

VAE-based latent-space classification of RNO-G data

Thorsten Glüsenkamp^{a,*} for the RNO-G Collaboration

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

^a*Uppsala University, Lägerhyddsvägen 1, 751 20 Uppsala, Sweden*

E-mail: thorsten.gluesenkamp@physics.uu.se

The Radio Neutrino Observatory in Greenland (RNO-G) is a radio-based ultra-high energy neutrino detector located at Summit Station, Greenland. It is still being constructed, with 7 stations currently operational. Neutrino detection works by measuring Askaryan radiation produced by neutrino-nucleon interactions. A neutrino candidate must be found amidst other backgrounds which are recorded at much higher rates—including cosmic-rays and anthropogenic noise—the origins of which are sometimes unknown. Here we describe a method to classify different noise classes using the latent space of a variational autoencoder. The latent space forms a compact representation that makes classification tractable. We analyze data from a noisy and a silent station. The method automatically detects and allows us to qualitatively separate multiple event classes, including physical wind-induced signals, for both the noisy and the quiet station.

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



*Speaker

1. Introduction

In the last decade, the field of high-energy neutrino astronomy has been established. At energies between 100 GeV and a few PeV, astrophysical neutrinos have been detected as extragalactic diffuse flux[1], from the galaxy as galactic diffuse emission[2] and from a few AGN as point source emission[3]. However, because the neutrino flux is steeply falling towards higher energies, current detectors like IceCube are not able to detect significant amounts of neutrinos above 10 PeV over its lifetime, let alone 100 PeV. The only viable way is to build a detector with a much bigger effective volume. One such effort is RNO-G (Radio Neutrino Observatory Greenland)[4], which is currently being constructed in Greenland. It measures high energy neutrinos via Askaryan radiation originating from the particle cascade following their interactions in the ice. The threshold for detection is about 100 PeV and extends into the EeV range, which allows it to complement current neutrino facilities to the highest energies, in particular for cosmogenic neutrino fluxes[5].

In order for RNO-G to make discoveries, it is vital to detect and understand different types of background events. In this contribution, we study the unsupervised detection of event classes in the shallow antennas of a RNO-G station in a 2-step process. First we, compress the data into a latent representation using a variational autoencoder (VAE)[6]. Second, we cluster the data in the latent space and search for distinct classes. For the study we use the shallow component of two stations: station 21, which has more noise from electronics and is closer to human activity, and station 23, which has less electronically induced noise and is more remote.

2. Station and data

RNO-G currently consists of seven stations, spread out within a few kilometers close to Summit Station, Greenland. A single station consists of 24 antennas, of which nine are located just below the surface. In this contribution, we will only focus on the shallow antennas. These are log-periodic dipole arrays (LPDAs), which are spread out in groups of three around the central data acquisition box. Three of the nine antennas are upwards facing, and six are downward facing. A threshold trigger algorithm is run separately on the downward facing and upward facing antennas, whose threshold is tuned to be 1 Hz. Once a trigger happens, an event consists of a 2048 samples long antenna trace for each antenna with a time resolution of roughly 0.3 ns. If there is more noise in the environment, the threshold is increased to keep the trigger rate constant. Additionally, there are fixed readouts of the whole detector in 10 second intervals. We focus on two stations. The first is station 21, which we call "noisy" station, and is located about one kilometer from the summit camp. The second is station 23, which we call "silent" station, as it is located about 4 km away from summit camp. The noisy station triggers significantly more events from human interference, and is also sensitive to background from battery charging and a long range wide area network (LoRaWAN) noise.

3. The clustering architecture

We perform the clustering as a two step process. In the first step, we employ a variational autoencoder (VAE) [6] to compress the antenna data of the shallow array (9 antennas) into a low-dimensional latent representation. A visual summary of the overall VAE processing is depicted in

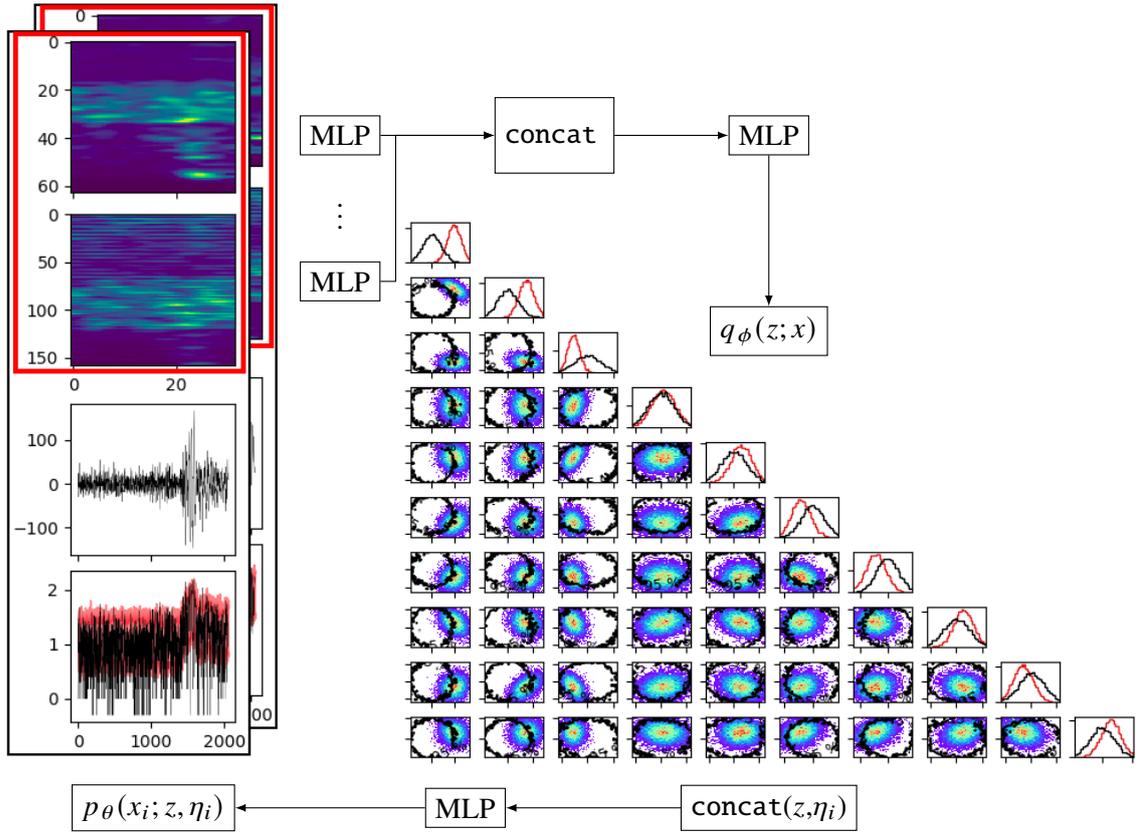


Figure 1: Overview of the VAE. The left part indicates a summary of antenna data: the upper two plots (red boxes) show the scattering coefficients of the voltage trace (see text), the lower two show the voltage trace of the antenna in linear scale and as the logarithm of the absolute value. The processing of the data moves clockwise following the arrows: the scattering coefficients for each antenna i are fed to a multi-layer perceptron (MLP) that is shared among all antennas. The output of all MLPs is concatenated and fed into a second MLP that predicts the parameters of the 10-d latent space posterior distribution $q_\phi(z; x)$. An example posterior distribution is shown by visualizing its probability mass projection in 2-d (histogram) and 1-d (red), i.e. plotting its marginal distributions. The prior marginal distributions are also shown (black) and in 2-d are indicated by their 95% contour. The data distribution $p_\theta(x_i; z, \eta_i)$ for a given antenna i is predicted by another MLP that takes as input a latent coordinate z concatenated with antenna orientation parameters η_i . The 95% envelope of samples from $p_\theta(x_i; z, \eta_i)$ is indicated in red on top of the logarithm of the absolute value of the data trace, which is the target space used for prediction.

fig. 1. The VAE broadly consists of two distributions: a stochastic mapping from the antenna data to latent space, in the following called the latent posterior distribution $q_\phi(z; x)$, and a stochastic mapping from latent to antenna data, in the following called the reconstruction term $p_\theta(x; z)$. As described further below and as depicted in fig. 1, we actually use different representations of the antenna data for the forward and backward mapping. The parameters ϕ and θ are the parameters of multi-layer perceptrons[7] (MLPs) that map from the conditional input to the parameters of the respective distribution. In order to learn $q_\phi(z; x)$ and $p_\theta(x; z)$, we minimize the negative evidence

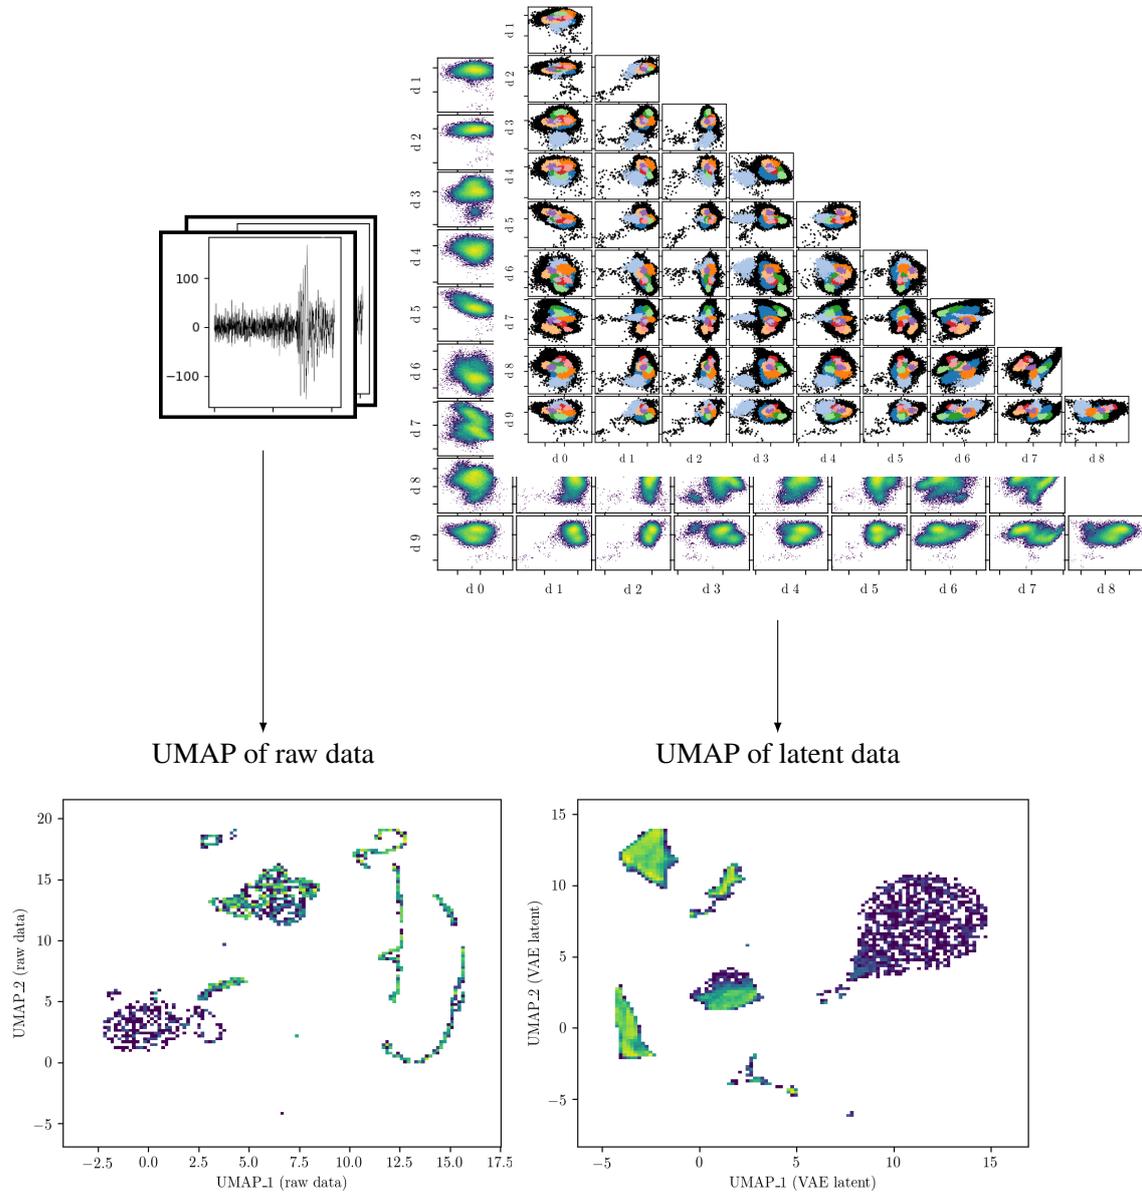


Figure 2: Overview of the clustering step. Every datum that is triggered with a voltage threshold trigger is projected into the latent space by the mean of its posterior distribution (upper right). These projections form a point cloud in the latent space onto which the clustering algorithm H-DBSCAN is applied (colored points). A UMAP projection of the raw data and of the latent representation are shown in the bottom.

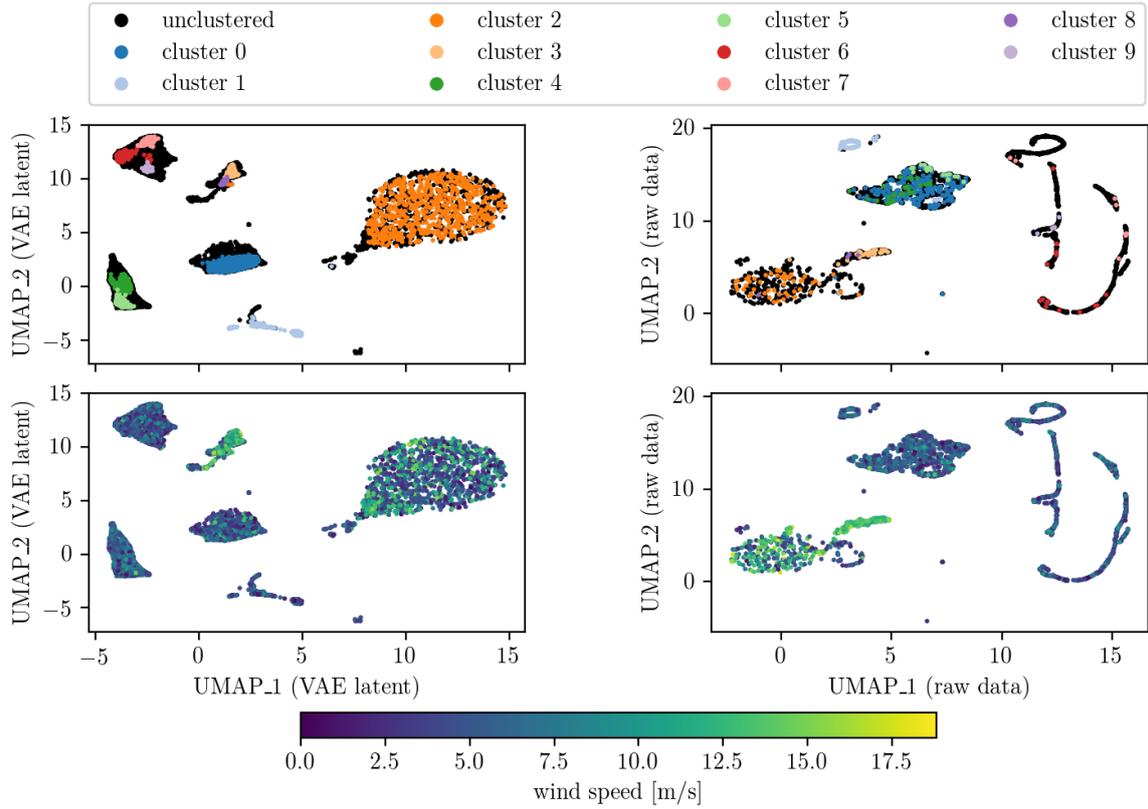


Figure 3: UMAP embeddings of (noisy) station 21 events. The left column shows a VAE latent space embedding, the right column an embedding of the raw antenna data of a smaller sub-sample (20000 events). The upper row shows corresponding clustering labels, the lower row overlays the wind speed measured simultaneously to an event.

lower bound (ELBO) as a loss function \mathcal{L} .

$$\mathcal{L} = -E_x \left[\underbrace{\int_z q_\phi(z; x) \cdot \left[\ln p_\theta(x; z) - \ln \frac{q_\phi(z; x)}{\mathcal{N}(z; \mathbf{0}, \mathbb{1})} \right] dz}_{\text{evidence lower bound (ELBO)}} \right] \quad (1)$$

The total reconstruction term $p_\theta(x; z)$ is defined as

$$p_\theta(x; z) = \prod_{i=1}^9 p_\theta(x_i; z, \eta_i), \quad (2)$$

which is a product of nine probability distributions (PDFs), one for each antenna. Each individual term is modeled via a 2048 dimensional Gaussian with diagonal covariance and a shared variance across all dimensions to model a probability distribution over the observable space. Instead of the normal antenna voltage trace, we predict the logarithm of the absolute trace, which is easier to model. In addition to the latent input, we also add position and rotation information η_i of a given antenna to describe the individual antenna responses. The target latent space z is 10-dimensional

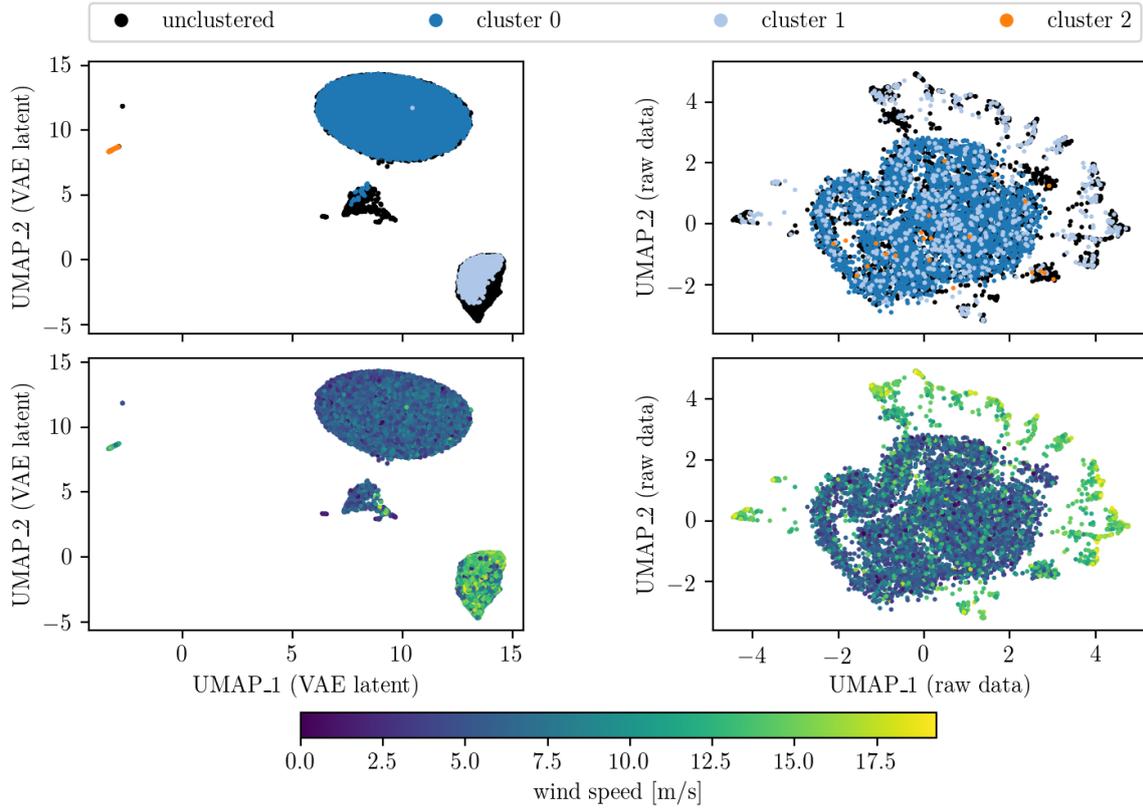


Figure 4: UMAP embeddings of (silent) station 23 events. The left column shows a VAE latent space embedding, the right column an embedding of the raw antenna data of a smaller sub-sample (20000 events). The upper row shows corresponding clustering labels, the lower row overlays the wind speed measured simultaneously to an event.

for all experiments. The MLP that predicts the parameters of $q_\phi(z; x)$ takes as input the antenna data in a frequency representation of the voltage trace. We use a scattering transform [8] for this frequency representation, because it allows us to consider time-dependence. We omit details for brevity but refer to [8] for details. The output of the scattering transform are a set of coefficients, starting at 0th order, and going up to a maximal order defined by the user. We choose to go to a maximum order of two, which for our setting yields 16 0th order coefficients, 1264 1st order coefficients, and 3840 2nd order coefficients. These are flattened to serve as a 5120 dimensional input to a MLP. This MLP is applied to all nine antennas and the output is concatenated and mapped via another MLP to $q_\phi(z; x)$. The prior distribution in the VAE is a standard normal distribution.

After training the VAE we perform clustering in the latent space. First, we generate points in the latent space by calculating the sample mean of the posterior distribution associated to each data point that has a voltage threshold trigger. The threshold trigger condition is used to filter out data that is just thermal noise, which mostly comes from other triggers. It is important to emphasize that the VAE training was performed on all triggered events, including an interval trigger that just takes data in fixed intervals. We then use H-DBSCAN [9], a common density based clustering algorithm, to cluster those physics-triggered events into N clusters in the latent space, where N is automatically determined by the algorithm. A visualization of the latent space and a colored representation

class description	cluster ids	
	station 21	station 23
thermal noise	2	0
wind induced	2,3,8	1,2
battery charging (+ CW)	6,7,9	
LoRaWAN (+ CW)	0,4,5	
human made (CW / pulsing)	1	

Table 1: Overview of the different event classes found in triggered data. The cluster ids correspond to the cluster ids in fig. 3 and fig. 4. CW stands for "continuous wave" signal, i.e. a single frequency signal.

including the associated cluster labels is shown in fig. 2 in the upper right. The clustering algorithm has a few tunable parameters, for which we take the standard settings. The only one we fix is the minimum cluster size, which we set to 30, and the final cluster selection method, which we set to leaf. Some data points also can not be associated to any cluster, which we call "unclustered". Finally, we apply UMAP [10] to a subsample of the raw antenna data (fig. 2 left) and to the latent representation (fig. 2 right) in order to obtain 2-d visualizations of the high-dimensional input. UMAP preserves local manifold structure and allows to qualitatively judge the clustering performance by super-imposing the cluster labels. Having two different UMAP projections allows us to better judge the performance of the VAE.

4. Application to RNO-G data and results

For both stations we train the setup on roughly 200000 triggered events taken between July and October in 2022. After applying the VAE, the clustering algorithm, and the UMAP projection, we obtain the visualizations in fig. 3 and fig. 4 for the noisy and silent station, respectively. In the noisy station the algorithm finds 10 clusters. To understand these clusters, we use the event time distributions, the measured wind speed and a visual inspection of individual events. Clusters 0, 4 and 5 can be identified as LoRaWAN noise, sometimes overlaid with continuous wave signals. Clusters 6, 7 and 9 can be identified as battery charging noise, which appears as a 1-d structure in the raw-data UMAP projection (fig. 3 top right). This comes from the fact that the underlying noise signal is always similar, just shifted due to a random triggering offset. There are two more 1-d periodic clusters which correspond to periodic trace signals from human communication (cluster 1) and the same but overlaid with LoRaWAN noise (cluster 0). Because the trigger threshold is rather high in this station, the thermal noise event rate of threshold-triggered events is rather low. We identify cluster 2 to partially contain noise events. Finally, clusters 3, 8 and also in parts cluster 2 are wind-induced events, which can be seen with the association to wind speeds (lower row in fig. 3).

For the silent station the algorithm detects only three clusters (fig. 4). In contrast to the noisy station there is much more thermal noise (cluster 0), because the trigger threshold is lower. Additionally, the algorithm identifies wind-events (cluster 1, 2), which can be associated due to higher wind speeds. Interestingly, in the raw-data UMAP representation, high wind speeds are situated further out in proportion to the speed, which can be seen in the lower right plot in fig. 4.

This structure is not seen in the VAE-based projection, and suggests that the VAE latent space representation still has lots of potential for improvement.

An overview of the classes and associated clusters is summarized in table 1. Ideally, the clustering settings are such that each class is associated only to a single cluster. We chose the Leaf option in H-DBSCAN, which usually captures all relevant clusters, with the downside that it sometimes splits a class into multiple subclasses. A better latent representation and finetuning of the clustering algorithm should help to obtain more quantitative results in this regard, which hopefully leads to better 1-to-1 class association in the future.

5. Conclusion and outlook

In this contribution, we applied a variational autoencoder to RNO-G shallow station data of a silent and noisy station. Subsequently, we performed a clustering in latent space using H-DBSCAN. UMAP projections of the final outcome show that the clustering makes sense, and identifies major event classes, including wind-induced background events, even on the noisy station. However, the results are qualitative, as we had to use loose clustering settings in order to capture all classes. Nonetheless, the results are promising. A better latent representation and a tuning of clustering parameters are avenues for further research to obtain a better classification, and thereby a better understanding of the detector and its various background events. The method has the potential to be developed into an online monitoring tool in the future.

References

- [1] The IceCube Collaboration *Physical Review Letters* **104** no. 2, (July, 2021) 022002.
- [2] The IceCube Collaboration *Science* **380** no. 6652, (2023) 1338–1343.
- [3] The IceCube Collaboration *Science* **378** no. 6619, (Nov., 2022) 538–543.
- [4] The RNO-G collaboration *Journal of Instrumentation* **16** no. 03, (Mar, 2021) P03025.
- [5] T. Stanev *Comptes Rendus Physique* **15** no. 4, (2014) 349–356. Ultra-high-energy cosmic rays: From the ankle to the tip of the spectrum.
- [6] D. Kingma and M. Welling, “Auto-Encoding Variational Bayes,” in *International Conference on Learning Representations*, ICLR. 2014. [arXiv:arXiv:1312.6114](https://arxiv.org/abs/1312.6114).
- [7] T. Hastie, R. Tibshirani, and J. Friedman, *The elements of statistical learning: data mining, inference and prediction*. Springer, 2 ed., 2009.
- [8] J. Bruna and S. Mallat *IEEE Transactions on Pattern Analysis and Machine Intelligence* **35** no. 08, (Aug, 2013) 1872–1886.
- [9] L. McInnes, J. Healy, and S. Astels *Journal of Open Source Software* **2** no. 11, (2017) 205.
- [10] L. McInnes, J. Healy, N. Saul, and L. Großberger *Journal of Open Source Software* **3** no. 29, (2018) 861.

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. Glüsenskamp⁴, 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. Viereggs⁸, 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, the European Research Council under the European Unions Horizon 2020 research and innovation programme (grant agreement No 805486) and the Carl Tryggers foundation.