Draft version July 6, 2018
Typeset using LaTeX **twocolumn** style in AASTeX62

The Crab Nebula spectrum extending past 100 TeV as seen by the HAWC Gamma-Ray Observatory

HAWC Collaboration[1]

[1] *Put affiliations here*

Submitted to The Astrophysical Journal

ABSTRACT

TeV observations of gamma-ray sources are essential in probing their acceleration mechanisms. We present very-high-energy observations of the Crab Nebula, the standard reference source in gamma-ray astronomy, taken with the High Altitude Water Cherenkov (HAWC) Gamma-Ray Observatory. The data are fit to a log parabola with emission extending past 100 TeV. This measurement serves as the highest-energy observation of a gamma-ray source to date. We describe the analysis techniques used, which are capable of reconstructing energies event-by-event and greatly improve HAWC's energy resolution as compared to previous analyses.

Keywords: acceleration of particles — astroparticle physics — gamma rays: general — ISM: individual objects (Crab Nebula)

1. INTRODUCTION

As one of the most studied gamma-ray sources, the Crab Nebula is firmly established as a leptonic source of gamma rays. The TeV emission arises from inverse-Compton (IC) scattering of electrons on the synchrotron component, with additional contributions from the far-IR dust and 2.7-K microwave background (Atoyan & Aharonian 1996). Measuring the highest-energy TeV IC flux and comparing it to the lower-energy synchrotron flux will allow us to learn about the magnetic field of the Crab Nebula (Amenomori et al. 2015).

Some have pointed out the possibility of an additional hadronic component to the emission (Atoyan & Aharonian 1996; Tanimori et al. 1998), becoming noticeable at no less than 10 TeV. While this is currently unproven, it is still a very intriguing thought, as a hadronic component could identify the Crab Nebula as an accelerator of Galactic cosmic rays.

The Crab Nebula has been observed by nearly all TeV gamma-ray experiments operating during the last three decades, starting with Weekes et al. (1989). However, due to the decreasing sensitivity of most gamma-ray experiments as they approach higher energies, observations in the > 50 TeV range are sparse. Two of the more notable papers in the literature are the HEGRA detection (Aharonian et al. 2004), which includes a 2.7σ detection above 50 TeV, and the limits set by the Tibet Air Shower Array above 100 TeV (Amenomori et al. 2015).

A detection extending past 100 TeV is presented here. Section 2 provides a basic description of the High Altitude Water Cherenkov (HAWC) Gamma-Ray Observatory. Section 3 describes two independent energy-estimation algorithms used to calculate the energies of the primary gamma rays. Section 4 shows the application of these energy estimation-algorithms to the Crab Nebula, while Section 5 discusses possible implications of these results.

Corresponding author: Kelly Malone, Samuel Marinelli
malone@psu.edu, marine20@msu.edu

2. DESCRIPTION OF HAWC

HAWC is a detector located in the state of Puebla, Mexico, at an elevation of 4100 meters. It consists of 300 water Cherenkov detectors, each outfitted with four PMTs (one of which is a high-quantum-efficiency PMT). High-energy gamma and cosmic rays hitting the atmosphere initiate air showers, consisting of relativistic electrons, positrons, and photons. The charged particles are then detected via the Cherenkov light they produce when they enter the water Cherenkov detectors. For more information on HAWC's design and operation, see Smith (2015). HAWC is optimized to detect gamma rays in the 100 GeV to 100 TeV range, although it can

in principle detect emission above this 100 TeV threshold.

HAWC's first published observations of the Crab Nebula (Abeysekara et al. 2017), as well as all subsequent publications up to this point, did not use an event-by-event energy estimator. Instead, the fraction of PMTs hit during an air shower event was used as a proxy for energy. Events were placed in analysis bins (indexed here by $\mathcal{B}$) depending on what fraction of the PMTs in the array were hit during the event. $\mathcal{B}$ is only weakly correlated with energy, as it does not take into account two variables that impact the energy: the zenith angle of the event and how well contained the air shower is within the array. Additionally, $\mathcal{B}$ is only capable of determining the mean energy in an analysis bin; it does not assign energies on an event-by-event basis. This creates a problem in the highest analysis bin, where nearly every PMT is hit. This bin includes nearly every event above ∼20 TeV, making it impossible to distinguish 20-TeV events from 100-TeV ones.

The Crab Nebula spectrum presented in Abeysekara et al. (2017) was only valid up to 37 TeV due to the saturation of the detector at high energies. A complete event-by-event energy analysis is presented here for the first time, allowing for the extension of the spectrum above 100 TeV.

3. ENERGY ESTIMATION

The HAWC Collaboration has developed two independent energy-estimation algorithms. This allows for cross-checking of results. The two methods, the ground parameter (GP) and the neural network (NN), are described below. Throughout this section, $\hat{E}$ will refer to estimated energy while E will refer to the true energy from simulation.

3.1. *Ground parameter*

The GP algorithm is based primarily on the charge density at some optimum distance from the shower axis. It is based on a method used by large cosmic-ray experiments such as the Pierre Auger Observatory and described in Newton et al. (2007), but with some modifications made for the much smaller size of the HAWC array.

The optimum radius to measure the charge at (hereafter known as the "optimal radius") must be far from the air-shower core due to the presence of large shower-to-shower fluctuations there that make energy estimation difficult, but it also must be close enough to the shower axis that the measured PMT signal is large enough that shower fluctuations are not a concern.

To determine this optimal radius, the lateral distribution is plotted on an event-by-event basis and a fit

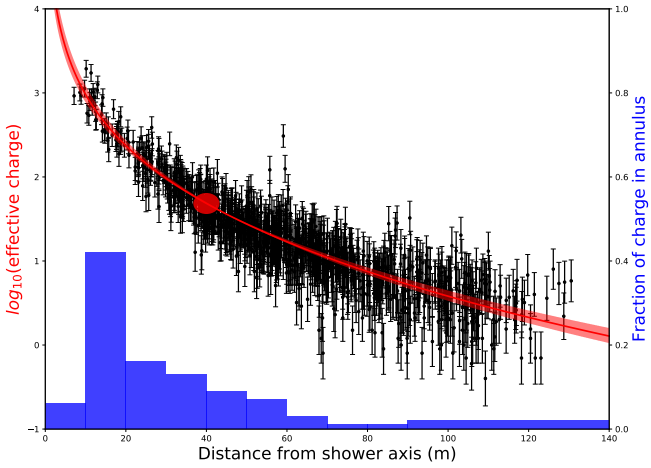

Figure 1. An event showing the information fed into the two energy estimators. The black points show the log of the effective charge, as measured by each PMT, as a function of the distance to the air-shower axis. The red line is the best fit to an NKG-like function, while the band was used to determine the optimal radius (see section 3.1). The circle at $r = 40$ meters visually denotes the location the charge is measured at in the GP method, which is used along with the zenith angle to calculate the estimated energy. The blue histogram is the fraction of charge in several rings, which are used as some of the inputs to the neural network.

to the NKG/r function, which measures energy density (Grieder 2010), is performed (see Figure 1):

$$\log_{10}(\mathrm{NKG}/r) = A + s\left[\log_{10}\frac{r}{r_m} + \log_{10}\left(1+\frac{r}{r_m}\right)\right] - 3\log_{10}\frac{r}{r_m} - 4.5\log_{10}\left(1+\frac{r}{r_m}\right). \quad (1)$$

Here, A is the amplitude of the fit, s is related to the shower age, and r_m is the Moliere radius, which is 124.21 meters at the HAWC site.

After obtaining the best fit, the shower age is varied by ±10%, and additional fits are performed, leaving the normalization free. The point where the fits vary by the least is the point where the uncertainty in not knowing the exact form of the lateral distribution function is minimized and is therefore the optimal radius. For HAWC, this value is found to be ∼40 meters irrespective of zenith angle or primary-particle energy. This mirrors findings in Newton et al. (2007), which notes that the optimal radius is almost solely a function of array geometry.

The log of the charge density at 40 meters ($\log_{10} s_{40}$) is then translated to energy. The simplest implementation of this would be to use Monte Carlo to perform a fit of the form $\log_{10}\hat{E} = f(s_{40})$, where $\hat{E}$ is the estimated energy. However, for a fixed primary energy, the signal

measured on the ground varies strongly with zenith angle due to the different amount of atmosphere air showers entering at different angles must travel through. The formula must be parameterized by zenith angle, becoming $\log_{10}\hat{E} = f(s_{40}, \theta)$. The exact functional form of the fit is chosen empirically to provide a good match to simulated events:

$$\log_{10}\hat{E} = m(\theta)\log_{10} s_{40} + b(\theta)\,. \tag{2}$$

For the performance of the GP on simulation, see Section 3.3.

3.2. *Neural network*

The NN energy-reconstruction algorithm employs an artifical neural network to estimate primary energies of photon events based on several quantities that are computed as part of HAWC's event reconstruction. The Toolkit for Multivariate Analysis (TMVA) NN implementation, described in Hoecker et al. (2007), is used.

The NN energy estimator uses a multilayer-perceptron architecture with two hidden layers and a logistic activation function. The hidden layers have 15 and 14 nodes respectively. The TMVA NN implementation is described in detail in Hoecker et al. (2007). The values of the 479 NN weights are chosen to minimize the error function

$$D(\mathbf{w}) \equiv \frac{1}{2}\sum_{i=1}^{n} u_i \left[\log_{10}\hat{E}(\mathbf{x}_i;\mathbf{w}) - \log_{10} E_i\right]^2 \tag{3}$$

evaluated using Monte Carlo events, where $\mathbf{w}$ is the vector of NN weights, n is the number of events, u_i is the relative importance of the ith event, $\mathbf{x}_i$ is the vector of input variables for the ith event, $\hat{E}$ is the function returning an energy estimate for a given vector of inputs and vector of weights, and E_i is the Monte Carlo-truth energy of the ith event. The u_is are chosen to reweight the Monte Carlo events into an E^{-2} power law, which was selected because a NN trained on such a spectrum was found to produce a relatively constant RMS error between 1 and 100 TeV, as shown in Figure 7. The minimization of the error function is performed via the Broyden-Fletcher-Goldfarb-Shanno algorithm, described in Hoecker et al. (2007).

For the performance of the neural network on simulation, see Section 3.3.

3.2.1. *Input variables*

The NN input variables are chosen to characterize three broad characteristics of the air shower: the amount of energy deposited in the detector, the extent to which the shower's footprint on the ground is contained within the detector, and the degree of attenuation of the shower by the atmosphere. The resulting algorithm can be thought of as a calorimetry measurement combined with corrections for the fraction of the shower not hitting the detector and for the atmospheric attenuation.

Three quantities are used to infer the amount of energy deposited in the detector: the fraction of PMTs hit within the event, the fraction of tanks hit, and the base-10 logarithm of the normalization from the fit of the lateral distribution function. These are all positively correlated with the shower's primary energy.

The fraction of the shower landing within the detector on the ground is inferred using the distance between the reconstructed core location of the shower and the center of the HAWC array.

The atmospheric attenuation of the shower is quantified in two ways: using the cosine of the reconstructed zenith angle of the shower and using the shower's lateral distribution, which contains information about the shower age. The lateral distribution is passed to the NN in the form of ten input variables. The first nine of these variables consist of the fraction of the charge deposited in all PMTs during the event that lands within each of nine concentric annuli around the reconstructed shower axis. Each annulus has a thickness of 10 m. The last of these ten input variables is the fraction of the event's charge landing more than 90 m from the shower axis (see Figure 1).

3.3. *Performance of the estimators*

3.3.1. *Bias, resolution, and RMS error*

Both energy estimators have been throughly studied with Monte Carlo simulations (see section 2.1 of Abeysekara et al. (2017) for details of the simulation chain).

The mixing matrices, which compare the performance of the event-by-event energy estimate to the true energy, can be seen in Figures 2 and 3. Figure 4 shows an event-by-event comparison of the two estimators. Both figures assume an isotropic E^{-2} spectrum of gamma rays. This spectrum assures that there are sufficient events at high energy to evaluate the performance. Only simulated gamma-ray events whose shower core is reconstructed on the HAWC array and have a zenith angle of $< 45^o$ are used. Additionally, the reconstructed zenith angle must be less than 0.75^o from the true Monte Carlo value.

Two quantities are used to evaluate the energy-dependent performance of the estimators. The first is the resolution: the standard deviation of the energy estimate in log-energy space. The second is the bias, defined as the average difference between the reconstructed and true energies in log space:

$$b \equiv \left\langle \log_{10}\hat{E} - \log_{10} E \right\rangle \tag{4}$$

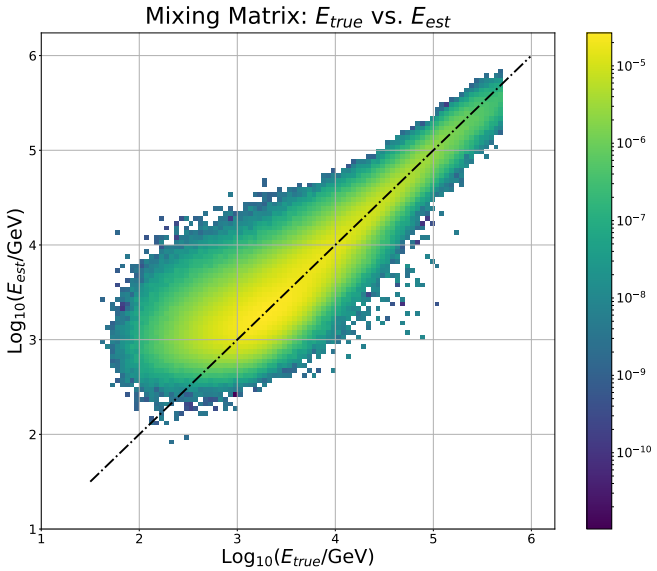

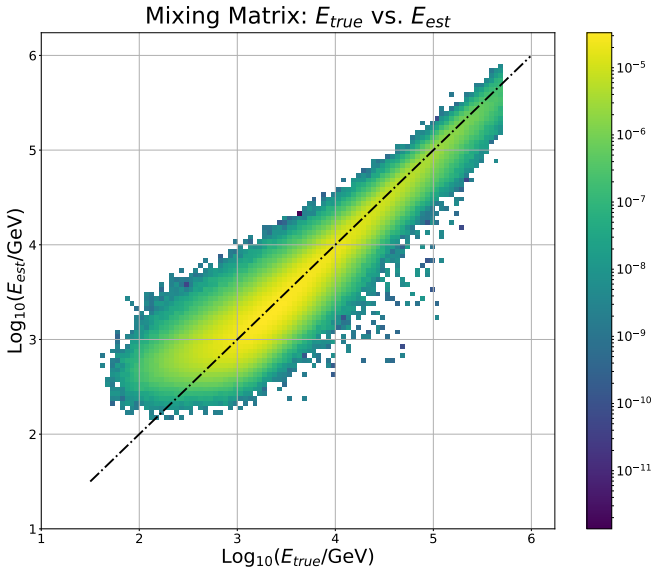

Figure 2. The mixing matrices for the GP (top) and NN (bottom) estimators, without background-rejection cuts. The dotted line is the 1:1 line, events that fall along this line are reconstructed perfectly.

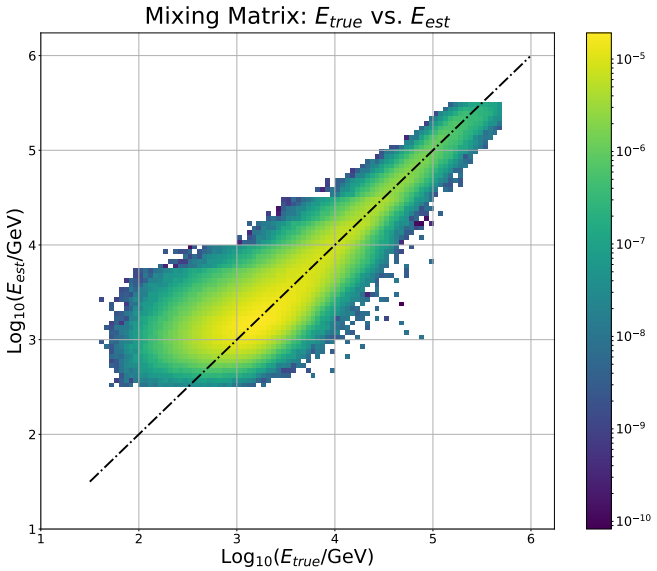

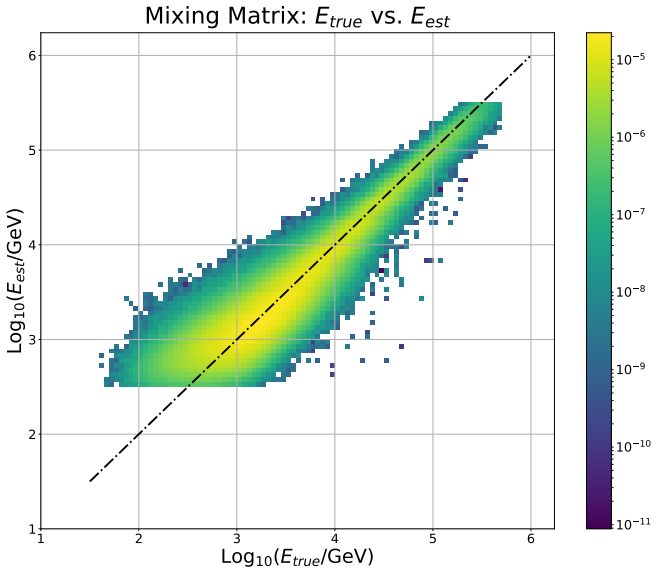

Figure 3. The mixing matrices for the GP (top) and NN (bottom) estimators, with background-rejection cuts applied.

The bias and resolution for both estimators can be seen in Figures 5 and 6. Both the NN and GP have a large bias below 1 TeV. This is due to data-quality cuts and the requirement that a minimum of 6.7% of the array be hit, which remove the vast majority of events below this energy; the only events left are from air showers with upward fluctuations in the number of PMTs hit. Due to the undesirable bias below 1 TeV, events with reconstructed energies below this threshold are excluded from the spectral fit.

Note that both estimators have very good resolution and almost no bias in the high-energy regime (10 TeV – 316 TeV). Above 316 TeV, the energy estimate is not defined due to a lack of available simulation at this energy.

The log RMS error is defined as

$$\rho \equiv \sqrt{\left\langle \left(\log_{10} \hat{E} - \log_{10} E \right)^2 \right\rangle}. \tag{5}$$

This is also equal to the bias and resolution added in quadrature. Figures 7 and 8 show the log RMS error for both energy estimators.

3.3.2. *Bin contamination in spectral fits*

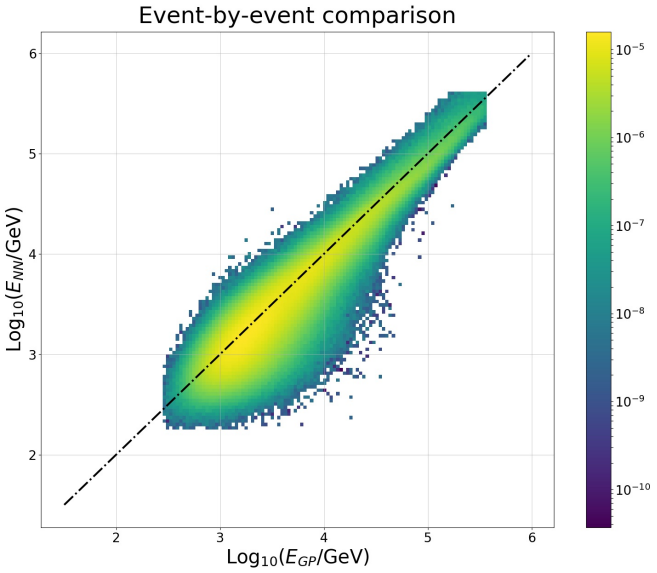

Figure 4. Event-by-event comparison of the two estimators, for gamma ray events that pass data quality cuts but before gamma/hadron sepration has been applied. The dotted line is the 1:1 line. Events falling on this line have the same energy estimate regardless of which method is used.

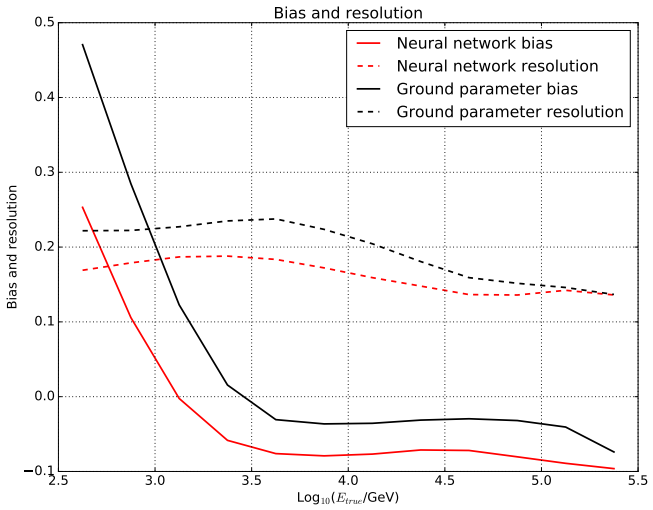

Figure 5. The bias and resolution for both energy estimates. No gamma/hadron separation has been applied yet.

Bin purity measures the contamination of a reconstructed-energy bin by mis-reconstructed events. It is defined here as the fraction of events in a half-decade reconstructed energy bin whose Monte Carlo-truth energy is also within that bin:

$$p \equiv P\left(E \in B \middle| \hat{E} \in B\right) \tag{6}$$

where B is a half-decade energy bin. Because astrophysical sources emit following spectra that are roughly power-law in shape, there are many more lower-energy gamma rays than higher energy ones. If even a small percentage of these low energy gamma rays are mis-reconstructed with a higher energy, the bin purity can be adversely affected. Bin purity is a function of spectral index. A softer spectrum will have worse bin purity. Figure 9 shows the bin purity for both estimators. For a power-law spectrum with index between 2 and 3, bin purity is $> 75\%$ above 100 TeV.

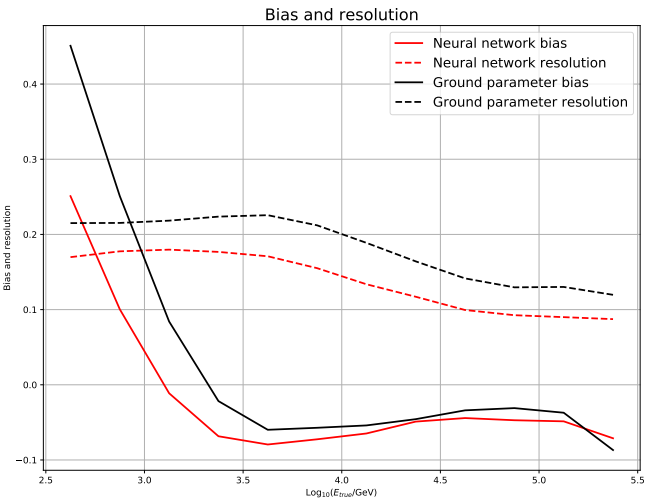

Figure 6. The bias and resolution for both energy estimates, with hadron-rejection cuts applied.

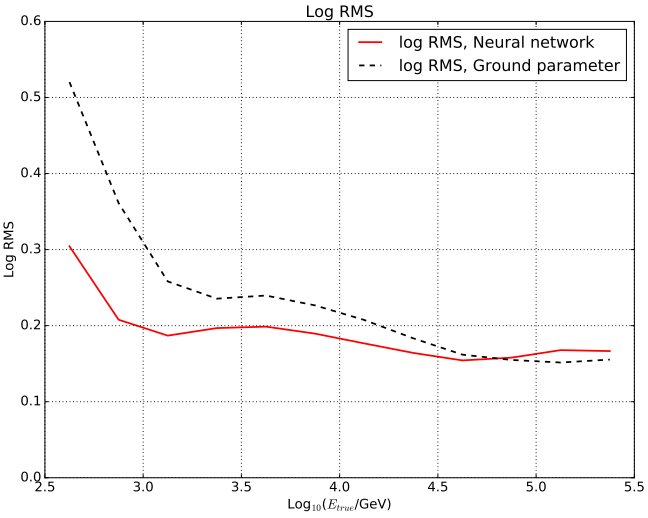

Figure 7. The log RMS error for the GP and NN estimators. No gamma/hadron separation has been applied yet.

4. CRAB NEBULA SPECTRAL ANALYSIS

4.1. *Dataset*

The data used in this analysis were collected between November, 2014, and December, 2017. The total live-time is 1017 days ($> 90\%$ uptime). The loss of livetime

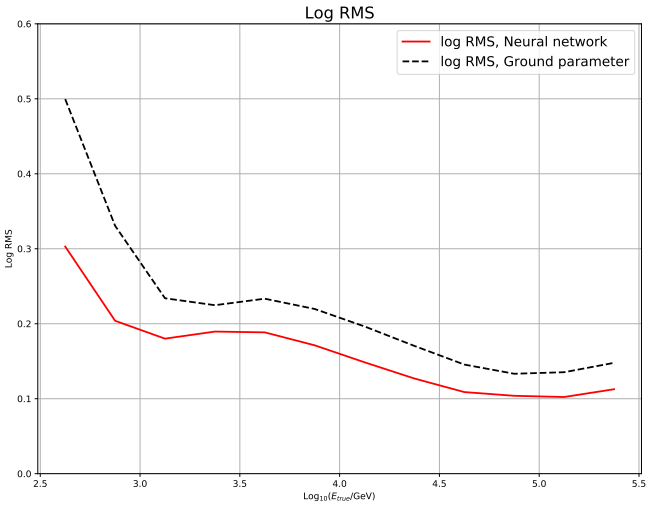

Figure 8. The log RMS error for the GP and NN estimators, with hadron-rejections cuts.

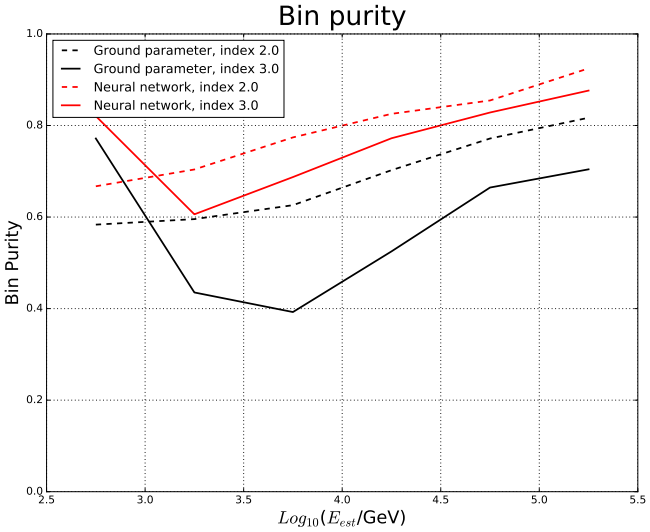

Figure 9. The bin purity for both estimators, for an extremely hard and an extremely soft spectrum.

comes from days where the detector was off for construction or due to operational difficulties. Additionally, a small amount of data was removed due to large variances in the zenith angle distribution, which is an indication that the detector was unstable during that time period. These instabilities make our background estimation method unusable, so they are removed.

4.2. *Event selection and binning*

The spectral fit is performed using a binned-likelihood technique. We use a 2D binning scheme based on $\mathcal{B}$ (described in Section 2) and the estimated energy. This 2D binning scheme was chosen instead of binning solely in energy because the gamma/hadron separation parameters depend on both the energy and size of the event. We use 9 $\mathcal{B}$ bins each subdivided into 12 energy bins, for a total of 108 bins. The energy bins consist of 9 bins each spanning a quarter decade in $\hat{E}$. See Tables 1 and 2 for the bin definitions. For example, bin 9k would be all of the events with $> 84\%$ of the array hit and energies between 100 TeV and 177 TeV.

Note that the definition of $\mathcal{B}$ has changed slightly from Abeysekara et al. (2017): there, $\mathcal{B}$ was defined as the number of PMTs seeing light divided by total number of PMTs that were operational at the time. Here, the numerator is changed to the fraction of PMTs seeing light that are within 20 ns of the shower front. This change reduces the impact of noise hits contributing to the size of the event.

The gamma/hadron separation variables in this analysis are unchanged from Abeysekara et al. (2017); two such parameters are used. Compactness, first described in Abeysekara et al. (2013), is effective at identifying air showers containing muons. Muons are only present in hadronic showers and appear as localized charge depositions far from the shower core. The second parameter is known as PINCness (Abeysekara et al. 2017) and measured the smoothness of the lateral distribution function. Gamma-ray showers have smoother profiles than hadronic air showers.

The gamma/hadron cuts are optimized in each 2D $\mathcal{B}$/energy bin. This allows for better identification of the highest energy events as compared to Abeysekara et al. (2017), where everything above 10 TeV was included in one analysis bin and had the same gamma/hadron cuts.

The cut values are determined a priori on simulation, with a requirement that each bin have at least 50% gamma ray efficiency. The efficiency to gamma rays in a given bin ranges from 50% to nearly 100%.

After events are binned in each the $\mathcal{B}$/energy bins, all sky maps for a given energy are combined. The result is 12 1-dimensional skymaps, corresponding to the energies given in Table 1.

The point spread function (PSF) is modeled as a linear combination of two two-dimensional Gaussians. Better modeling of this PSF at high energies is one of the more dramatic changes from the previous Crab analysis. Previously, all events above ∼20 TeV had the same average PSF. Figures 10 and 11 show that the angular resolution for these events can now be differentiated. Additionally, the energy estimators perform much better for events where the shower core is reconstructed on the HAWC array. For this reason, off-array events are removed from the analysis, also improving the angular resolution.

4.3. *Background estimation*

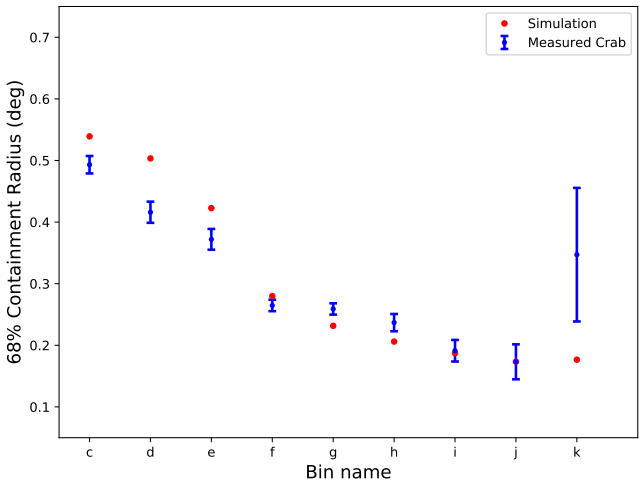

Figure 10. The 68% containment values in data and Monte Carlo for the ground parameter energy estimate. Only bins where the Crab Nebula is detected $> 3\sigma$ are shown. The data/MC discrepancy, described in the text, is treated in the systematic error analysis.

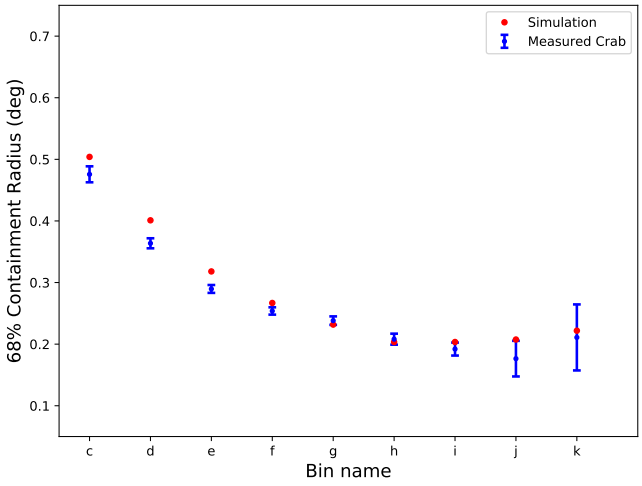

Figure 11. The 68% containment values in data and Monte Carlo for the neural network energy estimate. Only bins where the Crab Nebula is detected $> 3\sigma$ are shown. The data/MC discrepancy, described in the text, is treated in the systematic error analysis.

Even after gamma/hadron separation has been applied, hadronic cosmic rays dominate over the gamma rays. This is especially true for smaller showers. An estimate of this cosmic ray background is performed individually in each analysis bin. For the lower energy bins, where there are high statistics, the standard HAWC background estimation technique is applied. This is known as "direct integration." This algorithm was originally developed by the Milagro Collaboration (Atkins et al. 2003). Because HAWC's background rejection exceeds that of Milagro, the background must be smoothed by an 0.5° to compensate for the sparseness of the background.

In the highest energy bins, where the statistics are extremely low, direct integration fails to work. A new algorithm known as "background randomization" is used in those bins instead. For each bin, a 2D distribution of the local coordinates (zenith and azimuth) is constructed. A random (zenith, azimuth) pair is pulled from this distribution for each event and used with the time of the event to calculate a right ascension and declination, which is added to the background map. This process is repeated 10,000 times for each event and then the background map is normalized to the number of events in the map. This produces a background estimate much smoother than given by direct integration.

Table 1. Energy bins

Bin number	Energy range (TeV)
a	0.316-0.562
b	0.562-1.0
c	1.0 - 1.78
d	1.78 - 3.16
e	3.16 - 5.62
f	5.62 - 10.0
g	10.0 - 17.8
h	17.8 - 31.6
i	31.6 - 56.2
j	56.2 - 100
k	100 - 177
l	177 - 316

The energy bins. Each one spans one quarter decade. Note that the first two bins are not used in this analysis as the estimate is highly biased, as explained in Section 3.3

Table 2. $\mathcal{B}$ bins

Bin number	Fraction of array hit
1	6.7-10.5
2	10.5-16.2
3	16.2-24.7
4	24.7-35.6
5	35.6-48.5
6	48.5-61.8
7	61.8-74.0
8	74.0-84.0
9	84.0-1.00

The $\mathcal{B}$ (fraction hit) analysis bins used in this paper.

4.4. *Likelihood fit*

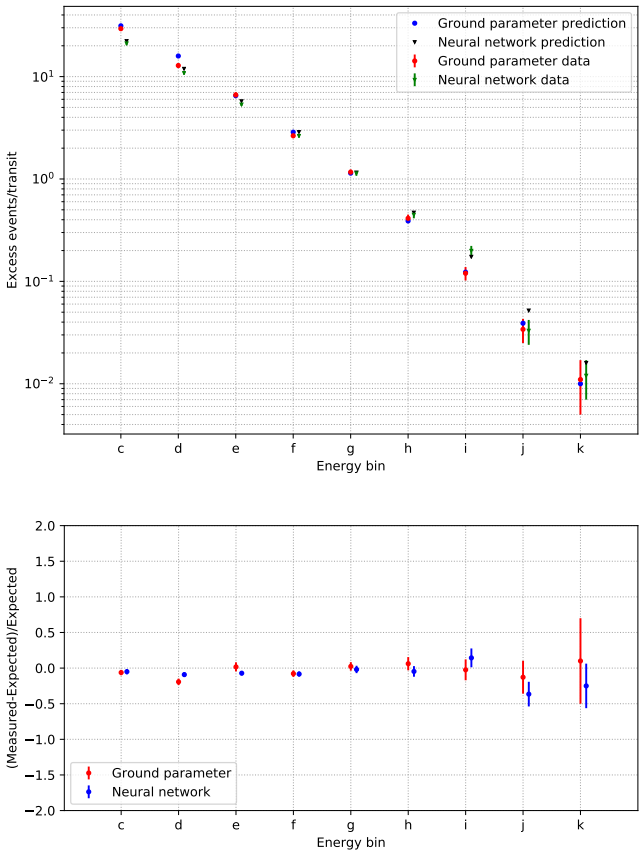

Figure 12. The Crab excess per transit, along with the residual (defined as [measured-expected]/expected). The two estimators have different numbers of events in some bins due to bias and different efficiencies to gamma rays.

The measured excess per transit, along with the expected value from simulation, can be seen in Figure 12. From simulation, we expect ∼59 gamma rays with a reconstructed energy above 1 TeV per day from the Crab Nebula using the ground parameter analysis chain and ∼ 45 using the neural network. The values are different because the gamma/hadron cuts were optimized separately for the two techniques and they therefore have slightly different efficiencies to gamma rays.

The fit is a forward-folded fit assuming a log-parabola spectrum:

$$\frac{dN}{dE} = \phi_0 \left(E/E_0\right)^{-\alpha-\beta \ln(E/E_0)} \tag{7}$$

The pivot energy, E_0, was chosen to be 7 TeV. The other parameters are kept free in the fit, which is performed using the HAWC plug-in to the Multi-Mission Maximum Likelihood framework (Younk et al. 2015; Vianello et al. 2015), an analysis pipeline that is capable of handling data from a wide variety of gamma-ray experiments.

The spectra of the Crab Nebula obtained using the two energy estimators can be seen in Figure 13. For the ground parameter, the best fit values are $\phi_0 = 2.55 \times 10^{-13} \pm 6.21 \times 10^{-15}$, $\alpha = 2.715 \pm 0.020$ and $\beta = 0.069 \pm 0.011$. For the neural network, the best fit values are $\phi_0 = 3.23 \times 10^{-13} \pm 5.33 \times 10^{-15}$, $\alpha = 2.598 \pm 0.014$ and $\beta = 0.084 \pm 0.010$. The log parabola fit is preferred over a simple power law by 6.40σ for the ground parameter and 9.13σ for the neural network. Errors quoted here are statistical only. Systematic errors are discussed in Section 4.5.

Flux points are calculated by holding α and β constant from the global fit and fitting the normalization individually for each reconstructed energy bin. While this is not a full unfolding prescription, it allows one to see if any energy bins are inconsistent with the fitted log parabola spectrum. Figure 13 shows the Crab Nebula spectrum computed in this manner. The error bars are the statistical and systematic errors added in quadrature. The location of each point on the x- axis is determined following Lafferty & Wyatt (1995), using the 10%-90% quantiles in true MC energy for a given reconstructed energy bin as the bin size.

The test statistic of the last bin (> 177 TeV) is -0.25 for the ground parameter and -2.88 for the neural network. Upper limits are set here using a 95% upper confidence interval following Feldman & Cousins (1998).

4.5. *Systematics*

The systematic errors described in section 4.3 of HAWC's previous Crab analysis (Abeysekara et al. 2017) are also present here, although improved modeling has decreased the amount of some of these systematic errors. When all of these sources of error are considered and added in quadrature, the systematic error on α is (-0.12, + 0.05) and the systematic error on β is (-0.03, + 0.04) for the ground parameter. The corresponding values for the neural network are (-xx,+xx) and (-xx,+xx). The systematic error on the flux points was studied as a function of energy, ranging from 20-25% below 32 TeV and increasing to ∼ 50% above that point. The > 100 TeV point has a large flux uncertainty but much of that is due to the time dependence of the spectrum (discussed below).

4.5.1. *Angular resolution discrepancy*

A discrepancy in the 68% containment between data and simulation from Figures 10 and 11. While the cause of this is not immediately clear, it is thought to be because the curvature model used during reconstruction does not yet have an energy-dependence. This is an avenue for future study.

The effect of this has been studied by artificially matching the point spread function in simulation to the

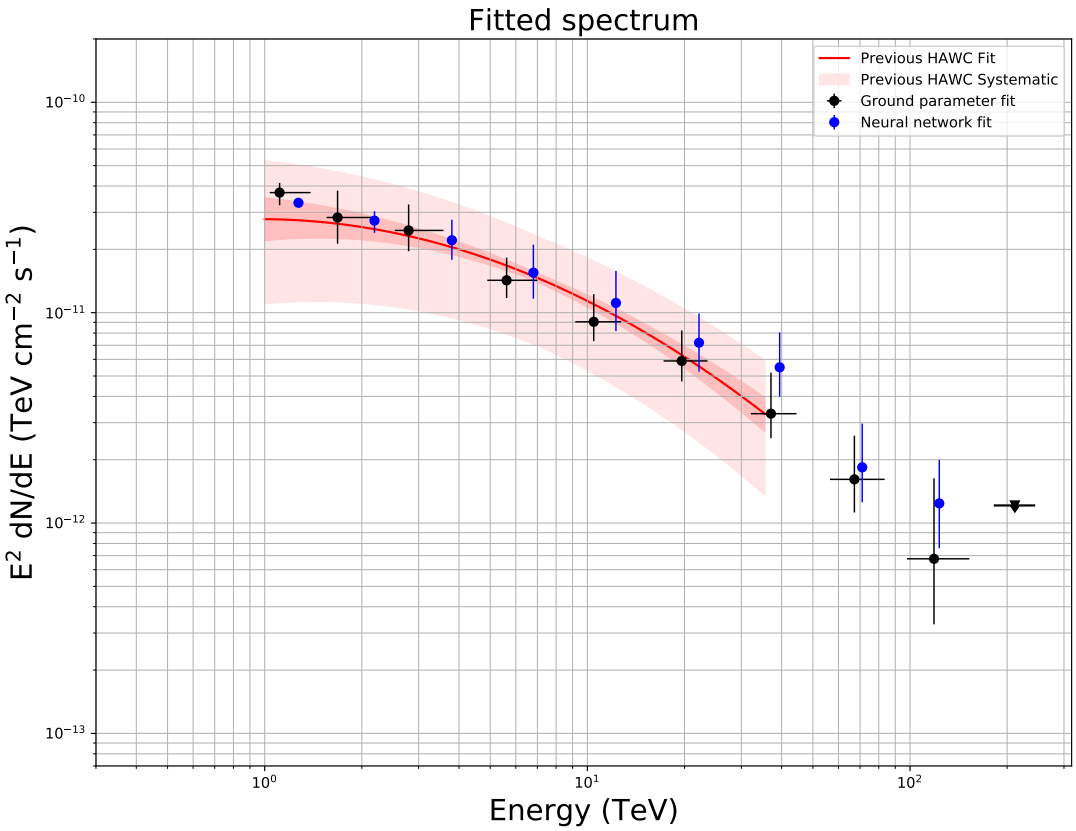

Figure 13. The Crab spectrum obtained with the ground parameter method (black) and neural network method (blue). The error bars are statistical and systematic errors added in quadrature. For comparison, the HAWC Crab fit from Abeysekara et al. (2017) is also shown. See the text for details of how the flux points were obtained.

data PSF and re-fitting the Crab spectrum. This systematic is at most ∼10-15% in E^2dN/dE for the ground parameter.

4.5.2. *Time dependence*

This dataset spans 1017 days, from November 2014 to December 2017. The detector layout changed during this time period; the detector was still under construction in 2014 and the early part of 2015. Additionally, PMTs are occasionally taken out of service for maintenance, and the calibration also changes as a function of time.

To investigate any effect on the changing detector layout, the data were divided into two chunks roughly equal in livetime and the fit was performed in both. The measured flux was found to be lower in the first half of the dataset than the second half, an effect of roughly ∼ 20% for the ground parameter.

During this study, it was found that all the of > 100 TeV gamma rays were observed during the second half of the livetime, making the uncertainty on that data point to be ∼100%. This not necessarily a cause for concern. Borione et al. (1997) noted in a paper about a CASA-MIA search for ultra-high-energy emission from the Crab Nebula that the >100 TeV flux may be expected to vary on timescales of ∼ 1 year, since the MeV flux varies on this timescale. If the MeV, created via synchrotron radiation, varies, then the parent electron population and therefore the 100 TeV gamma rays created via IC will also exhibit a time dependence.

Alternatively, this effect could be because the first half of the data was taken with a partially constructed array, and this prevented the detection of high-energy emission. A few more years of data should provide a conclusive answer to this question.

4.5.3. *Late light simulation*

The largest source of error (∼40% in flux) in Abeysekara et al. (2017), this systematic arises from a mismodeling of how late light in the air shower is treated. From simulation, it is expected that the arrival time distribution of PEs at the PMT should be <∼10 ns, but examining the raw PE distributions in data shows a discrepancy above 50 PEs.

Improved studies of the PMTs have decreased the size of this error here, affecting each flux point between 5 and 10 percent for the ground parameter.

4.5.4. *Charge uncertainty*

The charge uncertainty encapsulates how much a PMT measurement will vary for a fixed amount of light, and also the relative differences in photon detection efficiency from PMT to PMT. These two variables together are estimated to vary by ±15%. This number was estimated by selecting tanks with isolated downward going muons, plotting the resulting muon peak in PEs, and looking at the width of the distribution.

Propagating this charge uncertainty through to the flux, it is seen that this is a highly consequential parameter. It is the biggest contributor to the uncertainty on α and is also one of the dominant contributions to the flux systematic above 10 TeV ($\sim$ 30-50% for the ground parameter).

4.5.5. *Absolute PMT efficiency*

The absolute PMT efficiency cannot be precisely determined (see Abeysekara et al. (2017) for a discussion). The amount of uncertainty is estimated to be ±20%. This was determined by plotting the muon peak against the PE threshold required for a PMT to have a rate of 100 Hz and performing a linear fit. The width of the residuals are $\sim$ 20%.

Absolute PMT efficiency is the dominant contributor to the distribution of true MC energies, which affect the x-error bar (±20% for the ground parameter). It is not a large contributor to the overall flux systematic.

4.5.6. *PMT threshold*

The PMT threshold is set at 0.55 PEs in simulation so that the single PE peak matches the data. However, from looking at the cosmic ray rate, the true PMT threshold may be as low as 0.3 PEs.

Assuming a PMT threshold of 0.3 PEs, the first two data points for the ground parameter move down by 20%. There are small changes in the flux at the higher energies, but the amount is dwarfed by more consequential parameters such as charge uncertainty.

4.5.7. *Disagreement between the two estimators*

Looking at Figure 13, it is apparent that the spectral fits with the two methods do not agree within statistical errors above 10 TeV. This is likely due to potentially mismodeled parameters in the Monte Carlo that affect one estimator but not the other.

This disagreement does not affect the significance of ultra high energy emission from the Crab Nebula. Regardless of which analysis technique is used, $> 5\sigma$ emission above 100 TeV is observed.

5. DISCUSSION AND CONCLUSIONS

This observation is the highest-energy observation of the Crab Nebula to date. This is not to say that the results are unexpected. As early as 1992, de Jager & Harding (1992) predicted > 50 TeV gamma rays from the Crab simply by using the observed synchrotron emission and a MHD flow model. Atoyan & Aharonian (1996) later used a MHD flow model to self-consistently calculate both the synchrotron and IC emission and obtained predictions for > 30 TeV.

Horns & Aharonian (2004) have noted that photons up to at least 80 TeV are expected, coming from Inverse Compton scattering from a combination of synchrotron seed photons and soft photons with mm wavelengths. Interpreting > 50 TeV photons as IC emission implies the presence of electrons with PeV energies.

Now that a method to identify gamma rays above 50 TeV (and even above 100 TeV) has been developed, the door has been opened to do even more interesting analyses with HAWC. HAWC is the most sensitive currently operating TeV experiment above 10 TeV in the world(Abeysekara et al. 2017). High-energy surveys of the sky are already in progress. Extending spectra of previously discovered sources up to 100 TeV will allow us to distinguish between leptonic and hadronic gamma-ray emission mechanisms, as they have different signatures. This, in turn, may help determine if any Galactic gamma-ray sources are good candidates to be the source of the IceCube astrophysical neutrinos(Aartsen et al. 2013)[1] . Extending the spectra to high energies may even help identify some as PeVatron candidates and give glimpses into the origins of cosmic rays(Gabici & Aharonian 2007). Additionally, high energy observations also naturally lead to studies of Lorentz invariance violation(Martínez-Huerta 2017).

HAWC's sensitivity will get an even greater boost in the near future when its upgrade is completed. Currently under construction, this sparsely populated "outrigger" array will allow for better reconstruction of the largest, most energetic events(Joshi & Jardin-Blicq 2017).

We acknowledge the support from: the US National Science Foundation (NSF); the US Department of Energy Office of High-Energy Physics; the Laboratory Directed Research and Development program of Los Alamos National Laboratory; Consejo Nacional de Ciencia y Tecnologia, Mexico (grants 271051, 232656, 260378, 179588, 239762, 254964,

[1] Due to gamma-ray attenuation, it is expected that > 50 TeV gamma rays will only arrive at Earth from nearby sources

271737, 258865, 243290, 132197, and 281653) (Cátedras 873, 1563); Laboratorio Nacional HAWC de rayos gamma; L'OREAL Fellowship for Women in Science 2014; Red HAWC, Mexico; DGAPA-UNAM (Direccion General Asuntos del Personal Academico-Universidad Nacional Autonoma de Mexico; grants IG100317, IN111315, IN111716-3, IA102715, 109916, IA102917); VIEP-BUAP (Vicerrectoria de Investigacion y Estudios de Posgrado-Benemerita Universidad Autonoma de Puebla); PIFI (Programa Integral de Fortalecimiento Institucional) 2012 and 2013; PROFOCIE (Programa de Fortalecimiento de la Calidad en Instituciones Educativas) 2014 and 2015; the University of Wisconsin Alumni Research Foundation; the Institute of Geophysics, Planetary Physics, and Signatures at Los Alamos National Laboratory; Polish Science Centre grant DEC-2014/13/B/ST9/945; and Coordinacion de la Investigacion Cientifica de la Universidad Michoacana. Thanks to S. Delay, L. Diaz, and E. Murrieta for technical support.

REFERENCES

Aartsen, M. G., Abbasi, R., Abdou, Y., et al. 2013, Science, 342, 1. http://science.sciencemag.org/content/342/6161/1242856

Abeysekara, A. U., Alfaro, R., Alvarez, C., et al. 2013, Astropart. Phys., 50-52, 26

Abeysekara, A. U., Albert, A., Alfaro, R., et al. 2017, The Astrophysical Journal, 843, 39. http://iopscience.iop.org/article/10.3847/1538-4357/aa7555/meta

Aharonian, F. A., Akhperjanian, A. G., Beilicke, M., et al. 2004, Astrophys. J., 614, 897

Amenomori, M., Bi, X. J., Chen, D., et al. 2015, Astrophys. J., 813

Atkins, R., Benbow, W., Berley, D., et al. 2003, The Astrophysical Journal, 595, 803. http://iopscience.iop.org/article/10.1086/377498/meta

Atoyan, A. M., & Aharonian, F. A. 1996, Monthly Notices of the Royal Astronomical Society, 278, 525. https://academic.oup.com/mnras/article/278/2/525/951962

Borione, A., Catanese, M., Chantell, M., et al. 1997, The Astrophysical Journal, 481, 313. http://iopscience.iop.org/article/10.1086/304042/meta

de Jager, O. C., & Harding, A. K. 1992, The Astrophysical Journal, 396, 161

Feldman, G. J., & Cousins, R. D. 1998, Physical Review D, 57, 3873. https://doi.org/10.1103/PhysRevD.57.3873

Gabici, S., & Aharonian, F. A. 2007, The Astrophysical Journal Letters, 665, L131. http://iopscience.iop.org/article/10.1086/521047/meta

Grieder, P. K. 2010, Extensive Air Showers: High Energy Phenomena and Astrophysical Aspects (A Tutorial, Reference Manual, and Data Book: Volume I), 2nd edn. (Heidelberg: Springer), 190–191

Hoecker, A., Speckmayer, P., Stelzer, J., et al. 2007, PoS, ACAT, 040. https://arxiv.org/abs/physics/0703039

Horns, D., & Aharonian, F. A. 2004, in ESA Spec.Publ. No. 552, 439–445. https://arxiv.org/abs/astro-ph/0407119

Joshi, V., & Jardin-Blicq, A. 2017, Proceedings of Science (35th International Cosmic Ray Conference), 806. https://pos.sissa.it/301/806/pdf

Lafferty, G., & Wyatt, T. 1995, NIM A, 355, 541. https://doi.org/10.1016/0168-9002(94)01112-5

Martínez-Huerta, H. 2017, Proceedings of Science (35th International Cosmic Ray Conference), 301. https://pos.sissa.it/301/868/pdf

Newton, D., Knapp, J., & Watson, A. A. 2007, Astropart. Phys., 26, 414. http://www.sciencedirect.com/science/article/pii/S0927650506001216

Smith, A. J. 2015, Proceedings of Science(ICRC2015(966)). https://pos.sissa.it/236/966/pdf

Tanimori, T., Sakurazawa, K., Dazeley, S. A., et al. 1998, The Astrophysical Journal, 492, 33

Vianello, G., Lauer, R. J., Younk, P., et al. 2015, Proceedings of Science(ICRC2015(1042)), 1507.08343

Weekes, T. C., Cawley, M. F., Fegan, D. J., et al. 1989, Astrophys. J., 342, 379

Younk, P. W., Lauer, R. J., Vianello, G., et al. 2015, Proceedings of Science(ICRC2015(948))