

A Fast GRB Source Localization Pipeline for the Advanced Particle-astrophysics Telescope

Marion Sudvarg,^{a,*} Jeremy Buhler,^a James Buckley,^a Wenlei Chen,^{a,b} Zachary Hughes,^a Emily Ramey,^{a,c} Michael Cherry,^d Samer Alnussirat,^d Ryan Larm^a and Christofer Berruz Chungata^{a,e} on behalf of the ADAPT Collaboration
(a complete list of authors can be found at the end of the proceedings)

^aWashington University in St. Louis, St. Louis, MO, USA

^bUniversity of Minnesota, Minneapolis, MN, USA

^cUniversity of California, Berkeley, CA, USA

^dLouisiana State University, Baton Rouge, LA, USA

^eUniversity of Bridgeport, Bridgeport, CT, USA

E-mail: msudvarg@wustl.edu, jbuhler@wustl.edu, buckley@wustl.edu,
chen6339@umn.edu, zd Hughes@wustl.edu, emily_ramey@berkeley.edu,
cherry@lsu.edu, alnussirat1@lsu.edu, l.ryan@wustl.edu,
cberruzc@my.bridgeport.edu

We present a pipeline for fast GRB source localization for the Advanced Particle-astrophysics Telescope. APT records multiple Compton scatterings of incoming photons across 20 CsI detector layers, from which we infer the incident angle of each photon's first scattering to localize its source direction to a circle centered on the vector formed by its first two scatterings. Circles from multiple photons are then intersected to identify their common source direction. Our pipeline, which runs in real time on low-power hardware, uses an efficient tree search to determine the most likely ordering of scatterings for each photon (which cannot be measured due to the coarse time-scale of detection), followed by likelihood-weighted averaging and iterative least-squares refinement to combine all circles into an estimated source direction. Uncertainties in the scattering locations and energy deposits require that our pipeline be robust to high levels of noise.

To test our methods, we reconstructed GRB events produced by a Geant4 [1] simulation of APT's detectors paired with a second simulator that models measurement noise induced by the detector hardware. Our methods proved robust against noise and the effects of pair production, producing sub-degree localization for GRBs with fluence 0.3 MeV/cm². GRBs with fluence 0.03 MeV/cm² provided fewer photons for analysis but could still be localized within 2.5 degrees 68% of the time. Localization time for a 1-second 1.0 MeV/cm² GRB, measured on a quad-core, 1.4 GHz ARMv8 processor (Raspberry Pi 3B+), was consistently under 0.2 seconds — fast enough to permit real-time redirection of other instruments for follow-up observations.

37th International Cosmic Ray Conference (ICRC 2021)

July 12th – 23rd, 2021

Online – Berlin, Germany

*Presenter

1. Introduction

The Advanced Particle-astrophysics Telescope (APT) [2] is a space-based observatory, currently in development, to survey the entire sky for gamma-ray sources in the MeV to \sim TeV range. APT's goals include prompt detection of energetic transient events in the distant universe, such as gamma-ray bursts (GRBs), and rapid communication of these events to narrow-band instruments that can conduct follow-up observations in other spectral bands. Pursuant to this goal, we are developing analytical methods to perform real-time detection and localization of events, which will run on computing hardware on the orbiting APT platform.

In this work, we focus on detecting events for which most incident photons have energies in the Compton regime (511 keV to a few MeV). For such events, the photons of interest Compton-scatter one or more times within the detector, depositing energy with each scattering, until they are eventually photoabsorbed. All scatterings for one photon appear simultaneous at the time resolution of the detector. Our analytical tasks are twofold: first, to identify the first two scatterings for each photon in order of occurrence, which localizes the source's direction to an annulus centered on the line connecting these scatterings; and second, to combine the annuli from all detected photons to estimate the most likely direction in the sky for a distant point source (the GRB) that emitted them.

Several computational challenges arise in building a robust pipeline of algorithms for event detection and localization. The pipeline must accurately localize even low-fluence events (at most a few thousand total incident photons) while being efficient enough to keep up with high-fluence events (around 10^5 photons/sec) involving many photons that might scatter 5 or more times within the detector. It must yield results quickly enough to permit rapid retargeting of narrow-band instruments to the event – ideally in well under a second. And it must deliver this performance using a low-power processor of the type feasible for a space-based platform. All of these criteria must be met while also dealing robustly with the measurement limitations of the APT detector.

This work describes a computational pipeline for Compton-regime reconstruction and event localization. We build on the basic approach of Boggs and Jean [3] to reconstruct photon trajectories within the instrument by minimizing disagreement between the reconstructed angle of each scattering and the energy it deposited. To eliminate redundant computation and ensure rapid analysis even of photons with multiple scatterings, we implement a tree search with pruning over possible photon trajectories to find one with the best agreement. Event localization from reconstructed photons then follows a maximum-likelihood approach, with random sampling of reconstructed photons to guess a plausible source direction, followed by iterative refinement. We demonstrate that our pipeline can reliably localize events to within 2.53 degrees (68% containment) for low-fluence events and 0.42 degrees for high-fluence events while delivering results within 200 ms even on a low-power ARM Cortex-A53 processor.

2. Background

Current simulations of the APT instrument report a sequence of N gamma-ray interactions as a collection of pairs (\mathbf{x}_i, E_i) , where \mathbf{x}_i is a 3-vector denoting the interaction's coordinates and E_i is the energy deposited by the photon during the interaction. Currently, we assume that each reported

set of interactions is caused by a single incident gamma-ray photon, and that all interactions are Compton scatterings or a final photoabsorption.

Because the time resolution of the detector is not sufficiently fine to unambiguously order the photon interactions, our software reconstructs an ordering according to the approach of Boggs and Jean [3], which we now briefly review. For photons having three or more interactions with the telescope, we assume that there exists some true ordering of interactions such that the first $N - 1$ are Compton scatterings and the last is photoabsorption. For such an ordering, the photon's energy after interaction i , normalized to the electron mass, is defined as $W_i = \frac{1}{m_e c^2} \sum_{j=i+1}^N E_j$. The expected scattering angle ϕ'_i for interaction i is then given by the Compton formula: $\cos \phi'_i = 1 + \frac{1}{W_{i-1}} - \frac{1}{W_i}$.

To determine the true ordering, [3] considers all $N!$ possible orderings. For a given ordering, let \mathbf{r}_i be the vector pointing from interaction $i-1$ to i , with corresponding unit vector $\hat{\mathbf{r}}_i$. Then for the *triple* of consecutive interactions $(i-1, i, i+1)$, the measured scattering angle ϕ_i is given by $\cos \phi_i = \hat{\mathbf{r}}_i \cdot \hat{\mathbf{r}}_{i+1}$. Hereafter, we define $\eta_i = \cos \phi_i$ and $\eta'_i = \cos \phi'_i$. The following statistic measures the sum of disagreements between the measured angle of each scattering and the angle implied by its deposited energy:

$$\chi^2 = \sum_{i=2}^{N-1} \chi_i^2 = \frac{1}{N-2} \sum_{i=2}^{N-1} \frac{(\eta_i - \eta'_i)^2}{\delta\eta_i^2 + \delta\eta_i'^2}, \quad (1)$$

where $\delta\eta_i$ and $\delta\eta_i'$ are the estimated uncertainties in the measured and energy-derived angle cosines, respectively¹. For an event with N interactions, if $\eta_i - \eta'_i$ has an approximately Gaussian distribution, this statistic has a χ^2 distribution with $N - 2$ degrees of freedom.

Boggs and Jean compute χ^2 for every possible ordering of interactions and select the one that minimizes this statistic. From the inferred first and second interactions in order, an annulus containing the photon's initial direction can be derived using the Compton formula. The center of the annulus is defined by the vector $\mathbf{c} = -\hat{\mathbf{r}}_1$, while the angle ϕ_1 is defined by the Compton formula. (If the inferred ordering is such that $|\eta_i| > 1$, the event is discarded.) The effective half-width $\delta\phi_1$ of the annulus is described by the propagated uncertainty from interactions' coordinate and energy measurements [3, Equation 6].

3. Compton Reconstruction

Computing the χ^2 statistic of Equation 1 for all $N!$ possible orderings can be quite slow for high-fluence GRB events with energetic photons, which can have N of 6-10. To enable rapid reconstruction even in this extreme case, we implement a *tree traversal* over possible orderings. A node at depth k in the tree corresponds to a choice of the first k interactions in order; for all orderings that make this choice, the first $k - 2$ terms of χ^2 are the same and so need be computed only once. Overall, we need compute only one new term of the sum in Equation 1 per tree edge traversed. Total time to compute χ^2 for all orderings is thereby reduced by a factor of N versus naïve enumeration.

Organizing reconstruction as a tree search exposes further opportunities to reduce work through *pruning*. If at any node in the traversal, the sum of χ_i^2 values at this node exceeds the minimum χ^2

¹In this work, we assume that the value δE_i , which contributes to the calculation of $\delta\eta_i'$ in [3], is described by the function $\delta E = \sigma_E \sqrt{E}$. We estimated σ_E as 0.095 by fitting a Gaussian distribution to a histogram of simulated digitizer energy errors, normalized by the square-root of the true simulated energy.

computed at any leaf of the tree (i.e., for any complete ordering) so far, we may prune the tree at this node and omit any computations that would have been performed below it. To make pruning even more aggressive, we additionally check that the sum of χ_i^2 values at a node is not so large that its associated p -value (given the statistic's assumed χ^2 distribution, with $k - 2$ degrees of freedom at depth k in the tree) is implausible. We precompute χ^2 thresholds for each level of the tree corresponding to $p = 0.001$ and prune at a node if its sum exceeds the threshold for its depth. Finally, we prune at any node where the computed value η'_i significantly exceeds 1, which implies an energetically impossible reconstruction even given measurement uncertainty.

To further reduce computation cost, we first precompute η_i and $\delta\eta_i^2$ for all $P(N, 3)$ possible triples of consecutive interactions. These values are constant for a triple no matter where it occurs in the ordering. For a fixed triple $(i-1, i, i+1)$, η_i is computed from its coordinates, while $\delta\eta_i^2$ is calculated by propagation of uncertainty assuming a common value Δ for the three coordinate uncertainties² δx , δy , and δz :

$$\delta\eta_i^2 = 2\Delta^2 (1 - \eta_i^2) \left(\frac{1}{|\mathbf{r}_i|^2} + \frac{\eta_i}{|\mathbf{r}_i||\mathbf{r}_{i+1}|} + \frac{1}{|\mathbf{r}_{i+1}|^2} \right). \quad (2)$$

For $N \geq 3$, tree traversal yields some ordering with minimum χ^2 value, unless all possible orderings were pruned due to implausibility (in which case the photon is discarded). For this ordering, we propagate the annulus $(\mathbf{c}, \phi, \sigma)$, where $\sigma = \delta\phi_1$ is the propagated uncertainty in ϕ , to the localization stage of our pipeline. For events with $N = 2$, we cannot distinguish between the two possible orderings; instead, we propagate two annuli (assuming neither have $|\eta_i| > 1$) corresponding to the both possible orderings and let localization select between them.

4. Localization

For each detected photon, we must combine information across photons to infer a common direction vector \mathbf{s} in the sky from which they originated. Given three correctly reconstructed photons and no instrument noise, it is possible in principle to derive \mathbf{s} by trilateration on the unit sphere. However, uncertainty in the measured positions and energies of a photon's interactions result in uncertainty regarding its angle of arrival relative to the first two interactions. Moreover, we cannot assume that each event has been correctly reconstructed; indeed, for events with two interactions, at least one of the two inferred annuli is wholly incorrect. Hence, some annuli will not contain \mathbf{s} at all. Fortunately, for GRBs of reasonable fluence, we expect to detect many more than three photons and so have redundant information with which to robustly estimate \mathbf{s} even in the presence of noise.

Given N annuli, each described by a triple $(\mathbf{c}_i, \phi_i, \sigma_i)$, we seek to estimate the true source direction \mathbf{s} . This is done in two steps: first, we sample multiple annuli, estimate a likely source direction contained within each one, and then average these source estimates to compute a rough approximation of the source direction. Second, we use an iterative least-squares refinement of this initial approximation, identifying and using only those photons most likely to have been correctly reconstructed, to derive a final estimate of \mathbf{s} .

²For the APT detector, we estimate Δ as 0.7 mm by fitting a Gaussian distribution to a histogram of the coordinate errors produced by our simulated digitizer, then taking the standard deviation of the fit.

4.1 Identifying an Approximate Source Direction

To form an initial estimate of the source direction, we begin by selecting a single annulus from the input at random. If it was generated by a correctly reconstructed photon, we may assume that the source direction \mathbf{s} lies near the circle with radius ϕ about center \mathbf{c} . We therefore test a set of evenly spaced candidate source directions \mathbf{s}_j , $1 \leq j \leq m$, that lie on this circle. We compute the likelihood of s_j with respect to the i th input annulus as

$$\mathcal{L}(\mathbf{s}_j | (\mathbf{c}_i, \phi_i)) \approx \frac{1}{\sigma_i \sqrt{2\pi}} e^{-(\phi_i - \beta_i)^2 / 2\sigma_i^2}, \quad (3)$$

where $\beta_i = \arccos(\mathbf{s}_j \cdot \mathbf{c}_i)$. Each photon is assumed to be an independent event given the source direction \mathbf{s}_j , so we compute the joint log-likelihood of \mathbf{s}_j given all input annuli by summing the individual log-likelihood contributions for each. The candidate direction \mathbf{s}_j with the greatest log-likelihood is selected as an approximate source direction.

The above approach assumes that the circle on which the \mathbf{s}_j 's lie was derived from a correctly reconstructed photon. In practice, we cannot be sure of this, so we repeat this process for q randomly sampled annuli, where q is large enough that we can expect to sample annuli from at least one, and preferably several, correctly reconstructed photons. Let $\mathbf{s}_1 \dots \mathbf{s}_q$ be the directions reconstructed from each of these sampled annuli. Then our combined estimate \mathbf{s}_0 averages these q estimates, weighting each by its estimated likelihood. This weighted average suppresses estimates arising from incorrectly reconstructed annuli and tends to cancel out errors due to the true source direction not lying precisely on the corresponding circles. The estimate \mathbf{s}_0 (normalized to a unit vector) is passed on to iterative refinement as described below.

We implemented several performance optimizations to ensure rapid computation of \mathbf{s}_0 even for high-fluence events with many detected photons. First, when the number N of annuli passed to localization is large, we sample $n < N$ of them and use this reduced set, rather than the full set of N , to compute \mathbf{s}_0 . Because \mathbf{s} is highly overdetermined, sampling 10^2 to 10^3 photons from an initial set of 10^4 to 10^5 is unlikely to seriously compromise accuracy. Moreover, we still use *all* input photons in the subsequent refinement step. Second, we found empirically that the most time-consuming operation in the computation of Equation 3 is the arccosine needed to obtain β_i . We therefore introduced a faster, lower-precision approximation to the arccosine that reduced its cost by an order of magnitude compared to the single-precision version in the GNU C Library on the target platform. Finally, we parallelize the computation across the q choices of initial annuli, distributing work across the multiple cores of our processor.

4.2 Iterative Refinement of Source Direction

We next refine our initial estimate \mathbf{s}_0 using a least-squares approach. Given a set of input circles (\mathbf{c}_i, ϕ_i) from correctly reconstructed photons, if ϕ_i is known exactly, then for all i , the source direction \mathbf{s} lies on the circle; hence, $\mathbf{c}_i \cdot \mathbf{s} = \cos \phi_i$. This constraint is linear in the coordinates of \mathbf{s} . In addition to the constraints imposed by each circle, we require that \mathbf{s} be a direction, and hence a unit vector; this unit-norm constraint is quadratic in the coordinates of \mathbf{s} .

In practice, each ϕ_i has uncertainty σ_i , so the above system of equations may have no consistent solution for \mathbf{s} . We therefore find a least-squares solution to the linear constraints, subject to

enforcement of the quadratic unit-norm constraint. The linear constraint for the i th circle is weighted according to $1/\sigma_i^2$ to reflect the uncertainty in the corresponding measurement.

A linear-least squares problem subject to a norm constraint can be reduced to a quadratic eigenvalue problem [4]. The matrix for which the eigenvalues must be found has dimension proportional to the number of unknowns, i.e., the three coordinates of \mathbf{s} , independent of the number of photons. Quadratic eigenvalue problems are computationally tractable for small matrices [5], and forming the matrix for the problem has cost quadratic in the number of input photons.

The quality of a least-squares solution is in general quite sensitive to outliers. We therefore iteratively refine our source estimate in an attempt to discard annuli of incorrectly reconstructed events. For each refinement iteration $k = 0, 1, 2, \dots$, we test each input annulus i in parallel to determine whether it passes close to the current source estimate \mathbf{s}_k . In particular, we test whether the angle $\arccos(\mathbf{c}_i \cdot \mathbf{s}_i)$ lies within three standard deviations of ϕ_i . All such annuli are used as input to the least-squares computation, which produces a refined estimate \mathbf{s}_{i+1} that is used in the next refinement iteration. This process is iterated for a fixed number of steps, which empirically sufficed to converge to a stable final solution.

5. Evaluation

We simulated a long GRB with a typical spectrum using APTsoft (described in [6]), generating 10^6 uniformly distributed gamma-ray photons from a normally incident, collimated beam with a cross-section of 18 m^2 to fully cover the APT detector. The spectral-energy distribution of the simulated GRB was modeled according to a Band function [7] with parameters $\alpha = 0.6$, $\beta = -2.5$, $E_{\text{peak}} = 1 \text{ MeV}$, and incident energies ranging from 300 keV to 10 MeV . In this energy regime, Compton scattering effects dominate, though pair production does occasionally occur. We do not attempt to distinguish pair-production events, so these events cannot be correctly reconstructed and effectively act like noise for our reconstruction pipeline.

5.1 Localization Accuracy

To test the accuracy of our pipeline, we first parameterized the localization algorithms as follows. For the approximation stage, we selected $m = 720$ candidate source directions separated by half a degree. We additionally selected $q = 20$ randomly sampled annuli over which to average approximations, which we found to be large enough to sufficiently suppress bad reconstructions and cancel out errors. For the refinement stage, we performed 20 iterations of least-squares refinement.

We measured localization accuracy across a range of fluences, running our pipeline 1000 times for each fluence, each time randomly sampling photons from the corpus of 10^6 events. Where a tested fluence propagated $N > 1000$ annuli, we randomly selected a subset of $n = 1000$ annuli to compute the initial estimate of source direction.

Fluence	Mean Error	Std Dev	68% Containment	95% Containment
0.03 MeV/cm ²	2.15	1.22	2.53	4.42
0.1 MeV/cm ²	1.21	0.64	1.45	2.32
0.3 MeV/cm ²	0.70	0.36	0.87	1.32
1.0 MeV/cm ²	0.35	0.20	0.42	0.72

Table 1: Localization error (in degrees) measured over 1000 trials.

Table 1 quantifies our approach’s localization accuracy. For the lowest-fluence events, where typically fewer than 500 photons are available for reconstruction, our method is still able to infer the true source direction of a GRB with typical accuracy better than 2.5 degrees. At higher fluences, where more than 5000 photons are available, we can achieve consistent sub-degree localization. For 1.0 MeV/cm² events, where over 15,000 photons are available, we achieve localization well under 1 degree 95% of the time.

5.2 Execution Times

To test the efficiency of our pipeline on a low-power embedded platform, we ran it on a Raspberry Pi 3B+, which has a Cortex-A53 (ARMv8) quad-core, 1.4 GHz, 64-bit CPU and 1 GB of LPDDR2 DRAM. We used the same simulated set of 10⁶ photons, this time running 200 trials each for fluences from 0.01 to 0.05 MeV/cm² in increments of 0.01, and from 0.1 to 1.0 MeV/cm² in increments of 0.05. We used OpenMP to parallelize reconstruction of different photons, as well as localization as described above, across cores. We separated the computation times for event reconstruction, approximation of an initial source direction, and iterative least-squares refinement, then aggregated these into total pipeline execution times. Results are given in Figure 1.

Reconstruction times scale approximately linearly with fluence; this is to be expected since each photon is reconstructed independently. Additionally, this stage contributes only a minor portion of the total pipeline execution time. For the initial source approximation stage, execution time increases as a function of fluence until the number of propagated annuli exceeds $N = 1000$. For higher fluences, we sample a constant-size subset of $n = 1000 < N$ annuli for this pipeline stage, allowing times to remain approximately constant. For the iterative least-squares refinement stage, running times increase quadratically but are dominated by the cost of approximation in the fluence range considered.

Our measurements indicate that, even for high-fluence events, total pipeline execution times are well under a second, and the variance in measured times for individual fluences is low. Across all fluences considered, the maximum pipeline time measured was 197.7 ms, well under the 1-second target for total processing time. This result demonstrates that our pipeline has sufficient throughput to perform source localization for high-fluence GRB events in real time.

6. Conclusions and Future Work

We have demonstrated a pipeline for fast GRB source localization for APT. Even at high fluences, execution times are fast enough to process events in real time. Additionally, the software can




Figure 1: Mean execution times over 200 trials for each fluence tested, highlighting contributions from each pipeline stage. Standard deviation of total time is under 4.2 ms.

consistently estimate a true GRB source direction within a degree for high fluence events. At lower fluences, which provide fewer photons for analysis, localization is still accurate to within 2.5 degrees 68% of the time.

We plan to improve our methodology to increase accuracy without significantly impacting execution time. In our initial approximation of source direction, we do not yet model the probability that a photon arriving from direction \mathbf{s}_j would produce observed annulus i . Such a probability could be computed using, e.g., the Klein-Nishina model [8] to place a prior on the photon's angle with \mathbf{c}_i given its energy. Prior weighting of annuli could be used to down-weight those likely to have arisen through noise or error.

Currently, neither reconstruction nor localization assumes a particular model of background photons not due to GRB events. Once such a background model is developed for APT, we will investigate its impact on accuracy. We anticipate that the rate of background photons will be low enough that its impact on running time will be limited.

Finally, we are developing approaches to improve the input to our analysis pipeline, including improved estimates of interaction coordinates that can separate multiple, simultaneous photon arrivals at the detector, detection and rerouting of pair-production events to a separate analysis pipeline, and improved estimates of interaction positions and their errors on the detector's z-axis.

Acknowledgments

This work was supported by NASA awards 80NSSC19K0625 and 20-APRA20-0148, NSF award CNS-1763503, Washington U.'s McDonnell Center for the Space Sciences, and the Peggy and Steve Fossett Foundation. We are grateful to Washington U. technical staff including Richard Bose, Dana Braun, and Garry Simburger for contributions to the prototype detectors and the Antarctic flight of the APTlite instrument.

References

- [1] S. Agostinelli, J. Allison, K. Amako, J. Apostolakis, H. Araujo, P. Arce et al., *Geant4 — a simulation toolkit*, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **506** (2003) 250.
- [2] J. Buckley, *The Advanced Particle-astrophysics Telescope (APT) Project Status*, in *Proceedings of 37th International Cosmic Ray Conference — PoS(ICRC2021)*, vol. 395, p. 655, 2021, DOI.
- [3] S. Boggs and P. Jean, *Event reconstruction in high resolution Compton telescopes*, *Astronomy and Astrophys. Supp. Series* **145** (2000) 311.
- [4] W. Gander, G.H. Golub and U.V. Matt, *A constrained eigenvalue problem*, *Linear Algebra Appl.* **114-115** (1989) 815.
- [5] F. Tisseur and K. Meerbergen, *The quadratic eigenvalue problem*, *SIAM Rev.* **43** (2001) 235.
- [6] W. Chen, J. Buckley, S. Alnussirat, C. Altomare, R. Bose, D. Braun et al., *The Advanced Particle-astrophysics Telescope: Simulation of the Instrument Performance for Gamma-Ray Detection*, in *Proceedings of 37th International Cosmic Ray Conference — PoS(ICRC2021)*, vol. 395, p. 590, 2021, DOI.
- [7] D. Band, J. Matteson, L. Ford, B. Schaefer, D. Palmer, B. Teegarden et al., *BATSE observations of gamma-ray burst spectra. I. spectral diversity*, *Astrophys. J.* **413** (1993) 281.
- [8] O. Klein and Y. Nishina, *Über die Streuung von Strahlung durch freie Elektronen nach der neuen relativistischen Quantendynamik von Dirac*, *Zeitschrift für Physik* **52** (1929) 853.

Full Authors List: ADAPT Collaboration

S. Alnussirat¹, C. Altomare², R. G. Bose³, D. Braun³, J. H. Buckley³, J. D. Buhler⁴, E. Burns¹, R. D. Chamberlain⁴, W. Chen⁵, M. L. Cherry¹, L. Di Venere^{6,2}, J. Dumonthier⁷, M. Errando³, S. Funk⁸, F. Giordano^{6,2}, J. Hoffman³, Z. Hughes³, D. J. Huth³, P. L. Kelly⁵, J. F. Krizmanic^{9,10,11}, M. Kuwahara¹², F. Licciulli², G. Liu¹³, M. N. Mazziotta², J. G. Mitchell^{14,15}, J. W. Mitchell⁹, G. A. de Nolfo¹⁵, R. Paoletti¹⁶, R. Pillera^{17,2}, B. F. Rauch³, D. Serini², G. Simburger³, M. Sudvarg⁴, G. Suarez⁷, T. Tatoli^{15,18}, G. S. Varner¹², E. Wulf¹⁹, A. Zink⁸, W. V. Zober³

¹Department of Physics and Astronomy, Louisiana State University, Baton Rouge, Louisiana 70803, USA, ²Istituto Nazionale di Fisica Nucleare, Sezione di Bari, I-70126 Bari, Italy, ³Department of Physics and McDonnell Center for the Space Sciences, Washington University, St. Louis, MO 63130, USA, ⁴Department of Computer Science & Engineering, Washington University, St. Louis, MO 63130, USA, ⁵Department of Physics and Astronomy, University of Minnesota, Minneapolis, MN 55455, USA, ⁶Dipartimento di Fisica “M. Merlin” dell’Università e del Politecnico di Bari, I-70126 Bari, Italy, ⁷NASA Goddard Space Flight Center, Greenbelt, MD 20771, USA, ⁸Friedrich-Alexander-Universität Erlangen-Nürnberg, Erlangen Centre for Astroparticle Physics, D-91058 Erlangen, Germany, ⁹Center for Space Sciences and Technology, University of Maryland Baltimore County, Baltimore, MD 21250, USA, ¹⁰Astroparticle Physics Laboratory, NASA/GSFC, Greenbelt, MD 20771, USA, ¹¹Center for Research and Exploration in Space Sciences and Technology, NASA/GSFC, Greenbelt, MD 20771, USA, ¹²Department of Engineering, University of Hawai’i at Mānoa, Honolulu, HI 96822, USA, ¹³Department of Physics and Astronomy, University of Hawai’i at Mānoa, Honolulu, HI 96822, USA, ¹⁴Department of Physics, The George Washington University, Washington, DC 20052, USA, ¹⁵Heliospheric Physics Laboratory, NASA/GSFC, Greenbelt, MD 20771, USA, ¹⁶Università di Siena and INFN Pisa, I-53100 Siena, Italy, ¹⁷Politecnico di Bari, Department of Mechanics, Mathematics and Management, via Orabona, 4, I-70125 Bari, Italy, ¹⁸Department of Physics, Catholic University of America, Washington, DC 20064, USA, ¹⁹Naval Research Laboratory, Washington, DC 20375, USA