with 2σ significance (Plüschke et al. 2001), can probe the production and incorporation of ^{26}Al into young regions of stellar formation. A forward model of all massive stars in Upper Scorpius is developed in Forbes et al. (2021) to compute the probability that the observed abundance of ^{26}Al today originates from either W-R stars or supernovae. The stars evolve from an initial mass and birth time according to parameters of stellar evolution defined in the literature: priors

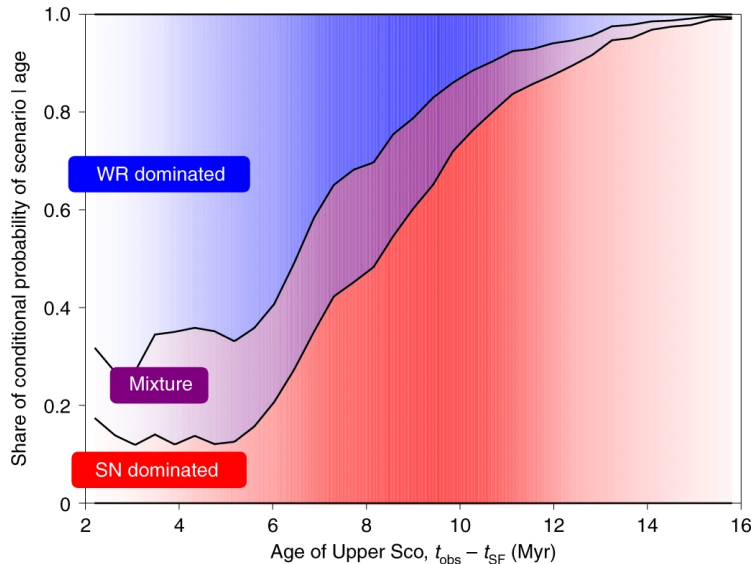


Figure 2.5: Probability from a population synthesis model that the present-day ^{26}Al in Upper Scorpius originated from Wolf-Rayet (WR) stars, supernovae (SN), or a mixture of the two, conditional on the uncertain age of Upper Scorpius. Color opacity scales with probability of age, such that Upper Scorpius is most likely ~ 10 Myr old. Image from Forbes et al. (2021).

on the initial mass function, the mass and age of Upper Scorpius, and various supernova explosion scenarios constrain the simulation. The quantity of ^{26}Al generated by each star is recorded with its ultimate fate (i.e. whether a star has exploded in a supernova or not, given by a model with priors on stellar mass and three different explosion mechanisms). The results are shown in Figure 2.5. The model indicates that the majority of living ^{26}Al in Upper Scorpius is most likely of supernova origin. However, the relative contributions of each source are dependent on the uncertain age of Upper Scorpius. Observations suggest that Upper Scorpius is likely ~ 10 Myr old, but if it is younger than this estimate, the model instead favors a W-R dominated scenario. Considering all possible ages of Upper Scorpius together favors supernovae as the dominant progenitor, with 90% of the observed ^{26}Al produced by supernovae in 59% (W-R in 27%) of realizations.

The above considerations support massive stars as the favored class of dominant ^{26}Al progenitors (see also Prantzos and Diehl 1996). Theoretical models of novae predict only a $0.1\text{--}0.4 M_{\odot}$ contribution to the observed mass of Galactic ^{26}Al (José et al. 1997). More recent hydrodynamic simulations of novae, though, using new precision measurements of the $^{25}\text{Al}(p, \gamma)^{26}\text{Si}$ reaction, indicate an increase in novae contribution from 20% to 30% of the Galactic budget (Bennett et al. 2013). Recent work to estimate the contribution from very massive stars with masses of $150\text{--}300 M_{\odot}$, exceeding that of the previously considered massive stars ($\sim 25\text{--}120 M_{\odot}$), suggests a 120–150% increase in the Galaxy’s ^{26}Al production (Martinet et al. 2022). Continued modeling of ^{26}Al production in stellar and explosive nucleosynthesis,

including in W-R stars, CCSN, novae, and AGB stars, remains a priority.

2.2.3 Complementary observations

Constraints on the dominant progenitors of ^{26}Al can be derived indirectly from complementary observations of other sources. Comparing γ -ray signatures of ^{26}Al to those of ^{60}Fe can inform models of stellar and explosive nucleosynthesis. Cross-correlations between the positron-electron annihilation morphology and that of ^{26}Al may constrain the mysterious origin of positrons in our Galaxy and their potential stellar origins. Geologic signatures from deep-sea sediments on Earth and meteoritic material are also used to understand how ^{26}Al may have been produced and incorporated into our Solar System near the time of its formation (see also theoretical discussion in Section 2.2.2 above). Each of these pursuits is discussed below.

2.2.3.1 Iron-60

Comparing maps of ^{60}Fe and ^{26}Al is an important goal of nucleosynthesis studies. While ^{26}Al is produced and released during the lifetime of massive stars and in CCSNe, ^{60}Fe is thought to be associated mostly with supernova events. Iron-60 is produced deep in the stellar interior during the helium- and carbon-shell burning stages of stellar evolution via slow neutron capture (s-process). As such, even if the two isotopes are co-produced in massive stars, the strong winds which dredge up ^{26}Al as it is synthesized in the hydrogen-burning phase cannot liberate ^{60}Fe , which is predominately ejected after the explosive nucleosynthesis in the supernova. This contrast is key to establishing a timeline of nucleosynthesis in the Galaxy. Because ^{60}Fe and ^{26}Al are thought to share massive stars as their main progenitors and ^{60}Fe has a longer half-life time than ^{26}Al (2.6 Myr vs. 0.7 Myr), comparing the distributions of ^{60}Fe and ^{26}Al will reveal the different evolutionary stages of regions across the sky. Regions with young massive stars and few to no CCSNe should be seen only in the ^{26}Al map, while regions with older massive star populations should have more CCSNe events and appear prominently in an image of ^{60}Fe emission.

However, there is no existing map of ^{60}Fe emission to date. Over 15 years of SPI data are insufficient to constrain the spatial morphology of ^{60}Fe and have only yielded a Galaxy-wide average of the flux ratio $^{60}\text{Fe}/^{26}\text{Al}$. The spectrum in Figure 2.6 shows spectral intensities of the isotopes' γ -ray signatures assuming emission over the best-fit spatial extent of the ^{26}Al distribution (an exponential disk of scale radius 7 kpc and scale height 0.8 kpc). The ^{60}Fe line flux is $(2.6 \pm 0.6) \times 10^{-4}$ ph cm $^{-2}$ s $^{-1}$ and the ^{26}Al line flux is $(14.4 \pm 0.7) \times 10^{-4}$ ph cm $^{-2}$ s $^{-1}$, giving $^{60}\text{Fe}/^{26}\text{Al} \sim (18.3 \pm 4.4)\%$ (Wang et al. 2020). Theoretical calculations predict a ratio of about 0.16 (Timmes et al. 1995), assuming again that ^{60}Fe and ^{26}Al originate from the same sources and that their emission traces similar spatial distributions. By contrast, a ratio of 0.29 or 0.86 is predicted from Galaxy-wide population synthesis simulations in Siegert et al. (2023), depending on the chosen models of stellar evolution and explodability; it is important to note that this simulation appears to underestimate overall yields of ^{26}Al .




Figure 2.6: Spectral intensities of the ^{60}Fe and ^{26}Al γ -ray signatures from SPI. The ^{60}Fe lines at 1.173 and 1.332 MeV are noticeably weaker than the 1.809 MeV line from ^{26}Al . Image from Wang et al. (2020).

Until a more detailed map of ^{60}Fe is produced, the validity of the assumptions in these studies will remain poorly constrained. The flux ratio, for example, is expected to vary between individual massive star regions (Scorpius-Centaurus and Cygnus; Voss et al. 2009) and ^{60}Fe signals to date have been too weak to resolve these localized sites.

It is challenging to obtain a strong ^{60}Fe measurement because of the weaker intensity of its lines and their spectral overlap with decays from instrumental activation. The longer lifetime of ^{60}Fe means that the isotope can propagate greater distances and spread throughout the ISM before decay, making its emission appear more diffuse. This is problematic for coded mask instruments like SPI which have limited sensitivity to weak flux gradients (Section 3.1.1). SPI’s sensitivity to ^{60}Fe was also diminished by the decay of activated ^{60}Co in the instrument, which emits γ -rays at the same energies as ^{60}Fe (the ^{60}Fe decay chain emits its 1.173 and 1.332 MeV signatures through an initial decay to ^{60}Co). A strong instrumental activation line from germanium at 1.337 MeV blends with the 1.332 MeV as well, further obscuring the signal (Wang et al. 2020).

Overall, sensitive observations of ^{60}Fe across the Galaxy and in specific regions of stellar activity are required to better compare its emission to that of ^{26}Al . Present analyses and theoretical predictions often predicate their constraints on an assumed similarity between the isotopes, yet the degree of this similarity remains largely unknown. Comparing detailed maps of ^{60}Fe and ^{26}Al could spatially associate their respective progenitors to specific classes of objects, like massive stars and supernovae, and inform the nucleosynthesis parameters of stellar evolution (e.g., stellar rotation, metallicity, neutron capture rates, main-sequence progenitor mass, final fate, binarity) required to produce the observed emission. Flux mea-

measurements of the two radioisotopes would additionally benchmark the yields predicted by theoretical models of nucleosynthesis.

2.2.3.2 Positrons

Radioactivity likely accounts for a sizable fraction of the Galaxy’s positron budget. Strong winds in massive stars ($> 25 M_{\odot}$) seed the ISM with heavy nuclei over the course of their lifetimes. Massive stars which terminate as CCSN can synthesize additional isotopes in the explosion itself. If these products of nucleosynthesis are β^+ unstable, they necessarily inject positrons into the ISM.

As one of the most abundant β^+ unstable isotopes in the Galaxy, ^{26}Al is important to understanding positron annihilation. The weaker 511 keV disk emission exhibits an annihilation rate of $10^{42}\text{--}10^{43} \text{ e}^+ \text{ s}^{-1}$, depending on the model used to describe its morphology (Knödlseeder et al. 2005; Prantzos 2006). Positrons from a combination of nucleosynthesis products such as ^{26}Al and ^{44}Ti may only fuel 10% of the observed rate in the stronger bulge emission ($\sim 10^{43} \text{ e}^+ \text{ s}^{-1}$); nucleosynthesis from ^{26}Al alone may account for $10.8 \pm 5.2\%$ of the annihilation rate in the disk and $6.8 \pm 2.1\%$ Galaxy-wide (Siegert 2017). Lingenfelter et al. (1993) expect that ^{26}Al can contribute $\sim 16 \pm 5\%$ of the observed annihilation emission. Ultimately, cross-correlations between maps of ^{26}Al and positron annihilation could constrain the uncertain origin of positrons which produce the annihilation signature at 511 keV. These correlations could in turn inform models of ^{26}Al production with expectations of positron yields from regions also associated with ^{26}Al .

Yet, to date there is no established correlation between the ^{26}Al and positron annihilation images. This could indicate that positrons propagate away from their progenitors before annihilating, divorcing the observed 511 keV emission from the sites of radioactivity which may have emitted the responsible positrons. Spectral analyses also support the significant slowdown of positrons as they propagate before annihilating (Leventhal et al. 1978; Ferrière 1998; Prantzos 2006; Beacom and Yüksel 2006). Nevertheless, further study with improved angular resolution may show in more detail where ^{26}Al is being produced, where it flows, and therefore where its positrons may be emitted. Performing new cross-correlation studies with enhanced maps could yield renewed insight. It is also possible that with more comprehensive models of positron behavior in the ISM, say with respect to magnetic fields along which positrons travel, a correlation between liberated positrons streaming through the Galaxy and their original stellar production sites could be achieved. Connecting observed positron yields to ^{26}Al abundances, in combination with localized observations of the nature of these sources, could then tie ^{26}Al production to specific classes of stellar objects.

2.2.3.3 Geologic searches

Measured abundances of stellar and explosive nucleosynthesis products on Earth can also probe the nature of their progenitors. As mentioned previously in discussion about theoretical models of early Solar System enrichment (Section 2.2.2), supernovae (or W-R) ejecta

are thought to have crossed paths with the early Solar System and deposited material into the neighborhood of the forming Earth (Ellis et al. 1996). Traces of these decaying isotopes are found in terrestrial materials. Deep-sea sediments from the Pacific Ocean show a lower limit of the local $^{60}\text{Fe}/^{26}\text{Al}$ flux ratio of $0.18_{-0.08}^{+0.15}$, consistent with estimates from SPI (Section 2.2.3.1), and indicate that a supernova between 1.7 and 3.2 Myr ago likely produced the terrestrial ^{60}Fe (Feige et al. 2018). Similar studies of ^{60}Fe and ^{53}Mn in ocean crusts point to a likely supernova explosion approximately 2.5 Myr ago from a progenitor of mass 11–25 M_{\odot} and solar metallicity (Korschinek et al. 2020). Measurements of ^{60}Fe in a sample of ferromanganese crust from the South Pacific also suggest two supernova events 0–2.8 Myr and 3.7–5.9 Myr ago which deposited ^{60}Fe on Earth (Knie et al. 1999). The measurement of ^{244}Pu in deep-sea manganese crusts that lack anthropogenic ^{239}Pu is consistent with the supernova origins suggested by ^{60}Fe data (Wallner et al. 2004). Furthermore, observations of ^{26}Al , ^{10}Be , ^{41}Ca , ^{59}Ni , and ^{60}Fe in Antarctic ice cores are presented as potential tracers of the Geminid supernova explosion (~ 300 kyr ago) in Ellis et al. (1996). Understanding the influence of nearby supernovae on geologic measurements of radioactivity therefore has clear relevance to astrophysical constraints on the dominant progenitors of these isotopes.

Note that atmospheric production of ^{26}Al on Earth complicates the assumption of a supernova origin. Cosmic ray spallation on argon in the atmosphere is the dominant source of an overall atmospheric ^{26}Al flux of ~ 1280 ^{26}Al atoms $\text{cm}^{-2} \text{yr}^{-1}$, of which only 5% is expected to come from extraterrestrial sources (Auer et al. 2009). This background rate must be accounted for in searches for a supernova-associated signal. Uncertainties in supernova distances, dynamics of the explosions, isotopic yields, and transport mechanisms of elements across the ISM complicate modeling of nearby supernovae as well.

Constraints on stellar and explosive nucleosynthesis can also be derived from radionuclide abundances in meteorites. Meteorites act as undisturbed time capsules of materials present in young circumstellar environments, thereby encoding products of nucleosynthesis from billions of years ago. Excess ^{26}Mg (from decayed ^{26}Al) in silicon carbide (SiC) grains of the Murchison CM2 chondritic meteorite may be compatible with W-R and supernovae origin of ^{26}Al , but higher ratios of ^{26}Mg in other SiC and graphite grains appear instead to favor an AGB origin (Zinner et al. 1991). Isotopic analyses of the Chervony Kut meteorite strongly suggest that live ^{60}Fe must have been present in the early Solar System as it formed (Shukolyukov and Lugmair 1992; Shukolyukov and Lugmair 1993), inviting the question of which source or sources were required to supply the abundance at that time. Evidence of a ^{26}Mg excess in an H4 chondrite (Zinner and Gopel 1992) also suggests that ^{26}Al was present in the early Solar System. Correlating these abundances of ^{26}Al and ^{60}Fe could help determine the likely progenitors, perhaps shared, of each. Observations of calcium-aluminum-rich inclusions in meteoritic material support a similar “pre-enrichment” scenario, whereby stars in our early Solar System were formed in the presence of and incorporated ^{26}Al synthesized by several earlier nearby supernovae. Detailed observations of Upper Scorpius and Ophiuchus, which is adopted as a model of our early Solar System in Forbes et al. (2021), can test this scenario of incorporation by revealing flows of ^{26}Al in the proximity of forming stars.

Therefore, measurements of radioisotopes in terrestrial materials comprise an alternative

tracer of historical nucleosynthesis activity in the Galaxy. Refined astrophysical γ -ray measurements of ^{26}Al , with new measurements of Galactic ^{60}Fe , are important counterparts to these studies.

2.2.4 Summary

The relative contributions of W-R stars and CCSNe to the observed abundance of ^{26}Al in the Milky Way Galaxy are historically uncertain. While these sources are thought to produce most of the observed yield, novae and AGB stars remain possible, if sub-dominant contributors. Advanced imaging with fine angular resolution can help to resolve the debate outlined in this section. Observing localized regions can pinpoint ^{26}Al emission to particular stars, stellar associations, or supernova ejecta. Conducting cross-correlation studies with astrophysical tracers at other wavelengths can also constrain likely sources of ^{26}Al . Gamma-ray observations of ^{60}Fe can differentiate between older sites of explosive nucleosynthesis and younger regions of the Galaxy characterized by stellar nucleosynthesis, i.e. more ^{26}Al emission. Comparing the ^{26}Al image to the 511 keV image may also shed light on the nature of candidate ^{26}Al sources through models of positron production in each. Geologic signatures of ^{26}Al in terrestrial materials and in CAIs provide yet another window into the possible supernova origin of ^{26}Al . Section 2.3 will explore how kinematic analysis with high-resolution spectroscopy complements imaging of localized regions to jointly develop a complete understanding of ^{26}Al emission.

2.3 Kinematics

High-resolution spectroscopy is fundamental to studies of nuclear lines because it reveals the underlying motion of radioisotopes in the Galaxy. Doppler shifts from the rest energy of a γ -ray line indicate bulk motion of the material. Broadening of the line width beyond instrumental resolution can point to the temperature, expansion velocity, and opacity effects influencing the emission. In the case of ^{44}Ti , for example, measurements of shifting and broadening together suggest asymmetries in CCSNe ejecta that can inform models of supernova explosion mechanisms (Grefenstette et al. 2014; Boggs et al. 2015). With respect to ^{26}Al , spectroscopic studies reveal the dynamics of its emission on Galaxy-wide scales, within localized regions of massive star activity, and understanding of its incorporation into the ISM. Recall that the 0.7 Myr half-life of ^{26}Al is what grants us this utility: ^{26}Al continues to decay after it is ejected from its progenitor sites and leaves behind the 1.809 MeV γ -ray signature as it moves through the ISM.

Interpretations of Galactic ^{26}Al motion have evolved with time. HEAO-3 found a line energy consistent with the rest energy of ^{26}Al and a line width consistent with instrumental resolution (≤ 3 keV, 1σ upper limit). This suggests that the observed ^{26}Al slows down over $\sim 10^5$ years as it is incorporated into the ISM, thereby tracing broader Galactic rotation or the random motion of the ISM of ≤ 250 km s $^{-1}$ (Mahoney et al. 1984). By contrast,

the line width measured by the balloon-borne Gamma-Ray Imaging Spectrometer (GRIS) was approximately three times broader than the Doppler broadening expected from Galactic rotation (Naya et al. 1996). This fast ($> 450 \text{ km s}^{-1}$) motion challenges HEAO-3's finding that ^{26}Al slows down before decay and instead suggests that ^{26}Al is being injected into the ISM at high speeds. W-R stars and CCSNe could accelerate ^{26}Al to these speeds in strong stellar winds, explosive ejecta, or through the creation of low-density, merging superbubbles into which ^{26}Al flows; the high speeds of the GRIS measurement can be interpreted as evidence that the majority of observed ^{26}Al in the Galaxy is produced by W-R stars and CCSNe. Measurements by SPI, discussed in the next section, were foundational in the development of the superbubble scenario.

2.3.1 Injection along Galactic rotation

High-resolution measurements by SPI also indicate fast motion, though not as fast as the speeds measured by GRIS. Figure 2.7a shows the evident Doppler shift in ^{26}Al spectra as a function of Galactic longitude (Kretschmer et al. 2013). The regions of interest were chosen as $12^\circ \times 10^\circ$ (Galactic longitude \times Galactic latitude) to optimize statistical and systematic energy uncertainties. For comparison, the nearly Gaussian instrumental line at 1764.49 keV from ^{214}Bi in SPI's anticoincidence shielding does not exhibit the same shift with Galactic longitude. Converting the energy shifts to velocities yields the velocity vs. longitude plot in Figure 2.7b, where positive velocities indicate red-shifted emission and negative indicate blue-shifted. The measured velocities (crosses with error bars) exceed those expected from Galactic rotation, especially at longitudes $\pm 30^\circ$, by $\sim 200 \text{ km s}^{-1}$ (Kretschmer et al. 2013). The blue shading shows measurements from smaller, overlapping regions of interest (3° longitude spacing). The solid blue line follows the expected distribution of molecular CO gas tracing broader Galactic rotation. Clearly, the hot gas carrying ^{26}Al does not follow this model. The dotted red line assumes that one-third of the observed ^{26}Al is distributed along the inner bar of the Milky Way and the remainder throughout the disk and spiral arms. The ^{26}Al velocities agree with the model along the bar but depart from expectations outside of it, suggesting that ^{26}Al is likely influenced by motion in the inner region of the Galaxy. Note that this is consistent with the 1.8 MeV images (Figures 2.2 and 2.3) that show concentrated ^{26}Al in the Inner Galaxy. The dashed green line, which adds a $\sim 200 \text{ km s}^{-1}$ boost in the direction of motion to spiral arm sources tracing broader Galactic rotation, fits the measurements most closely. An injection of ^{26}Al at high velocities and association with spiral arms is thus required to explain the data.

Additional studies by SPI refine this picture of fast-moving ^{26}Al with a scenario by which ^{26}Al streams at $\sim 300 \text{ km s}^{-1}$ into hydrogen (HI) superbubbles blown out from the spiral arms by stellar winds and supernovae ejecta (Krause et al. 2015). A population synthesis model of massive stars and a 3-D hydrodynamic model of superbubble expansion show that HI superbubbles in the Milky Way merge frequently. The acts of merging and ejection from progenitor sites together accelerate ^{26}Al to high velocities as it preferentially streams into the low-density superbubbles. Figure 2.8 illustrates this model of ^{26}Al kinematics.

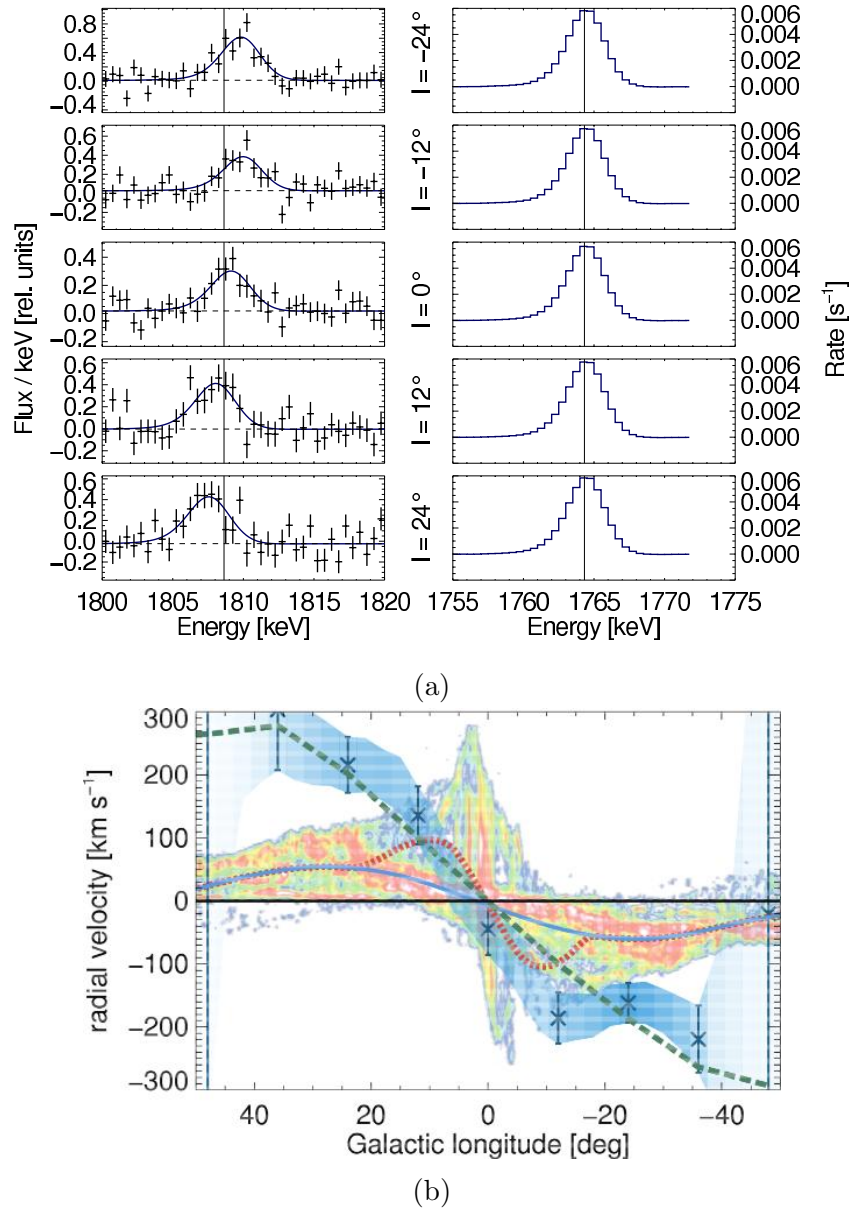


Figure 2.7: Longitude-resolved measurements of Galactic ^{26}Al by SPI. (a) The left panel shows the Doppler shift of ^{26}Al as a function of Galactic longitude. For comparison, the right panel shows that the instrumental line at 1764.49 keV from ^{214}Bi in SPI’s anticoincidence shielding does not exhibit the same shift with Galactic longitude. (b) Energy shifts are converted to velocities as a function of Galactic longitude (crosses with error bars). Refer to Section 2.3.1 for an explanation of the color shading. Both images from Kretschmer et al. (2013).




Figure 2.8: SPI model of how ^{26}Al flows into superbubbles from spiral arms. A spiral arm (solid line) rotates counterclockwise. Large superbubbles were blown out of the disk when the arm was at the dashed location. ^{26}Al produced in the young star clusters streams into the low-density regions of old superbubbles. Image from Krause et al. (2015).

2.3.2 Motion within localized regions

Other regions of the Galaxy can also test the validity of this kinematic model and determine how ^{26}Al moves near massive star activity. For example, observations of different regions within the Orion-Eridanus superbubble could discern internal “sloshing” or flows of ^{26}Al that is distributed non-uniformly within the bubble (Krause et al. 2014). Similar observations of the Scorpius-Centaurus region could reveal the interaction of ^{26}Al with supershells around Upper Scorpius and Ophiuchus (Krause et al. 2018). As shown in Figure 2.9, Upper Scorpius is approximately 20° in diameter and is surrounded by $\sim 10^{-4} M_\odot$ of ^{26}Al . Identifying flows of ^{26}Al within this $\sim 20^\circ$ region may show motion towards Ophiuchus, providing a direct observational counterpart to the dynamical picture of ^{26}Al presented in Forbes et al. (2021) (Section 2.2.2). As a supplementary benefit, new observational estimates of the age of Upper Scorpius would also constrain model predictions in Forbes et al. (2021) of a W-R or supernova origin of ^{26}Al (Figure 2.5).

Finally, future observations could enable studies of extragalactic motion of ^{26}Al in the Large Magellanic Cloud (LMC). The LMC is a particularly interesting target of kinematic studies because its face-on orientation is expected to exhibit a different velocity profile than that of the Milky Way, which is necessarily viewed edge-on from Earth. Models of superbubbles in the LMC indicate velocity flows of $\sim 2000 \text{ km s}^{-1}$ towards and away from the face of the galaxy (Figure 2.10). The contrast between these $\sim 2000 \text{ km s}^{-1}$ flows and expected ^{26}Al

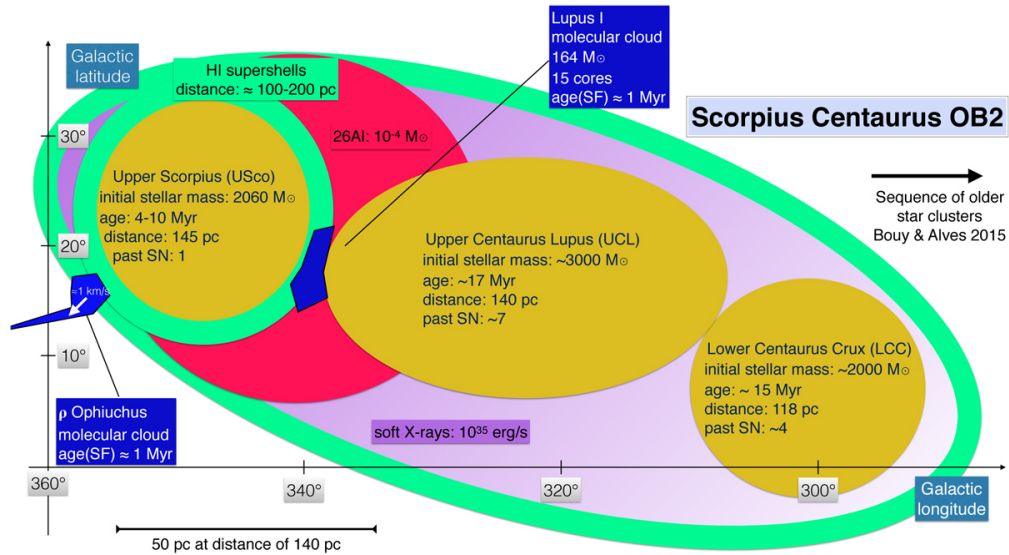


Figure 2.9: The Scorpius-Centaurus region as a potential target of ^{26}Al kinematic observations. HI supershells and ^{26}Al surround Upper Scorpius, the youngest OB subgroup in the region. Ophiuchus is an active site of stellar formation that may be inundated by flows of the neighboring ^{26}Al . Image from Krause et al. (2018).

speeds above and below the plane of the Milky Way of $\sim 1000 \text{ km s}^{-1}$ (Figure 2.11; Rodgers-Lee et al. 2019) may be an interesting tool for modeling ^{26}Al motion. However, more detailed observations of the Milky Way will be required for comparison to the LMC. While ^{26}Al imaging and correlations with spiral arms emphasize the Galactic Plane as the primary target of its emission, understanding of ^{26}Al at high Galactic latitudes, i.e. extending above and below the Galactic Plane, is largely unconstrained. Hydrodynamic models of merging superbubbles in the Milky Way suggest that “chimneys,” or vertical plumes, of ^{26}Al extend above and below the plane and blow out into the Galactic halo (Rodgers-Lee et al. 2019; Krause et al. 2021). The COMPTEL 1.8 MeV image (Figure 2.2) may show signs of this chimney structure, but an all-sky survey of ^{26}Al with fine angular resolution will be critical to testing the model.

2.3.3 Summary

Observations by SPI have informed a model of ^{26}Al dynamics closely associated with the spiral arms of the Milky Way Galaxy. Stellar activity creates low-density, merging HI superbubbles into which newly synthesized ^{26}Al flows. The injection of ^{26}Al into these regions at speeds greater than that of broader Galactic rotation is supported as well by GRIS data, albeit with slightly slower speeds, and disfavors HEAO-3’s result that ^{26}Al slows down to broader Galactic speeds before decay in the ISM.

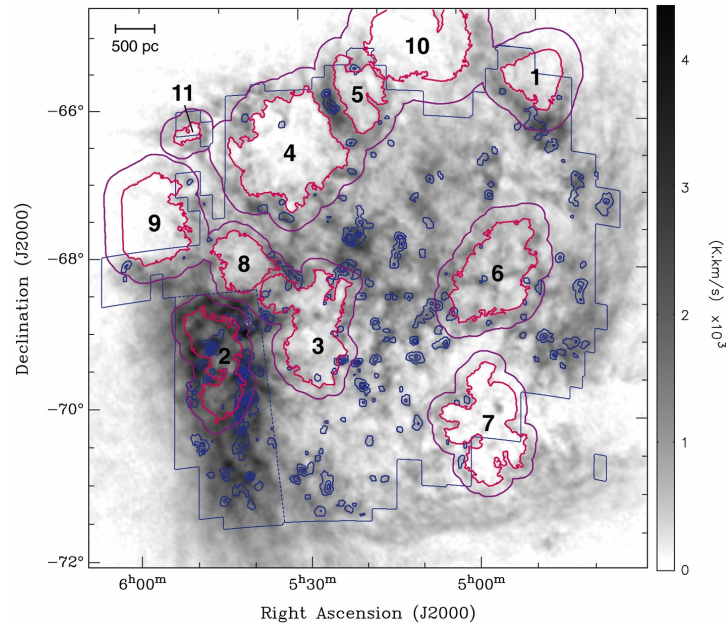


Figure 2.10: Velocity flows (grayscale) in the Large Magellanic Cloud. Dark pink and purple lines mark the inner and outer edges of HI supershells, respectively. Image from Dawson et al. (2013).

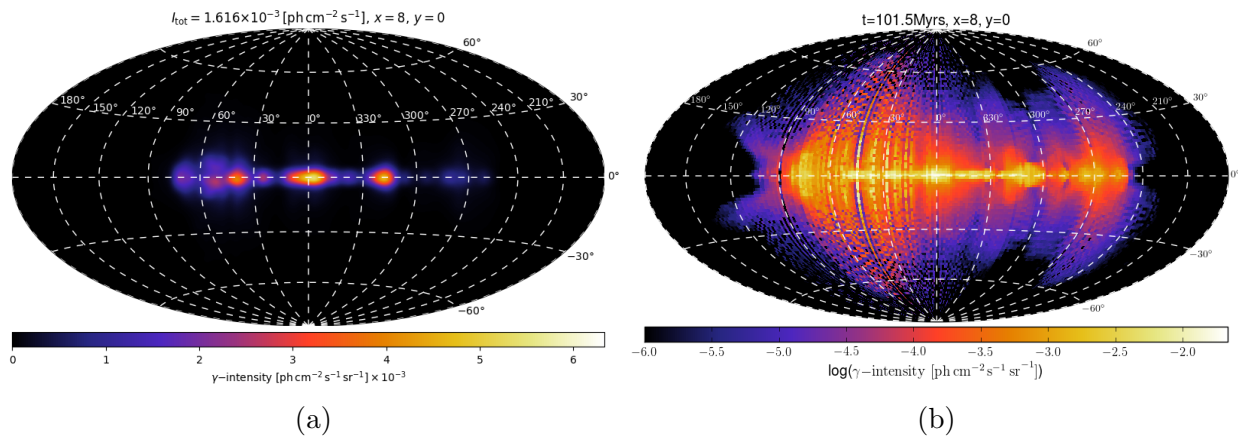


Figure 2.11: Simulated observations of ²⁶Al assuming the sensitivities and resolutions of (a) COMPTEL and (b) the satellite version of COSI, which may be able to see the orange tones of the logarithmic color scale. The image in (a) is from Rodgers-Lee et al. (2019). The image in (b) is a rescaling of that in (a) and is provided by Martin Krause, with permission.

Applications of this or similar models in localized areas of emission are an active area of research. The Orion-Eridanus region may serve as an additional test case of the superbubble scenario, and observations of Upper Scorpius with fine angular resolution may be sensitive to the motion of ^{26}Al around regions of young star formation like Ophiuchus. Simulations also indicate the potential to discern velocity profiles of ^{26}Al outside of the Milky Way in the LMC. The next section will summarize prospects for obtaining the necessary observational counterparts to these models.

2.4 Summary of current and future observations

This chapter established that spectroscopy and imaging of ^{26}Al emission from several astrophysics observatories, including HEAO-3, COMPTEL, and SPI, yield valuable insight into massive star activity in the Milky Way Galaxy. Imaging by COMPTEL and SPI found concentrated emission in the Inner Galaxy ($|\ell| \leq 30^\circ, |b| \leq 10^\circ$) and additional signs of emission in localized regions along the Galactic Plane. The Inner Galaxy flux as reported by COMPTEL is $3.3 \times 10^{-4} \text{ ph cm}^{-2} \text{ s}^{-1}$ and that from SPI is $(2.89 \pm 0.07) \times 10^{-4} \text{ ph cm}^{-2} \text{ s}^{-1}$ (Siegert 2017). The full sky flux from SPI is $(1.84 \pm 0.03) \times 10^{-3} \text{ ph cm}^{-2} \text{ s}^{-1}$ (Pleintinger 2020).

While the dominant sources of ^{26}Al are thought to be W-R stars and CCSNe, the relative contributions of each to the total observed mass of ^{26}Al remain uncertain. Other sources, including novae and AGB stars, have fallen out of favor as major contributors but remain viable progenitors. Comparing 1.8 MeV maps to those of other wavelengths provides strong evidence that ^{26}Al is correlated with young star-forming regions and massive stars that can support ionization of the ISM. Modeling of ^{26}Al production in Ophiuchus as an analog to the presence of ^{26}Al in our early Solar System endorses CCSNe as the dominant progenitor without ruling out a significant contribution from W-R stars. Complementary observations from other γ -ray emitters like ^{60}Fe and positron-electron annihilation may also help identify the origins of ^{26}Al , in conjunction with non- γ -ray studies of isotopic abundances in geologic material.

Measuring the Doppler shift and broadening of the 1.809 MeV line is also necessary to understand the nature of ^{26}Al emission. Additional observations are required to disentangle conflicting velocity measurements from HEAO-3, GRIS, and SPI, though SPI's model of preferential streaming into supershells at speeds exceeding broader Galactic rotation is generally accepted today. In particular, confirmation of SPI's longitude-resolved Doppler shift measurements by another observatory would lend credence to the models of ^{26}Al kinematics that associate its emission with spiral arms and superbubbles. Conducting these spectral studies with advanced imaging can reveal the dynamics of ^{26}Al in localized regions including Orion-Eridanus, Upper Scorpius and Ophiuchus, and the LMC. Excellent sensitivity is required to probe the faint fluxes of $\sim 10^{-5} \text{ ph cm}^{-2} \text{ s}^{-1}$ (Table 2.1) of these regions.

To date, there have been no measurements of ^{26}Al from individual stars. These measurements could constrain models of stellar nucleosynthesis by correlating stellar properties like

Region	Cygnus	Scorpius-Centaurus	Orion	Upper Scorpius
Flux [10^{-5} ph cm $^{-2}$ s $^{-1}$]	9.28 ± 1.75	7.58 ± 1.43	3.65 ± 1.19	3.17 ± 0.95

Table 2.1: Flux of ^{26}Al in massive star regions as measured by SPI and reported in Siegert (2017).

mass, metallicity, and rotation to measured yields of ^{26}Al . Although emission from individual stars is likely very faint, the binary system WR 11 (γ^2 Vel) in Vela contains the closest W-R star to the Sun (~ 377 pc) and may be a promising target (Oberlack et al. 1994). It is possible, however, that γ^2 Vel is too close to identify it as a single source; from Earth’s perspective, the ^{26}Al emission may diffuse so broadly that it instead appears as part of the greater Galaxy-wide emission.

The need for advanced instrumentation in pursuit of these observations is clear. The excellent energy resolution, improved sensitivity, fine angular resolution, and wide field of view of COSI as a Small Explorer satellite are designed to fill this prescription. The energy resolution (required 0.4% FWHM, expected 0.3%) will test SPI’s longitude-resolved measurements of ^{26}Al emission. Measuring ^{26}Al in localized regions (Table 2.1) is achievable given COSI’s required sensitivity of $< 3 \times 10^{-6}$ ph cm $^{-2}$ s $^{-1}$ (3σ) at 1.8 MeV (Tomsick et al. 2019) and can help identify other emission sites linked to particular progenitors. COSI’s 1.5° angular resolution (Tomsick et al. 2019), compared to the final image resolutions of $\sim 3.8^\circ$ by COMPTEL and $\sim 3^\circ$ by SPI, will also facilitate these measurements. Daily full-sky coverage will provide a comprehensive view of ^{26}Al emission across the Galaxy. Observing the entire sky is required to affirm the concentration of ^{26}Al along the Galactic Plane or reveal, for the first time, “chimney” structures extending towards higher latitudes. Hydrodynamic simulations of these structures indicate that these observations may be newly possible with COSI’s sensitivity (Figure 2.11). Note also that COSI’s improved sensitivity and all-sky coverage will produce the first-ever image of ^{60}Fe in the Galaxy, leading a great advancement in complementary nucleosynthesis observations.

Population synthesis simulations in Siegert et al. (2023) also underscore the importance of sensitive high-latitude observations. A full, bottom-up Population SYnthesis CODE (PSYCO) simulates the evolution of galaxies given parameters defining star formation rate, stellar metallicity, rotation, and explodability, initial mass function, superbubble expansion, nucleosynthetic yields, and more. Of all tested parameter combinations, the PSYCO all-sky map that most closely agrees with INTEGRAL/SPI data (Pleintinger et al. 2019) exhibits noticeable emission at high latitudes (Figure 2.12a with sensitivity 10^{-6} ph cm $^{-2}$ s $^{-1}$ sr $^{-1}$). However, when adjusted to match the $\sim 3^\circ$ resolution of the COMPTEL and SPI images and scaled with a minimum intensity of 5×10^{-5} ph cm $^{-2}$ s $^{-1}$ sr $^{-1}$ for more meaningful comparison with observations, the high-latitude features ($|b| \gtrsim 30^\circ$) in the PSYCO map are no longer visible (Figure 2.12b). Instruments with better sensitivity than COMPTEL and SPI

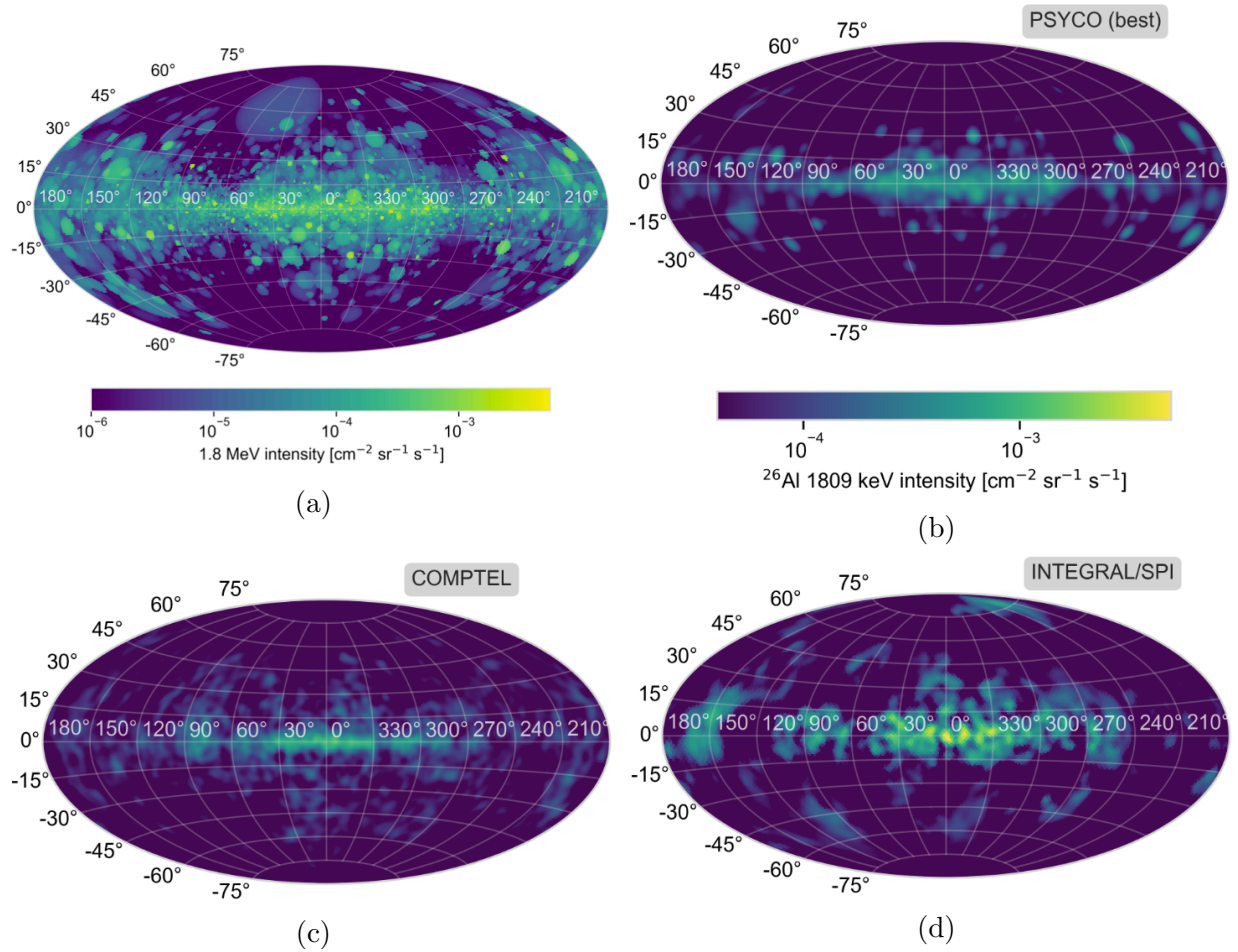


Figure 2.12: Maps of ^{26}Al from Population SYnthesis COde (PSYCO) simulations compared to COMPTEL and SPI observations. (a) The best-fit PSYCO map to SPI observations (Pleintinger et al. 2019) shows high-latitude emission, assuming a sensitivity of $10^{-6} \text{ ph cm}^{-2} \text{ s}^{-1} \text{ sr}^{-1}$. (b) Adapting the PSYCO map to have 3° resolution and minimum intensity of $5 \times 10^{-5} \text{ ph cm}^{-2} \text{ s}^{-1} \text{ sr}^{-1}$, similar to the capabilities of COMPTEL and SPI (maps of which are shown in (c) and (d), respectively, with the same color scale as in (b)), demonstrates that the high-latitude emission in (a) will remain obscured without improvements to current instrument sensitivity. All images from Siegert et al. (2023).

are therefore necessary to reveal the high-latitude emission of ^{26}Al predicted by theoretical models.

Motivated by these scientific goals, the remainder of this dissertation is dedicated to the principles of γ -ray detection, the development of hardware and analysis tools in the form of the COSI balloon mission, and the culmination of these efforts in the measurement of ^{26}Al in the COSI 2016 balloon flight.

Chapter 3

Soft Gamma-ray Instrumentation and the Compton Spectrometer and Imager

The Compton Spectrometer and Imager (COSI) is a compact Compton telescope designed to survey the soft γ -ray (0.2–5 MeV) sky. With its wide field of view, excellent energy resolution, intrinsic sensitivity to polarization, and imaging capabilities, COSI is uniquely equipped to perform studies of stellar and explosive nucleosynthesis, the origin of Galactic positrons, and polarization of compact objects in its historically under-explored energy range. The COSI instrument was originally designed as a balloon-borne telescope and at the time of writing, is being upgraded to a NASA Small Explorer satellite mission slated for launch in 2027. In this document, “COSI” refers to the balloon iteration of the instrument unless noted otherwise.

This chapter explores the rationale for COSI’s design with an introduction to several types of γ -ray telescopes. The historical usage of these telescopes in γ -ray astrophysics reveals measurement difficulties in the MeV range which weaken sensitivity in this bandpass. The advantages of using compact Compton telescopes with germanium detectors like COSI to close the MeV gap are presented, followed by an explanation of the operating principle of compact Compton telescopes. Finally, the COSI instrument itself is described in detail. The numerous scientific accomplishments of COSI’s successful 2016 balloon flight are proof of the instrument’s potency as a γ -ray explorer, and the analysis methods used to realize these scientific accomplishments set the stage for the analysis described in Chapters 5 and 6.

3.1 Gamma-ray telescopes

There are several typical γ -ray telescope designs, each with their own merits and drawbacks, that enable researchers to choose the type of telescope best suited to their research goals. Four designs are discussed below: the coded mask, classic Compton telescope, compact Compton telescope, and Laue lens.




Figure 3.1: Schematic view of the operating principle of SPI on INTEGRAL, a coded mask imager. Incident photons either pass through or are blocked by the tungsten mask (upper plane). Photons that pass through the mask hit the germanium detectors (lower plane) in a “shadowgram” pattern cast by the mask. Deconvolution of the pattern gives the source distribution of photons. Image from the European Space Agency ([INTEGRAL](#)).

3.1.1 Coded mask

A coded mask imager collects photons on a detector array positioned below a high atomic number (high-Z) “mask.” The mask blocks a subset of incident γ -rays from reaching the detectors and casts a “shadowgram” of the mask on the detectors. It is difficult to assign an origin to single photons and rather the full shadowgram formed by a collection of photons is deconvolved. This deconvolution technique transcribes the pattern of photons seen on the detectors (those unobstructed by the known geometry of the mask) into an origin on the sky.

Primary examples of coded mask instruments in the MeV γ -ray bandpass are SPI and the Imager on Board the INTEGRAL Satellite (IBIS). SPI is a coded mask instrument comprised of 19 high-purity germanium semiconductor detectors obscured by a tungsten mask (Vedrenne et al. 2003). A schematic of the SPI aperture which illustrates the coded mask operating principle is shown in Figure 3.1. IBIS also uses a coded mask to obscure cadmium telluride (CdTe) and cesium iodide (CsI) detectors (Winkler et al. 2003).

An advantage of coded mask instruments is that the instrumentation for the mask and detector array is relatively straightforward compared to the multi-hit reconstruction required for compact Compton telescopes (Section 3.1.3). If a photon hits the detector array, it is

recorded as an event. The downside to this approach is that coded mask telescopes are susceptible to high background with no intrinsic method of discriminating between source photons of interest and background photons. Event reconstruction in compact Compton telescopes, by contrast, permits such discrimination. For this reason, background determination in coded masks often relies on taking “empty field” observations of regions of the sky thought not to contain signal. This approach can be problematic when the spatial extent of a source is not well-constrained and it is unclear where to expect zero emission. Ultimately, the fine angular resolution of coded masks and need to observe away from a known source optimizes them for point source observations. Large-scale diffuse imaging is also achievable, though with diminished sensitivity to emission with low gradients in flux.

3.1.2 Classic Compton Telescope

A classic Compton telescope induces one Compton scatter of an incident photon in a detector plane mounted above a secondary plane. The scattered photon deposits all of its energy into the secondary plane in a single photoabsorption event. The first detector plane is called the “low-Z scatterer” and the secondary plane is called the “high-Z absorber” because the comparatively smaller atomic number in the former is low enough to induce a Compton scatter while the higher atomic number in the latter is sufficiently high to induce photoabsorption. The vertical distance between the detectors is large enough to resolve the time of flight between the scatter and photoabsorption and deduce the correct order of interactions. The locations and energies of each hit are recorded so that the initial energy and direction of the incident photon can be determined with the Compton equation.

COMPTEL was a “classic” Compton telescope with a low-Z liquid scintillator scatterer and a high-Z NaI(Tl) absorber separated by 1.5 m. It detected 1–30 MeV photons with a wide field of view of $\sim 1\pi$ sr and an angular resolution of $\sim 1\text{--}2^\circ$ (1σ width of the angular resolution measure distribution; see Section 4.8.2 for more information on this measure) (Schönfelder et al. 1993). COMPTEL could not track the recoil electron in the Compton scatter and thus reconstructed the origin of each incident photon to a circle on the sky (“event circle”). The overlap of event circles from several source photons marks the true source location (refer to Section 3.3 for more details).

3.1.3 Compact Compton Telescope

As in the “classic” version, the compact Compton telescope (CCT) relies on Compton scattering to reconstruct the energy and direction of incident photons. However, CCTs are comprised of several large volume ($\sim 10^2$ cm³) solid-state semiconductor detectors arranged in one detector volume. Rather than one Compton scatter and one photoabsorption event, the incident photons often undergo up to eight interactions inside a CCT. The techniques required to reconstruct these interactions into full events help discriminate signal from background photons. For example, events with eight or more interactions are likely of sufficient initial energy to pair produce an electron and positron; these events cannot be reconstructed

and are rejected from analysis (Boggs and Jean 2000). This intrinsic background suppression is one of the primary advantages of CCTs compared to coded mask instruments. For a detailed explanation of the operating principle and event reconstruction of Compton telescopes, refer to Section 3.3.

An additional advantage of CCTs is their enhanced sensitivity to photon polarization. While classic Compton telescopes can in principle measure polarization, the large distance between scattering planes is less than ideal for detecting polarized photons, which preferentially Compton scatter at 90° . In COMPTEL, incident photons could not scatter through an angle larger than that subtended by the high-Z absorber below the low-Z scatterer. Only photons which hit the absorber after scattering out of the scatterer could be analyzed. Conversely, photons incident upon the large detector volume of a CCT remain detectable after Compton scattering at 90° and over a range of angles much larger than those allowed in a classic Compton telescope. This flexibility increases the overall efficiency of CCTs compared to the classic design.

Finally, a practical advantage of CCTs is that the instruments are, as the name suggests, physically compact. The detectors are stacked close together as a composite active volume and no substantial separation between detector planes is required, as is the case for classic Compton telescopes. Compact instruments can more readily be launched as science payloads on a balloon platform and can simplify the design of a satellite mission. The Compton Spectrometer and Imager (COSI; Section 3.4) is a notable realization of the CCT design. Like COMPTEL, COSI has a wide field of view of $\sim 1\pi$ sr and constrains the origin of a photon to an event circle on the sky (no electron tracking). Its germanium semiconductor detectors improve upon the energy resolution of COMPTEL's scintillator detectors by an order of magnitude (refer to Section 3.2.2 for a discussion about the advantages of germanium detectors). The combination of its wide field of view, excellent energy resolution, and polarization capabilities distinguish COSI as a powerful γ -ray observatory.

3.1.4 Laue lens

The coded mask, classic Compton telescope, and CCT designs discussed previously collect γ -rays on a large detector plane. They do not focus incident photons to a single point, as is commonplace in optical astronomy, for example. Focusing the γ -rays to a detector plane much smaller than the area of the collecting “lens” would dramatically reduce the high background levels which plague these telescopes. Using Laue diffraction to focus γ -rays from rings of crystals to a focal point is thus an active area of study. The basic principle of Laue lens instruments is that if the Bragg relation ($2d\sin\theta = n\lambda$) is satisfied, γ -rays can pass through the lens. The distinction between Bragg diffraction and Laue diffraction is that in Bragg diffraction, the diffraction occurs near the surface of the crystal rather than inside of the crystal as with Laue. In addition to their reduced background, Laue lenses have excellent angular resolution. The costs of these advantages, though, are a very small field of view and a severely restricted energy bandpass.

CLAIRE is a balloon-borne γ -ray telescope which uses Ge(Si) crystals to focus incident γ -rays onto a high-purity germanium detector array. The chosen diffraction energy is 170 keV with a bandpass of about 3 keV. Two balloon flights in 2000 and 2001 did not reveal the desired Gaussian peak signal of the Crab Nebula at 170 keV, but the flights demonstrated that launching Laue lenses for γ -rays is possible on a balloon platform (Halloin et al. 2003).

Efforts to develop a “tunable” Laue lens, whose crystals could be tilted during orbit, are underway. These tunable instruments would expand the traditionally limited energy bandpass of a few keV to about 0.2–2.5 MeV (Lund 2021). Successful implementation of the tunable lens would provide the community with MeV measurements of excellent angular resolution to complement those of other instruments.

3.2 Exploring the MeV gap with germanium CCTs

The strengths and weaknesses of the γ -ray instruments presented above tailor them to particular scientific tasks. Coded mask instruments are relatively simple in design and have fine angular resolution for point sources, but they cannot easily resolve low-gradient emission or reconstruct the origin of an individual incident photon. Classic Compton telescopes are more sensitive to low-gradient emission with single-photon reconstruction, but historically have had poorer energy resolution (e.g., scintillator detectors in COMPTEL) than germanium coded masks and CCTs (e.g., SPI and COSI, respectively). They also do not have the higher efficiency, compact design, and polarization sensitivity of CCTs. Laue lenses have excellent angular resolution but a limited energy bandpass and small field of view. Thus, this section¹ serves specifically to convey the unique instrumental needs of MeV γ -ray astrophysics. After establishing the historical measurement difficulties in this bandpass, high-purity germanium is presented as a preferred detector medium for pursuing scientific goals in the MeV range.

3.2.1 The MeV gap

The MeV bandpass (~ 0.1 –100 MeV) of γ -ray astrophysics is rich in scientific potential. Understanding the nature of Galactic positron-electron annihilation, studying emission from various sources of nucleosynthesis across the Galaxy, performing novel polarization measurements of GRBs and compact objects on cosmological scales, and conducting follow-up measurements of multi-messenger sources have motivated decades of astrophysics experiments in this energy regime. However, these experiments have historically suffered from poorer sensitivity than experiments that study photons of lower and higher energy. The 0.1–100 MeV range is consequentially referred to as the “MeV (sensitivity) gap” (e.g., Siegert

¹This section draws from the contributed chapter Beechert et al. (2023, in review). “Germanium detectors for MeV gamma-ray astrophysics with the Compton Spectrometer and Imager”. In: “X-ray Photon Processing Detectors: Medical Imaging and Industrial Applications.” Ed. by Conny Hansson and Krzysztof Iniewski. ©Springer Nature Switzerland AG 2023.

et al. 2022). This diminished instrument sensitivity (Figure 3.2a) is attributed to low overall photon-matter interaction cross-sections, high instrumental background, and instrumental constraints. Each of these factors is explored below.

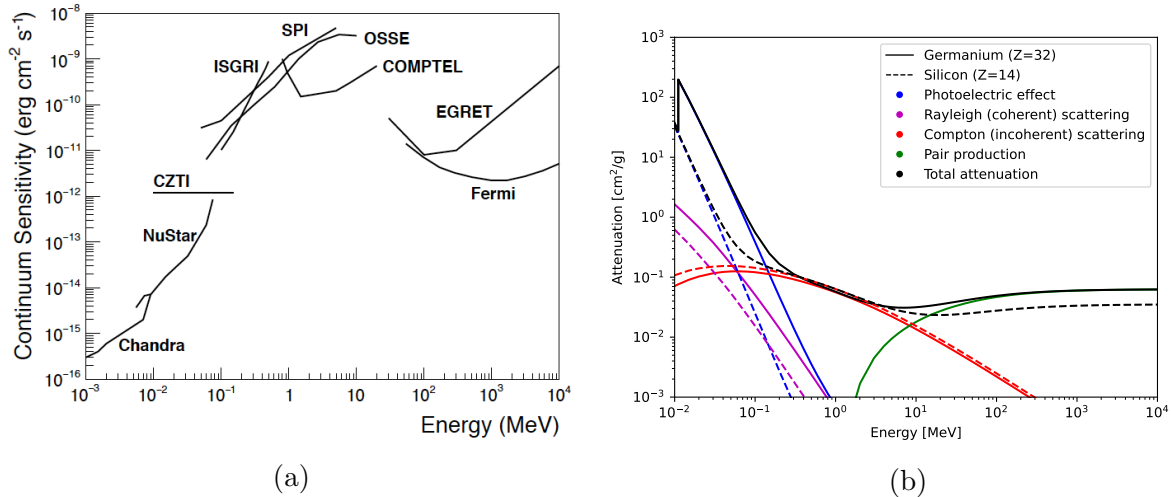


Figure 3.2: Instrument sensitivities and attenuation of photons in the MeV gap. (a) Instruments in the “MeV gap” (0.1–100 MeV) have poorer continuum sensitivity than those in adjacent energy ranges. Continuum sensitivity is a measure of minimum detectable flux. Image from Kierans (2018). (b) Attenuation of photons in germanium and silicon, two common semiconductors used for γ -ray detection. Compton scattering is the dominant interaction mechanism in the MeV gap and reaches an overall minimum cross-section near a few MeV. Attenuation data are from NIST XCOM (Berger et al. 2010).

Low interaction cross-sections

Photons are measured via interactions with matter, i.e. with the detector material. The interaction cross-section, or probability, depends on the energy of the photon and the properties of the material. Thus, it is desirable to choose materials with high cross-sections for detection of photons in a given energy range. However, photons between $\sim 0.1 - 10$ MeV have lower overall interaction cross-sections in common detector materials like silicon (atomic number $Z = 14$) and germanium ($Z = 32$) than photons in adjacent energy ranges (Figure 3.2b). Compounding the disadvantage of lower overall cross-sections is that the dominant interaction mechanism in this range is Compton scattering. MeV γ -ray detectors must maximize detection efficiency by inducing and tracking Compton scatters, the process of which is inherently complex. Rather than depositing all energy at a single interaction location, as is the case in photoelectric interactions, photons which Compton scatter typically interact between two and seven times in the detector volume, each at a different location and with changing energy. Sophisticated event reconstruction algorithms are required to disentangle

gle this sequence of Compton scatters. An overview of the operating principle of Compton telescopes is provided in Section 3.3.

High instrumental background

MeV astrophysics missions must combat a unique source of background called “instrumental activation.” Instrumental activation refers to the nuclear excitation (or “activation”) of instrument materials to higher energy states which then de-excite on the time scale of desired measurements via the emission of an MeV γ -ray. High-energy particles like cosmic rays often induce the excitation in MeV missions deployed on high-altitude balloons or as satellites. The nuclear lines from instrumental activation are problematic because they must be distinguished from signal photons; these background γ -ray lines can blend with nearby astrophysical lines of interest or even overlap with them completely, hindering high-resolution spectroscopic studies. Characterizing these background γ -rays is a challenging task. Substantial simulation time, estimations of high-energy particle fluxes in the instrument’s orbit, and dedicated observation time of background irradiation on γ -ray instruments are all useful methods of disentangling activation lines from astrophysical γ -ray sources.

Additional instrumental constraints

Effects intrinsic to Compton scattering, detector manufacturing, and high-altitude orbits pose additional constraints on MeV photon detection. Compton scattering of MeV photons produces long-range secondaries. Scattered photons and electrons travel non-negligible distances in detector material and can even reach sufficient energy to pair produce. These several centimeter-long tracks must be fully contained in the detector volume. Scintillators can be manufactured in large volumes that can contain these tracks, but they do not have the fine energy resolution of a segmented solid state detector that can reconstruct events with accurate energy measurements at each interaction location. It is difficult, though, to manufacture large solid state detectors and stacking large quantities of smaller detectors is not necessarily a feasible option. Limitations on payload weight in high-altitude balloon and satellite platforms place instrument mass at a premium.

Additionally, instruments at these high-altitude orbits are susceptible to dominating atmospheric background contamination. Particle and photon interactions in the Earth’s atmosphere emit γ -rays that cumulatively glow as strong “Earth albedo radiation.” Heavy shielding is required to reject, if imperfectly, Earth albedo events emanating from below the instrument orbit. Balloon altitudes below and at nominal float altitudes of $\gtrsim 33$ km are particularly vulnerable to atmospheric background contamination. Moving to further displaced satellite orbits (e.g., low-Earth orbit of ~ 500 km) can help limit the influence of the atmosphere, though difficulties from instrumental activation remain.

Thus, MeV astrophysics relies upon the detection of photons via Compton scattering, an inherently complex sequence of interactions which has the lowest overall probability of occurrence compared to interaction mechanisms in adjacent energy ranges. The background in the MeV regime is perhaps even more complicated. Combating instrumental activation

and atmospheric background, the latter of which is an unavoidable consequence of the need to place the detectors at high altitudes (Earth’s atmosphere is opaque to MeV γ -rays), requires the development of sophisticated analysis techniques and physical shielding in the instrument design. This instrument design must also fully contain the Compton-scattered photons (and pair production events, if those are desired) without exceeding the allowed mass and volume budgets of the launch platform. The germanium detectors described below help satisfy these instrument requirements.

3.2.2 High-purity germanium detectors

The ability to manufacture germanium semiconductor detectors of thicknesses up to several centimeters, induce Compton scattering by nature of germanium’s atomic number, and perform spectroscopy with excellent energy resolution make germanium a popular detector material that can mitigate some of the challenges described above. The sections below elaborate on these advantages with comparisons to other solid-state detectors and scintillators.

Large volume, high purity: Modern fabrication techniques

An incident photon must deposit all of its energy in the detector volume for complete Compton event reconstruction (see Section 3.3). As mentioned previously, Compton scattering produces long-range secondaries that can travel non-negligible distances, motivating the need for detectors which are sufficiently large to fully contain all Compton scattering interactions and the final photoelectric absorption. Fabrication techniques developed in the 1960s produced germanium detectors with lithium ion drifting (called Ge(Li) detectors; Freck and Wakefield 1962; Tavendale and Ewan 1963; Henck et al. 1968; Malm 1967) and achieved thicknesses of up to ~ 2 cm. Advanced techniques in the 1970s also succeeded in fabricating large-volume germanium detectors with a procedure called zone refining (Knoll 2010). In zone refining, impurities are compensated through local heating of the bulk germanium instead of with donated lithium ions. The impurities are attracted to the melted germanium in the heated zones; sequentially moving the heat source moves the melted material throughout the volume of germanium, picking up impurities along the way and eventually removing them from the sample. Larger high-purity germanium crystals are grown from this seed, creating the large-volume detectors necessary for MeV detection. Furthermore, these modern techniques reduced initial crystal impurity concentrations from $\sim 10^{11} - 10^{12}$ atoms cm^{-3} (Hall and Soltys 1971; Hansen 1971; Llacer 1972) to very low impurity levels of about 10^{10} atoms cm^{-3} (Amman 2020). The low impurity levels in high-purity germanium detectors reduce the risk of detector effects like charge trapping, for example, which impair signal readout.

Atomic number

Germanium’s atomic number ($Z = 32$) is well-suited to MeV γ -ray detection with Compton telescopes because it preferentially Compton scatters photons in this energy range, rather than inducing a premature photoabsorption or a pair production event which cannot be reconstructed (Figure 3.3). The probability of a Compton interaction scales with atomic

number and energy as $\sim Z/E$, while that for photoelectric absorption scales as $\sim Z^4/E^{3.5}$. Pair production scales as $\sim Z^2 \ln(E - 2m_e c^2)$. An additional physical advantage of germanium is that its density of 5.32 g cm^{-3} (at 300 K) is high enough to fully contain scattering events in compact volumes. Silicon, for example, is less dense at 2.33 g cm^{-3} and requires more material to stop photons of the same energy (Knoll 2010). This can be further illustrated by the half-value layer, or the material required to stop half of the incident photons: that of germanium at 511 keV is $\sim 1.6 \text{ cm}$ and that of silicon is $\sim 3.4 \text{ cm}$. Note that scintillators with higher atomic number, like iodine ($Z = 53$) in sodium-iodide detectors, may have higher stopping power than germanium, but will not have the excellent energy resolution discussed in the next section.




Figure 3.3: Dominant interaction mechanisms as a function of detector atomic number Z and photon energy $E = h\nu$. The solid lines indicate equally probable interaction mechanisms. Compton scattering is the dominant interaction mechanism of photons between ~ 0.1 – 10 MeV in germanium ($Z = 32$). Image from Evans (1955).

Energy resolution

Germanium has excellent energy resolution because of its low ionization energy (2.96 eV at 77 K) and small Fano factor (0.11; Croft and Bond 1991). Ionization energy is loosely defined as the minimum energy required to generate one electron-hole pair in an interaction. As energy resolution scales inversely with the square-root of the number of charge carriers, the statistical energy resolution of a detector improves with smaller ionization energy. Silicon and cadmium zinc telluride (CZT) semiconductor detectors have coarser energy resolution than germanium with ionization energies of 3.6 eV and 4.6 eV, respectively. Gas detectors have even greater ionization energies of around 30 eV.

Furthermore, the $\sim 0.5\%$ energy resolution of HPGe detectors at 662 keV (Beechert et al. 2022a) outperforms the 6–7% resolution at 662 keV of CsI(Na) and NaI(Tl) scintillators

(Syntfeld-Kazuch et al. 2010). The Fano factor is largely responsible for this difference. The limiting Poisson (statistical) resolution is governed by $1/\sqrt{N}$, for N liberated charge carriers. In practice, however, energy resolutions 3 to 4 times finer than the Poisson limit are observed (Knoll 2010), suggesting that charge carriers are not necessarily generated independently in a pure Poisson process. The Fano factor quantifies this departure from statistical performance as the ratio between the observed variance in N and the Poisson-predicted variance ($= N$). Materials with small Fano factors like germanium and other semiconductor detectors thus outperform their limiting statistical resolution, while scintillators with Fano factors of ~ 1 are Poisson-limited to resolutions coarser than in germanium.

Overall, the energy resolution of HPGe detectors is paramount to closing the MeV gap. High-resolution spectroscopy is critical to studying the radioactive signatures and dynamics of MeV γ -ray sources across the Milky Way Galaxy and beyond. Distinguishing between γ -ray lines from instrumental activation and nearby astrophysical γ -ray lines is also only achievable with high energy resolution.

Additional considerations

There are a few complications to using HPGe worthy of consideration. Scintillators are more easily manufactured than HPGe detectors and can be operated at room temperature. Germanium must be maintained at cryogenic temperatures of approximately 77–84 K. Higher temperatures would excite electrons across the small 0.7 eV bandgap, creating prohibitive leakage current that obscures real signals. Although cooling is routinely achieved in near-vacuum ($\lesssim 10^{-6}$ Torr) cryostats with refillable liquid nitrogen dewars or mechanical cryocoolers, working within the constraints of these safe conditions is still a potential hindrance. Warming germanium detectors for long-term storage or transport is a delicate procedure often performed over the course of several days. Additionally, long-term room temperature storage and repeated thermal cycling of amorphous germanium contacts may worsen leakage current and long-term stability (Looker 2014).

Room temperature detectors like scintillators and CZT, on the other hand, can be operated outside of cumbersome cryostats and be readily stored or transported without gradual acclimation to standard air temperature and pressure. This makes room-temperature detectors attractive to technologies that are designed for hand-held or mobile use (e.g., Haefner et al. 2017). Thus, in applications where high-resolution spectroscopy is not the highest priority, detector materials like gas, scintillators, or room-temperature semiconductors may be preferred.

3.3 Principles of gamma-ray detection in Compton telescopes

The previous sections outlined a variety of telescope designs used to study γ -rays, the unique difficulties of studying astrophysical MeV γ -rays in the so-called “MeV gap,” and why HPGe

is a suitable medium for detecting γ -rays in this energy regime. Fundamentally, the dominant interaction mechanism of MeV photons interacting with germanium is Compton scattering and it is therefore prudent to use instruments whose measurement techniques are sensitive to the Compton scattering process. This section¹ explains the operating principle of Compton telescopes as an introduction to the realization of a HPGe Compton telescope in COSI (Section 3.4).

In an ideal event, an incident photon Compton scatters throughout the detector volume and undergoes a final interaction as a photoelectric absorption. The deposited energy and position of each interaction site are recorded. The Compton equation (derived in Appendix D) defines the energy $E_{\text{scattered}}$ of a photon after it Compton scatters at angle ϕ with the detector material:

$$E_{\text{scattered}} = \frac{E_0}{1 + \frac{E_0}{m_e c^2} (1 - \cos \phi)}. \quad (3.1)$$

Figure 3.4 illustrates a typical 3-site event: a photon Compton scatters at locations $\vec{r}_{1,2}$ and deposits energy $E_{1,2}$. It deposits energy E_3 in a final photoabsorption interaction at \vec{r}_3 . The energy of the incident photon is $E_0 = E_1 + E_2 + E_3$ and its origin is constrained to a circle on the sky, called an “event circle,” defined by the cosine of the first Compton scattering angle ϕ :

$$\cos \phi = 1 - \frac{m_e c^2}{E_2 + E_3} + \frac{m_e c^2}{E_1 + E_2 + E_3}. \quad (3.2)$$

The overlap of multiple photons’ event circles marks the source location. Some Compton telescopes are able to track the recoil electron in each Compton scatter, thereby recovering additional kinematic information that can reduce the circles to arcs on the sky. Arcs limit the number of possible photon origins and as such can improve background discrimination. Note that electron tracking does not improve angular resolution because the radial distance from the true source location to the arc is the same as that to the event circle (see Section 4.8.2 for more information about the angular resolution of Compton telescopes).

Recall that in classic Compton telescopes, an incident photon Compton scatters once in a scattering plane that is spatially displaced from an absorbing plane. The displacement is large enough to temporally resolve the order of interactions and easily determine the Compton scattering angle that defines the event circle. Event reconstruction in CCTs is more complicated because the timing resolution is too coarse to resolve the sequence of interactions within the compact detector volume. In COSI, the ~ 10 ns resolution is insufficient to resolve sequential interactions that occur on sub-nanosecond scales. As such, instead of unambiguously determining the order of interactions with absolute time tags, the most likely γ -ray path must be determined through interaction probabilities and kinematics. The Klein-Nishina differential cross-section (Equation 3.3; Klein and Nishina 1929) defines the

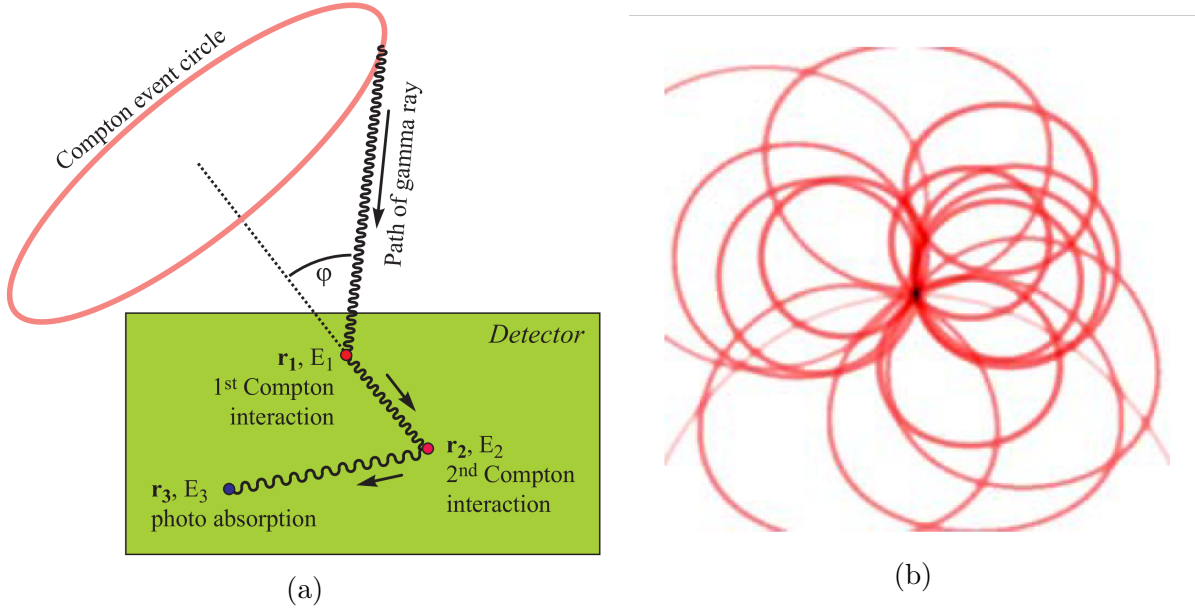


Figure 3.4: Illustration of a typical γ -ray event and source localization in a CCT. (a) An incident photon of energy $E_0 = E_1 + E_2 + E_3$ Compton scatters at \vec{r}_1 and \vec{r}_2 , deposits energy E_1 and E_2 , and deposits energy E_3 in a photoabsorption at \vec{r}_3 . The Compton scattering angle (ϕ) of the first Compton scatter restricts the source location to an event circle on the sky. Image from Zoglauer et al. (2021). (b) The intersection of overlapping event circles marks the source location.

probability of a Compton interaction given its measured scattering angles (Compton: ϕ , azimuthal: η) and energy (incident: E_0 , scattered: $E_{\text{scattered}}$):

$$\frac{d\sigma}{d\Omega} = \frac{r_0^2}{2} \left(\frac{E_{\text{scattered}}}{E_0} \right)^2 \left(\frac{E_0}{E_{\text{scattered}}} + \frac{E_{\text{scattered}}}{E_0} - 2 \sin^2 \phi \cos^2 \eta \right), \quad (3.3)$$

where r_0 is the classical electron radius. Known properties of the detector material, γ -ray energy, and probable distances between subsequent interactions further constrain the scattering path of the photon. The order of interactions can also be derived from redundant measurements of the Compton scattering angle: kinematic and geometric determinations of the angle will only be equal, within uncertainties, for the correct order of interactions (Boggs and Jean 2000; Zoglauer 2005). More information about advanced approaches to event reconstruction, including a Bayesian approach, a random forest of decision trees, and shallow neural networks, can be found in Zoglauer (2005) and Zoglauer and Boggs (2007). Comprehensive reviews of Compton telescopes are provided in Zoglauer (2005); Zoglauer et al. (2021); Siegert et al. (2022); Kierans et al. (2022).

3.4 The Compton Spectrometer and Imager

3.4.1 Detectors

COSI is a CCT (Section 3.1.3) comprised of twelve double-sided cross-strip HPGe semiconductor detectors (GeDs; Section 3.2.2). Each GeD is approximately $8 \times 8 \times 1.5 \text{ cm}^3$ (Figure 3.5a). The top and bottom faces of each detector are instrumented with 37 aluminum strips of 2 mm strip pitch, defined as the distance between the center of each strip, oriented such that the strips on one side of each detector are orthogonal to those on the opposing side. A gap between strips of 0.25 mm was chosen to balance the desire for minimal charge loss between strips with stringent requirements on energy resolution, the latter of which improves with greater spacing between strips (decreased capacitance). A 2 mm-wide guard ring surrounds the active area of each face of the detectors to reject events with interactions close to the edge of the crystal and to minimize leakage current. The COSI GeDs were developed using the Lawrence Berkeley National Laboratory’s amorphous germanium contact technology (Amman 2020).

The two-dimensional x - y position of an interaction is defined by the intersection of orthogonal triggered strips. The strip pitch of 2 mm defines the x and y position resolutions. The z -position, i.e. the depth of the interaction in the detector, is determined through a depth calibration that maps the timing difference between the collection of electrons on the anode side of the detector and the holes at the cathode (Section 4.6). The z resolution is approximately 0.5 mm and the three-dimensional position resolution is $2 \text{ mm} \times 2 \text{ mm} \times 0.5 \text{ mm} \sim 2 \text{ mm}^3$.

The detectors are stacked in a $2 \times 2 \times 3$ configuration (Figure 3.5b). They operate as over-depleted p-i-n junctions with bias voltages between 1000 and 1500 V. The high-voltage side of each detector is AC-coupled and accordingly referred to as the “AC side” (anode). The low-voltage side (grounded, 0 V) is called the “DC side” (cathode). To reduce the influence of electronic noise, the opposing AC- and DC-side strips must trigger within a coincidence window of 360 ns.

3.4.2 Cryostat

The GeDs are held at cryogenic temperatures because germanium has a small bandgap of $\sim 0.7 \text{ eV}$. Higher temperatures would thermally excite electrons in the valence band to the conduction band and produce significant leakage current. To this end, the stack of detectors shown in Figure 3.5b is surrounded by an infrared shield and thermal insulation and is housed inside an anodized aluminum cryostat (Figure 3.6). The cryostat is evacuated to very low pressure, $\sim 10^{-6}$ Torr, to minimize thermal energy imparted to the detectors by surrounding air. A mechanical cryocooler is mounted to the outside of the cryostat to maintain the GeDs at cryogenic temperatures of $\sim 83\text{--}84 \text{ K}$.




Figure 3.5: COSI's cross-strip germanium detectors (GeDs) and full detector stack. (a) One COSI GeD. The 37 vertical strip electrodes of 2 mm pitch are shown in the foreground. A mirror shows the orthogonal strips on the opposite detector face. Image from Sleator et al. (2019). (b) The $2 \times 2 \times 3$ stack of twelve GeDs which form the COSI active volume. Each detector is $\sim 8 \times 8 \times 1.5 \text{ cm}^3$. Image from Kierans et al. (2016).




Figure 3.6: The COSI cryostat. The cryostat houses the twelve GeDs in cryogenic low-pressure ($\sim 10^{-6}$ Torr) conditions. (a) A mechanical cryocooler maintains the GeDs at cryogenic temperatures of $\sim 83\text{--}84 \text{ K}$ and is visible in the foreground. Image from Beechert et al. (2022a). (b) Preamplifier boxes for electronic readout are visible on the sides of the cryostat. Image courtesy of the COSI collaboration.

3.4.3 Cryocooler and thermal control

The ultra-long duration superpressure balloon on which COSI flies (Section 3.4.7) can support a science payload of approximately 2000 pounds (~ 900 kg). In its full flight configuration, COSI weighs about this maximum limit and cannot carry large reserves of liquid nitrogen to cool the GeDs for an extended period of time. As NASA's nominal goal is to maintain a 100+ day flight, deploying large dewars of liquid nitrogen is not a viable option. Using a mechanical cryocooler is preferred.

COSI uses a Sunpower CryoTel CT mechanical cryocooler (Figures 3.6a and 3.7a). The cryocooler is compact at 10.2 inches long, including the passive balancer installed on one end, 3.3 inches in diameter, and weighs 3.1 kg. It operates via the Stirling cycle between 60 and 160 W with a lift capability of 11 W at 77 K (35°C heat reject) (CryoTel). The copper cold tip of the cryocooler is thermally coupled to the copper coldfinger inside the COSI cryostat, which is in turn coupled to all twelve GeDs (Figure 3.7b).

The cryocooler can operate in two modes: constant temperature and constant power. In constant temperature mode, the cryocooler maintains its cold tip at a specified temperature by frequently adjusting its expended power. In constant power mode, the cryocooler is instead held at a specified power. During the COSI 2016 balloon flight, the cryocooler was operated in constant temperature mode. The cold tip temperature was set to 77 K and the cryocooler expended ~ 100 W in maintaining this target. During COSI 2020 ground calibrations (Section 3.5.3), the cryocooler was operated in constant power mode at 95 W in an attempt to stabilize power-dependent cryocooler vibrations. In practice, this change did not significantly impact instrument performance and there was consistency between the 2016 and 2020 calibrations (see Chapter 4). Both modes resulted in a detector temperature of ~ 83 – 84 K. Inevitable thermal losses between the cold tip, coldfinger, and GeDs result in a temperature which is slightly higher than that specified in constant temperature mode.

As the cryocooler expends power, it heats up and becomes less efficient with increasing temperature, driving it to higher power and further increasing temperatures. The decreased efficiency with increasing cryocooler temperature is seen in Figure 3.8. The cryocooler's lift, which is the maximum load in Watts that it can cool to a certain temperature, is plotted against the desired detector temperature for a variety of cryocooler temperatures. At a detector temperature of 77 K, for example, the cryocooler's lift decreases as the cryocooler temperature increases from 35°C to 45°C .

The positive feedback loop of increasing power and temperature is dangerous because internal temperatures in excess of 70°C can damage the cryocooler material. It is therefore imperative to keep the cryocooler itself cool as it works to maintain steady cryogenic temperatures at its cold tip. COSI employs a liquid cooling system to transport excess heat away from the cryocooler.




Figure 3.7: The COSI cryocooler and coldfinger. (a) The Sunpower CryoTel CT mechanical cryocooler keeps the germanium detectors inside the COSI cryostat at cryogenic temperatures. The cold tip is labeled. Image from [CryoTel](#). (b) The COSI cryostat opened to show the U-shaped copper coldfinger coupled to the GeDs, one of which is removed for clearer visibility in this photograph. Image courtesy of the COSI 2016 team.




Figure 3.8: Sunpower CryoTel CT cryocooler lift vs. cold tip temperature. As the “reject” temperature of the cryocooler’s collar, which is the cryocooler’s primary site of heat rejection, increases, the lift decreases at a given cold tip temperature. Hence, the cryocooler can only maintain smaller loads at higher reject temperatures. This is interpreted as a decrease in efficiency with increasing reject temperature and motivates the need for COSI’s liquid cooling system, which prevents the cryocooler from reaching temperatures above the 30–40°C range. Image from [CryoTel](#).




Figure 3.9: The COSI liquid cooling system as seen (a) on the cryocooler itself and (b) from the top of the gondola. Image from Sleator (2019).

3.4.3.1 The Liquid Cooling System

Figure 3.9a shows the liquid cooling system (LCS) as installed on the cryocooler itself. The cryocooler’s primary site of heat rejection is called the “collar.” The copper heat sink around the collar removes heat from the site. The temperature of the collar is the thermostat of the liquid cooling system: if the collar gets too hot, the liquid cooling system works to lower the temperature. The copper cooling fluid jacket around the body of the cryocooler also regulates the temperature of the cryocooler. Though the copper of the collar’s heat sink and body jacket is intrinsically conducive to heat transfer, the cylindrical pipes attached to both remove the most heat. High-pressure rubber Tygon tubes connect to the barbed fittings and deliver the liquid of the LCS (3M Novec 7200) to the pipes (McMaster tubing; McMaster hose). This cooling fluid is optimized for heat transfer, remains in liquid form over a wide temperature range to avoid freezing during flight, is not conductive, is non-flammable, and is of low toxicity (3M). The latter three characteristics are important to prevent damage to COSI in the event of an accidental spill. Also visible in Figure 3.9a are the bellows and copper foils which dampen vibrations induced by the oscillating cryocooler.

Driven by one of two Fluido-Tech FG209 magnetic drive pumps (Fluid-o-Tech) (two for redundancy, should one fail), the fluid circulates around the cryocooler, picks up expended heat, and then travels away from the cryocooler to the underside of the cooler radiator plate, to which it transfers the expended heat (Figure 3.9b). The radiator plate is a solid copper structure mounted to the top of the gondola. The gondola is oriented such that the radiator plate, which is painted white for maximum emissivity, points away from the Sun during daytime hours. This ensures that the radiator plate sees the coldest air temperatures around the gondola and can efficiently radiate away the heat delivered to it by the fluid which circulated the cryocooler. A schematic of this fluid flow in the LCS is shown in Figure 3.10.




Figure 3.10: A schematic of fluid flow in the LCS. 3M Novec 7200 fluid is held in a reservoir and pumped with one of two Fluido-Tech FG209 magnetic drive pumps through Tygon rubber hoses towards the cryocooler. The fluid delivers heat rejected by the cryocooler to the copper radiator plate, which radiates the heat away into the cold atmosphere.

3.4.3.2 Thermal environment during the COSI 2016 flight

Recalling the ultra-long duration balloon environment, COSI must maintain sufficient thermal control to operate the cryocooler at safe temperatures for 100+ days. At high altitudes, convection cooling cannot prevent the cryocooler from overheating and the LCS is required to transfer away heat. It is also important, however, to house the cryostat, cryocooler, LCS, and gondola electronics as a whole in a thermally stable shelter within which these systems can function properly. For this reason, COSI insulates the electronics bay of the gondola with aluminum side panels, foam panels, and aluminized mylar. Although not part of the LCS itself, this insulation has a substantial impact on the thermal environment. Figure 3.11 shows that during the 2016 flight, the insulation kept the interior of the high bay between $\sim 20\text{--}30^\circ\text{C}$ while the outside temperature changed by 90°C in a day-night cycle.

Figure 3.12a demonstrates the thermal control provided by the LCS. The radiator plate, which remains exposed to the atmosphere unlike the insulated components of the LCS, shows a 50°C change in ambient temperature outside the gondola. The cryocooler collar remains within a much smaller range of $30\text{--}40^\circ\text{C}$. Figure 3.12b highlights the oscillations in cryocooler collar temperature. These oscillations were purposely induced by the COSI team. The LCS performed better than expected during the flight, cooling the cryocooler to temperatures lower than originally planned. Out of a desire to maintain broader thermal stability, the team turned off the LCS when the collar temperature reached 30°C and back on when it reached 40°C . The oscillation period decreases with decreasing external temperature because lower outside temperatures allowed the collar to cool to its minimum 30°C more quickly,




Figure 3.11: A 24-hour portion of the 2016 flight demonstrating the dramatic change in exterior temperature and the relatively stable temperature inside the electronics bay. The temperature was stabilized by aluminum side panels, foam panels, and aluminized mylar insulating the gondola frame.

prompting the team to turn off the LCS with greater frequency. This on-off procedure was automated in the flight computer software.

Figure 3.13 shows that the ultimate goal of the LCS was achieved: through maintaining a healthy cryocooler collar temperature, the cryocooler was able to keep its cold tip at the specified 77 K in constant temperature mode and the resulting detector temperature was 83–84 K. The power expended remained at approximately 100 W and the visible oscillations follow those of the cryocooler collar temperature. As the collar temperature increased (decreased), the cryocooler became less (more) efficient and increased (decreased) its power. Thus, the cryocooler, LCS, and physical insulation on the outside of the gondola worked in concert to ensure proper thermal conditions during the 2016 flight.

3.4.4 Shields

COSI rejects background events through passive and active means. Six anti-coincidence cesium iodide (CsI) shields surrounding the four sides and bottom of the cryostat define the $\sim 1\pi$ sr field of view. Each shield module is comprised of a $40 \times 20 \times 4$ cm³ block of CsI. Scintillation light is collected with two photomultiplier tubes (PMTs) per shield. The shields suppress the Earth albedo radiation by passively absorbing γ -rays incident from below the instrument. This suppression is critical at balloon altitudes which are especially vulnerable to Earth’s albedo. Events of energy greater than the shields’ ~ 80 keV threshold are also actively vetoed if the shields trigger within 0.7–1.1 s of a trigger in the GeDs (Sleator 2019). This “anti-coincidence” shielding reduces background by limiting the instantaneous field of view to 25% of the sky and rejecting events which are incompletely absorbed in the detector

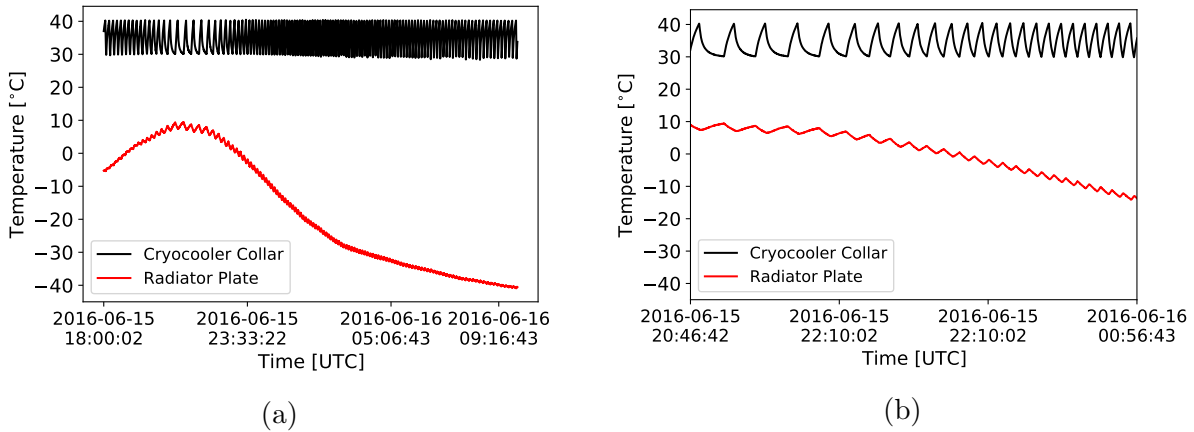


Figure 3.12: Cryocooler collar temperature during the 2016 flight as compared to the external temperature seen by the radiator plate. The stability of the collar temperature, despite changing outside temperatures, is evidence of a functioning LCS. The image in (b) is a time slice of that in (a) to show the deliberate, automated oscillation in collar temperature between the minimum 30°C and maximum 40°C set by the COSI team during flight.

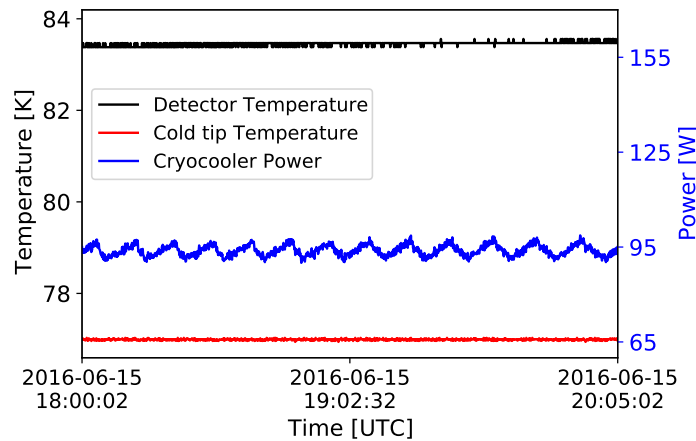


Figure 3.13: Proper detector temperature, cold tip temperature, and cryocooler power during the 2016 flight. The LCS ensured that the cryocooler could work efficiently to maintain its cold tip temperature of 77 K and resulting detector temperature of 83–84 K. The oscillations in power follow the oscillations in collar temperature (Figure 3.12b) because of the temperature dependence of cryocooler efficiency.

volume. Incompletely absorbed events do not deposit their full energy in the GeDs and are rejected because they cannot be correctly reconstructed (see the denominator of the second term in Equation 3.2, which is the sum of all deposited energy).

Another form of passive shielding is the graded-Z shield on top of the cryostat. The graded-Z shield is designed to attenuate incident photons of energy below COSI’s 0.2–5 MeV bandpass. The upper layer (furthest from the cryostat) is 0.22 mm of tin and the lower layer is 0.44 mm of copper. Tin, which has a higher atomic number than copper, induces photoabsorption of low-energy photons and blocks them from reaching the cryostat. X-ray fluorescence from interactions in the tin are absorbed by the copper layer below. The aluminum lid of the cryostat itself, ~ 0.75 mm thick, functions as a third layer that absorbs any residual emission (Kierans 2018).

3.4.5 Electronics

All of COSI’s strip electrodes are read out individually by analog electronics. The signals from these 888 strips (12 detectors \times 2 sides per detector \times 37 strips per side) pass through the cryostat walls via Kapton-Manganin flex circuits. The flex circuits feed the signals to charge-sensitive preamplifiers. There are two preamplifier boards per detector (one for each side). The preamplifiers are housed in preamplifier boxes that are visible on the sides of the cryostat in Figure 3.6b. Each box contains four preamplifier boards. One preamplifier board consumes ~ 0.5 W of power.

The preamplifiers output an amplified voltage pulse which is subsequently shaped by circuitry in analog boards. The analog boards, eight per detector, are housed inside “card cages.” The twelve card cages (one per detector) hold the analog boards corresponding to their respective detectors and several additional boards: a high-voltage power supply biases the GeD to its appropriate voltage of either 1000, 1200, or 1500 V, a low-voltage power supply powers the preamplifier boxes and the rail voltages of the analog boards, and a DSP board executes the trigger logic and communicates with the flight computer. The card cages are mounted on the balloon gondola individually and are electrically isolated from each other so that damage or a short, for example, in one does not compromise another.

Separate channels in the pulse-shaping analog circuits define the time and pulse height of an interaction. The trigger sequence proceeds as follows: A trigger is initiated when the fast shaper of any strip crosses the fast threshold of ~ 40 –50 keV. This 200 MHz fast bipolar shaper defines the time of an interaction as the zero-crossing time of the shaped pulse, i.e. 5 ns timing resolution. If the fast shaper on the opposite side of the detector occurs within 360 ns of the first trigger, the sequence proceeds. This 360 ns coincidence window mitigates the influence of false noise triggers on only one side of the detector. The readout then waits $2.4 \mu\text{s}$ for veto signals from the CsI anti-coincidence shields or the guard ring. If the CsI shields or detector guard rings trigger within this $2.4 \mu\text{s}$ period, the event is rejected; a coincident trigger in the shield may indicate an incompletely absorbed event, and a trigger in the guard ring indicates that the event occurred too close to the edge of the detector for accurate reconstruction. If no veto arrives, the trigger sequence concludes by measuring the

pulse heights of the slow shaper channels. The pulse-height channel is a unipolar shaper with $6\ \mu\text{s}$ rise time and an energy threshold of $\sim 20\ \text{keV}$. The longer rise time in the pulse-height channel ensures complete collection of charge carriers generated in an interaction. The pulse height is time-stamped at 50% of its maximum amplitude. The timing and pulse-height information is relayed to the COSI flight computer via Ethernet.

A 10 MHz clock from the flight computer is shared between the 12 card cages. This clock defines a 100 ns coincidence window for hits across all detectors. Hits within the 100 ns window are grouped together as an “event.” The flight computer also controls all gondola subsystems (e.g., card cage power, GeD high voltage, strip thresholds, cryocooler temperature, etc.). During flight, COSI’s telemetry system sends commands to gondola subsystems through the flight computer. When the balloon gondola is within the line of sight (LOS) of the launch location, typically in the first few hours of flight, communications are transmitted over two radio transmitters with combined 1 Mbps downlink speeds. This is the fastest line of communication. Telemetry during the rest of the flight, after the balloon has left the LOS, uses Iridium OpenPort (OP) technology. The flight computer communicates with the two OP transceivers over Ethernet and relays data to the ground via an internet connection. The telemetry system does not telemeter all data to the ground during flight. To keep the size of the data stream within the capabilities of the telemetry system, the flight computer parses raw data in real time to distinguish Compton events from single-site events. Compton events contain at least two interactions (one Compton scatter and a final photoabsorption). These are prioritized and transmitted to the ground in real time. Single-site events, while informative in offline analyses of overall detector rates and studies of lower energy interactions, for example, are of lower priority. These events and all other raw data, including the telemetered Compton events, are stored on the gondola’s three 1 TB hard drives. In the 2016 flight, the gondola landed safely and all data stored on the hard drives were recovered.

3.4.6 Aspect system

During the 2016 balloon flight, a Magellan ADU5 differential Global Positioning System (dGPS) communicated the position and orientation of the gondola. The position is recorded in terms of Galactic latitude, longitude, and balloon altitude. The orientation of the gondola is defined by heading, pitch, and roll angles. These angles define the rotation of the instrument about its coordinate axes. A Trimble BX982 GPS was installed as a backup to the dGPS, along with a Model 544 magnetometer from Applied Physics Systems.

3.4.7 NASA Balloon Program

Scientific ballooning is a desirable method of astrophysical exploration because of its flexibility in launch location, decreased cost and preparation time relative to ambitious satellite missions, and ability to conduct near-space experiments that can be safely recovered and deployed repeatedly. The NASA Balloon Program has facilitated many such scientific

and technological efforts over the last several decades from launch locations spanning the Columbia Scientific Balloon Facility in Palestine, Texas, the Scientific Balloon Flight Facility in Fort Sumner, New Mexico, McMurdo Station in Antarctica, Esrange Space Center in Kiruna, Sweden, and the Wanaka Airport in Wanaka, New Zealand. Once a safe landing location is identified, the flights are terminated with a radio signal that tears the balloon, separating it from the scientific payload and freeing both to drift to the ground for recovery.

There are three types of balloons that each fulfill specific flight requirements. Conventional balloons support missions lasting 2 hours to 3 days and long-duration balloons can support flights lasting ~ 55 days. Both conventional and long-duration balloons are filled with helium, made of polyethylene film approximately 0.02 mm thick, and can suspend payloads weighing up to about 6000 pounds (~ 2700 kg) at altitudes up to approximately 39 km. The bottoms of the balloons are open to and at equilibrium with the atmosphere and are accordingly called “zero pressure” balloons ([NASA CSBF](#); [NASA Scientific Balloons](#)).

The ultra long duration balloon (ULDB) is designed to fly at constant altitudes of up to 36 km for 100 days. This pumpkin-shaped “superpressure” balloon is closed to the atmosphere and floats at near-constant altitude by maintaining a slightly higher pressure than that of the surrounding environment. The balloon is designed not to drop in altitude with decreasing temperature each day-night cycle, as is the case with the “zero pressure” balloons, and can stay afloat for much longer. The helium-filled, 0.04 mm-thick polyethylene ULDB used by COSI in 2016 was 18.8 million ft^3 ($\sim 530,000 \text{ m}^3$) in volume and 100 m in diameter. It achieved a maximum float altitude of 34 km ([Kierans 2018](#); [NASA Scientific Balloons](#)).

The ability to maintain constant altitude is of great interest to experiments like COSI, whose analysis is simplified by minimal change in its altitude-dependent background rate. Furthermore, in the pursuit of scientific goals dominated by background, COSI benefits enormously from the increase in the amount of data collected over a 100-day mission. The 46-day COSI flight from Wanaka, New Zealand in 2016 was deemed a successful test of the ULDB technology, despite exhibiting altitude variations primarily in the latter half of the flight ([Figure 3.14](#)).

3.4.8 Flight heritage

The success of the 2016 mission drew upon instructive flight heritage dating back to COSI’s predecessor, the Nuclear Compton Telescope (NCT). NCT first demonstrated the astrophysical potential of CCTs through measurements of atmospheric and instrumental γ -ray background in a conventional balloon flight from Fort Sumner, New Mexico in 2005 ([Bowen et al. 2006](#)). The instrument had two GeDs and was key proof-of-concept for the subsequent iteration of NCT, which was launched from Fort Sumner on a conventional zero-pressure balloon in 2009 with 10 GeDs. The 2009 campaign detected the Crab Nebula and made history as the first CCT to detect an astrophysical source ([Bowen 2009](#)). Unfortunately, an attempted re-flight in 2010 from Alice Springs, Australia failed during launch. The gondola suffered severe damage but the electronics and detectors were recovered, allowing the NCT

team to build upon the previous years' success and develop a new instrument, called COSI, for flight on the ULDB platform.

The COSI team prepared the upgraded instrument for launch from McMurdo Station, Antarctica on December 28, 2014. The instrument cooled its 12 GeDs with a mechanical cryocooler rather than the expendable liquid nitrogen carried by the NCT payload, and was mounted in a lightweight aluminum gondola frame compatible with the ULDB. A leak in the balloon prematurely terminated the flight after ~ 40 hours. No astrophysical sources were detected over the short flight, which was nevertheless long enough for the team to diagnose problems with and improve thermal control of the cryocooler and the monitoring software (Sleator 2019). These improvements culminated in the 2016 iteration of COSI, discussed in the following section.

3.5 The COSI 2016 balloon flight

3.5.1 Flight overview

On May 17, 2016, COSI was launched as a science payload on a NASA ULDB from the Wanaka Airport in Wanaka, New Zealand. The launch site from New Zealand was chosen to maximize exposure of the Galactic Center, observations of which are important for COSI's science goals to measure nuclear lines and positron-electron annihilation. Launching from the Pacific-bound New Zealand also minimizes the safety risk of landing unintentionally over developed land. A complete summary of the 46-day COSI 2016 flight is provided in Kierans et al. (2016).

COSI floated freely beneath the balloon and was always pointed straight above towards zenith. There is no dedicated pointing to specific sources; it swept the sky through the Earth's rotation during flight. A rotator kept the instrument's solar panels oriented towards the Sun during daytime to charge the payload's batteries. These batteries powered the instrument and balloon electronics during nighttime, when the rotator was inactive and the instrument was free to rotate about its azimuth. Refer to Figure F.4 for a visualization of the changing instrument azimuth during nighttime compared to its relative stability during daytime.

Nine of COSI's twelve detectors operated continuously throughout the flight. Detectors 8 and 5 were turned off within the first 48 hours of the flight, on May 17, 2016 at $\sim 23:03$ UTC and May 18, 2016 at $\sim 18:35$ UTC, respectively. Detector 0 was turned off on June 6, 2016 at $\sim 16:45$ UTC. All three detectors were turned off due to high voltage (HV) malfunctions. A probable HV breakdown between detectors 8 and 5 and their HV supplies was identified from oscillating spikes in preamplifier currents above their nominal values. The problem was reproduced after the flight in vacuum tests conducted at the Space Sciences Laboratory in Berkeley, CA and attributed to faulty HV filters. The HV problem in detector 0, however, originated from a different electronic component: a resistor in the HV supply was operated outside of its current range and failed, sending the HV reading in detector 0 abruptly to

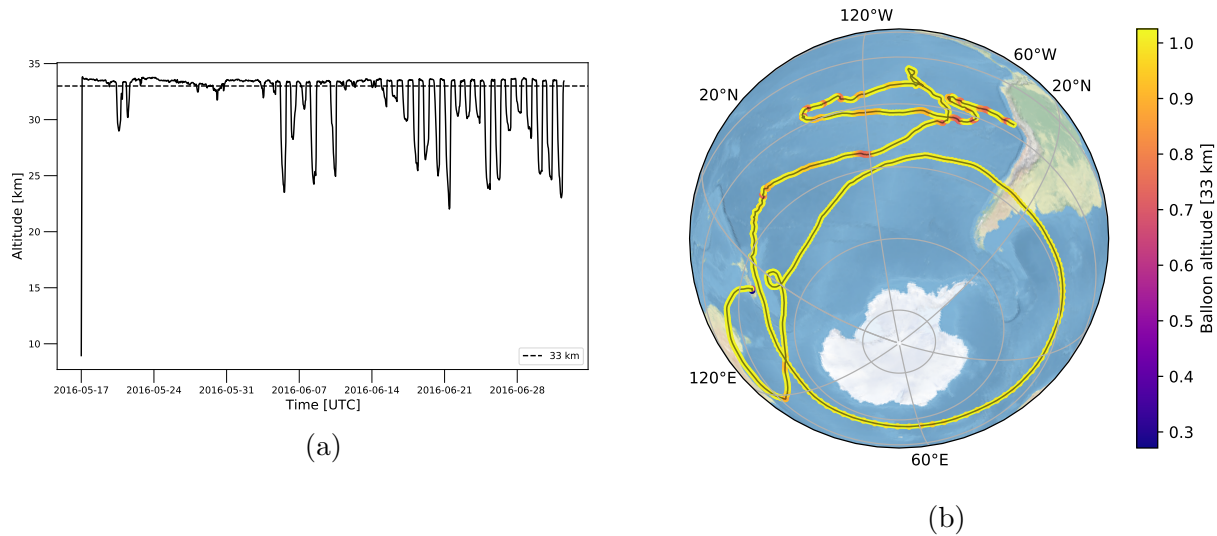


Figure 3.14: Balloon altitude and flight path of the COSI 2016 balloon flight. (a) The nominal flight altitude of 33 km is shown with a dashed line. Issues with the balloon in the second half of the flight resulted in altitude variations between 33 and 22 km with the day-night cycle. (b) COSI’s flight path as the balloon circumnavigated the globe. The color scale illustrates the drops in altitude.

zero. This problem was also reproduced in vacuum tests in Berkeley after the flight and the resistor was upgraded to accommodate higher currents (Sleator 2019).

The nominal flight altitude was 33 km, though the balloon experienced altitude variations between 33 and 22 km with the day-night cycle (Figure 3.14a). Remaining at high altitude is preferable because the strong background from Earth’s albedo and atmospheric absorption decreases with increasing altitude. Additionally, modeling the background at constant altitude simplifies the analysis. The altitude drops in the latter half of the flight are plausibly attributed to a small leak in the balloon, though the cause has not been definitively determined. Cold weather caused the smaller drops shown towards the beginning of the flight. The instrument circumnavigated the globe within the first 14 days of the flight and then remained largely above the South Pacific Ocean before the mission was safely terminated on July 2, 2016 (Figure 3.14b). The exposure map is shown in Figure 3.15. The instrument was recovered from its landing site in Peru with no signs of consequential damage.

3.5.2 Scientific achievements

A brief summary of scientific achievements in the 2016 flight is provided below. The measurement of Galactic ^{26}Al , the primary analysis of this dissertation, is described in detail in Chapter 5.




Figure 3.15: Exposure map of the 2016 balloon flight, assuming 20 cm^2 effective area at 356 keV. Image from Kierans et al. (2016).

3.5.2.1 Positron-electron annihilation

In the 1970s, a balloon-borne astrophysics mission detected a strong 511 keV signature of positron-electron annihilation emanating from the center of the Milky Way Galaxy (Johnson III et al. 1972; Leventhal et al. 1978). Decades later, the morphology of the emission and the origin of the positrons fueling the annihilation remain poorly constrained. The emission is characterized by two components: a narrow “disk” emission that extends along the Galactic Plane and an extended “bulge” emission that is concentrated around the Galactic Center (Figure 3.16). It is unclear whether these features are produced by the same annihilation mechanism or if they trace different sites of annihilation. The annihilation may occur at the sites of positron production or displaced from the positron sources, should the positrons propagate away from their production sites before annihilating. The nature of the source or collection of sources supplying the positrons is additionally uncertain. Massive stars, core-collapse supernovae, thermonuclear supernovae, microquasars, black holes, and dark matter are all potential positron sources which could contribute to the observed emission. Spectroscopic and imaging observations are required to further constrain the morphology and the “source, transport, and sink” of the annihilating positrons.

Doppler shifts in the 511 keV spectrum inform the kinematics of the positrons as they annihilate in the interstellar medium (ISM). High-resolution spectroscopy can also disentangle the para-Positronium (511 keV photopeak) and ortho-Positronium (≤ 511 keV continuum; Ore and Powell 1949) components of the spectrum, whose relative contributions can reveal the conditions (e.g., ISM temperature) in which the majority of the positrons are annihilating. For example, spectral analyses have indicated that the positrons likely propagate away from their hot (stellar) production sites and annihilate in warm (~ 8000 K) regions of the ISM (Jean et al. 2006). Additional high-resolution spectroscopic studies from SPI are detailed in Jean et al. (2006); Churazov et al. (2011); Weidenspointner et al. (2008); Siegert et al. (2016a); Siegert et al. (2019).

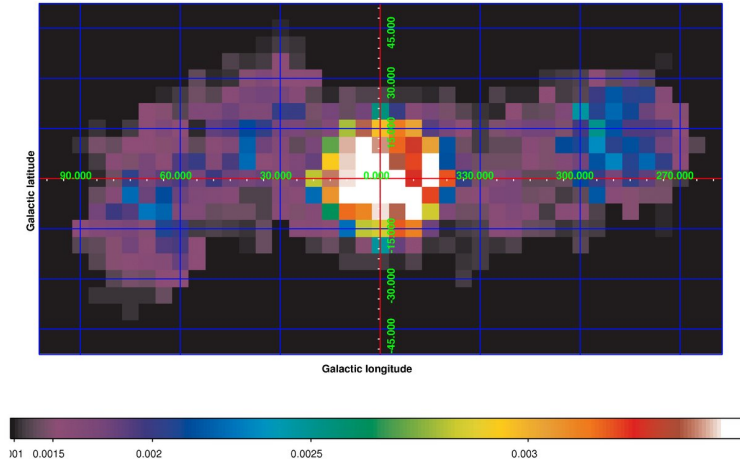


Figure 3.16: Image of 511 keV emission from positron-electron annihilation in the Milky Way Galaxy, as observed by SPI. The contrast between the bright “bulge” and extended “disk” emission is not understood. Image from Bouchet et al. (2010).

COSI detected the 511 keV line with $\sim 7\sigma$ significance in the 2016 flight (Figure 3.17a). COSI’s measured positronium fraction of 0.76 ± 0.12 is less than previous measurements of ~ 1 , while the integrated flux of $(3.9 \pm 0.4) \times 10^{-3} \text{ ph cm}^{-2} \text{ s}^{-1}$ exceeds that from SPI of $(2.74 \pm 0.03) \times 10^{-3} \text{ ph cm}^{-2} \text{ s}^{-1}$ (Siegert et al. 2016a). The centroid of the fitted Gaussian is $511.8 \pm 0.3 \text{ keV}$ and the width, having subtracted out the inherent broadening of the line from finite instrumental resolution, is $\sigma = 1.7 \pm 0.4 \text{ keV}$ (Kierans 2018; Kierans et al. 2020).

Imaging can reveal whether the emission comes from many individual point sources or if it is truly diffuse. SPI’s coded mask technology produced an image with fine angular resolution of $\sim 2.7^\circ$ (Figure 3.16; Bouchet et al. 2010; Skinner et al. 2014; Siegert et al. 2016b). COSI’s 511 keV image ($\sim 5^\circ$ resolution; Figure 3.17b) shows a central 511 keV bulge (FWHM $28_{-12}^{+19} \text{ }^\circ$; Siegert et al. 2020) that is 2–3 times larger than that reported by SPI (FWHM 8° ; Knödlseider et al. 2005). There is no observed disk emission, though this is expected given COSI’s limited 46 days of observation time. The disk has low surface brightness compared to the significantly brighter bulge component. COSI’s enhanced sensitivity to low-gradient diffuse emission, as compared to SPI’s coded mask design, could possibly account for its greater observed flux and broader bulge measurements. Systematic uncertainties in COSI’s analysis pipeline may also have an effect. Refer to Section 5.5.2 in the ^{26}Al analysis for further discussion about potential systematic uncertainties. The tension in these results motivates further study and highlights the importance of observations from COSI as a satellite mission, which is projected to have an angular resolution of $\sim 3.2^\circ$ and energy resolution of $\sim 0.4\%$ at 511 keV (Tomsick et al. 2019).

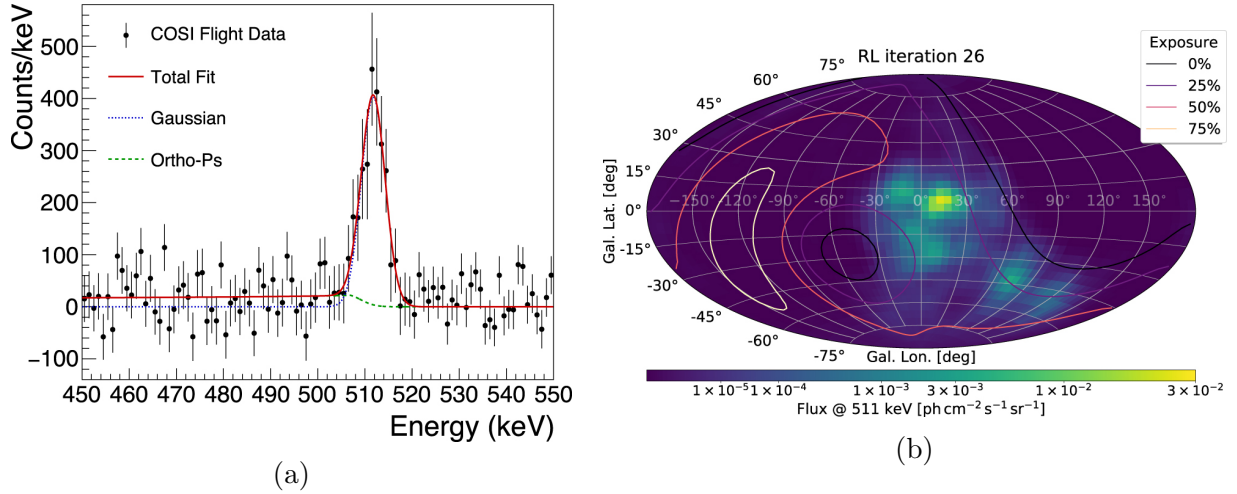


Figure 3.17: COSI observations of positron-electron annihilation during the 2016 balloon flight. (a) Spectrum with relative contributions of the para-Positronium (Gaussian) and ortho-Positronium (Ortho-Ps) components. The centroid of the fitted spectrum is 511.8 ± 0.3 keV and the width, subtracting out instrumental resolution, is $\sigma = 1.7 \pm 0.4$ keV. Image from Kierans (2018); Kierans et al. (2020). (b) Image of the emission. COSI detected the bright “bulge” emission, but not the faint “disk” emission seen by SPI. Image from Siebert et al. (2020).

3.5.2.2 GRB detection and polarization

Gamma-ray bursts (GRBs) are extremely energetic (10^{51} erg) explosions that occur across the universe. The initial “prompt” phase of these transient events emits γ -rays between 0.1–1 MeV. The subsequent “afterglow” emission emits more strongly across lower energies spanning X-rays to radio waves. Despite thousands of GRB observations, remarkable diversity in measurements of GRB spectra and timing presents conflicting evidence of dominant emission mechanisms and the geometry of outflowing material. A promising tool for disentangling these complications is polarization, a powerful probe of the magnetic environment of GRBs. For example, synchrotron emission from a coherent magnetic field that powers the prompt emission can generate highly polarized photons. Measuring the polarization and jet geometry of prompt emission in a large sample of GRBs can therefore discriminate between theoretical models of energy dissipation, including synchrotron in ordered magnetic fields (Toma et al. 2009), synchrotron in random magnetic fields (Medvedev and Loeb 1999), and inverse Compton scattering (Lazzati et al. 2004; Toma et al. 2009).

There are far fewer GRB polarization measurements (only a few dozen in the 0.1–1 MeV energy range; Chattopadhyay 2021; Kole et al. 2020) than there are GRB observations, which number in the thousands. The first GRB polarization measurement was made by RHESSI on GRB 021206. The high significance ($> 5.7\sigma$), high polarization level ($\Pi =$

80% \pm 20%) measurement (Coburn and Boggs 2003) remains controversial, as analyses by other researchers were unable to reproduce the result (Rutledge and Fox 2004; Wigger et al. 2004). It also stands in contrast to more recent measurements which have had a maximum measurement significance of $\approx 3.7\sigma$ (Gamma-ray Burst Polarimeter; Yonetoku et al. 2012) and several measurements which indicate no significant photon polarization. The POLAR collaboration detected 55 GRBs in 2016–2017 and reports unpolarized or mostly unpolarized emission in a time-integrated analysis between ~ 30 –750 keV of 14 events; moderate linear polarization is observed in a time-resolved analysis (Kole et al. 2020).

More GRB polarization measurements are necessary to realize the potential of this tool. COSI detected GRB 160530A during the 2016 flight. Subsequent analysis by the COSI team found a polarization level of $\Pi = 16_{-16}^{+27}\%$, below the minimum detectable polarization of $57.5 \pm 0.8\%$ (Lowell 2017; Lowell et al. 2017). The analysis technique developed for this measurement will be used in expected polarization measurements by COSI as a satellite mission, which will aim to add to a growing catalog of GRB polarization measurements across the astrophysics community.

3.5.2.3 Detection of the Crab Nebula

The Crab Nebula is the brightest persistent γ -ray source in the sky. Located inside the Milky Way’s Taurus constellation within the remnant of supernova 1054, this pulsar wind nebula surrounds the Crab Pulsar left behind by the supernova explosion. Pulsar wind nebulae are powered by strong magnetized winds that accelerate charged particles away from the pulsar to very high energies, emitting broadband synchrotron emission from radio waves to γ -rays (Gaensler and Slane 2006).

COSI detected the Crab Nebula during the 2016 flight, despite only 12 days of observation time in its vicinity; the Crab Nebula is more easily viewed from Earth’s Northern Hemisphere and COSI was launched from and largely remained in the Southern Hemisphere. Dominant background complicated a detailed spectral analysis and polarization measurement (Sleator 2019). Nevertheless, the detection is further proof of COSI’s ability to study known γ -ray sources in the MeV bandpass.

3.5.3 Attempted 2020 mission

The COSI team attempted to re-fly the instrument on NASA’s ULDB platform in the spring of 2020. Over the course of six weeks in January and February 2020, COSI was assembled, calibrated, and tested for compatibility with NASA’s balloon technology at the Columbia Scientific Balloon Facility (CSBF) in Palestine, Texas. Testing at CSBF also included a critical three-day thermal vacuum test of the cryostat, which confirmed that the repaired high voltage filters were functioning properly at flight pressures and temperatures for longer than their time to failure in the 2016 flight.

All equipment and essential personnel traveled to the Wanaka Airport in Wanaka, New Zealand immediately following activities at CSBF. The launch window was open from April

15, 2020 to June 6, 2020. The cryostat was pumped down to vacuum pressures and cooled to its required $\sim 83\text{K}$ during the first week of work and the instrument was assembled for calibration by March 5. After ten days of successful data-taking and smooth operations, the campaign was canceled due to the emerging COVID-19 pandemic. Numerous analyses of the calibration data taken before cancellation are detailed in Chapter 4 and serve as a valuable check of instrument performance. The methods described in that chapter will be used to calibrate and benchmark the performance of COSI as a Small Explorer satellite.

3.6 Analysis methods

3.6.1 MEGAlib

The Medium-Energy Gamma-ray Astronomy library (MEGAlib) is a software package designed for the calibration and analysis of hard X-ray and γ -ray instruments (Zoglauer et al. 2006). MEGAlib is written primarily in C++ and provides a comprehensive interface for specifying the geometry and materials of an instrument, calibrating real data, running simulations of calibration measurements and astrophysical sources, performing event reconstruction, benchmarking instrument capabilities, and producing images. The use of MEGAlib in COSI calibrations is detailed in Chapter 4 and much of the analysis in Chapters 5 and 6 is facilitated by the MEGAlib framework. Both real data and simulations are processed through this same analysis pipeline. A brief overview of the programs within MEGAlib is provided below:

Geomega (“Geometry for MEGAlib”) defines the physical properties of the detector, including specific information about its geometry, constituent materials, trigger criteria, and more. A detailed and accurate “mass model” is important for subsequent stages of the analysis.

Cosima (“Cosmic simulator”) generates Monte Carlo simulations of photon interactions inside Geomega’s instrument mass model. Cosima is based on Geant4 (Allison et al. 2006; Allison et al. 2016; Agostinelli et al. 2003).

Detector Effects Engine The “DEE” converts simulation data, originally free of imperfections intrinsic to instrumentation and the measurement process, into events which mimic those measured by the instrument. Thus, the simulated events are converted from physical units to electronic units. Applying measurement effects to the simulations allows for meaningful comparisons of instrument performance to expectations. A detailed review of the development and validation of COSI’s DEE is provided in Sleator et al. (2019).

Nuclearizer calibrates real instrument data by converting measured parameters from electronic (pulse height, strip number) to physical units (energy, position). It calibrates

simulations by converting the DEE output back to physical units of position and energy, ensuring that simulations share the imperfections seen in real data. Nuclearizer then applies corrections to both to mitigate these imperfections.

Revan (“Real event analyzer”) combines individual hits in real data and simulations into events. It performs Compton reconstruction to recover the initial energy and direction of each incident photon.

Mimrec (“MEGALib image reconstruction”) performs high-level analysis of the events reconstructed by Revan. It can generate energy spectra and light curves, define energy and angular resolution, inform polarization analysis, and reconstruct images using an iterative list-mode-likelihood method. A wide variety of event selections are applied at this stage, the tuning of which is a key component of COSI analysis.

3.6.2 `cosipy` and COSI Data Challenges

In recognition of Python as an increasingly popular programming language in the astrophysics community, Dr. Thomas Siegert of the COSI team began to develop a Python-based COSI analysis toolkit called “`cosipy`” in 2019. After calibrating and processing raw data with MEGALib, the final event lists are passed to the `cosipy` library to study basic properties of the data, including energy spectra, light curves, instrument pointing, sky exposure, and more. Advanced tasks like defining background models and performing Richardson-Lucy image deconvolution are also possible in `cosipy`, as demonstrated in COSI’s 511 keV imaging analysis (Figure 3.17b; Siegert et al. 2020). MEGALib and `cosipy` together comprise a broader COSItools software package².

The COSI team began releasing `cosipy` to the public in 2023 in the form of “data challenges.” In the first Data Challenge, Jupyter notebook tutorials of COSI balloon analysis were published on GitHub² to introduce members of the broader astrophysics community to spectroscopy and imaging with Compton telescopes. Chapter 6 details much of the work conducted by the author of this dissertation to test, debug, and develop `cosipy`’s imaging capabilities for release in Data Challenge 1. Future Data Challenges will focus on the COSI satellite, rather than the balloon instrument, with advanced tools to study simulated sources spanning all of COSI’s main science objectives. By the time COSI is launched, `cosipy` will be a mature and sophisticated analysis framework capable of rapidly analyzing satellite data. Efforts to coordinate `cosipy` development in tandem with existing γ -ray analysis frameworks, including the Multi-Mission Maximum Likelihood framework (3ML; Vianello et al. 2015), are underway.

²COSItools GitHub link: <https://github.com/cositools>; Data Challenge 1 GitHub link: <https://github.com/cositools/cosi-data-challenge-1>

Chapter 4

Calibrations of the Compton Spectrometer and Imager

Unlocking the scientific potential of the MeV gap is possible only with sensitive instruments. COSI, described in Chapter 3, is one such instrument with notable advantages over previous γ -ray observatories in the MeV bandpass: COSI has an energy resolution $\sim 10\times$ better than that of COMPTEL, it is more sensitive to low-gradient diffuse emission than coded mask instruments like SPI, and it is a surveying instrument with an instantaneous field of view spanning $\sim 25\%$ of the sky at any given time. This wide field of view matches that of COMPTEL ($\sim 1\pi$ sr) and exceeds that of SPI ($\sim 16^\circ$). The combination of fine energy resolution, sensitivity to diffuse emission, and wide field of view is critical to studying transient and persistent sources across the Galaxy, positron-electron annihilation, and diffuse emission from ^{26}Al in the well-constrained Inner Galaxy and at poorly constrained high Galactic latitudes.

Analyses such as those presented in Chapter 5 of this dissertation hinge on instrument calibrations. Calibrations translate raw data recorded in electronic units to units of the physical world. The conversion from pulse height amplitude to energy must be empirically determined. The positions of the interactions in the detectors are labeled by the triggered strip electrodes rather than Cartesian coordinates (x - y - z) in units of centimeters. Furthermore, the z -coordinate of an interaction is not directly stored in the electronics at all and can only be determined through instrument calibrations.

Thus, robust calibrations permit the accurate event reconstruction needed to fully exploit COSI's high-resolution spectroscopy, direct imaging capabilities, polarization sensitivity, and intrinsic suppression of background events. They are fundamental to the interpretation of raw data as scientific data. This chapter¹ provides a detailed overview of the methods used to calibrate the instrument prior to the 2016 and attempted 2020 balloon campaigns. Expected consistency is observed between the calibration and benchmarking results from

¹This chapter largely follows the published paper “Calibrations of the Compton Spectrometer and Imager” by Bechert et al. (2022a). ©Nuclear Instruments and Methods in Physics Research Section A 2022, reprinted with permission.

these measurements, all of which were taken at the launch site in Wanaka, New Zealand. The consistency between 2016 and 2020 calibrations confirms that the instrument remained stable over time and that the calibration procedures are well-documented, achievable on the timescale of a balloon campaign, and yield reproducible results when carried out by different scientists (the author of this dissertation co-coordinated the 2020 campaign effort).

These procedures were developed in the years preceding both campaigns at the Space Sciences Laboratory (SSL) in Berkeley, CA. Refining the methods on laboratory measurements is essential to executing them properly before launch. They will also be integral to the calibration and benchmarking of the COSI satellite mission. Other experiments that use double-sided cross-strip high-purity germanium detectors, including the Gamma-ray Imager/Polarimeter for Solar Flares (GRIPS) balloon instrument (Duncan et al. 2016), which is scheduled to fly for a second time in 2024, can also benefit from the calibration techniques described in this chapter.

4.1 Calibration pipeline

Calibrations are performed within Nuclearizer (Section 3.6.1), which converts electronic parameters (either directly from measured data or from simulations processed by the DEE) into physical parameters. This chapter pertains to the calibration of real data from the 2016 and 2020 balloon campaigns. As the fidelity of these calibrations governs the reliable interpretation of data, the highest possible accuracy in the following steps is critical to achieving COSI’s scientific goals:

1. **Load data** For each triggered strip, the strip ID, detector ID, pulse height, and timing are recorded. One event can contain several triggered strips over multiple detectors.
2. **Energy calibration** The pulse height of an interaction, recorded in the units of the analog-to-digital converter (ADC), is converted to a deposited energy in keV (Section 4.3).
3. **Cross-talk correction** Recorded energies are enhanced by interactions on neighboring strips. This “cross-talk” enhancement scales linearly with energy and is removed with a linear correction (Section 4.4).
4. **Strip pairing** The intersection of triggered strips on opposing sides of a detector (Section 3.4.1) defines the x - y position of an interaction. If more than one interaction occurs inside the detector and triggers multiple strips per side, the strip pairing algorithm determines the most likely interaction position given the energies deposited on all triggered strips (Section 4.5).
5. **Depth calibration** While strip pairing finds the x - y position of an interaction, the depth calibration finds the z -position by converting the difference between electron and

hole collection times (the “collection time difference” or “CTD”) into an interaction depth in physical units (Section 4.6).

6. **Save calibrated data** Each event is saved with the energies and positions of its constituent interactions across multiple detectors.

4.2 Data collection

Calibration data are collected with Isotrak Eckert & Ziegler Type D sealed γ -ray sources (^{241}Am , ^{57}Co , ^{133}Ba , ^{22}Na , ^{137}Cs , ^{88}Y , ^{60}Co) that emit fifteen nuclear lines within COSI’s energy range (Table 4.1). These point-like sources are mounted in a variety of positions surrounding the instrument in order to illuminate the entire field of view. Data are collected in three configurations: low-energy (LE; < 511 keV), high-energy (HE; ≥ 511 keV), and polarized radiation data collection. This chapter will focus on calibrations using LE and HE data. For a detailed explanation of polarized data collection and the process of benchmarking COSI’s polarization response, refer to Section 4.8.4.

Source	Line energy [keV] (BR)
^{241}Am	59.5 (35.9%)
^{57}Co	122.1 (85.6%), 136.5 (10.7%)
^{133}Ba	81.0 (34.1%), 276.4 (7.1%), 302.9 (18.3%), 356.0 (62.1%), 383.85 (8.9%)
^{22}Na	511.0 (180.7%), 1274.5 (99.9%)
^{137}Cs	661.7 (85.1%)
^{88}Y	898.0 (93.7%), 1836.0 (99.2%)
^{60}Co	1173.2 (99.97%), 1332.5 (99.99%)

Table 4.1: The seven radioactive isotopes used to calibrate COSI, each listed with their γ -ray line energies and respective branching ratios (BR).

4.2.1 Low-energy configuration

Photons from low-energy sources (^{241}Am , ^{57}Co , and ^{133}Ba) are quickly attenuated and are unable to penetrate the full depth of a single GeD. As such, the low-energy sources are placed in numerous positions in the immediate vicinity of the cryostat until all strips on the twelve detectors are sufficiently exposed. Moving the source to ten or more positions for collection times ranging from ten minutes to several hours in each is often necessary. The required collection time depends on the activity of the source, the branching ratio of the decay, and the physical accessibility of the strips in question. Given that source activities varied between 2016 and 2020, the precise positions and integration times for the low-energy sources changed

between campaigns. The procedure of moving the sources around the cryostat, however, remained the same.

All three low-energy sources are used in the energy calibration to fit the lower end of COSI's energy range. The temperature correction in 2016 also required low-energy data from ^{241}Am (and high-energy data from ^{137}Cs ; see Section 4.3.2). The cross-talk corrections in 2016 and 2020 used a range of low-energy sources together with the high-energy sources described in the next section.

4.2.2 High-energy configuration

In 2016, an aluminum calibration structure was designed to hold the high-energy radioactive sources at fixed positions around the cryostat. The calibration structure is described in Kierans (2018) and is shown in Figure 4.1a. The HE energy calibration data were obtained by placing the HE sources at a height of ~ 63 cm above the center of the detector stack for a minimum of five hours per source. The ^{137}Cs data collected in this position were also used for the 2016 temperature correction and depth calibration. To study the angular resolution and effective area as a function of energy across COSI's entire field of view, the HE sources were held in numerous positions on the calibration structure. These calibration runs spanned several hours each in order to collect data with ample statistics. For example, sources placed towards the edge of the field of view, i.e. significantly off-axis from the zenith of the instrument, required longer calibration runs.

Data for the 2020 HE energy calibration, temperature correction, cross-talk correction, depth calibration, angular resolution, and effective area studies were collected with a new plywood calibration structure. Built at the SSL machine shop, the structure was designed to secure sources across COSI's field of view in reproducible positions that could be mimicked in simulations. Figure 4.1b shows the structure attached to the top of the gondola with four bolts on each corner, one of which is visible between the "X" marks in the bottom left corner. It is only affixed to the top of the gondola when collecting calibration data and is mounted onto and lifted off the gondola by hand.

A Delrin plastic source holder secures radioactive sources to the protruding arch at any polar angle spanning 0° at zenith to approximately 60° on either side. The radius of the arch is ~ 63 cm and the arch rotates freely in the azimuthal direction. With these polar and azimuthal degrees of freedom, data collected with the structure are used to characterize COSI's performance over the entire field of view.

In the 2020 HE energy calibration, the COSI team placed ^{137}Cs , ^{60}Co , and ^{88}Y at zenith but closer to the cryostat (~ 14.6 cm, ~ 14.6 cm, and resting atop, respectively) for expedited data collection. The ^{22}Na source was mounted at the zenith of the wooden calibration structure. These data runs lasted at least 7 and up to 24 hours.

The temperature correction in 2020 differed from that in 2016 (Section 4.3.2) and used ^{22}Na data collected from the zenith of the calibration structure. ^{137}Cs data taken from the zenith of the calibration structure were used in the depth calibration. Angular resolution and effective area data were also collected using the calibration structure, but limited calibration

time in 2020 restricted measurements to ^{60}Co , ^{137}Cs , and ^{22}Na at the zenith of the calibration structure. The campaign was terminated before the team could take measurements across the entire field of view. As in 2016, these data runs spanned several hours.




Figure 4.1: COSI calibration structures in 2016 and 2020. (a) The structure used in 2016 calibrations. The radioactive source is affixed to the structure above the cryostat. Image courtesy of the COSI team in 2016. (b) The custom-built calibration structure used in 2020. Radioactive sources are secured to the arch above the cryostat. Both structures are mounted to the top of the gondola.

4.3 Energy calibration

Each of COSI’s strips is read out individually by the data acquisition system and uniquely calibrated. These 888 calibrations (recall: 12 detectors \times 37 strips per side \times 2 sides per detector) are determined by illuminating the strips with radioactive sources of known γ -ray line energies (Table 4.1) and measuring the resulting pulse height signal. Fitting a polynomial to the relationship between pulse height and energy from multiple γ -ray lines yields the desired conversion from electronic to physical units for each strip. The energy calibration data also define the single-strip spectral resolution of COSI. This section details the energy calibration procedure and results from data taken prior to the 2016 launch and data taken prior to the intended 2020 launch.

Energy calibration is performed in “MEGAlib’s line calibrator,” or “Melinator.” The summed, raw spectra of all collected data are loaded into Melinator such that the total fifteen photopeaks across COSI’s energy range are simultaneously identified (see Figure 4.2 for an example in the 2020 energy calibration). To reduce computational strain, one detector is calibrated at a time. Melinator fits the photopeaks seen on each strip in ADC space (the

pulse heights) with a Gaussian (convolved with a delta function for energy loss) and a linear background model. The ADC centroid of this fitted peak is matched with the corresponding true photopeak energy in keV. Melinator then fits the relationship between energy and pulse height with a user-specified function. A third-order polynomial is used to account for non-linearities at low energies (see Section 4.3.1 for a discussion about choosing which polynomial to use). The parameters of the function, returned for each strip individually, define the conversion from ADC to energy. Melinator also returns the full-width half maximum (FWHM) of each peak in keV. This FWHM is the primary metric of COSI’s single-strip, energy-dependent spectral resolution.




Figure 4.2: The Melinator window showing the 2020 energy calibration of strip 8 on the DC side of detector 11. Melinator identifies the known photopeaks of all isotopes in each strip’s cumulative spectrum (top center plot). A third-order polynomial is fit to the plot of energy [keV] vs. read-out units [ADC] (bottom right plot). The ¹³⁷Cs 661.7 keV photopeak on this strip is identified at 2921.9 read-out units (bottom left plot).

Repeated calibrations before each campaign are necessary to account for slight variations in the instrument’s configuration with time. Routine detector maintenance, including electronic repairs and gain adjustments to the read out system, for example, changes the energy-ADC relationship on affected strips. Thus, minor differences in the energy calibrations between years are expected and motivate repeated calibrations before each campaign.

Accordingly, COSI’s single-strip spectral resolution as calculated with the FWHMs from Melinator in 2016 and 2020 are compared. The resolution is defined as the ratio of the FWHM of the identified ^{137}Cs γ -ray line to its known photopeak energy of 661.7 keV. By this definition, the single-strip spectral resolution of COSI in 2016 was $0.453 \pm 0.004\%$ on the AC side and $0.45 \pm 0.01\%$ on the DC side. In 2020, it was $0.52 \pm 0.01\%$ on the AC side and $0.48 \pm 0.01\%$ on the DC side (Table 4.2).

Side, energy	2020 [%]	2020 [keV]	2016 [%]	2016 [keV]
AC, 661.7 keV	0.52 ± 0.01	3.42 ± 0.06	0.453 ± 0.004	3.00 ± 0.03
DC, 661.7 keV	0.48 ± 0.01	3.17 ± 0.04	0.45 ± 0.01	3.00 ± 0.05

Table 4.2: Single-strip energy resolution (FWHM) in 2020 and 2016 calibration data, averaged over all strips on the AC and DC sides of COSI’s 12 GeDs. The resolution is defined as the ratio of the FWHM of the measured ^{137}Cs photopeak to 661.7 keV.

Electronic noise is expected to dominate the single-strip energy resolution of COSI up to about 1 MeV. This effect is visible in measured resolution as a function of photon energy E . Electronic noise scales as N^{-1} , where N is the number of charge carriers generated in an interaction. The energy required to produce an electron-hole pair in germanium is $W = 2.96$ eV, such that $N = E/W$. Thus, single-strip energy resolution is expected to scale as $\propto E^{-1}$. Figure 4.3 shows a best-fit power law relationship $\propto E^{-0.95}$ and $E^{-0.96}$ in 2016 and 2020 data, matching expectations derived from electronic noise.

4.3.1 Linearity investigations

Though the energy calibration is predominately linear with small coefficients on higher-order terms, non-linear behavior is observed across the energy range. Figure 4.4 shows the deviation from linearity in three arbitrary strips in the 2020 energy calibration. The centroid of a photopeak is identified in ADC and converted to energy with a best-fit linear model. The deviation is calculated as the difference between the linear model’s prediction and the true photopeak energy in keV. The percent deviation is the quotient of this difference and the true photopeak energy. The lower end of the range, in particular, exhibits a noticeable percent deviation from linearity.

Testing a variety of functions for the energy vs. ADC relationship is necessary to ensure that the model captures this non-linear behavior. The Aikake Information Criterion (AIC; Akaike 1974) is a statistical metric used to evaluate the balance between each model’s agreement with calibration data and the model’s predictive power at interpolated energies. This is important because overfitting the data with a higher-order function than required may closely match calibration points but inaccurately predict energies outside of the known photopeaks (e.g., when interpolating at energies higher than the 1.836 MeV peak of ^{88}Y).

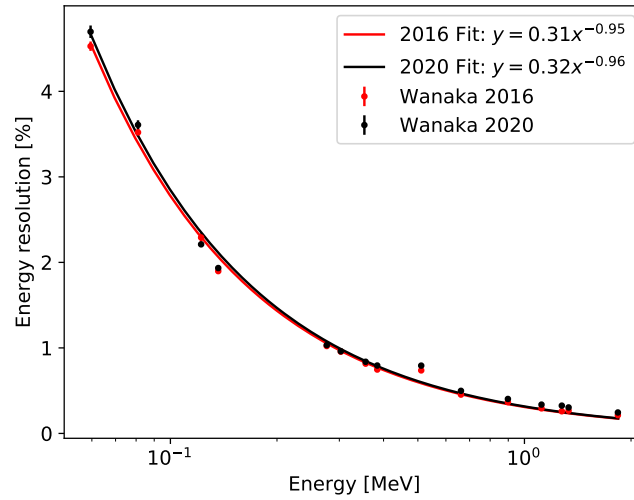


Figure 4.3: Single-strip energy resolution (FWHM) as a function of energy in 2020 and 2016 energy calibrations. The relationship between resolution and energy is consistent with expectations of dominant electronic noise.




Figure 4.4: Deviation from linearity in the 2020 energy calibration of three COSI strip electrodes (detector 0, strip number, detector side n (DC) or p (AC)).

As an example, a Markov-Chain Monte Carlo (MCMC) method was used to find the best-fit parameters of linear, second-order, third-order, and fourth-order polynomials to the measured energy vs. ADC data on one strip of detector 0 in 2020. The deviations from linearity in units of standard deviation σ (difference between true energy and median model prediction, divided by the median model prediction) are shown in Figure 4.5. The deviation from linearity at low energies is again noticeable when fitting the energy vs. ADC points with a linear model (top left).

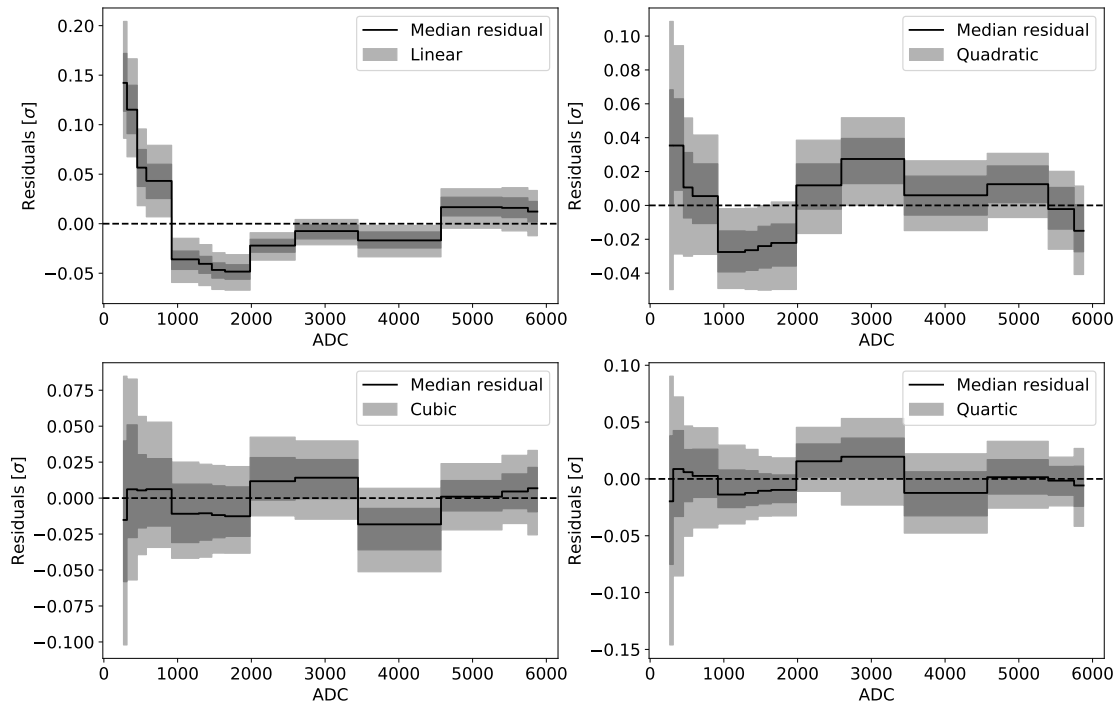


Figure 4.5: Residuals of MCMC linear, second-order, third-order, and fourth-order polynomial fits to 2020 energy vs. ADC calibration data on strip 1 of the DC side of detector 0. The shaded contours indicate 1σ and 2σ intervals.

Fitting with more parameters will always increase the likelihood of the fit. Figure 4.6a shows that the likelihood is indeed maximized in the quartic model (five fitted parameters, including the ordinate; the calibration does not assume 0 keV at 0 ADC). Figure 4.6b shows the change in AIC from the model with the minimum AIC. The change is evaluated because AIC values are only meaningful relative to each other. The third-order polynomial (four fitted parameters) has the minimum AIC and is therefore the best model for this strip. This was the chosen model for 2020 calibrations. Both third- and fourth-order polynomials were used when calibrating COSI 2016 data, though an analogous statistical study of overfitting was not conducted at that time.




Figure 4.6: Results of MCMC polynomial fits to 2020 energy vs. ADC calibration data. (a) The fourth-order model (five fitted parameters) maximizes the likelihood of the fit. (b) The AIC test indicates that the third-order model (four fitted parameters) gives the best results without overfitting.

4.3.2 Temperature correction

It was discovered during preparation for the 2016 flight that COSI's preamplifier boards are sensitive to environmental temperature changes. Temperature sensors placed near the preamplifiers and around the gondola monitor the temperature during calibrations and during flight. The temperature dependence shifts recorded pulse height spectra such that the identified photopeak energies are displaced from their true values. Shifts of up to $0.5 \text{ keV}/^\circ\text{C}$ at 661.7 keV were observed in 2016 calibration data (Kierans 2018). It is necessary to correct this shift before proceeding with subsequent calibrations and benchmarking because these steps require accurate energy determination. Additionally, proper calibration of flight data relies upon this correction because the balloon gondola experiences drastic changes in temperature with the day-night cycle.

To correct the 2016 spectra, ^{137}Cs and ^{241}Am energy calibration data were collected over a wide range of temperatures, approximately 12°C to 34°C , meant to mimic temperatures seen during flight. The energy-ADC relationship for each isotope was determined for five periods of constant temperature and the resulting linear relationship between preamplifier temperature and ADC peak location was applied as a simple offset correction to the measured ADC. The linear relationship was determined for each strip, yielding a precise correction tailored to each strip's individual readout. Applying the correction reduced the average difference between measured and true line energy from 0.5% to 0.1% (Kierans 2018).

The 2020 mission was canceled before the COSI team could take calibration data over the wide temperature range required for the 2016 temperature correction method. Temperatures in 2020 data only range from $\sim 27^\circ\text{C}$ to 34°C and an alternative method of correcting for the temperature dependence was developed. Rather than correct for the temperature strip-




Figure 4.7: Temperature correction: deviation of fitted photopeak energy from the true photopeak energy in (a) keV and (b) percentage of the true photopeak energy. Averaged across photopeaks, the new temperature correction method in 2020 brings uncorrected energies to within an average of $\sim 0.1\%$ of the true energies, as did the 2016 method on 2016 calibration data.

by-strip, one linear correction is applied to all strips. The correction is derived from the centroid shifts of the ^{22}Na 511 keV photopeak. This choice of line was motivated by a ~ 1.3 keV underestimation of the true 511 keV photopeak energy in the uncorrected ^{22}Na spectrum, among the largest discrepancies across calibration sources (Figure 4.7).

All events in the 2020 ^{22}Na calibration data are first binned by temperature into 0.35°C bins. This bin size is empirically chosen because it is fine enough to observe a shift in the line across the limited range of temperatures and it preserves enough events to sufficiently populate an energy spectrum in each bin. The energy spectrum in each bin is fit with a Gaussian and a linear estimation of the background. A linear fit to the centroid of the Gaussian in each bin against its central temperature yields the linear correction applied to all strips. As shown in Figure 4.8, the uncorrected photopeak energies exhibit a dependence on temperature (note the slope of ~ -0.05 keV/ $^\circ\text{C}$). Analogous plots for other calibration photopeaks are provided in Appendix A for reference; these have not been used to correct calibration data and are only provided to show the temperature dependence across COSI's energy bandpass. Before the correction, the fitted energies are $\sim 0.5 - 1$ keV below the true photopeak energy. After the correction, the fitted energy is less dependent on temperature (the slope is now ~ 0 keV/ $^\circ\text{C}$) and the energies are much closer to the true 511 keV photopeak. Overall, prior to the correction, the average difference between measured and true line energy across all 2020 photopeaks was 0.2%. The correction mitigates the average offset to 0.1%, as did the strip-by-strip correction in 2016.




Figure 4.8: Temperature correction in 2020 ^{22}Na (511.0 keV) calibration data. (a) Before the temperature correction, fitted energy decreases with increasing temperature. (b) After the temperature correction, the fitted energy is much less dependent on temperature (slope ~ 0 keV/ $^{\circ}\text{C}$) and is closer to the true photopeak energy of 511.0 keV.

Further investigation of the temperature dependence is underway, including efforts to reconcile the two methods and encode a potential energy dependence in the correction. It is clear that not all photopeaks see the same benefit from the ^{22}Na correction in the 2020 method (e.g., > 1 MeV lines in Figure 4.7, save the 1274.5 keV line from ^{22}Na). An alternative method of temperature correction is also presented in Mitra et al. (2016). The primary advantage of this method is that it is agnostic to the absolute temperature of the calibration and does not require a pre-existing energy calibration. Instead, one spectrum obtained at an arbitrary temperature is chosen as a “reference spectrum” to which all other spectra are shifted. Future analyses could test this method on COSI data. In summary, thorough characterization of temperature dependence in detector readout is critical to the success of the energy calibration and dedicated calibration time at controlled temperatures would greatly benefit this effort.

4.4 Cross-talk correction

The proximity of a strip electrode to another triggered strip can increase the energy recorded by the former. This electronic influence of the neighboring strip on the measured energy is referred to as “cross-talk.” Correcting for cross-talk in the data is necessary to obtain accurate energy measurements.

Cross-talk is diagnosed by visually examining energy spectra of Compton events. It manifests as a bump in the spectrum at energies higher than the photopeak energy; this




Figure 4.9: The cross-talk effect in 2016 ^{137}Cs calibration data. (a) Prior to the cross-talk correction, the induced enhancement in energy is seen at ~ 670 keV. (b) The spectrum after applying the linear cross-talk correction appears as expected, with a single photopeak at 661.7 keV. Images from Kierans (2018).

is the enhancement caused by the neighboring strip. An example of cross-talk in ^{137}Cs calibration data is shown in Figure 4.9. A single photopeak is expected at 661.7 keV but there is an additional feature at ~ 670 keV. The enhancement scales linearly with energy such that cross-talk is corrected via a simple linear offset to the energy spectra. The parameters of the correction are determined by examining the degree of the offset as a function of energy in events with neighboring and next-to-neighboring activated strips. The parameters derived from 2016 and 2020 calibration data are consistent with each other (Beechert et al. 2022a). After applying the cross-talk correction (Figure 4.9b), the enhancement in Figure 4.9a is removed and the photopeak appears as expected.

4.5 Strip pairing algorithm

Strip pairing localizes the two-dimensional (x - y) position of an interaction to the intersection of triggered AC- and DC-side strips. This procedure is straightforward when there is only one interaction in a detector. If there are multiple interactions in a detector, however, there are several possible interaction locations and the strip pairing algorithm must find the most likely solution by comparing the energies on all candidate strips (Figure 4.10). The strip pairing algorithm is detailed in Sleator (2019).

Intrinsic detector effects that compromise the energy measurements on candidate strips impair the ability of the algorithm to identify the correct solution. These effects include finite




Figure 4.10: Illustration of the strip pairing algorithm. (a) If there is only one interaction in a detector, the interaction location is the intersection of the triggered AC and DC strips. (b) If there are two interactions in a detector, the strip pairing algorithm identifies the most likely (green) of the two possible solutions as that with comparable energies on the AC- and DC-side strips. Images from Sleator (2019).

energy resolution, charge sharing between adjacent strips, multiple interactions occurring on a single strip, and charge loss (either from charge carriers falling between strips or from sub-threshold deposits on strips). The fraction of events deemed too complicated to pair is energy-dependent and ranges from $\sim 17\%$ at 511 keV to $\sim 28\%$ at 1274 keV (Sleator 2019). Future improvements to the strip pairing algorithm that can overcome these effects are underway.

4.6 Depth calibration

While the two-dimensional position of an interaction is derived from the intersection of orthogonal strips on the AC (anode, y) and DC (cathode, x) sides of the GeDs (Section 4.5), the third dimension, depth, is calculated from the collection time difference (CTD) of electrons and holes generated in the interaction.

Photon interactions in the GeDs generate charge clouds of electron-hole pairs. These electrons and holes drift in opposite directions to the AC and DC sides of the detector, respectively, under the influence of the detector's external voltage bias. The high detector voltage bias, ranging from 1000 V to 1500 V, ensures rapid charge carrier collection at high drift velocities along its field lines towards the electrodes. Suppose that a charge cloud is

produced close to the AC side of the detector. The electrons in this charge cloud traverse a shorter distance along the z -axis to the AC side than the holes traverse to the opposing DC side. Hence, the collection time of the electron is less than that of the hole and the CTD serves as a proxy for localizing the depth of the interaction. Extracting this depth from the CTD is referred to as depth calibration.

4.6.1 “Classic” depth calibration

Each of the GeDs’ “pixels” (regions segmented by the grid of orthogonal strips on either side of the detectors) is calibrated individually to maintain unique treatment of each individually-instrumented electronic channel. Small inhomogeneities in detector composition can change drift velocities slightly across the detector volume and vary collection times. The Nuclear Compton Telescope (NCT), a predecessor to COSI, successfully implemented a method of depth calibration to calibrate the $37 \text{ strips} \times 37 \text{ strips} \times 12 \text{ detectors} = 16,428$ pixels (Bowen et al. 2007). This method is referred to as the “classic” approach. Note that because several of the 888 strips (typically fewer than 10) often appear “dead” due to damaged electronics channels, for example, fewer than 16,428 pixels are calibrated.

Detailed outlines of this approach are presented in Lowell et al. (2016) and Lowell (2017). The explanations are adapted here for clarity. There are three main ingredients:

1. **Charge transport simulations** Charge transport simulations create a look-up table of $z_n(\tau_{\text{sim}})$ for each detector n that relates the CTD τ_{sim} to depth z . The simulation models a simplified COSI GeD with a 5×5 strip configuration, detector biases imposed as boundary conditions, and various physical detector characteristics including impurity concentration and thickness. The electrostatic potential inside the active volume of this GeD is solved via Poisson’s equation. The weighting field of each strip is calculated with an alternative boundary condition that sets the bias on that strip to 1 V and that on the rest to 0 V. The Shockley-Ramo theorem (Shockley 1938; Ramo 1939) uses the weighting field to calculate the current induced on the electrodes by moving charge carriers. The look-up table is created from these quantities.
2. **Real calibration run** A real, measured CTD for each pixel is recorded by collecting data with a $\sim 0.08 \text{ mCi } ^{137}\text{Cs}$ calibration source placed above the top of the cryostat (Section 4.2). Only events in the ^{137}Cs photopeak (650–670 keV) and in the 200–477 keV continuum are used. Several hours of collection time yield at least one hundred counts in each bin of the per-pixel CTDs (Figure 4.11a).
3. **Simulated calibration run** A Cosima simulation of the real calibration run produces a histogram of interaction depths, i.e. simulated depth distributions, for each detector. This simulation does not include background events because the ^{137}Cs source in the real calibration run is close enough to the detectors such that the real spectra are source-dominated.

Using the twelve look-up tables from Step 1, the simulated depth distributions from Step 3 (in units of centimeters) are converted to simulated CTD distributions (in units of nanoseconds) for each detector. Each detector’s simulated CTD distribution is called its “CTD template,” which is subsequently used to calibrate its constituent pixels.

A fit for the “stretching” λ and “offset” Δ factors finds those that most closely transform measured CTDs from each pixel in Step 2 into the CTD template of the detector that contains the pixel of interest. Figure 4.11a shows an example of a successful fit. The transformation is defined by

$$\tau_{\text{meas}} = \lambda\tau_{\text{sim}} + \Delta. \quad (4.1)$$

With the resulting λ and Δ for each pixel, Equation 4.1 is solved for a per-pixel τ_{sim} that is converted to a depth via the look-up table in Step 1. The depth calibration thus relates the CTD in each pixel to a z -coordinate of the interaction.


4.6.2 Comparison to 2016

The transformation factors λ and Δ and the reduced χ^2 of the CTD templates to measured CTDs from depth calibrations in 2016 and 2020 are compared. This comparison serves to evaluate the consistency of the calibration approach and the stability of the detectors over time. The averages of the above quantities over all pixels in each detector are shown in Figure 4.11.


The stretching and offset factors in 2016 and 2020 are largely consistent and the reduced χ^2 values are slightly offset from each other. As averaged across all pixels and detectors, the mean 2016 reduced χ^2 is 2.2 and the mean 2020 reduced χ^2 is 1.7. Improvements made to the DEE (Sleator et al. 2019) after the 2016 campaign are likely responsible for the evident improvement in CTD template fits to the data in 2020. Continuous improvements to Cosima and MEGALib may have also enhanced our ability to mimic experimental conditions more accurately. A consistent trend is that the detectors at the top of the COSI cryostat (0, 5, 6, 11) have greater mean reduced χ^2 values than the detectors in the middle and bottom layers. Detectors in the bottom layer of the cryostat (2, 3, 8, 9) have among the lowest mean reduced χ^2 values. One possible reason for this discrepancy is that the higher count rate in the top detectors suppresses the contribution of statistical noise in favor of higher systematic uncertainties.

4.6.3 Future improvements


Work is underway to improve the depth calibration for future studies. A new approach outlined in Lowell et al. (2016) and Lowell (2017) can better account for detector inhomogeneities. An advanced charge transport simulation which fully models COSI’s 37×37 strip configuration rather than the simplified 5×5 configuration described above can help elucidate the impact of these inhomogeneities. Incorporating recent advancements in the




(a) The CTD template yields a good match to the data in pixel 12508 of detector 1.



(b) Average λ



(c) Average Δ



(d) Average reduced χ^2

Figure 4.11: Comparison of depth calibration in 2020 and 2016 calibration data. (a) An example of a measured CTD (blue) and its corresponding fit (red) in 2020 calibration data. The CTD template which generates the fit yields a good match to the data. Comparisons of the 2020 and 2016 average (b) stretching factor λ , (c) offset factor Δ , and (d) reduced χ^2 are also shown. The averages are taken over all calibrated pixels in the specified detector. Symmetric error bars indicate one standard deviation spread in value as averaged across all pixels in each detector.

DEE will also facilitate exploration of small-scale detector effects beyond the scope of the “classic” approach. Furthermore, future work to incorporate depth-dependent energy measurements (Appendix C) may aid the depth calibration, and by extension, associated efforts to understand how the timing resolution of COSI varies with interaction energy.

4.7 Detector Effects Engine

Simulations are routinely conducted to study COSI’s expected response to calibration and flight data. These predictions inform important metrics of instrument performance, including angular resolution and effective area. Simulations, however, are generated in physical units and are free from imperfections intrinsic to measurements in the COSI detectors. As such, they do not accurately reflect COSI’s response to real events. The detector effects engine (DEE) converts these simulated events, which represent events that would have been measured in an ideal detector, to events which mimic real data.

The DEE inverts the calibrations discussed in previous sections of this chapter to store the simulated hits in electronic, rather than physical, units (detector ID, timing, and ADC). Then, the DEE artificially applies effects such as charge sharing, charge loss, and cross-talk to the events. It also vetoes GeD events coincident with those in the shields, removes events on dead GeD strips, and removes events which would have occurred during the dead time of COSI’s electronics. These treated simulations then run through the same event calibration pipeline as real data. Nuclearizer calibrates the simulations and real data identically; comparisons of the results benchmark the ability of the DEE to accurately transform simulations and help identify inaccuracies in the calibration pipeline. Extensive tests that demonstrate the ability of the DEE to successfully model real detector effects are detailed in Sleator et al. (2019).

4.8 Instrument performance

The COSI team collected approximately 1500 hours of calibration data over more than 200 separate measurements prior to the 2016 balloon launch. The analogous calibration program prior to the attempted 2020 launch was limited to 10 days. Though more narrow in scope, the calibration data from 2020 remain a valuable check of instrument performance over years of operation. It is important to study these calibration data repeatedly as analysis techniques are developed with time. The benchmarking analyses described in this section were conducted after the cancellation of the 2020 campaign and are a cumulative reflection of COSI’s most recently determined capabilities. As there were no major design changes, the instrument’s performance between the 2016 and 2020 campaigns was expected and proved to be consistent.

4.8.1 Energy resolution

COSI’s spectral performance is quantified as the energy resolution of fully reconstructed Compton events. Note the distinction between this energy resolution and the single-strip energy resolution in Section 4.3. The single-strip resolution considers hits on individual strips rather than the total energy resolution of a fully reconstructed event with several energy depositions across multiple detectors. Fully reconstructed Compton events thus inform the spectral performance of the instrument as a whole, while the single-strip resolution measures the spectral performance of the individual GeDs. Studying the fully reconstructed energy resolution can also gauge the fidelity of the combined calibration procedures discussed previously. For example, much coarser energy resolution than expected could indicate an insufficient cross-talk correction which failed to remove high-energy enhancements.

The full event reconstruction is performed in MEGAlib (Figure 4.12). Energy calibration converts the ADC of individual hits to units of energy. Temperature and cross-talk corrections remove problematic effects in the spectra. The strip pairing algorithm and depth calibration then localize each interaction to a 3-D coordinate in the detectors. Revan constructs events from these calibrated hits and the energy resolution is calculated as the ratio of a reconstructed photopeak’s FWHM to its true line energy.

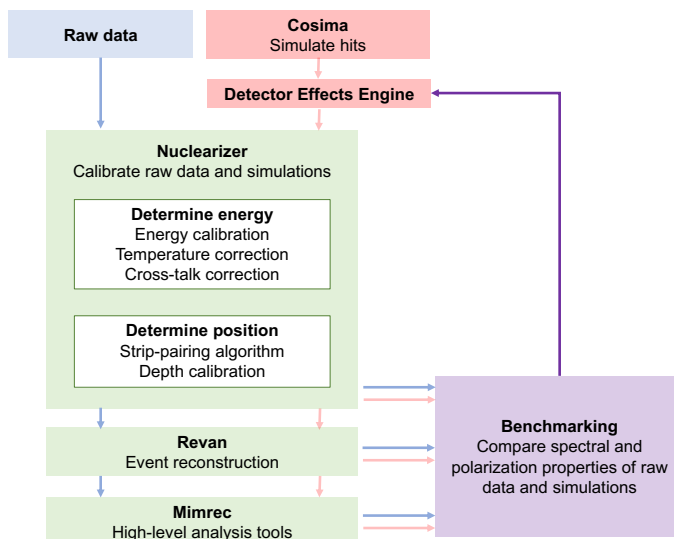


Figure 4.12: Flowchart of calibration and event reconstruction in MEGAlib. Raw data and DEE-processed simulations from Cosima are identically calibrated in Nuclearizer. Revan combines these calibrated hits into reconstructed events that are loaded into Mimrec for high-level analysis tasks including energy resolution, angular resolution, effective area, and polarization. The DEE is iteratively tuned to match these metrics in real data and simulations as closely as possible.

The energy resolution as a function of energy in 2020 data is shown in Table 4.3. Only

Isotope	Line energy [keV]	2020 [keV]	2020 [%]	2016 [keV]	2016 [%]
²² Na	511.0	5.78 ± 0.01	1.1	5.56 ± 0.04	1.1
¹³⁷ Cs	661.7	5.27 ± 0.01	0.8	5.1 ± 0.02	0.8
⁶⁰ Co	1173.2	6.80 ± 0.02	0.6	7.36 ± 0.05	0.6
²² Na	1274.5	7.04 ± 0.03	0.6	6.42 ± 0.1	0.5
⁶⁰ Co	1332.5	6.97 ± 0.02	0.5	6.95 ± 0.05	0.5

Table 4.3: Fully-reconstructed energy resolution (FWHM) in 2020 and 2016 (Sleator et al. 2019) calibration data.

¹³⁷Cs, ⁶⁰Co, and ²²Na data were collected prior to the launch cancellation. The resolutions in 2020 data are comparable to those reported in Sleator et al. (2019) from 2016 data.

4.8.2 Angular resolution

The angular resolution measure (ARM) of a Compton telescope is the smallest angular distance between a source’s known location and the event circle of each event (Figure 4.13a). Given a sample of Compton events, the distribution of ARM values represents the effective width of the telescope’s point spread function (Figure 4.13b). The angular resolution of the instrument is defined as the FWHM of this ARM distribution. The maximum achievable (lowest FWHM) angular resolution is fundamentally limited by the Doppler broadening of the scattering electron (Du Mond 1929), which is neither free nor at rest as assumed in the Compton equation. The bound and moving electron necessarily introduces some spread in the ARM distribution, giving an angular resolution of $\sim 1^\circ$ (Zoglauer and Kanbach 2003). In practice, the angular resolution is limited by the 3-D position resolution of the instrument, which is primarily governed by the strip pitch of its detectors. The accuracy of this position measurement (and the energy measurement) affects the quality of the event reconstruction.

The angular resolution of COSI as measured in 2020 and 2016 calibration data is compared. The 2020 data consist of ⁶⁰Co, ¹³⁷Cs, and ²²Na measurements taken with sources directly overhead the cryostat (at zenith). Analogous measurements from 2016 calibration data are used and include points at 898 keV and 1836 keV from ⁸⁸Y data as additional points of reference; the COSI team in 2016 had ample time to collect calibration data from more sources. Photopeak events with a $\pm 1.5\sigma$ energy cut on each line emission, where σ is the Gaussian width of the line measured in Section 4.8.1, are selected. Additional event selections require a Compton scattering angle less than 90° , a minimum distance between any two interactions of 1 cm, and that events originate from within 90° of COSI’s zenith. These event selections are empirically designed to optimize the angular resolution. The ARM FWHM is highly dependent on these cuts and changes with different event selections. The central peak of the ARM distribution (e.g., $\pm 6^\circ$ at 662 keV) was fitted with a double Lorentzian plus asymmetric Gaussian function to determine the FWHM (see Figure 4.13b for an example of

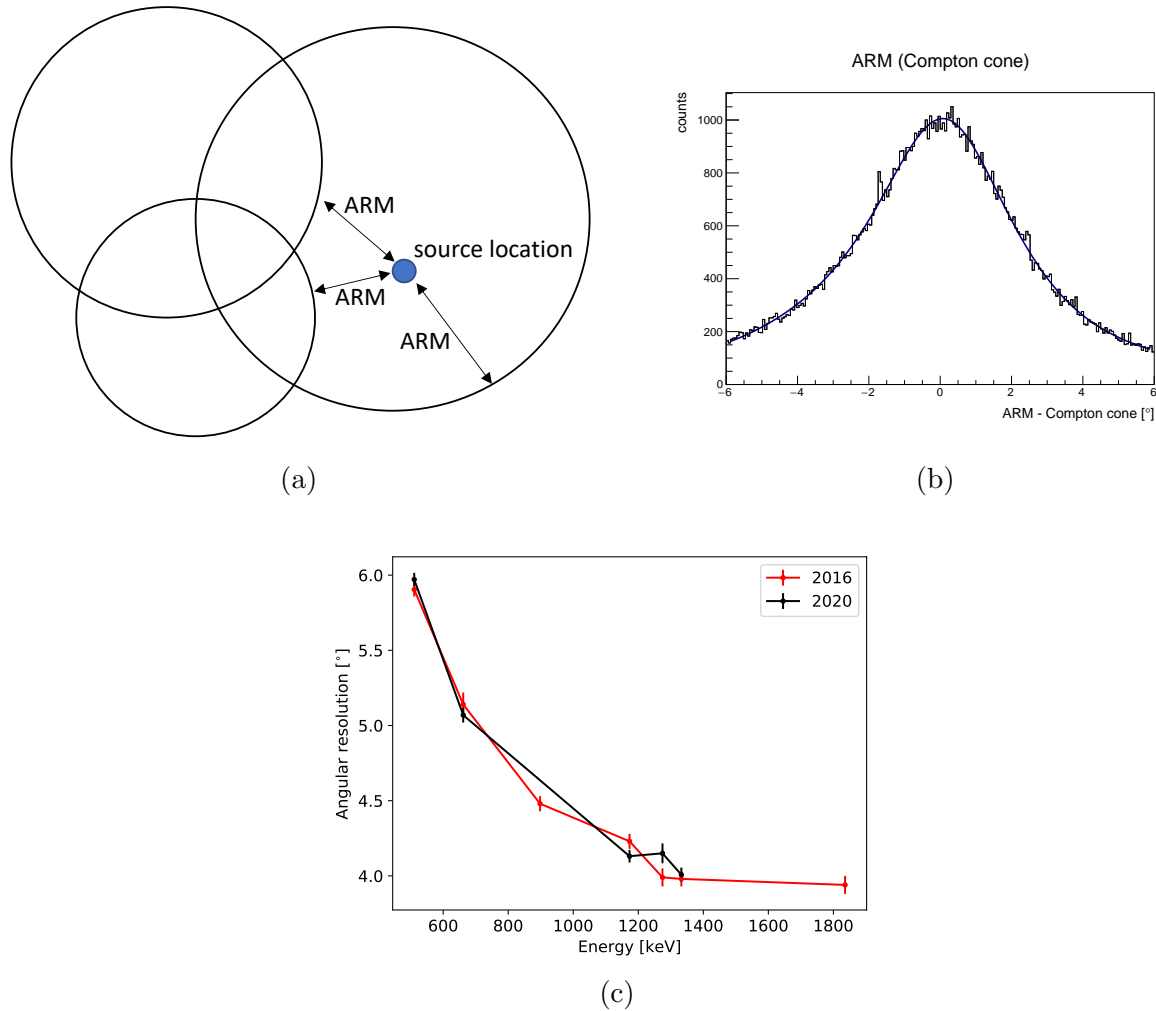


Figure 4.13: Illustrations of the angular resolution of a Compton telescope. (a) The angular resolution measure (ARM) of each event is the smallest angular distance between the known source location (blue dot) and the event circle (black circle). (b) The distribution of ARM values for 2020 ^{137}Cs data. The FWHM of the fit to the central peak of the ARM distribution defines the angular resolution. (c) The angular resolution as a function of energy in 2016 and 2020 calibration data.

the distribution and fit to 2020 ^{137}Cs data).

As shown in Figure 4.13c and Table 4.4, the angular resolution measurements in 2016 and 2020 are consistent over the tested energy range. The results from both years are new analyses; a similar analysis on 2016 data was originally published in Sleator et al. (2019). The expected relationship between angular resolution and photon energy is observed. Angular resolution improves (smooth, monotonically decreasing ARM values) with increasing energy because higher incident photon energy increases the distance between Compton interactions. Greater distance between Compton interactions improves the accuracy of the event reconstruction. Thus, the angular resolution analyses of 2020 and 2016 data are consistent with each other and with instrumental expectations.

Isotope	Line energy [keV]	2020 [$^{\circ}$]	2016 [$^{\circ}$]
^{22}Na	511.0	5.97 ± 0.04	5.9 ± 0.1
^{137}Cs	661.7	5.1 ± 0.1	5.1 ± 0.1
^{88}Y	898.0	–	4.5 ± 0.1
^{60}Co	1173.2	4.13 ± 0.04	4.2 ± 0.1
^{22}Na	1274.5	6.1 ± 0.3	6.5 ± 0.3
^{60}Co	1332.5	4.2 ± 0.1	4.0 ± 0.1
^{88}Y	1836.0	–	3.9 ± 0.1

Table 4.4: Angular resolution (FWHM) in 2020 and 2016 calibration data. Event selections: Compton events with incident photon energy within $\pm 1.5\sigma$ of the photopeak line energy, Compton scattering angle 0° to 90° , 2–7 total interactions, minimum distance between any two interactions of 1 cm.

4.8.3 Effective area

The geometric size and detection efficiency of a telescope are quantified by its effective area. The effective area is the area seen by an incident photon as it interacts with the instrument, weighted by effects including interaction probabilities, detector geometry, and event selections applied to the data. Understanding the sensitivity of the instrument to photons incident from across its field of view and energy bandpass relies upon effective area calculations.

Effective area is defined as the product of the telescope’s true geometric collecting area and its efficiency ϵ at a given photon energy. The efficiency is the fraction of all incident photons that are detected by the instrument. Equation 4.2 formulates this fraction using measured and incident photon rates in calibration data. The measured photon rate, L_{meas} , is the number of photons in the measured photopeak, N , divided by the exposure time t_{meas} . This number N is affected by the dead time of the instrument’s electronics. Interactions which occur while the analog boards and coincidence logic are “high” ($\sim 10 \mu\text{s}$ after previously

triggered events) are not recorded. As the calibration sources are of sufficient activity to induce non-negligible dead time, N is scaled by a multiplicative dead time factor δ .

The incident photon rate, L_{inc} , is determined through known parameters of the data collection. It is the product of the radioactive source activity S , the branching ratio BR of the photopeak decay, and the detector area A_{geo} divided by $4\pi z^2$ for a source placed distance z from the detector. The latter term is the surface area subtended by the instrument of the $4\pi z^2$ sphere surrounding the source. Though photon attenuation in air is negligible in COSI's γ -ray energy bandpass and near-field calibration distances z , it can be modeled in the calculation of L_{inc} as an exponential decay with linear attenuation coefficient, λ , of air at a given photopeak energy. For example, the exponential air attenuation term for a typical COSI source calibration distance of $z \sim 60$ cm is ~ 0.996 at 1.5 MeV.

Given these definitions, the effective area A_{eff} reads as follows:

$$\begin{aligned} A_{\text{eff}} &= A_{\text{geo}} \cdot \epsilon \\ \epsilon &= L_{\text{meas}}/L_{\text{inc}} \\ L_{\text{meas}} &= \frac{N \cdot \delta}{t_{\text{meas}}}, \text{ for } \delta = \frac{100}{100 - \text{dead time } [\%]} \\ L_{\text{inc}} &= S \cdot \text{BR} \cdot \frac{A_{\text{geo}}}{4\pi z^2} \cdot e^{-\lambda z} \end{aligned}$$

Then,

$$A_{\text{eff}} = 4\pi z^2 \frac{N \cdot \delta}{t_{\text{meas}} \cdot S \cdot \text{BR} \cdot e^{-\lambda z}} \quad (4.2)$$

The effective area of COSI is calculated with 2020 and 2016 calibration data. Only Compton events which pass the following event selections are considered: incident photon energy within $\pm 2\sigma$ of the photopeak line energy, Compton scattering angle of 0° – 180° , 2–7 total interactions, minimum distance between the first two interactions of 0.5 cm, and minimum distance between any two interactions of 0.3 cm. As shown in Table 4.5 and Figure 4.14, the effective area values from 2016 and 2020 are consistent. Relaxing the minimum distances between interactions to 0 cm in 2020 data demonstrates that the calculated effective area is dependent on event selections (which dictate N in the above formulation). This new set of “open” selections increases the effective area by $\sim 8.1\%$ at 511 keV and $\sim 6.4\%$ averaged across all measured photopeaks, with decreasing improvement as the photon energy increases. Note that this example of decreasing the minimum distance between interactions worsens angular resolution; curating event selections to the analysis task at hand often requires careful consideration of competing performance metrics. Overall, both years show consistent trends of decreasing effective area with increasing energy. This relationship between effective area and energy is expected because as incident photon energy increases, so does the probability that a scattered photon escapes the instrument without depositing all of its energy in the detector volume. These incompletely absorbed events cannot be correctly reconstructed and thereby reduce ϵ through N , limiting COSI's effective area.




Figure 4.14: The effective area as a function of energy in 2020 and 2016 calibration data. The error bars include statistical uncertainties and systematic uncertainties in source activity and distance from the detector.

Isotope	Line energy [keV]	2020 [cm ²]	2020 “Open” [cm ²]	2016 [cm ²]
²² Na	511.0	12.0 ± 0.5	13.0 ± 0.6	11.9 ± 0.5
¹³⁷ Cs	661.7	9.6 ± 0.4	10.3 ± 0.5	9.7 ± 0.4
⁸⁸ Y	898.0	–	–	8.2 ± 0.4
⁶⁰ Co	1173.2	6.4 ± 0.3	6.7 ± 0.3	6.3 ± 0.3
²² Na	1274.5	6.1 ± 0.3	6.5 ± 0.3	5.8 ± 0.3
⁶⁰ Co	1332.5	5.3 ± 0.2	5.6 ± 0.3	5.3 ± 0.2
⁸⁸ Y	1836.0	–	–	4.0 ± 0.2

Table 4.5: Effective area in 2020 and 2016 calibration data. Event selections: Compton events of incident photon energy within $\pm 2\sigma$ of the photopeak energy, Compton scattering angle 0° – 180° , 2–7 total interactions, minimum distance between the first two interactions of 0.5 cm, and minimum distance between any two interactions of 0.3 cm. “Open” event selections on 2020 data relax the minimum interaction distance to 0 cm. The error bars include statistical uncertainties and systematic uncertainties in source activity and distance from the detector.

4.8.4 Polarization response

The probability that a photon undergoes Compton scattering is given by the Klein-Nishina differential cross-section. Compton telescopes are sensitive to polarization because the Klein-Nishina equation is dependent on the azimuthal scattering direction η , which encodes information about the linear polarization of the scattered photon. COSI’s polarization response was determined with data taken in 2019 at SSL using a custom-built calibration structure that produced partially polarized beams of photons. The process of extracting the polarization response through comparisons of these data to corresponding simulations is detailed in Lowell et al. (2016); Beechert et al. (2022a); Tomsick et al. (2022). The COSI team intended to repeat the procedure by collecting new calibration data in 2020 but did not have time to do so prior to the cancellation of the campaign. A brief summary of the polarization response as determined with 2019 data is presented here for completeness.

Algebraically manipulating the Klein-Nishina equation gives the probability density function of a photon with initial energy E to Compton scatter at azimuthal angle η :

$$p(\eta; E, \phi, \Pi, \eta_0) = \frac{1}{2\pi}[1 - \Pi\mu(E, \phi)\cos(2(\eta - \eta_0))], \quad (4.3)$$

where ϕ is the Compton scattering angle, Π is the polarization level, η_0 is the polarization angle, and $\mu(E, \phi)$ is the “modulation,” defined as follows for scattered photon energy E' :

$$\mu(E, \phi) = \frac{\sin^2\phi}{\frac{E'}{E} + \frac{E}{E'} - \sin^2\phi}. \quad (4.4)$$

The polarization level $\Pi \in [0, 1]$ is defined as the fraction of photons whose electric field points in a specific direction, i.e. the fraction of photons that are linearly polarized at polarization angle η_0 . The probability in Equation 4.3 obeys a sinusoidal modulation with $(\eta - \eta_0)$ and is maximized at $\eta - \eta_0 = \pm 90^\circ$. This modulation is visible in real data, where the distribution of observed counts as a function of η is called the measured azimuthal scattering angle distribution (ASAD). The polarization properties of these data are extracted by fitting the ASAD with a generalized form of Equation 4.3, where A , B , and η_0 are free parameters:

$$A - B\cos(2(\eta - \eta_0)). \quad (4.5)$$

The offset A and amplitude B give the measured modulation $\hat{\mu} = B/A$, which is in turn used to calculate the polarization level $\Pi = \hat{\mu}/\mu_{100}$. The “modulation factor” μ_{100} is a simulated estimate of the modulation when $\Pi = 1$.

Calibrations at SSL in 2019 produced ASADs with a partially polarized γ -ray beam, a background ASAD, and an unpolarized ASAD generated with simulations. The background-subtracted source ASAD, after applying a geometrical correction derived from the unpolarized ASAD, was fit with Equation 4.5 (Figure 4.15). This corrected and fitted ASAD validates COSI’s capabilities as a polarimeter: for a known polarization angle, COSI observes the expected sinusoidal modulation.




Figure 4.15: The measured ASAD (blue) from partially polarized calibration data collected at SSL in 2019. The ASAD is corrected for background and geometric effects intrinsic to the instrument. The fit of Equation 4.5 to the data is shown in red. The expected modulation is clearly visible and is well-matched by simulations.

4.9 Summary

This chapter detailed the procedures used to calibrate COSI prior to its balloon campaign in 2016 and its attempted re-flight in 2020 from Wanaka, New Zealand. Data from both campaigns indicate consistent instrument performance between the years and that the calibration procedures are repeatable. The author of this dissertation performed the 2020 energy calibration, single-strip energy resolution, temperature correction, depth calibration, fully-reconstructed energy resolution, angular resolution, and effective area analyses presented in this chapter. The temperature correction for 2020 was specifically developed by the author for the purposes of this analysis. The author also conducted the reported single-strip energy resolution, angular resolution, and effective area analyses of 2016 data; the energy calibration, temperature correction, depth calibration, and fully-reconstructed energy resolution results from 2016 data are from previous work (Kierans 2018; Lowell 2017; Sleator et al. 2019).

The 2016 calibrations are of particular relevance to the analysis of Galactic ^{26}Al discussed in Chapter 5 of this dissertation. Per the above analyses of instrument performance, the 1809 keV signature of ^{26}Al in balloon flight data is measurable by COSI with $\sim 0.2\%$ en-

ergy resolution and $\sim 4^\circ$ angular resolution. Interpretation of the flight data relies on these calibrations. The simulations in Chapter 5 also rely upon the Detector Effects Engine and calibration pipeline, underscoring the importance of these calibration analyses to COSI's science goals.

Chapter 5

Measurement of Galactic Aluminum-26 in the COSI 2016 Flight

The preceding chapters of this document outlined the scientific wealth of MeV γ -ray astrophysics and explained why COSI is a capable observatory in this regime. Chapter 2 established ^{26}Al as a valuable tracer of nucleosynthesis in the Milky Way Galaxy. Chapter 3 detailed the variety of instrument designs available to experimentalists and presented COSI as a CCT developed for MeV science. The COSI calibration procedures and metrics of instrument performance described in Chapter 4 facilitate scientific study. Analyses of COSI’s successful balloon flight in 2016 include the detection (Kierans 2018; Kierans et al. 2020) and imaging (Siegert et al. 2020) of the 511 keV positron-electron annihilation excess, detection of the Crab Nebula (Sleator 2019), and an upper limit on the polarization of GRB 160530A, which was detected mid-flight (Lowell 2017; Lowell et al. 2017). This chapter¹ adds the first measurement of ^{26}Al in the COSI 2016 balloon flight to this list.

This analysis is the first demonstration of COSI’s ability to study nuclear lines from decaying radioisotopes and is key proof-of-concept for anticipated advancements in nucleosynthesis studies by the COSI satellite mission. As mentioned previously, COSI’s germanium detectors will have an energy resolution comparable to that of SPI and a wide field of view ($\sim 60^\circ$) that exceeds SPI’s $\sim 16^\circ$. Similarly, COMPTEL had a field of view comparable to COSI but COSI’s energy resolution will exceed COMPTEL’s by an order of magnitude. COSI’s $\sim 2^\circ$ angular resolution at 1809 keV will also improve upon those of SPI and COMPTEL ($\sim 2.7^\circ$ and $\sim 3.8^\circ$, respectively). Overall, the analysis in this chapter sets the stage for the COSI satellite’s future contributions to ^{26}Al science. It reports a 3.7σ significance measurement of ^{26}Al in the Inner Galaxy and places upper limits on the dynamics of this emission. The measured Inner Galaxy flux is $(8.6 \pm 2.5) \times 10^{-4} \text{ ph cm}^{-2} \text{ s}^{-1}$. Systematic uncertainties, some of which may be mitigated on the satellite platform, are discussed at length.

¹This chapter largely follows the published paper “Measurement of ^{26}Al by the Compton Spectrometer and Imager” by Beechert et al. (2022b). ©Astrophysical Journal 2022, reprinted with permission.

5.1 Data selection

We select data from the 2016 flight (Section 3.5) based on previous observations of ^{26}Al and through cuts in the Compton Data Space (CDS; Section 6.1). The event time (UTC) and energy of each incident photon are also recorded. We integrate over the scattered γ -ray direction (ψ, χ) since we are not performing imaging; these quantities are not relevant to the analysis described in this chapter. We use the recorded photon energy for spectral analysis and use the event time to select data from the signal and background regions of the flight.

Previous studies of Galactic ^{26}Al by COMPTEL and SPI indicate that its emission is concentrated in the Inner Galaxy ($|\ell| \leq 30^\circ$, $|b| \leq 10^\circ$). Following these observations, we define the Inner Galaxy as our signal region and the area of the sky exclusive of the Inner Galaxy as the background region. We note that this is a conservative approach to the search for ^{26}Al that emphasizes the well-constrained Inner Galaxy emission compared to hints of emission inside our chosen background region (see Section 5.5.2 for further discussion about the systematic uncertainties associated with this approach). The signal and background region data therefore comprise observations taken when COSI’s zenith pointing fell inside the respective regions.

The Compton scattering angle effectively broadens the observation region. A zero-degree Compton scattering angle points back at the source location in image space, and an increase in the accepted Compton scattering angle broadens this image space region by the same angle in the CDS. We therefore expect photons from a region extending beyond the Inner Galaxy out to a maximum Compton scattering angle ϕ_{max} to contribute to the signal region spectrum. To prevent overlap between the signal and background regions, the pointing cuts for the background region are chosen such that the ϕ_{max} extensions beyond the borders of the signal and background regions fall tangential to each other (Table 5.1). We define ϕ_{max} using an optimization procedure (Appendix E) which reveals that $\phi_{\text{max}} = 35^\circ$ yields an acceptable signal-to-noise ratio and preserves a fraction of the sky outside of the signal region large enough for sufficient background statistics. A minimum $\phi_{\text{min}} = 10^\circ$ removes more atmospheric background (Ling 1975) than ^{26}Al signal events. Thus, we apply a cut in the CDS on the Compton scattering angle ϕ as an optimized event selection which aims to reduce the background in the selected data. The signal and background regions, with their 35° extensions, are superimposed on the SPI 1.8 MeV image in Figure 5.1. Figure 5.2 shows the COSI 2016 flight path displayed over the signal and background regions.

Additional event selections common to both the signal and background regions focus the analysis to ^{26}Al . Only Compton events with incident energy of 1750–1850 keV and incident angle $\leq 90^\circ$ from COSI’s zenith are considered. This restriction in incident angle, called the “Earth Horizon Cut,” reduces the dominant Earth albedo background. COSI’s six anticoincidence CsI shields actively suppress the Earth albedo radiation by vetoing γ -rays incident from below the instrument. The shield veto system reduces atmospheric background levels by $\sim 1 - 2$ orders of magnitude above 1750 keV. Installing these shields for atmospheric background rejection introduces the potential for instrumental activation of the shield materials, the γ -ray lines from which are accounted for empirically in the presented analysis. The


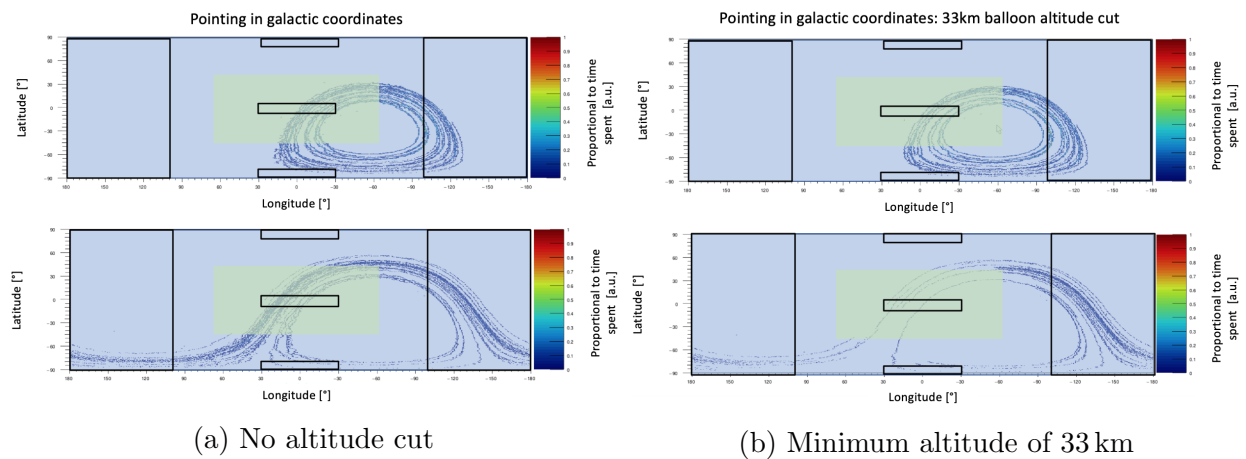


Figure 5.1: The signal and background regions (Table 5.1) of the COSI balloon ^{26}Al analysis displayed over the SPI 1.8 MeV image (Bouchet et al. 2015). The signal region is defined by the Inner Galaxy (black rectangular outline) and the surrounding hatched green shading maps the effective broadening of this region by the maximum Compton scattering angle $\phi_{\text{max}} = 35^\circ$. The gray and hatched gray shadings map the background region and its effective 35° broadening, respectively. There is no overlap between the broadened signal and broadened background regions.



(a) No altitude cut

(b) Minimum altitude of 33 km

Figure 5.2: The COSI 2016 flight path displayed over the signal and background regions of the ^{26}Al analysis. (a) Flight pointings over all balloon altitudes. (b) Flight pointings at minimum balloon altitude of 33 km. Top row: Flight path from May 17, 2016 to June 5, 2016, when COSI had 10 operational detectors. Bottom row: Flight path from June 6, 2016 to July 2, 2016, when COSI had 9 operational detectors. Overlay on all: The signal region and its effective 35° broadening are marked in green, with the central Inner Galaxy pointing cut marked by a black rectangle. The background region is analogously shown in blue with its three constituent regions marked by black rectangles (Table 5.1).

Region	Galactic (ℓ, b) [$^{\circ}$]
Signal	$(0 \pm 30, 0 \pm 10)$
Background Region 1	$(-180 \pm 80, 0 \pm 90)$
Background Region 2	$(0 \pm 30, 85 \pm 5)$
Background Region 3	$(0 \pm 30, -85 \pm 5)$

Table 5.1: The longitudinal (ℓ) and latitudinal (b) definitions of the signal and background regions in the COSI 2016 ^{26}Al analysis. The latter is comprised of three constituent pointing cuts.

Parameter	Permitted values
Altitude in signal, background regions	≥ 33 km, all
Energy	1750–1850 keV
Compton scattering angle ϕ	10° – 35°
Number of Compton scatters	2–7
Minimum distance between the first two, any interactions	0.5, 0.3 cm
Earth Horizon Cut	Accept only events originating above the Earth’s horizon

Table 5.2: Event selections on flight data in the ^{26}Al signal and background regions. The resulting observation time in the signal region is 156 ks and that in the background region is 1356 ks.

minimum distance between the first two interactions is 0.5 cm and that between subsequent interactions of all events is 0.3 cm. These minimum distances improve angular resolution. Overall, we permit 2–7 total interactions per event; a minimum of two interactions is required for the reconstruction of Compton events and events with greater than seven interactions are likely to be pair-production events that cannot be reconstructed (Boggs and Jean 2000).

To combat worsening atmospheric background and attenuation with decreasing balloon altitude, we require a minimum balloon altitude of 33 km in the signal region. All balloon altitudes are considered in the background region to preserve more statistics for improved determination of the spectral shape of the background, which is not expected to change with altitude. The only observations removed from the background region data are those taken before the balloon reached float altitude and those with high shield rates. These event selections (Table 5.2) yield a total observation time of $T_{\text{SR}} \approx 156$ ks in the signal region and $T_{\text{BR}} \approx 1356$ ks in the background region. Data and simulations from before and after June 6, 2016 are processed with 10- and 9-detector mass models, respectively.

A full spectrum of the flight containing events which pass the signal and background region event selections is shown in Figure 5.3. The spectra are normalized by the observation

time in each region. The bottom panel is the difference between the signal and background region spectra and the result is smoothed with a Gaussian filter of width $\sigma = 5$ keV for clarity. In addition to the strong 511 keV line and a general continuum, a peak near 1809 keV is visible.




Figure 5.3: The full COSI 2016 flight spectrum of events which pass the signal and background region event selections (top) and the background-subtracted spectrum (bottom). The background-subtracted spectrum is smoothed by a Gaussian filter of width $\sigma = 5$ keV. Error bars are $\sqrt{\text{counts}}$.

5.2 Residual method

As an initial attempt, a simple scaling and residual method is employed to gauge the feasibility of an ²⁶Al measurement. Flight data from the signal and background regions are considered with the event selections listed in Table 5.2 and are divided into 10- and 9-detector portions of the flight. These four data sets are binned into one time bin encompassing the full duration of each and are integrated over all scattering angles in the CDS (ϕ, ψ, χ). The data are binned into three energy bands: 1750–1803 keV, 1803–1817 keV, and 1817–1850 keV. The first and third bins comprise an effective continuum (background) band around the central

bin of interest, which contains the 1809 keV signature of ^{26}Al . The desired result of this procedure is to observe an excess of counts in the central energy bin after subtracting the appropriately scaled counts in the continuum bins.

We estimate the source counts $N_S(E)$ by subtracting the background region measurement, $N_{\text{BR}}(E)$, from the signal region measurement, $N_{\text{SR}}(E)$, with an estimated scaling factor f such that

$$N_S(E) = N_{\text{SR}}(E) - f \cdot N_{\text{BR}}(E). \quad (5.1)$$

The scaling factor f scales the background region counts N_{BR}^C to the signal region counts N_{SR}^C in each of the continuum bins $C = 1, 2$: $f^C = N_{\text{SR}}^C/N_{\text{BR}}^C$. This yields two scaling factors f^1 and f^2 which are averaged to obtain the overall scaling factor $f = (f^1 + f^2)/2$. The background region counts are scaled to the signal region counts to preserve the true number of candidate ^{26}Al events in the latter. Scaling the signal instead of the background could artificially enhance the former, effectively creating photons that are not present in the real data set.

For an ideal scaling factor f , the background-subtracted spectrum should yield approximately zero counts in the continuum bins (perfect background subtraction). A non-zero excess of counts in the central 1803–1817 keV bin would then suggest the presence of ^{26}Al in the signal region. Two different methods are employed to estimate the scaling factor f :

1. Independent scaling: Spectra from the 9- and 10-detector portions of the flight are scaled separately by their own, independently calculated scaling factors. These scaled spectra are subsequently summed together to obtain one spectrum representing the entire flight:

$$\begin{aligned} f_{9 \text{ dets}}^C &= N_{\text{SR}, 9 \text{ dets}}^C / N_{\text{BR}, 9 \text{ dets}}^C \\ f_{10 \text{ dets}}^C &= N_{\text{SR}, 10 \text{ dets}}^C / N_{\text{BR}, 10 \text{ dets}}^C \\ f_{9 \text{ dets}} &= \frac{f_{9 \text{ dets}}^1 + f_{9 \text{ dets}}^2}{2} \\ f_{10 \text{ dets}} &= \frac{f_{10 \text{ dets}}^1 + f_{10 \text{ dets}}^2}{2} \\ N_{\text{BR}, \text{ indep, entire flight}} &= f_{9 \text{ dets}} \cdot N_{\text{BR}, 9 \text{ dets}} + f_{10 \text{ dets}} \cdot N_{\text{BR}, 10 \text{ dets}} \\ N_{\text{SR}, \text{ entire flight}} &= N_{\text{SR}, 9 \text{ dets}} + N_{\text{SR}, 10 \text{ dets}} \\ N_S, \text{ indep} &= N_{\text{SR}, \text{ entire flight}} - N_{\text{BR}, \text{ indep, entire flight}} \end{aligned}$$

2. Total Scaling: Spectra from the 9- and 10-detector portions of the flight are first summed together to obtain one spectrum of the entire flight. A total scaling factor is subsequently calculated and applied to the background:




Figure 5.4: Residual counts between signal and background region measurements indicate an excess contribution of signal region counts of energy 1803–1817 keV. The background region measurements are scaled via “independent” and “total” methods. The data points are plotted in the center of each energy bin. Horizontal error bars span the width of each energy bin. Vertical error bars are calculated assuming Poisson statistics.

$$\begin{aligned}
 N_{\text{BR, entire flight}} &= N_{\text{BR, 9 dets}} + N_{\text{BR, 10 dets}} \\
 f_{\text{total}}^C &= N_{\text{SR, entire flight}}^C / N_{\text{BR, entire flight}}^C \\
 f_{\text{total}} &= \frac{f_{\text{total}}^1 + f_{\text{total}}^2}{2} \\
 N_{\text{BR, total, entire flight}} &= f_{\text{total}} \cdot N_{\text{BR, entire flight}} \\
 N_{\text{SR, entire flight}} &= N_{\text{SR, 9 dets}} + N_{\text{SR, 10 dets}} \\
 N_{\text{S, total}} &= N_{\text{SR, entire flight}} - N_{\text{BR, total, entire flight}}
 \end{aligned}$$

The residual plot comparing $N_{\text{S, indep}}$ and $N_{\text{S, total}}$ is shown in Figure 5.4. Both the independent and total scaling methods yield an excess contribution of signal region counts in the central 1803–1817 keV energy bin. The discrepancy between the two methods can be understood by examining their scaling factors. The independent case relies on the application of two scaling factors $f_{9 \text{ dets}}$ and $f_{10 \text{ dets}}$. The total case uses one scaling factor f_{total} . Though the average of $f_{9 \text{ dets}}$ and $f_{10 \text{ dets}}$, approximately 0.14, is very close to $f_{\text{total}} \sim 0.15$, $f_{10 \text{ dets}} \sim 0.26$ is much greater than $f_{9 \text{ dets}} \sim 0.03$. Thus, $N_{\text{BR, 10 dets}}$, which is already greater than $N_{\text{BR, 9 dets}}$, is multiplied by a factor $f_{10 \text{ dets}}$ that exceeds $f_{9 \text{ dets}}$ and the total

scaling factor. This means that the independent scaling method estimates more background region counts in the 1803–1817 keV bin than does the total scaling method. As a result, the background-subtracted spectrum in the total scaling case shows more residual counts in the 1803–1817 keV bin.

In summary, the excess emission in Figure 5.4 suggests an enhanced number of photons with energy 1803–1817 keV emanating from the signal region, which is precisely the expectation for ^{26}Al emission of 1809 keV γ -rays from the Inner Galaxy. This finding supports the potential residual emission visually identified in the full flight spectrum of Figure 5.3. However, this relatively coarse residual method must be interpreted cautiously, given the low statistics of the data set. The excess number of counts keV^{-1} in Figure 5.4 is less than 10. Small changes in the scaling factor could meaningfully alter the residual signal. Furthermore, because the scaling factor is not a whole number, the residual method is un-physical in its subtraction of a decimal number of background counts from an integer number of real, observed counts. With these limitations in mind, we treat the scaling method and the background-subtracted spectrum in Figure 5.3 as encouraging indications of a potential ^{26}Al measurement and employ a maximum likelihood technique to formalize the analysis.

5.3 Maximum likelihood method

The residual method in Section 5.2 derives a background model from higher and lower energy bins to estimate the number of background counts in the signal energy bin. The maximum likelihood method in this section instead finds fit parameters which describe the most likely contributions of sky (signal) and background models to 1750–1850 keV photons in the signal region data.

COSI data, d , are modeled as a linear combination of a sky model, s , and a background model, b , with unknown amplitudes α and β , respectively. The data are binned in 1-keV bins, i , spanning 1750 to 1850 keV, such that the model reads

$$m_i = \alpha s_i + \beta b_i. \quad (5.2)$$

The following sections describe model templates s and b in detail. Photon counting is a Poisson process and the likelihood that data d are produced by model m is given by the Poisson distribution

$$L(d|m) = \prod_{i=1}^N \frac{m_i^{d_i} e^{-m_i}}{d_i!}, \quad (5.3)$$

where $N = 100$ energy bins. We fit for the scaling factors α and β in the signal region data d_i by minimizing the Cash statistic (Cash 1979), which is the negative logarithm of the likelihood in Equation 5.3, agnostic to model-independent terms:

$$\mathcal{C}(d|m) := - \sum_{i=1}^N [m_i - d_i \ln(m_i)]. \quad (5.4)$$

The measured data from the signal and background regions are shown in Figure 5.5.




Figure 5.5: COSI 2016 flight spectra in the signal and background regions of the ^{26}Al analysis.

5.3.1 Sky model

In order to construct an absolute spectral response, we simulate multiple potential realizations of the COSI 2016 measurements using the far-infrared Diffuse Infrared Background Experiment (DIRBE) $240\ \mu\text{m}$ map (Hauser et al. 1998) as an image template. We find that the expected number of photons from the signal region between 1750 and 1850 keV is about 41. We therefore generate 50 simulations to obtain sufficient statistics for a smooth sky model spectrum. The flux in this bandpass is heavily dominated by ^{26}Al emission ($\sim 95\%$) and we expect only a $\sim 5\%$ contribution from the Galactic continuum (Wang et al. 2020).

We use the DIRBE $240\ \mu\text{m}$ image because it is a good tracer of Galactic ^{26}Al emission that has been measured by COMPTEL and SPI (Section 2.2.1; Knödlseider et al. 1999; Bouchet et al. 2015). It also does not exhibit the weak artifacts of emission found in the COMPTEL and SPI 1.8 MeV maps which are not easily distinguishable from true ^{26}Al emission (see Plüschke et al. 2001; Bouchet et al. 2015). Furthermore, with the DIRBE $240\ \mu\text{m}$ image we can probe structures of emission finer than those granted by the 3° resolutions of the SPI and COMPTEL images. The Inner Galaxy flux of the DIRBE $240\ \mu\text{m}$ image is normalized to the COMPTEL ^{26}Al Inner Galaxy flux of $3.3 \times 10^{-4}\ \text{ph cm}^{-2}\ \text{s}^{-1}$. The total flux in the image is $1.2 \times 10^{-3}\ \text{ph cm}^{-2}\ \text{s}^{-1}$. The simulated photopeak energy is chosen as the laboratory rest energy of 1808.72 keV. Each of the 50 realizations is simulated in two parts, the first with a 10-detector mass model and the second with a 9-detector mass model, to ensure consistency with the measurements. The transmission probability of γ -rays through the atmosphere is assumed to be constant at the selected flight altitude of 33 km.

Figure 5.6 shows the energy spectrum of events simulated over 50 realizations of the




Figure 5.6: The spectral sky model defined by COSI’s response to the DIRBE 240 μm map (inset image) over 50 2016 flights.

DIRBE 240 μm map which pass the event selections described in Section 5.1. This spectrum defines the sky model. The tailing above 1809 keV is possibly a consequence of increased electronic cross-talk between strips at high energies, which enhances the recorded energy of an event and complicates event reconstruction at high energies. Recall that the strip pairing algorithm fails to reconstruct $\sim 28\%$ of events at 1275 keV (Section 4.5). Applying this same reconstruction check to an ^{26}Al all-sky simulation reveals that $\sim 30\%$ of events at 1809 keV are too complicated to reconstruct. However, this complication does not prohibit ^{26}Al analysis of real flight data because COSI’s complete simulated spectral response is generated using the same event reconstruction algorithm. The complication is thus represented in the sky model and simulations.

5.3.2 Background model

As a data-driven approach to background modeling that draws upon the expectation of concentrated ^{26}Al emission in the Inner Galaxy, we infer a background model from high Galactic latitudes. Recent discussion in the literature about high-latitude emission of ^{26}Al (Pleintinger et al. 2019; Rodgers-Lee et al. 2019) competes with this assumption of concentrated Inner Galactic emission. However, high-latitude emission of ^{26}Al remains unconstrained against the well-established emission from the Inner Galaxy. Additionally, if the high-latitude emission is of extragalactic origin, then it will also be present behind the Inner Galaxy. In that case it is necessary to account for it as background in a measurement of the Inner Galaxy. We therefore proceed with our expectation of dominant Inner Galactic emission. Regions outside of the Inner Galaxy remain valid contributors to our estimation of the background

spectrum. Systematic uncertainties from this assumption are discussed in Section 5.5.2.

We probe the underlying shape of the background spectrum in Figure 5.5 with an empirical fit to data in the background region. For enhanced statistics, these data are considered with minimal event selections compared to those outlined in Section 5.1, limited only to Compton events of incident energy 1750–1850 keV and Compton scattering angles $\phi \leq 90^\circ$. We use a power law plus $N_l = 3$ Gaussian-shaped lines to provide a smooth description of and evaluate uncertainties in the measured background:

$$b(E) = C_0 \left(\frac{E}{E_c} \right)^\gamma + \sum_{l=1}^{N=3} \frac{A_l}{\sqrt{2\pi}\sigma_l} \exp \left(-\frac{1}{2} \left(\frac{E - E_l}{\sigma_l} \right)^2 \right). \quad (5.5)$$

The first term of Equation 5.5 describes the continuum emission from atmospheric background with a power law of amplitude C_0 , pivotal energy $E_c = 1.8$ MeV, and index γ . The three Gaussian-shaped lines l are parameterized by their rates A_l , centroids E_l , and widths σ_l .

The fit of Equation 5.5 to the background spectrum is shown in Figure 5.7 and the fitted parameters are listed in Table 5.3. As a measure of the goodness-of-fit, the reduced- χ^2 is calculated after the maximum likelihood fit and is found to be ~ 1.5 . The Gaussian-shaped lines are due to excitation of materials in the instrument payload which decay on the timescale of the flight. The exact origins of these instrumental lines are uncertain but appear in various other experiments with similar instrument materials (Mahoney et al. 1984; Malet et al. 1991; Naya et al. 1997; Ayre et al. 1984; Boggs and Jean 2000; Weidenspointner et al. 2003). The line near 1764 keV is commonly identified as the decay of natural ^{238}U . The 1779 keV line is likely from the neutron capture process $^{27}\text{Al}(n, \gamma)^{28}\text{Al}$ followed by the 1779 keV γ -ray emission from $^{28}\text{Al}(\beta^-)^{28}\text{Si}$. The line near 1808 keV is likely a blend of activation lines, for example $^{27}\text{Al}(n, np)^{26}\text{Mg}^*$ and $^{26}\text{Na}(\beta^-)^{26}\text{Mg}^*$ which then de-excite to ^{26}Mg . The decay of $^{56}\text{Mn}(\beta^-)^{56}\text{Fe}^*$, which produces a line at 1810.9 keV of similar intensity to the signal 1808.7 keV line in the background spectrum of SPI (Weidenspointner et al. 2003), could also contribute to the blend. The empirical approach to modeling the background attempts to capture these lines, whose centroids differ by less than the instrumental energy resolution. The spectral shapes and uncertainties of the fit shown in Figure 5.7 are then included as normal priors to the simultaneous fit of the background and signal regions, discussed in the next section.

	C_0	γ	A_1	E_1	σ_1	A_2	E_2	σ_2	A_3	E_3	σ_3
Value	2.32	-5.8	2.0	1763.8	3.8	5.2	1779.2	7.1	6.6	1808.0	6.6
Uncertainty	0.03	0.3	0.7	0.8	1.0	0.8	0.6	1.2	0.6	1.0	0.5

Table 5.3: Fit parameters of a power law plus three Gaussian fit to the flight data in the background region with minimal event selections (Figure 5.7). Units: $[C_0] = 10^{-3}$ cnts s^{-1} keV^{-1} , $[A_l] = 10^{-3}$ cnts s^{-1} , $[E_l] = [\sigma_l] = \text{keV}$.




Figure 5.7: Empirical fit to COSI flight data in the background region, with minimal event selections, which provides a smooth description of the background template shape. The fitted parameters are listed in Table 5.3.




Figure 5.8: Posterior distributions of the sky amplitude α and background amplitude β in the COSI 2016 signal region. The green and black lines indicate the median α and β , respectively.

5.3.3 Propagating background uncertainties in a joint fit

We mitigate the potential for bias introduced by the noisy background spectrum in Figure 5.5 by including the spectral features of the fit to the minimally constrained background spectrum (Figure 5.7) in a subsequent, simultaneous fit of the sky and background models. We do not expect the spectral shape of the background to vary significantly during the 46-day flight and allow the complete background model $b(E)$ to vary only within the uncertainties of the parameters from the background region fit (Section 5.3.2). The continuum slope and amplitude are left variable to account for possible continuum emission in the signal region; this procedure only detects γ -ray lines and suppresses any instrumental as well as celestial continuum contribution. We note that the extended Galactic Plane continuum emission from Inverse Compton scattering might readily be visible with COSI in a separate analysis which does not suppress the continuum as background (see continuum emission in Figure 5.3 and further discussion in Section 5.5.3). Thus, by using Equation 5.2, we fit for α and β accounting for the 11 known but uncertain background parameters. The only constraint (prior) for α and β is to be positive definite. The likelihood (Equation 5.3) is therefore used to construct a joint posterior distribution by including the uncertainties in Table 5.3 as normal priors. We use `emcee` (Foreman-Mackey et al. 2013) to estimate the posterior distribution by Monte Carlo sampling. The final fit values of the continuum are $C_0 = (1.13 \pm 0.02) \times 10^{-3}$ cts s^{-1} keV^{-1} and $\gamma = -4.1 \pm 0.6$. This is considerably different from the background-only region (Table 5.3), suggesting that the celestial continuum is absorbed in the background model fit and that COSI can measure the extended Galactic Plane continuum (Section 5.5.3).

As a check of consistency, we compare the amplitudes of the three Gaussian-shaped lines in the empirical fit to the background region data (Figure 5.7; Table 5.3) and the amplitudes returned by this simultaneous fit to the signal region data in Figure 5.5. We call the ~ 1764 keV, ~ 1779 keV, and ~ 1808 keV peak amplitudes A1, A2, and A3, respectively, per the notation in Table 5.3. Normalizing all amplitudes to A1, we find amplitude ratios in the empirical background fit of $A1/A1 \sim 1.0 \pm 0.4$, $A2/A1 \sim 2.6 \pm 0.4$, and $A3/A1 \sim 3.3 \pm 0.3$. Those in the simultaneous fit are $A1/A1 \sim 1.0 \pm 0.4$, $A2/A1 \sim 2.2 \pm 0.5$, and $A3/A1 \sim 2.4 \pm 0.5$. The ratios are consistent within 1σ uncertainties.

5.3.4 Results

5.3.4.1 Signal region

We find an expected dominance of background with best-fit values of $\alpha = 1.1 \pm 0.3$ and $\beta = 28.1 \pm 0.6$ (Figure 5.8). Amplitudes α and β represent the number of photons per keV emitted by the sky and background, respectively. An α value consistent with zero would imply that the signal region data are entirely explained by the background model. Hence, from α we derive a signal-to-noise ratio, as estimated by the best-fit amplitude compared to its uncertainty (negative error value), of $1.1/0.3 \sim 3.7$.

A maximum likelihood ratio calculation (Li and Ma 1983) formalizes the significance of

the measurement above background. This ratio λ is defined as

$$\lambda = \ln L(D|\alpha, \beta) - \ln L(D|\alpha = 0, \beta), \quad (5.6)$$

where $L(D|\alpha, \beta)$ is the likelihood of the simultaneous fit including non-zero sky and background model contributions. The second term, $L(D|\alpha = 0, \beta)$, is the likelihood that the signal region data are explained solely by the background (the null hypothesis). The significance σ of the measurement above background is then calculated as the square-root of the test statistic $TS = 2\lambda$, such that

$$\sigma = \sqrt{TS} = \sqrt{2\lambda}. \quad (5.7)$$

This calculation yields a 3.7σ significance above background measurement of the 1809 keV ^{26}Al peak in the COSI 2016 flight data. Systematic uncertainties do not enter this significance calculation and are summarized in Table 5.6 (Section 5.3.5); the significance ranges only $\sim 2\sigma$ between different model assumptions and as such, the method is dominated by statistical uncertainties. Multiplying the measured rate of 6.8×10^{-4} cts s^{-1} between 1750 and 1850 keV by the exposure time T_{SR} gives ~ 106 ^{26}Al photons. The background rate of 3.0×10^{-4} cts s^{-1} between 1803 and 1817 keV gives ~ 407 background photons.

The background-subtracted spectrum is provided in Figure 5.9. Note that the count rates near the prominent background lines at 1764 and 1779 keV (Figure 5.7) are consistent with zero. This is validation of our background handling method.

5.3.4.2 Line parameters

A summary of line parameters from the COSI 2016 flight is provided in Table 5.4. We use the ratio of fitted ^{26}Al counts in the signal region to the number of ^{26}Al counts expected from DIRBE 240 μm all-sky simulations to calculate COSI's measured ^{26}Al flux. The ratio between the fitted flight counts and simulated counts is ~ 2.6 .

Using atmospheric transmission data from NRLMSISE-00 (NRLMSISE), we find that the response of COSI near 1.8 MeV at 33 km altitude exhibits a sharp decrease in the number of photons beyond a zenith angle of 35–40° (Figure 5.10). As such, we expect COSI to be sensitive to photons out to $\sim 35^\circ$ beyond the specified Inner Galaxy pointing cut. We also defined the maximum Compton scattering angle as 35° (Appendix E). Assuming that the true flux follows the DIRBE 240 μm image, we report a measured COSI 2016 ^{26}Al flux of $(1.70 \pm 0.49) \times 10^{-3}$ ph $\text{cm}^{-2} \text{s}^{-1}$ in the broadened signal region ($|\ell| \leq 65^\circ$, $|b| \leq 45^\circ$). The COSI 2016 measurement of ^{26}Al flux from the Inner Galaxy ($|\ell| \leq 30^\circ$, $|b| \leq 10^\circ$) is $(8.6 \pm 2.5) \times 10^{-4}$ ph $\text{cm}^{-2} \text{s}^{-1}$.

Next, we fit for a shift in the line centroid from the ^{26}Al laboratory energy of 1808.72 keV to probe the dynamics of the emission. Kretschmer et al. (2013) measure a maximum shift of $\sim 300 \text{ km s}^{-1}$, corresponding to $\sim 1.8 \text{ keV}$ at 1809 keV. Including systematic uncertainties from instrument calibrations, the line shift could be at most 3 keV, or $\sim 500 \text{ km s}^{-1}$. To estimate the line centroid in the flight data, we assume that the spectral response within




Figure 5.9: Background-subtracted spectrum and fitted sky models from the COSI 2016 flight. The 1σ contours of the sky models when we fit for line shift ($\Delta E = 2.5 \pm 1.8$ keV) and combined shift and broadening ($\Delta E = 2.9 \pm 1.4$ keV, intrinsic sky broadening < 9.7 keV (2σ upper limit)) are also shown.




Figure 5.10: Zenith response of COSI to 2 MeV photons at a flight altitude of 33 km, indicating strongest sensitivity to photons originating from within $\leq 35\text{--}40^\circ$.

our 1750–1850 keV energy window is constant. We use a spline interpolation of the sky model template and invoke a scale parameter ΔE that shifts the total spectrum along the energy axis. Since at small velocities the Doppler shift is proportional to the difference in centroid energy, ΔE provides a direct measure of the line shift. By including ΔE as a free parameter in our model, we find a shift of $\Delta E = 2.5 \pm 1.8$ keV for a centroid energy of $E_{\text{sky}} = 1811.2 \pm 1.8$ keV, and a line flux in the Inner Galaxy of $(8.8 \pm 2.5) \times 10^{-4}$ ph cm $^{-2}$ s $^{-1}$. The 1σ contour of this shifted sky model is plotted over the background-subtracted spectrum in Figure 5.9.

We also include a free parameter to estimate the broadening of the line. Fitting for both the line shift and broadening, we obtain a shift of $\Delta E = 2.9 \pm 1.4$ keV and a 2σ upper limit on the intrinsic sky broadening of 9.7 keV. The 2σ upper limit on the turbulent velocity of the ^{26}Al ejecta is ~ 2800 km s $^{-1}$. The fit of the total model to the data, with the shifted and broadened sky model, is shown in Figure 5.11. The 1σ contour of the shifted and broadened sky model is also shown in Figure 5.9 and the line flux is enhanced by $\sim 30\%$.




Figure 5.11: Summed (Sky+BG) and individual sky and background models plotted over the flight signal region spectrum. The sky model shown here includes the fitted energy shift and broadening parameters. The medians of the models are shown as solid lines with their 1σ and 2σ uncertainties as shaded contours. Normalized residuals of the fit are displayed in the bottom panel.

5.3.5 Method validation

We repeat the flight data analysis under a variety of assumptions in order to validate the method and define systematic uncertainties (Section 5.5.2). Section 5.3.5.1 tests the method

Line parameter	Value
Measurement significance	3.7σ
Inner Galaxy flux	$(8.6 \pm 2.5) \times 10^{-4} \text{ ph cm}^{-2} \text{ s}^{-1}$
Centroid	$1811.2 \pm 1.8 \text{ keV}$
Intrinsic sky broadening (2σ)	$< 9.7 \text{ keV}$
Turbulent velocity (2σ)	$< 2800 \text{ km s}^{-1}$

Table 5.4: Aluminum-26 line parameters from the COSI 2016 flight. The chosen template map is the DIRBE $240 \mu\text{m}$ image and the quoted uncertainties are statistical.

with the COMPTEL 1.8 MeV and SPI 1.8 MeV images as template maps. The subsequent tests use the DIRBE $240 \mu\text{m}$ image.

5.3.5.1 Different template maps

Using the COMPTEL 1.8 MeV image as a template map instead of the DIRBE $240 \mu\text{m}$ image yields an Inner Galaxy flux of $(6.6 \pm 1.9) \times 10^{-4} \text{ ph cm}^{-2} \text{ s}^{-1}$ with 3.6σ significance. Using the SPI 1.8 MeV image gives $(7.3 \pm 2.1) \times 10^{-4} \text{ ph cm}^{-2} \text{ s}^{-1}$ with 3.7σ significance. The COSI 2016 Inner Galaxy flux values across template maps are therefore consistent with each other within uncertainties.

5.3.5.2 Signal region altitude

As a check on the consistency of our maximum likelihood framework, we repeat the analysis considering flight data in the signal region from decreasing minimum altitudes. We observe an expected decrease in measurement significance as atmospheric background and absorption increase (black points in Figure 5.12). To estimate a spread in the significance, we generate simulated data sets by drawing 25 Poisson samples from the signal region flight spectrum at each altitude. These simulated realizations of the real data contain different numbers of photons, resulting in significance values with some scatter. The mean and standard deviation of these 25 scattered significance values per altitude define the gray 1σ contour in Figure 5.12. The severity of background contamination at balloon altitudes is especially clear, given that the observation time gained by permitting lower altitude observations cannot compensate for the worsening background environment.

We also record the Inner Galaxy flux for each minimum altitude, corresponding to each black point in Figure 5.12. The minimum flux is $(6.8 \pm 2.9) \times 10^{-4} \text{ ph cm}^{-2} \text{ s}^{-1}$ at a minimum altitude of 30 km and the maximum is the $(8.6 \pm 2.5) \times 10^{-4} \text{ ph cm}^{-2} \text{ s}^{-1}$ measurement at a minimum altitude of 33 km in the signal region. The flux values therefore range from $(3.9 - 11.1) \times 10^{-4} \text{ ph cm}^{-2} \text{ s}^{-1}$.




Figure 5.12: Significance above background of the ^{26}Al measurement as a function of minimum flight altitude in the signal region. Black points: significance from flight data. Gray contour: 1σ uncertainties from 25 Poisson samples of the flight data signal region spectrum. Red points: signal region observation time from flight data.

5.3.5.3 Background region altitude

To conform with the event selections of the signal region, we apply a 33 km minimum altitude cut in the background region and repeat the analysis. We measure ^{26}Al with 3.6σ significance above background and find an Inner Galaxy flux of $(8.3 \pm 2.5) \times 10^{-4} \text{ ph cm}^{-2} \text{ s}^{-1}$. This is consistent with the originally presented results.

5.3.5.4 Separate 10-, 9-detector portions

We separate the data from the first half (10-detector portion) and second half (9-detector portion) of the flight and repeat the analysis procedure on each subset. The spectra from each portion of the flight are shown in Figure 5.13 and the posterior distributions of the sky and background amplitudes are shown in Figure 5.14. Using only 10-detector data, we measure ^{26}Al with 2.3σ significance above background and find an Inner Galaxy flux of $(6.8 \pm 3.0) \times 10^{-4} \text{ ph cm}^{-2} \text{ s}^{-1}$. Using only 9-detector data, we find 2.0σ significance above background and an Inner Galaxy flux of $(8.1 \pm 4.1) \times 10^{-4} \text{ ph cm}^{-2} \text{ s}^{-1}$. Within uncertainties, these results are consistent with those of the combined data set. The significance of the measurement in the first part of the flight is slightly greater than that in the second part of the flight because COSI had more exposure to the signal region in the former. Thus, despite the lower background during the second part of the flight (Figure 5.15), we see a stronger signal in the higher background conditions of the first half. Combining the data from both parts of the flight gives the strongest signal.




Figure 5.13: COSI flight spectra in the signal and background regions for separate (a) 10- and (b) 9-detector portions of the flight.

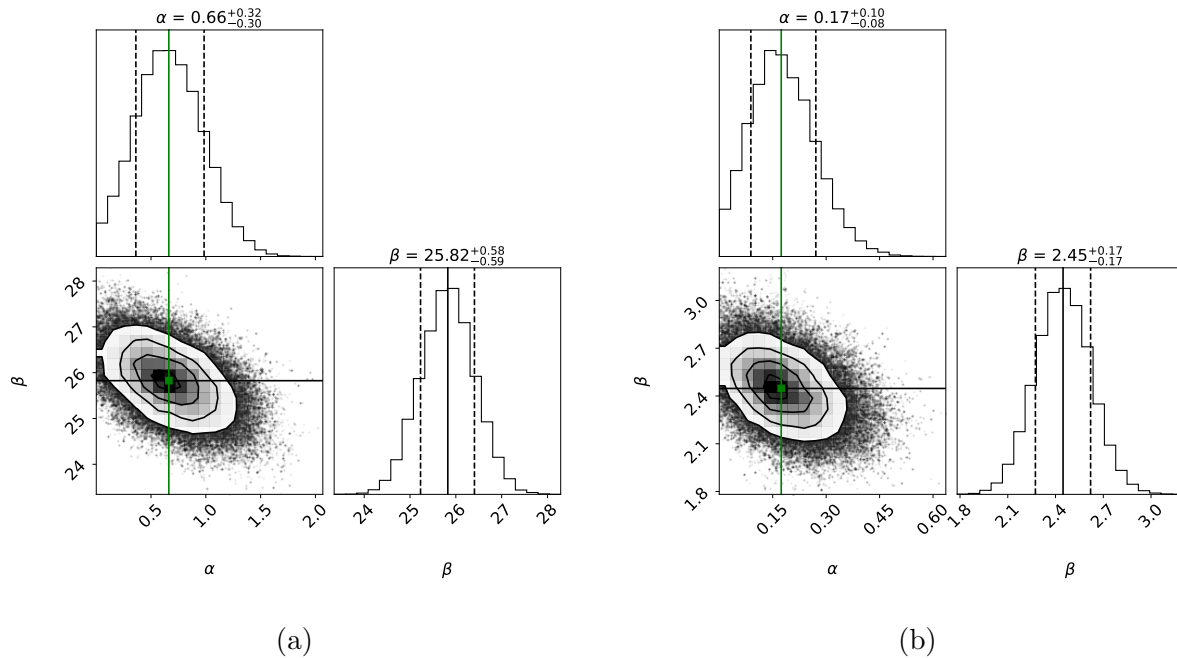


Figure 5.14: Posterior distributions of the sky amplitude α and background amplitude β in the COSI 2016 signal region for separate (a) 10- and (b) 9-detector portions of the flight. The green and black lines indicate the median α and β , respectively.




Figure 5.15: Empirical fit to COSI flight data for separate (a) 10- and (b) 9-detector portions of the flight in the background region, with minimal event selections, which provides a smooth description of the background template shape.

5.3.5.5 Rigidity

In Figures 5.3 and 5.5, we are agnostic to changes in geomagnetic rigidity over the course of the flight. Although the final fit to the flight data accounts for variations in the continuum spectra with changing rigidity, here we manually consider different rigidity regions.

Rigidity R and latitude from Earth's magnetic equator λ are related by

$$R = 14.5 \cos^4(\lambda) / r^2$$

(Smart and Shea 2005) for distance from Earth's dipole center r , regarded here as a constant. As such, to account for rigidity we bin the signal region and background region flight data, each divided between the 10- and 9-detector portions of the flight, into four latitude bins (Figure 5.16). We generate four energy spectra, each corresponding to one latitude bin, in the signal and background regions' 10- and 9-detector portions of the flight, i.e. 16 spectra total. We then re-weight the photon counts in the eight latitude spectra of the background region by the fraction of time COSI observed in the corresponding latitudes of the signal region. After weighting, the four latitude spectra in each of the signal and background data sets are summed to form one energy spectrum, integrated over latitude, and combined over the 10- and 9-detector portions of the flight. Both spectra are normalized by the observation time in each region.

The subtracted spectrum of the signal and weighted background region data is shown in Figure 5.17. After weighting by latitude (and thus rigidity), the 1809 keV signature of ^{26}Al is clearly visible. Some of the line features in the full flight spectrum (Figure 5.3) disappear

Line energy [keV]	Integrated count rate [10^{-4} cnts s^{-1}]	Significance
511	32 ± 11	2.9σ
662	48 ± 42	1.1σ
847	1.4 ± 5.8	0.2σ
1809	8.3 ± 2.1	4.0σ
2223	2.0 ± 1.2	1.7σ

Table 5.5: Line rates and uncertainties after the rigidity-weighted background subtraction.

and the continuum is more suppressed. In particular, the ~ 847 keV line seen in Figure 5.3 is no longer visible. We fit the spectrum to estimate the count rates of the remaining lines (Table 5.5); those of instrumental origin are interpreted as systematic uncertainties in the analysis. The 511 keV significance is smaller than that of 1809 keV because the analysis is optimized to identify the 1809 keV line. Overall, the instrumental lines at 662 keV, 847 keV, and 2223 keV are insignificant compared to 511 keV and 1809 keV.




Figure 5.16: COSI 2016 flight data in the signal (SR) and background (BR) regions from the 10- and 9-detector portions of the flight, binned by Earth latitude. The area under each distribution is normalized to 1.

After weighting by rigidity, we measure the ^{26}Al signal with 3.9σ significance and find an Inner Galaxy flux of $(10.7 \pm 3.0) \times 10^{-4}$ ph $\text{cm}^{-2} \text{s}^{-1}$. This is consistent with previous iterations of the analysis.




Figure 5.17: Background-subtracted COSI 2016 spectrum of the signal and background regions after weighting the flight data by latitude, i.e. geomagnetic rigidity. Error bars are $\sqrt{\text{counts}}$.

5.3.5.6 Broader energy range

To demonstrate that our method can accommodate the continuum background independent of line emission, we expand the energy range of the analysis to 1650–1950 keV. The flight spectra are shown in Figure 5.18a. We simulate the sky model over this new energy range as described in Section 5.3.1 and empirically fit the new background region spectrum with a power law and five Gaussian-shaped lines (Figure 5.19).

We simultaneously fit these new models to the signal region data between 1650–1950 keV and measure the ^{26}Al signal with 4.1σ significance and an Inner Galaxy flux of $(8.9 \pm 2.4) \times 10^{-4} \text{ ph cm}^{-2} \text{ s}^{-1}$. The posterior distributions of the sky and background amplitudes are shown in Figure 5.18b. The slightly higher significance may indicate that by expanding the energy range, we are able to more strongly constrain the continuum in favor of the line signal. The consistency with the results in Sections 5.3.4.1 and 5.3.4.2 is affirmation of our method.

5.3.5.7 Systematic uncertainties in flight data analysis

The results from the previous tests of method validation are summarized in Table 5.6. All Inner Galaxy flux values are consistent with each other within 1σ statistical uncertainties. They range from $(3.8\text{--}13.7) \times 10^{-4} \text{ ph cm}^{-2} \text{ s}^{-1}$, placing a $\sim 57\%$ systematic uncertainty on the $(8.6 \pm 2.5) \times 10^{-4} \text{ ph cm}^{-2} \text{ s}^{-1}$ measurement reported in Section 5.3.4.2. Instrumental lines of less than 2σ significance (Table 5.5) indicate that the instrumental background is noticeably, if imperfectly, suppressed compared to lines of interest. Additional considerations of systematic uncertainties are derived from simulations in Section 5.4 and a cumulative discussion of these uncertainties is presented in Section 5.5.2.

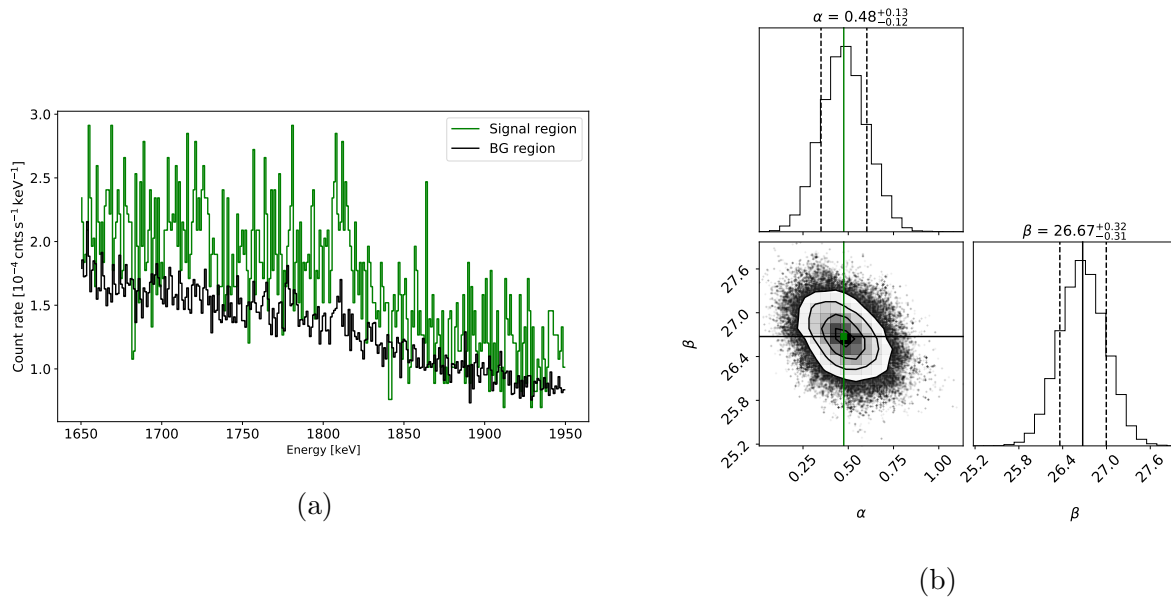


Figure 5.18: (a) COSI flight spectra in the signal and background regions between 1650–1950 keV. (b) Posterior distributions of the sky amplitude α and background amplitude β in this expanded energy range. The green and black lines indicate the median α and β , respectively.




Figure 5.19: Empirical fit to COSI flight data in the background region, with minimal event selections, using an expanded energy range of 1650–1950 keV.

Test	Measurement significance	Inner Galaxy flux [10 ⁻⁴ ph cm ⁻² s ⁻¹]
COMPTEL 1.8 MeV	3.6σ	6.6 ± 1.9
SPI 1.8 MeV	3.7σ	7.3 ± 2.1
M.A. Signal	2.4–3.6σ	3.9–11.1
M.A. Background	3.7σ	8.3 ± 2.5
Only 10-detector data	2.3σ	6.8 ± 3.0
Only 9-detector data	2.0σ	8.1 ± 4.1
Rigidity	3.9σ	10.7 ± 3.0
1650–1950 keV	4.1σ	8.9 ± 2.4

Table 5.6: Summary of flight data results from various tests of method validation (Section 5.3.5). “M.A. Signal:” Minimum 27–33 km altitudes in the signal region. “M.A. Background:” Minimum 33 km altitude in the background region.

5.3.6 Flux implications

The Inner Galaxy flux measurements of ²⁶Al reported in the previous sections are calculated by scaling the DIRBE 240 μm flux in the Inner Galaxy, f_{IG} , by the ratio of total fitted counts, Y , to the number of expected ²⁶Al photons, X , in the Inner Galaxy. The flux f_{IG} of the DIRBE 240 μm map is 3.3×10^{-4} ph cm⁻² s⁻¹. We find $Y \sim 106$ ²⁶Al photons; in Section 5.3.4, this is calculated as the product of the measured ²⁶Al rate and the signal region observation time $T_{\text{SR}} \approx 156$ ks. Equivalently, $Y \sim 106$ is obtained by integrating the normalized sky model over its energy range and multiplying the resulting number of counts by the fitted sky amplitude α . A simulation of the COSI 2016 flight over the DIRBE 240 μm map gives the number of expected ²⁶Al photons, X , in the Inner Galaxy between 1750–1850 keV. We find $X = 41$ (Section 5.3.1). Hence, we write that the measured flux M is

$$M = \frac{Y}{X} \cdot f_{\text{IG}} \quad (5.8)$$

and $M = (8.6 \pm 2.5) \times 10^{-4}$ ph cm⁻² s⁻¹ (Table 5.4). This measured flux of ²⁶Al can be used to estimate the mass of ²⁶Al in the Inner Galaxy, the rate of positron production from its decay, and the rate of core-collapse supernova events which may have ejected the measured mass of ²⁶Al.

5.3.6.1 Mass of Aluminum-26

The mass of ²⁶Al in the Inner Galaxy can be estimated from the reported flux. For distance R_0 between COSI and the source of ²⁶Al, we convert the flux to a luminosity L :

$$L = M \cdot 4\pi R_0^2, \quad (5.9)$$

where R_0 is approximated as the distance from Earth to the center of the Galaxy. Here, we adopt the distance to the Galactic Center as $R_0 = 8178 \pm 13_{\text{stat}} \pm 22_{\text{sys}}$ pc, found by tracking the orbit of star S2 around Sagittarius A* (The GRAVITY Collaboration et al. 2019). Following the approach of Siegert (2017), we divide L by the luminosity of one solar mass of ^{26}Al . This luminosity (L_\odot) is calculated as the product of the probability per unit time that ^{26}Al decays ($\lambda(^{26}\text{Al})$, i.e. the decay constant), the probability that this decay emits the desired 1809 keV γ -ray (p_{1809}^{26}), and the number of ^{26}Al atoms in one solar mass (N_\odot^{26}):

$$L_\odot = \lambda(^{26}\text{Al})p_{1809}^{26}N_\odot^{26} \quad (5.10)$$

$$= \frac{\ln 2}{T_{1/2}(^{26}\text{Al})}p_{1809}^{26}\frac{M_\odot}{M(^{26}\text{Al})}. \quad (5.11)$$

The half-life $T_{1/2}(^{26}\text{Al})$ of ^{26}Al is 7.17×10^5 years, the probability p_{1809}^{26} of 1809 keV emission is 99.7%, the solar mass $M_\odot = 1.99 \times 10^{30}$ kg, and the atomic mass of one atom $M(^{26}\text{Al})$ is $25.98\text{u} \approx 4.32 \times 10^{-26}$ kg. These quantities give $L_\odot = 1.41 \times 10^{42}$ ph s^{-1} (Siegert 2017). Dividing L in Equation 5.9 by L_\odot yields an Inner Galactic mass of ^{26}Al of $M_{\text{IG}} = 4.8 \pm 1.4 M_\odot$. COMPTEL found a total Galactic mass of approximately $2.5 M_\odot$, depending on models of spiral arm structure, disk structure, and free-electron models (Knödlseider et al. 1996). SPI found a similar total Galactic mass of $2.8 \pm 0.8 M_\odot$ (Diehl et al. 2006). The COSI mass estimate exceeds these previous findings because its flux measurement is nearly double that of COMPTEL and SPI (see Section 5.5 for further discussion).

5.3.6.2 Positron production rate

The positron yield from the measured ^{26}Al flux is the product of the number of ^{26}Al atoms in the above-calculated mass and the probability that these atoms emit a positron upon decay ($p_{e^+}^{26}$). The positron production rate over 1 Myr is then written as:

$$e^+ = \frac{M_{\text{IG}}}{M(^{26}\text{Al})}p_{e^+}^{26}\frac{1}{1 \text{ Myr}}. \quad (5.12)$$

The probability of positron emission $p_{e^+}^{26}$ is 81.73% (Siegert 2017). It follows that the calculated positron production rate from ^{26}Al in the Inner Galaxy is $(5.7 \pm 1.7) \times 10^{42} e^+ \text{ s}^{-1}$. Siegert (2017) estimates $(3.35 \pm 0.07) \times 10^{42} e^+ \text{ s}^{-1}$ in the Galaxy from SPI observations.

5.3.6.3 Core-collapse supernova rate

Finally, we estimate the core-collapse supernova rate from the observed mass of ^{26}Al . Dividing the COSI Inner Galaxy mass of ^{26}Al by the estimated quantity of ^{26}Al ejected per CCSN event gives the number of ^{26}Al -producing stars which explode as a CCSN. The latter quantity

is the initial mass function-averaged ejected mass of ^{26}Al per star: $\langle Y(^{26}\text{Al}) \rangle = 1.4 \times 10^{-4} M_{\odot}$ (Diehl et al. 2006). Dividing by the lifetime of ^{26}Al ($\tau(^{26}\text{Al})$), i.e. the timescale on which we may detect such supernovae) gives a CCSN rate (Siegert 2017):

$$R_{\text{SN}} = \frac{M_{\text{IG}}}{\tau(^{26}\text{Al}) \langle Y(^{26}\text{Al}) \rangle}. \quad (5.13)$$

For $\tau(^{26}\text{Al}) = 1.04 \text{ Myr}$, we obtain a CCSN rate of $R_{\text{SN}} = 3.3 \pm 1.0 \text{ CCSNe century}^{-1}$. SPI observations indicate a CCSN rate of $1.9 \pm 1.1 \text{ CCSNe century}^{-1}$ in the Galaxy (Diehl et al. 2006). A Galaxy-wide population synthesis model suggests a SN rate in the Milky Way of $1.8\text{--}2.8 \text{ century}^{-1}$ (Siegert et al. 2023). Again, the Inner Galaxy COSI result exceeds the full-Galaxy estimates because the COSI Inner Galaxy mass is greater than previous results in the literature.

5.4 Validating the method with simulations

To further validate our method and results, the analysis outlined above is repeated on purely simulated data sets using four different template maps to model the ^{26}Al signal: DIRBE $240 \mu\text{m}$, SPI 1.8 MeV , COMPTEL 1.8 MeV , and ROSAT 0.25 keV (Snowden et al. 1997). The latter is included as a map which traces high-latitude rather than Galactic Plane emission, and serves as a test of the sensitivity of our method. We develop a simulated background model (Section 5.4.2) and simulate COSI 2016 flights at different flux levels above this same background. We cross-check our results with statistical expectations (Section 5.4.4). Finally, in Section 5.4.5, we perform an analysis on a data set comprised entirely of simulated background as a measure of systematic uncertainty and validation of the real signal significance.

5.4.1 Simulated data sets

The simulations of the template maps are conducted assuming a constant transmission probability of $\sim 69.5\%$ at zenith (Figure 5.10), corresponding to a flight altitude of 33 km . The 10- and 9-detector portions of the flight are simulated separately with appropriate mass models. These simulations are run using MEGAlib’s *cosima* (Section 3.6.1). The template map simulations are combined with a *cosima* simulation of instrumental activation over 46 days of cosmic ray and atmospheric particle irradiation (Zoglauer et al. 2008) and an atmospheric (photonic) background model to account for the Earth albedo (Ling 1975). We scale the level of our background simulations to the best possible match with our flight observations. We maintain the spectral shape of the simulated background. The activation and atmospheric simulations together comprise the total simulated background and are discussed in more detail in Appendix F. We apply the same pointing cuts and event selections from the flight data (Tables 5.1 and 5.2) to the combined signal and background simulations. This yields representative realizations of the COSI 2016 flight in the signal and background regions with a response to different ^{26}Al tracers.




Figure 5.20: Combined simulations of one 2016 flight over the DIRBE $240 \mu\text{m}$ template image, instrumental activation background, and atmospheric background in the signal and background regions, similar to Figure 5.5.

5.4.2 Complete flight simulation

The simulated spectra in the signal and background regions of the DIRBE $240 \mu\text{m}$ template image are shown in Figure 5.20. These simulated spectra are similar to the flight spectra in Figure 5.5, suggesting a sufficiently accurate description of the data. The background model is informed by applying minimal event selections to the combined activation and atmospheric simulations and fitting them with a power law and three Gaussian-shaped lines (Equation 5.5). This procedure is analogous to that with real flight data in Section 5.3.2. The simulated spectrum and fit parameters are shown in Figure 5.21 and Table 5.7.

	C_0	γ	A_1	E_1	σ_1	A_2	E_2	σ_2	A_3	E_3	σ_3
Value	2.69	-3.7	3.2	1778.4	2.4	2.7	1784.1	6.7	0.6	1808.5	2.0
Uncertainty	0.01	0.1	0.3	0.1	0.2	0.4	1.2	0.8	0.1	0.3	0.5

Table 5.7: Fit parameters of a power law plus three Gaussian fit to the instrumental activation and atmospheric background simulations in the background region with minimal event selections (Figure 5.21). Units: $[C_0] = 10^{-3} \text{ cnts s}^{-1} \text{ keV}^{-1}$, $[A_l] = 10^{-3} \text{ cnts s}^{-1}$, $[E_l] = [\sigma_l] = \text{keV}$.

The largest differences between the real flight and simulated background spectra are the count rates of the 1764 and 1779 keV lines: The 1764 keV line is prominently seen in the flight data, yet the activation simulation appears to show no 1764 keV line at all. This may be expected, however, because it is likely a line originating from the natural




Figure 5.21: Power law plus three Gaussian empirical fit to the instrumental activation and atmospheric background simulations with minimal event selections, similar to Figure 5.7. The parameters of the fit are listed in Table 5.7.

^{238}U decay series, i.e. it is not due to local activation by cosmic rays (Appendix F.1). The simulated 1779 keV line appears as a blend of two lines at 1778 and 1784 keV. The slope of the background continuum is less steep with $\gamma_{\text{sim}} \sim -3.7$ compared to $\gamma_{\text{flight}} \sim -5.8$. These differences motivate our empirical approach in the analysis of real flight data and underscore the difficulty of modeling the MeV background in a balloon environment. As with the real flight data, the fitted spectral parameters of the simulated background and its uncertainties are fed as normal priors to a simultaneous fit of the sky and background models to the simulated signal region spectrum.

The best-fit sky amplitude $\alpha = 0.7 \pm 0.3$ and the background amplitude $\beta = 28.7 \pm 0.6$. The signal-to-noise ratio is $0.7/0.3 \sim 2.3$. We note that this is less than the measured signal-to-noise ratio of ~ 3.7 in the real flight data. We calculate a 2.8σ significance over the background compared to 3.7σ significance in the flight data. The simulated signal rate between 1750 and 1850 keV is $4.5 \times 10^{-4} \text{ cnts s}^{-1}$, corresponding to ~ 70 ^{26}Al photons. The simulated background region rate between 1803 and 1817 keV is $2.9 \times 10^{-4} \text{ cnts s}^{-1}$, corresponding to ~ 392 background photons. The simulated and flight background counts are comparable to those in the real flight data and the simulated sky photons are lower by a factor of ~ 1.5 . This difference suggests a systematic uncertainty in the absolute calibration of COSI’s effective area (see Section 5.5.2).

We plot the fitted total, sky, and background models for this simulation in Figure 5.22. The background-subtracted spectrum is shown in Figure 5.23. The estimated ^{26}Al Inner




Figure 5.22: Summed (Sky+BG) and individual sky and background models plotted over the signal region spectrum in the complete flight simulation, similar to Figure 5.11. Energy shift and broadening parameters are not considered in this figure, as we do not expect these astrophysical effects in the simulations.

Galaxy flux from this simulated data set is $(2.4 \pm 1.0) \times 10^{-4} \text{ ph cm}^{-2} \text{ s}^{-1}$. Within uncertainties, this flux appears to be about 1.8 times smaller than that of the flight data. We also see a similar factor in the significance estimate, again suggesting a systematic offset.

As with the flight data (Section 5.3.4.2), we fit for an energy shift in the line. We expect ΔE to be consistent with zero because the simulations do not include the intrinsic (astrophysical) broadening of the sky seen in real flight data. Indeed we find a shift of $\Delta E = -0.2 \pm 2.2 \text{ keV}$ and the Inner Galaxy flux is unchanged. Including free parameters for shifting and broadening gives a shift of $\Delta E = 1.5 \pm 1.7 \text{ keV}$ and a 2σ upper limit on the intrinsic sky broadening of 13.7 keV. The 1σ contours of these shifted and broadened sky models are shown in Figure 5.23.

5.4.3 Simulations with different template maps

We repeat the analysis of Section 5.4.2 using the SPI 1.8 MeV, COMPTEL 1.8 MeV, and ROSAT 0.25 keV images as template maps. Comparing the results across multiple template maps is both a check of the flight data measurement and a check of the consistency of our analysis pipeline.




Figure 5.23: Background-subtracted spectrum and 1σ sky model contours from the complete flight simulation, similar to Figure 5.9. Fitting for an energy shift gives $\Delta E = -0.2 \pm 2.2$ keV. Fitting for line shift and broadening gives $\Delta E = 1.5 \pm 1.7$ keV and an intrinsic sky broadening < 13.7 keV (2σ upper limit).

Template map	Significance [σ]	Measured IG flux	Map IG flux	Sky amp. α	BG amp. β
DIRBE 240 μm	2.8 ± 0.5	2.5 ± 0.4	3.3	0.7 ± 0.1	28.7 ± 0.1
SPI 1.8 MeV	2.8 ± 0.4	1.9 ± 0.3	2.7	0.8 ± 0.1	28.8 ± 0.1
COMPTEL 1.8 MeV	3.2 ± 0.5	2.5 ± 0.4	3.3	0.9 ± 0.1	28.8 ± 0.1
ROSAT 0.25 keV	—	0.2 ± 0.1	0.3	0.3 ± 0.1	28.7 ± 0.1

Table 5.8: Summary of results from complete flight simulations of various tracer maps. The listed values are the mean significance above background, measured ^{26}Al Inner Galaxy (IG) flux [10^{-4} ph cm^{-2} s^{-1}], true simulated ^{26}Al IG flux [10^{-4} ph cm^{-2} s^{-1}], sky amplitude α , and background amplitude β over 50 independent complete flight simulations of each tested map (Section 5.4.3).

Table 5.8 shows the signal significance, measured ^{26}Al Inner Galaxy flux, true ^{26}Al Inner Galaxy flux in the template map, and the best-fit α and β averaged over 50 independent realizations of flight simulations per template map. We find that the DIRBE, SPI, and COMPTEL template maps return Inner Galaxy fluxes consistent within two standard deviations of the true expected values. The ROSAT map, which is not a tracer of ^{26}Al given its dearth of emission in the Inner Galaxy, yields a flux measurement nearly consistent with zero, as expected. This is affirmation of the null hypothesis: the likelihood that COSI’s signal region emission is traced by the ROSAT map is accounted for entirely by the background model ($\alpha \approx 0$, $\beta > 0$).

The analysis pipeline underestimates the ^{26}Al flux in the Inner Galaxy of each template map by about a factor of 1.5. This is probably due to the fact that the high-latitude emission in the template maps is significantly different from zero. The background model then absorbs some portion (10–30%) of the total flux outside of the Inner Galaxy. In addition to the absolute effective area calibration, this value can be considered a systematic uncertainty in our method’s definition of all emission outside of the Inner Galaxy as background (see Section 5.5.2 for further discussion). More comprehensive observations of the ^{26}Al sky are necessary to constrain high-latitude emission and the resulting uncertainty.

5.4.4 Increasing the signal

To assess the validity of our simulation, we conduct additional iterations of the analysis outlined above by simulating different flux levels of the sky above the simulated background. To obtain an objective measure that our method works, we increase the flux in our simulations while keeping the background level constant. That is, we pick at random n out of 50 sky simulations and perform the same analysis as above to benchmark the simulation results against expectations.

For each case, we run 25 realizations by randomly selecting n out of 50 simulations. The background in each case is the simulated instrumental activation and atmospheric background described in Section 5.4.2. Figure 5.24 shows the estimated significance against the estimated flux for the DIRBE, SPI, and COMPTEL maps. We find the expected square-root-like behavior of increasing flux or, equivalently, exposure time.

As expected, using the ROSAT map as a template of ^{26}Al emission does not yield estimates of significant positive excess above the background. This is further validation of the pipeline because the ROSAT map shows strong emission only at high latitudes.

5.4.5 Background-only simulations

Finally, we repeat the analysis on simulated data sets devoid of any sky signal. In this way, we obtain a distribution of test statistic (TS) values that follows a χ^2 -distribution with one degree of freedom, i.e. $\alpha = 0$ versus $\alpha \neq 0$ (Wilks’ theorem; Wilks 1938). We fit the background region spectrum from the flight data (Section 5.3.2) 1000 times. In each iteration,




Figure 5.24: Significance vs. estimated Inner Galaxy flux for simulated data sets containing n DIRBE, SPI, and COMPTEL simulations of the flight combined with activation and atmospheric background simulations. The analysis is performed 25 times per n simulations, creating the scatter of different realizations.

we define the signal region spectrum as a Poisson sample of the flight data background model defined by the fit parameters describing the background spectrum.

Figure 5.25 demonstrates that the TS indeed follows a χ_1^2 -distribution. The 3.7σ (equivalent to p value = 0.00022) measurement from the real flight analysis clearly exceeds the significance returned by 1000 assumptions of the null hypothesis. Thus, we verify that the TS calculated in our analysis method is a reliable proxy of the likelihood that the flight data d are described by our model description m .

5.5 Discussion

5.5.1 Comparison to previous measurements

Depending on the template map used, we find an ^{26}Al flux in the Inner Galaxy between $4.7 \times 10^{-4} \text{ ph cm}^{-2} \text{ s}^{-1}$ (-1σ uncertainty from the COMPTEL map) and $11.1 \times 10^{-4} \text{ ph cm}^{-2} \text{ s}^{-1}$ ($+1\sigma$ uncertainty from the DIRBE map). Our measured flux is consistent with previous measurements from SPI and COMPTEL of $2.8\text{--}3.3 \times 10^{-4} \text{ ph cm}^{-2} \text{ s}^{-1}$ within 2σ uncertainties. We find a line centroid of $1811.2 \pm 1.8 \text{ keV}$ using the DIRBE, SPI, and COMPTEL template maps. This is consistent with previous measurements and in particular with the laboratory energy of 1808.72 keV within 2σ uncertainties. While SPI measured a Doppler shift of




Figure 5.25: Distribution of the test statistic (TS) from 1000 simulated data sets. The signal in each is defined as a Poisson realization of the fitted background model from flight data.

1809.02 ± 0.04 keV in the Inner Galaxy (Siegert 2017), the systematic uncertainties in these measurements due to calibration, detector degradation, and line shape are about one order of magnitude larger than the statistical uncertainties. We repeat the COSI flight analysis in Section 5.3.4.2 with the line shift fixed to 0.3 keV (to the SPI centroid of 1809.02 ± 0.04 keV). This gives an Inner Galaxy flux of $(9.9 \pm 2.8) \times 10^{-4} \text{ ph cm}^{-2} \text{ s}^{-1}$, which is fully consistent with the results when the line shift is left as a free parameter. Overall, the absolute line shift in the Inner Galaxy is difficult to model because individual stellar groups, the large-scale Galactic rotation, and preferential streaming directed along Galactic rotation (Section 2.3; e.g., Kretschmer et al. 2013) all contribute to the total line shift.

Our line width places a 2σ upper limit on the turbulent motion of ^{26}Al ejecta in the Inner Galaxy of $\lesssim 2800 \text{ km s}^{-1}$. Accounting for the large-scale motion as measured in Kretschmer et al. (2013), the intrinsic velocity broadening is limited to $\lesssim 2400 \text{ km s}^{-1}$. This is about one order of magnitude greater than the expected turbulent motion of hot gas in the ISM, where a line width of 1 keV corresponds to a velocity of 122 km s^{-1} (Diehl et al. 2006; Wang et al. 2009). Measuring the broadening is a difficult task, despite the excellent energy resolution of germanium detectors. Adding an instrumental resolution of $\sim 3 \text{ keV}$ at 1809 keV in quadrature with an intrinsic sky broadening of 1 keV, for example, gives a measured line width of $\sim 3.2 \text{ keV}$. The measured width in this toy example is only $\sim 7\%$ larger than the instrumental resolution, even though the intrinsic sky broadening is 33% as wide as the instrumental resolution.

A measurement of the Galaxy-wide Doppler broadening of the 1.8 MeV emission also remains an open issue because measuring the broadening, rather than the shift, requires considerably longer integration times. Detectors degrade over these long integration times,

changing the instrumental line response and complicating the analysis. However, as a satellite mission, COSI’s enhanced line sensitivity of 1.7×10^{-6} ph cm⁻² s⁻¹ at 1809 keV (3σ over 24-month survey; Tomsick et al. 2019) compared to SPI may expedite a Doppler broadening measurement of the ²⁶Al line. Additionally, the satellite’s improved angular resolution of 1.5° (Tomsick et al. 2019) has potential to advance explorations of ²⁶Al dynamics (Krause et al. 2015; Fujimoto et al. 2020) and those of recently created elements (Forbes et al. 2021).

5.5.2 Systematic uncertainties

The ²⁶Al flux value measured in the COSI 2016 flight is approximately two times greater than expected. This enhancement is similar to that seen in analyses of the 511 keV positron annihilation line during the COSI flight (Kierans et al. 2020; Siegert et al. 2020). Applying this systematic factor to the ²⁶Al measurement gives an Inner Galaxy flux of $(4.3 \pm 1.3) \times 10^{-4}$ ph cm⁻² s⁻¹, consistent with previous measurements from SPI and COMPTEL. Thus, we see a systematic uncertainty on the overall flux normalization of $\sim 50\%$, probably owing to the absolute calibration of the effective area, independent of energy. This uncertainty may also be attributed to possible imperfections in the atmospheric model assumed by MEGALib when simulating COSI’s spectral sky model at a minimum altitude of 33 km. Repeating the flight data analysis under a variety of conditions (Section 5.3.5) also indicates a systematic uncertainty on the flux of $\sim 57\%$.

Additional systematics arise from the analysis method itself. Our approach relies on the assumption that at high latitudes ($|b| \gtrsim 45^\circ$) and longitudes ($|\ell| \gtrsim 105^\circ$), the sky is devoid of any ²⁶Al signal. The template maps used for the signal (DIRBE 240 μ m, SPI 1.8 MeV, and COMPTEL 1.8 MeV) all show a non-zero contribution in these background regions. While we can estimate the flux contribution from regions like Orion, Perseus, Taurus, Carina, or Vela from previous studies to account for at most 15% of the total ²⁶Al emission (see Bouchet et al. 2015; Siegert 2017; Pleintinger 2020), the emission at high latitudes is essentially unknown. The COMPTEL map shows nearly homogeneous diffuse emission at these latitudes, which is likely residual emission from the reconstruction algorithm. Likewise, the SPI 1.8 MeV image shows one particularly bright spot at $(\ell, b) = (226^\circ, 76^\circ)$, which is almost certainly an artifact in the image reconstruction because no ²⁶Al source is known at this position with a flux of $5\text{--}9 \times 10^{-5}$ ph cm⁻² s⁻¹ (Bouchet et al. 2015). Finally, because the DIRBE 240 μ m map performs well in a fit to raw γ -ray data from SPI and COMPTEL, it only traces, rather than shows directly, the true distribution of ²⁶Al. We estimate the systematic uncertainties in the template map as 10–30%, given the DIRBE 240 μ m simulated flux of $(2.5 \pm 0.4) \times 10^{-4}$ ph cm⁻² s⁻¹ compared to the true map flux of 3.3×10^{-4} ph cm⁻² s⁻¹ (Table 5.8).

We perform an additional check of this systematic by modifying the DIRBE 240 μ m template image to contain zero flux outside of the 35°-broadened Inner Galaxy ($|\ell| \leq 65^\circ$, $|b| \leq 45^\circ$) and repeating the flight data analysis. This artificial map, which contains ²⁶Al only in the broadened signal region, yields an Inner Galaxy flux of $(9.3 \pm 2.7) \times 10^{-4}$ ph cm⁻² s⁻¹. The enhanced flux confirms that defining unconstrained emission of ²⁶Al at higher latitudes

as background introduces systematic uncertainty. We also note that its consistency with the flight measurement of $(8.6 \pm 2.5) \times 10^{-4} \text{ ph cm}^{-2} \text{ s}^{-1}$ is validation of the claim that COSI is sensitive to photons $\sim 35^\circ$ beyond the Inner Galaxy.

This test may clarify the factor of ~ 1.5 seen in Section 5.4 and clearly illustrates the need to constrain this systematic with a more detailed description of ^{26}Al across the entire sky. The improved sensitivity of CCTs to shallow emission gradients and isotropic emission compared to that of coded mask instruments like SPI, combined with better spectral resolution compared to NaI scintillators (e.g., COMPTEL), make imaging high-latitude emission an achievable goal for the COSI satellite. Constrained high-latitude emission will provide valuable insight to the open problem of the true ^{26}Al morphology in the Milky Way (Pleintinger et al. 2019).

5.5.3 Galactic diffuse continuum emission

The Galactic Plane exhibits Galactic diffuse continuum emission (GDCE) from inverse Compton (IC) scattering of MeV and GeV electrons (Porter et al. 2008). Emission from IC scattering is expected to be brightest along the Galactic Plane and Galactic Center, with sub-dominant contributions at higher Galactic latitudes; it is observationally more difficult to detect the emission at high latitudes because of lower signal-to-noise. As a similar dichotomy between low and high Galactic latitudes is central to the analysis of this chapter, ^{26}Al emission is an interesting environment within which to test COSI's sensitivity to the GDCE.

The full flight spectrum in Figure 5.3 has an underlying power law shape which is likely continuum emission from the sky. If there were no detected continuum, the spectrum would appear with a flatter underlying spectrum and more prominent γ -ray lines. Figure 5.26a shows predicted count rates from MEGALib simulations of the GDCE during the 2016 flight, assuming three different GALPROP models of the underlying emission. The simulated rates compare favorably to the measured rates shown in the rigidity-corrected, background-subtracted spectrum (black points, taken from Figure 5.17) of the ^{26}Al analysis, especially above $\sim 1 \text{ MeV}$.

The rates are converted to flux in Figure 5.26b and show notable agreement with INTEGRAL measurements, save a systematically higher flux in COSI below $\sim 1 \text{ MeV}$. The COSI measurements include point source contributions, while the INTEGRAL measurements are of the total diffuse emission with point sources subtracted. The INTEGRAL data thus show the GDCE as comprised of IC emission, positron-electron annihilation at 511 keV and the lower energy ortho-positronium continuum, the ^{26}Al line at 1809 keV, and unresolved point sources below 100 keV. Importantly, the COSI measurements (light blue) do not consider the scattered component of the atmospheric response. A preliminary correction that accounts for atmospheric scattering is shown in dark blue to bring the COSI measurements in closer alignment with INTEGRAL. Additionally, we note that the COSI measurements are shown only with 1σ statistical uncertainties. Future work will define the systematic uncertainties of this measurement for more thorough comparison to the INTEGRAL data which have

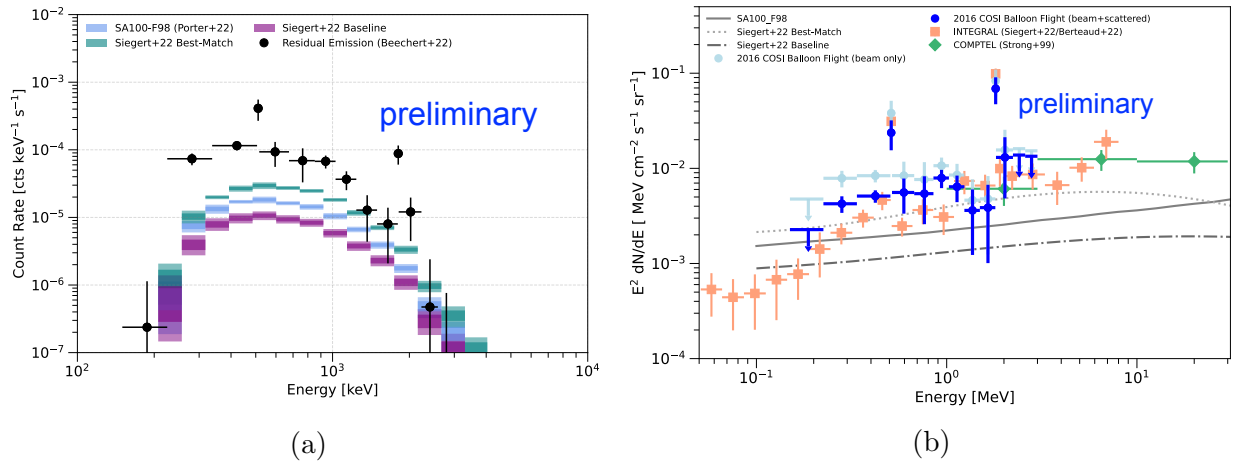


Figure 5.26: The Galactic diffuse continuum emission (GDCE) in COSI 2016 balloon flight measurements. (a) The rates extracted from the ^{26}Al analysis (black points, taken from Figure 5.17) compare favorably with simulated predictions of flight measurements assuming three different GALPROP models (color bands) of the GDCE. The dark and light regions of each color band show the 1σ and 2σ statistical uncertainties. (b) COSI’s background-subtracted measured flux towards the Inner Galaxy (light blue), shown with 1σ statistical uncertainties, agrees well with INTEGRAL measurements (orange) above ~ 1 MeV. A preliminary atmospheric correction to the COSI measurements (dark blue) mitigates some of the discrepancy below ~ 1 MeV, which may also be due in part to the fact that the COSI measurements do not have point source contributions removed, while the INTEGRAL measurements (statistical and systematic uncertainties shown) are of the total diffuse emission with point sources subtracted. GALPROP models are shown in gray. Images from Karwin et al. (2023, in prep.).

both statistical and systematic uncertainties included. Overall, however, Figure 5.26 contains promising evidence that the continuum observed in the background-subtracted spectra of the COSI balloon ^{26}Al analysis reflects that of the GDCE. Estimations that consider the point source contribution and correct for atmospheric energy dispersion in COSI are underway and may help reconcile the discrepancies below ~ 1 MeV. For a complete discussion of Galactic diffuse continuum emission in the COSI 2016 balloon flight, refer to Karwin et al. (2023, in prep.).

We derive additional confidence that the continuum in the background-subtracted spectra of Figures 5.3 and 5.17 reflects the sky continuum in the signal region (Inner Galaxy, i.e. along the Galactic Plane) from the ROSAT 0.25 keV test in Section 5.4.3. The null detection of ^{26}Al as traced by the high-latitude emission of the ROSAT 0.25 keV map verifies that the pointing cuts which define the background region in this analysis do in fact select data from high latitudes, which again are expected to contain weaker continuum emission than

observed along the Galactic Plane. Hence, agreement between observations and simulations indicates that COSI can observe the expected continuum emission from the Galactic Plane.

The analysis in Wang et al. (2020) also points to continuum emission in the COSI data. Figure 5 of Wang et al. (2020) (Figure 2.6 in this document) shows $\sim 5\%$ contribution from the continuum in the vicinity of the 1809 keV γ -ray line from ^{26}Al : the flux of the former is $\sim 1 \times 10^{-5}$ cnts $\text{cm}^{-2} \text{s}^{-1} \text{keV}^{-1}$ and that of the latter is $\sim 2 \times 10^{-4}$ cnts $\text{cm}^{-2} \text{s}^{-1} \text{keV}^{-1}$. The analogous rates in Figure 5.17 are similar, with a flux of $\sim 10^{-5}$ ph $\text{cm}^{-2} \text{s}^{-1}$ in the energy bins adjacent to the flux of $\sim 10^{-4}$ ph $\text{cm}^{-2} \text{s}^{-1}$ in the 1809 keV γ -ray bin.

Moreover, the analysis in Wang et al. (2020) suggests that the continuum in COSI's background-region spectra traces instrumental continuum rather than continuum from the Galactic Plane. The continuum in Wang et al. (2020) is characterized by a power law of $(2.4 \pm 0.2) \times (E/1000 \text{ keV})^{-1.3 \pm 0.2} \times 10^{-5}$ cnts $\text{cm}^{-2} \text{s}^{-1} \text{keV}^{-1}$, while the fitted continuum in COSI's background region data (Figure 5.7 and Table 5.3) is $(2.32 \pm 0.03) \times (E/1809 \text{ keV})^{-5.8 \pm 0.3} \times 10^{-3}$ cnts $\text{cm}^{-2} \text{s}^{-1} \text{keV}^{-1}$. The significant deviation in power law parameters points to a different origin, likely instrumental, of the continuum emission in COSI's background region.

5.6 Summary

We report a 3.7σ measurement of Galactic ^{26}Al in the COSI 2016 balloon flight. The Inner Galaxy ($|\ell| \leq 30^\circ$, $|b| \leq 10^\circ$) flux is estimated as $(8.6 \pm 2.5_{\text{stat}} \pm 4.9_{\text{sys}}) \times 10^{-4}$ ph $\text{cm}^{-2} \text{s}^{-1}$. Within 2σ uncertainties, this value is consistent with previous measurements by SPI and COMPTEL. Systematic uncertainties seen in previous COSI analyses of the 511 keV positron annihilation line and those intrinsic to the assumption of no ^{26}Al emission at high latitudes may account for the discrepancy. We find a total line shift of 2.5 ± 1.8 keV, an intrinsic line broadening of 9.7 keV (2σ upper limit), and limit the turbulent velocity of ^{26}Al ejecta to ~ 2800 km s^{-1} (2σ upper limit). Extensive simulations of the flight with several template maps affirm the consistency of the analysis pipeline with expectations. Overall, the framework behaves as expected and returns a 3.7σ measurement above background, consistent with previous measurements within $\sim 2\sigma$ statistical uncertainties.

The COSI 2016 balloon flight's measurement of ^{26}Al is key proof-of-concept for future studies of nucleosynthesis. Preserving the advantages of high resolution germanium Compton telescopes as demonstrated in the balloon iteration, the COSI satellite's low-Earth orbit presents a much more favorable background environment than the dominant atmospheric background and atmospheric attenuation seen at balloon altitudes (Cumani et al. 2019). These preferred background conditions and additional layer of four germanium detectors will increase the effective area. With improved angular resolution as well, the observational capabilities of the satellite platform are expected to strengthen COSI's ^{26}Al balloon measurement and probe unsolved questions about its origin, distribution, dynamics, and influence on the early Solar System.

5.7 Future work

The analysis presented in this chapter is the first analysis of ^{26}Al in the COSI 2016 balloon flight data. Repeated analyses employing different methods could verify, improve upon, and refine the results further. An alternative approach of modeling the background using adjacent energy bins on either side of the central 1803–1817 keV band, for example, would mimic the approach taken by COMPTEL (Oberlack 1997; Plüschke et al. 2001). A time-dependent background model could also improve the analysis, but modeling the temporal variability of the pernicious balloon background environment is notoriously difficult. The atmospheric model (Ling 1975) employed in this chapter uses the present knowledge available to the scientific community and is used in other analyses in the field (e.g., Lowell 2017; Sleator 2019; Takada et al. 2011), but it has its shortcomings. In particular, it does not incorporate the effects of cutoff rigidity and solar modulation which vary over the course of a balloon flight; note the effort to correct for rigidity manually in Section 5.3.5.5 of this chapter. With respect to instrumental γ -rays, we include the entire mass model of COSI in a MEGALib simulation which aims to reproduce the cosmic ray activation of balloon instrument material. Despite the proven reliability of MEGALib (which is based on the widely used Geant4 package, has been used in Sleator et al. (2019) to reproduce real instrument data, and has been further verified in the space radiation environment of the Transient Gamma-Ray Spectrometer (Zoglauer et al. 2008)), we observe largely the same spectral signatures seen in flight data with different amplitudes. To this end, future advancements in simulations may inform a more complete understanding of how to properly model the background environment.

Ultimately, however, it is important to emphasize that differences between real data and simulated spectra in this chapter’s analysis are not consequential to the presented flight results. All characteristics of the real flight data and simulations are captured in the empirical models generated for each. The simulations are created as representative “observations” which are analyzed identically to the real data, i.e. the simulations test the validity of the method under different assumptions. Results from the flight data are independent of the atmospheric and instrumental background simulations. This approach was specifically chosen in recognition of the difficulties presented by the balloon environment.

Chapter 6

Imaging with the Compton Spectrometer and Imager

This chapter pivots from the spectroscopic analysis in Chapter 5 to imaging with COSI. COSI’s imaging capabilities for line emission have been previously demonstrated with the 511 keV positron-electron annihilation excess (Siegert et al. 2020). The Crab Nebula was also correctly localized as a demonstration of COSI’s ability to image point sources (Zoglauer et al. 2021). These studies were facilitated by the development of `cosipy`, which was extensively tested by the author of this document to adapt the software for imaging in the COSI Data Challenges (Section 3.6.2). These preparatory efforts for the Data Challenge are presented here and culminate in the imaging of simulated ^{26}Al emission in Data Challenge 1. Limited statistics prevent imaging of the ^{26}Al emission measured in the COSI 2016 balloon flight data. Overviews of the Compton Data Space (CDS) and Richardson-Lucy deconvolution (Richardson 1972; Lucy 1974), the iterative algorithm used to reconstruct images, are also provided.

6.1 The Compton Data Space

The Compton Data Space (CDS; Schönfelder et al. 1993; Zoglauer et al. 2021) is a data space developed specifically for Compton telescopes which parameterizes photon interactions by their scattering angles. The CDS was first conceptualized by the COMPTEL collaboration as the “COMPTEL Data Space” and is more broadly used today in Compton telescope experiments as the “Compton Data Space.” The CDS is spanned by three parameters that specify the observed Compton scattering process as well as the measured changed state of the incident γ -ray: the Compton scattering angle ($\phi \in [0^\circ, 180^\circ]$) describes the former and the latter is given by the polar ($\psi \in [0^\circ, 180^\circ]$) and azimuthal ($\chi \in [-180^\circ, 180^\circ]$) directions of the scattered γ -ray. These three parameters describe the arrival (defined by ϕ) and scattered (defined by ψ and χ) directions of the γ -ray.

Figure 6.1 illustrates the connection between the Compton scattering process and the

CDS. In the left panel, the origin, i.e. true (point) source location, of an incident γ -ray is given by (Ψ_0, ξ_0) . The γ -ray undergoes an initial Compton interaction with the detector material (red point) at angle ϕ and subsequently moves in a different direction. The path of the scattered γ -ray is defined by the coordinates of its final interaction (Ψ_g, ξ_g) in the 2-D plane spanned by Galactic longitude Ψ and Galactic latitude ξ . Thus, the coordinates (Ψ, ξ) encode information about the directions of the incident and scattered γ -rays and span a “scattering plane.” Note that the photons are originally measured in terms of the instrument’s local detector coordinate system. The conversion between detector coordinates and Galactic coordinates is a rotation in spherical coordinates of the instantaneous instrument frame and the Galactic coordinate system frame (Sleator 2019).

The CDS is spanned by this scattering plane and the initial Compton scattering angle (ϕ) . The scattering plane can be visualized as the x - y axes of the CDS and ϕ as the z -axis. In an ideal case of unpolarized photons, the orientation of the scattering plane is random. This creates a hollow cone shape in the CDS, commonly referred to as the “Compton cone” (right image in Figure 6.1). The deviation of the scattered γ -ray direction (Ψ_g, ξ_g) from the incident direction (Ψ_0, ξ_0) is the Compton scattering angle ϕ , as the scattering act itself is what changes the path of the incident γ -ray. When the Compton scattering angle is zero, there is no deviation from the incident direction and the apex of the cone lies on the surface of the (Ψ, ξ) scattering plane at the true source position (Ψ_0, ξ_0) . As ϕ (i.e. the deviation from the incident direction) increases, the (Ψ, ξ) coordinates change at the same rate. As a result, the opening angle of the cone is 90° .

The notion of an “event circle,” as introduced in Section 3.3, becomes clear within the context of the Compton cone. Say that the photon scatters at Compton scattering angle ϕ_{scatt} . As ϕ_{scatt} increases, the photon deviates more from its original path, moving up the z -axis in the right panel of Figure 6.1, and lands at any (Ψ, ξ) on the circle which traces the flat face of the cone bisected by the 2-D plane $(\phi = \phi_{\text{scatt}}, \Psi, \xi)$. These are the “event circles” or “circles on the sky” often mentioned in explanations of Compton event reconstruction.

It is important to note that this is a simplified depiction of the scattering process and as such, the CDS is a simplified data space. The CDS does not encode the redundant kinematic and geometric information gleaned from real events which Compton scatter multiple times throughout the instrument. This information is key to Compton event reconstruction (e.g., Compton Kinematic Discrimination; Boggs and Jean 2000). Furthermore, the Compton cone in Figure 6.1 is not uniformly populated because certain scattering angles are more probable than others, given interaction probabilities governed by energy and polarization in the Klein-Nishina equation. The geometry of the detector itself also influences the probability of observing certain scattering angles. Comparing COMPTEL’s “classic” Compton telescope design to the “compact” Compton telescope design of COSI elucidates this point: COMPTEL was comprised of a scattering plane and an absorbing plane placed > 2 m apart. Photons scattered at large Compton scattering angles would not hit the absorbing plane and would consequently escape detection. COSI’s compact stack of detectors can detect photons scattered at a much broader range of scattering angles.

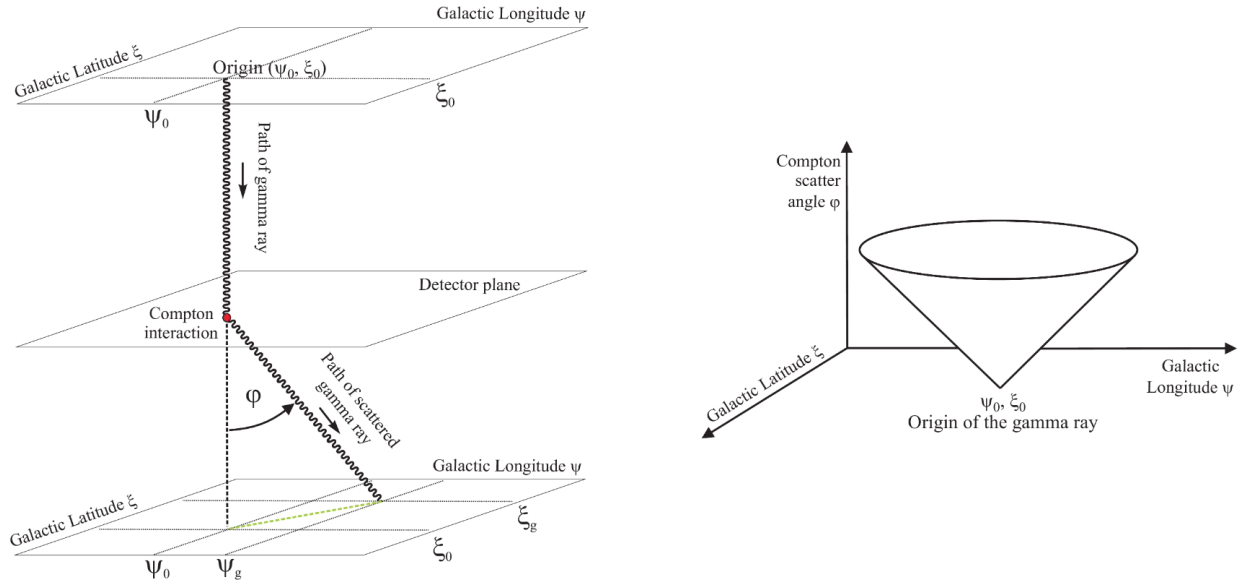


Figure 6.1: The connection between the Compton scattering process in a detector and the Compton Data Space (CDS). Image from Zoglauer et al. (2021). An interactive version of this figure is provided in a Jupyter notebook at <https://github.com/JBeechert/cds-visualization>.

6.2 Instrument response

The instrument response encodes how the instrument responds to, or measures, incident photons. It relates a photon's origin in the sky (Galactic longitude ℓ , latitude b) to photon interactions recorded in local detector coordinates (energy E , ϕ , ψ , χ). In doing so, it translates the data from physical to measured units. This is critical for imaging analyses which seek to project data recorded by the instrument in the CDS back onto the 2-D sky in Galactic coordinates. Visualizing source distributions in this way can reveal the true morphologies of γ -ray emission. Of particular interest are the strong signature of positron-electron annihilation (511 keV) near the Galactic Center and sites of nucleosynthesis traced by ^{26}Al (1.809 MeV), ^{60}Fe (1.173 and 1.332 MeV), and ^{44}Ti (1.157 MeV; also 68 and 78 keV in hard X-rays).

Mathematically, the response is written as a function of coordinates:

$$\text{Response } R(Z, A; \phi, \psi, \chi)$$

Z = zenith (angle in instrument coordinates)

A = azimuth (angle in instrument coordinates),

where (ϕ, ψ, χ) are the measured CDS angles (data space). In principle, the response is

energy-dependent but the energy dimension is omitted here because the presented analysis focuses only on the 1.809 MeV photopeak. Note that Z and A point to a location on the sky (image space) via a unique mapping of $(Z, A) \leftrightarrow (\ell, b)$ at each instance in time.

The measurement process implicitly sorts photons from the sky into the CDS. This can be visualized as a cube spanned by the three parameters of the CDS. Figure 6.2 shows the CDS (arbitrarily) binned into 4 bins in each of ϕ , ψ , and χ . Each measured photon falls into one of these bins, such that the instrument records 0, 1, 2, 3, etc. photons in each. The response is akin to an effective area (units: cm^2) which translates sky counts to the CDS:

$$\left[\int R_{\text{CDS}}^{ZA} S(\ell, b; \vec{\alpha})_{ZA} d\Omega \right] = \text{cm}^2 \cdot \text{ph cm}^{-2} \text{s}^{-1} \text{sr}^{-1} \cdot \text{sr} = \text{ph s}^{-1}, \quad (6.1)$$

i.e. a rate, which is expected because Compton telescopes are photon-counting instruments and the Poisson distribution refers to rates. The sky distribution S describing the source of photons is dependent on Galactic Coordinates (ℓ, b) and a vector of source parameters $\vec{\alpha}$ (amplitude of the source in the sky, for example). The ZA subscript denotes that the sky distribution is specified in the image space. The response takes this ZA distribution to the CDS, the measured data space. By integrating over solid angle, we obtain a photon rate in units of ph s^{-1} .

6.3 Model fitting

We use this understanding of the CDS and the response to image the γ -ray sky. One such imaging technique is called “model fitting.” This section presents a toy example of model fitting to motivate Richardson-Lucy deconvolution, the technique implemented in `cosipy`, described in the next section.

Consider a simple point source:

$$\text{Sky } S = F \cdot \delta(\ell - \ell_0) \delta(b - b_0)$$

In words, a point source S with arbitrary flux F is located at Galactic coordinates (ℓ_0, b_0) . When S is folded into the response (Equation 6.1), the delta function selects only the elements of the response with $\ell = \ell_0$ and $b = b_0$. Assuming a known “true” source location (ℓ_0, b_0) from complementary observations simplifies the imaging procedure by leaving the flux F as the only parameter to be fitted, though in principle one could fit for both flux and source location. Since the response R is not an infinitesimal function, it is binned (Section 6.2) and the bins of R which are closest to (ℓ_0, b_0) are weighted more heavily than distant bins.

The preceding discussion considered the source distribution (the sky) as a function of longitude and latitude: $S(\ell, b; \vec{\alpha})_{ZA}$. The background, by contrast, is assumed to appear identically across the sky. It does not change with ℓ and b . It can, though, depend on parameters $\vec{\beta}$, such as the background amplitude, such that:




Figure 6.2: Visualization of the measurement process, which sorts detected photons into the three-dimensional Compton Data Space (CDS). The CDS is parameterized by the scattering angles (ϕ, ψ, χ) of the photon path. For example, two photons are detected with $(\phi \in [2.0^\circ, 3.0^\circ], \psi \in [1.0^\circ, 2.0^\circ], \chi \in [2.0^\circ, 3.0^\circ])$.

$$\text{Background } B = B(\vec{\beta})_{\text{CDS}}$$

Extensive background simulations define B . The instrument is irradiated isotropically with simulated background photons, specified by a chosen model (e.g., Ling 1975), until the CDS is populated with sufficient statistics to comprise an approximation of typical background measurements. Summing the sky model (originally specified in Galactic coordinates and subsequently convolved with the response to obtain its representation in the CDS) and the simulated background (as defined in the CDS), yields a total model M_{CDS} which can be fitted to real data that are measured in terms of the CDS:

$$M_{\text{CDS}} = R_{\text{CDS}}^{ZA} S(\ell, b; \vec{\alpha})_{ZA} + B(\vec{\beta})_{\text{CDS}} \quad (6.2)$$

The returned fit parameters $\vec{\alpha}$ and $\vec{\beta}$ can then be interpreted with associated statistical uncertainties.




Figure 6.3: A simplified illustration of the “image space:” the two-dimensional sky is spanned by Galactic longitude ℓ and Galactic latitude b and divided into pixels (bins) of specified size. Richardson-Lucy deconvolution finds the most likely number of photons in each pixel given the measured data in the CDS.

6.4 Richardson-Lucy deconvolution

6.4.1 Overview

The above fitting method is predicated on knowledge (an assumption) of the sky distribution S . For example, the DIRBE $240\ \mu\text{m}$ map is commonly assumed as a tracer of Galactic ^{26}Al emission. If the sky distribution S is unknown, however, fitting for the most likely parameters of S and B within the CDS is technically possible but not very meaningful without constraints. Thus, instead of maximizing the likelihood in the CDS with the fitting method, Richardson-Lucy (RL) deconvolution maximizes the likelihood in image space. RL deconvolution seeks to answer the question: “what is the most likely location (origin) of recorded photons (data space, i.e. CDS) in the sky (image space)?”

Figure 6.3 shows 50 photons distributed in image space across the sky, as projected onto a simple 2-D plane spanned by Galactic longitude ℓ and Galactic latitude b . Note that this is an image space analogy to Figure 6.2, which shows photons populating the three dimensions of the CDS. The RL algorithm finds the most likely (ℓ, b) bin in image space from which a photon originated given its position as recorded by the instrument in the CDS.

Recalling that the response transforms physical to measured units, a simple method of answering the posed question would be to apply the inverse of the response to the measured

data and immediately find the corresponding location on the sky. This would constitute a simple mapping of the CDS directly onto the sky (image space). However, the response matrix R is not invertible. Instead of applying the inverse (R^{-1}), we apply the adjoint of the response (R_{CDS}^{ZA}) † to the data D_{CDS} and obtain a sky distribution $S(\ell, b)$, which is equivalent to back-projecting the circles of the CDS onto the sky:

$$S(\ell, b) = (R_{\text{CDS}}^{ZA})^\dagger D_{\text{CDS}} \quad (6.3)$$

The algorithm is typically initialized with a flat, uniform map as an unbiased starting assumption. This map is folded through the response, i.e. translated to the CDS, and added to the simulated background ($B(\vec{\beta})_{\text{CDS}}$) described previously. This becomes the complete model expectation, whereby the adjoint of the response projects the background-subtracted data (where the background is that defined by simulations) from the CDS onto the image space:

$$(R_{\text{CDS}}^{ZA})^\dagger (D_{\text{CDS}} - B(\vec{\beta})_{\text{CDS}}) = S(\ell, b), \text{ the most likely image of sky photons.} \quad (6.4)$$

6.4.2 Richardson-Lucy Algorithm

This section presents a brief, more formal explanation of the Richardson-Lucy (RL) algorithm. Successive images f are updated according to the following equations, as adapted from Knödlseeder et al. (1999) to reflect the notation used in this chapter:

$$f_j^{k+1} = f_j^k + \delta f_j^k, \quad (6.5)$$

where

$$\delta f_j^k = f_j^k \left(\frac{\sum_{i=1}^N (\frac{D_i}{M_i^k} - 1) R_{ij}}{\sum_{i=1}^N R_{ij}} \right). \quad (6.6)$$

More compactly,

$$f_j^{k+1} = f_j^k \left(\frac{\sum_{i=1}^N \frac{D_i}{M_i^k} R_{ij}}{\sum_{i=1}^N R_{ij}} \right). \quad (6.7)$$

The data space is indexed by i (above, by ‘‘CDS’’), the image space by j (above, by ‘‘ZA’’), and the iteration by k . The instrument response is R_{ij} and the number of observed photons in data space bin i is denoted by D_i . The expected number of photons in data space bin i is the model M , where $M_i^k = \sum_{j=1}^m R_{ij} f_j^k + B_i$ for background B_i . Equation 6.7 then demonstrates that image f_j^k is corrected by a multiplicative factor which encapsulates the ratio of measured to expected counts. In practice, this image f_j^k is corrected by the additive delta image defined in Equation 6.6.

We emphasize that this deconvolution occurs in the image space: the ratio of counts in the data space (D_i/M_i^k) is applied to the response R_{ij} , which moves from the data space (CDS) to the image space. This is normalized by the sum over the instrument response in the CDS (i.e. the exposure) and the resulting factor is applied to image f_j^k , yielding the new image f_j^{k+1} .

Note that the sum over the instrument response in the denominator of Equation 6.6 can lead to edge effects in the images. Poor exposure (small $\sum_{i=1}^N R_{ij}$) increases the magnitude of the corrective factor. The algorithm assumes that if the instrument detects a photon from a minimally observed region, it is likely to have originated from that point given the relative lack of opportunity to detect it. Weighting these photons more heavily manifests as bright artifacts in the images near the edges of regions of poor exposure. In the following sections, regions of bad exposure are manually masked out to mitigate this unwanted effect.

The `cosipy` implementation of the RL algorithm is slightly modified from the standard algorithm above. First, the fourth-root of the exposure, $(\sum_{i=1}^N R_{ij})^{0.25}$, is applied as a noise-dampening term to the delta image. This weight reduces the influence of artifacts from poor exposure. Second, a multiplicative parameter is also applied to the delta image to accelerate the convergence of the algorithm, which in its original form iterates slowly through small changes to the delta image. This parameter is determined by a maximum likelihood fit in each iteration and preferentially weights photons in adjacent pixels (if several photons are detected in a given pixel, it is more likely for an adjacent pixel to detect photons than a distant pixel). Third, the `cosipy` RL algorithm accommodates variations in the background, unlike the original algorithm which assumes it to be constant. Here, the background $B = B(\vec{\beta}_{\text{CDS}})$ is re-scaled by fitting parameters $\vec{\beta}_{\text{CDS}}$ simultaneously with the acceleration parameter in each iteration, thereby minimizing the dependence on the initial guess of the image through intermediate adjustments.

Section 6.5 summarizes initial tests of the `cosipy` RL algorithm developed to check basic imaging capabilities. Establishing confidence in the RL algorithm facilitated its deployment in COSI Data Challenge 1, detailed in Section 6.6. An important consideration when interpreting the images in Sections 6.5 and 6.6 is that those from `cosipy` have arbitrary flux normalization. To this end, in the following sections we focus on the qualitative comparison of point sources and extended sources in images produced by `mimrec` in MEGALib and by `cosipy`.

6.5 Testing the Richardson-Lucy algorithm

A series of tests were conducted to examine the behavior of the RL algorithm in the `cosipy` framework. Simulations of point sources and extended sources in `cosima` with varying degrees of complexity in instrument orientation probe the extent to which the algorithm can recover typical astrophysical sources. All sources are simulated at the monoenergetic ^{26}Al decay signature of 1808.74 ± 1.6 keV (rest energy). No background is included, the 9-detector mass model of the COSI balloon instrument is used, and the sky is divided into $6 \times 6 \text{ deg}^2$

pixels. The transmission probability of the atmosphere at 33 km is assumed in order to resemble the nominal atmospheric conditions during flight. Images from `mimrec` are generated using an iterative list-mode-likelihood algorithm and those from `cosipy` are generated using the RL deconvolution algorithm. A $6 \times 6 \text{ deg}^2$ 9-detector instrument response to 1809 keV emission was simulated specifically for imaging purposes. The response is set to zero beyond 60° from COSI’s zenith in order to mimic the extent of COSI’s real field of view (see further investigation in Section 6.5.4 and Figure 6.17) and to avoid sparsely populated response elements at higher zenith angles.

Caution is advised against over-interpreting comparisons between the imaging cases below; this is good practice when analyzing individual images as well, given inevitable artifacts and reconstruction effects. Due to their Poisson nature, potentially low-number statistics, etc., the imaging algorithms sometimes require fine-tuning of constituent parameters and adjustments to color scaling in order to obtain discernible images. As such, the following sections are not presented as definitive case studies benchmarking the quantitative limits of COSI observations with respect to observable fluxes and instrument orientations. Rather, they are presented as documentation of basic debugging checks which confirm qualitative functionality of the new `cosipy` imaging capability compared to that of `mimrec`, which has been extensively tested over many years of MEGALib development.

6.5.1 “Fixed” orientation

6.5.1.1 Point sources

The simplest observing case is that of, for example, a traditional ground-based telescope: the instrument stares with fixed orientation at a point source of known location. To test this simple case, a point source at Galactic longitude and latitude $(\ell, b) = (3^\circ, 21^\circ)$ with arbitrary flux of $0.1 \text{ ph cm}^{-2} \text{ s}^{-1}$ is simulated in `cosima`. This location is chosen because it falls in the middle of a $6 \times 6 \text{ deg}^2$ pixel. The total observation time is 10^4 s and all data are binned into one time bin encompassing the full observation time. The reconstructed images from `mimrec` and `cosipy` are shown in Figures 6.4a and 6.4b, respectively. There are 509 photons in the simulation. Both algorithms correctly localize the source to the expected pixel.

The simulation is repeated at four additional locations to test how well the algorithm can simultaneously localize multiple point sources across the sky: $(\ell, b) = (3^\circ, 21^\circ)$, $(21^\circ, -3^\circ)$, $(-3^\circ, -3^\circ)$, $(21^\circ, 3^\circ)$, and $(21^\circ, 33^\circ)$. There are 509, 530, 552, 541, and 436 photons in the simulated sources, respectively. Each source contains approximately the same number of photons, within statistical fluctuations, because each is simulated with the same flux and observed by COSI with fixed orientation over identical observation time. The `mimrec` (Figure 6.5a) and `cosipy` (Figure 6.5b) images are qualitatively similar and recover the five simulated sources. The true image of these five point sources, plotted directly as illuminated pixels, i.e. not simulated as viewed by COSI and reconstructed via an algorithm, is shown




Figure 6.4: Images in (a) `mimrec` and (b) `cosipy` of a simulated point source at $(\ell, b) = (3^\circ, 21^\circ)$. The instrument stares directly at the source with fixed orientation for 10^4 s (Section 6.5.1.1).

in Figure 6.5c. These tests confirm that the RL algorithm in `cosipy` can image multiple simulated point sources in the field of view when COSI stares at them with fixed orientation.

6.5.1.2 Extended sources

As a test of the algorithm’s ability to image diffuse emission, four two-dimensional Gaussian distributions of flux $0.1 \text{ ph cm}^{-2} \text{ s}^{-1}$ are simulated over an observation time of 10^4 seconds. The 1σ extent of the Gaussians in longitude is $\pm 30^\circ$ ($\sigma_\ell = 30^\circ$) and that in latitude is $\pm 1^\circ$ ($\sigma_b = 1^\circ$). Each Gaussian is centered at $(\ell = 0^\circ, b = 0^\circ)$. The four Gaussians are rotated 45° , 90° , 135° , and 180° from their nominal orientation along the Galactic Plane. The number of photons in the 45° -, 90° -, 135° -, and 180° -rotated sources are 489, 449, 478, and 580, respectively. True images of the sources, i.e. not simulated through COSI’s field of view and reconstructed in an algorithm, are plotted directly as two-dimensional Gaussians in Figure 6.6. Images of the sources, as simulated in `cosima` and reconstructed in `mimrec`, are shown in Figure 6.7. The combination of the 45° - and 135° -rotated Gaussians creates a clearly visible “X” in Figure 6.7e. Figure 6.7f shows a “plus” created by the 90° - and 180° -rotated Gaussians and Figure 6.7g shows a “star” formed by combining all four sources together. The “plus” and the “star” are less clearly resolved than the individual sources and the “X.” The $6^\circ \times 6^\circ$ binning blends the sources together and limits the achievable resolution of the image.

The images are subsequently generated in `cosipy`’s RL deconvolution algorithm. Several iterations of the “X” are shown in Figure 6.8 to demonstrate the progression of the algorithm. The convergence of the fitted background amplitude, fitted acceleration parameter, and likelihood of the image in each iteration are also shown in Figure 6.9. There is no background included in the simulations and the initial guess for the background parameter is set to 10^{-6} . It largely remains at this value throughout the iterations. The spike at an early iteration is

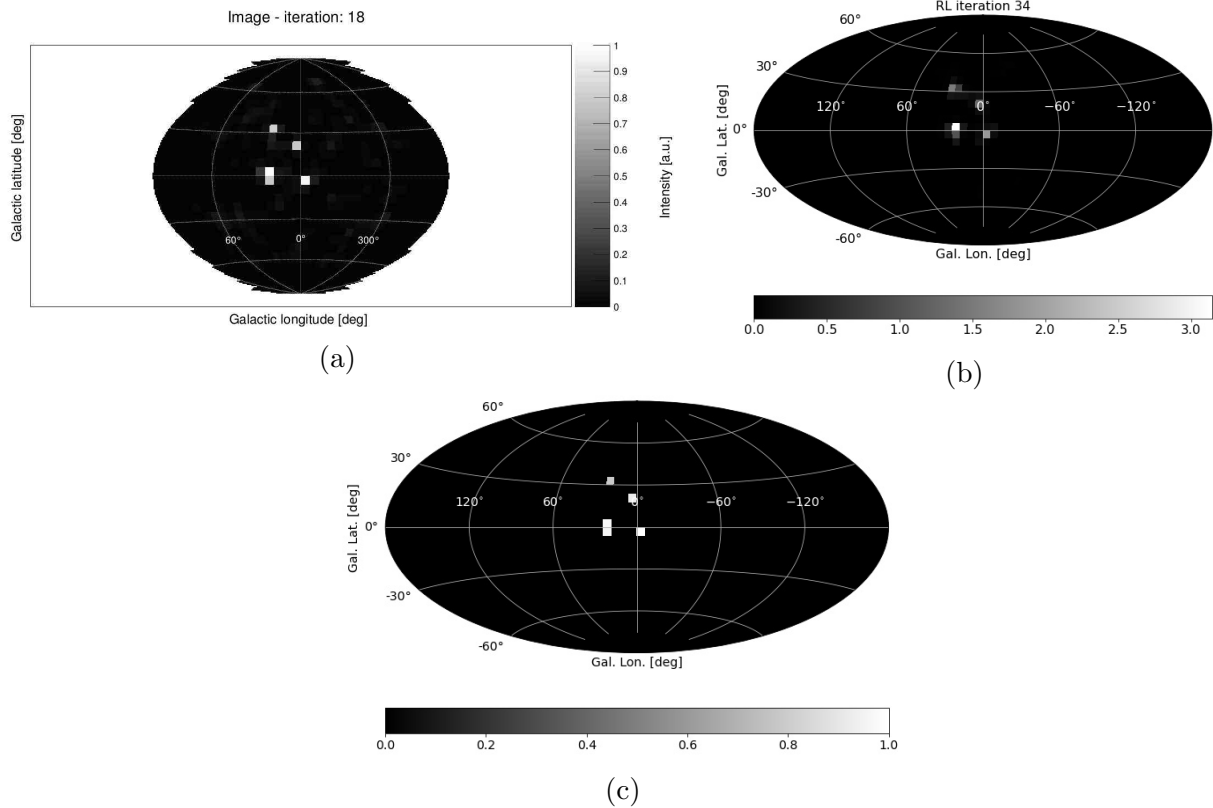


Figure 6.5: Images in (a) `mimrec` and (b) `cosipy` of five simulated point sources: $(\ell, b) = (3^\circ, 21^\circ)$, $(21^\circ, -3^\circ)$, $(-3^\circ, -3^\circ)$, $(21^\circ, 3^\circ)$, and $(21^\circ, 33^\circ)$. The instrument stares directly at each source with fixed orientation for 10^4 s (Section 6.5.1.1). The true sources, not simulated through COSI’s field of view and reconstructed by an algorithm, are plotted directly on the sky in (c).

a temporary aberration in the algorithm’s fitting procedure as it finds the right scale. The acceleration parameter is initially set to 2000 and converges at later iterations. Both the “X” and “plus” are discernible in their “final” iteration 34 (Figure 6.10). The choice of final iteration is flexible and is selected here at iteration 34 because the fitted parameters and likelihood converged and the image stopped changing appreciably in appearance with each iteration. Overall, the imaging algorithms behave as expected. Two-dimensional Gaussians are recognizably reconstructed in `mimrec` and `cosipy`, indicating that COSI can image diffuse emission when staring with fixed orientation at the simulated sources.

6.5.1.3 Few photons

As an additional check of imaging performance, we select only the first ten photons of the fixed orientation point source simulation at $(\ell = -3^\circ, b = -3^\circ)$ in Section 6.5.1.1. Figure 6.11

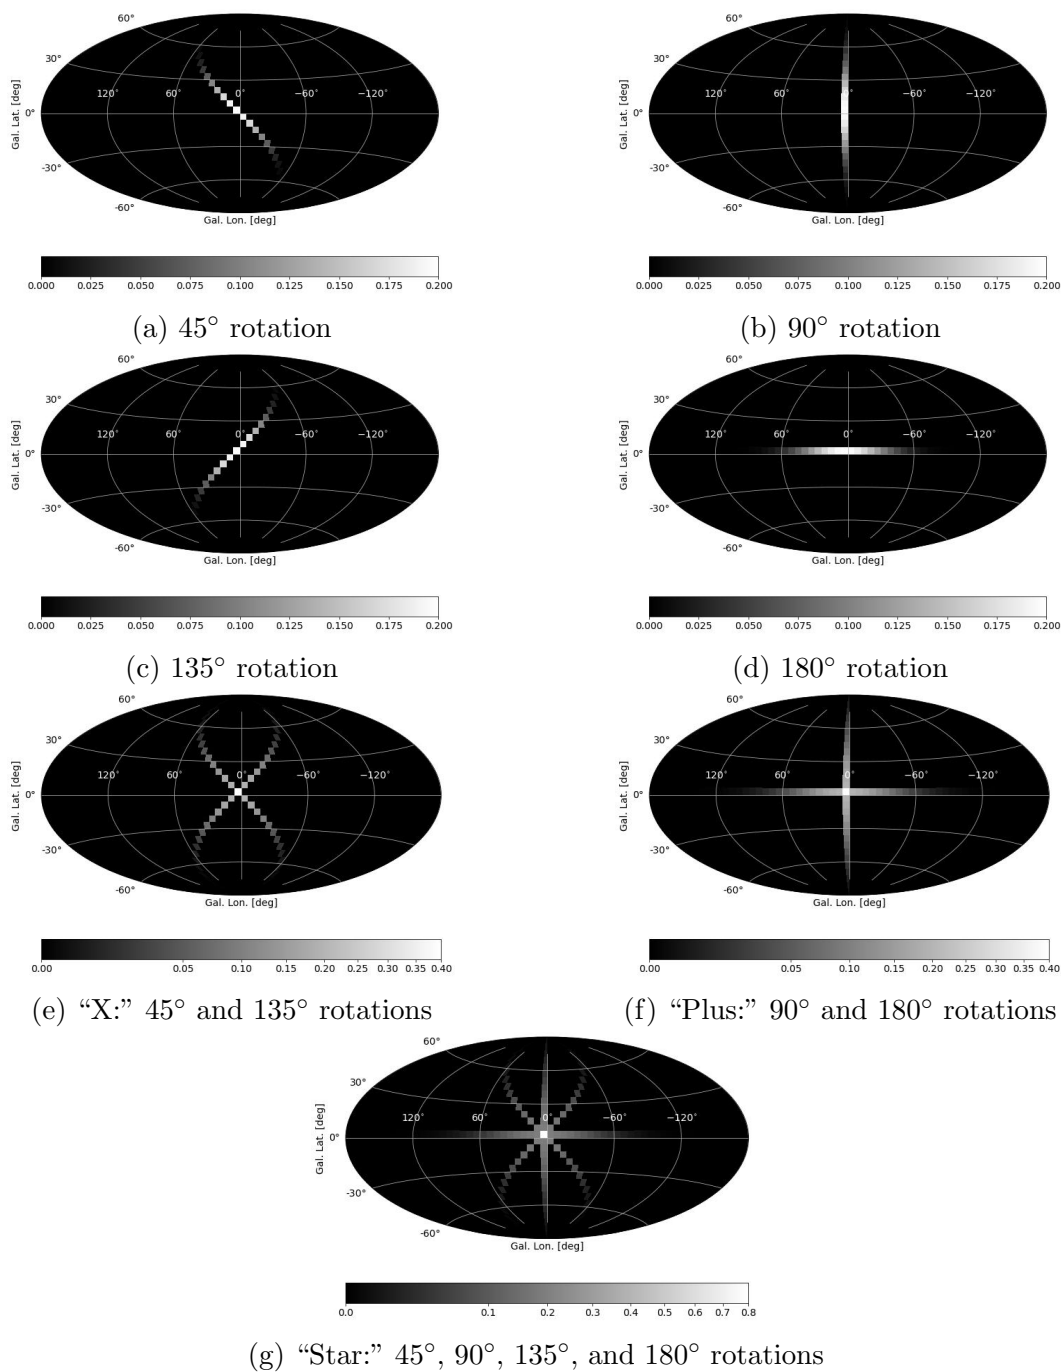


Figure 6.6: True images of the two-dimensional Gaussian sources, i.e. plotted directly from their functional forms (not reconstructed). These distributions are simulated separately in *cosima* and reconstructed in *mimrec* and *cosipy* throughout this section.

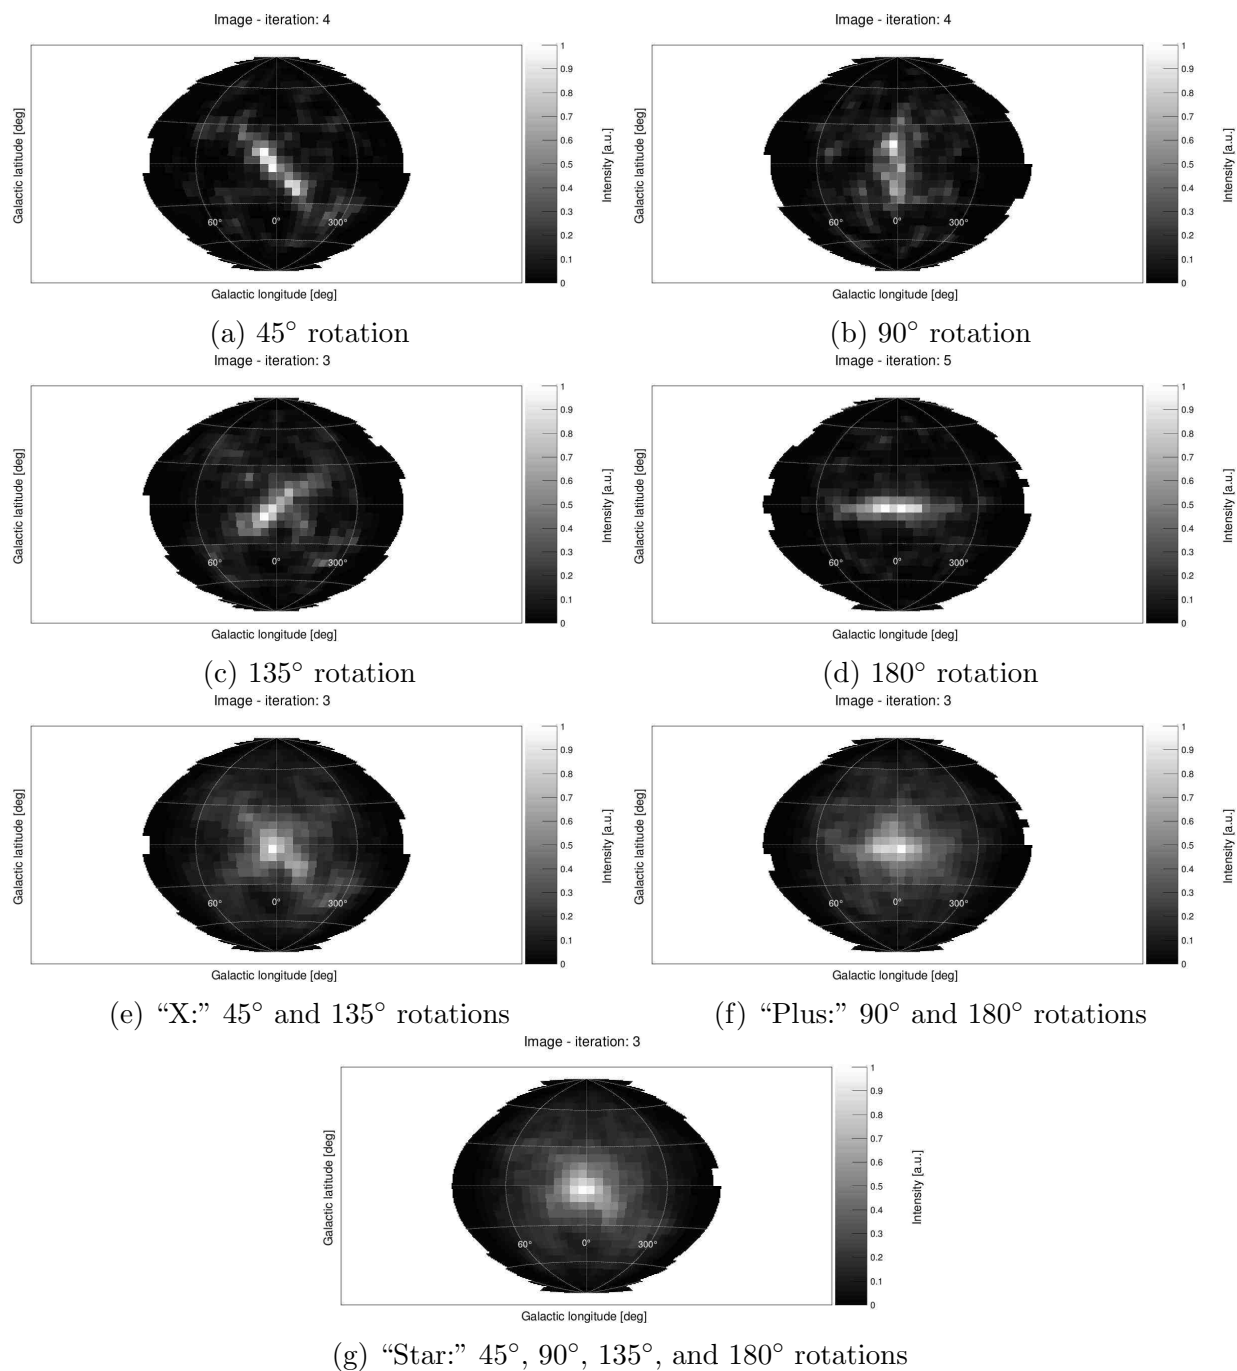


Figure 6.7: Images in *mimrec* of extended emission when COSI stares with fixed orientation at two-dimensional Gaussian sources (Section 6.5.1.2).




Figure 6.8: Successive iterations of the *cosipy* RL algorithm when COSI stares with fixed orientation at a composite “X”-shaped source (Section 6.5.1.2).

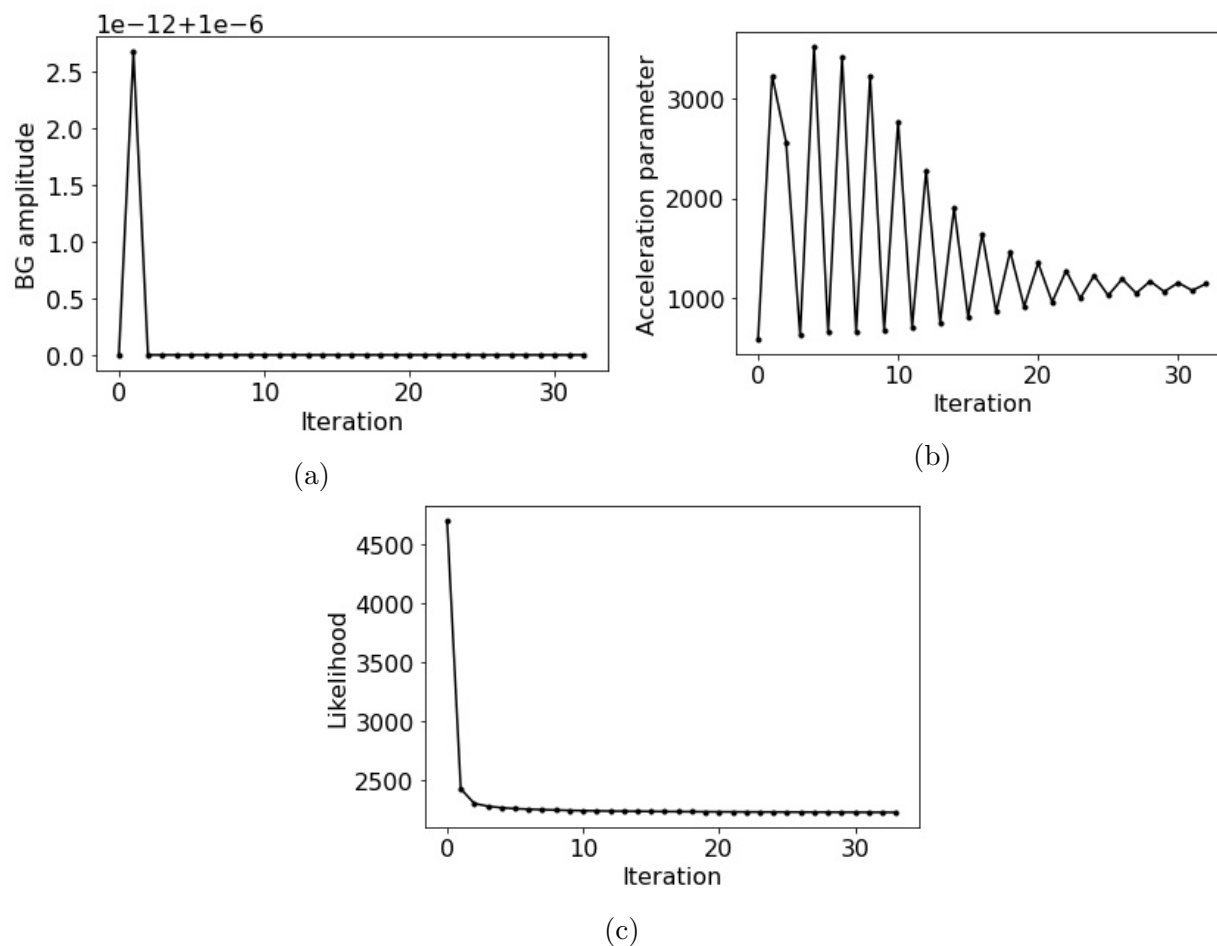


Figure 6.9: Convergence of the background amplitude, acceleration parameter, and likelihood of the “X” shape in the `cosipy` RL algorithm when COSI stares with fixed orientation (Section 6.5.1.2; Figure 6.8). The negative likelihood is plotted.

shows that the image created by the RL algorithm in `cosipy` agrees with that created by the list-mode algorithm in `mimrec`. This is further qualitative validation of the imaging algorithms.

6.5.2 “Simple” orientation

6.5.2.1 Point sources

A layer of complexity is added to the imaging analysis by allowing COSI to move with a simple, well-controlled orientation. This orientation stands in contrast to the “fixed” case which holds the source directly in COSI’s field of view for the entire observation time. The previous 10^4 s observation time is divided into two 5000s parts. In the first 5000s, the




Figure 6.10: The “X” and “plus” as imaged by `cosipy` when COSI stares with fixed orientation at the constituent two-dimensional Gaussian sources (Section 6.5.1.2). The image in (a) is that from Figure 6.8f, shown again here for completeness.




Figure 6.11: Images in (a) `mimrec` and (b) `cosipy` of the first ten photons of a simulated point source at $(\ell = -3^\circ, b = -3^\circ)$ (Section 6.5.1.1). Even with few photons, the list-mode imaging algorithm in `mimrec` and RL algorithm in `cosipy` produce similar images.

instrument stares directly at the source. In the second 5000 s, the instrument rotates 60° off the previous pointing. Each source is again simulated with a flux of $0.1 \text{ ph cm}^{-2} \text{ s}^{-1}$.

Predictably, fewer photons are recorded in this simulation because the instrument does not observe each source for the entire simulation time. Instead, it moves away from the source halfway through the total exposure. The number of photons at $(\ell, b) = (3^\circ, 21^\circ)$, $(21^\circ, -3^\circ)$, $(-3^\circ, -3^\circ)$, $(21^\circ, 3^\circ)$, and $(21^\circ, 33^\circ)$ are 182, 408, 101, 425, and 325, respectively. Recording fewer overall photons slightly degrades the appearance of the sources in the `cosipy` RL algorithm (compare Figure 6.12b to the fixed orientation case in Figure 6.5b). The source at $(21^\circ, 33^\circ)$ is spread across three pixels in the former compared to two pixels in the latter. Additionally, the “simple” case appears to blur the $(-3^\circ, -3^\circ)$ source at positive longitude (the color scaling in Figure 6.12c is adjusted to highlight fainter emission) compared to the expected localization in the fixed case (Figure 6.5b).

6.5.2.2 Extended sources

Simulating the same 2-D Gaussian sources as in Section 6.5.1.2 with the “simple” orientation more clearly demonstrates how recording fewer photons degrades source recovery. For brevity, only the “X” and “plus” shapes are shown. A consequence of moving away from the source halfway through the observation time, the instrument does not measure enough photons from the newly off-axis regions to image them fully and only half of the “X” and “plus” are visible on the sky in `minrec` and `cosipy` (Figure 6.13). This example highlights the basic yet important principle that we cannot see what we do not observe. Our present understanding of the sky is necessarily restricted to knowledge only of regions from which we have collected photons. Conclusions about a source must be drawn cautiously in the potential absence of comprehensive observations.

6.5.3 “Fine” orientation

6.5.3.1 Point sources

We further increase the complexity of the orientation file by changing the zenith pointing of the instrument by $\sim 10^\circ$ every ~ 2000 s:

- 0 s, 0°
- 1000 s, 10°
- 2000 s, 10°
- 3000 s, 20°
- 4000 s, 20°
- 5000 s, 30°
- 6000 s, 30°
- 7000 s, 40°
- 8000 s, 40°
- 9000 s, 60°
- 10000 s, 60°

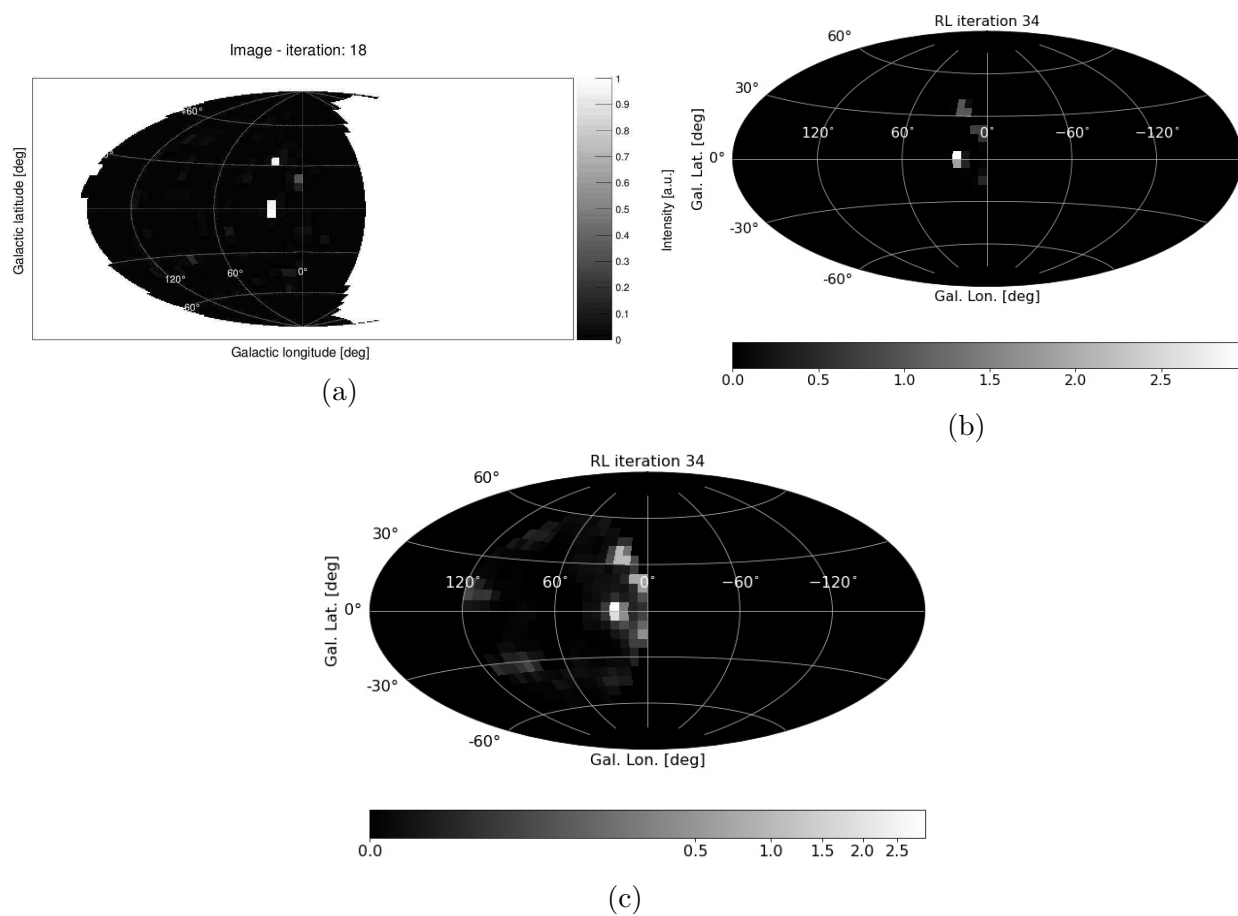


Figure 6.12: Images in (a) *mimrec* and (b) *cosipy* of five simulated point sources: $(\ell, b) = (3^\circ, 21^\circ)$, $(21^\circ, -3^\circ)$, $(-3^\circ, -3^\circ)$, $(21^\circ, 3^\circ)$, and $(21^\circ, 33^\circ)$. The instrument orientation shifts 60° after 5000 s for a total observation time of 10^4 s (“simple” orientation, Section 6.5.2.1). The *cosipy* image in (c) is the same as that in (b) but with a different color scale to enhance smaller flux gradients.

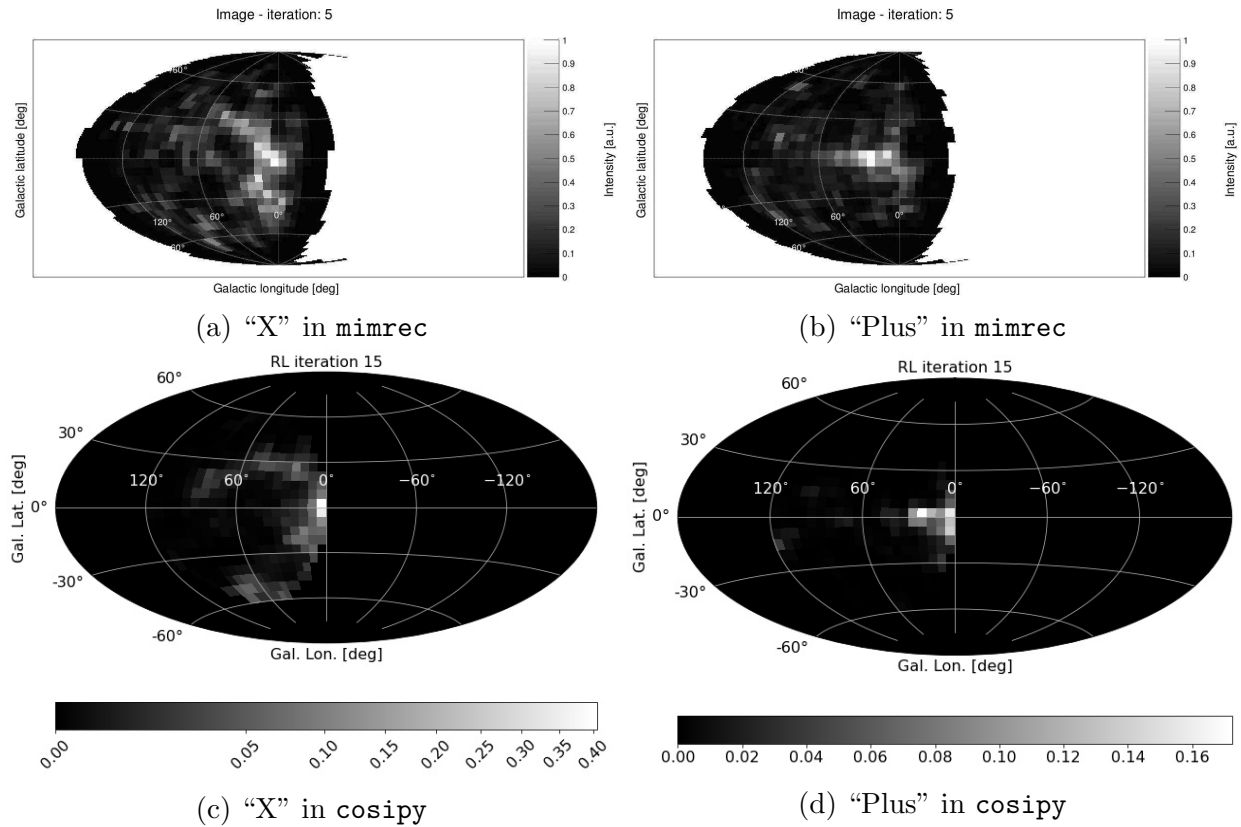


Figure 6.13: Images in (a, b) *mimrec* and (c, d) *cosipy* of “X”- and “plus”-shaped sources. The instrument orientation shifts 60° after 5000 s for a total observation time of 10^4 s (“simple” orientation, Section 6.5.2.1). The dramatic change in pointing obfuscates half of the underlying source morphology.

Fewer counts are recorded in each source than in the case of the fixed orientation, when the instrument stared at each source. The number of photons at $(\ell, b) = (3^\circ, 21^\circ)$, $(21^\circ, -3^\circ)$, $(-3^\circ, -3^\circ)$, $(21^\circ, 3^\circ)$, and $(21^\circ, 33^\circ)$ is 424, 547, 394, 525, and 423, respectively. More counts are recorded in this orientation scheme than in the previous “simple” case because all observations after the sole orientation change in the latter are highly off-axis. The more gradual change in orientation here yields more observation time with the sources inside the instrument’s field of view.

As such, the sources are more clearly resolved in the “fine” *cosipy* image (Figure 6.14b) than in the *cosipy* image of the “simple” case (Figure 6.12b). The simulations are binned in 1000 s time bins with little observed change in results for larger and smaller time bins. In keeping with the previous tests, the five point sources are generally more well-defined in the *mimrec* image (Figure 6.14a) than in the *cosipy* images.




Figure 6.14: Images in (a) `mimrec` and (b) `cosipy` of five simulated point sources: $(\ell, b) = (3^\circ, 21^\circ)$, $(21^\circ, -3^\circ)$, $(-3^\circ, -3^\circ)$, $(21^\circ, 3^\circ)$, and $(21^\circ, 33^\circ)$. The instrument orientation shifts 10° off-axis every ~ 2000 s (“fine” orientation, Section 6.5.3.1).

6.5.3.2 Extended sources

The extended two-dimensional Gaussian sources from Section 6.5.1.2 are simulated assuming the same 1809 keV energy and 10^4 s observation time, except in this case the flux is doubled to $0.2 \text{ ph cm}^{-2} \text{ s}^{-1}$ and the orientation follows the more complex “fine” scheme described above. There are 681, 701, 705, and 774 photons in the 45° -, 90° -, 135° -, and 180° -rotated Gaussian distributions. The “X” and “plus” sources for this orientation are shown in Figure 6.15. The `cosipy` images are generated with 1000 s time bins to follow the changes in instrument orientation.

Even with the more complex instrument orientation, the “X” and “plus” shapes are visible, if slightly contaminated by errant bright pixels (artifacts). The “plus” image is more difficult to resolve than the “X” in both the `mimrec` in `cosipy` images. Unlike in previous tests, the “plus” is slightly more clear in the `cosipy` image than in the `mimrec` image, though this should not be over-interpreted as a direct reflection of algorithm performance given the variations in appearance with iteration number, color scaling, etc. This test instead demonstrates that COSI can image simulated sources when its pointing moves more rapidly through its field of view than it did in the case of one change in orientation.

6.5.4 “Step 6 degrees” orientation

6.5.4.1 Point sources

To study COSI’s response over a smoother scan of its $\sim 60^\circ$ field of view, a new orientation file is created. The pointing of the instrument is changed in 6° increments every 1000 s from 0° to 78° for a total observation time of 13000 s. The 6° spacing mimics the angular resolu-

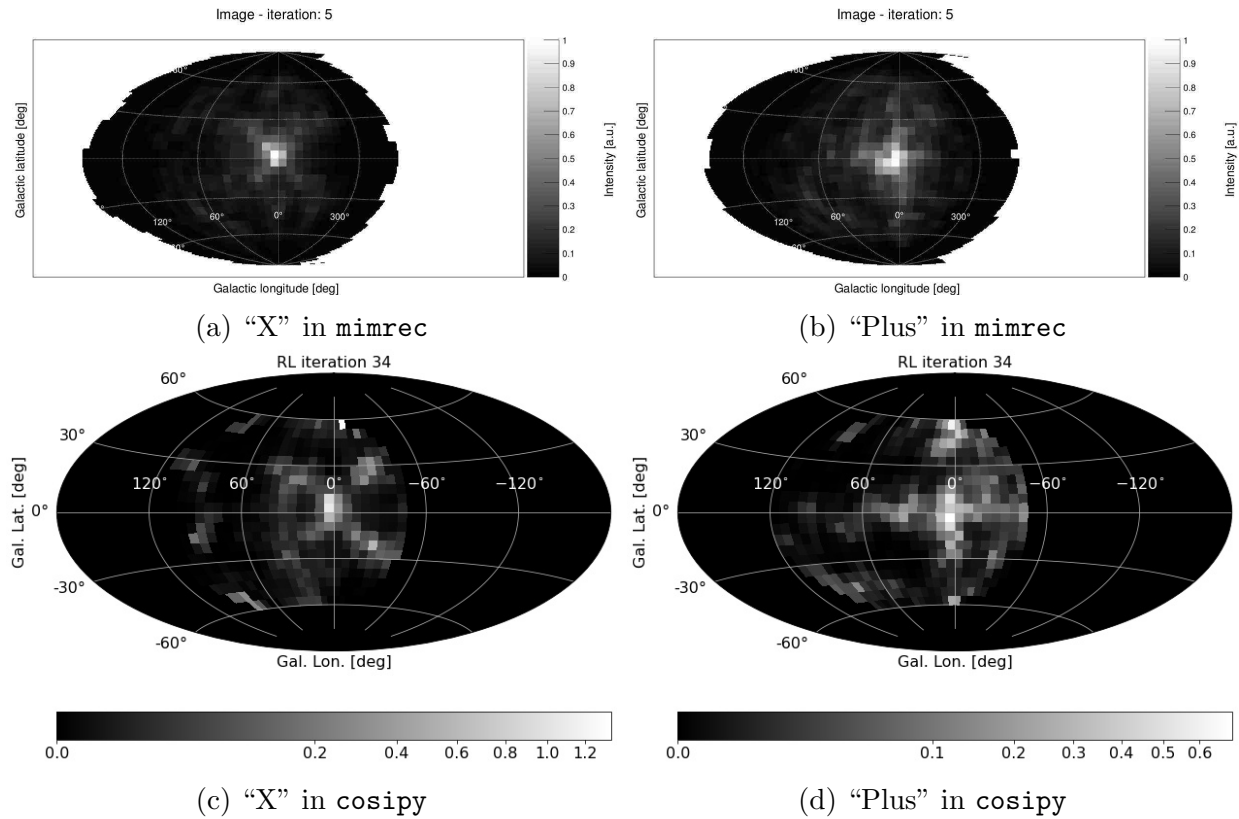


Figure 6.15: Images in (a, b) *mimrec* and (c, d) *cosipy* of “X”- and “plus”-shaped sources. The instrument orientation shifts 10° off-axis every ~ 2000 s (“fine” orientation, Section 6.5.3.2).

tion assumed in the simulation of COSI’s response matrix and the binning of the generated images. The pointing changes at the following times and degrees from zenith, respectively:

- 0 s, 0°
- 1000 s, 6°
- 2000 s, 12°
- 3000 s, 18°
- 4000 s, 24°
- 5000 s, 30°
- 6000 s, 36°
- 7000 s, 42°
- 8000 s, 48°
- 9000 s, 54°
- 10000 s, 60°
- 11000 s, 66°
- 12000 s, 72°
- 13000 s, 78°





Figure 6.16: Images in (a) `mimrec` and (b) `cosipy` of five simulated point sources: $(\ell, b) = (3^\circ, 21^\circ)$, $(21^\circ, -3^\circ)$, $(-3^\circ, -3^\circ)$, $(21^\circ, 3^\circ)$, and $(21^\circ, 33^\circ)$. The instrument orientation shifts 6° off-axis every 1000 s (“step 6 degrees” orientation, Section 6.5.4.1).


Point sources at $(\ell, b) = (3^\circ, 21^\circ)$, $(21^\circ, -3^\circ)$, $(-3^\circ, -3^\circ)$, $(21^\circ, 3^\circ)$, and $(21^\circ, 33^\circ)$ are simulated with a flux of $0.2 \text{ ph cm}^{-2} \text{ s}^{-1}$; 872, 1197, 817, 1173, and 991 photons are recorded in each, respectively. For the `cosipy` RL image, the data are binned in 1000 s time bins to follow the change in orientation. Both the `mimrec` (Figure 6.16a) and `cosipy` (Figure 6.16b) images clearly show the five point sources.

This orientation scheme also serves as an informative test of `cosipy`’s response handling method. The (unitless) response is normalized such that the sum of the response over all sources equals one. As such, weighting the response at each source location by the total number of simulated counts in that source should equal the true (raw) number of simulated counts in the source. Hence, we compare the product of the simulated counts and the normalized response to the true number of simulated counts (Figure 6.17). We perform this test under two different conditions. First, the response of the instrument is set to zero at a zenith angle of 60° to unambiguously restrict instrument sensitivity to within its $\sim \pm 60^\circ$ field of view. All images in this and previous sections are generated with this setting unless otherwise noted. The second condition sets the cut to 90° , allowing a non-zero, if severely diminished, instrument response to photons incident from up to 90° beyond zenith.

As shown in Figure 6.17a, when the response is truncated at 60° , the instrument does not record all of the simulated counts. This is most clear in the point sources at $(\ell, b) = (3^\circ, 21^\circ)$ and $(-3^\circ, -3^\circ)$. The solid line, which shows the weighted response, falls to zero in later time bins when the instrument pointing places the sources $> 60^\circ$ off-axis. However, the dashed line shows non-zero counts at these times, indicating that truncating the response loses counts which are simulated as having been recorded by COSI in this orientation scheme. By contrast, these off-axis photons are recorded in the case of a 90° cut (Figure 6.17b) because



(a)



(b)

Figure 6.17: COSI’s normalized response, weighted by the number of simulated counts in each source (solid lines), compared to the true (raw) number of simulated counts in each source (dashed lines). (a) The response is truncated at 60° such that no photons incident from $\pm 60^\circ$ beyond zenith are recorded. This reflects COSI’s field of view. (b) The response is truncated at 90° . More photons are recorded in the latter case because the response is sensitive to photons from a wider spatial origin.

the response is sensitive to photons from a wider spatial origin.

In summary, we observe an expected drop-off in response as point sources move off-axis in the 60° field of view. The close agreement between the weighted response and true simulated counts indicates that the response handling in `cosipy` works as expected. Setting the response cut to 90° gives the most complete representation of sources located at the boundary of COSI’s field of view. Note that setting the cut to 180° yields the same results as 90° . This is expected because the Earth Horizon cut applied to the simulations removes all photons originating beyond 90° from zenith.

6.5.4.2 Extended sources

Extended two-dimensional Gaussian sources are again simulated with a flux of $0.2 \text{ ph cm}^{-2} \text{ s}^{-1}$ over the “step 6 degrees” orientation file. The 45° -, 90° -, 135° -, and 180° -rotated Gaussians contain 789, 714, 800, and 835 photons. The “X” and “plus” sources are shown in Figure 6.18 and the `cosipy` images are generated with 1000 s time bins. Qualitatively, the “X” `cosipy` image appears more poorly resolved in this case than in the “fine” orientation (Figure 6.15c). The lower right leg, in particular, is not as visible in this orientation. The “plus,” on the other hand, is arguably more clear, if a bit dimmer, than in the “fine” orientation (Figure 6.15d), though this judgment subjectively depends on the precise color scaling and the chosen “final” iteration. The “X” and “plus” shapes appear more well-defined in the `mimrec` images (Figures 6.18a and 6.18b) than in the `cosipy` images (Figures 6.18c and 6.18d), as was the case in previous tests. Overall, we find that COSI can image structures of diffuse emission as its orientation shifts steadily across the field of view.

6.5.5 Scanning the Galactic Plane

6.5.5.1 Point sources

As a final preliminary test of imaging, we create an orientation file that scans the Galactic Plane from $\ell = -100^\circ$ to 100° at $b = 0^\circ$ and then moves back from $\ell = 100^\circ$ to -100° at $b = 15^\circ$. The pointing changes $\sim 4^\circ$ every 100 s. Five point sources are simulated along this flight path: $(-99^\circ, -3^\circ)$, $(-51^\circ, -3^\circ)$, $(-3^\circ, -3^\circ)$, $(51^\circ, -3^\circ)$, $(99^\circ, -3^\circ)$. The total simulation time is 10^4 s and each source is simulated with a flux of $0.2 \text{ ph cm}^{-2} \text{ s}^{-1}$. The number of photons recorded in each source is 259, 491, 525, 506, and 283, respectively.

The simulation is binned into 100 s time bins and examined in `cosipy`. Each source moves in and out of the field of view as expected, given the scan from negative to positive longitude and reversal of course at slightly higher latitude (Figure 6.19). For example, the source at $(\ell = -99^\circ, b = -3^\circ)$ is maximally in the field of view (close to zenith) at the beginning of the simulation, falls out of the field of view and disappears as the instrument moves towards more positive longitudes, and later re-enters the field of view as the instrument returns to more negative longitudes. Note also that the source is more off-axis on its second appearance because the instrument is at latitude $b = 15^\circ$ rather than $b = 0^\circ$, placing it further from the

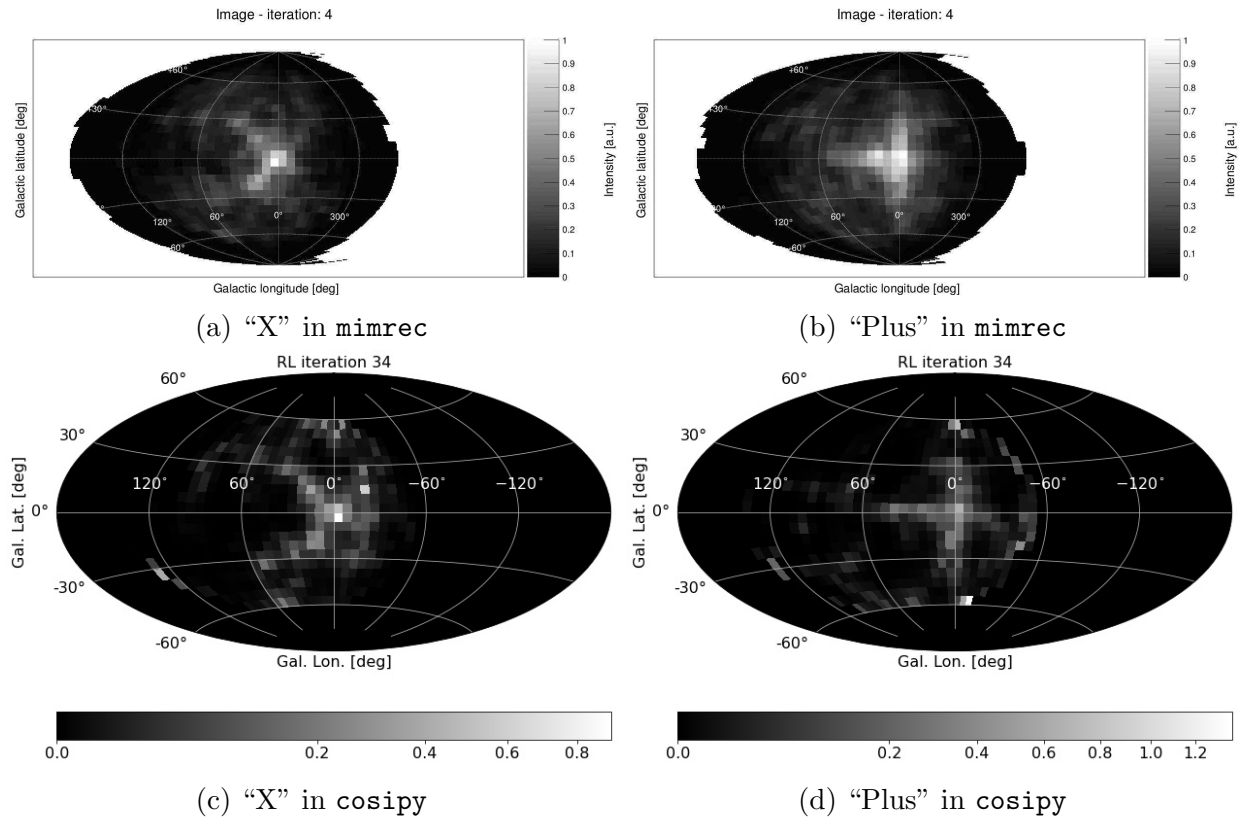


Figure 6.18: Images in (a, b) `mimrec` and (c, d) `cosipy` of “X”- and “plus”-shaped sources. The instrument orientation shifts 6° off-axis every 1000 s (“step 6 degrees” orientation, Section 6.5.4.2).

source. Similarly, the source at $(\ell = 99^\circ, b = -3^\circ)$ is maximally in the field of view about halfway through the observation time, when the instrument reaches its turn-around point at $\ell = 100^\circ$, and then immediately recedes from view as the instrument moves up to $b = 15^\circ$ and heads back to more negative longitudes.

All five point sources are clearly visible in `mimrec` and `cosipy` (Figure 6.20), despite an emerging artifact in the `cosipy` image near $\sim (170^\circ, -15^\circ)$. Generally, the relative brightness of the sources follows the number of simulated photons in each. The edge pixels are dimmer than the central three pixels, though in the `cosipy` image the source with the fewest photons (259) at $(-99^\circ, -3)$ appears noticeably brighter than the source at $(99^\circ, -3)$, which has 283 photons. These differences aside, the `mimrec` and `cosipy` images are in broad agreement and reconstruct the sources as expected.






(a) $\ell = -99^\circ$, $b = -3^\circ$ (b) $\ell = -51^\circ$, $b = -3^\circ$ (c) $\ell = -3^\circ$, $b = -3^\circ$ (d) $\ell = 51^\circ$, $b = -3^\circ$ (e) $\ell = 99^\circ$, $b = -3^\circ$

Figure 6.19: Appearance of point sources in COSI’s field of view as the instrument moves along a custom orientation file that scans the Galactic Plane (Section 6.5.5.1). The instrument moves from $\ell = -100^\circ$ to 100° at $b = 0^\circ$ and then moves back from $\ell = 100^\circ$ to -100° at $b = 15^\circ$.

6.5.5.2 Extended sources

Figure 6.21 shows images of the reconstructed “X” and “plus” sources as simulated with the scanning orientation of the instrument across the Galactic Plane. The simulated flux of each constituent Gaussian is $0.2 \text{ ph cm}^{-2} \text{ s}^{-1}$. The 45° -, 90° -, 135° -, and 180° -rotated Gaussians contain 494, 446, 491, and 496 photons. In *mimrec*, the “X” remains identifiable while the “plus” is more difficult to discern, compared to the “plus” in the “step 6 degrees” orientation (Figure 6.18b). As with the point source imaging, the *cosipy* images here are generated with 100 s time bins. The “X” is only visible at low iterations (iteration 10 is shown in Figure 6.21c), after which it devolves into disparate bright pixels. It is also necessary to mask artifacts in regions of poor exposure in the final image. At late iterations, the “plus” simulation develops a strong artifact near $(-120^\circ, -60^\circ)$. Though the “plus” is difficult to identify, iteration 15 is the clearest image in *cosipy* (Figure 6.21d). It is evident that diffuse




Figure 6.20: Images in (a) `mimrec` and (b) `cosipy` of five simulated point sources: $(-99^\circ, -3^\circ)$, $(-51^\circ, -3^\circ)$, $(-3^\circ, -3^\circ)$, $(51^\circ, -3^\circ)$, $(99^\circ, -3^\circ)$. The instrument orientation scans the Galactic Plane (Section 6.5.5.1).

imaging becomes more difficult with increasingly complex changes in instrument orientation.

6.5.6 Summary

Reconstruction algorithms in `mimrec` and `cosipy` are able to image simulated point sources and extended sources at ~ 1.8 MeV. The list-mode algorithm in `mimrec` appears to more stably reproduce the morphology of the underlying simulations compared to the RL algorithm in `cosipy`, which in these tests tends to produce slightly more blurred images. An advantage to using `cosipy`, though, is that its “bin-mode” imaging can accommodate exposure corrections and flux determination. The latter is under development. Nevertheless, the images from both algorithms are qualitatively in agreement and establish confidence in `cosipy`’s RL algorithm.

Several custom-made orientation files illustrate the potential difficulties of observing on a freely-floating balloon platform. Altitude changes and rotations about its azimuth at nighttime, when the rotator does not keep the instrument oriented towards the Sun, may complicate image reconstruction and blur the generated images. Even if making exact comparisons between the “simple,” “fine,” “step 6 degrees,” and “scanning the Galactic Plane” orientations is difficult due to somewhat subjective choices of a final iteration, it is clear that all cases of instrument motion are more blurred than when the instrument stares at sources with fixed orientation. This is an important consideration when studying flight data and is investigated further by running the RL algorithm in `cosipy` with finer time bins. Finer time bins better capture the motion of the instrument and mitigate smearing induced by binning together measurements taken at different orientations; however, running the above tests with smaller time binning does not appreciably change the results, probably due to the

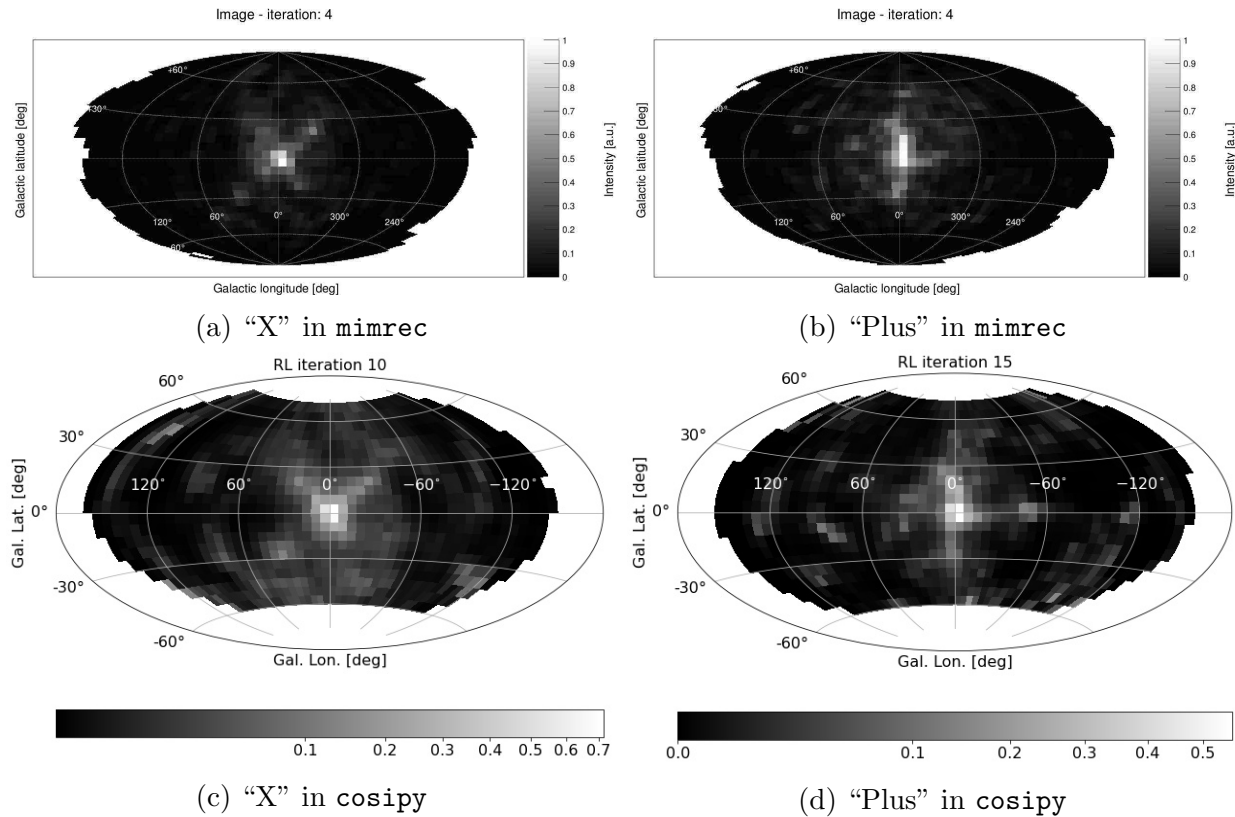


Figure 6.21: Images in (a, b) `mimrec` and (c, d) `cosipy` of “X”- and “plus”-shaped sources. The instrument orientation scans the Galactic Plane (Section 6.5.5.2).

coarse binning of the imaging response itself. Additionally, arbitrarily small time bins may not contain sufficient statistics for reliable imaging. Determining an ideal time bin width should be performed on a case-by-case basis. Taken together, the lessons and conclusions from these basic orientation tests inform the development of imaging with `cosipy` in the COSI Data Challenge 1, discussed in the next section.

6.6 Data Challenge 1

In 2022, the COSI team began developing a series of “data challenges” (Section 3.6.2) to release `cosipy` in stages to the astrophysics community and prepare for the COSI satellite launch in 2027. Data Challenge 1, published in early 2023 on GitHub, focuses on simulations of four point sources (the Crab Nebula, Centaurus A, Cygnus X-1, and Vela) and diffuse emission from Galactic positron-electron annihilation and ^{26}Al . All simulations mimic observations of these sources over the 46-day flight path of the COSI 2016 mission. For convenience, only the 9-detector mass model is used to account for the shut-off of detectors 0,

5, and 8 during flight and the simulations assume 33 km transmission probability (nominal balloon altitude). All sources are simulated at 10 times (denoted below by “10×”) their true flux values for increased statistics; the priority of Data Challenge 1 is to demonstrate analysis procedures on sources which can be readily interpreted, rather than benchmark sensitive flight analyses with more realistic simulations. Simulated atmospheric background, approximately scaled to flight levels, is also included. To maintain a relatively simple background environment in Data Challenge 1, simulations of instrumental activation are not included.

This section focuses on the imaging component of the data challenge, the notebooks for which were refined by the author of this dissertation after conducting the previous sections’ qualitative tests of `cosipy`’s RL algorithm. Preparing these notebooks for release required extensive fine tuning of specific parameters in the RL algorithm: the initial values of the fitted background parameter, fitted acceleration parameter, and the flux value of the initial isotropic map were determined in advance to ensure that the algorithm runs smoothly and produces reasonable images in each case. All simulations are binned in 1800 s time bins to strike a balance between preserving sufficient statistics in each bin, blurring over changes in instrument orientation, and computational strain. The response is set to zero at 90° off-axis from zenith to maximize sensitivity to photons across COSI’s field of view (see discussion in Section 6.5.4.1). This effort lays the groundwork for imaging in upcoming releases of `cosipy` and community analyses of COSI data.

6.6.1 Point sources

The Crab Nebula, Cygnus X-1, Centaurus A, and Vela are simulated with 10× their true flux values. The Crab Nebula ($\ell = 184.6^\circ, b = -5.8^\circ$) surrounds the Crab Pulsar produced by SN 1054. Though it is located in the Milky Way constellation Taurus and largely observable from Earth’s Northern Hemisphere, the Crab was detected during the largely Southern COSI 2016 flight with only 12 days of observation time (Sleator 2019). Its flux from other measurements between 100 keV–10 MeV is $0.049 \text{ ph cm}^{-2} \text{ s}^{-1}$. Cygnus X-1, likely a black hole in an X-ray binary system, is located at ($\ell = 71.3^\circ, b = 3.1^\circ$) in the Cygnus constellation of the Milky Way and is visible in γ -rays with a flux of $0.041 \text{ ph cm}^{-2} \text{ s}^{-1}$ between 100 keV–10 MeV. The supermassive black hole within Centaurus A (NGC 5128, a galaxy in the Centaurus constellation at ($\ell = 309.5^\circ, b = 19.4^\circ$)) has a flux of $0.0036 \text{ ph cm}^{-2} \text{ s}^{-1}$ in this energy range. Lastly, the fainter Vela pulsar ($\ell = 263.6^\circ, b = -2.8^\circ$) is simulated by extrapolating Fermi-LAT measurements to lower energies consistent with previous observations by the OSSE, COMPTEL, and EGRET missions (Pavlov et al. 2001). The measured flux between 100 keV–10 MeV is $1.2 \times 10^{-4} \text{ ph cm}^{-2} \text{ s}^{-1}$.

Each source is simulated individually between 100 keV–10 MeV and studied up to 5 MeV, in accordance with the upper bound of COSI’s energy range. A $6 \times 6 \text{ deg}^2$ continuum response matrix with energy bins defined by [150, 220, 325, 480, 520, 765, 1120, 1650, 2350, 3450, and 5000] keV is also provided for the analysis. All images are generated using data from the 325–480 keV energy bin, which has the most counts. Future iterations of `cosipy` will adapt the algorithm to accommodate multiple energy bins.




Figure 6.22: The combined (a) lightcurve and (b) energy spectrum in `cosipy` of the Crab Nebula ($\ell = 184.6^\circ, b = -5.8^\circ$), Cygnus X-1 ($\ell = 71.3^\circ, b = 3.1^\circ$), Centaurus A ($\ell = 309.5^\circ, b = 19.4^\circ$), and Vela ($\ell = 263.6^\circ, b = -2.8^\circ$), each at $10\times$ their true flux with no background (Section 6.6.1.1).

6.6.1.1 No background

The preferred initial background parameter, acceleration parameter, and flux value of the isotropic map for RL imaging of all four point sources together, without background, are 10^{-6} , 2000, and 1.0. These are the same parameters used in the algorithm tests from Section 6.5. They appear to work well in cases of strong sources with no background. The light curve and energy spectrum are shown in Figure 6.22. Dips in the lightcurve as the sources, dominated by the Crab Nebula, move in and out of the field of view are visible. The 325–480 keV energy bin has the most counts, hence the choice to perform imaging in this range. The images in `mimrec` and `cosipy` are shown in Figure 6.23 and the convergence of the background parameter, acceleration parameter, and likelihood is shown in Figure 6.24. The `mimrec` and `cosipy` images qualitatively agree and localize the sources to their expected positions. The Crab Nebula correctly appears on the right-hand side of the image as the brightest source. It appears to bleed into a few pixels on the left-hand side of the `cosipy` image, but not in the `mimrec` image. Cygnus X-1 and Centaurus A are also visible, though Vela is too dim to be seen. The acceleration parameter largely converges at late iterations, and the likelihood converges cleanly.

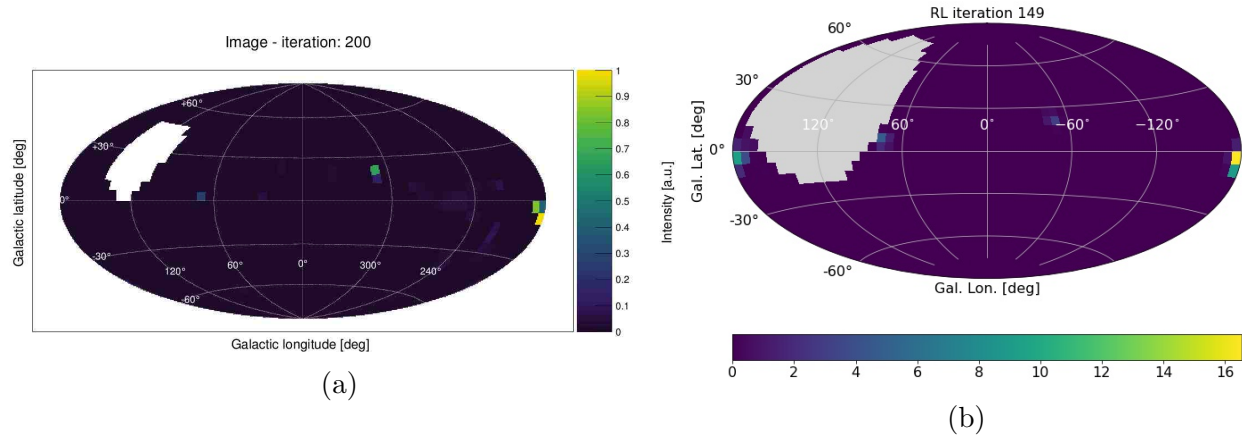


Figure 6.23: Images in (a) `mimrec` and (b) `cosipy` of the Crab Nebula ($\ell = 184.6^\circ, b = -5.8^\circ$), Cygnus X-1 ($\ell = 71.3^\circ, b = 3.1^\circ$), Centaurus A ($\ell = 309.5^\circ, b = 19.4^\circ$), and Vela ($\ell = 263.6^\circ, b = -2.8^\circ$), each at $10\times$ their true flux with no background (Section 6.6.1.1).

6.6.1.2 Including background

The preferred initial background parameter, acceleration parameter, and flux value of the isotropic map for imaging all four point sources together, including an atmospheric (Ling model) background simulation, are 0.9, 2000, and 0.01. The higher background parameter is roughly akin to the fractional contribution of background to the total simulation, i.e. the sources are heavily dominated by atmospheric background. Lowering the initial flux of the isotropic map is also necessary to run the RL algorithm successfully in the presence of high background. Choosing a higher initial value, like 1.0 as in the case without background, causes the likelihood fit in the RL algorithm to fail. The light curve and energy spectrum are shown in Figure 6.25. The lightcurve remains at a more consistent level because atmospheric background is constantly present in the observations, in contrast to the varying presence of the point sources moving in and out of the field of view. The energy spectrum now peaks in the 480–520 keV bin because much of the atmospheric background emission is at 511 keV. To avoid this strong background line, we again image the sky between 325–480 keV.

As seen in Figure 6.26a, the `mimrec` algorithm struggles to find the point sources and instead images a large area around the region of maximum exposure, suggesting that it is imaging the dominant atmospheric background. There is, though, a potential bright spot near the location of the Crab Nebula. The `cosipy` RL algorithm is more successful. In the presence of heavy background, only the Crab Nebula is readily visible (Figure 6.26b). Changing the color scaling as in Figure 6.26c reveals hints of emission near the other point sources. Figure 6.26d is included as an illustration of the final image when regions of bad exposure are not masked. The algorithm settles into a single pixel artifact with ever-increasing flux. Overall, it is more difficult to image the sources in the presence of strong atmospheric background, but the Crab Nebula is bright enough to be recovered in this setting.

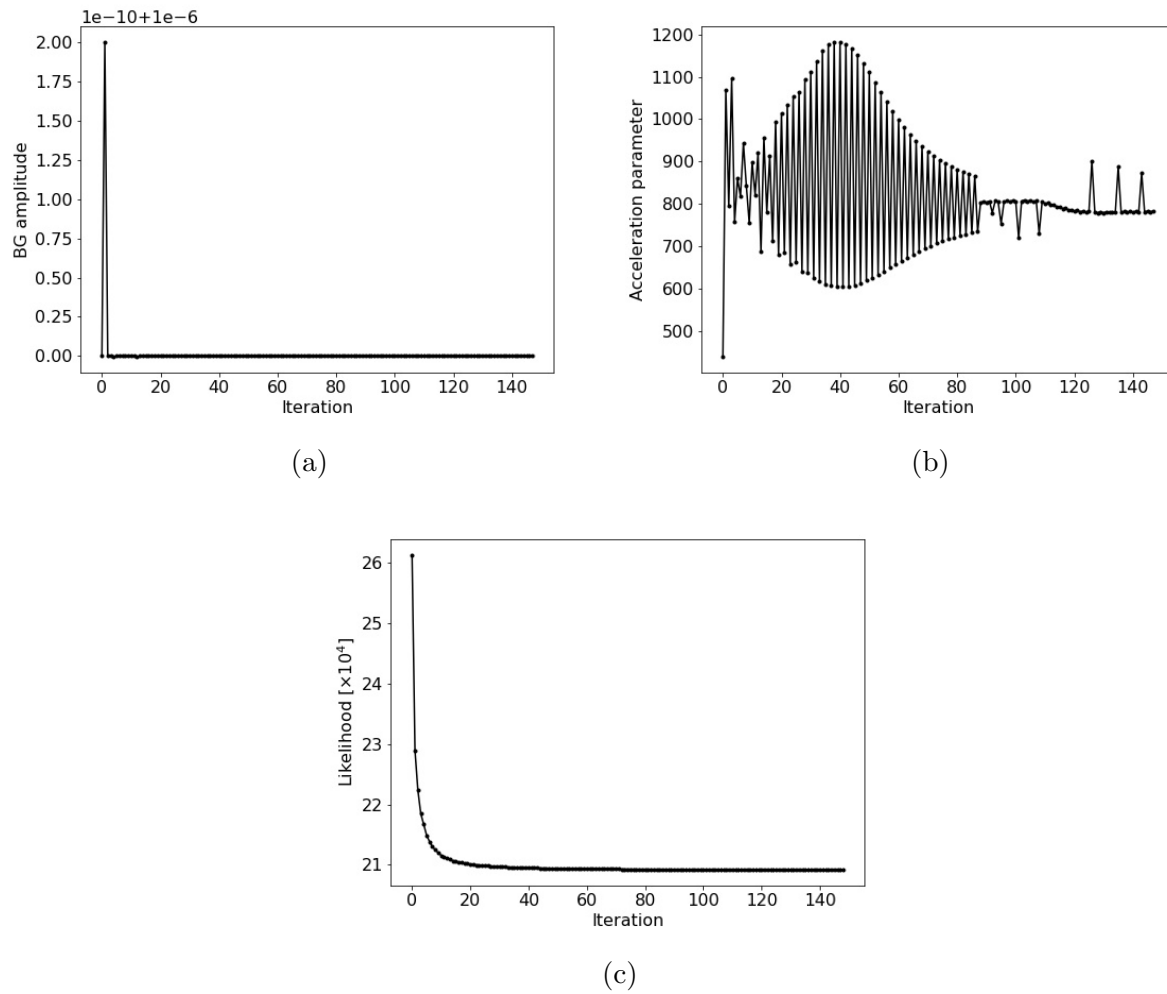


Figure 6.24: Convergence of the background amplitude, acceleration parameter, and likelihood of the Crab Nebula ($\ell = 184.6^\circ, b = -5.8^\circ$), Cygnus X-1 ($\ell = 71.3^\circ, b = 3.1^\circ$), Centaurus A ($\ell = 309.5^\circ, b = 19.4^\circ$), and Vela ($\ell = 263.6^\circ, b = -2.8^\circ$) without background in the *cosipy* RL algorithm (Section 6.6.1.1). The negative likelihood is plotted.

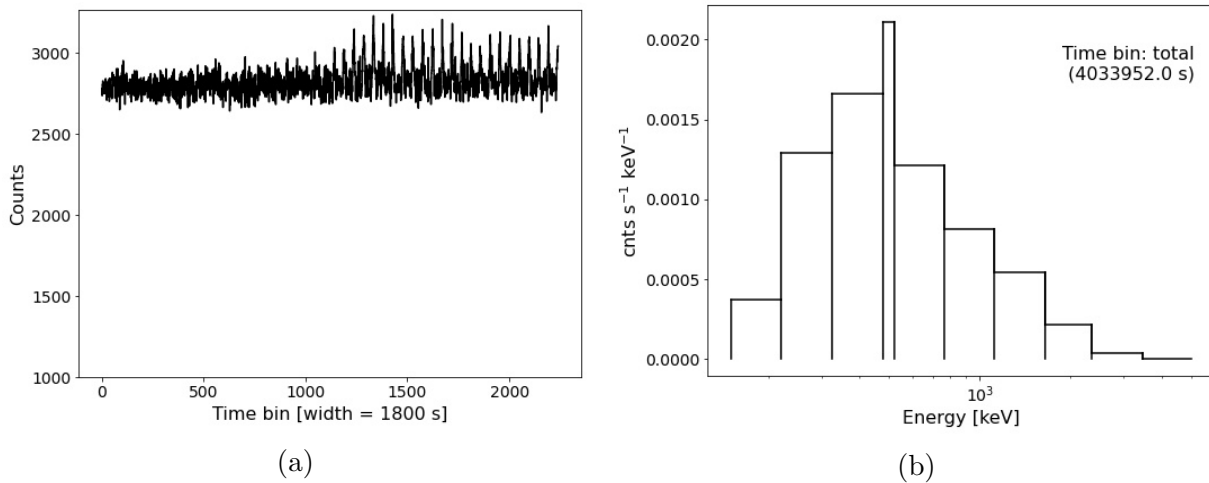


Figure 6.25: The combined (a) lightcurve and (b) energy spectrum in `cosipy` of the Crab Nebula ($\ell = 184.6^\circ, b = -5.8^\circ$), Cygnus X-1 ($\ell = 71.3^\circ, b = 3.1^\circ$), Centaurus A ($\ell = 309.5^\circ, b = 19.4^\circ$), and Vela ($\ell = 263.6^\circ, b = -2.8^\circ$), each at $10\times$ their true flux with atmospheric background (Section 6.6.1.2).

Figure 6.27 illustrates the convergence of the fitted background parameter, acceleration parameter, and likelihood. The background parameter settles into a value near 0.9–1, indicating the expected dominance of background in the simulation. The acceleration parameter undergoes oscillations but settles into a stable value in the latter half of the iterations, and the likelihood clearly converges.

6.6.2 Positron-electron annihilation

Next, we image positron-electron annihilation emission at 511 keV in the vicinity of the Galactic Center. The emission is simulated following the two-dimensional Gaussian surface brightness distribution of INTEGRAL/SPI observations in Knödlseher et al. (2005): The Gaussian is centered at ($\ell = -0.6^\circ, b = 0.1^\circ$) and extends in longitude with $\sigma_\ell = 3.5^\circ$ and in latitude with $\sigma_b = 3.1^\circ$. The flux is reported as $1.1 \times 10^{-3} \text{ ph cm}^{-2} \text{ s}^{-1}$ and is simulated for these imaging tests as a monoenergetic source at 511.0 keV with $1.1 \times 10^{-2} \text{ ph cm}^{-2} \text{ s}^{-1}$. The imaging is performed with a simulated $6 \times 6 \text{ deg}^2$ response matrix with one energy bin spanning 501–521 keV.

6.6.2.1 No background

The initial background parameter, acceleration parameter, and flux value of the isotropic map are set to 10^{-6} , 1000, and 0.01. Lowering the acceleration parameter stabilized the progression through iterations, and reducing the flux of the initial isotropic map may have

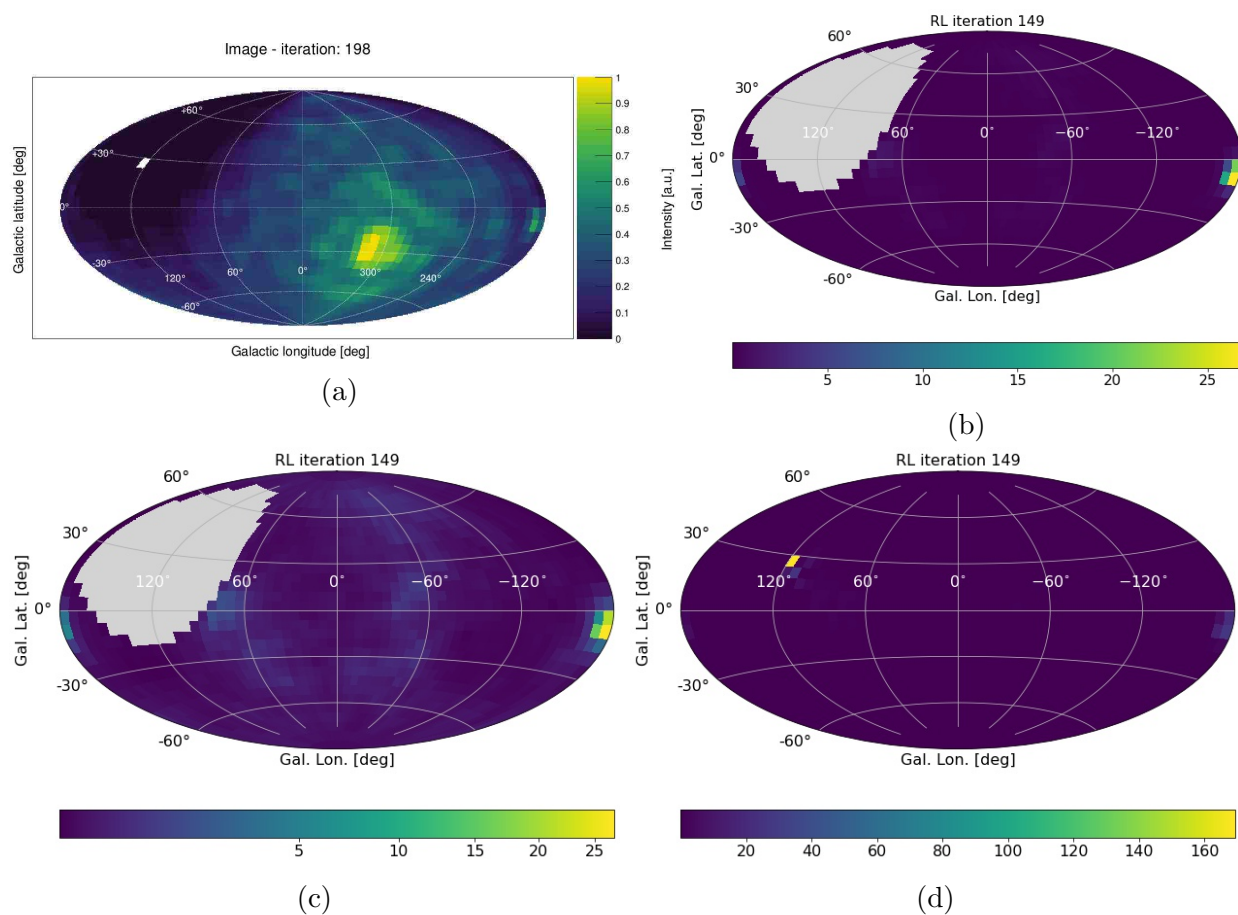


Figure 6.26: Images in (a) *mimrec* and (b) *cosipy* of the Crab Nebula ($\ell = 184.6^\circ, b = -5.8^\circ$), Cygnus X-1 ($\ell = 71.3^\circ, b = 3.1^\circ$), Centaurus A ($\ell = 309.5^\circ, b = 19.4^\circ$), and Vela ($\ell = 263.6^\circ, b = -2.8^\circ$), each at $10\times$ their true flux with atmospheric background (Section 6.6.1.2). The image in (c) is the same as in (b) with different color scaling. The image in (d) shows that the final iteration settles into a single artifact of excessively high flux without masking regions of poor exposure.

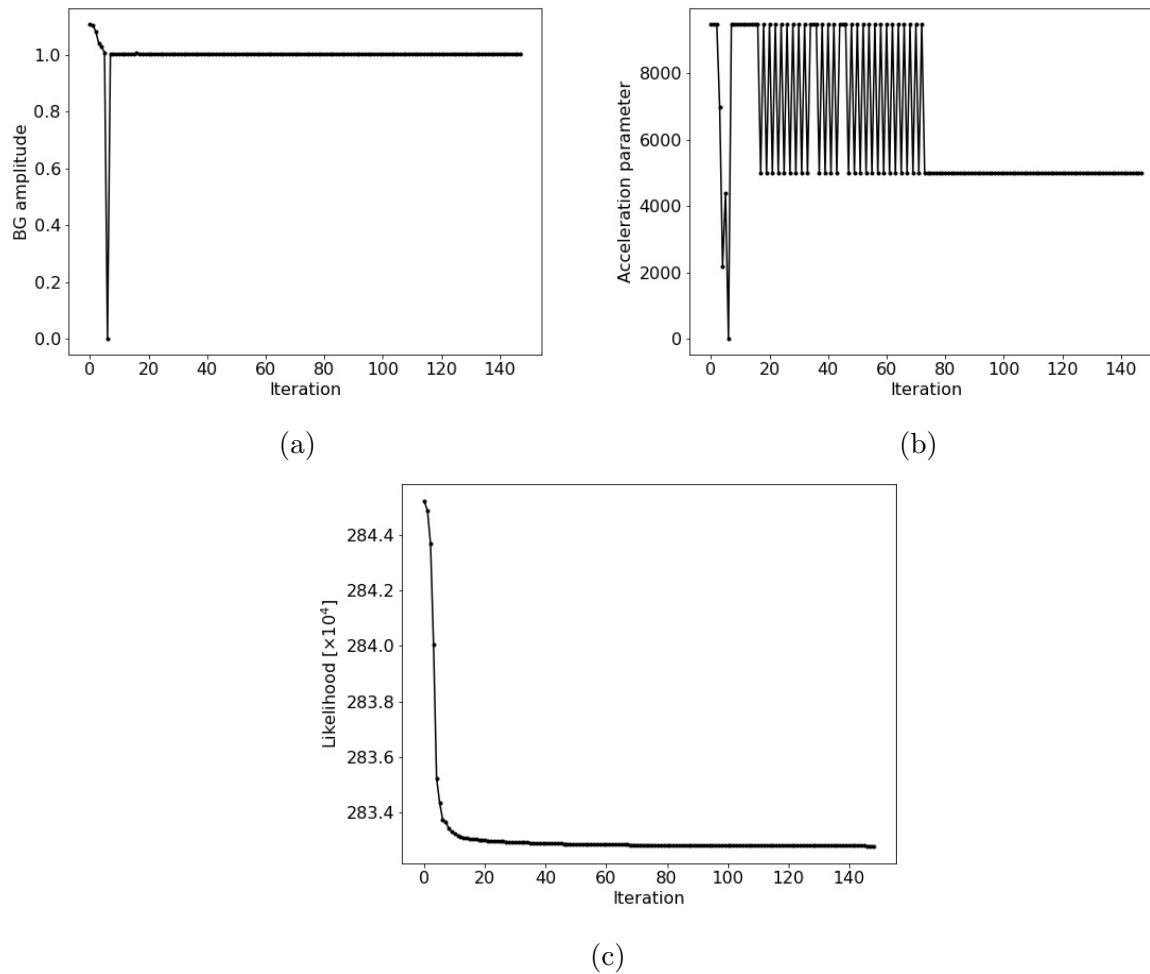


Figure 6.27: Convergence of the background amplitude, acceleration parameter, and likelihood of the Crab Nebula ($\ell = 184.6^\circ, b = -5.8^\circ$), Cygnus X-1 ($\ell = 71.3^\circ, b = 3.1^\circ$), Centaurus A ($\ell = 309.5^\circ, b = 19.4^\circ$), and Vela ($\ell = 263.6^\circ, b = -2.8^\circ$) with atmospheric background in the *cosipy* RL algorithm (Section 6.6.1.2). The negative likelihood is plotted.

been necessary to capture the fainter gradients of the diffuse morphology. The simulated count rate ($10\times$ true 511 keV flux) is $\sim 0.006 \text{ ph s}^{-1}$. Kierans (2018) reports a background-subtracted 511 keV rate of $\sim (8.2 \pm 0.8) \times 10^{-4} \text{ ph s}^{-1}$ in the COSI 2016 balloon flight, consistent within about 2σ uncertainties with $0.1X$ the simulated rate. The discrepancy may be attributed to an imperfect imaging response, atmospheric broadening of the image, or other systematic uncertainties. The images in `mimrec` and `cosipy` are shown in Figure 6.28 and the convergence of the background parameter, acceleration parameter, and likelihood is shown in Figure 6.29. The `mimrec` and `cosipy` images qualitatively agree, localizing the emission to the strong “bulge” at the center of the Galaxy. Figure 6.28c fits the emission with a two-dimensional Gaussian and shows contour levels at 1%, 40%, and 80% of the maximum flux of the best-fit Gaussian. The close proximity of the contours indicates that the flux is strongly localized to the central few pixels. Indeed, the fitted parameters of the best-fit Gaussian are $\ell_0 \sim -2.5^\circ$, $b_0 \sim -2.7^\circ$, $\sigma_\ell \sim 2.2^\circ$ (FWHM $\sim 5.1^\circ$), and $\sigma_b \sim 2.5^\circ$ (FWHM $\sim 5.8^\circ$). The FWHM of this distribution is smaller than the simulated input of $\Delta\ell$ (FWHM) $\sim 8^\circ$, Δb (FWHM) $\sim 7^\circ$ from Knödlseeder et al. (2005).

6.6.2.2 Including background

The initial background parameter, acceleration parameter, and flux value of the isotropic map for imaging the 511 keV distribution with background are set to 0.99, 1000, and 0.01. As in the point source test with background (Section 6.6.1.2), the initial background parameter is increased to reflect dominant background in the simulation. The simulated count rate ($10\times$ true 511 keV flux plus background scaled to flight level between 501–521 keV) is $\sim 0.06 \text{ ph s}^{-1}$. Siegert et al. (2020) also reports a rate of $\sim 0.06 \text{ ph s}^{-1}$ between 506–516 keV in the balloon flight, including background photons. The consistency in these values suggests that the background rate dominates the emission regardless of true or $10\times$ true positron-electron annihilation flux. Figure 6.30 shows the images in `mimrec` and `cosipy` with a fit to the Gaussian distribution. The convergence of the background parameter, acceleration parameter, and likelihood is shown in Figure 6.31. The temporary sharp drops which deviate from the broader curves are from failed fits in the RL algorithm which then recovered in the next iteration. As in the case without background, the `mimrec` and `cosipy` images are in agreement. The “bulge” is localized to the center of the Galaxy. The distribution appears more diffuse in the `cosipy` image with background (Figure 6.30b) than without background (Figure 6.28b). This is corroborated by a two-dimensional Gaussian fit to the emission with background (Figure 6.30c). The 1%, 40%, and 80% levels are more separated, indicating a broader distribution of flux. The fitted parameters are $\ell_0 \sim -2.7^\circ$, $b_0 \sim -1.8^\circ$, $\sigma_\ell \sim 5.2^\circ$ (FWHM $\sim 12.3^\circ$), and $\sigma_b \sim 4.9^\circ$ (FWHM $\sim 11.5^\circ$).

Overall, imaging positron-electron annihilation with `cosipy` in Data Challenge 1 reveals the bright bulge of 511 keV emission around the Galactic Center. The extended disk emission seen in SPI analyses is not visible here, though this is expected given COSI’s limited 46-day observation time; SPI saw about 1 photon per week from the disk and has over a decade of

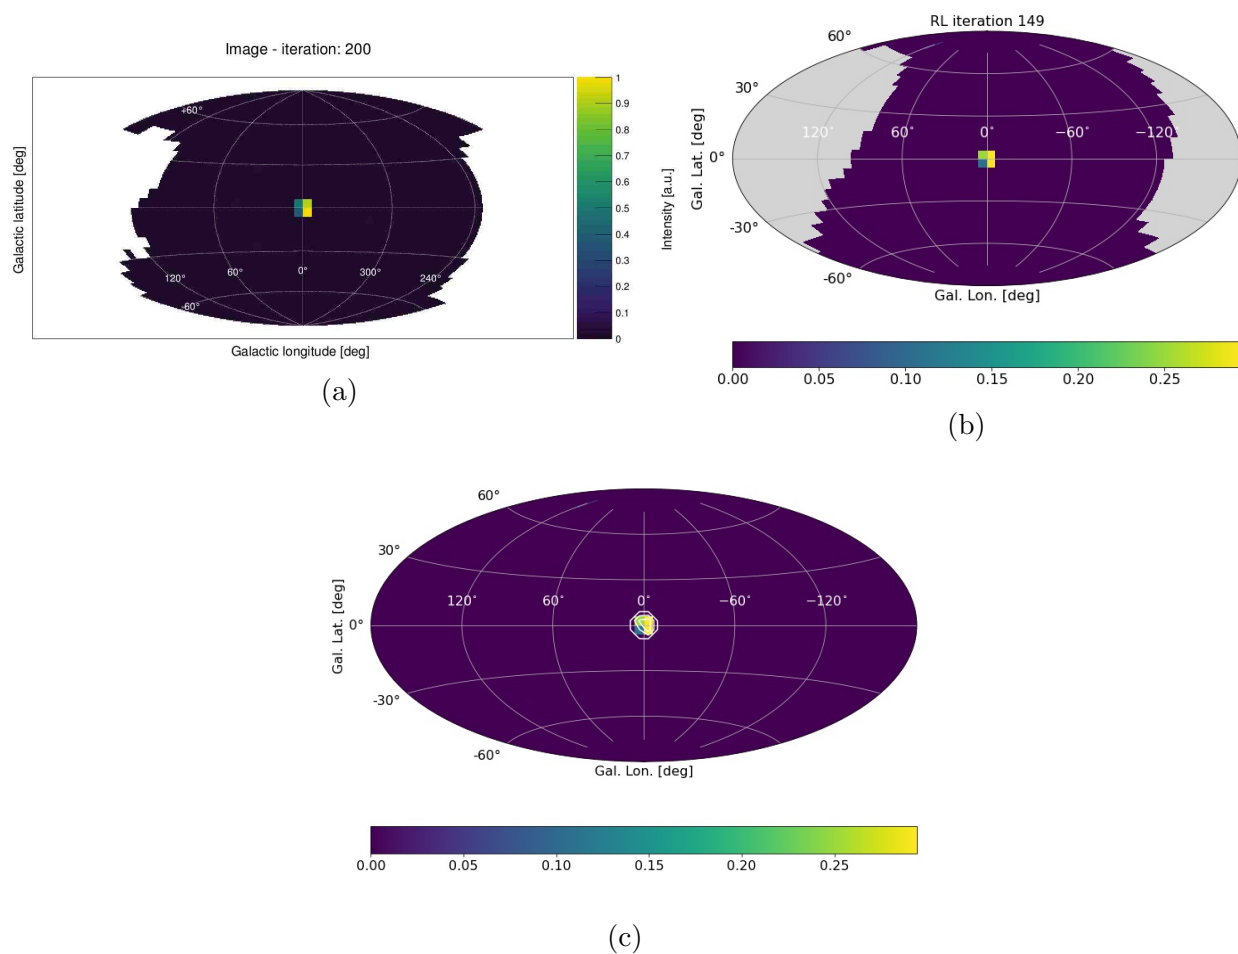


Figure 6.28: Images in (a) *mimrec* and (b) *cosipy* of 511 keV emission, simulated at $10\times$ its true flux with no background (Section 6.6.2.1). The image in (c) fits the emission with a two-dimensional Gaussian. Contour levels at 1%, 40%, and 80% of the maximum flux of the best-fit Gaussian are shown.

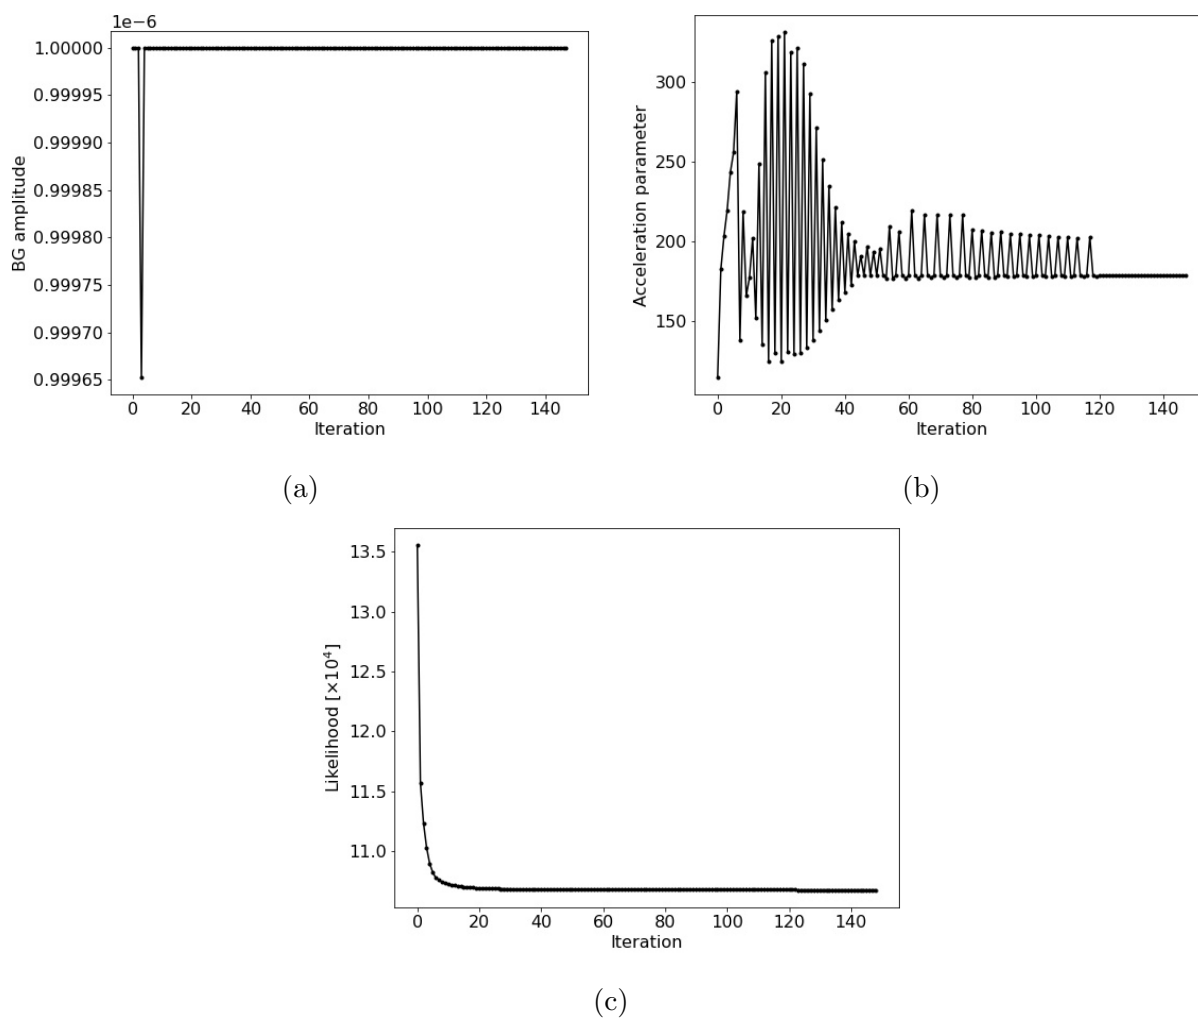


Figure 6.29: Convergence of the background amplitude, acceleration parameter, and likelihood of 511 keV emission without background in the `cosipy` RL algorithm (Section 6.6.2.1). The negative likelihood is plotted.

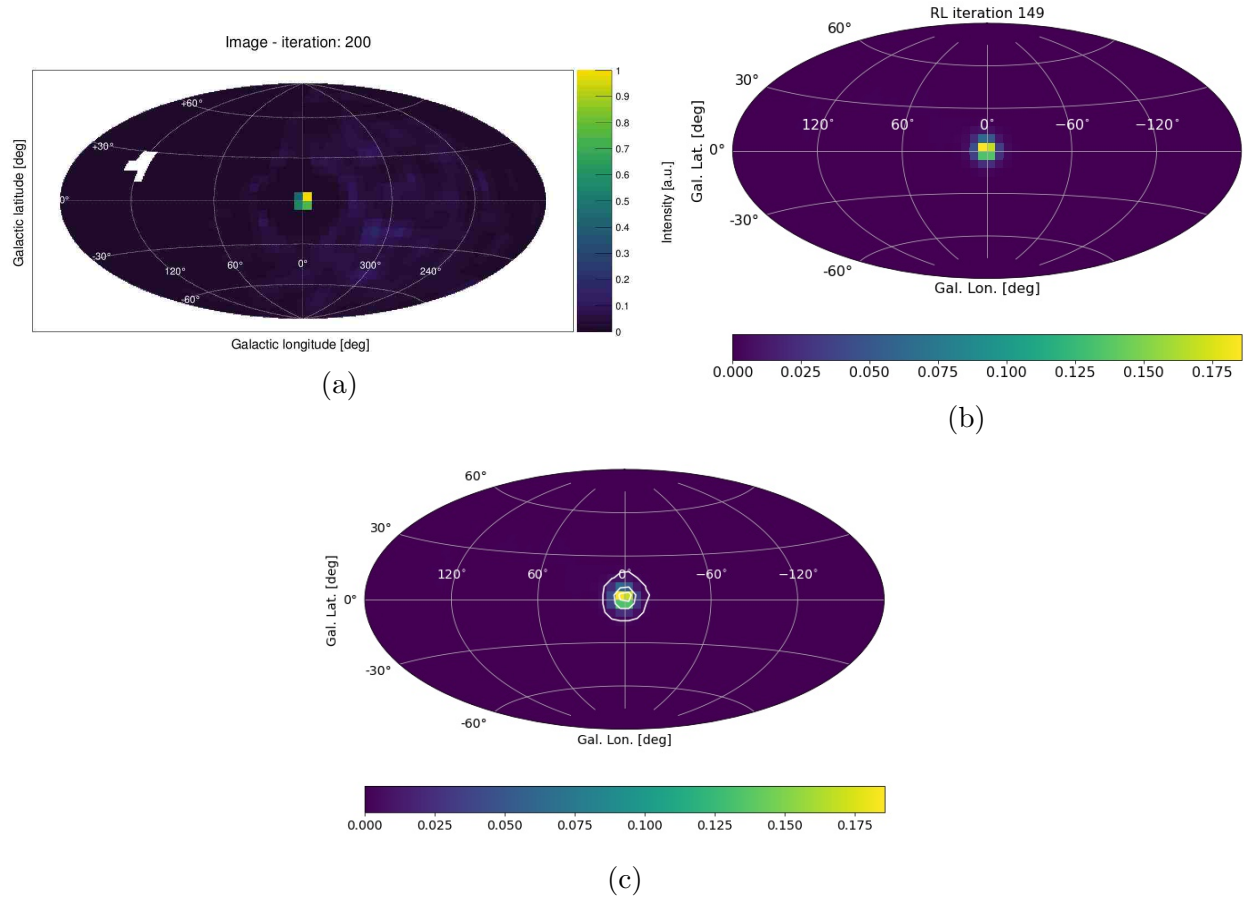


Figure 6.30: Images in (a) *mimrec* and (b) *cosipy* of 511 keV emission, simulated at $10\times$ its true flux with atmospheric background (Section 6.6.2.2). The image in (c) fits the emission with a two-dimensional Gaussian. Contour levels at 1%, 40%, and 80% of the maximum flux of the best-fit Gaussian are shown.

observation time. This test demonstrates that *cosipy* can image known diffuse sources with and without background over real flight orientations.

6.6.3 Aluminum-26

Finally, we image ^{26}Al emission at 1809 keV along the Galactic Plane. As in the spectroscopic analysis of Chapter 5, the DIRBE $240\ \mu\text{m}$ image is assumed as a template map. The flux of the DIRBE map is scaled to $10\times$ the observed value from COMPTEL ($10\times$ Inner Galaxy flux = $3.3 \times 10^{-3}\ \text{ph cm}^{-2}\ \text{s}^{-1}$; $10\times$ total map flux = $1.1 \times 10^{-2}\ \text{ph cm}^{-2}\ \text{s}^{-1}$) and the source is defined to be monoenergetic at 1809 keV. The response matrix is $6 \times 6\ \text{deg}^2$ with one energy bin spanning 1803–1817 keV.

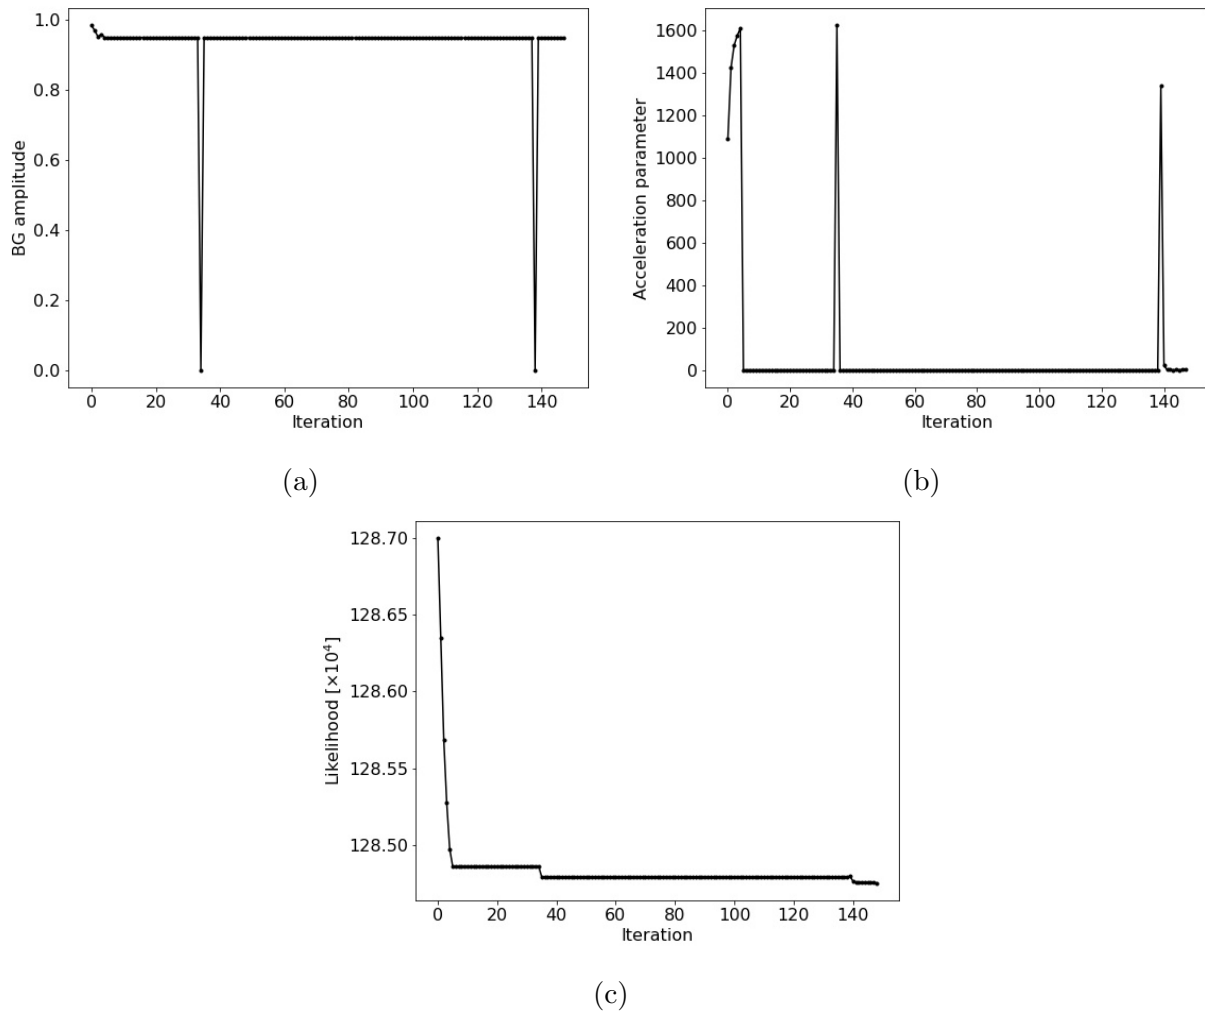


Figure 6.31: Convergence of the background amplitude, acceleration parameter, and likelihood of 511 keV emission with atmospheric background in the *cosipy* RL algorithm (Section 6.6.2.2). The negative likelihood is plotted.

6.6.3.1 No background

The initial background parameter, acceleration parameter, and flux value of the isotropic map are set to 10^{-6} , 2000, and 0.1. The simulated count rate ($10\times$ true ^{26}Al flux) is $\sim 1.8 \times 10^{-3}$ ph s $^{-1}$. Note that the background-subtracted rate in the COSI balloon flight of 6.8×10^{-4} ph s $^{-1}$ spans 1750–1850 keV, but most of the emission is concentrated in the 1.8 MeV line (Chapter 5). The background parameter plot (Figure 6.33) is shown starting at iteration 10 to mask a single failure in the RL algorithm in an early iteration. Again, the background parameter remains near its initial guess and the acceleration parameter and likelihood converge.

The `mimrec` and `cosipy` images in Figure 6.32 both reveal low-latitude emission extending along the Galactic Plane, as expected. In both images, there is a bright pixel near ($\ell \sim 22^\circ, b \sim 3^\circ$). An artifact emerges in the `cosipy` image near ($\ell \sim 70^\circ, b \sim 30^\circ$) and becomes stronger at later iterations. Iteration 50 (Figure 6.32b) is chosen as a middle-ground between sufficient iterations for convergence and excessive iterations which break up the image into disparate bright pixels across the sky. DIRBE 240 μm contours (5%, 10%, 50%, and 80% levels) are also plotted to show that the `cosipy` image reasonably recovers its underlying tracer. Figure 6.32c fits the emission with a two-dimensional Gaussian. As expected, the contours (1%, 40%, and 80% of the maximum flux of the best-fit Gaussian) extend to greater longitudes than in the case of concentrated 511 keV emission (Figure 6.28c), indicating brighter ^{26}Al emission along the Galactic Plane than from the dimmer 511 keV disk. The fitted parameters of the best-fit Gaussian are $\ell_0 \sim -1.5^\circ$, $b_0 \sim -3.0^\circ$, $\sigma_\ell \sim 46.4^\circ$ (FWHM $\sim 109.3^\circ$), and $\sigma_b \sim 3.3^\circ$ (FWHM $\sim 7.7^\circ$).

6.6.3.2 Including background

The initial background parameter, acceleration parameter, and flux value of the isotropic map of the ^{26}Al distribution with atmospheric background are set to 0.9, 2000, and 0.01. The convergence of these parameters is shown in Figure 6.35. The simulated count rate ($10\times$ true ^{26}Al flux plus background scaled to flight level between 1803–1817 keV) is ~ 0.006 ph s $^{-1}$.

As in the case without background, the `mimrec` and `cosipy` images in Figure 6.34 reveal emission along the Galactic Plane, though the images with background have a smoother distribution with fewer disparate pixels than the images without background (Figure 6.32). The background photons are likely populating the region and making it appear more diffuse. The bright pixel near ($\ell \sim 22^\circ, b \sim 3^\circ$) seen in the images without background is visible with background as well. Again we see that even when including background, the DIRBE 240 μm contours (5%, 10%, 50%, and 80% levels) trace the `cosipy` image reasonably well. However, a simple two-dimensional fit to the emission when including atmospheric background gives a FWHM in longitude of approximately 116° and a FWHM in latitude of about 26° . The extent in latitude is much larger than in the case without background, probably owing to the covariance between source and background. The two-dimensional fit also appears tilted

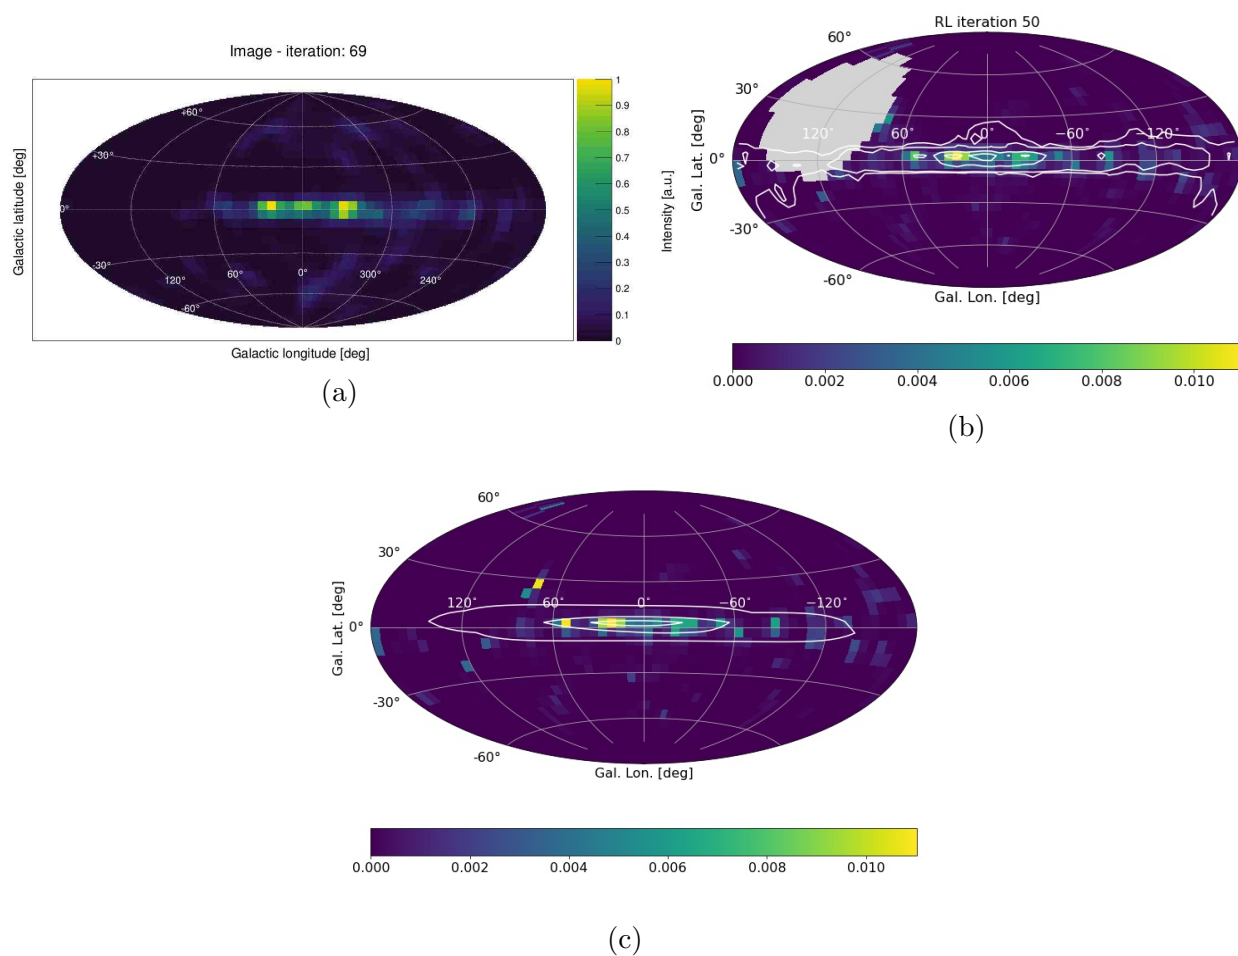


Figure 6.32: Images in (a) *mimrec* and (b) *cosipy* of the 1809 keV signature of ^{26}Al , simulated at $10\times$ its true flux with no background (Section 6.6.3.1). The image in (b) is shown with contours of the DIRBE 240 μm map at 5%, 10%, 50%, and 80% levels. The image in (c) fits the emission with a two-dimensional Gaussian. Contour levels at 1%, 40%, and 80% of the maximum flux of the best-fit Gaussian are shown.

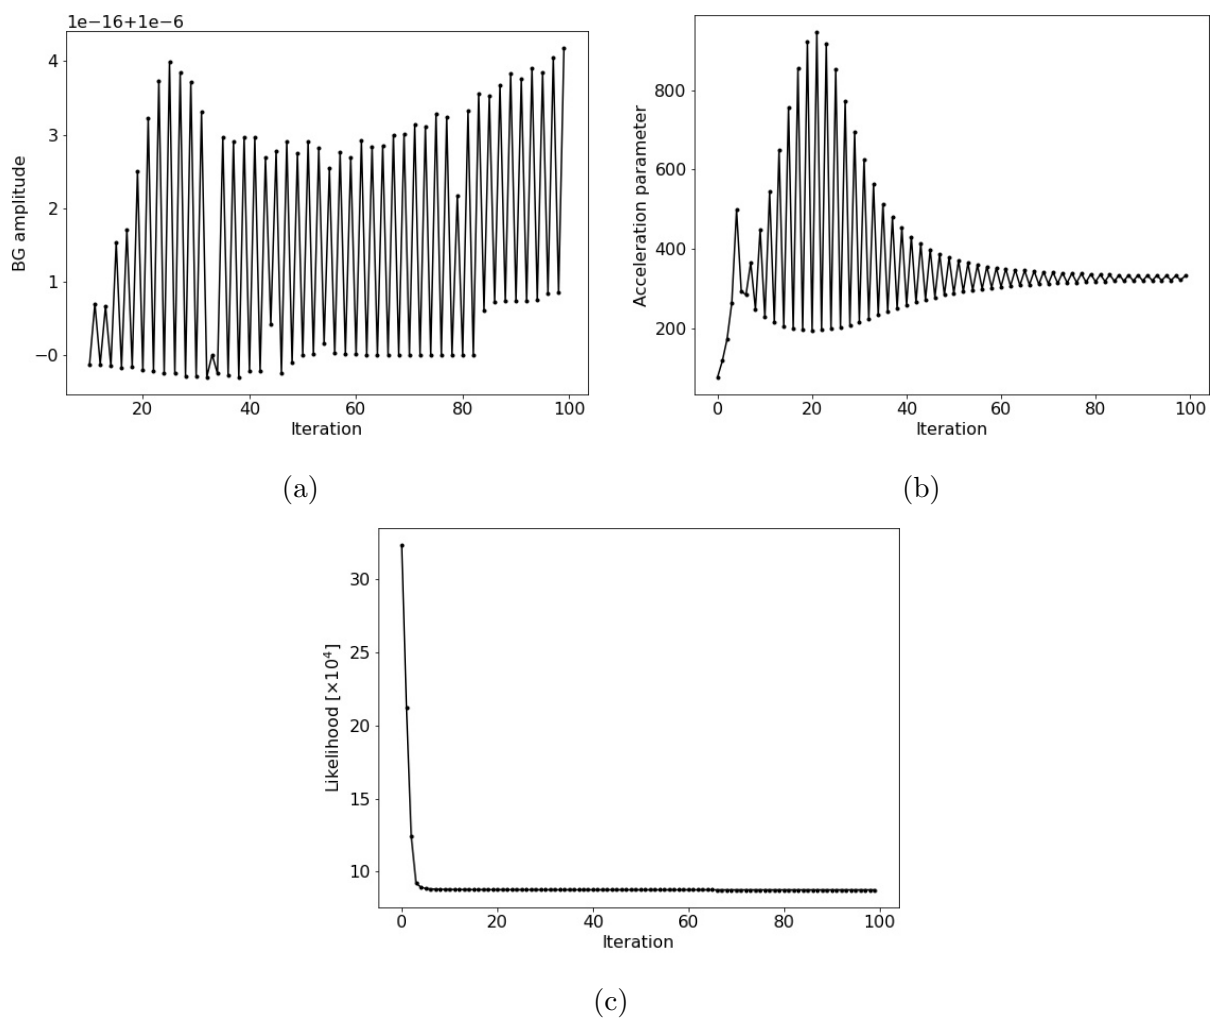


Figure 6.33: Convergence of the background amplitude, acceleration parameter, and likelihood of ^{26}Al emission without background in the *cosipy* RL algorithm (Section 6.6.3.1). The negative likelihood is plotted.

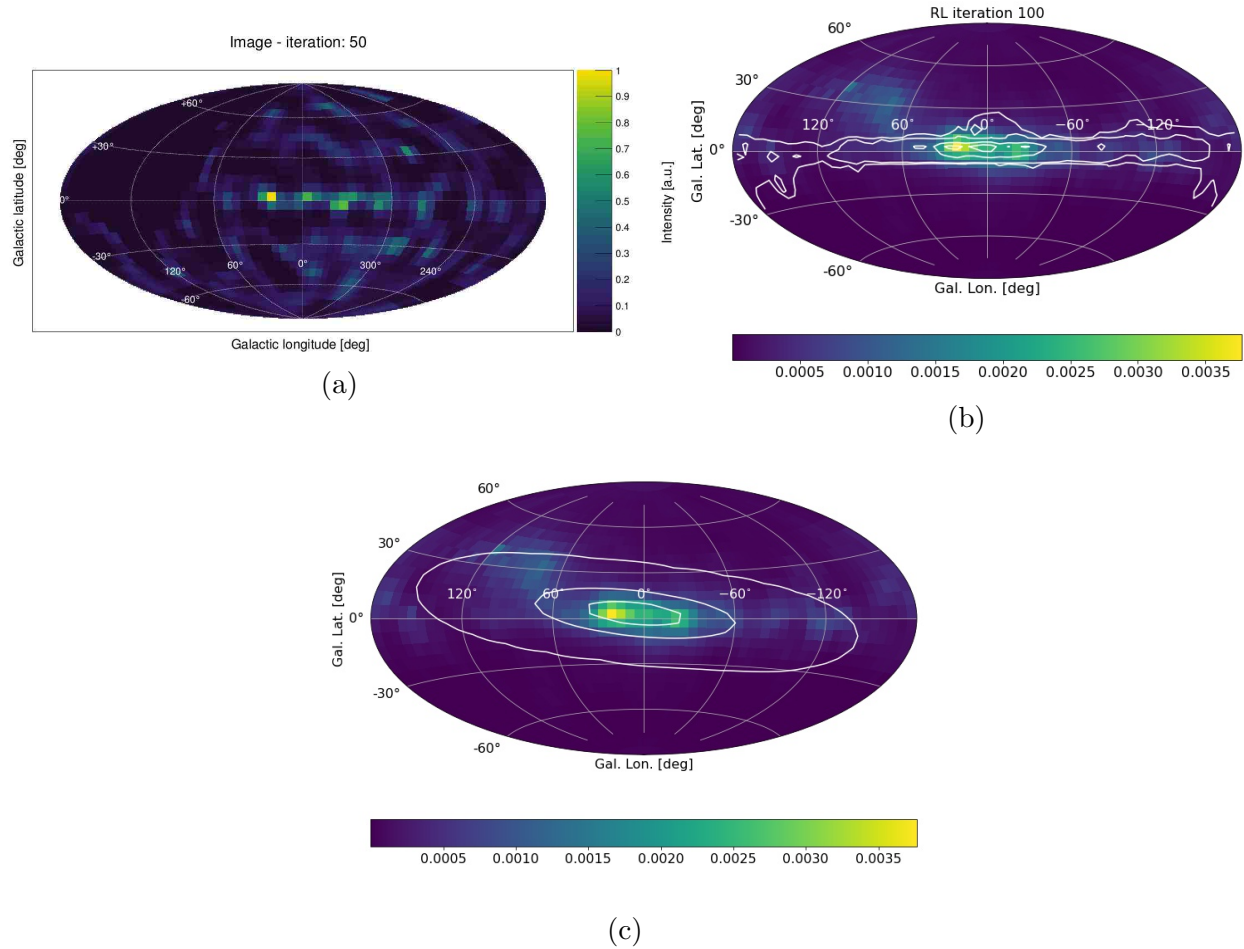


Figure 6.34: Images in (a) *mimrec* and (b) *cosipy* of the 1809 keV signature of ^{26}Al , simulated at $10\times$ its true flux with atmospheric background (Section 6.6.3.2). The image in (b) is shown with contours of the DIRBE 240 μm map at 5%, 10%, 50%, and 80% levels. The image in (c) fits the emission with a two-dimensional Gaussian. Contour levels at 1%, 40%, and 80% of the maximum flux of the best-fit Gaussian are shown.

off of the Galactic Plane; it may be influenced by a diffuse artifact near $\ell \sim 100^\circ$, $b \sim 30^\circ$.

For this reason, we do not over-interpret the fit to this image and instead emphasize the goal of Data Challenge 1: the RL algorithm in *cosipy* gives qualitatively sensible results that introduce users to the primary targets of MeV γ -ray science. With and without background, reconstructed ^{26}Al emission is visible along the Galactic Plane and with the right shape, acknowledging the effect of background variations and that this section imaged only one realization of the simulation.

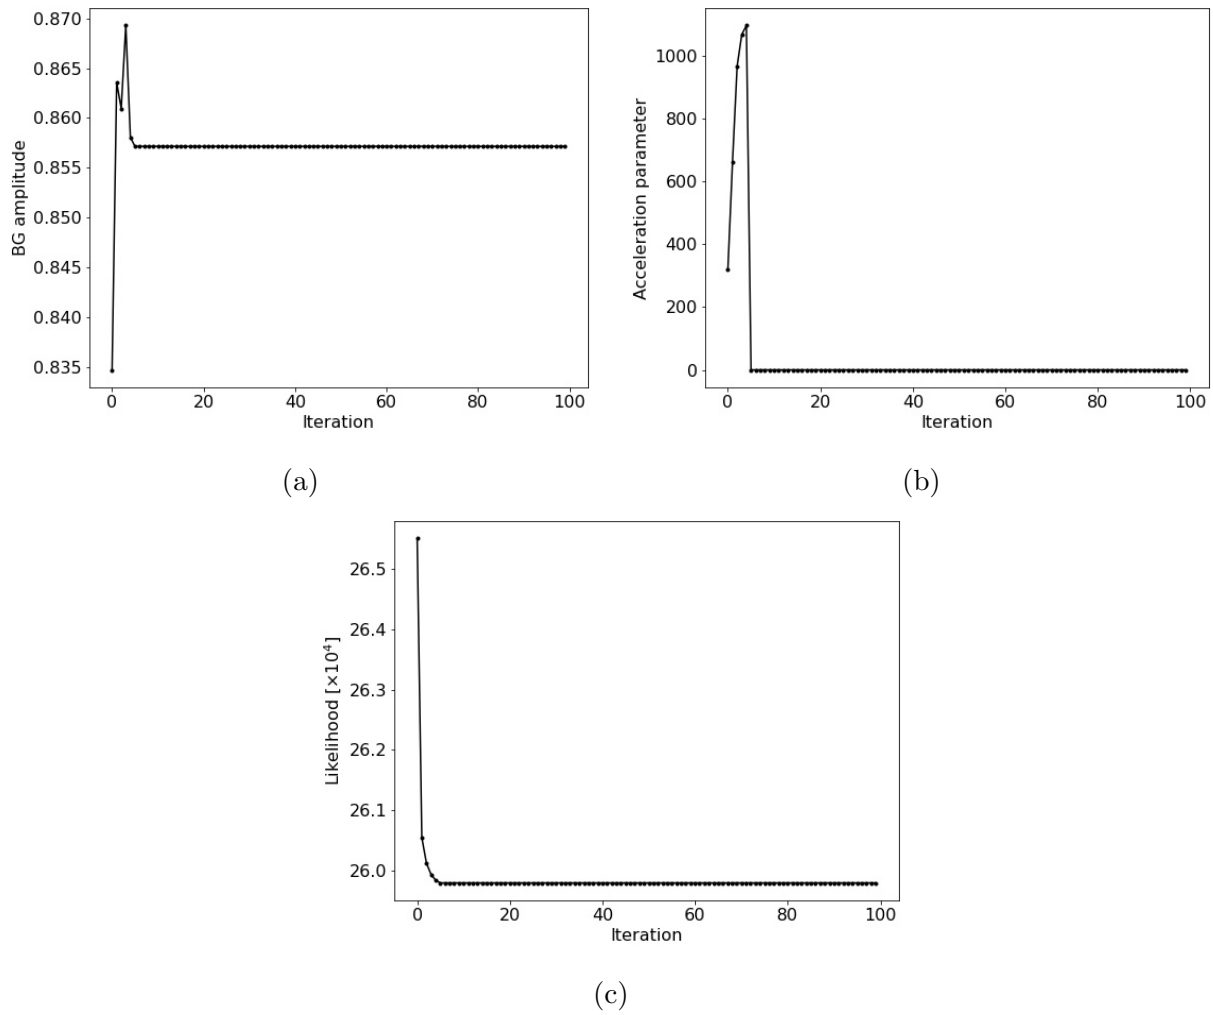


Figure 6.35: Convergence of the background amplitude, acceleration parameter, and likelihood of ^{26}Al emission with atmospheric background in the *cosipy* RL algorithm (Section 6.6.3.2). The negative likelihood is plotted.

6.7 Summary

In this chapter we introduced the fundamentals of imaging the MeV γ -ray sky. The instrument response, generated with computationally expensive simulations, relates locations on the sky (in “image space”) to measurements recorded in the CDS (in “data space”). The response matrix is not invertible and iterative techniques are required to find the most likely distribution in image space that describes the measured data. We adapt the Richardson-Lucy deconvolution algorithm for use in `cosipy`, a Python toolkit for COSI analysis currently being developed in advance of the COSI satellite launch in 2027. Section 6.5 detailed several tests of increasing complexity designed to confirm qualitative functionality of the algorithm. The images generated by `cosipy` are in good qualitative agreement with physical expectations and with the images generated by `mimrec`’s well-tested list-mode method. Having established confidence in its ability to image toy point sources and diffuse sources with changing instrument orientation, the algorithm was packaged into user-friendly Jupyter notebook tutorials for COSI Data Challenge 1. Participants of the data challenge are guided through imaging of the Crab Nebula, Cygnus X-1, Centaurus A, and Vela (Section 6.6.1), the positron-electron annihilation excess at the center of the Milky Way (Section 6.6.2), and Galactic ^{26}Al (Section 6.6.3).

Despite good qualitative agreement between the produced images and expectations, an important caveat is that additional realizations of the simulations may yield slightly different images, given that the number of photons above background would vary between each. Consider, for example, the images from 50 separate realizations of a 511 keV simulation (Figure 6.36). The images have different appearances, as expected from the different statistics in each. Running an ensemble of simulations is therefore necessary to draw quantitative conclusions about simulated source morphology. Overall, though, successfully implementing the RL algorithm in Data Challenge 1, coupled with the published image of 511 keV emission from the 2016 balloon flight (Siegert et al. 2020), together mark important first steps in `cosipy` development.

However, imaging the ^{26}Al emission from the 2016 balloon flight, the measurement of which was described in Chapter 5, is not achievable given the limited statistics (only ~ 100 measured ^{26}Al photons). This is far fewer than the ~ 7000 photons in the $10\times$ flux ^{26}Al simulation without background (Figure 6.32b), which already exhibits difficulties that manifest in visibly disparate pixels at late iterations of the RL algorithm. Following a similar argument in Mahoney et al. (1984) about the inability to generate a map from the HEAO-3 measurements of ^{26}Al , assume that the total significance s of a hypothetical map from the COSI flight is written as:

$$s = \left(\sum_{i=1}^n s_i^2 \right)^{1/2} . \quad (6.8)$$

The measurement in Chapter 5 is $s = 3.7\sigma$ significance. Distributing this significance over individual pixels i of, for example, a minimum of $\sim 3\sigma$ significance for a meaningful image,

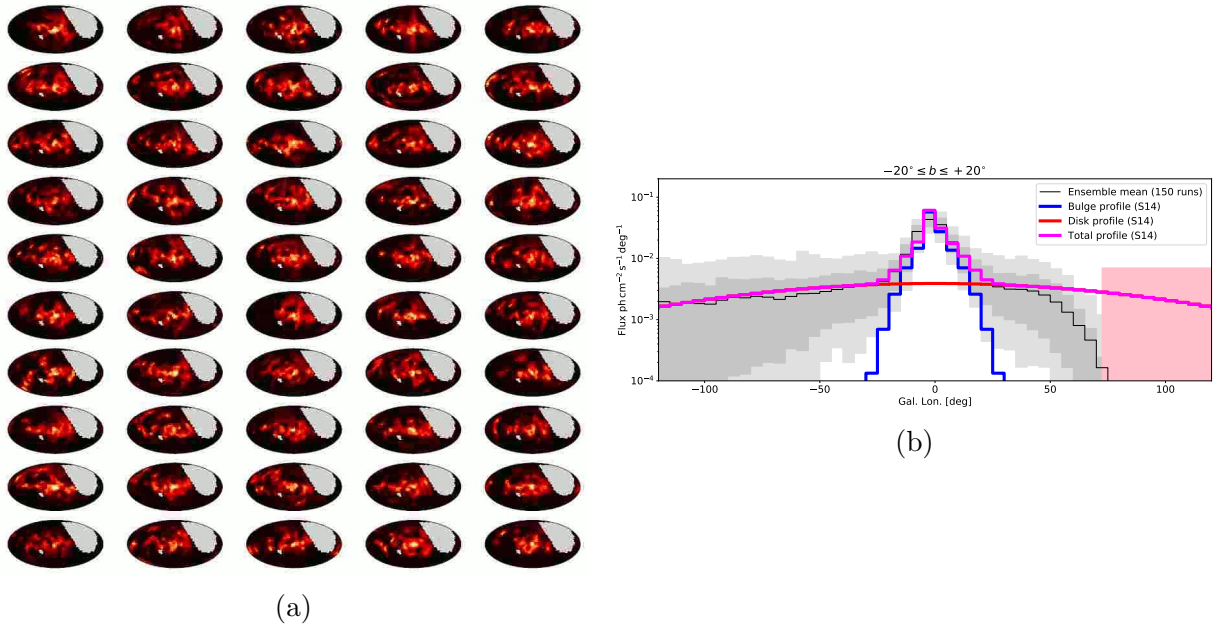


Figure 6.36: An ensemble of simulations of 511 keV emission as observed by COSI-SP, a proposed balloon instrument very similar to COSI with a 120-day flight, illustrates the expected variation in source morphology with statistical differences between realizations. The 511 keV morphology is based on the Skinner et al. (2014) model. Atmospheric background is included. Taking a slice of each image in (a) and calculating the variance in each pixel yields the distribution in (b). Images from Siegert, private communication, 2023.

would yield fewer than $n = 2$ significant pixels. The 2016 data set is therefore insufficient to make an all-sky map of ^{26}Al , and indeed, attempting to feed the real data through the `cosipy` RL algorithm fails. The COSI satellite mission’s two-year observation time, wide field of view, and fine angular resolution (Tomsick et al. 2019) have great potential to image the emission with sophisticated `cosipy` tools and build upon the measurement of ^{26}Al from the 2016 balloon flight.

Chapter 7

Conclusion

Products of stellar nucleosynthesis emit MeV γ -rays upon decay. Observing these γ -rays with high-resolution spectroscopy and imaging reveals the life cycle of elements in the Milky Way Galaxy as they are synthesized, ejected, and distributed from their progenitor sites and into the interstellar medium. The 1.809 MeV γ -ray signature from ^{26}Al traces massive star activity over millions of years (Chapter 2). The majority of its observed abundance is likely produced by Wolf-Rayet stars and core-collapse supernovae with sub-dominant contributions from novae and AGB stars. The relative yields from each source type, however, are not well-constrained. Studies of data from SPI's high resolution germanium detectors paint a kinematic picture of ^{26}Al motion by which it preferentially flows into low-density superbubbles blown out by earlier stellar activity, thereby injecting it at velocities which exceed that of broader Galactic rotation. Historically, measurements of its speed have varied and are complicated by the fundamental limits of instrumental resolution.

Imaging of ^{26}Al by COMPTEL and SPI has revealed concentrated emission in the Inner Galaxy and along the Galactic Plane. Enhanced emission in more localized regions of massive star activity may help identify likely progenitors and the dynamics of ^{26}Al on smaller scales. Emission at high Galactic latitudes remains poorly understood and additional observations are needed to properly model the dynamics of ^{26}Al as it flows out of the Galactic Plane. Thus, answering open questions about the dominant progenitors of ^{26}Al and the kinematics of its motion are important to refining models of stellar nucleosynthesis and the chemical evolution of the Milky Way.

The Compton Spectrometer and Imager (COSI; Chapter 3) is a compact Compton telescope with excellent energy resolution and imaging capabilities that can aid in this pursuit. COSI was launched on a NASA ultra-long duration balloon in 2016 and flew for 46 days. This document detailed the calibration procedures developed before the 2016 launch and calibrations before the intended launch in 2020 (Chapter 4). Calibrations are used to validate instrument performance, inform studies of harmful detector effects like charge sharing and charge trapping (Appendices B and C), and are critical to correctly interpreting flight data. They facilitated the primary analysis of this dissertation, which reports the first measurement of Galactic ^{26}Al in the COSI 2016 balloon flight (Chapter 5). The isotope is measured with

3.7σ significance above background and an Inner Galaxy flux of $(8.6 \pm 2.5) \times 10^{-4} \text{ ph cm}^{-2} \text{ s}^{-1}$ (uncertainty is statistical). The flux exceeds previous measurements by COMPTEL and SPI of $\sim 3.3 \times 10^{-4} \text{ ph cm}^{-2} \text{ s}^{-1}$ but is consistent within 2σ statistical uncertainties. Extensive tests of the analysis method with both flight data and simulations place a $\sim 57\%$ systematic uncertainty on the COSI measurement. The ~ 100 photons in the measurement are insufficient to create an image of the emission. Instead, the imaging capabilities of the COSI balloon instrument are qualitatively assessed with simulations of toy point and diffuse sources, real astrophysical point sources, 511 keV positron-electron annihilation, and ^{26}Al (Chapter 6). These tests were advertised to the astrophysics community in a public “data challenge” using COSI’s new `cosipy` analysis framework.

The calibrations, measurement of ^{26}Al , and imaging tests with `cosipy` will be integral to the success of COSI as a NASA Small Explorer satellite. Launching in 2027, the satellite draws significant heritage from the balloon mission and will utilize the same calibration techniques outlined in this document. Improvements to the calibrations derived from the studies of charge sharing and charge trapping are underway. Among other benefits, incorporating the former can refine the satellite’s detector effects engine. Incorporating the latter’s first demonstration of intrinsic electron and hole trapping in COSI’s detectors can improve the already excellent energy resolution.

With these developments, the COSI satellite is anticipated to reveal element formation in the Galaxy with a strong detection of ^{26}Al and the most detailed image of its emission to date. It is also expected to make the first all-sky image of ^{60}Fe , which when combined with that of ^{26}Al , will help constrain models of stellar and explosive nucleosynthesis. By advancing the proven potency of COSI as a balloon-borne MeV telescope to a satellite mission in low-Earth orbit with improved effective area, angular resolution, and increased observation time, COSI has great potential to shed light on the historically under-explored “MeV gap” of γ -ray astrophysics.

Bibliography

- 3M. *3M Novec 7200 Engineered Fluid*. URL: https://www.3m.com/3M/en_US/p/d/b40045142/.
- Abbott, D. C. (1982). “The theory of radiatively driven stellar winds. II-The line acceleration”. In: *The Astrophysical Journal* 259, pp. 282–301.
- Abbott, D. C. and P. S. Conti (1987). “Wolf-rayet stars”. In: *Annual review of astronomy and astrophysics* 25.1, pp. 113–150.
- Abt, I. et al. (2021). “Simulation of semiconductor detectors in 3D with SolidStateDetectors.jl”. In: *Journal of Instrumentation* 16.08, P08007. DOI: [10.1088/1748-0221/16/08/p08007](https://doi.org/10.1088/1748-0221/16/08/p08007).
- Adkins, J. (2017). *Scientific Balloons*. URL: <https://www.nasa.gov/scientific-balloons/types-of-balloons>.
- Agostinelli, S. et al. (2003). “Geant4—a simulation toolkit”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 506.3, pp. 250–303. ISSN: 0168-9002. DOI: [https://doi.org/10.1016/S0168-9002\(03\)01368-8](https://doi.org/10.1016/S0168-9002(03)01368-8).
- Akaike, H. (1974). “A new look at the statistical model identification”. In: *IEEE Transactions on Automatic Control* 19.6, pp. 716–723. DOI: [10.1109/TAC.1974.1100705](https://doi.org/10.1109/TAC.1974.1100705).
- Allison, J. et al. (2006). “Geant4 developments and applications”. In: *IEEE Transactions on Nuclear Science* 53.1, pp. 270–278. DOI: [10.1109/TNS.2006.869826](https://doi.org/10.1109/TNS.2006.869826).
- Allison, J. et al. (2016). “Recent developments in Geant4”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 835, pp. 186–225. ISSN: 0168-9002. DOI: <https://doi.org/10.1016/j.nima.2016.06.125>.
- Amman, M. (2020). “High Purity Germanium Based Radiation Detectors with Segmented Amorphous Semiconductor Electrical Contacts: Fabrication Procedures”. In: *arXiv preprint arXiv:2006.05471*.
- Auer, M. et al. (2009). “Cosmogenic ^{26}Al in the atmosphere and the prospect of a $^{26}\text{Al}/^{10}\text{Be}$ chronometer to date old ice”. In: *Earth and Planetary Science Letters* 287.3, pp. 453–462. ISSN: 0012-821X. DOI: <https://doi.org/10.1016/j.epsl.2009.08.030>.
- Ayre, C. A. et al. (1984). “Line features in the background γ -ray spectrum in the range 0.05–10 MeV at balloon altitudes”. In: *Nuclear Instruments and Methods in Physics Research* 220.2-3, pp. 549–560.

- Bandstra, M. S. (2010). “Observation of the Crab Nebula in Soft Gamma Rays with the Nuclear Compton Telescope”. PhD thesis, p. 187. ISBN: 978-1-124-55117-3. URL: <https://www.proquest.com/dissertations-theses/observation-crab-nebula-soft-gamma-rays-with/docview/861483974/se-2>.
- Beacom, J. F. and H. Yüksel (2006). “Stringent Constraint on Galactic Positron Production”. In: *Phys. Rev. Lett.* 97 (7), p. 071102. DOI: [10.1103/PhysRevLett.97.071102](https://doi.org/10.1103/PhysRevLett.97.071102).
- Beechert, J. et al. (2022a). “Calibrations of the Compton Spectrometer and Imager”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 1031, p. 166510. ISSN: 0168-9002. DOI: <https://doi.org/10.1016/j.nima.2022.166510>.
- Beechert, J. et al. (2022b). “Measurement of Galactic ^{26}Al with the Compton Spectrometer and Imager”. In: *The Astrophysical Journal* 928.2, p. 119.
- Bennett, M. B. et al. (2013). “Classical-nova contribution to the Milky Way’s ^{26}Al abundance: Exit channel of the key $^{25}\text{Al}(p, \gamma)^{26}\text{Si}$ resonance”. In: *Physical review letters* 111.23, p. 232503.
- Berger, M. J. et al. (2010). *XCOM: Photon Cross Sections Database. NIST Standard Reference Database 8 (XGAM)*. NIST, PML, Radiation Physics Division.
- Berger, M. J. et al. (2017). *Stopping-Power & Range Tables for Electrons, Protons, and Helium Ions*. NIST Standard Reference Database 124. Last update to data content: July 2017. DOI: <https://dx.doi.org/10.18434/T4NC7P>.
- Boggs, S. E. (2023). “Modeling charge cloud dynamics in cross strip semiconductor detectors”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, p. 168310. ISSN: 0168-9002. DOI: <https://doi.org/10.1016/j.nima.2023.168310>.
- Boggs, S. E. and P. Jean (2000). “Event reconstruction in high resolution Compton telescopes”. In: *Astronomy and Astrophysics Supplement Series* 145.2, pp. 311–321.
- Boggs, S. E. et al. (2015). “ ^{44}Ti gamma-ray emission lines from SN1987A reveal an asymmetric explosion”. In: *Science* 348.6235, pp. 670–671.
- Boss, A. P. (2017). “Triggering collapse of the presolar dense cloud core and injecting short-lived radioisotopes with a shock wave. V. Nonisothermal collapse regime”. In: *The Astrophysical Journal* 844.2, p. 113.
- Boss, A. P. and S. A. Keiser (2014). “Triggering collapse of the presolar dense cloud core and injecting short-lived radioisotopes with a shock wave. III. Rotating three-dimensional cloud cores”. In: *The Astrophysical Journal* 788.1, p. 20.
- Bouchet, L., E. Jourdain, and J.-P. Roques (2015). “The Galactic ^{26}Al emission map as revealed by SPI”. In: *The Astrophysical Journal* 801.2, p. 142. DOI: [10.1088/0004-637x/801/2/142](https://doi.org/10.1088/0004-637x/801/2/142).
- Bouchet, L., J.-P. Roques, and E. Jourdain (2010). “On the morphology of the electron-positron annihilation emission as seen by SPI/INTEGRAL”. In: *The Astrophysical Journal* 720.2, p. 1772.
- Bowen, J. D. (2009). “The Prototype Nuclear Compton Telescope: Observations of the Galactic Anticenter Region”. PhD thesis, p. 192. ISBN: 978-1-109-44839-9. URL: <https://www>

- [.proquest.com/dissertations-theses/prototype-nuclear-compton-telescope-observations/docview/304837165/se-2?accountid=14496](https://www.proquest.com/dissertations-theses/prototype-nuclear-compton-telescope-observations/docview/304837165/se-2?accountid=14496).
- Bowen, J. D. et al. (2006). “Background Measurements With the Prototype Nuclear Compton Telescope”. In: *AAS/High Energy Astrophysics Division #9*. Vol. 9. AAS/High Energy Astrophysics Division, p. 18.23.
- Bowen, J. D. et al. (2007). “Depth dependent background measurements with NCT”. In: *2007 IEEE Nuclear Science Symposium Conference Record*. Vol. 1, pp. 436–444. DOI: [10.1109/NSSMIC.2007.4436364](https://doi.org/10.1109/NSSMIC.2007.4436364).
- Cameron, A. G. W. and J. W. Truran (1977). “The supernova trigger for formation of the solar system”. In: *Icarus* 30.3, pp. 447–461.
- Cash, W. (1979). “Parameter estimation in astronomy through application of the likelihood ratio.” In: *Astrophysical Journal* 228, pp. 939–947. DOI: [10.1086/156922](https://doi.org/10.1086/156922).
- Chattopadhyay, T. (2021). “Hard X-ray polarimetry—an overview of the method, science drivers, and recent findings”. In: *Journal of Astrophysics and Astronomy* 42.2, p. 106.
- Chomiuk, L., B. D. Metzger, and K. J. Shen (2021). “New Insights into Classical Novae”. In: *Annual Review of Astronomy and Astrophysics* 59.1, pp. 391–444. DOI: [10.1146/annurev-astro-112420-114502](https://doi.org/10.1146/annurev-astro-112420-114502).
- Churazov, E. et al. (2011). “Positron annihilation spectrum from the Galactic Centre region observed by SPI/INTEGRAL revisited: annihilation in a cooling ISM?” In: *Monthly Notices of the Royal Astronomical Society* 411.3, pp. 1727–1743.
- Coburn, W. and S. E. Boggs (2003). “Polarization of the prompt γ -ray emission from the γ -ray burst of 6 December 2002”. In: *Nature* 423.6938, pp. 415–417.
- Croft, S. and D. S. Bond (1991). “A determination of the Fano factor for germanium at 77.4 K from measurements of the energy resolution of a 113 cm³ HPGe gamma-ray spectrometer taken over the energy range from 14 to 6129 keV”. In: *International journal of radiation applications and instrumentation. Part A. Applied radiation and isotopes* 42.11, pp. 1009–1014.
- Cumani, P. et al. (2019). “Background for a gamma-ray satellite on a low-Earth orbit”. In: *Experimental Astronomy* 47.3, pp. 273–302.
- Dale, J. E. (2015). “The modelling of feedback in star formation simulations”. In: *New Astronomy Reviews* 68, pp. 1–33.
- Dawson, J. R. et al. (2013). “Supergiant shells and molecular cloud formation in the large magellanic cloud”. In: *The Astrophysical Journal* 763.1, p. 56.
- Diehl, R. et al. (2006). “Radioactive ²⁶Al from massive stars in the Galaxy”. In: *Nature* 439.7072, pp. 45–47.
- Du Mond, J. W. M. (1929). “Compton modified line structure and its relation to the electron theory of solid bodies”. In: *Physical Review* 33.5, p. 643.
- Duncan, N. (2017). “Gamma-Ray Observations of Solar Flares with RHESSI Imaging Spectroscopy and the GRIPS instrument”. PhD thesis, p. 213. ISBN: 978-0-355-57253-7. URL: <https://www.proquest.com/dissertations-theses/gamma-ray-observations-solar-flares-with-rhessi/docview/2001493923/se-2>.

- Duncan, N. et al. (2016). “First flight of the Gamma-Ray Imager/Polarimeter for Solar flares (GRIPS) instrument”. In: *Space Telescopes and Instrumentation 2016: Ultraviolet to Gamma Ray*. Vol. 9905, 99052Q. DOI: [10.1117/12.2233859](https://doi.org/10.1117/12.2233859).
- Ellis, J., B. D. Fields, and D. N. Schramm (1996). “Geological Isotope Anomalies as Signatures of Nearby Supernovae”. In: *Astrophysical Journal* 470, p. 1227. DOI: [10.1086/177945](https://doi.org/10.1086/177945).
- European Space Agency. *INTEGRAL pictures and videos*. URL: http://integral.esac.esa.int/integ_pictures.html.
- Evans, R. D. (1955). *The atomic nucleus*. Vol. 582. McGraw-Hill New York.
- Feige, J. et al. (2018). “Limits on supernova-associated $^{60}\text{Fe}/^{26}\text{Al}$ nucleosynthesis ratios from accelerator mass spectrometry measurements of deep-sea sediments”. In: *Phys. Rev. Lett.* 121, p. 221103.
- Ferrière, K. (1998). “The Hot Gas Filling Factor in Our Galaxy”. In: *The Astrophysical Journal* 503.2, p. 700. DOI: [10.1086/306003](https://doi.org/10.1086/306003).
- Fluid-o-Tech. *Fluid-o-Tech FG200-300-400*. URL: <https://www.fluidotech.it/en/products/technologies/external-gear-pumps/fg200-300-series/>.
- Forbes, J. C., J. Alves, and D. N. C. Lin (2021). “A Solar System formation analogue in the Ophiuchus star-forming complex”. In: *Nature Astronomy*, pp. 1–8.
- Foreman-Mackey, D. et al. (2013). “emcee: The MCMC Hammer”. In: *PASP* 125, pp. 306–312. DOI: [10.1086/670067](https://doi.org/10.1086/670067).
- Freck, D. V. and J. Wakefield (1962). “Gamma-Ray Spectrum obtained with a Lithium-drifted p–i–n Junction in Germanium”. In: *Nature* 193.4816, pp. 669–669.
- Fujimoto, Y., M. R. Krumholz, and S-i. Inutsuka (2020). “Distribution and kinematics of ^{26}Al in the Galactic disc”. In: *Monthly Notices of the Royal Astronomical Society* 497.2, pp. 2442–2454.
- Fujimoto, Y., M. R. Krumholz, and S. Tachibana (2018). “Short-lived radioisotopes in meteorites from Galactic-scale correlated star formation”. In: *Monthly Notices of the Royal Astronomical Society* 480.3, pp. 4025–4039.
- Gaensler, B. M. and P. O. Slane (2006). “The Evolution and Structure of Pulsar Wind Nebulae”. In: *Annual Review of Astronomy and Astrophysics* 44.1, pp. 17–47. DOI: [10.1146/annurev.astro.44.051905.092528](https://doi.org/10.1146/annurev.astro.44.051905.092528).
- Gallagher, J. S. and S. Starrfield (1978). “Theory and observations of classical novae”. In: *Annual Review of Astronomy and Astrophysics* 16.1, pp. 171–214.
- Grefenstette, B. W. et al. (2014). “Asymmetries in core-collapse supernovae from maps of radioactive ^{44}Ti in Cassiopeia A”. In: *Nature* 506.7488, pp. 339–342.
- Haefner, A. et al. (2017). “Handheld real-time volumetric 3-D gamma-ray imaging”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 857, pp. 42–49.
- Hall, R. N. and T. J. Soltys (1971). “High purity germanium for detector fabrication”. In: *IEEE Transactions on Nuclear Science* 18.1, pp. 160–165.
- Halloin, H. et al. (2003). “Performance of CLAIRE, the first balloon-borne gamma-ray lens telescope”. In: *Nuclear Instruments and Methods in Physics Research Section A: Acceler-*

- ators, *Spectrometers, Detectors and Associated Equipment* 504.1. Proceedings of the 3rd International Conference on New Developments in Photodetection, pp. 120–125. ISSN: 0168-9002. DOI: [https://doi.org/10.1016/S0168-9002\(03\)00807-6](https://doi.org/10.1016/S0168-9002(03)00807-6).
- Hansen, W. L. (1971). “High-purity germanium crystal growing”. In: *Nuclear Instruments and Methods* 94.2, pp. 377–380.
- Hauser, M. et al. (1998). “COBE Ref Pub No. 98-A Greenbelt”. In: *MD: NASA/GSPC*. Available in electronic form from the NSSDC.
- Henck, R., P. Siffert, and A. Coche (1968). *Characteristics of a 85 cm³ Ge(Li) gamma-ray detector*. Tech. rep. Centre de Recherches Nucleaires, Strasbourg.
- Höfner, S. and H. Olofsson (2018). “Mass loss of stars on the asymptotic giant branch: Mechanisms, models and measurements”. In: *The Astronomy and Astrophysics Review* 26, pp. 1–92.
- Howarth, I. D. and R. K. Prinja (1989). “The stellar winds of 203 Galactic O stars: A quantitative ultraviolet survey”. In: *Astrophysical Journal Supplement Series (ISSN 0067-0049)* 69, pp. 527–592.
- Hoyle, F. and W. A. Fowler (1960). “Nucleosynthesis in Supernovae”. In: *The Astrophysical Journal* 132, p. 565.
- Huang, W. and D. R. Gies (2006). “Stellar rotation in young clusters. I. Evolution of projected rotational velocity distributions”. In: *The Astrophysical Journal* 648.1, p. 580.
- Hull, E. L. et al. (2014). “Charge-trap correction and radiation damage in orthogonal-strip planar germanium detectors”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 762, pp. 125–129. ISSN: 0168-9002. DOI: <https://doi.org/10.1016/j.nima.2014.05.086>.
- Iyudin, A. F. et al. (1994). “COMPTEL observations of ⁴⁴Ti gamma-ray line emission from Cas A.” In: *Astronomy and Astrophysics* 284, pp. L1–L4.
- Iyudin, A. F. et al. (1997). “Cas A in the Light of the ⁴⁴Ti 1.15 MeV Gamma-Ray Line Emission”. In: *The Transparent Universe*. Vol. 382, p. 37.
- Janka, H.-T. (2012). “Explosion mechanisms of core-collapse supernovae”. In: *Annual Review of Nuclear and Particle Science* 62, pp. 407–451.
- Jean, P. et al. (2006). “Spectral analysis of the Galactic e⁺e⁻ annihilation emission”. In: *Astronomy & Astrophysics* 445.2, pp. 579–589.
- Johnson, M. (2020). *Columbia Scientific Balloon Facility*. URL: <https://www.csbf.nasa.gov/balloons.html>.
- Johnson III, W. N., F. R. Harnden Jr, and R. C. Haymes (1972). “The spectrum of low-energy gamma radiation from the galactic-center region.” In: *The Astrophysical Journal* 172, p. L1.
- José, J., M. Hernanz, and A. Coc (1997). “New results on ²⁶Al production in classical novae”. In: *The Astrophysical Journal* 479.1, p. L55.
- Kaur, T. and S. Sahijpal (2019). “Heterogeneous evolution of the Galaxy and the origin of the short-lived nuclides in the early solar system”. In: *Monthly Notices of the Royal Astronomical Society* 490.2, pp. 1620–1637. ISSN: 0035-8711. DOI: [10.1093/mnras/stz2720](https://doi.org/10.1093/mnras/stz2720).

- Kierans, C. (2018). “Detection of the 511 keV Positron Annihilation Line with the Compton Spectrometer and Imager”. PhD thesis, p. 273. ISBN: 978-0-438-64291-1. URL: <https://www.proquest.com/dissertations-theses/detection-511-kev-positron-annihilation-line-with/docview/2135808024/se-2>.
- Kierans, C., T. Takahashi, and G. Kanbach (2022). “Compton Telescopes for Gamma-ray Astrophysics”. In: *arXiv e-prints*. DOI: [10.48550/arXiv.2208.07819](https://doi.org/10.48550/arXiv.2208.07819).
- Kierans, C. et al. (2016). “The 2016 Super Pressure Balloon flight of the Compton Spectrometer and Imager”. In: *11th INTEGRAL Conference Gamma-Ray Astrophysics in Multi-Wavelength Perspective*, p. 75.
- Kierans, C. et al. (2020). “Detection of the 511 keV Galactic positron annihilation line with COSI”. In: *The Astrophysical Journal* 895.1, p. 44.
- Klein, O. and Y. Nishina (1929). “Über die Streuung von Strahlung durch freie Elektronen nach der neuen relativistischen Quantendynamik von Dirac”. In: *Zeitschrift für Physik* 52.11-12, pp. 853–868.
- Knie, K. et al. (1999). “Indication for Supernova Produced ^{60}Fe Activity on Earth”. In: *Phys. Rev. Lett.* 83 (1), pp. 18–21. DOI: [10.1103/PhysRevLett.83.18](https://doi.org/10.1103/PhysRevLett.83.18).
- Knödseder, J. (1999). “Implications of 1.8 MeV gamma-ray observations for the origin of ^{26}Al ”. In: *The Astrophysical Journal* 510.2, p. 915.
- Knödseder, J. et al. (1996). “Modelling the 1.8 MeV sky: Tests for spiral structure.” In: *Astronomy and Astrophysics Supplement* 120, pp. 335–338. arXiv: [astro-ph/9604053](https://arxiv.org/abs/astro-ph/9604053) [[astro-ph](https://arxiv.org/abs/astro-ph)].
- Knödseder, J. et al. (1999). “A multiwavelength comparison of COMPTEL 1.8 MeV ^{26}Al line data”. In: *Astronomy and Astrophysics* 344, pp. 68–82.
- Knödseder, J. et al. (1999). “Image reconstruction of COMPTEL 1.8 MeV ^{26}Al line data”. In: *Astronomy and Astrophysics* 345, pp. 813–825. arXiv: [astro-ph/9903172](https://arxiv.org/abs/astro-ph/9903172) [[astro-ph](https://arxiv.org/abs/astro-ph)].
- Knödseder, J. et al. (2005). “The all-sky distribution of 511 keV electron-positron annihilation emission”. In: *Astronomy & Astrophysics* 441.2, pp. 513–532.
- Knoll, G. F. (2010). *Radiation detection and measurement*. John Wiley & Sons.
- Kole, M. et al. (2020). “The POLAR gamma-ray burst polarization catalog”. In: *Astronomy & Astrophysics* 644, A124.
- Korschinek, G. et al. (2020). “Supernova-produced ^{53}Mn on Earth”. In: *Physical Review Letters* 125.3, p. 031101.
- Krause, M. G. H. et al. (2014). “Feedback by massive stars and the emergence of superbubbles-II. X-ray properties”. In: *Astronomy & Astrophysics* 566, A94.
- Krause, M. G. H. et al. (2015). “ ^{26}Al kinematics: superbubbles following the spiral arms?-Constraints from the statistics of star clusters and HI supershells”. In: *Astronomy & Astrophysics* 578, A113.
- Krause, M. G. H. et al. (2018). “Surround and Squash: the impact of superbubbles on the interstellar medium in Scorpius–Centaurus OB2”. In: *Astronomy & Astrophysics* 619, A120.

- Krause, M. G. H. et al. (2021). “Galactic ^{26}Al traces metal loss through hot chimneys”. In: *Monthly Notices of the Royal Astronomical Society* 501.1, pp. 210–218. DOI: [10.1093/mnras/staa3612](https://doi.org/10.1093/mnras/staa3612).
- Kretschmer, K. et al. (2013). “Kinematics of massive star ejecta in the Milky Way as traced by ^{26}Al ”. In: *Astronomy & Astrophysics* 559, A99. ISSN: 1432-0746. DOI: [10.1051/0004-6361/201322563](https://doi.org/10.1051/0004-6361/201322563).
- Lawson, T. V. et al. (2022). “Radioactive nuclei in the early Solar System: Analysis of the 15 isotopes produced by core-collapse supernovae”. In: *Monthly Notices of the Royal Astronomical Society* 511.1, pp. 886–902.
- Lazzati, D. et al. (2004). “Compton drag as a mechanism for very high linear polarization in gamma-ray bursts”. In: *Monthly Notices of the Royal Astronomical Society* 347.1, pp. L1–L5.
- Leventhal, M., C. J. MacCallum, and P. D. Stang (1978). “Detection of 511 keV positron annihilation radiation from the galactic center direction”. In: *The Astrophysical Journal* 225, pp. L11–L14.
- Li, T.-P. and Y.-Q. Ma (1983). “Analysis methods for results in gamma-ray astronomy”. In: *The Astrophysical Journal* 272, pp. 317–324.
- Ling, J. C. (1975). “A semiempirical model for atmospheric γ rays from 0.3 to 10 MeV at $\lambda = 40^\circ$ ”. In: *Journal of Geophysical Research* 80.22, pp. 3241–3252.
- Lingenfelter, R., K.-W. Chan, and R. Ramaty (1993). “Diffuse galactic annihilation radiation and the present rate of galactic nucleosynthesis”. In: *Physics Reports* 227.1, pp. 133–142. ISSN: 0370-1573. DOI: [https://doi.org/10.1016/0370-1573\(93\)90061-H](https://doi.org/10.1016/0370-1573(93)90061-H).
- Llacer, J. (1972). “Planar and coaxial high purity germanium radiation detectors”. In: *Nuclear Instruments and Methods* 98.2, pp. 259–268.
- Looker, Q. (2014). “Fabrication process development for high-purity germanium radiation detectors with amorphous semiconductor contacts”. PhD thesis, p. 177. ISBN: 978-1-321-26013-7. URL: <https://www.proquest.com/dissertations-theses/fabrication-process-development-high-purity/docview/1627787065/se-2>.
- Lowell, A. (2017). “Polarimetric Studies of the Long Duration Gamma-Ray Burst GRB 160530A with the Compton Spectrometer and Imager”. PhD thesis, p. 181. ISBN: 978-0-355-83142-9. URL: <https://www.proquest.com/dissertations-theses/polarimetric-studies-long-duration-gamma-ray/docview/2031543342/se-2>.
- Lowell, A. et al. (2016). “Positional calibrations of the germanium double-sided strip detectors for the Compton Spectrometer and Imager”. In: *High Energy, Optical, and Infrared Detectors for Astronomy VII*. Ed. by Andrew D. Holland and James Beletic. Vol. 9915. International Society for Optics and Photonics. SPIE, pp. 797–806. DOI: [10.1117/12.2233145](https://doi.org/10.1117/12.2233145).
- Lowell, A. et al. (2017). “Maximum likelihood Compton polarimetry with the Compton Spectrometer and Imager”. In: *The Astrophysical Journal* 848.2, p. 120.
- Lucy, L. B. (1974). “An iterative technique for the rectification of observed distributions”. In: *The Astronomical Journal* 79, p. 745.

- Lund, N. (2021). “Wide band, tunable gamma-ray lenses”. In: *Experimental Astronomy* 51.1, pp. 153–163.
- Mahoney, W. A. et al. (1984). “HEAO 3 discovery of ^{26}Al in the interstellar medium”. In: *The Astrophysical Journal* 286, pp. 578–585.
- Malet, I. et al. (1991). “Observation of the galactic 1809 keV gamma-ray line with the HEXAGONE spectrometer”. In: *AIP Conference Proceedings* 232.1, pp. 123–128. DOI: [10.1063/1.40952](https://doi.org/10.1063/1.40952).
- Malm, H. L. (1967). “Encapsulated Coaxial Ge(Li) Detectors Operated in Parallel”. In: *IEEE Transactions on Nuclear Science* 14.1, pp. 521–522.
- Martinet, S. et al. (2022). “Very massive star winds as sources of the short-lived radioactive isotope ^{26}Al ”. In: *Astronomy and Astrophysics* 664, A181, A181. DOI: [10.1051/0004-6361/202243474](https://doi.org/10.1051/0004-6361/202243474).
- McMaster. *McMaster Barbed Hose Fitting for Air and Water*. URL: <https://www.mcmaster.com/91355K47/>.
- *McMaster High-Pressure Tygon Rubber Tubing*. URL: <https://www.mcmaster.com/1489T11/>.
- Medvedev, M. V. and A. Loeb (1999). “Generation of magnetic fields in the relativistic shock of gamma-ray burst sources”. In: *The Astrophysical Journal* 526.2, p. 697.
- Mitra, P. et al. (2016). “Application of spectrum shifting methodology to restore NaI(Tl)-recorded gamma spectra, shifted due to temperature variations in the environment”. In: *Applied Radiation and Isotopes* 107, pp. 133–137. DOI: [10.1016/j.apradiso.2015.10.002](https://doi.org/10.1016/j.apradiso.2015.10.002).
- Mowlavi, N. and G. Meynet (2000). “Aluminum-26 production in asymptotic giant branch stars”. In: *Astronomy and Astrophysics* 361, pp. 959–976.
- Naya, J. E. et al. (1996). “Detection of high-velocity ^{26}Al towards the Galactic Centre”. In: *Nature* 384.6604, pp. 44–46.
- (1997). “GRIS Detection of ^{26}Al 1809 keV Line Emission from the Galactic Center Region as a Broad Line”. In: *The Transparent Universe, Proc. 2nd INTEGRAL Workshop*. Ed. by T. J.-L. Courvoisier C. Winkler and Ph. Durouchoux (The Netherlands: European Space Agency). Vol. 382, p. 59.
- Oberlack, U. (1997). “Über die Natur der galaktischen ^{26}Al -Quellen-Untersuchung des 1.8-MeV-Himmels mit COMPTEL”. PhD thesis. Dissertation, Technische Universität München, Germany.
- Oberlack, U. et al. (1994). “Implications of ^{26}Al emission at 1.8 MeV from the VELA region”. In: *The Astrophysical Journal Supplement Series* 92, pp. 433–439.
- Oberlack, U. et al. (1996). “The COMPTEL 1.809 MeV all-sky image.” In: *Astronomy and Astrophysics Supplement Series* 120, pp. 311–314.
- Ore, A. and J. L. Powell (1949). “Three-photon annihilation of an electron-positron pair”. In: *Physical Review* 75.11, p. 1696.
- Pavlov, G. G. et al. (2001). “The X-Ray Spectrum of the Vela Pulsar Resolved with the Chandra X-Ray Observatory”. In: *The Astrophysical Journal* 552.2, p. L129. DOI: [10.1086/320342](https://doi.org/10.1086/320342).

- Peterson, L. E., R. M. Pelling, and J. L. Matteson (1972). “Techniques in balloon X-ray astronomy”. In: *Space Science Reviews* 13.2, pp. 320–336.
- Peterson, L. E., D. A. Schwartz, and J. C. Ling (1973). “Spectrum of atmospheric gamma rays to 10 MeV at $\lambda = 40^\circ$ ”. In: *Journal of Geophysical Research* 78.34, pp. 7942–7958.
- Picone, J. M. et al. (2002). “NRLMSISE-00 empirical model of the atmosphere: Statistical comparisons and scientific issues”. In: *Journal of Geophysical Research: Space Physics* 107.A12, SIA–15. DOI: [10.1029/2002JA009430](https://doi.org/10.1029/2002JA009430).
- Pleintinger, M. M. M. (2020). “Star Groups and their Nucleosynthesis”. PhD thesis. Max Planck Institute for Extraterrestrial Physics.
- Pleintinger, M. M. M. et al. (2019). “Comparing simulated ^{26}Al maps to gamma-ray measurements”. In: *Astronomy & Astrophysics* 632, A73.
- Plüschke, S. et al. (2001). “The COMPTEL 1.809 MeV survey”. In: *Exploring the Gamma-Ray Universe*. Ed. by A. Gimenez, V. Reglero, and C. Winkler. Vol. 459. ESA Special Publication, pp. 55–58. arXiv: [astro-ph/0104047](https://arxiv.org/abs/astro-ph/0104047) [[astro-ph](https://arxiv.org/abs/astro-ph)].
- Porter, Troy A. et al. (2008). “Inverse Compton Origin of the Hard X-Ray and Soft Gamma-Ray Emission from the Galactic Ridge”. In: *The Astrophysical Journal* 682.1, p. 400. DOI: [10.1086/589615](https://doi.org/10.1086/589615). URL: <https://dx.doi.org/10.1086/589615>.
- Prantzos, N. (2006). “On the intensity and spatial morphology of the 511 keV emission in the Milky Way”. In: *Astronomy & Astrophysics* 449.3, pp. 869–878.
- Prantzos, N. and R. Diehl (1996). “Radioactive ^{26}Al in the galaxy: observations versus theory”. In: *Physics Reports* 267.1, pp. 1–69. ISSN: 0370-1573. DOI: [https://doi.org/10.1016/0370-1573\(95\)00055-0](https://doi.org/10.1016/0370-1573(95)00055-0). URL: <https://www.sciencedirect.com/science/article/pii/0370157395000550>.
- Ramo, S. (1939). “Currents induced by electron motion”. In: *Proceedings of the IRE* 27.9, pp. 584–585.
- Rastaetter, L. *NRLMSISE-00 Atmosphere Model*. URL: <https://ccmc.gsfc.nasa.gov/modelweb/models/nrlmsise00.php>.
- Richardson, W. H. (1972). “Bayesian-based iterative method of image restoration”. In: *JoSA* 62.1, pp. 55–59.
- Rodgers-Lee, D. et al. (2019). “Synthetic ^{26}Al emission from galactic-scale superbubble simulations”. In: *Monthly Notices of the Royal Astronomical Society* 490.2, pp. 1894–1912.
- Rutledge, R. E. and D. B. Fox (2004). “Re-analysis of polarization in the γ -ray flux of GRB 021206”. In: *Monthly Notices of the Royal Astronomical Society* 350.4, pp. 1288–1300.
- Schaller, G. et al. (1992). “New grids of stellar models from 0.8 to 120 solar masses at $Z=0.020$ and $Z=0.001$ ”. In: *Astronomy and Astrophysics Supplement Series* 96, pp. 269–331.
- Schönfelder, V. et al. (1993). “Instrument Description and Performance of the Imaging Gamma-Ray Telescope COMPTEL aboard the Compton Gamma-Ray Observatory”. In: *Astrophysical Journal, Supplement* 86, p. 657. DOI: [10.1086/191794](https://doi.org/10.1086/191794).
- Section, IAEA Nuclear Data (2010). *Live Chart of Nuclides: nuclear structure and decay data*. URL: <https://www-nds.iaea.org/relnsd/vcharthtml/VChartHTML.html>.
- Shockley, W. (1938). “Currents to conductors induced by a moving point charge”. In: *Journal of Applied Physics* 9.10, pp. 635–636.

- Shukolyukov, A. and G. W. Lugmair (1992). “First evidence for live ^{60}Fe in the early solar system”. In: *Lunar and Planetary Science Conference*. Vol. 23.
- (1993). “Live Iron-60 in the early solar system”. In: *Science* 259.5098, pp. 1138–1142.
- Siegert, T. (2017). “Positron-annihilation spectroscopy throughout the Milky Way”. In: *Ph. D. Thesis*.
- Siegert, T. and R. Diehl (2015). “Gamma-ray lines from SN2014J”. In: *arXiv e-prints*. arXiv: [1501.05648](https://arxiv.org/abs/1501.05648) [astro-ph.HE].
- Siegert, T., D. Horan, and G. Kanbach (2022). “Telescope Concepts in Gamma-Ray Astronomy”. In: *arXiv preprint arXiv:2207.02248*.
- Siegert, T. et al. (2016a). “Gamma-ray spectroscopy of positron annihilation in the Milky Way”. In: *Astronomy & Astrophysics* 586, A84.
- (2016b). “Gamma-ray spectroscopy of positron annihilation in the Milky Way”. In: *Astronomy & Astrophysics* 586, A84.
- Siegert, T. et al. (2019). “Constraints on positron annihilation kinematics in the inner Galaxy”. In: *Astronomy & Astrophysics* 627, A126.
- Siegert, T. et al. (2020). “Imaging the 511 keV positron annihilation sky with COSI”. In: *The Astrophysical Journal* 897.1, p. 45.
- Siegert, T. et al. (2023). “Galactic Population Synthesis of Radioactive Nucleosynthesis Ejecta”. In: *arXiv preprint arXiv:2301.10192*.
- Skinner, G. et al. (2014). “The Galactic distribution of the 511 keV e^+/e^- annihilation radiation”. In: *10th INTEGRAL Workshop: A Synergistic View of the High-Energy Sky-Integral 2014*. SISSA-International School for Advanced Studies, pp. 1–10.
- Sleator, C. (2019). “Measuring the Polarization of Compact Objects with the Compton Spectrometer and Imager”. PhD thesis, p. 237. URL: <https://www.proquest.com/dissertations-theses/measuring-polarization-compact-objects-with/docview/2313733159/se-2?accountid=14496>.
- Sleator, C. et al. (2019). “Benchmarking simulations of the Compton Spectrometer and Imager with calibrations”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 946, p. 162643. ISSN: 0168-9002. DOI: <https://doi.org/10.1016/j.nima.2019.162643>.
- Smart, D. F. and M. A. Shea (2005). “A review of geomagnetic cutoff rigidities for earth-orbiting spacecraft”. In: *Advances in Space Research* 36.10, pp. 2012–2020.
- Snowden, S. L. et al. (1997). “ROSAT survey diffuse X-ray background maps. II.” In: *The Astrophysical Journal* 485.1, p. 125.
- Sunpower. *CryoTel Cryocoolers*. URL: <https://www.sunpowerinc.com/products/stirling-cryocoolers/cryotel-cryocoolers/ct>.
- Syntfeld-Każuch, A. et al. (2010). “Energy resolution of CsI (Na) scintillators”. In: *Radiation Measurements* 45.3-6, pp. 377–379.
- Takada, A. et al. (2011). “Observation of diffuse cosmic and atmospheric gamma rays at balloon altitudes with an electron-tracking Compton camera”. In: *The Astrophysical Journal* 733.1, p. 13.

- Tavendale, A. J. and G. T. Ewan (1963). “A High-Resolution Lithium-Drift Germanium Gamma-Ray Spectrometer”. In: *Nucl. Instr. Methods* 25.
- The GRAVITY Collaboration et al. (2019). “A geometric distance measurement to the Galactic center black hole with 0.3% uncertainty”. In: *Astronomy and Astrophysics* 625, p. L10. DOI: [10.1051/0004-6361/201935656](https://doi.org/10.1051/0004-6361/201935656).
- Timmer, F. X. et al. (1995). “ ^{26}Al and ^{60}Fe from Supernova Explosions”. In: *arXiv preprint astro-ph/9503120*.
- Toma, K. et al. (2009). “Statistical properties of gamma-ray burst polarization”. In: *The Astrophysical Journal* 698.2, p. 1042.
- Tomsick, J. A. et al. (2019). “The Compton Spectrometer and Imager”. In: *Bulletin of the American Astronomical Society*. Vol. 51, p. 98. arXiv: [1908.04334](https://arxiv.org/abs/1908.04334) [[astro-ph.IM](#)].
- Tomsick, J. A. et al. (2022). “Soft gamma-ray polarimetry with COSI using maximum likelihood analysis”. In: *arXiv preprint arXiv:2204.00027*.
- Troja, E. et al. (2014). “Swift/BAT Detection of Hard X-Rays from Tycho’s Supernova Remnant: Evidence for Titanium-44”. In: *The Astrophysical Journal Letters* 797.1, p. L6.
- Vedrenne, G. et al. (2003). “SPI: The spectrometer aboard INTEGRAL”. In: *Astronomy & Astrophysics* 411.1, pp. L63–L70.
- Vianello, G. et al. (2015). “The multi-mission maximum likelihood framework (3ML)”. In: *arXiv preprint arXiv:1507.08343*.
- Vink, J. S. and A. de Koter (2005). “On the metallicity dependence of Wolf-Rayet winds”. In: *Astronomy & Astrophysics* 442.2, pp. 587–596.
- Voss, R. et al. (2009). “Using population synthesis of massive stars to study the interstellar medium near OB associations”. In: *Astronomy & Astrophysics* 504.2, pp. 531–542.
- Wallner, C. et al. (2004). “Supernova produced and anthropogenic ^{244}Pu in deep sea manganese encrustations”. In: *New Astronomy Reviews* 48.1-4, pp. 145–150.
- Wang, W. et al. (2009). “Spectral and intensity variations of Galactic ^{26}Al emission”. In: *Astronomy and Astrophysics* 496.3, pp. 713–724. DOI: [10.1051/0004-6361/200811175](https://doi.org/10.1051/0004-6361/200811175). arXiv: [0902.0211](https://arxiv.org/abs/0902.0211) [[astro-ph.HE](#)].
- Wang, W. et al. (2020). “Gamma-Ray Emission of ^{60}Fe and ^{26}Al Radioactivity in Our Galaxy”. In: *Astrophysical Journal* 889.2, 169, p. 169. DOI: [10.3847/1538-4357/ab6336](https://doi.org/10.3847/1538-4357/ab6336). arXiv: [1912.07874](https://arxiv.org/abs/1912.07874) [[astro-ph.HE](#)].
- Weidenspointner, G. et al. (2003). “First identification and modelling of SPI background lines”. In: *Astronomy and Astrophysics* 411, pp. L113–L116. DOI: [10.1051/0004-6361:20031209](https://doi.org/10.1051/0004-6361:20031209).
- Weidenspointner, G. et al. (2008). “Positron astronomy with SPI/INTEGRAL”. In: *New Astronomy Reviews* 52.7-10, pp. 454–456.
- Wigger, C. et al. (2004). “Gamma-ray burst polarization: limits from RHESSI measurements”. In: *The Astrophysical Journal* 613.2, p. 1088.
- Wilks, S. S. (1938). “The large-sample distribution of the likelihood ratio for testing composite hypotheses”. In: *The Annals of Mathematical Statistics* 9.1, pp. 60–62.
- Winkler, C. et al. (2003). “The INTEGRAL mission”. In: *Astronomy & Astrophysics* 411.1, pp. L1–L6.


- Yonetoku, D. et al. (2012). “Magnetic structures in gamma-ray burst jets probed by gamma-ray polarization”. In: *The Astrophysical Journal Letters* 758.1, p. L1.
- Zinner, E. and C. Gopel (1992). “Evidence for ^{26}Al in feldspars from the H4 chondrite Ste. Marguerite”. In: *Meteoritics* 27.
- Zinner, E. et al. (1991). “Large amounts of extinct ^{26}Al in interstellar grains from the Murchison meteorite”. In: *Nature* 349.6304, pp. 51–54.
- Zoglauer, A. (2005). “First Light for the next Generation of Compton and Pair telescopes”. PhD thesis. Technische Universität München. URL: <https://mediatum.ub.tum.de/node?id=603105>.
- Zoglauer, A., R. Andritschke, and F. Schopper (2006). “MEGALib The Medium Energy Gamma-ray Astronomy Library”. In: *New Astronomy Reviews* 50.7-8, pp. 629–632. DOI: [10.1016/j.newar.2006.06.049](https://doi.org/10.1016/j.newar.2006.06.049).
- Zoglauer, A. and S. E. Boggs (2007). “Application of neural networks to the identification of the Compton interaction sequence in Compton imagers”. In: *2007 IEEE Nuclear Science Symposium Conference Record*. Vol. 6. IEEE, pp. 4436–4441.
- Zoglauer, A. and G. Kanbach (2003). “Doppler broadening as a lower limit to the angular resolution of next-generation Compton telescopes”. In: *X-Ray and Gamma-Ray Telescopes and Instruments for Astronomy*. Ed. by Joachim E. Truemper and Harvey D. Tananbaum. Vol. 4851. International Society for Optics and Photonics. SPIE, pp. 1302–1309. DOI: [10.1117/12.461177](https://doi.org/10.1117/12.461177).
- Zoglauer, A. et al. (2008). “Status of instrumental background simulations for gamma-ray telescopes with Geant4”. In: *2008 IEEE Nuclear Science Symposium Conference Record*, pp. 2859–2864. DOI: [10.1109/NSSMIC.2008.4774966](https://doi.org/10.1109/NSSMIC.2008.4774966).
- Zoglauer, A. et al. (2021). “COSI: From Calibrations and Observations to All-sky Images”. In: *arXiv e-prints*. arXiv: [2102.13158](https://arxiv.org/abs/2102.13158) [astro-ph.IM].

Appendix A


Temperature correction

The dependence of COSI's recorded energies on ambient temperature was examined using calibration data collected in Wanaka, New Zealand for the attempted COSI 2020 balloon flight. This appendix contains plots analogous to those in Figure 4.8 of Section 4.3.2, which explains how the temperature dependence observed in the ^{22}Na 511 keV photopeak was used to mitigate the temperature dependence of recorded energies in 2020 calibration data. The left figures show the fitted photopeak energy as a function of temperature. The right figures show the dependence after applying the linear correction in the left figure. The correction restores the photopeaks closer to their true energies.


As discussed in Section 4.3.2, limited time for calibration in the 2020 campaign restricted the temperature range over which data could be collected. The limited calibration time also restricted the quantity of data in each measured temperature bin (see, for example, the sparsely populated ^{57}Co plots in Figure A.1). Nevertheless, in performing this study, the author of this dissertation developed a method of temperature correction for 2020 calibration data which corrected the photopeaks to the same average accuracy ($\sim 0.1\%$) achieved by the 2016 correction method in Kierans (2018).




(a) ^{133}Ba : Uncorrected 81.0 keV




(b) ^{133}Ba : Corrected 81.0 keV




(c) ^{57}Co : Uncorrected 122.0 keV



(d) ^{57}Co : Corrected 122.0 keV




(e) ^{57}Co : Uncorrected 136.0 keV




(f) ^{57}Co : Corrected 136.0 keV


Figure A.1: Temperature correction in 2020 data: 81.0 keV, 122.0 keV, and 136.0 keV. Analogous to Figure 4.8.




(a) ^{133}Ba : Uncorrected 276.4 keV




(b) ^{133}Ba : Corrected 276.4 keV




(c) ^{133}Ba : Uncorrected 302.9 keV



(d) ^{133}Ba : Corrected 302.9 keV




(e) ^{133}Ba : Uncorrected 356.0 keV




(f) ^{133}Ba : Corrected 356.0 keV


Figure A.2: Temperature correction in 2020 data: 276.4 keV, 302.9 keV, and 356.0 keV. Analogous to Figure 4.8.




(a) ^{133}Ba : Uncorrected 383.0 keV




(b) ^{133}Ba : Corrected 383.0 keV




(c) ^{137}Cs : Uncorrected 661.7 keV



(d) ^{137}Cs : Corrected 661.7 keV




(e) ^{88}Y : Uncorrected 898.0 keV




(f) ^{88}Y : Corrected 898.0 keV


Figure A.3: Temperature correction in 2020 data: 383.0 keV, 661.7 keV, and 898.0 keV. Analogous to Figure 4.8.




(a) ^{60}Co : Uncorrected 1173.2 keV




(b) ^{60}Co : Corrected 1173.2 keV




(c) ^{22}Na : Uncorrected 1274.5 keV



(d) ^{22}Na : Corrected 1274.5 keV



(e) ^{60}Co : Uncorrected 1332.5 keV



(f) ^{60}Co : Corrected 1332.5 keV

Figure A.4: Temperature correction in 2020 data: 1173.2 keV, 1274.5 keV, and 1332.5 keV. Analogous to Figure 4.8.



(a) ^{88}Y : Uncorrected 1836.0 keV(b) ^{88}Y : Corrected 1836.0 keV

Figure A.5: Temperature correction in 2020 data: 1836.0 keV. Analogous to Figure 4.8.

Appendix B

Charge sharing in GRIPS detectors

B.1 Overview of charge sharing

Photon interactions in semiconductor detectors liberate charge clouds of electrons and holes. Under high voltage bias, the electrons (holes) drift towards (away from) the positively biased side of the detector. The drift of charge induces a current defined by the Shockley-Ramo theorem (Shockley 1938; Ramo 1939) on the detector’s strip electrodes. This current is integrated to measure the total charge (deposited energy) of the interaction. The energy measurement is critical to accurate event reconstruction. As the charge carriers drift towards the strip electrodes, however, they spread laterally due to charge carrier repulsion and thermal diffusion. The spreading gives rise to a deleterious effect called “charge sharing,” which refers to the triggering of multiple adjacent strips by a single interaction. Refer to Figure B.1 for a schematic of this effect.

Charge deposition over multiple strips complicates the measurement in several ways.

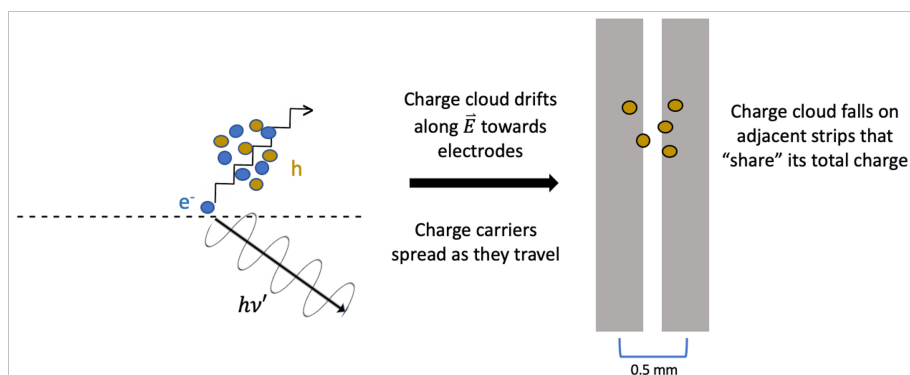


Figure B.1: Schematic of charge sharing in a detector with strip electrodes. Charge carriers liberated by an interaction have the potential to spread laterally as they drift under high voltage bias, “sharing” the charge across triggered adjacent strips.

Charge loss, or the loss of charge that falls between adjacent strips and is not measured, distorts the energy measurement on the affected side of the detector. Because the strip pairing algorithm constrains the interaction location by searching for coincident triggers on either side of a detector with equivalent energies (Section 4.5), charge sharing events with non-negligible charge loss are complicated to reconstruct. An additional correction for charge loss can be applied to help mitigate this effect (Bandstra 2010). Furthermore, depositing charge across multiple strips may split the total interaction energy into hits which fall below a triggered strip’s energy threshold. These sub-threshold hits may be reconstructed into an event with a misleading deficit in the total energy. A compounding difficulty is that the measurement uncertainty in deposited energy on each strip adds in quadrature, worsening the overall energy resolution of a charge sharing event that triggers multiple strips. Finally, charge sharing creates another class of events that is hard to distinguish from closely spaced multi-site events. Thus, it is important to study how charge sharing affects cross-strip detectors like those in COSI, properly model the effect in simulations for reliable comparison to real data (see Section 4.7), and implement understanding of charge sharing in advanced strip pairing and event reconstruction algorithms that can accommodate the challenges described above.

To this end, four adjacent strips of a cross-strip germanium detector were illuminated with a variety of radioactive sources to study the degree to which charge is deposited across multiple strips. This section presents descriptions of the spare detector, the data-taking procedure, and analysis of the measurements with implications for ongoing efforts to understand charge sharing.

B.2 The GRIPS detector

The Gamma-Ray Imager/Polarimeter for Solar flares (GRIPS) is a balloon-borne instrument for solar physics (Duncan et al. 2016). GRIPS is comprised of high-purity germanium detectors (GeDs) very similar to COSI’s, but with 149 strip electrodes per side of 0.5 mm strip pitch (Figure B.2a). The COSI balloon GeDs have 37 strips per side of 2 mm strip pitch and the COSI satellite GeDs will have 64 strips per side of 1.16 mm strip pitch. The GRIPS detectors therefore present an ideal test-bed for studying charge sharing: with very fine strip pitch, GRIPS detectors may be more susceptible to charge sharing than COSI detectors and provide detailed insight into charge sharing behavior.

B.3 Measurements

A spare GRIPS GeD (named HP50838-1) was instrumented in an aluminum cryostat at SSL in October 2021. It was cooled with a dewar of liquid nitrogen and biased to 600 V on the AC side. Note that the depletion voltage of the detector was determined to be ~ 800 V when manufactured, but raising the bias above 600 V at the time of these measurements increased




Figure B.2: Charge sharing measurements. (a) Spare GRIPS GeD with 149 strips per side of 0.5 mm strip pitch (distance between strip centers). Image from Duncan (2017). (b) Experimental setup. The spare GRIPS GeD is located inside the aluminum cryostat. Electronics and high voltage connections are housed in aluminum boxes above and to the left of the cryostat. The liquid nitrogen dewar, Picoscope 5444, and ^{241}Am sealed radioactive source are also labeled. The source is facing the DC side of the detector.

the leakage current to prohibitive levels. The recommended operating bias is 1500 V. To collect complete waveforms of induced signals, four adjacent strips were instrumented with discrete charge sensitive preamplifiers, the outputs of which were digitized and processed with a Picoscope 5444. The shaper was a digital CR-RC⁶ with 10 μs peaking time. An image of this setup is shown in Figure B.2b.

The energy threshold of each strip was determined by illuminating the detector with ^{241}Am . The transition between a large noise peak at low energies to an expected continuum marks an acceptable threshold. Setting thresholds to low energies maximizes sensitivity to charge sharing events; recording small energy depositions on adjacent strips is required for a complete picture of how charge carriers move throughout the detector. The thresholds (in arbitrary ADC units) on channels 0, 1, and 2 were set to 150 and that on channel 3 was set to 100. With the same procedure, these thresholds were confirmed as reasonable for ^{57}Co and ^{137}Cs sources as well. An energy calibration for each strip was performed in Melinator (see Section 4.3) using data from ^{241}Am , ^{57}Co , and ^{137}Cs sources placed on the DC side of the detector. Using this calibration, the channel thresholds were determined to be 5.56 keV, 5.53 keV, 5.74 keV, and 3.83 keV.

To study charge sharing as a function of energy and distance of the interactions from the electrodes, ^{241}Am , ^{57}Co , and ^{137}Cs were individually placed on both the DC and AC sides of the detector. The interactions inside the detector occur closer to (farther from) the DC-side channel readout with the radioactive sources positioned on the DC (AC) side. Collection




Figure B.3: Sample spectra collected on November 10, 2021 from channel 2 on the DC side of the spare GRIPS GeD. The radioactive sources were approximately centered on and placed ~ 6 inches from the DC side (see Figure B.2b). The ^{241}Am 59.5 keV photopeak is at 1722 ADC. The ^{57}Co 122.1 keV photopeak is at 3532 ADC and its 136.5 keV photopeak is at 3951 ADC. The ^{137}Cs 661.7 keV photopeak is at 19240 ADC.

times for ^{241}Am were about 30 minutes, for ^{57}Co about 25 minutes, and for ^{137}Cs about 1 hour, sufficient for pronounced photopeaks in each spectrum. The 59.5 keV line from ^{241}Am and 122.1 keV line from ^{57}Co had a maximum amplitude of several thousand counts, the 136.5 keV line from ^{57}Co had close to 1000, and the 661.7 keV line from ^{137}Cs peaked at around 100 counts. A subset of collected spectra is shown in Figure B.3.

B.4 Results

As a simple search for potential charge sharing events, we examine the number of triggered strips as a function of energy and DC-/AC-side placement. All hits recorded within a 10 sample (320 ns time step) coincidence window are grouped together as an event. Note that an event is not necessarily a photopeak event. Here, events are defined only by the coincidence window with no additional energy selection and are comprised of all hits which exceed the strips' energy thresholds. The results are shown in Figure B.4.

There is a higher proportion of single-strip events (events which trigger only one strip) in lower energy source data. This makes intuitive sense, as the size of the generated charge cloud increases with increasing photon energy. Additionally, there are more multi-strip events when the sources are placed on AC side. This behavior is expected because the liberated




Figure B.4: Percentage of triggered strips as a function of energy and (a) DC- and (b) AC-side placement. Total number of events on the DC side: $^{241}\text{Am} = 777850$, $^{57}\text{Co} = 3303679$, $^{137}\text{Cs} = 5503009$. Total number of events on the AC side: $^{241}\text{Am} = 2074656$, $^{57}\text{Co} = 3966138$, $^{137}\text{Cs} = 5305066$.

holes must travel a greater distance to the readout DC strips than they do when the source is, and therefore the interactions are, closer to the DC side. The spreading effect, or change in number of triggered strips as a function of DC-/AC-side placement, is most visible in ^{241}Am . This is also expected because lower energy photons from ^{241}Am will undergo rapid attenuation and interact close to the detector surface. Thus, they undergo more (less) drift when placed on the AC (DC) side than higher energy photons from ^{57}Co and ^{137}Cs , which interact deeper in the detector volume and likely experience a smaller change in drift distance with source location. It is also possible that the charge clouds from the higher energy ^{57}Co and ^{137}Cs are large enough to hit multiple strips without the extra diffusion incurred by traveling from the AC side of the detector, and so the distribution across strips remains more consistent with source location. Finally, very few events trigger 3–4 strips, and those that do are from the larger charge cloud of ^{137}Cs . An important caveat is that the multi-strip events seen here are not necessarily from charge sharing. Compton scattering, especially at higher energies, could trigger adjacent strips within the defined coincidence window. Nevertheless, the conclusions derived from these plots are consistent with expectations of charge carrier motion.

Next, charge sharing is examined in the context of potential charge loss. These studies are motivated by Chapter 7 of Looker (2014), which explores inter-electrode charge collection of a double-sided strip GeD with 5 strips of 1 mm pitch and 0.5 mm gap. Plotting the sum of charge on two strips against the charge deposited on one of those strips can reveal loss of charge falling between the two. In the case of no charge loss, events would fall on a

horizontal line such that the sum of the two channels' energies equals the incident photon energy. Events falling below this line are indicative of charge loss, as the two strips failed to recover the full energy of the photon.

Refer to Figure B.5 for plots which sum the charge deposited by ^{241}Am and ^{57}Co (near its 122.1 keV photopeak) on channels 1 and 2 (the middle two channels), including single-strip and shared events. Zero charge is deposited on channels 0 and 3 in these events. Plots are shown for both AC- and DC-side source placement and E_i denotes the energy (charge) on channel i .

The red boxes in Figure B.5 indicate potential charge loss in summed strips when the source is moved to the AC side. When the source is on the DC side, the majority of events have $E_1 + E_2 \sim 59.5$ keV and 122.1 keV, i.e. most of the photon's incident energy is recovered by the adjacent channels. When the source illuminates the AC side, however, there are more events with $E_1 + E_2 < 59.5$ keV and 122.1 keV. This may be evidence of lost charge between strips as the holes spread while traveling a greater distance to the DC-side channel readout, though it could also be evidence of lost charge to a neighboring strip in a sub-threshold deposit. The cyan boxes indicate more events with charge shared between channels 1 and 2 when the source is placed on the AC side than on the DC side: there are more events with $E_2 < 59.5$ keV and 122.1 keV when $E_1 + E_2 \sim 59.5$ keV and 122.1 keV. The effect is stronger in ^{241}Am likely due to reasons discussed previously, namely that there is a greater change in drift distance with lower energy ^{241}Am photons.

Figure B.6 is analogous to Figure B.5 with the 136.5 keV photopeak of ^{57}Co and 661.7 keV photopeak of ^{137}Cs . At these higher energies of 136.5 keV and 661.7 keV, the change in drift distance with source location is less than that at lower energies and it is more difficult to identify the clear changes in energy distribution seen at 59.5 keV and 122.1 keV. However, though the enhancement is less obvious, there are more events with $E_2 < 136.5$ keV when $E_1 + E_2 \sim 136.5$ keV with ^{57}Co on the AC side than on the DC side. The ^{137}Cs distributions appear very similar, suggesting comparable movement of charge through the detector regardless of DC-/AC-side placement.

As a point of comparison and additional sanity check, the analysis is repeated with ^{241}Am for the sum of channels 0 and 2 (one strip between the channels of interest, such that $E_1 = E_3 = 0$) and the sum of channels 0 and 3 (two strips between, such that $E_1 = E_2 = 0$). Figure B.7 shows little evidence of charge sharing when there are one or two strips between channels of non-zero charge. This is expected, as it is unlikely that an interaction would trigger two distanced strips without triggering the strips between them. The evidence of charge loss with increased drift distance remains, though, as seen in the diminished recovery of the full 59.5 keV photopeak energy when ^{241}Am is placed on the AC side.

Finally, we examine the distribution of charge across all channels. For each event, we find the percentage of total event energy measured on each channel. All events triggering any number of channels, i.e. ranging from single-strip events to events which trigger all four channels, are considered. No restrictions on total energy are applied. The results are shown in Figure B.8. The vertical line at 0% is from <4-strip events, as these events will always have at least one channel with zero recorded energy. The vertical line at 100% is from single-




Figure B.5: Summed charge on middle channels 1 and 2 plotted against the charge on channel 2. Only events which trigger channels 1 and 2 (zero energy on channels 0 and 3) are considered. The top (bottom) row shows ^{241}Am and ^{57}Co , near its 122.1 keV photopeak, when placed on the DC (AC) side of the detector. The red boxes indicate potential charge loss in summed strips when the source is moved to the AC side. The cyan boxes indicate more events with charge shared between channels 1 and 2 when the source is placed on AC side than on DC side, though the effect is stronger in ^{241}Am .

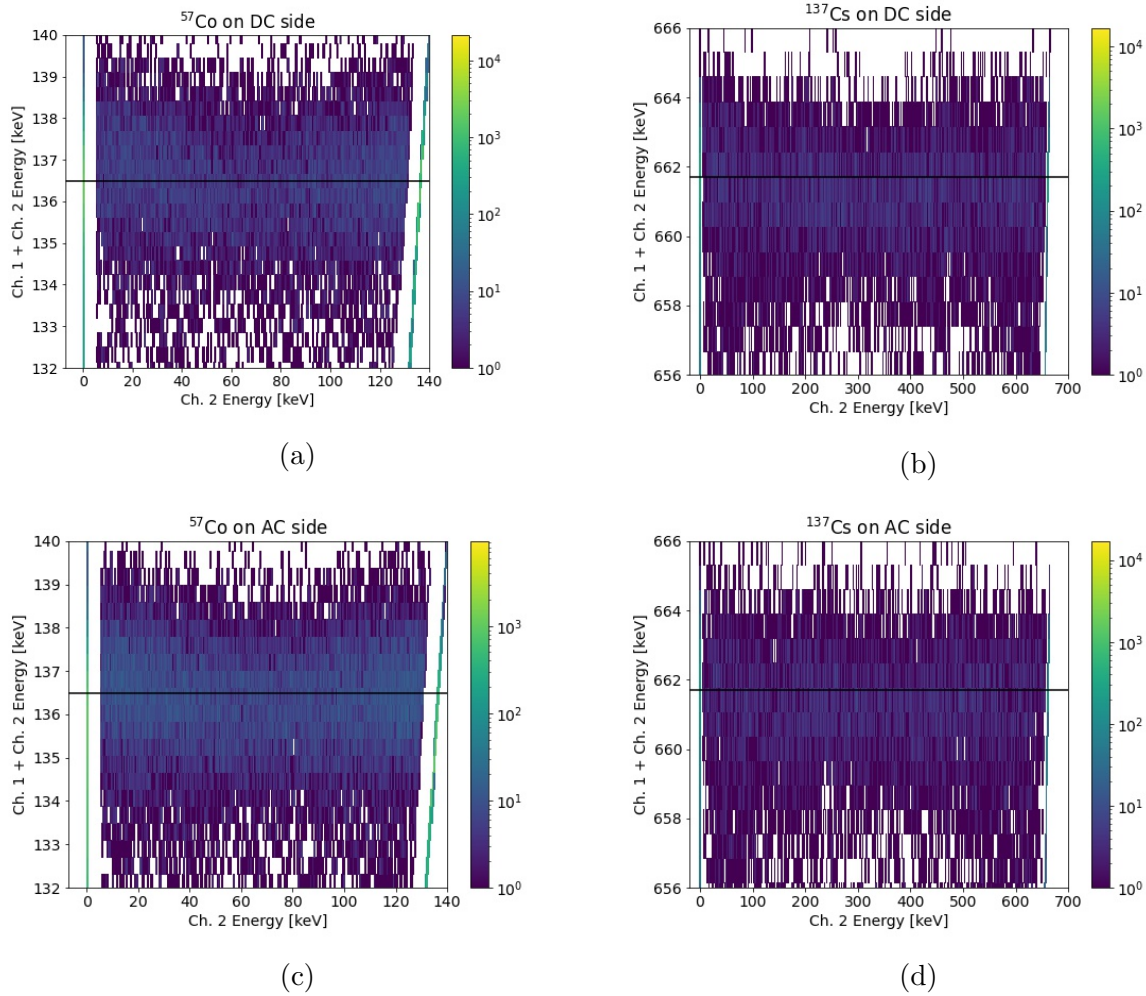


Figure B.6: Summed charge on middle channels 1 and 2 plotted against the charge on channel 2. Only events which trigger channels 1 and 2 (zero energy on channels 0 and 3) are considered. Analogous to Figure B.5. The top (bottom) row shows ^{57}Co , near its 136.5 keV photopeak, and ^{137}Cs when placed on the DC (AC) side of the detector.

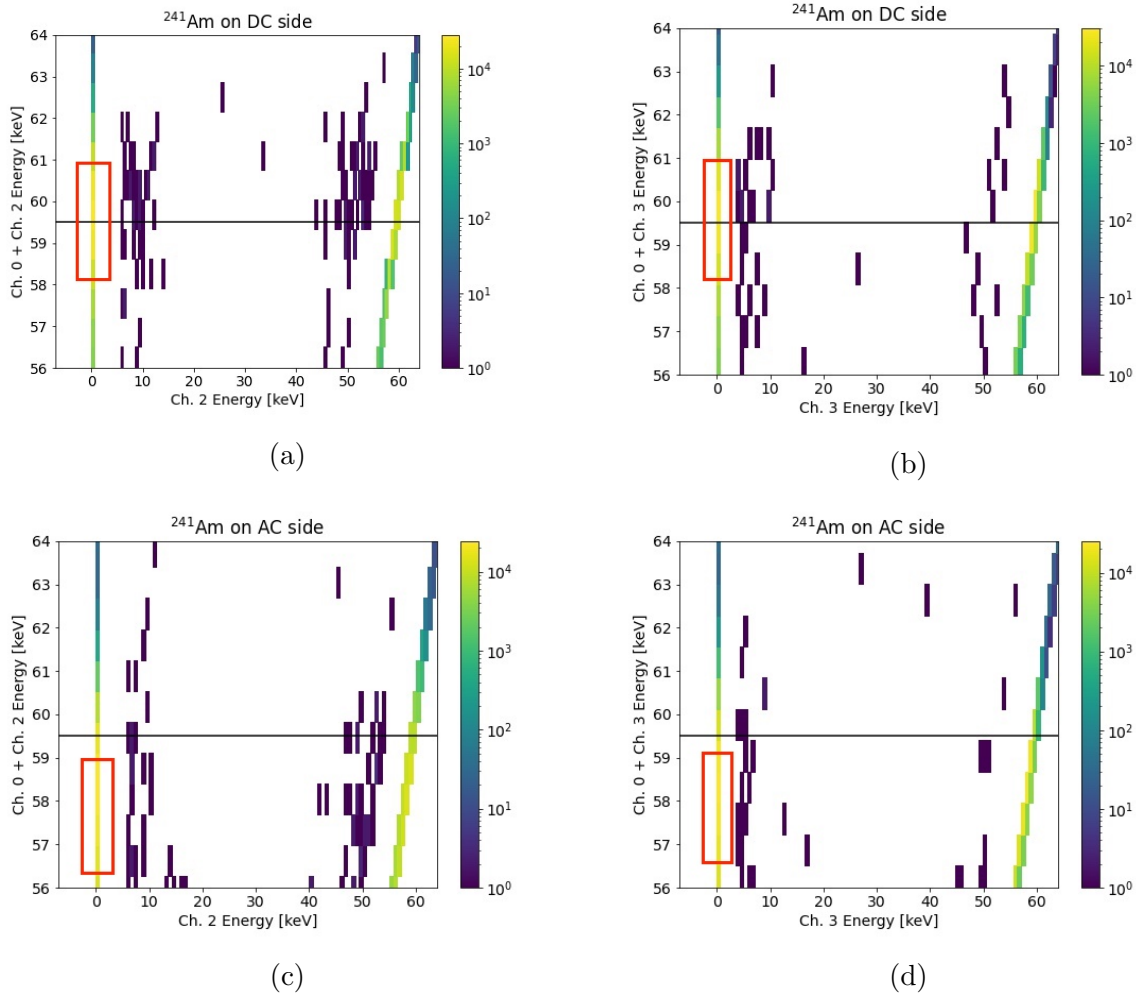


Figure B.7: Left column: Summed charge on channels 0 and 2 for ^{241}Am placed on the DC and AC side of the detector. Events deposit zero energy above threshold on channels 1 and 3. Right column: Summed charge on channels 0 and 3, where events deposit zero energy above threshold on channels 1 and 2.

strip events in which 100% of the energy is deposited in one channel. Note the broader distribution of energy with higher energy sources: the bigger charge clouds of higher energy photons are more likely to share charge on adjacent channels. This broader distribution is most easily visible in the gap between the vertical 0% bin and the uptick in counts at non-zero percentages to the right. The gap is widest for ^{241}Am and smallest for ^{137}Cs . Additionally, the dip in counts near 50% of total energy may be indicative of charge sharing or charge cloud size. Higher energy events from ^{57}Co and ^{137}Cs , compared to the flatter distribution in ^{241}Am , are less likely to deposit half of their total charge on one channel.

B.5 Conclusions

Overall, the presented measurements are consistent with basic expectations of charge carrier motion in semiconductor detectors. Charge clouds that undergo greater drift distances show more evidence of charge sharing, as do charge carriers liberated by photons of higher energy sources. Supplementary measurements are required to discern whether the hits recorded on adjacent strips are necessarily a consequence of charge sharing or, for example, a consequence of Compton scattering. Simultaneously reading out orthogonal strips on the AC side of the detector with the DC-side strip readout will be an important step in advancing these studies. Time constraints imposed by rapid development of COSI satellite electronics prohibited these additional measurements, along with analogous measurements on a spare 37-strip COSI GeD for comparison to a wider strip pitch. It will also be important to use a properly biased detector in future measurements. The measurements in this appendix, taken with a severely under-biased detector with a correspondingly weak electric field and slow drift velocities, are potentially complicated by the effect of recombination of the carriers due to finite carrier lifetime.

A simplified calculation of charge cloud size, however, may suggest that charge sharing is a reasonable mechanism of the observed behavior. The 600 V bias across the 1.5 cm thickness of the GRIPS GeD gives an electric field of roughly $E = 400 \text{ V cm}^{-1}$. Figure 11.2c of Knoll (2010) indicates that the drift velocity of holes in germanium is $\sim 4 \times 10^6 \text{ cm s}^{-1}$. Assuming that an interaction on the AC side of the detector drifts the full 1.5 cm to the DC-side readout, the drift time t is $\sim 375 \text{ ns}$ (note that the assumed drift velocity is a coarse approximation, given that the detector was under-biased at only 600 V). Hole mobility in germanium at 77 K is $\mu = 4.2 \times 10^4 \text{ cm}^2 \text{ V}^{-1} \text{ s}$ (Table 11.1 of Knoll 2010). Then, for Boltzmann constant $k = 1.38 \times 10^{-23} \text{ J K}^{-1}$, temperature $T \sim 80 \text{ K}$, and electronic charge $e = 1.6 \times 10^{-19} \text{ C}$, the lateral diffusion coefficient $D = \mu kT/e = 290 \text{ cm}^2 \text{ s}^{-1}$. The standard deviation of charge cloud size due to diffusion is given by $r_d = \sqrt{Dt} \sim 0.01 \text{ cm}$, and the initial charge cloud size r_i from electron range data is $1.4 \times 10^{-3} \text{ cm}$ at 60 keV and $4.4 \times 10^{-3} \text{ cm}$ at 122 keV (NIST 2017). Thus, the final charge cloud size $r_f = \sqrt{r_d^2 + r_i^2}$ is approximately 0.10 mm for 60 keV photons and 0.11 mm for 122 keV photons. Both sizes exceed the 0.06 mm gap between strips in the tested GRIPS detector, preserving charge sharing as a possible actor in the measurements. Refer to Boggs (2023) for more detailed modeling of charge cloud size.

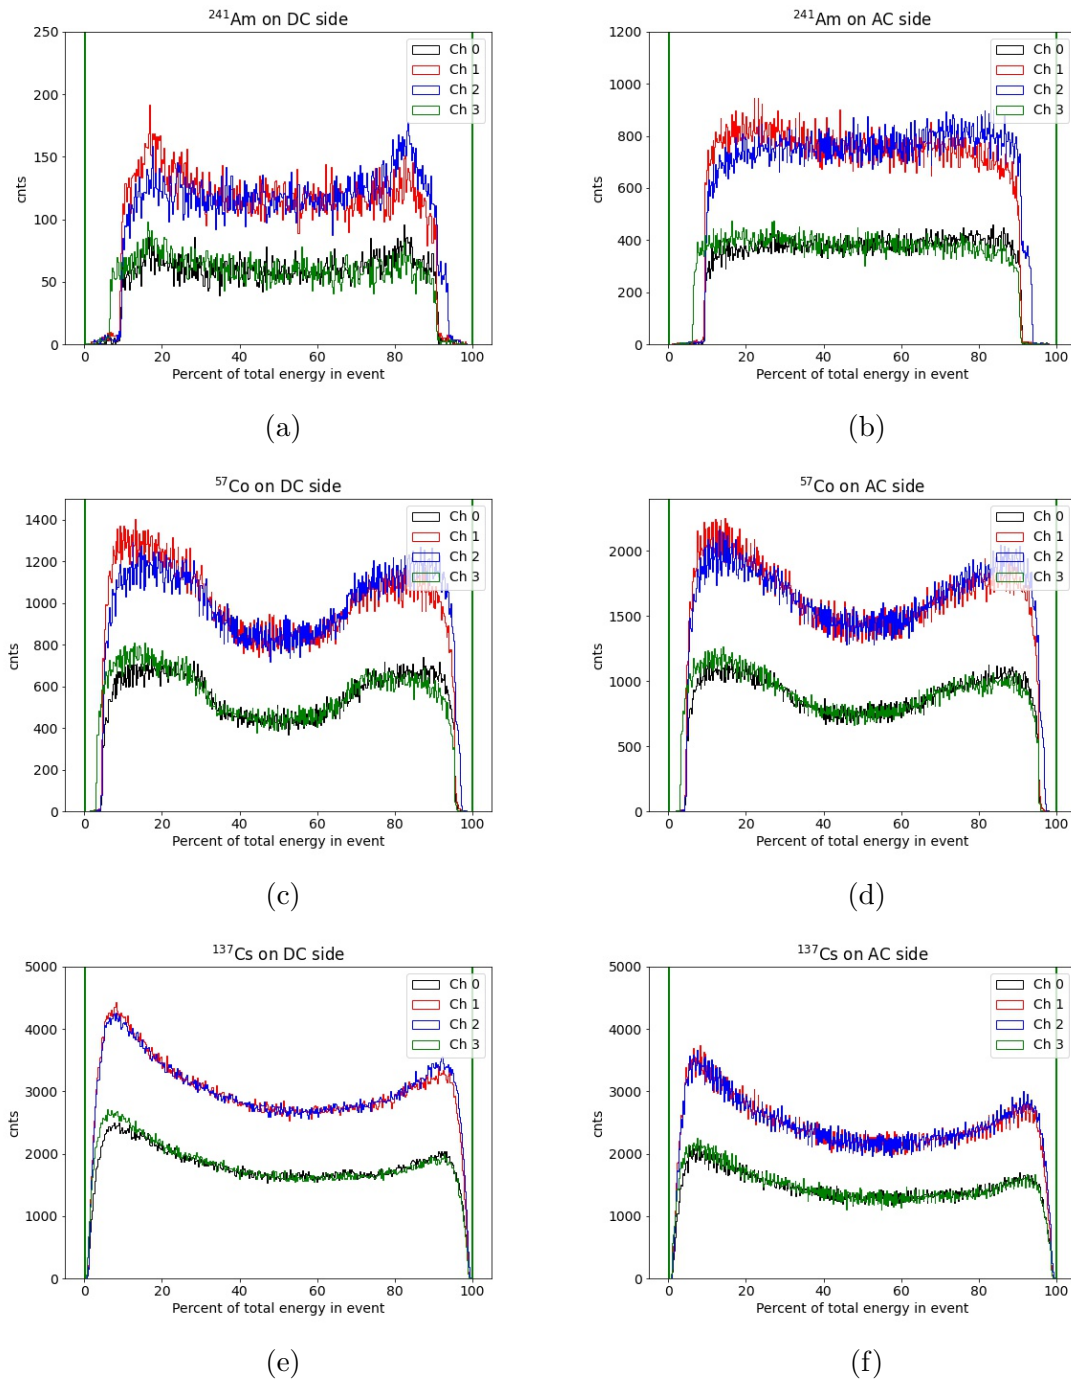


Figure B.8: Percentage of total event energy measured on each channel for ^{241}Am , ^{57}Co , and ^{137}Cs placed on the DC (left) and AC (right) sides of the detector. No restrictions on event energy or number of triggered channels are applied.

Furthermore, comparing these measurements to simulated waveforms will test physical understanding of charge sharing that can be implemented in models of detector effects. Robust simulations will also inform future mission designs that must strike a balance between fine position resolution with small strip pitch and problematic charge sharing. To this end, efforts are underway to use a Julia package, `SolidStateDetectors.jl` (Abt et al. 2021), to model charge sharing in GRIPS detectors. This package simulates the generation and drift of charge carriers through a user-defined detector volume and produces the resulting waveforms on the detector's electrodes. Future work will compare these simulated results to the measurements discussed in this appendix.

The conclusions derived from this appendix and ongoing efforts will be used to improve COSI's detector effects engine (DEE; Section 4.7) and strip pairing algorithm (Section 4.5), with similar usage in GRIPS analysis as well. Note that because the influence of charge sharing increases with increasing photon energy, understanding this effect has important implications for the primary science topic of this dissertation: measuring Galactic ^{26}Al via its 1.809 MeV γ -ray signature (Chapter 5), which lies at the upper end of COSI's energy range. Improving the DEE and strip pairing algorithm in the COSI satellite mission will help COSI advance its studies of ^{26}Al , especially because the COSI satellite detectors will have a finer strip pitch than that of the balloon detectors.

Appendix C

Charge trapping in COSI detectors

C.1 Background

Charge carriers produced by the interaction of a γ -ray with a semiconductor detector ideally travel uninhibited to collecting electrodes. In this case, the full energy of the incident γ -ray is recovered. In practice, small concentrations of impurities, typically gold, zinc, cadmium, or other metallic atoms, in the detector volume can capture transient electrons and holes as they move towards the electrodes (Knoll 2010). These sites trap the carriers for some time, removing their influence from the induced signal; even if a trapped carrier is released from the site, too much time may have elapsed for it to register on an electrode as part of the original interaction. Thus, charge trapping in semiconductor detectors can result in pulse height deficits which falsely underestimate the energy of an event. Understanding how trapping sites affect the signal is consequently critical to maximizing spectroscopic performance.

The radiation environment of balloon missions presents an interesting laboratory in which to study charge trapping: radiation from heavy charged particles can transfer damaging energy to the atoms of the detector and disrupt its crystalline lattice structure. The GeDs in COSI are accordingly examined for signs of charge trapping that manifest as lower pulse heights in 2016 and 2020 calibration data. The calibration data from 2016 were taken prior to the 46-day balloon flight and are a useful point of comparison against the condition of the detectors in 2020. However, there is no significant observed change in overall instrument performance before and after the balloon flight, giving little indication that the detectors sustained appreciable radiation damage.

This appendix presents the first analysis of COSI spectra, generated separately from signals on the individual sides of each detector, as a function of distance from the electrodes. There is strong evidence of trapping for both charge carriers. Given that the detectors did not suffer radiation damage in the 2016 flight, the effect we see on the spectral performance of the detectors is a distinct indication of intrinsic, rather than radiation-induced, charge trapping in the germanium. This is a significant finding uniquely enabled by COSI's ability to measure the depth of interaction for individual events. To date, charge trapping in GeDs has

been widely considered negligible unless the detectors have been exposed to a high radiation environment. Evidence of intrinsic electron trapping in Hull et al. (2014) similarly competes with the expectation of minimal trapping without radiation damage. It is clear that intrinsic trapping should no longer be considered a negligible effect and that understanding charge trapping is fundamental to achieving the optimal resolution of these detectors.

C.2 Manifestation in COSI

A change in recorded energy with interaction depth may indicate trapping of charge carriers as they drift through the detector volume. Note that the electronics actually measure pulse heights that are interpreted as energies in the calibration process; because COSI's energy calibration does not consider the depth of the interaction, the trapping effects will appear as lower energies. To this end, single-site events (events in which the entire photon's energy is deposited in one interaction) in each detector are parsed for recorded energy and interaction location. To ensure sufficient statistics across the full depth of the detectors, the higher energy photopeak of 661.7 keV from ^{137}Cs data is chosen for the analysis. Lower energy sources are unable to penetrate the full detector volume and instead produce interactions largely at the top face of the detector. The recorded energies on the p-side (HV, electron-collecting) and n-side (LV, hole-collecting) of each detector are restricted to 652–672 keV and the events are binned by interaction depth (z -coordinate) into 0.2 cm bins, in accordance with COSI's ~ 2 mm depth resolution.

The expected distribution is visible when all events in the cryostat are binned by depth (Figure C.1). Exponential attenuation decreases the number of recorded events with increasing interaction depth. The three distinct sections correspond to COSI's three detector stacks separated in z : the top stack spans $z \sim [-0.4, 1.3]$ cm, the middle stack $\sim [-2.8, -1.4]$ cm, and the bottom stack $\sim [-5.4, -4]$ cm.

With this affirmation of expected behavior, the nine most populated depth bins in each detector are chosen to identify clear trends in the data. Some detectors in the middle and bottom stacks only have sufficient statistics for seven or eight populated depth bins. The spectrum of events from each bin is generated using energies recorded on both the p- and n-sides of each detector. The spectra are then fitted with a Gaussian. Figure C.2 shows the spectra from detector 0. To highlight the shift in spectra rather than absolute energy measurements that depend on each year's particular calibration, the spectra in this analysis are normalized to their mean energy.

If the detectors had no intrinsic charge trapping or if the energy calibration accounted for interaction depth when translating the misleadingly-small pulse heights to energy, the fitted spectra would have precisely the same centroid and indicate charge carrier recovery independent of interaction depth. On the contrary, the centroids of the Gaussian fits in Figure C.2 exhibit a dependence on the interaction depth of the incident photon. Plotting the centroid in each bin as a function of depth illustrates the dependence in detector 0




Figure C.1: Depth distribution of all 652–672 keV single-site events from COSI 2016 (left) and 2020 (right) ^{137}Cs calibration data. The expected exponential attenuation with increasing interaction depth is observed. Depth (z) bins are 0.2 cm wide.

more clearly (Figure C.3). Analogous plots for the other COSI detectors are provided in Figures C.5, C.6, C.7, and C.8.

Overall, the trends in these plots are similar to the evidence of electron trapping presented in Hull et al. (2014) (Figure C.4). In a detector with known electron trapping, Hull et al. (2014) see underestimated energies from diminished pulse heights in depth bins farther from the positively biased contact. The electrons in these bins must travel greater distances to reach the HV contact and as such are more likely to encounter traps in the detector volume. Depth bin 9 is the farthest from the HV contact and yields the lowest pulse heights of all bins.

Taking COSI’s detector 0 in Figure C.3 as an example, we find analogous behavior. Detector 0 is oriented with its n-side (LV, hole-collecting) upwards. Positive z points upwards towards the zenith of the COSI instrument, where the calibration source was placed in both years. Hence, the rightmost z value in the x -axes of Figures C.3, C.5, C.6, C.7, and C.8 corresponds to the side of the detector facing directly up towards the calibration source. Detectors 1–2 and 6–8 were also oriented with their n-sides facing upwards in the COSI cryostat. Detectors 3–5 and 9–11 were instead oriented with their p-sides facing upwards. It is clear in Figure C.3 that the p-side (electron collecting) of detector 0 records lower pulse heights for events interacting farther from that side. The same logic applies to COSI detectors of opposite orientation. Detector 3 in Figure C.5 is oriented with the p-side upwards. The




Figure C.2: Spectra in the nine most populated depth bins of detector 0 in COSI (a) 2016 and (b) 2020 ^{137}Cs calibration data. Separate spectra are shown for readout on the p- (left) and n- (right) sides of the detector and are normalized to their respective means. The centroid of the fitted Gaussian exhibits a dependence on the interaction depth of the incident photon.




Figure C.3: The centroid of the Gaussian fit to COSI ^{137}Cs spectra as a function of interaction depth in detector 0 (n-side facing upwards, i.e. closest to the calibration source and on the rightmost part of the x -axis). Results from 2016 (left) and 2020 (right) calibration data are shown for readout on the p- and n-sides of the detector. The error bar in y is the standard deviation of the fitted centroid. The error bar in x spans the 0.2 cm width of each depth bin.

electrons originating from events farthest from this contact, namely from the most negative z , induce the lowest centroids.

Hole trapping is also visible in the COSI detectors. Holes drift in the direction opposite the marked electron drift. Again taking detector 0 in Figure C.3 as an example, we envision holes traveling from left to right along the x -axis, towards positive z (n-side facing upwards). The pulse height recorded on the n-side (hole-collecting) of the detector is lowest for interactions which occur farthest from it, reflecting possible trapping of holes as they move greater distances through the germanium crystal. Without proper correction for the trapping effects, this is interpreted as lower energy. Note that the curve in Figure C.4 from Hull et al. (2014) is not an exact comparison to the COSI curves in this appendix. Hull et al. (2014) averaged the individually read-out p-side and n-side pulse heights to produce a single measured energy. This averaging reduces the visibility of trapping effects in the charge carrier with less trapping, which in this case is the holes, and preferentially reveals the effects of dominant electron trapping in this detector. Preserving the individual read-out without averaging would improve sensitivity to minor levels of intrinsic hole trapping.

Detailed comparison between the COSI 2016 and 2020 results is beyond the scope of this work. Understanding the differences may require extensive review of small changes in




Figure C.4: Evidence from Hull et al. (2014) of electron trapping in a segmented planar HPGe detector. The centroid energy of ^{137}Cs data exhibits a clear dependence on depth. The greatest charge loss due to trapping is seen in depth bin 9, which is farthest from the positively-biased (electron-collecting) contact. This bin contains events whose liberated electrons must traverse the greatest distance across the detector volume. Image from Hull et al. (2014).

instrument configuration between the years, including adjustments to electronic readout or calibrations, or even a study of the condition of individual detectors. It is feasible that low-level charge trapping analyses could more readily expose the effects of these changes than higher-level metrics of instrument performance. This appendix instead serves to document the first observation of clear intrinsic charge trapping in COSI’s GeDs.

C.3 Future work

The COSI group is currently conducting additional work to understand the severity of this effect in its HPGe detectors. This is especially relevant for COSI as a satellite mission, which will likely endure radiation damage dominated by > 20 MeV trapped protons over its 2-year prime mission in low-Earth orbit. Radiation damage preferentially creates hole traps which, though they have a minor effect on the electron signal, can significantly impact the hole signal. Reliable measurements of both signals are required to exploit the highest-attainable spectral resolution of the instrument: the motion of electrons and holes together contributes to the total induced signal on the strip electrodes. This preliminary analysis shows that COSI calibration data can be utilized to benchmark intrinsic charge trapping in germanium detectors and set a baseline for experiments that will study trapping in spare HPGe detectors deliberately bombarded with damaging radiation.

Furthermore, efforts are underway to empirically model charge loss in COSI measurements. The chosen Gaussian fit in the above analysis does not trace evident tailing at low energies and as such is only an approximation of the recorded peak. While this approximation is sufficient to identify the influence of charge trapping, it cannot correct for it. A tailored fit function that better matches the observed spectra would define an empirical correction capable of manually recovering lost charge in real data. Correcting the spectra would also reduce the width of measured photopeaks, currently broadened by the tailing, and thereby improve spectral resolution. Lastly, incorporating the empirical model in the Detector Effects Engine (Section 4.7) would reproduce charge trapping effects in simulations and inform more reliable expectations of instrument performance.

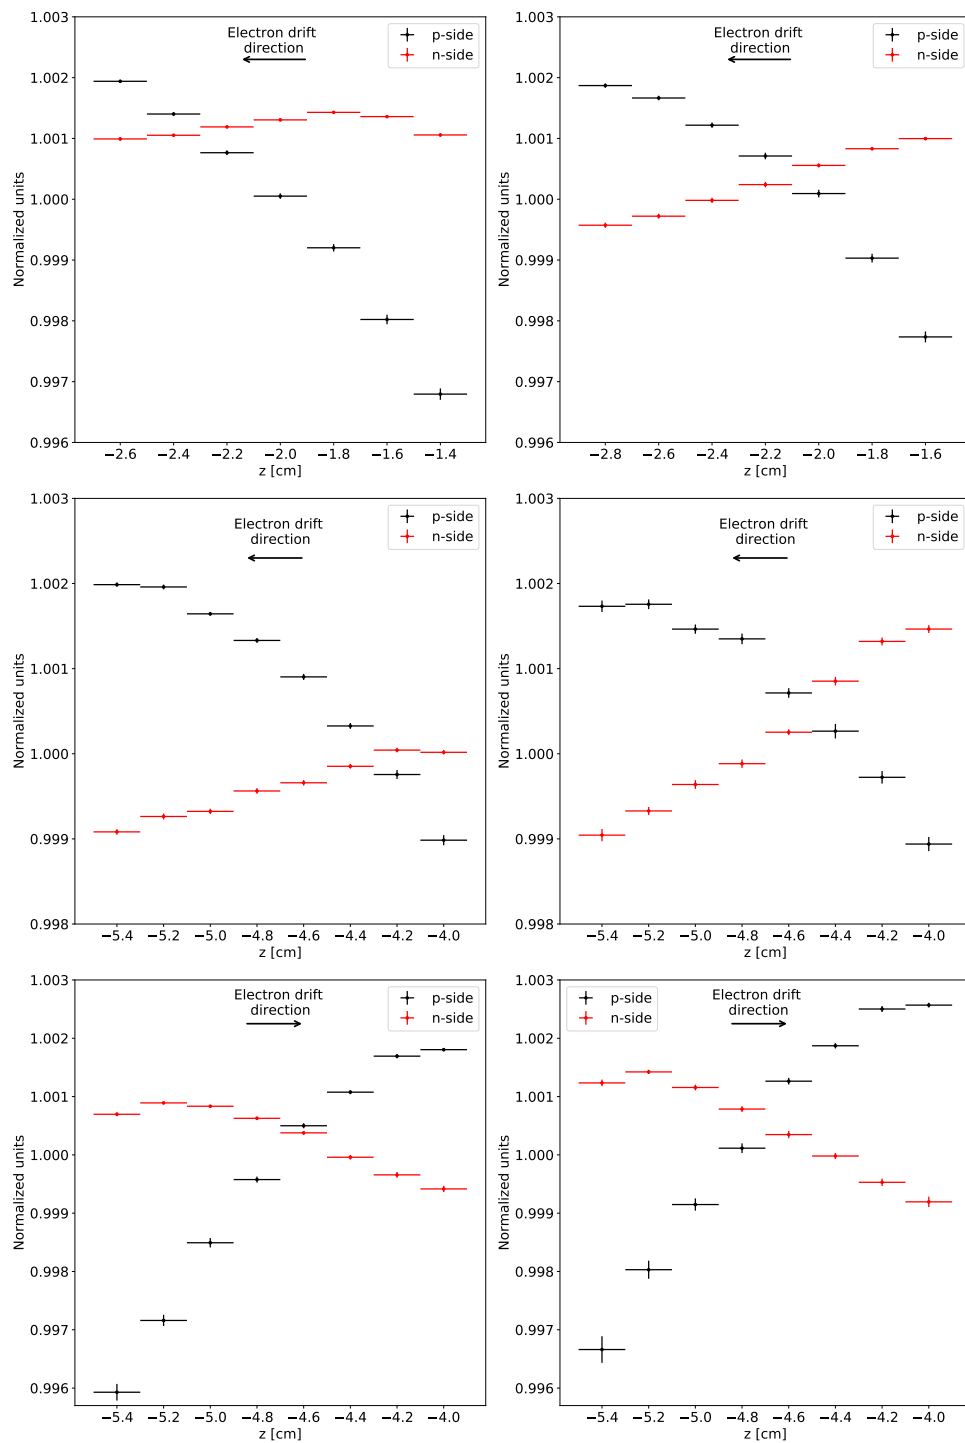


Figure C.5: Fitted centroid vs. depth in COSI detectors 1 (top, n-side up), 2 (middle, n-side up), and 3 (bottom, p-side up). Left: 2016 data. Right: 2020 data. Analogous to Figure C.3.

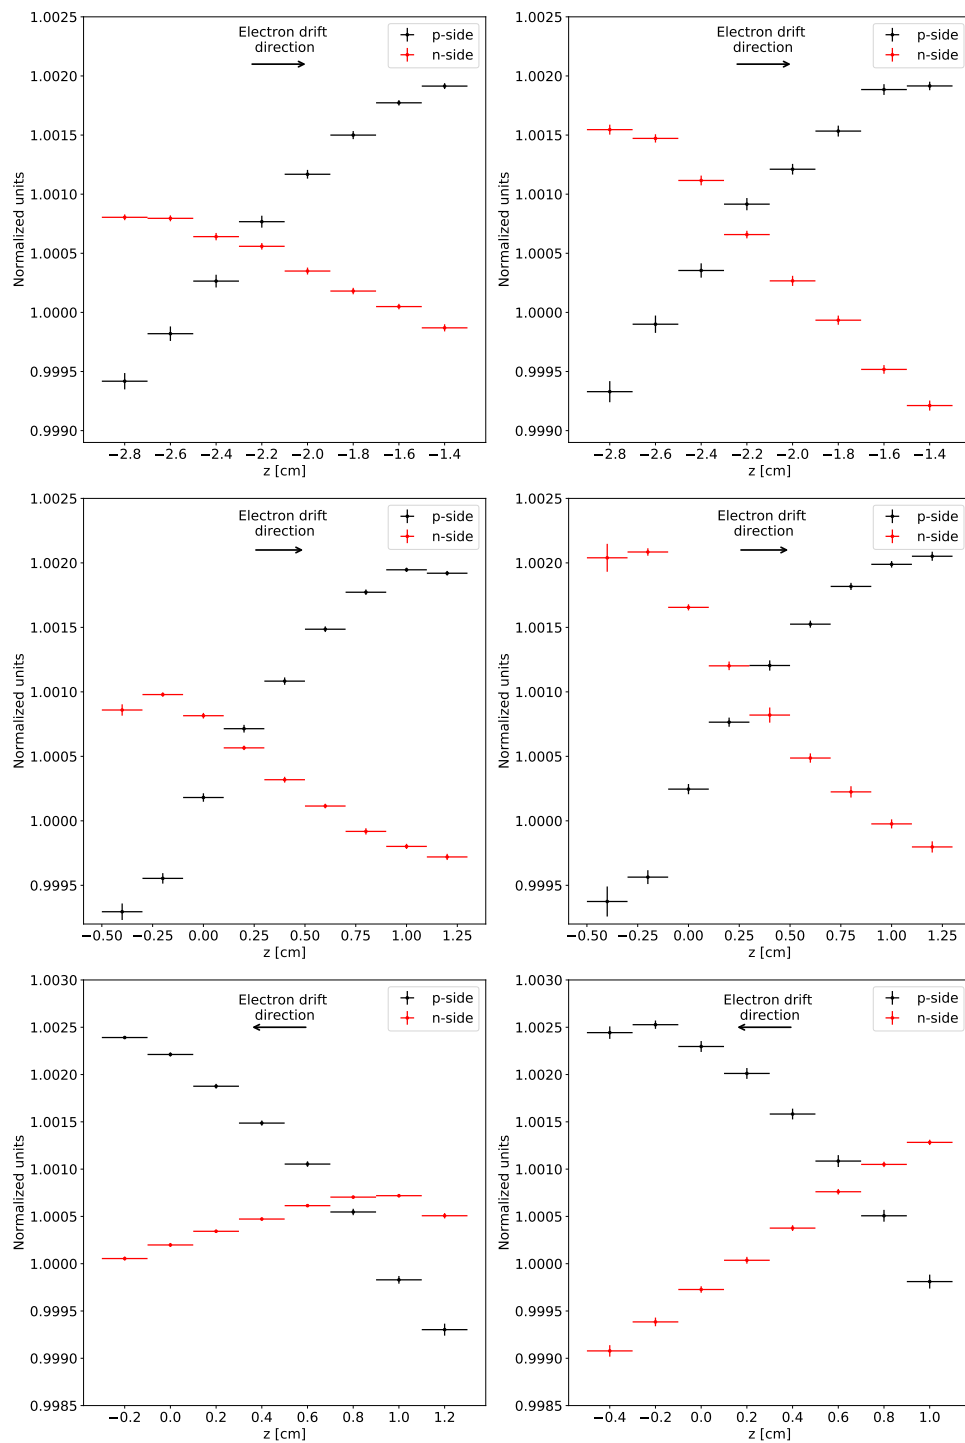


Figure C.6: Fitted centroid vs. depth in COSI detectors 4 (top, p-side up), 5 (middle, p-side up), and 6 (bottom, n-side up). Left: 2016 data. Right: 2020 data. Analogous to Figure C.3.

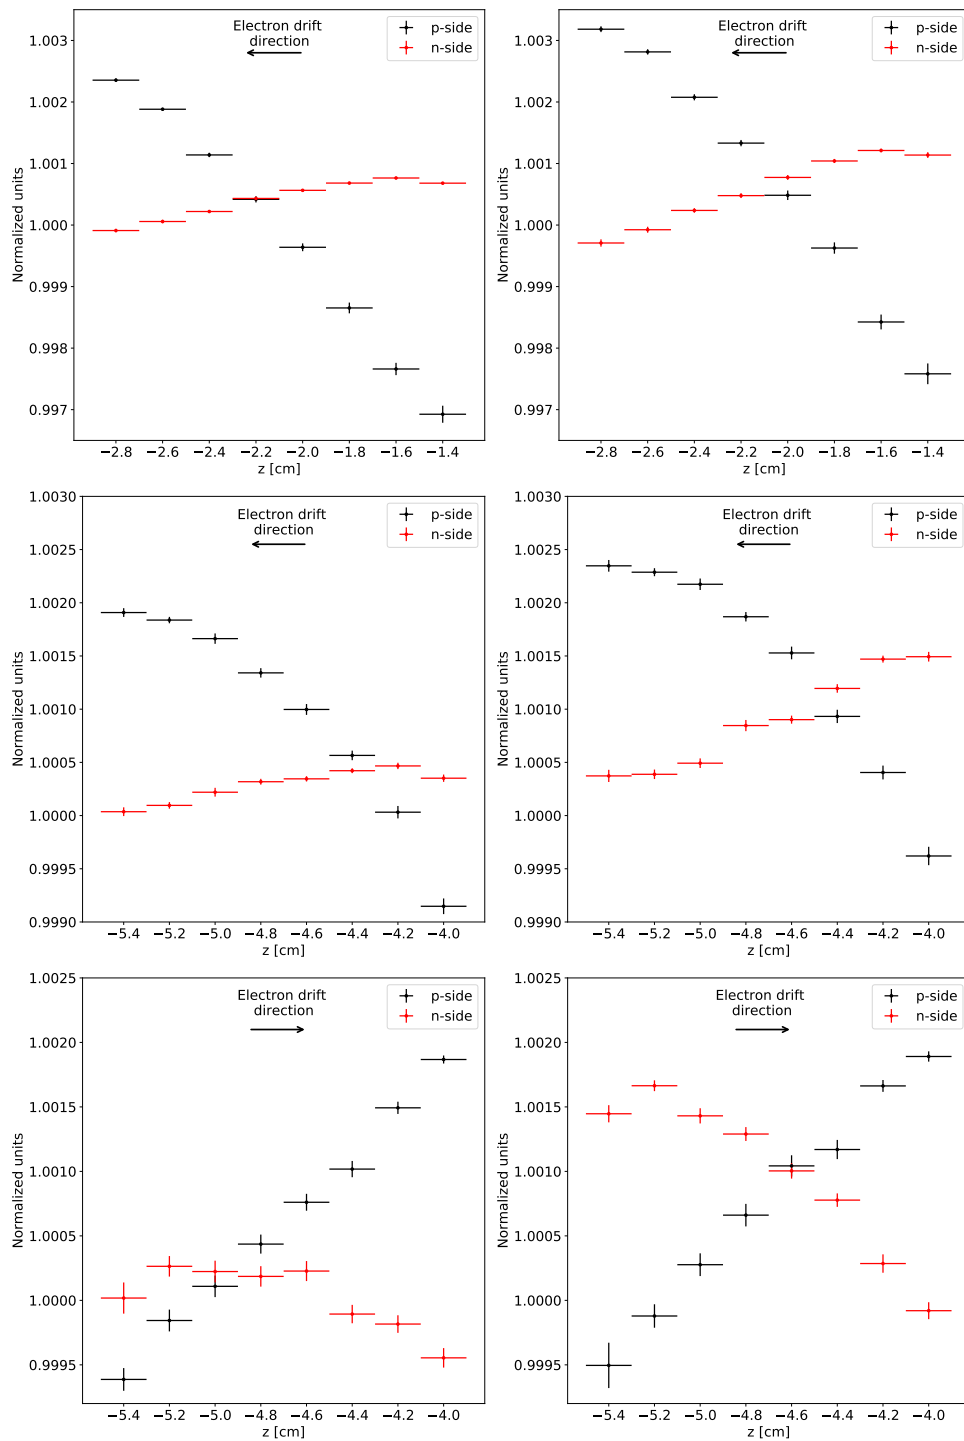


Figure C.7: Fitted centroid vs. depth in COSI detectors 7 (top, n-side up), 8 (middle, n-side up), and 9 (bottom, p-side up). Left: 2016 data. Right: 2020 data. Analogous to Figure C.3.

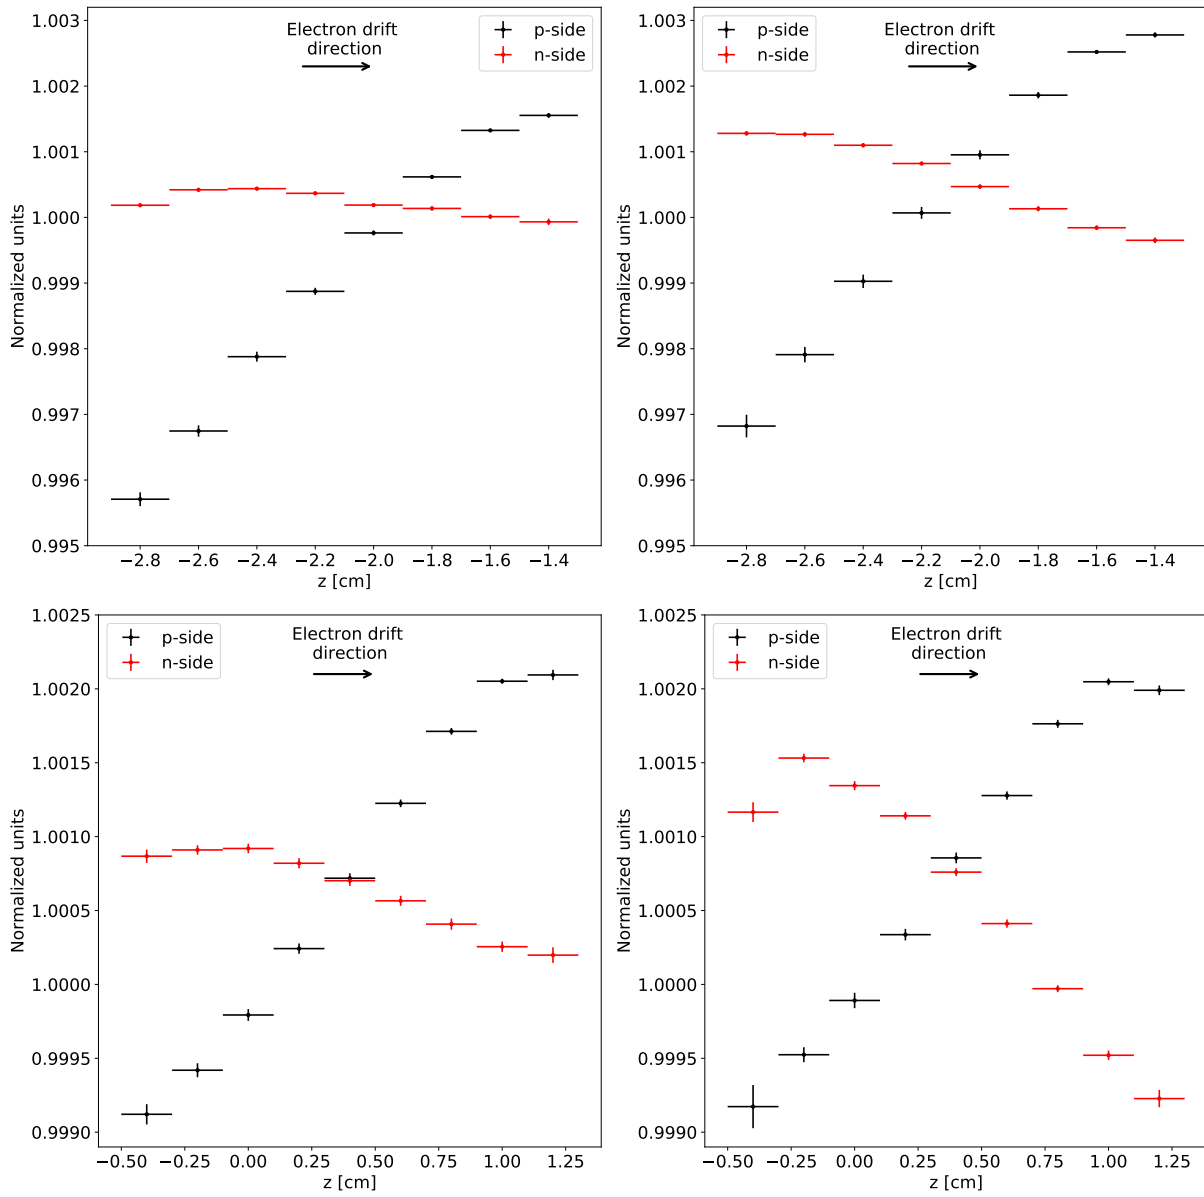


Figure C.8: Fitted centroid vs. depth in COSI detectors 10 (top, p-side up) and 11 (bottom, p-side up). Left: 2016 data. Right: 2020 data. Analogous to Figure C.3.

Appendix D

Derivation of the Compton equation

A basic derivation of the Compton equation is provided here for completeness.

Consider a photon with initial energy $E_{\gamma i} = h\nu$ traveling along the horizontal axis towards an electron that is both free (unbound) and at rest. The electron has rest energy E_{ei} . The photon collides with the electron and scatters at angle ϕ from the horizontal with an energy of $E_{\gamma f} = h\nu'$ after the collision. The electron scatters at an arbitrary angle with energy E_{ef} . The Compton equation describes the energy of the photon after the collision, $h\nu'$, as a function of the Compton scattering angle ϕ .

Per conservation of energy:

$$\begin{aligned} E_{\gamma i} + E_{ei} &= E_{\gamma f} + E_{ef} \\ E_{\gamma i} + m_e c^2 &= E_{\gamma f} + \sqrt{(p_{ef}c)^2 + (m_e c^2)^2} \\ (E_{\gamma i} + m_e c^2 - E_{\gamma f})^2 &= p_{ef}^2 c^2 + m_e^2 c^4 \end{aligned}$$

Call the final equation in the above sequence “Equation 1.”

Per conservation of momentum:

$$\begin{aligned} \vec{p}_{\gamma i} + \vec{p}_{ei} &= \vec{p}_{\gamma f} + \vec{p}_{ef} \\ \vec{p}_{\gamma i} + 0 &= \vec{p}_{\gamma f} + \vec{p}_{ef} \\ (\vec{p}_{\gamma i} - \vec{p}_{\gamma f}) \cdot (\vec{p}_{\gamma i} - \vec{p}_{\gamma f}) &= (\vec{p}_{ef})^2 \\ |\vec{p}_{\gamma i}|^2 + |\vec{p}_{\gamma f}|^2 - 2\vec{p}_{\gamma i} \cdot \vec{p}_{\gamma f} &= |\vec{p}_{ef}|^2 \\ p_{\gamma i}^2 + p_{\gamma f}^2 - 2p_{\gamma i}p_{\gamma f}\cos\phi &= p_{ef}^2 \\ c^2[p_{\gamma i}^2 + p_{\gamma f}^2 - 2p_{\gamma i}p_{\gamma f}\cos\phi] &= c^2[p_{ef}^2] \\ E_{\gamma i}^2 + E_{\gamma f}^2 - 2E_{\gamma i}E_{\gamma f}\cos\phi &= p_{ef}^2 c^2 \end{aligned}$$

Call the final equation in the above sequence “Equation 2” and plug Equation 2 into Equation 1:

$$(E_{\gamma i} + m_e c^2 - E_{\gamma f})^2 = E_{\gamma i}^2 + E_{\gamma f}^2 - 2E_{\gamma i}E_{\gamma f}\cos\phi + m_e^2 c^4$$

Recalling that $E_{\gamma i} = h\nu$ and $E_{\gamma f} = h\nu'$, algebraic simplification results in the Compton equation:

$$h\nu' = \frac{h\nu}{1 + \frac{h\nu}{m_e c^2}(1 - \cos\phi)}$$

The Compton equation is also commonly written in terms of the wavelength of the photon. Using $\nu = c/\lambda$ and rearranging terms, we have

$$\lambda' = \lambda + \frac{h}{m_e c}(1 - \cos\phi)$$

The quantity $\frac{h}{mc}$ is referred to as the “Compton wavelength.” The Compton wavelength of a particle with rest mass m (here, an electron with mass m_e) equals the wavelength of a photon with energy equal to the rest mass of the particle.

Appendix E

Optimization of the Compton scattering angle in Aluminum-26 analysis

To preferentially select ^{26}Al events over the abundant background events in both the signal and background regions, we employ a scanning procedure over the Compton scattering angle ϕ to identify an ideal range of allowed ϕ -values in the signal and background spectra. Identifying the maximum value also informs selection of the pointing cuts listed in Table 5.1 that define the signal and background regions. This ϕ_{max} effectively broadens the region of the sky included in each pointing cut because photons recorded in each region may originate up to ϕ_{max} outside of that region. The signal region (the Inner Galaxy) is broadened by ϕ_{max} to $(|\ell| \leq 30^\circ + \phi_{\text{max}}, |b| \leq 10^\circ + \phi_{\text{max}})$. To avoid overlap between the signal and background regions, the latter is defined such that the extent of its ϕ_{max} -broadened border encloses everywhere outside of the broadened signal region. Identifying the ideal minimum and maximum ϕ is discussed in this appendix¹.

The representative signal region data for this procedure are ^{26}Al events generated with an all-sky simulation of COSI’s response to the DIRBE 240 μm map. We run the simulation for the 9- and 10-detector flight configurations of the instrument. The representative background data set is comprised of atmospheric photons, simulated with the Ling model (Ling 1975), on June 12, 2016. COSI’s altitude on this day remained fairly stable at its nominal flight altitude of $\sim 33\text{ km}$ and it had nine active detectors. The 33 km altitude here conforms with the 33 km minimum altitude in the signal region event selections of the ^{26}Al analysis presented in Chapter 5. Note that we use simulations for this optimization study rather than real data because real data are subject to uncertainties and are always background-dominated, prohibiting a clean comparison of Compton scattering angles from ^{26}Al and background photons.

¹This appendix largely follows Appendix B of the published paper “Measurement of ^{26}Al by the Compton Spectrometer and Imager” by Beechert et al. (2022b). ©Astrophysical Journal 2022, reprinted with permission.

The ^{26}Al simulation is binned into one time bin for each configuration of the flight (9 detectors: May 17, 2016 to June 5, 2016; 10 detectors: June 6, 2016 to July 2, 2016). Likewise, the background simulation is binned into one time bin spanning June 12, 2016. Only Compton events with incident energy of 1803–1817 keV are considered to restrict the analysis to the ^{26}Al line band. We seek the optimal range of allowed Compton scattering angles which removes more background than ^{26}Al events from the simulated data set.

Figure E.1a shows that in both the ^{26}Al and background simulations, the large majority of events have ϕ less than 60° . The sharp drop in events, all comprised of 2–7 interactions, after $\sim 15^\circ$ is a consequence of 2-site events in this regime which have two possible flight paths and cannot be reconstructed (Zoglauer 2005). These events are rejected by the event reconstruction algorithm, creating the cumulative decrease which is also visible in flight data (Figure E.1b).

Overall, the background events scatter at smaller Compton scattering angles than the simulated ^{26}Al events, even though both distributions are limited to events of incident energy 1803–1817 keV. A plausible physical explanation for why the background events are more forward-scattered is that a high-energy background photon of, for example, 5 MeV may impart only 1.8 MeV in COSI’s detectors and escape the instrument without a final photoabsorption interaction. The photon would then exit the detector with 3.2 MeV of energy having deposited a false 1.8 MeV signature. Because the hypothetical photon’s true energy is greater than that recorded by COSI, it would Compton scatter at smaller angles than the ^{26}Al events simulated at 1.8 MeV and skew the distribution to smaller values.

Having identified a starting minimum, maximum $\phi = (0^\circ, 60^\circ)$ in Figure E.1, we next examine the impact of changing the minimum and maximum allowable values of ϕ on the signal and background events. Changing the maximum value of ϕ and holding the minimum value at 0° is referred to below as the “step down” procedure. Changing the minimum value of ϕ and holding the maximum value at 60° is referred to as the “step up” procedure. Changing both is referred to as “simultaneous” optimization.

E.1 “Step down”

Figure E.2a shows the percentage of events removed by cuts in ϕ that keep only events that Compton scatter at an angle below the maximum ϕ on the x -axis. The minimum ϕ is 0° and the maximum ϕ is stepped down in 1° increments. Note that the percentage of events is plotted against the center of each 1° bin in ϕ and is calculated with respect to the total number of events in each data set, not to the sum of signal and background events. For example, the red curve demonstrates that 69.7% of all background events are removed from the background data set when the maximum allowed angle is $\phi = 11^\circ$.

The percentage of total background events removed plummets between approximately 6° and 15° and decreases at a slowing rate for maximum $\phi > 15^\circ$. In other words, increasing the maximum allowed ϕ beyond 15° does not remove significantly more background events than does a smaller maximum. Thus, most background events have Compton scattering




Figure E.1: Compton scattering angles of (a) simulated ^{26}Al (9- and 10-detector portions of the COSI 2016 flight) and atmospheric background events with incident energies 1803–1817 keV, and (b) real flight data (1803–1817 keV; 10 detectors). The majority of ^{26}Al and background events have $\phi < 60^\circ$. As such, the maximum Compton scattering angle ϕ considered in this optimization procedure is 60° .

angles between approximately 6° and 15° . This conclusion corroborates the dominance of background events shown in the background distribution between 6° and 15° in Figure E.1a.

On the other hand, the percentage of total signal events removed beyond a maximum allowed ϕ begins to plummet at a greater maximum ϕ than does the percentage of background events, at approximately 10.5° , and continues to drop fairly rapidly as maximum ϕ increases beyond that point. The rapid drop around 10.5° suggests that more signal events have Compton scattering angles closer to 10.5° than 6° , as was the case with the background simulation. These trends are also corroborated by the distributions plotted in Figure E.1a.

These conclusions initially suggest a beneficial minimum $\phi \sim 10^\circ$, since imposing this minimum would reject the background events between 6° and 10° and reject minimal signal, the presence of which becomes significant starting at around 10° . The goal is to minimize background with respect to signal without removing a detrimental number of potential signal photons.

Because the difference between the 9-detector ^{26}Al and the 10-detector ^{26}Al simulations is small, we proceed with only the 9-detector configuration to match the geometry of the 9-detector background simulation. Figure E.3 conveys the results in Figure E.2a in another way and underscores the competition between background minimization and costly rejection of signal events. Colors closer to yellow are desirable as they indicate the removal of more background than signal events. By that metric alone, a maximum allowed ϕ of approximately 6° is ideal. Indeed this value follows the previous conclusion from Figure E.2a

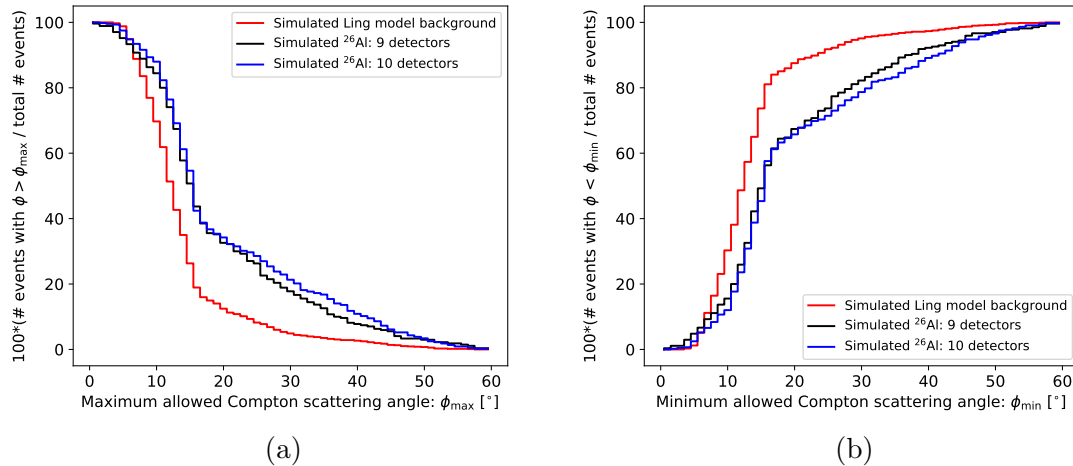


Figure E.2: (a) “Step down” optimization procedure. Percentage of simulated ^{26}Al and background events rejected by a Compton scattering angle cut which permits ϕ between a minimum $\phi_{\min} = 0^\circ$ and the maximum ϕ_{\max} on the x -axis. (b) “Step up” optimization procedure. Percentage of simulated ^{26}Al and background events rejected by a Compton scattering angle cut which permits ϕ between the minimum ϕ_{\min} on the x -axis and a maximum $\phi_{\max} = 60^\circ$.

that background events dominate the signal events near this value. However, because 93.3% of all signal events are removed by allowing only events with ϕ less than 6° , it is not an ideal maximum ϕ . This figure suggests that maximum ϕ values $\gtrsim 40^\circ$ may be a desirable alternative.

E.2 “Step up”

The procedure of adjusting ϕ in 1° increments is repeated at the lower end of allowed ϕ values. Instead of “stepping down” cuts as previously described ($(0^\circ, 60^\circ)$, $(0^\circ, 59^\circ)$, $(0^\circ, 58^\circ)$, etc.), we “step up” the cuts from $\phi_{\min} = 0^\circ$ to explore the impact of cuts at lower ϕ values ($(0^\circ, 60^\circ)$, $(1^\circ, 60^\circ)$, $(2^\circ, 60^\circ)$, etc).

Figure E.2b shows the percentage of events removed by cuts in ϕ that keep only events which Compton scatter at an angle above the minimum ϕ on the x -axis. The maximum ϕ is 60° and the minimum ϕ is stepped up in 1° increments. For example, the red curve demonstrates that 38.1% of all background events are removed from the background data set when the minimum allowed $\phi = 11^\circ$.

The percentage of total background events below the minimum allowed ϕ increases rapidly between approximately 6° and 15° and increases more slowly beyond there. Thus, as con-

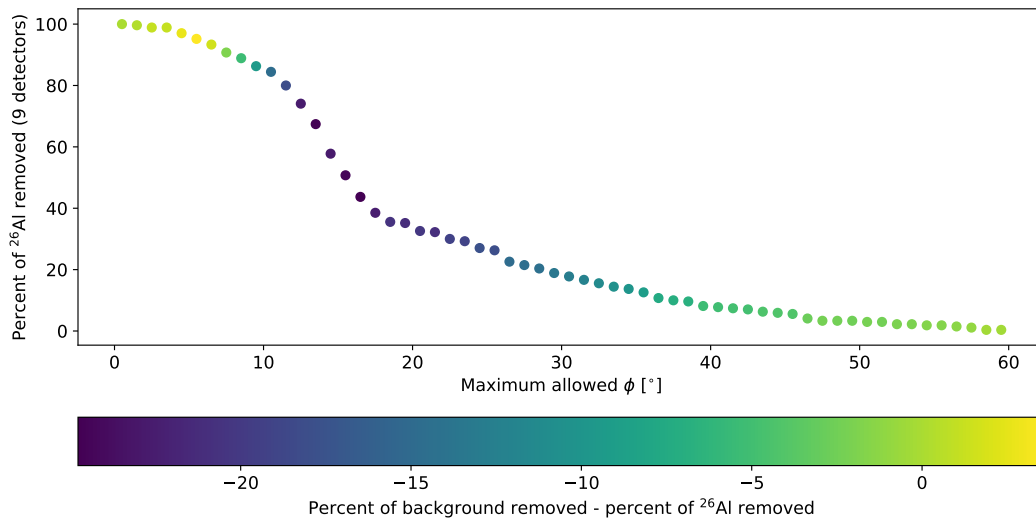


Figure E.3: The percentage of signal events removed as a function of maximum Compton scattering angle ϕ . It is necessary to balance the preferred removal of more background than signal with the associated consequence of removing too many raw signal events. A maximum ϕ of 6° initially appears to remove the most background and least signal, but this cut removes over 90% of all signal events. A maximum cut $\gtrsim 40^\circ$ may be a worthy compromise.

cluded from Figures E.1a and E.2a, most background events lie between 6° and 15° and including background events with $\phi > 15^\circ - 20^\circ$ does not add significant background to the total number of events.

The percentage of total ^{26}Al events below the minimum allowed ϕ begins to increase around 10.5° and continues to increase fairly rapidly as minimum ϕ increases. This implies that most signal events lie between about 10.5° and beyond, with a slightly greater spread towards larger values of ϕ than the background.

Again, these plots suggest that imposing a minimum ϕ near 10° could be beneficial, since doing so would reject the background in between 6° and 10° before the percentage of total signal counts becomes more significant near 10° . Because the percentage of ^{26}Al events removed does not plateau as quickly as the background events when increasing minimum ϕ , we recognize a more steady contribution of ^{26}Al events at greater values of ϕ and do not yet place constraints on maximum ϕ .

E.3 Simultaneous optimization

We retain the greatest number of ^{26}Al events with $\phi_{\text{max}} = 60^\circ$ because it allows the broadest possible range of ϕ values, and therefore the greatest number of overall events. However, this high maximum ϕ effectively broadens the signal region to subtend a significant fraction of the sky, limiting the size of the background region and the number of background events with which to develop a robust background model. A well-determined background is important for minimizing uncertainties in the search for ^{26}Al in flight data.

For a more complete visualization of the impact of ϕ cuts on the ^{26}Al and background simulations, we probe every acceptable range of ϕ by “stepping down” ϕ_{max} and “stepping up” ϕ_{min} simultaneously. Figure E.4 shows the percentage of events with ϕ between the minimum and maximum values on the x - and y -axes, respectively. The loosest cut of 0° to 60° retains the most events, as expected by nature of its wide range. The background events Compton scatter at smaller angles than the ^{26}Al events.

To estimate signal-to-background significance, we scale the raw numbers of events used to calculate the percentages in Figure E.4 to match the background and ^{26}Al simulations in flux. The ^{26}Al counts are multiplied by a factor of 1.6 because the full-sky DIRBE $240\ \mu\text{m}$ map flux used in the simulations is $1.1 \times 10^{-3}\ \text{ph cm}^{-2}\ \text{s}^{-1}$ and the most recent value from the literature is $1.7 - 1.8 \times 10^{-3}\ \text{ph cm}^{-2}\ \text{s}^{-1}$. Then, an estimated significance of signal $S / \sqrt{\text{background } B}$ is calculated for every possible ϕ cut.

The significance estimates are shown in Figure E.5 and illustrate the initial conclusions of Figure E.2 more simply. The maximum significance is ~ 2.6 for $\phi_{\text{min}} = 12^\circ$ and $\phi_{\text{max}} = 60^\circ$. Figure E.2 indicates a good candidate minimum value near 10° and as explained previously, the maximum of 60° is always preferred because it permits the most overall ^{26}Al events. Note that even though the significance plot in Figure E.5 grants effectively the same conclusion as Figure E.2, Figure E.2 helps demonstrate why the ideal ϕ_{min} is greater than 0° : setting the $\phi_{\text{min}} \sim 10^\circ$ rejects the domain of approximately 6° to 10° where the fraction of background events dominates that of ^{26}Al events.

To identify the optimal ϕ cut, we consider the conclusions from the “step up,” “step down,” and significance plots in concert with the desire to preserve a large portion of the sky for a well-determined background estimation. We choose to allow events with $\phi \in [10^\circ, 35^\circ]$. The minimum of 10° removes the background-dominated range of 6 – 10° and accepts more signal events than $\phi_{\text{min}} = 12^\circ$. The maximum of 35° preserves a fraction of the sky large enough for a well-determined background and has an acceptable signal-to-background significance. With a better description of the instrumental background, the maximum Compton scattering angle could be relaxed to its optimal value and the estimated significance of the ^{26}Al events would increase by $\sim 20\%$. In Appendix F, we attempt to build a more comprehensive background model with both instrumental and atmospheric components. Although the continuum shape of the atmospheric background can be matched to some extent, the instrumental lines in this energy range are difficult to model precisely.




Figure E.4: Simultaneous optimization procedure. The color scale indicates the percentage of events (left: atmospheric background, right: ^{26}Al signal) which have ϕ between the minimum (x -axis) and maximum (y -axis) limits. The broadest cut of 0° to 60° accepts the most events. The enhanced presence of background events towards smaller scattering angles than in ^{26}Al events is evidence of increased forward scattering in the former.




Figure E.5: Estimated significance (signal / $\sqrt{\text{background}}$) as a function of cuts in ϕ defined by the minimum and maximum values indicated on the axes. The greatest $S/\sqrt{B} \sim 2.6$ is achieved at $\phi_{\min} \sim 12^\circ$ and $\phi_{\max} = 60^\circ$. We choose $\phi_{\min} = 10^\circ$ to accept more overall signal events than at $\phi_{\min} = 12^\circ$. To preserve sufficient background statistics, we choose $\phi_{\max} = 35^\circ$.

Appendix F

Activation and atmospheric background simulations

This appendix¹ describes the activation and atmospheric (Ling model) simulations generated for the ^{26}Al analysis in Chapter 5. A subsequent discussion details the effects of applying various event selections to the simulations.

F.1 Activation background

It is important to simulate γ -rays from activation in order to understand the presence of instrumental background in data and simulations. Activation simulations of various cosmic ray and atmospheric particles are performed in MEGALib in three steps. The dominant contributors are protons p , neutrons n , and α -particles; emission from other particles, including muons, electrons, and positrons, was found to constitute a much smaller fraction of the background ($\sim 0.1\%$) in previous activation simulations (Kierans 2018). The first step (1) simulates the prompt emission (emission from excitations that decay on a timescale less than the detector timing resolution of $5\ \mu\text{s}$) from initial particles generated in the bombardment. All isotopes produced in step (1) are stored and fed to step (2) of the simulations, which calculates the activation of each isotope after a specified irradiation time. The final step (3) yields the delayed emission from the decays and de-excitations of extended irradiation encoded in step (2).

Step (1) of each particle type was performed by Kierans (2018). For the purposes of this study, we examine activation halfway through the COSI 2016 flight with a simulated irradiation time of 23 days in step (2). Step (3) is then run for 46 days to approximate the activation background over the entire COSI 2016 flight. We focus on activation lines between 1750–1850 keV to model background photopeaks near the signature ^{26}Al emission

¹Sections F.1 and F.2 of this appendix largely follow Appendix C of the published paper “Measurement of ^{26}Al by the Compton Spectrometer and Imager” by Beechert et al. (2022b). ©Astrophysical Journal 2022, reprinted with permission.




Figure F.1: Simulated spectra of delayed emission from instrumental activation due to protons, neutrons, and α -particles. The summed contribution of all components is shown in red. All Compton events between 1750–1850 keV with Compton scattering angle between 0 – 90° are included.

at 1809 keV. The simulations are conducted with a 12-detector mass model to account for all material in the COSI instrument.

Spectra of the step (3) delayed emission from each of the dominant particles are shown in Figure F.1. The events are subject to minimal restrictions: we select Compton events from all times between 1750–1850 keV with Compton scattering angles from 0 – 90° , no minimum distance between subsequent interactions, no Earth Horizon cut, and no pointing cut on the sky. Additional cuts are used in the ^{26}Al analysis to further restrict the events in this “initial” data set to, for example, the signal and background regions (Section 5.1).

Protons constitute the significant majority of simulated instrumental background in the COSI 2016 flight. The general shape of the activation spectra largely follows that of the background region flight data with minimal event selections (Figure 5.7). The peaks at ~ 1779 keV and ~ 1809 keV are easily identifiable and their likely origins are documented in the literature as captures on ^{27}Al (see Section 5.3.2). The total count rates of both peaks, summed over particle type, are $\sim 3.0 \times 10^{-3}$ cnts $\text{keV}^{-1} \text{ s}^{-1}$ and $\sim 2.1 \times 10^{-3}$ cnts $\text{keV}^{-1} \text{ s}^{-1}$, comparable to those seen in Figure 5.7 within an order of magnitude.

The simulated instrumental activation spectra do not contain the ~ 1764 keV peak seen in real flight background data (Figure 5.7). Its absence from the simulation might be expected because the literature widely attributes this line to the natural decay of ^{238}U in instrument materials, rather than a signature of de-excitation after activation. Uncertainty in the true origin of the line in real data motivates the empirical description of the flight background in

Chapter 5, which accounts for this line regardless of origin.

F.2 Atmospheric background

Rejecting atmospheric background from the glow of Earth albedo radiation is essential to measuring astrophysical γ -rays on balloon-borne platforms. Many instruments, including COSI, use anti-coincidence shielding to reject events emanating from below the instrument. COSI analyses also often choose to reject events incident from greater than 90° beyond the instrument's zenith (the "Earth Horizon Cut"). However, these methods do not guarantee complete background rejection. It is necessary to model the contamination of flight data by atmospheric background.

The atmospheric γ -ray background model by Ling (1975) describes the 0.3–10 MeV energy range at geomagnetic latitude $\lambda = 40^\circ$. It derives an isotropic, semi-empirical source function which models the production of γ -ray continuum and lines per unit air mass. The continuum is dominated by bremsstrahlung of primary and secondary cosmic ray electrons, neutral pion decays, and the scattering of incident photons to lower energies. The primary discrete contribution is from positron-electron annihilation at 511 keV; other lines from particle excitation and decay are also possible. The intensity of photons with incident energy E' and incident angle θ (measured from zenith) seen by a detector at atmospheric depth h [g cm^{-2}], as measured from the top of the atmosphere, is given by

$$\frac{dF(E', h)}{d\Omega} = \left(\int_r S(E', x) \rho(x) \exp\left[-\int_0^r \mu(E') \rho(r) dr\right] \frac{dr}{4\pi} + \frac{dF_c(E')}{d\Omega} \exp\left[-\int_0^\infty \mu(E') \rho(r) dr\right] \right) \text{ph cm}^{-2} \text{s}^{-1} \text{sr}^{-1} \text{MeV}^{-1}, \quad (\text{F.1})$$

where $\rho(x)$ is the air density for depth x and $\mu(E')$ is the mass absorption coefficient. Equation F.1 is reproduced from Equation 1 of Ling (1975); refer to the paper for a complete explanation.

As a quantitative estimate of the presence of atmospheric background in the ^{26}Al band-pass, and a demonstration of the anticoincidence shields' effectiveness in reducing this background, we consider the flux above 1750 keV predicted by the Ling model. The atmospheric background spectrum behaves as a power law ($\sim E^{-2}$) between approximately 0.1 and 10 MeV (Figure F.2). These photons can be detected and subsequently vetoed by the shields and would otherwise hit the detectors directly.

Given that COSI's CsI shields have an energy threshold of ~ 80 keV, integrating over the total shield flux and that above 1750 keV yields an estimate of the reduction in counts in




Figure F.2: Strength of the source function as a function of energy in three models (Peterson et al. 1972; Ling 1975; Peterson et al. 1973). Image from Ling (1975).

the ^{26}Al energy bandpass:

$$\begin{aligned}
 A &= \int_{80 \text{ keV}}^{\infty} E^{-2} = 12.5 \\
 B &= \int_{1750 \text{ keV}}^{\infty} E^{-2} = 0.57
 \end{aligned}
 \tag{F.2}$$

We find $B/A \sim 0.046$. Thus, $\sim 4.6\%$ of the total atmospheric background, roughly estimated, can be detected by the shields at energies greater than 1750 keV. By hitting the shields, these events are vetoed and the count rate associated with the atmospheric background is reduced by approximately 1–2 orders of magnitude.

F.2.1 Atmospheric quantities for the Ling model

While Ling (1975) provides expressions for the source functions $S(E', x)$ for both the continuum and line contributions, in this work we adopt a description of air density and mass absorption coefficient $\mu(E')$ given by Picone et al. (2002). We choose one day of the 2016 flight to represent the atmospheric conditions over the entire flight, as a longer background model simulation would be too computationally intense. Given that the focus of the analysis

is ^{26}Al from the Inner Galaxy, a day with maximum exposure of the Galactic Plane, corresponding to negative Earth latitudes, is chosen. The following flight conditions on May 22, 2016 00:00:00 UTC are fed to the NRLMSISE-00 atmospheric model (NRLMSISE): flight altitude = 33.6 km, latitude = -56.2° , longitude = 161° . The model returns the densities of atmospheric atomic and molecular oxygen and nitrogen, as well as helium, argon, and hydrogen in units of cm^{-3} , the total mass density in g cm^{-3} , and the atmospheric temperature in Kelvin for heights of 0–100 km.

F.2.2 Generating an orientation file for the Ling model

The atmospheric background simulation runs in MEGALib, using the above atmospheric quantities and an orientation file as inputs. The balloon orientation is required in the simulation to point to the coordinates which trace the 2016 flight. Five quantities define the orientation: time, the longitude and latitude of COSI’s x -axis, and the longitude and latitude of COSI’s z -axis. Here, the z -axis defines the instrument’s optical axis (zenith = 0), and the x -axis defines its azimuthal rotation.

As mentioned previously, atmospheric simulations are very computationally expensive: a simulation of the atmospheric background in one day of the balloon flight can take 24 hours to run. A comprehensive simulation of the 46-day flight is impractical. Instead, we develop a method of condensing all of COSI’s pointings over the whole flight to one day (86400 s). The entire sky is first divided into a grid of Galactic latitude and longitude. The number of pointings that fall into each pixel of the grid, namely the number of entries in the orientation file with z -axis angles within those spanned by each pixel, is plotted as an exposure map. The first entry of the flight orientation file found in each pixel is chosen as a representative coordinate (defined by all four angles) for that pixel. Then, the exposure time in each pixel is defined as the product of 86400 s and the fraction of total flight pointings found in that pixel. A new orientation file is then written using the representative coordinate and the pointing-weighted exposure time of each pixel. Hence, an orientation file containing all flight pointings is condensed to one day such that each orientation spans a time proportional to the number of pointings at that orientation in the 2016 flight.

We divide the sky into 100,000 FISBEL pixels of size $\sim 0.64 \text{ deg}^2$ (Figure F.3a). Note that this binning is much smaller than COSI’s $\sim 5^\circ$ (1650 FISBEL pixels; Figure F.3b) angular resolution. Sampling the original flight orientation file more frequently to determine which z -axis pointings fall inside each $\sim 0.64 \text{ deg}^2$ pixel limits the variation in x which would otherwise be present in a larger pixel containing more z -axis pointings. To see this effect, consider the x - and z -coordinates shown in Figure F.4. When the sky is binned coarsely into 100 FISBEL pixels and we consider only orientations in one of these large sky pixels (size $\sim 20.3 \text{ deg}^2$), the change in x -coordinate is much more dramatic compared to that in z . COSI’s z -coordinate is stable because the instrument is always oriented upwards and only changes as the instrument drifts with the balloon. Stability in the x -coordinate, on the other hand, is governed by a rotator that keeps the instrument’s solar panels pointed towards the Sun during the daytime. At nighttime, the rotator does not maintain this orientation and the instrument is free to




Figure F.3: Representative exposure map of the 2016 flight, scaled down to one day, generated by dividing the sky into (a) 100,000 pixels of size $\sim 0.64 \text{ deg}^2$ and plotting the percentage of total pointings in each. The horizontal band of blue points across the top of the figure is a plotting artifact. (b) The same effective exposure map with 1650 pixels of size 5 deg^2 , included here for visual clarity.

rotate azimuthally. The dark blue bulge in Figure F.4 shows the relative azimuthal stability from the rotator during daytime and the broader band across longitudes shows the times at night when the instrument freely rotated. These are the variations in x which would be ignored by binning too coarsely in z . Overall, we see that unless the sky binning is sufficiently fine, there are many degenerate azimuthal orientations that map to each zenith pointing. Dividing the sky into 100,000 pixels implies 100,000 pointings over 86400 s, or about 1.2 pointings per second. This is a sufficiently fine probe of instrument orientation, as the azimuthal rotation of the instrument does not change dramatically per second.

Given this one-day, pointing-weighted orientation file and the NRLMSISE-00 atmospheric conditions from May 22, 2016, we run the Ling model simulation and process it using 10- and 9-detector mass models. Concatenating the 10- and 9-detector portions thus yields a representation of atmospheric background over the COSI 2016 flight.

F.3 Event selections

A study of event selections on the activation and Ling model atmospheric simulations is performed to understand how the event selections may affect COSI's background.

First, we determine the fraction of activation and atmospheric events that are characterized as Compton events and as such, contribute to COSI data analysis. Single-site events (photoelectric absorption) are typically excluded. Figure F.5 shows that on average, 69.7% of all simulated events are Compton events and 21.8% are single-site events. The split between Compton and single-site events is consistent across event types.




Figure F.4: The change in azimuthal (x , blue) and zenith (z , red) orientations during the COSI 2016 flight for one arbitrary pixel of size $\sim 20.3 \text{ deg}^2$. COSI always points upwards during flight; the zenith pointing is relatively stable and changes only with the drift of the balloon. During the daytime, COSI’s rotator orients the solar panels towards the Sun and stabilizes the azimuthal orientation (central blue bulge) compared to that at nighttime, when the instrument is free to rotate azimuthally (band of blue points across longitude).

The effects of additional selections on the Compton events are examined. Event selections are applied to simulations of each activation particle in the 12-, 10-, and 9-detector configurations. Similarly, event selections are applied to the atmospheric background in the 10- and 9-detector configurations. These event selections are applied individually, meaning that the percentage shown on a given point is of the events which pass or fail only the event selection in question and not a combination of that event selection with another. The results from these mass model-specific tests are found in Figures F.6 and F.7. More events fail the 33 km minimum altitude cut in the 9-detector configuration than in the 10-detector configuration. This is expected because the balloon experienced altitude drops in the 9-detector part of the 2016 flight. The effect of altitude on the atmospheric simulation is not examined because the simulation is performed at 33.6 km and timing cuts excluding or including certain altitudes are not relevant. Overall, within each particle type, the effect of event selections is largely independent of the mass model used.

Given the minimal influence of the mass model, we summarize the study more concisely in Figure F.8, which shows the 12-detector activation and 10-detector atmospheric numbers from Figures F.6 and F.7 on one plot. Below is a brief explanation of the event selections in these three figures:






Figure F.5: The fraction of Compton events and single-site events present in each of the simulated background data sets (activation proton, neutron, and α -particles: 12-detector mass model; atmospheric: 10-detector mass model). The events not characterized as Compton or single-site events are unable to be reconstructed. The mean fraction of Compton events is 69.7% and that of single-site events is 21.8%.

- The Earth Horizon Cut (EHC), which is used in nearly all COSI analyses to mitigate background contamination, removes approximately 84% of activation and atmospheric events.
- The “Flight GTI” selection indicates the “good time intervals” during the 2016 flight. Applying this selection excludes data taken before the balloon achieved float altitude and excludes data taken during system malfunctions or high shield rates.
- The “33 km GTI” applies an additional cut to the “Flight GTI” that excludes all events collected below 33 km float altitude.
- The Compton scattering angle cut removes events outside of the $10^\circ - 35^\circ$ range optimized for ^{26}Al analysis (Appendix E). It eliminates approximately 60% of events.
- The minimum distance between the first and second interactions of at least 0.5 cm removes far fewer events.
- The signal region (SR) and three background region (BR) pointing cuts are as listed in Table 5.1.




(a) α -particles




(b) Neutrons

Figure F.6: Individual event selections applied to (a) α -particles and (b) neutrons in the activation simulations. Results with 12-, 10-, and 9-detector instrument configurations are shown. Refer to Section F.3 for an explanation of each event selection.



(a) Protons



(b) Atmospheric events

Figure F.7: Individual event selections applied to (a) protons in the activation simulations and (b) atmospheric photons in the Ling model simulation. Refer to Section F.3 for an explanation of each event selection.

- The three rightmost event selections in Figure F.8 demonstrate the severe cut to Compton events imposed by ^{26}Al event selections, which are a combination of the individual event selections listed to their left on the x -axis. The “Pass ^{26}Al (no pointing)” selection accepts all Compton events which pass the 33 km minimum altitude cut, have energy 1750–1850 keV, pass the allowed Compton scattering angle range of $10^\circ - 35^\circ$, pass the 0.5–30 cm distance cut, and pass the EHC. The “Pass ^{26}Al (SR)” cut adds the signal region pointing cut to this list of requirements. The “Pass ^{26}Al (BR)” is identical to the “Pass ^{26}Al (SR)” except that it requires passage of the Flight GTI instead of the 33 km GTI and passage of one of the three background region pointing cuts rather than the signal region pointing cut. The atmospheric points in these three rightmost fields consider all event times rather than specific GTIs since the atmospheric simulation is conducted at constant altitude.

It is clear that the event selections used in the ^{26}Al analysis, namely the EHC, 1750–1850 keV energy cut, and signal and background region pointing cuts, remove over 99% of the background Compton events simulated for this analysis.




Figure F.8: Percentage of simulated activation and atmospheric Compton events which pass or fail the event selections listed on the x -axis. These are the 12 (10)-detector numbers for the activation (atmospheric) simulations in Figures F.6 and F.7. The results are consistent across types of background. Refer to Section F.3 for an explanation of each event selection.

Next, we cumulatively apply the event selections described above to the 12-detector α -particle simulation (chosen as a representative case given the similarity between activation

particles) and the 10-detector atmospheric simulation. Figure F.9 shows the results for both simulations. Read from left to right, the event selections on the x -axis are applied in succession to events which already passed the nominal (minimal) selections listed in the figure caption. Thus, the percentages at each point give the percentage of Compton events which pass the cut listed at that point and all cuts listed to the left of that point. An exception to this scheme is the rightmost “Pass BR pointing” cut, which employs all of the cuts to its left except for the 33 km GTI. The Flight GTI is used instead. Another exception is the exclusion of timing considerations in the atmospheric simulations. Overall, the activation and atmospheric results are quite similar. The simple EHC is an effective method of background rejection, removing over 80% of background Compton events. The negligible percentage of events remaining after the ^{26}Al event selections underscores both the power of these event selections and the difficulty of the ^{26}Al analysis at its core; in attempting to remove background, the large majority of COSI’s entire data set is eliminated from the analysis.




Figure F.9: Application of event selections in succession to the 12-detector α -particle simulation (3016057 total simulated Compton events) and 10-detector atmospheric simulation (3058353 total simulated Compton events), both initially limited by the following nominal event selections: only Compton events of energy 0–2000 keV, Compton scattering angle $\phi \in 0 - 90^\circ$, and distance between first two interactions 0–30 cm are considered. No EHC and no pointing cuts are initially applied. Displayed is the percentage of these events which pass the cut listed on the x -axis and all cuts listed to the left. Note that in the α -particle simulation, the background region pointing cut uses the “Flight GTI” event selection rather than the “33 km GTI” selection. There is no timing consideration in the atmospheric simulation.