

Anomaly Detection in Early Data From the Radio Neutrino Observatory Greenland

Z.S. Meyers^{a,b} and A. Nelles^{a,b} for the RNO-G Collaboration

(a complete list of authors can be found at the end of the proceedings)

^a*Deutsches Elektronen-Synchrotron DESY,
Platanenallee 6, 15738 Zeuthen, Germany*

^b*Erlangen Center for Astroparticle Physics (ECAP), Friedrich-Alexander-Universität Erlangen-Nürnberg
91058 Erlangen, Germany*

E-mail: zachary.meyers@desy.de

After two seasons of installation, 7 stations built and many lessons learned, the Radio Neutrino Observatory Greenland (RNO-G) is now operational. In the coming years, the construction of another 28+ stations will bring the array to full capacity as an instrument with an eye towards the ultra-high energy neutrino (>10 PeV) regime, creating another link in the fast paced and rapidly changing landscape of multi-messenger astronomy. Until now, the data volume of our two initial seasons has remained manageable. However, as the array continues to grow, we need to develop faster and more efficient processes regarding how to filter our data; we must throw away the noise and identify the most promising events. Data reduction tools become crucial for anthropogenic, environmental and local noise identification/removal in order to test and monitor our instrument as we scale up. We present a convolutional encoder-decoder network that assigns an anomaly ranking to events, helping to classify different categories of background and signal.

The 38th International Cosmic Ray Conference (ICRC2023)
26 July – 3 August, 2023
Nagoya, Japan



1. Introduction

The Radio Neutrino Observatory Greenland (RNO-G) is a neutrino telescope currently under construction in the Greenlandic Ice Sheet, with support and infrastructure provided by Summit Station [1]. When fully operational, the array will consist of 35 independent "hybrid" stations, each with a deep (three 100 meter holes drilled into the ice containing vertically polarized (Vpol) and horizontally polarized (Hpol) antennas) and a shallow (9 log-periodic dipole antennas (LPDAs)) component. This hardware aims to observe the Askaryan radio emission induced by ultra-high energy neutrinos interacting in the ice. There are currently four constantly running trigger schemes. The deep **Low-Threshold Trigger (LT)** is a 2/4 coincidence in the four *phased array* Vpol antennas at the bottom of the detector. Near the surface, there is the **Upward Facing Radiant Trigger (RF0)** requiring a 2/3 coincidence in the upward pointing LPDAs in order to be sensitive to Cosmic Ray Air Showers, and the **Downward Facing Radiant Trigger (RF1)** requiring a 2/6 coincidence of LPDAs directed towards the ice. **Forced Triggers** are taken at an interval of every 10 seconds in order to record the local noise environment.

Should RNO-G report a neutrino detection, we have shown the ability to reconstruct the simulated neutrino energy [2] and direction [3]. However, with a designed trigger rate of 1 Hz, it would not be feasible to carry out these reconstruction algorithms in real time on all events. Furthermore, due to the remoteness of Summit Station and the lack of computing infrastructure, we only receive a random 10% of our recorded data during the year (we are only able to retrieve the full data sets while physically in Greenland). And once the full array is complete, this number is expected to drop to under 3%.

There are two bottlenecks in the data flow: the throughput of our LTE network and the limitations of our satellite transfer connection. Regarding the former, the LTE network allows each individual station to send data back to our small data center at Summit Station. While this limit of data transfer is constant per station, there are limited resources on board each station for discrimination if we want to achieve future trigger rates $> 1\text{Hz}$. Please see [4] for a deep learning approach for improving our trigger. Other computationally simple Level 1 (L1) cuts are also being considered. Unfortunately, the second bottleneck, our satellite connection, will scale as the detector grows. Our on site data center offers more computational power than the stations alone, but efficiency and speed are still vitally important. In the modern era of Multi-Messenger Astronomy, being able to rapidly react to, as well as send our own alerts, becomes essential for coincident observations of transient sources. A single metric, as calculated via the currently under development anomaly detection pipeline, can be helpful in deciding which events to transfer south for further analysis.

2. Anomaly Detection

Anomaly Detection in signal processing involves identifying patterns or events within a signal or dataset that do not conform to expected behavior, signifying potential outliers, signal, abnormalities, or other items of interest. In the following we describe how to use a convolutional autoencoder-decoder for anomaly detection.

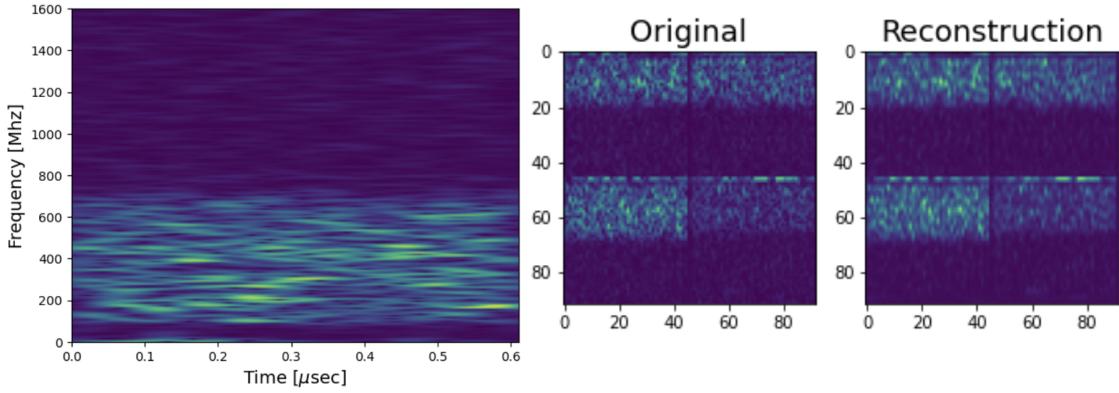


Figure 1: Left: A high resolution STFT of a forced trigger event. Right: A mosaic of four STFTs from the deep phased array channels as actually used in the anomaly detection pipeline. The frequency axis is flipped in order to reduce operations, and the resolution has been optimized for speed. The original input STFT and the network’s reconstruction are shown.

2.1 Short Time Fourier Transform

RNO-G data consists of *events*. When certain trigger criteria are fulfilled, the information from each antenna, or *channel*, is digitized and recorded. The resulting product is a 24 channel waveform or *traces* from the various antennas. For the 2022 season, every trace has 2048 samples at 3.2 GHz, corresponding to 640 nanoseconds. One could use a Fourier Transform to extract the frequency content of the traces. However, when this is applied for the whole trace, we lose the timing information. Alternatively, we can employ a discrete Short Time Fourier Transform (STFT), where the data is broken up into different windows in order to get a two-dimensional snapshot of the frequency content evolving over the trace:

$$\text{STFT}x[n](m, \omega) \equiv X(m, \omega) = \sum_{k=-\frac{n_w}{2}}^{\frac{n_w}{2}} x[m+k]w[k]e^{-i\omega(m+k)\Delta t} \quad (1)$$

where $X(m, \omega)$ is the time-frequency representation of the signal $x[n]$, showing how the frequency content changes over time. m represents the center of the window function $w[k]$, used to taper the signal within each window to mitigate spectral leakage effects caused by the finite length of the window. We are then able to make a spectrogram by taking the squared absolute value as seen in Figure 1:

$$\text{spectrogram}\{x[n]\}(m, \omega) \equiv |X(m, \omega)|^2 \quad (2)$$

2.2 Autoencoder - Decoder Network

One powerful set of tools for image processing are Autoencoders. This specific type of feedforward neural network is well developed in common software packages such as *PyTorch* and *TensorFlow*, the latter being used for this analysis. We train an encoder (with two convolutional layers and one dropout) and a parallel decoder that reduces our STFTs into a latent space with lower dimensionality, and then reconstructs the original signal out of the compressed space. By comparing

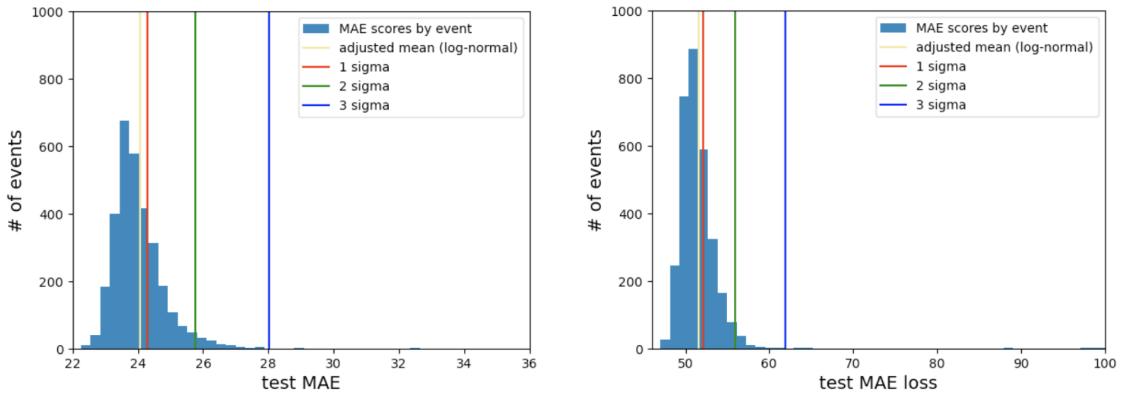


Figure 2: Comparing 2D and 3D network performance and architectures for station 21 "Amaroq". While the 3D case eliminates any artificial biases introduced from the borders of the mosaic image and appears to be more sensitive to timing correlations, it has greater than a factor of 10 more trainable parameters and the operations with data are much more computationally intensive. The 1, 2, and 3 sigma lines correspond to the 68, 95, and 99.7 containment of the heavy tailed log-normal distribution.

the original image to the reconstruction pixel by pixel, we can compute a Mean Absolute Error (MAE) score that quantifies how anomalous the event is. For this study, the network was trained on 13,600 unlabeled forced trigger events from the 2022 season. The data is randomly shuffled and 10% is saved as a validation set for measurement every epoch. A model is trained for every individual station in order to account for subtle differences in signal chain and installation variables. The forced triggers show a flat power spectral density over the primary passband of our hardware response, confirming that they are mostly dominated by thermal fluctuations. Therefore, the pipeline will be most efficient at reconstructing events that resemble the signal-less noise environment.

As seen in Figure 1, for this analysis, we construct a mosaic image of the data of the four phased array antennas, those being used for the deep low threshold trigger. It is also possible to run a 3D version of this network, where one can layer an arbitrary number of channels on top of each other in a tensor, and have the kernel run over the three dimensional space. However, while this reveals a greater sensitivity to more varying trace content as seen in Figure 2, it is currently too time consuming for real-time deployment.

3. Anomalies

The autoencoder-decoder pipeline is designed to identify events or periods in which "unusual" signals appear, i.e. anything other than the environmental noise fluctuations from the forced triggers that the network was trained on. This includes not only neutrino and cosmic ray signals, but also human-made and astrophysical backgrounds. These are also important to understand, and can be used to help monitor and calibrate our instrument.

3.1 Anthropogenic Noise and Hardware Glitches

One common type of background in events is the presence of continuous wave (CW) signals. These signals are typically human-made radio communications transmitted on a single channel

or frequency. Addressing this contamination can be accomplished using a notch filter, either in hardware or software, targeting known CW frequencies. However, this approach filters out not only the CW but also the information from those frequencies in all other data, including clean signals. If we only wanted to exclude one frequency, this would not have a large impact. But if we employ this method for all observed frequencies we quickly start losing a significant amount of information. Notch filter implementations are also complex, with trade-offs in stability and response shape. Alternatively, a basic Convolutional Neural Network (CNN) can be used for CW classification and discrimination [5]. In this analysis, since strong CW events are highly anomalous, we remove them by using a simple function that checks the Fourier space for an excessively high concentration in any single frequency bin. It's worth noting that CW signals can be advantageous in some cases, such as when weather balloons are launched from Summit Station, transmitting at 403 MHz. By accessing GPS data from these weather balloons, their sounding can serve as a calibration tool for the stations [6].

Snowmobiles can also cause impulsive events when they drive by our stations. However, Summit Station does not keep daily logs of what the staff does every day, nor when they travel. Fortunately, once a month they do take part in the "IceSat Traverse", a part of NASA's IceSat-2 [7] mission where Summit Station personnel take snow depth measurements in preparation for the satellite to pass overhead. In these instances, we do have GPS data of where and when the snowmobiles passed. The anomaly detection pipeline is able to pick up on these events, both helping to understand what snowmobile triggered events look like, and when combined with the GPS data, can also be used as another calibration source similar to the weather balloons.

Anomaly detection also identifies events where unexpected failure modes happen, known as *glitches*. This allows for direct investigation and to rapidly respond to previously unknown issues that have now been remedied, such as offsets caused by instabilities in station power or sample scrambling while digitizing.

3.2 Solar Radio Bursts

Another impulsive background is of astrophysical origin: the Sun. Solar Radio Bursts occur when the the radio emission from the Sun is elevated over its typical background, lasting from seconds to minutes as a combination of impulsive and continuum events determined by the underlying physical processes [8]. Identified by anomaly detection and initially referred to as "anomalous low threshold events", it was later verified with e-Callisto data [9] that we had observed a strong Type III burst on September 29th, 2022, concurrent in all operating stations.

Given that this background is only likely to increase as we approach a solar maximum, a 'Solar Model' was trained using a butterworth filter up to 250 MHz on the same training set of data, as Solar Radio Bursts tend to be strongest in the lower frequencies. By removing the higher frequency content, we were able to identify other solar bursts in the 2022 data, as seen in Figure 3. This model is now actively being used to find more examples of this kind of event, because while it could be considered a contamination, it again can be used as a calibration source across the entire array [10].

3.3 Cosmic Ray Air Showers

Until now, all analysis has been conducted using only the four deep phased array Vpol Antennas. However, if one takes the same network structure and instead train on a mosaic of the 9 LPDAs,

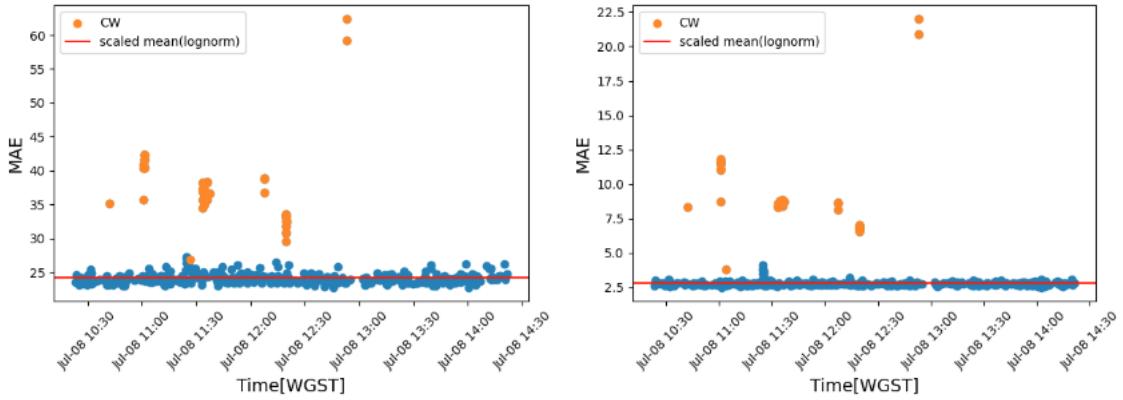


Figure 3: MAE event scores for the same period. Left: the standard forced trigger model is applied. Right: The encoder-decoder network is trained with a butterworth filter up to 250 MHz, in order to be more sensitive to Solar Radio Bursts, which are strongest in the kHz regime and up to this limit. Comparing the two models, the burst on July 8th just before 11:30 becomes easily identifiable.

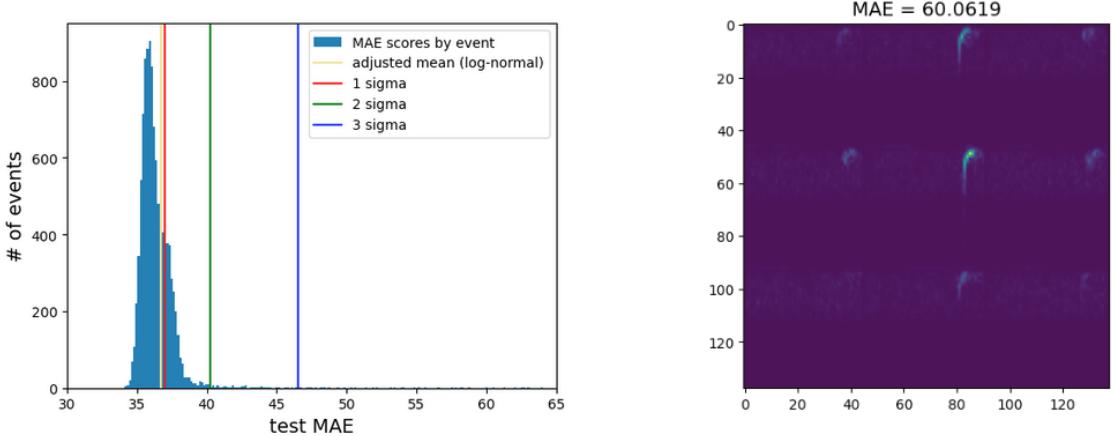


Figure 4: Left: One month distribution of shallow channel anomaly scores for station 23 "Ukaliatsiaq". Right: Mosaic STFT of the 9 LPDA channels for the *Armstrong* cosmic ray air shower candidate [11]. The MAE value >60 as determined by the Shallow Model is highly anomalous.

one can run anomaly detection on the shallow channels. The surface is generally a much noisier environment, due to high wind periods and other anthropogenic effects. Work is ongoing to better understand and classify these features. Even so, when running the anomaly detection pipeline, our cosmic ray air shower candidate *Armstrong* [11], as seen in Figure 4, has one of the highest anomaly values (>60) in the sample.

4. Neutrino Precision

In order to determine the effectiveness of the anomaly detection pipeline on actual neutrino signal, the same simulation set as described in [3] is used. However, while it is relatively simple

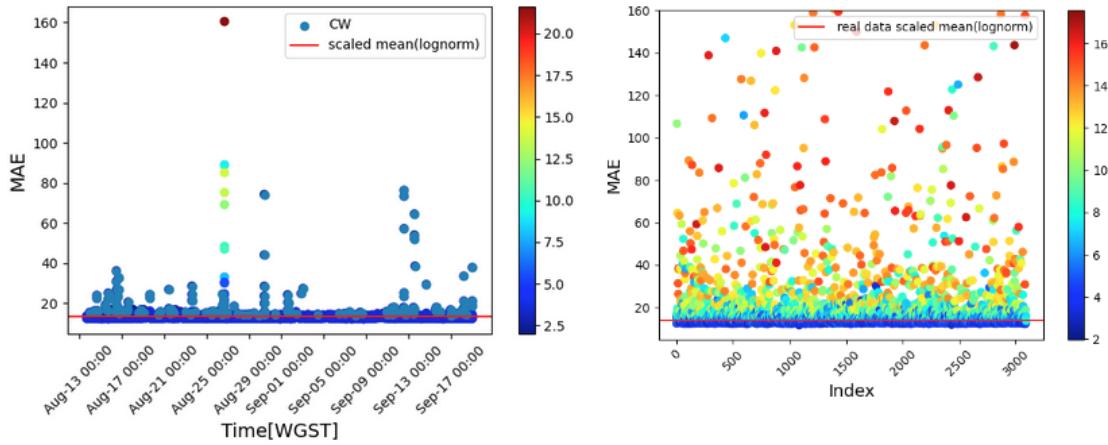


Figure 5: Mean Absolute Error (MAE) anomaly scores from signal in the phased array on an event by event basis. Left: Sample of ~ 1 month of data from Station 23 "Ukaliatsiaq". Events containing strong CW are labeled. Right: A simulation set of 3092 neutrino events. The simulations have realistic channel by channel full signal chain noise added as measured in 2022 data. The colorbar shows the signal to noise ratio of the events.

to simulate the neutrinos, it is challenging to simulate all the different hardware response effects of our signal chain. Therefore, we add the correct cable delays and use measured ADC counts to calculate the relative average gain on a channel by channel basis in order to have the most realistic simulation set that mirrors our current detector status. Only the deep channels are simulated.

We use one month of transferred Station 23 data from the 2022 season (68,500 triggered events), and inject over 3000 simulated neutrinos into the sample with an energy range of 10^{16} to 10^{22} eV (astrophysical + GZK), with the most events in the 10^{17} decade. The MAE scores of the real data and the simulations are shown in Figure 5. If we make a cut on the adjusted mean of the real data MAE distribution, we can remove 58% of the noise traces while retaining 76% of the neutrino sample. This results in a more neutrino pure sample, and doubles our effective transfer rate. We could make a more aggressive cut, for example to maximize both precision and efficiency as seen in Figure 6. However, this would cut away too much data, such that we would not be able to functionally monitor the stations. Furthermore, as a discovery experiment, we only expect on the order of one neutrino per year upon full operation, so we want to minimize any neutrino false negatives. In the event that a neutrino candidate does not make the selection cut and the event is not transferred south, it would still be accessible via deep neutrino searches when the full dataset is recovered from Summit Station.

5. Conclusion

RNO-G is an operational neutrino telescope located at Summit Station, Greenland. As the detector nears completion, data rates will continue to increase and fast filtering techniques must be applied in order to give us the best chance to quickly make a discovery and to actively participate in the Multi-Messenger astronomical community. An autoencoder-decoder based anomaly detection pipeline is under development that helps discover and mitigate anthropogenic and astrophysical

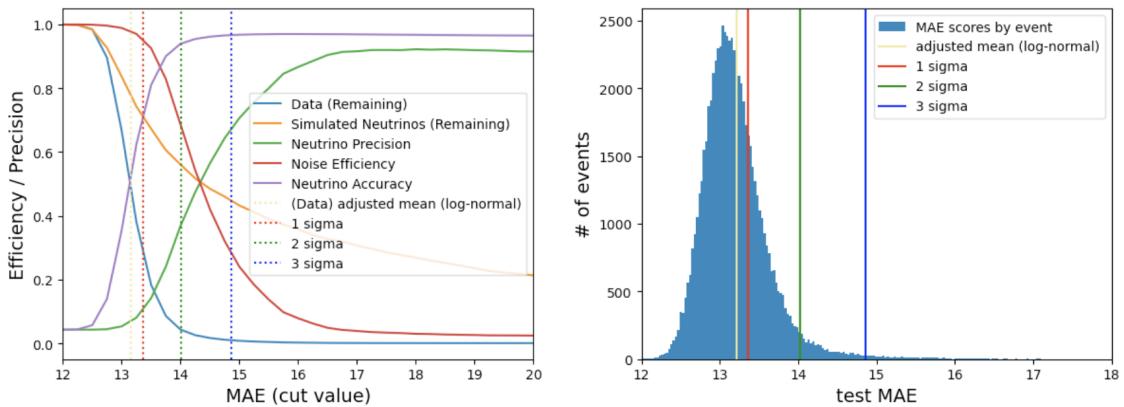


Figure 6: Left: Efficiency, Accuracy and Precision plotted against cut value for MAE. Right: The underlying log-normal distribution of MAE scores for our test sample data.

backgrounds, monitor the detector health, and detect neutrino signal. When the same network structure is trained on the 9 shallow antennas, cosmic ray air shower candidates are also shown to be highly anomalous. Many improvements, such as efficiently running in 3D, further refining the simple network architecture, and training new models in real time can still be made. Currently, this single metric can be used to help prioritize which events to transfer by satellite for further analysis, and in the future can assist with our real-time response strategy.

References

- [1] **RNO-G** Collaboration, J. A. Aguilar *et al.* *JINST* **16** no. 03, (2021) P03025. [Erratum: *JINST* **18**, E03001 (2023)].
- [2] **RNO-G** Collaboration, J. A. Aguilar *et al.* *The European Physical Journal C* **82** no. 2, (2022) 147.
- [3] I. Plaisier, S. Bouma, and A. Nelles *The European Physical Journal C* **83** no. 5, (2023) 443.
- [4] **RNO-G** Collaboration, A. Coleman *PoS ICRC2023* (these proceedings) 1043.
- [5] Z. S. Meyers for the RNO-G Collaboration, *Data Analysis for the Radio Neutrino Observatory Greenland (RNO-G)*. Zenodo, Jun, 2022.
- [6] **RNO-G** Collaboration, B. Oeyen *PoS ICRC2023* (these proceedings) 1042.
- [7] T. Markus *et al.* *Remote sensing of environment* **190** (2017) 260–273.
- [8] G. A. Dulk *Annual review of astronomy and astrophysics* **23** no. 1, (1985) 169–224.
- [9] A. Benz *et al.* *Earth, Moon, and Planets* **104** (2009) 277–285.
- [10] **RNO-G** Collaboration, S. Hallmann *PoS ICRC2023* (these proceedings) 1043.
- [11] **RNO-G** Collaboration, J. Henrichs *PoS ICRC2023* (these proceedings) 259.

Full Author List: RNO-G Collaboration

J. A. Aguilar¹, P. Allison², D. Besson³, A. Bishop¹⁰, O. Botner⁴, S. Bouma⁵, S. Buitink⁶, W. Castiglioni⁸, M. Cataldo⁵, B. A. Clark⁷, A. Coleman⁴, K. Couberly³, P. Dasgupta¹, S. de Kockere⁹, K. D. de Vries⁹, C. Deaconu⁸, M. A. DuVernois¹⁰, A. Eimer⁵, C. Glaser⁴, T. Glüsenkamp⁴, A. Hallgren⁴, S. Hallmann¹¹, J. C. Hanson¹², B. Hendricks¹⁴, J. Henrichs^{11,5}, N. Heyer⁴, C. Hornhuber³, K. Hughes⁸, T. Karg¹¹, A. Karle¹⁰, J. L. Kelley¹⁰, M. Korntheuer¹, M. Kowalski^{11,15}, I. Kravchenko¹⁶, R. Krebs¹⁴, R. Lahmann⁵, P. Lehmann⁵, U. Latif⁹, P. Laub⁵, C.-H. Liu¹⁶, J. Mammo¹⁶, M. J. Marsee¹⁷, Z. S. Meyers^{11,5}, M. Mikhailova³, K. Michaels⁸, K. Mulrey¹³, M. Muzio¹⁴, A. Nelles^{11,5}, A. Novikov¹⁹, A. Nozdrina³, E. Oberla⁸, B. Oeyen¹⁸, I. Plaisier^{5,11}, N. Punsuebsay¹⁹, L. Pyras^{11,5}, D. Ryckbosch¹⁸, F. Schlüter¹, O. Scholten^{9,20}, D. Seckel¹⁹, M. F. H. Seikh³, D. Smith⁸, J. Stoffels⁹, D. Southall⁸, K. Terveer⁵, S. Toscano¹, D. Tosi¹⁰, D. J. Van Den Broeck^{9,6}, N. van Eijndhoven⁹, A. G. Vieregg⁸, J. Z. Vischer⁵, C. Welling⁸, D. R. Williams¹⁷, S. Wissel¹⁴, R. Young³, A. Zink⁵

¹ Université Libre de Bruxelles, Science Faculty CP230, B-1050 Brussels, Belgium

² Dept. of Physics, Center for Cosmology and AstroParticle Physics, Ohio State University, Columbus, OH 43210, USA

³ University of Kansas, Dept. of Physics and Astronomy, Lawrence, KS 66045, USA

⁴ Uppsala University, Dept. of Physics and Astronomy, Uppsala, SE-752 37, Sweden

⁵ Erlangen Center for Astroparticle Physics (ECAP), Friedrich-Alexander-Universität Erlangen-Nürnberg, 91058 Erlangen, Germany

⁶ Vrije Universiteit Brussel, Astrophysical Institute, Pleinlaan 2, 1050 Brussels, Belgium

⁷ Department of Physics, University of Maryland, College Park, MD 20742, USA

⁸ Dept. of Physics, Enrico Fermi Inst., Kavli Inst. for Cosmological Physics, University of Chicago, Chicago, IL 60637, USA

⁹ Vrije Universiteit Brussel, Dienst ELEM, B-1050 Brussels, Belgium

¹⁰ Wisconsin IceCube Particle Astrophysics Center (WIPAC) and Dept. of Physics, University of Wisconsin-Madison, Madison, WI 53703, USA

¹¹ Deutsches Elektronen-Synchrotron DESY, Platanenallee 6, 15738 Zeuthen, Germany

¹² Whittier College, Whittier, CA 90602, USA

¹³ Dept. of Astrophysics/IMAPP, Radboud University, PO Box 9010, 6500 GL, The Netherlands

¹⁴ Dept. of Physics, Dept. of Astronomy & Astrophysics, Penn State University, University Park, PA 16802, USA

¹⁵ Institut für Physik, Humboldt-Universität zu Berlin, 12489 Berlin, Germany

¹⁶ Dept. of Physics and Astronomy, Univ. of Nebraska-Lincoln, NE, 68588, USA

¹⁷ Dept. of Physics and Astronomy, University of Alabama, Tuscaloosa, AL 35487, USA

¹⁸ Ghent University, Dept. of Physics and Astronomy, B-9000 Gent, Belgium

¹⁹ Dept. of Physics and Astronomy, University of Delaware, Newark, DE 19716, USA

²⁰ Kapteyn Institute, University of Groningen, Groningen, The Netherlands

Acknowledgments

We are thankful to the staff at Summit Station for supporting our deployment work in every way possible. We also acknowledge our colleagues from the British Antarctic Survey for embarking on the journey of building and operating the BigRAID drill for our project. We would like to acknowledge our home institutions and funding agencies for supporting the RNO-G work; in particular the Belgian Funds for Scientific Research (FRS-FNRS and FWO) and the FWO programme for International Research Infrastructure (IRI), the National Science Foundation (NSF Award IDs 2118315, 2112352, 211232, 2111410) and the IceCube EPSCoR Initiative (Award ID 2019597), the German research foundation (DFG, Grant NE 2031/2-1), the Helmholtz Association (Initiative and Networking Fund, W2/W3 Program), the University of Chicago Research Computing Center, and the European Research Council under the European Unions Horizon 2020 research and innovation programme (grant agreement No 805486).