





A neural-network framework for tracking and identifying cosmic-ray nuclei in the RadMap Telescope

Luise Meyer-Hetling^{1,2,*} , Martin J. Losekamm^{1,2,3} , Stephan Paul^{1,2} , and Thomas Pöschl⁴ 

¹ Technical University of Munich, School of Natural Sciences, Garching, Germany

² Excellence Cluster ORIGINS, Garching, Germany

³ European Space Agency, Noordwijk, Netherlands

⁴ European Organization for Nuclear Research (CERN), Geneva, Switzerland

Received 22 July 2025 / Accepted 21 January 2026

Abstract – The detailed characterization of the radiation environment aboard spacecraft is a prerequisite for assessing shielding requirements and for minimizing the exposure of crew and equipment during future deep-space missions. The scintillating-fiber tracking calorimeter at the heart of the RadMap Telescope is designed for detailed studies of cosmic rays within the resource constraints of an operational radiation monitor. We present a neural-network framework that can reconstruct the properties of cosmic-ray nuclei traversing the instrument. Employing the GEANT4 simulation toolkit and a simplified model of the detector to generate training and test data, we achieve the spectroscopic capabilities required for an accurate determination of the biologically relevant dose that astronauts receive in space. We can reconstruct the trajectory of a particle with an angular resolution of better than 1.4° and achieve a charge separation of better than 95% for nuclei with $Z \leq 8$; specifically, we reach an accuracy of 99.8% for hydrogen. The energy resolution is $< 20\%$ for energies below 1 GeV/n and elements up to iron. We also discuss the limitations of our detector, the reconstruction framework, and this feasibility study, as well as possible improvements.

Keywords: Cosmic-ray nuclei / Radiation monitoring / Tracking calorimeter / Neural networks

1 Introduction

Mitigating the effects of exposure to cosmic and solar radiation on the human body is one of the major challenges of future missions to the Moon, Mars, and other deep-space destinations (Chancellor et al., 2014). Detailed knowledge of the radiation environment in space – including its spectral, temporal, and spatial variations – is a prerequisite for developing medical and operational radiation-protection measures (Montesinos et al., 2021). It is also essential for the design of new spacecraft, habitats, surface vehicles, and spacesuits (Durante & Cucinotta, 2011; Barthel & Sarigul-Klijn, 2019).

Multiple instruments have provided, or are still providing, data on the radiation environment in deep space, for example in lunar orbit (Looper et al., 2020; George et al., 2024), on the surface of the Moon (Wimmer-Schweingruber et al., 2020), and on the surface of Mars (Hassler et al., 2014). Long-term observations have also been made by radiation monitors on the International Space Station (ISS) (e.g., Stoffle et al., 2015; Zeitlin et al., 2023), and by space-weather observatories and astrophysics experiments (e.g., Rodríguez-Pacheco et al.,

2020; Aguilar et al., 2021). Despite these and other valuable measurements, we still lack comprehensive long-term observations with instruments capable of resolving the various components of the deep-space radiation environment.

We also have an incomplete understanding of how the complex radiation field in space interacts with spacecraft shielding and the human organism (Chancellor et al., 2018; Walsh et al., 2019; Vozenin et al., 2024). We know, however, that exposure to it results in short- and long-term health risks, including cancer (Chancellor et al., 2014; National Academies of Sciences, Engineering, and Medicine, 2021; Guo et al., 2022) and cardiovascular diseases (Delp et al., 2016; Rikhi et al., 2020). Recent studies suggest that (prolonged) exposure may also permanently impair cognitive functions such as learning and memory (Parihar et al., 2015; Klein et al., 2021; Alagband et al., 2023). Altogether, health risks due to radiation exposure may be one of the most limiting factors in determining the maximum permissible length of deep-space missions (Cucinotta et al., 2015).

We are developing a new class of radiation monitors to help address the shortage of in-situ data on the space radiation environment (Walsh et al., 2019; Fogtman et al., 2023). The charged-particle detectors at the heart of these instruments are

*Corresponding author: luise.meyer-hetling@tum.de

capable of differentiating between nuclei of varying biological effectiveness – which to leading order depends on their nuclear charge, Z – and can measure their energy. The first instrument to operationally demonstrate our detectors is the RadMap Telescope (Losekamm et al., 2021), which so far has collected data on the radiation environment inside the ISS between April 2023 and January 2024 (Losekamm et al., 2023, 2025).

In this article, we present a neural-network-based framework for analyzing the data generated by the RadMap Telescope. We explain the challenges of interpreting the measurements of our detector and how the architecture of the framework addresses them. The results we show are based on benchmarked simulation data and provide best-case limits on the performance we can expect to achieve. We also discuss the advantages and challenges of relying on deep-learning algorithms, as well as our approaches to tackle the latter.

2 The space radiation environment

The space radiation environment is a complex field of charged particles and atomic nuclei. It is dominated by two major components: galactic cosmic rays (GCR) and solar energetic particles (SEP) 98% of GCR are fully ionized nuclei of all naturally occurring elements, with protons (87%) and helium nuclei (12%) being most abundant; only about 1% are nuclei of heavier elements (Longair, 2012; Rankin et al., 2022). The remaining 2% are mostly electrons and positrons. GCR nuclei have energies between a few MeV per nucleon (MeV/n) and hundreds of EeV. Most relevant to radiation protection is the MeV-to-GeV range. Though the flux of protons is orders of magnitude larger than that of heavier nuclei, the higher biological effectiveness of the latter results in them contributing more than half of the radiation dose astronauts receive (Naito & Kodaira, 2022).

SEP are protons and light nuclei with energies up to several GeV that are released by the Sun in sudden energetic outflows of coronal matter (Reames, 2021, 2022). Such SEP bursts deliver high, quasi-instantaneous doses that can cause acute radiation sickness. In lightly shielded environments, the strongest events may even be fatal (Hellweg et al., 2020; Mao et al., 2021). Understanding the temporal and spatial evolution of SEP bursts is thus crucial for establishing operational procedures to limit the exposure of the crew. Close to Earth and other planets with a magnetic field, particles (mostly protons and electrons) trapped in planetary radiation belts constitute a third major source of radiation exposure.

Uncertainties in the assessment and prediction of the risk posed by exposure to the space radiation environment arise mainly from our limited knowledge of the biological effectiveness of cosmic-ray nuclei (Cucinotta et al., 2017; Chancellor et al., 2018). Another source of uncertainty is that instruments used for operational radiation monitoring, with few exceptions (e.g., Hassler et al., 2012; Looper et al., 2020; Di Fino et al., 2023), cannot resolve the charge, Z , and (kinetic) energy, E_{kin} , of a particle or nucleus (we use the terms interchangeably throughout this article) but instead measure only the linear energy transfer (LET) or, worse, the energy deposition in a planar detector (e.g., a silicon diode). Though the composition of the radiation field can be reconstructed from a recorded LET

or energy-deposition spectrum, the required post-processing suffers from uncertainties associated with assumptions (e.g., about the directionality of the particle flux), models, and transport simulations (Norbury et al., 2017; Slaba et al., 2017). The accurate determination of the biologically relevant dose therefore requires the development of new detector systems with spectroscopic capabilities, especially when relying on the Z and E_{kin} -dependent quality factors now adopted by NASA (Dietze et al., 2013; Cucinotta, 2015). Ideally, new detectors should also have a field of view covering the full solid angle to be sensitive to flux anisotropies.

3 Measurement principle and reconstruction methodology

Instruments employed for the continuous characterization of the radiation environment that astronauts are exposed to must be small and consume little power. In contrast to those built for astrophysical investigations, they are hence often limited in the choice of technologies and in the complexity of their layout. The use of magnetic spectrometers for momentum measurement and charge determination, for example, is often not possible due to mass and power constraints. Other limitations arise because of environmental conditions or safety considerations.

3.1 Detector geometry

We are attempting to overcome the limitations of many current-generation detectors by using a novel detection principle. Figure 1 shows a schematic representation of the main detector of the RadMap Telescope, the Active Detection Unit (ADU). Its radiation-sensitive volume consists of 1024 scintillating-plastic fibers with a square cross-section of $(2 \times 2) \text{ mm}^2$ and a length of 80 mm. They are arranged in a stack of 32 layers of alternating orientation. The layers, and neighboring fibers within them, are spaced about 100 μm apart. The intensity of the scintillation light created in each fiber is, to leading order, proportional to the energy loss of a particle traversing it. We detect the light using silicon photomultipliers (SiPMs) attached to one end of each fiber. This layout results in a detector volume of about $(65 \times 65 \times 70) \text{ mm}^3$ and allows recording the total energy deposition of individual particles and the energy-deposition profile along their tracks with nearly omnidirectional acceptance.

To suppress optical crosstalk, each fiber is sputter-coated with aluminum (with a layer thickness of about 100 nm–150 nm) except at the end where it is read out. This coating, as well as other intrinsic (e.g., self-absorption) and extrinsic (e.g., reflection losses) effects lead to the attenuation of scintillation light and hence to a strongly position-dependent signal amplitude along the fibers (Losekamm et al., 2024). Such losses are, however, not relevant to the work presented here because we do not trace individual photons in our simulation.

Figure 1 also shows the coordinate system and the spherical coordinates $\phi \in [-180^\circ, 180^\circ)$ and $\theta \in [0^\circ, 180^\circ]$ we use to parametrize the orientation and direction of particle tracks. The origin of the coordinate system is at the center of the x and z -dimensions of the detector. We project the signal amplitudes along the two fiber orientations onto the y - x - and y - z -planes (see red and blue arrows in the figure), obtaining two gray-scale images with 16×32 pixels each. For illustration, Figure 2

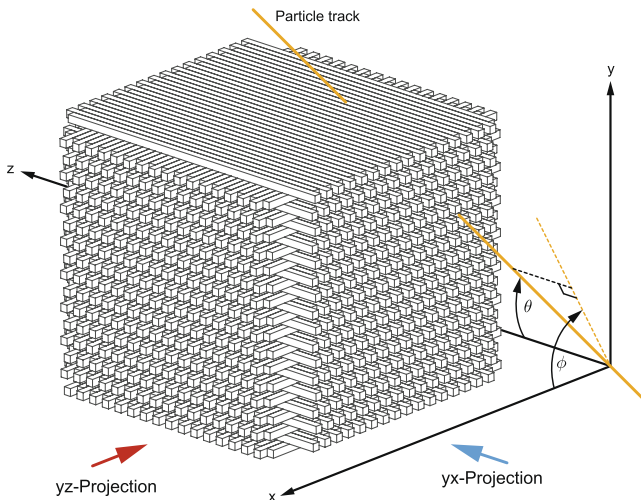


Figure 1. Schematic representation of the main detector of the RadMap Telescope, which consists of a stack of 1024 scintillating plastic fibers. They are arranged in a stack of 32 layers of alternating orientation. Also shown are the coordinate system and the spherical coordinates ϕ and θ we use to parametrize the orientation of particle tracks. The red and blue arrows indicate the two-dimensional projections we use to visualize and analyze the data from the detector.

shows the simulated event signatures (see Section 3.4 below for a description of our GEANT4 setup) of a 540-MeV proton and a 1.5-GeV/n iron nucleus. The latter is much broader due to the emission of high-energy δ electrons and hadronic fragments along the track.

3.2 Measurement principle and challenges

The energy-deposition profile of cosmic-ray nuclei stopping (or losing a significant portion of their energy) in the detector encodes all information required for determining their identity and energy (see Bethe formula in Navas et al., 2024). Gruhn et al. first described the corresponding reconstruction technique, which they called *Bragg curve spectroscopy* (Gruhn et al., 1982). Figure 3 shows how a single measurement of a Bragg curve in an ideal detector (dashed blue curve) yields information about the nuclear charge of a particle, Z , its mass number, A , and its kinetic energy, E_{kin} (the latter being a function of A and the velocity, β). Since there is little difference in the biological effectiveness of isotopes of the same element, only Z and E_{kin} are relevant for radiation dosimetry. E_{kin} can be inferred from Z and β using the average atomic mass. Thus, even for particles passing through the detector (see orange, red, and purple curves in Fig. 3), the charge and velocity dependence of the energy loss contains sufficient information for identifying nuclei and measuring their energy. At higher energies, however, the resolution intrinsically decreases because the velocity dependence becomes less pronounced (see red and purple curves).

There are three effects that complicate this simplified picture. First, the statistical nature of the electronic interactions between nuclei and the detector material causes fluctuations of the formers' energy loss (called *straggling*). For clarity, we show in Figure 3 only the mean energy deposition (dE/dx), of an ensemble of particles; the profiles for individual particles

of the same charge and energy can differ significantly from this average. Second, the spallation of incident nuclei into fragments with lower Z (and consequently smaller energy loss) leads to further variation in the energy-deposition profiles. Nuclei may fragment immediately upon entering the detector, somewhere along their path through it, or not at all. Our particle-identification and energy-reconstruction algorithms must therefore be able to cope with a broad range of potential energy-loss profiles for nuclei with identical Z and E_{kin} . This effect becomes stronger with increasing Z . Third, the scintillation efficiency – i.e., the fraction of energy deposited in a fiber that is converted into detectable light (see solid curves in Fig. 3) – decreases with increasing energy-deposition density due to ionization quenching (Birks, 1951; Pöschl et al., 2021). This effectively results in an upper limit on the measurable energy-deposition density, reducing the separation power for heavier nuclei because of the quadratic Z -dependence of the energy loss.

3.3 Reconstruction methodology

Our goal is to perform an event-by-event analysis of the data gathered by the RadMap Telescope, i.e., to determine the properties of each individual particle detected by the ADU. For this, we need to (1) reconstruct the track of the particle through the detector, (2) determine its charge, and (3) calculate its initial kinetic energy. In the past, we tested multiple approaches for an automated event reconstruction, including a Bayesian particle filter (Losekamm et al., 2017), simulated annealing (van Laarhoven & Aarts, 1987), and a Markov Chain Monte Carlo (Caldwell et al., 2009). Though some of these methods produced acceptable results in a limited parameter space (Hollender, 2019; Pöschl, 2022), their biggest drawback was the computing effort required, which resulted in processing times upwards of 15 min per event (Milde, 2016; Losekamm et al., 2017). Since our objective is to build a system that can operate in real time, such single-particle execution times are orders of magnitude too long.

We thus developed an alternative approach for event reconstruction based on neural networks that we train on simulated interactions of cosmic-ray nuclei with the ADU. Because each of the reconstruction steps depends on the result of the previous one, we developed three sequentially-applied frameworks. For each of them, we use the unaltered projections of the event signatures (see Fig. 2) as input.

The networks at the heart of these frameworks have architectures with a similar basic structure (see Fig. 4). Their main building blocks are convolutional layers, a layer type designed specifically for recognizing patterns in image-like data (Lecun et al., 1998). Convolutions divide the input images into smaller areas and attempt to extract local features – like edges or endpoints – before merging the areas back together to find a global structure – e.g., a straight line. They are well suited for detecting translation- and rotation-invariant patterns, for example the straight lines of particle tracks or the energy-deposition profiles along such tracks. The size of the areas – the filter size – is a key parameter of a convolutional layer, since it determines the initial scale at which the network searches for features. The scale of the track left by a particle in our detector, however, cannot be generally defined because it depends, among other things, on its orientation, its entrance point, and on the properties of the particle. We thus rely on inception layers that combine multiple

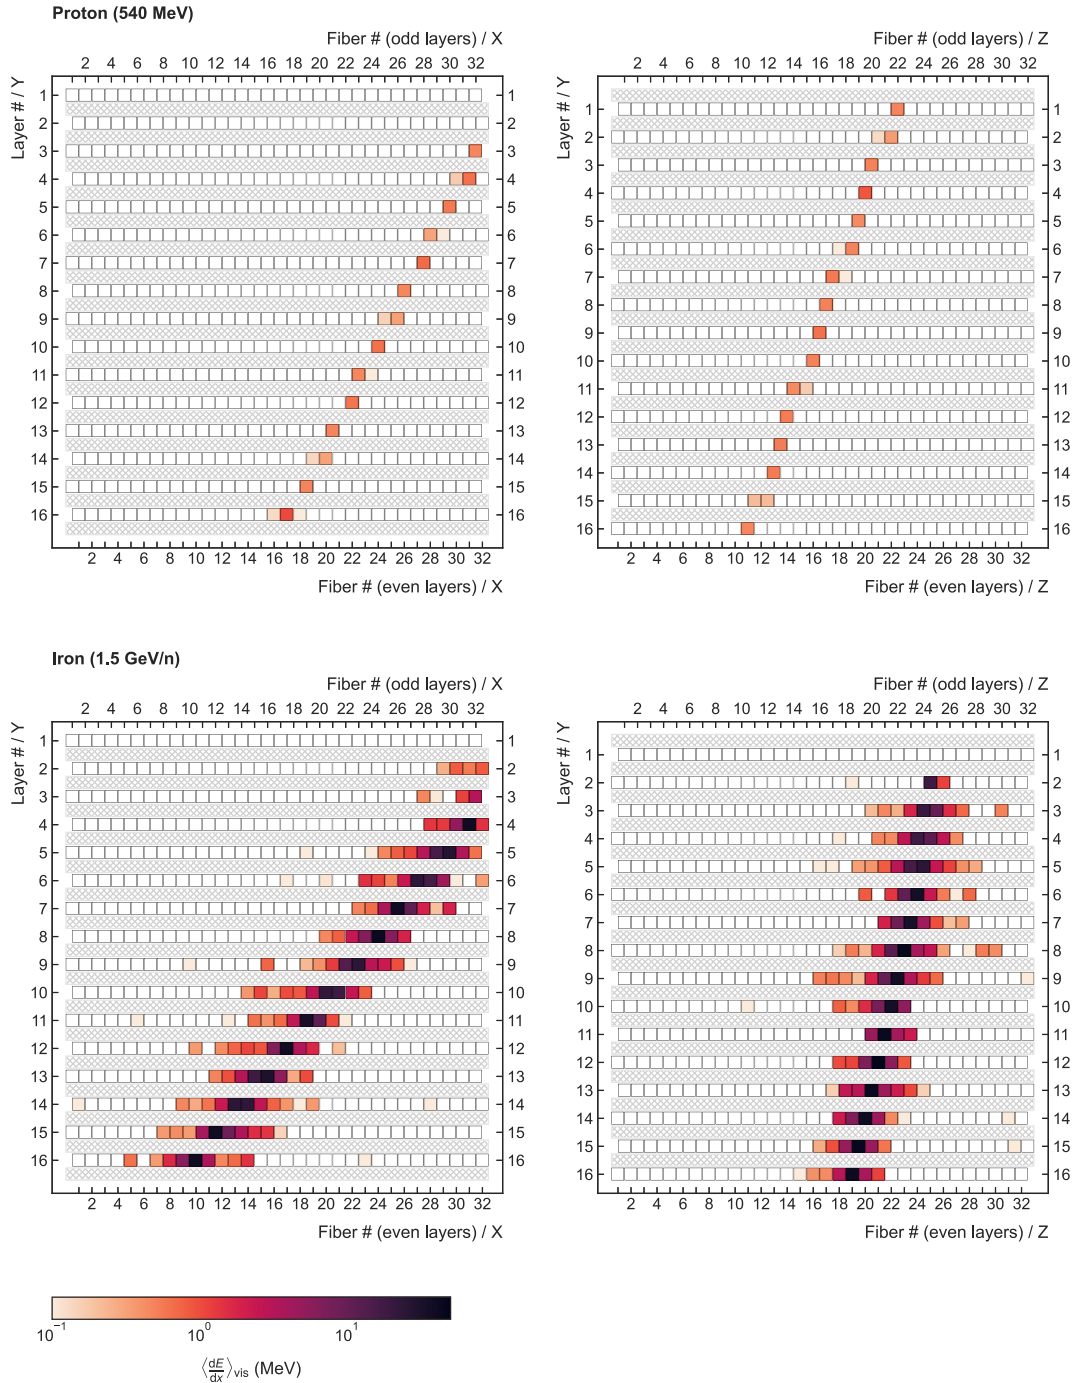


Figure 2. Simulated event signatures of a 540-MeV proton (top) and a 1.5-GeV/n iron nucleus (bottom), shown in the yx - (left) and yz -projections (right) illustrated in Figure 1. The color indicates how much energy is converted into scintillation light in each fiber.

parallel convolutions with different filter sizes (Szegegy et al., 2015). This architecture allows the networks to learn by themselves which initial scale is best suited to finding the patterns we are looking for.

3.4 Training data

We generate the data used for the training, validation, and testing of the neural networks with the GEANT4 Monte Carlo

simulation (MC) simulation toolkit (version 11.2) using its standard physics list, FTFP_BERT (Agostinelli et al., 2003; Allison et al., 2016). We implemented detector-specific effects like ionization quenching for a realistic representation of the physical processes leading to signal formation in the ADU but did not include any noise effects. For the purposes of this article, the detector model consisted of only the stack of 1024 scintillating fibers. We deliberately did not include any material around the detector because it is a significant source of scattering and frag-

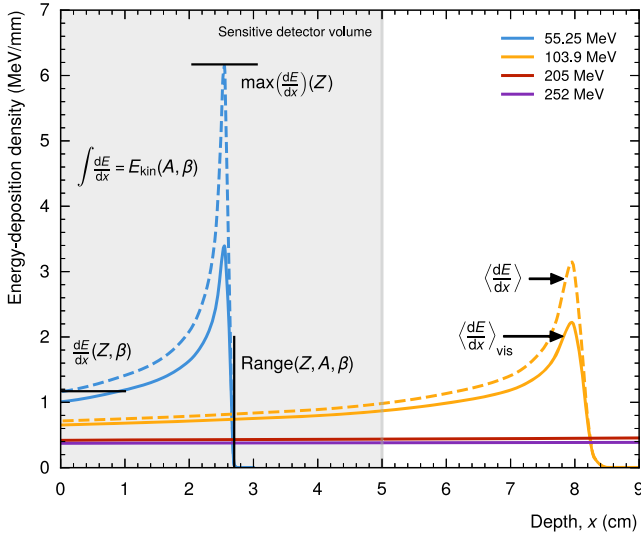


Figure 3. Illustration of the parameters that can be determined from the energy-deposition profiles of cosmic-ray nuclei. The dashed curves show the mean energy deposition, (dE/dx) , of protons with different energies in high-density polyethylene¹. The energy-deposition profile of a stopping particle encodes its nuclear charge, Z , mass number, A , velocity, β , and kinetic energy, E_{kin} . For penetrating particles, only Z and β can be determined. The solid curves illustrate how ionization quenching reduces the measurable energy-deposition density, which we denote as $(dE/dx)_{\text{vis}}$. For clarity, fluctuations due to energy-loss straggling are not shown.

mentation, and we wanted to assess the performance of the ADU in the most unbiased way possible. In the extreme case of thick shielding – as provided by the ISS, for example (Koontz et al., 2005) – the fragmentation probability for heavy nuclei is nearly 100%, which means they can almost never be observed inside the spacecraft (Zeitlin & La Tessa, 2016). Recreating this environment would not have allowed us to study the reconstruction performance for all but the lightest elements.

We concentrated on the GCR component of the deep-space environment and chose particle and energy distributions that ensured an unbiased training of the networks. We simulated the GCR spectrum using the most abundant isotope of each naturally occurring element from hydrogen to iron (i.e., ^1H , ^4He , ^7Li , ..., ^{58}Fe). We did not include electrons and gamma rays because they are not relevant to our work. Furthermore, we assumed all elements to be equally abundant to prevent the networks from being biased towards certain elements.

Likewise, we did not use a fully realistic power-law spectrum for modeling the GCR energy distribution to avoid introducing a significant bias towards lower energies. Instead, we used a log-uniform distribution with task-specific energy ranges, which does take into account that smaller energies are more probable but still ensures that our training data includes a sufficiently large number of nuclei with higher energies. The nuclei were created on a spherical surface surrounding the detector, with an angular distribution that follows a cosine law to make the incident flux isotropic. Unless stated otherwise, we included only events for which a particle was detected in at least three fibers of each orientation ($N_{\text{sig}}^{\text{Jx}} \geq 3$ and $N_{\text{sig}}^{\text{Jz}} \geq 3$).

¹ <https://www.bnl.gov/nsrl/userguide/bragg-curves-and-peaks.php>

4 Track reconstruction

Full track reconstruction requires determining the angles θ and ϕ (see Fig. 1) and the position of the track. We are, however, primarily interested in the two angles because they are all that is required to determine the angular distribution of the incident particle flux. ϕ can be determined directly from the yx -projection; θ must be calculated from its projection onto the yz -plane, θ_{proj} , via $\theta = \arctan(\tan \theta_{\text{proj}}/\sin \phi)$. Note that $\theta \in [0^\circ, 180^\circ]$ and that we choose the value range of the arctan accordingly.

The neural network for track reconstruction has two output layers, one for θ_{proj} and one for ϕ (see Fig. 4). It performs a dual classification task with a bin width of 0.2° . Although using a regression approach seemed more intuitive to determine the two real-valued angles at first, a classification yielded significantly better results with an architecture of comparable complexity. The output of the network therefore consists of two discrete pseudo-probability distributions over 900 and 1800 classes, respectively. The reconstructed value for each angle is given by the output class with the highest attributed probability.

We trained the network on 0.7 million simulated events, reading the full data set in each iteration of the training process. We used the AdamW optimizer (Loshchilov & Hutter, 2019) to update the network weights and selected the final model via early stopping to avoid overfitting (Prechelt, 1998): After each iteration, the performance of the network was validated on a distinct data set of 0.2 million events, and the result was saved only if the reconstruction performance improved with respect to the previous result.

For track reconstruction, we require a more restrictive event selection than outlined above, with $N_{\text{sig}}^{\text{Jx}} \geq 5$ and $N_{\text{sig}}^{\text{Jz}} \geq 5$. In addition, we selected only particles that reached the central region of the detector, a cuboid measuring (24 layers \times 24 fibers \times 24 fibers).

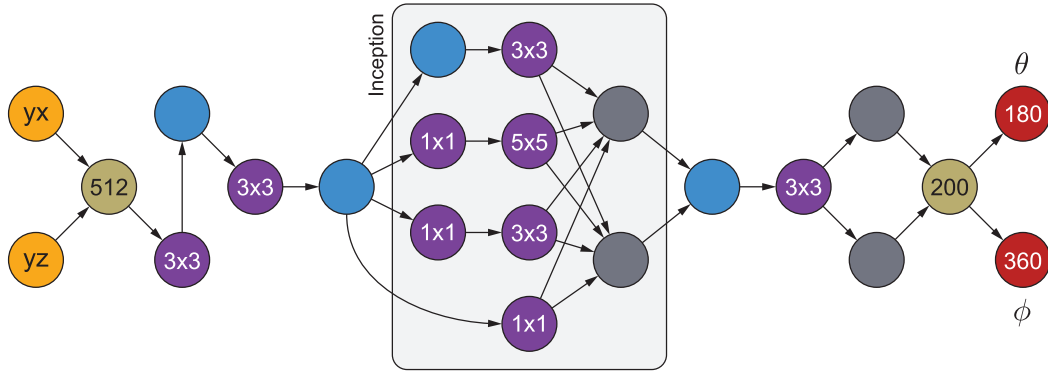
4.1 Results

We assessed the capabilities of the network architecture by training it on three different data sets each for both protons and iron nuclei. The first consisted of particles that stopped in the detector, which we selected using knowledge from the GEANT4 simulation. The second comprised monoenergetic, penetrating particles with energies of 120 MeV for protons and 500 MeV/n for iron nuclei. This ensured that the particles' energy-deposition density increased significantly along their track (see yellow curve in Fig. 3). Finally, the third data set contained minimum-ionizing particles (MIPs) with an energy of 3 GeV/n. Each data set contained 0.1 million events.

As figure of merit we use the differences between the reconstructed angles (θ_{rec} and ϕ_{rec}) and the true angles from the GEANT4 simulation (θ_{mc} and ϕ_{mc}), i.e., $\Delta\theta = \theta_{\text{rec}} - \theta_{\text{mc}}$ and $\Delta\phi = \phi_{\text{rec}} - \phi_{\text{mc}}$. Figure 5 shows the resulting distributions of $\Delta\theta$ and $\Delta\phi$. They are well centered around zero for all data sets, exhibiting only minimal shifts, μ , of their mean values: $|\mu_{\Delta\phi}| \leq 0.07^\circ$ and $|\mu_{\Delta\theta}| \leq 0.1^\circ$ for protons, and $|\mu_{\Delta\phi}| \leq 0.08^\circ$ and $|\mu_{\Delta\theta}| \leq 0.12^\circ$ for iron nuclei. All μ are smaller than the bin width of 0.2° , showing that our reconstruction is not significantly biased.

In the distribution of $\Delta\phi$ for minimum-ionizing protons, we observe that about 28% of events are reconstructed with

Track Reconstruction



Charge Determination & Energy Measurement

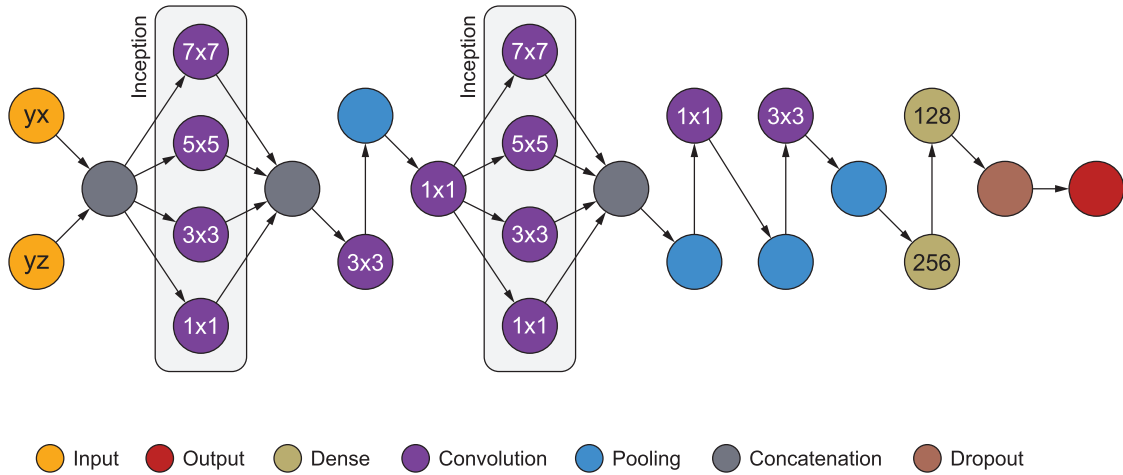


Figure 4. Architectures of the neural networks we use. Each circle represents a network layer and the color indicates the layer type. Dense and convolutional layers are labeled with their total size and their filter size, respectively. The network for track reconstruction has about 2.8 million trainable parameters, the one for charge determination and energy measurement about 2.1 million.

$|\Delta\phi| \approx 180^\circ$. That is because the network is able to correctly determine the orientation of these tracks but not the direction in which the particle is traveling. For nuclei whose energy-deposition density changes very little throughout the detector, the network does not have sufficient information to correctly determine the direction. Consequently, the fraction of events with an inverted reconstructed direction is significantly reduced for 120-MeV and for stopping protons. In the former case, the energy deposition changes substantially along their track while in the latter case it exhibits a Bragg peak and the corresponding track has an entry but no exit point.

To account for our inability to reconstruct the track direction for minimum-ionizing protons, we examined the distributions of $\Delta\phi$ and $\Delta\theta$ for direction-independent tracks. The definition of the track orientation is then ambiguous, with two combinations of ϕ and θ describing the same track. The two allowed values for ϕ are off by 180° . This translates into two possible values for the other angle, θ and $180^\circ - \theta$. To account for this ambiguity, we manually determined the direction of the track based on our knowledge from the MC data. This allowed us to select the correct (ϕ, θ) combination for the computation of $\Delta\phi$ and $\Delta\theta$. We assume their distributions to be Gaussian and centered

at zero, such that we can calculate their standard deviation, σ , from the width of the interval containing 68% of events.

For minimum-ionizing protons, we achieve $t(\sigma_{\Delta\phi}^{\text{mip}} = 0.6^\circ$ and $\sigma_{\Delta\theta}^{\text{mip}} = 0.7^\circ)$. This demonstrates that the network has learned correctly that adjacent output classes correspond to small angular differences for both angles. For stopping ($\sigma_{\Delta\phi}^{\text{stop}} = 1.4^\circ$ and $\sigma_{\Delta\theta}^{\text{stop}} = 1.3^\circ$) and 120-MeV protons ($t(\sigma_{\Delta\phi}^{\text{mono}} = 1.2^\circ$ and $\sigma_{\Delta\theta}^{\text{mono}} = 1.2^\circ)$), we obtain slightly larger values. The event signatures of protons with lower energies on average contain a smaller number of fibers with signal, making the track reconstruction more difficult for the network (or any other algorithm). Conversely, we show in Figure 6 that the overall tracking performance improves with the number of fibers with signal. However, since only a very small fraction of stopping protons generate a signal in more than 40 fibers, the network does not learn to reconstruct them well and the angular resolution worsens.

In contrast, the reconstruction performance for iron nuclei does not exhibit a notable energy dependence and we obtain very similar values for all three examined cases: $\sigma_{\Delta\phi} = 0.8^\circ$ and $\sigma_{\Delta\theta} \leq 1.1^\circ$. This is because iron nuclei create a broader energy-deposition signature than protons and we count all fibers

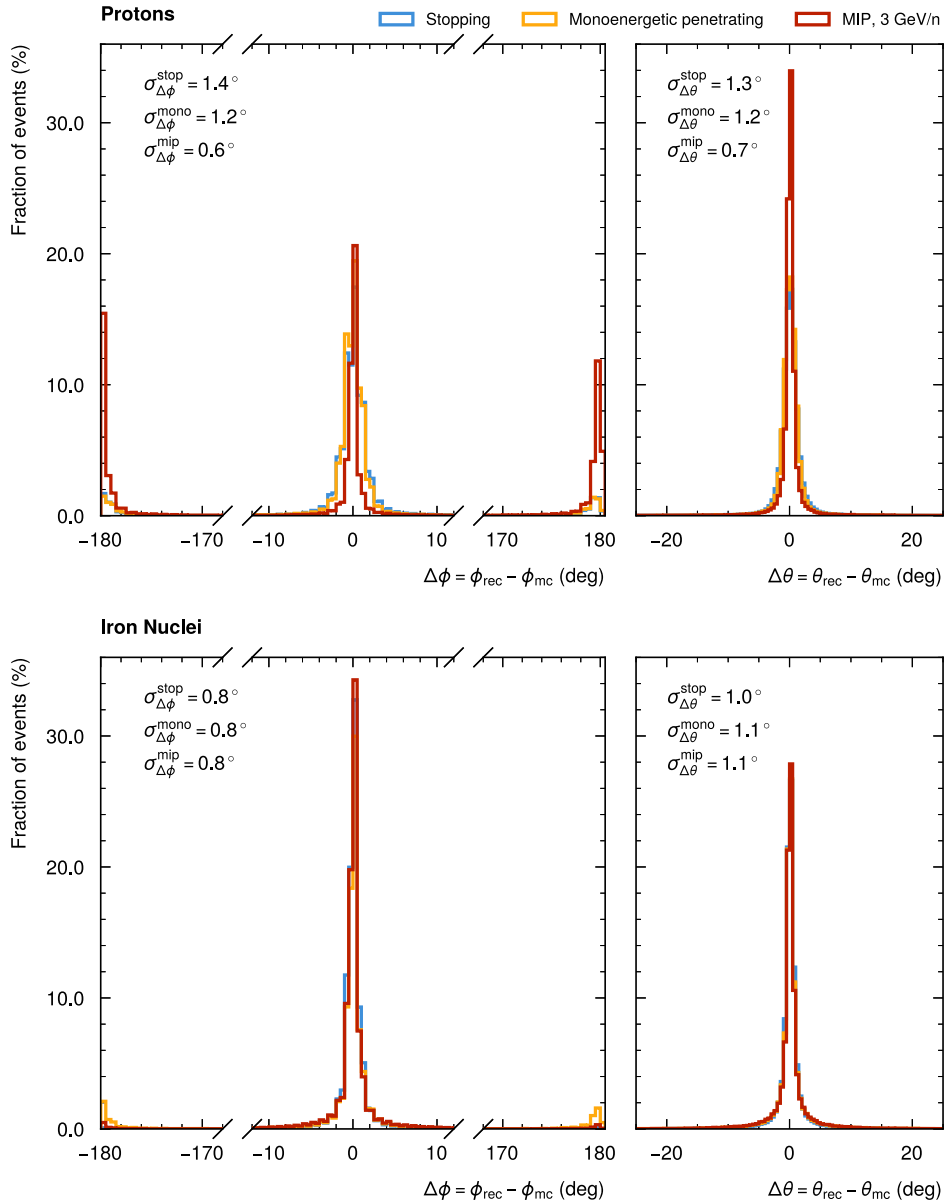


Figure 5. Reconstruction performance for the track angles ϕ (left) and θ (right) for protons and for iron nuclei, shown as the differences ($\Delta\phi$ and $\Delta\theta$) between the reconstructed angles (θ_{rec} and ϕ_{rec}) and the true angles (θ_{mc} and ϕ_{mc}). The blue histograms show the performance using networks trained to reconstruct the tracks of stopping particles. The yellow histograms show the corresponding results for monoenergetic particles (120 MeV for protons and 500 MeV/n for iron), the red histograms for minimum-ionizing particles with an energy of 3 GeV/n. For each case, we quote the σ of the direction-independent $\Delta\phi$ and $\Delta\theta$ distributions.

in which either the primary nucleus or any secondary particle deposits energy as generating a signal (see Fig. 2). The resulting higher information content of the events makes it easier for the network to reconstruct the particle tracks. Notably, it has significantly less difficulty in determining the direction because the track structures become wider along the direction of motion.

4.2 Discussion

The angular resolutions of our track reconstruction for protons and iron nuclei ($\sigma_{\Delta\phi} \leq 1.4^\circ$ and $\sigma_{\Delta\theta} \leq 1.3^\circ$) are certainly

adequate for directionally resolved radiation monitoring. Using a simple geometric approximation, we estimated that the performance of our network is in fact close to the ideal limit imposed by the detector's effective pixel size of $2 \times 2 \text{ mm}^2$ (Pöschl, 2022). Nonetheless, the values must be interpreted as a limit on the achievable resolution. In the real detector, deviations in the placement of the scintillating fibers lead to reconstruction uncertainties. Other effects, for example a non-uniform light yield of the fibers and variations in the SiPM response, contribute as well. The study of these uncertainties remains future work and is beyond the scope of this paper.

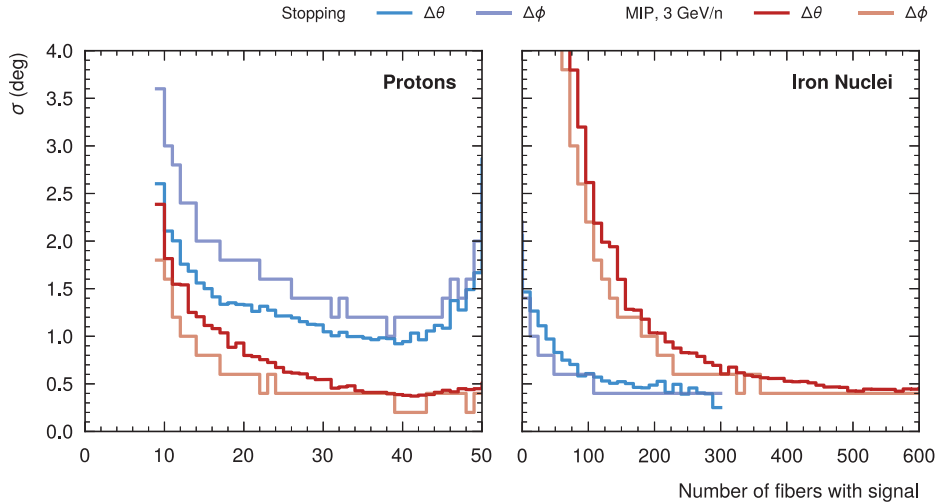


Figure 6. Dependence of σ on the number of fibers with signal per event for stopping particles (blue) and minimum-ionizing particles (red). Values of $\sigma_{\Delta\phi}$ are multiples of the output resolution of the network (0.2°); $\sigma_{\Delta\theta}$ is calculated from ϕ and θ_{proj} . The number of fibers with signal can be larger than the depth of the ADU (32 fibers/layers) due to the emission of secondary particles (see Fig. 2).

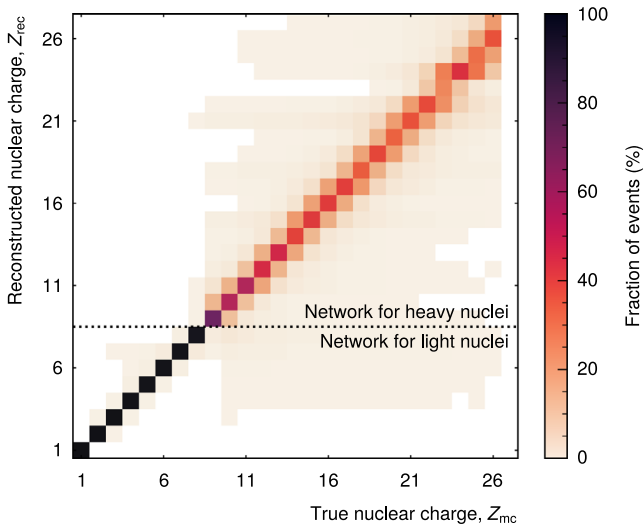


Figure 7. Confusion matrix of the charge determination through the consecutive application of two neural networks, trained to identify light nuclei through oxygen ($Z = 8$) and heavy nuclei through iron ($Z = 26$), respectively. The second network was trained (but not tested) on a data set containing cobalt ($Z = 27$). Each column is normalized and shows the distribution of reconstructed nuclear charge, Z_{rec} , for all events with true nuclear charge, Z_{mc} . Cells with values smaller than 0.1% are drawn in white.

5 Charge determination

The ADU can identify nuclei only via their specific energy loss, which is a function of their (nuclear) charge and velocity (see Fig. 3). For stopping particles, the energy-deposition profile (whose reconstruction requires knowledge of the track parameters) would in principle also allow to determine a nucleus' mass

number, A . We are, however, not interested in A because it (a) is largely irrelevant for radiation dosimetry and (b) does not have a noticeable effect (within the resolution of our detector) for all but the lightest nuclei. In the simplified case of an unshielded detector that is subjected to a modified GCR environment as described in Section 3.4, we can also assume that all primary particles with $|Z| = 1$ are protons. Other relevant particles with this charge (e.g., deuterons, muons, pions, electrons, and positrons) are only created as secondaries in the detector volume. In the context of the work presented here, determining the charge of particles therefore allows us to unambiguously identify nuclei of all elements from hydrogen to iron.

As described in Section 3.3, we found that adapting the architecture of the neural network used for track reconstruction is advantageous for charge determination. Instead of centering the architecture around a single inception constituted of stacked convolutional layers, we used multiple flat inceptions. Each of the inceptions contains four convolutional layers arranged in parallel (see Fig. 4). The number of classes in the output layer of the network corresponds to the range of Z we attempt to identify. For each event, the network computes a pseudo-probability distribution over all possible Z , and we select the class with the highest attributed probability as the reconstructed value.

We initially attempted to identify all Z with a single network, achieving acceptable results for light elements but poor performance for heavier ones. This is plausible because the effects leading to deviations from the ideal energy-deposition profile (straggling, fragmentation, and quenching) become stronger for increasing Z . We therefore perform a two-step charge determination: A first network attempts to identify nuclei from hydrogen ($Z = 1$) through oxygen ($Z = 8$), sorting all particles it believes to be of higher charge into an additional ‘overflow’ class. The events in this class are passed on to a second network that tries to identify fluorine ($Z = 9$) through iron ($Z = 26$). This approach greatly improves the reconstruction performance. We used the same training methodology as for

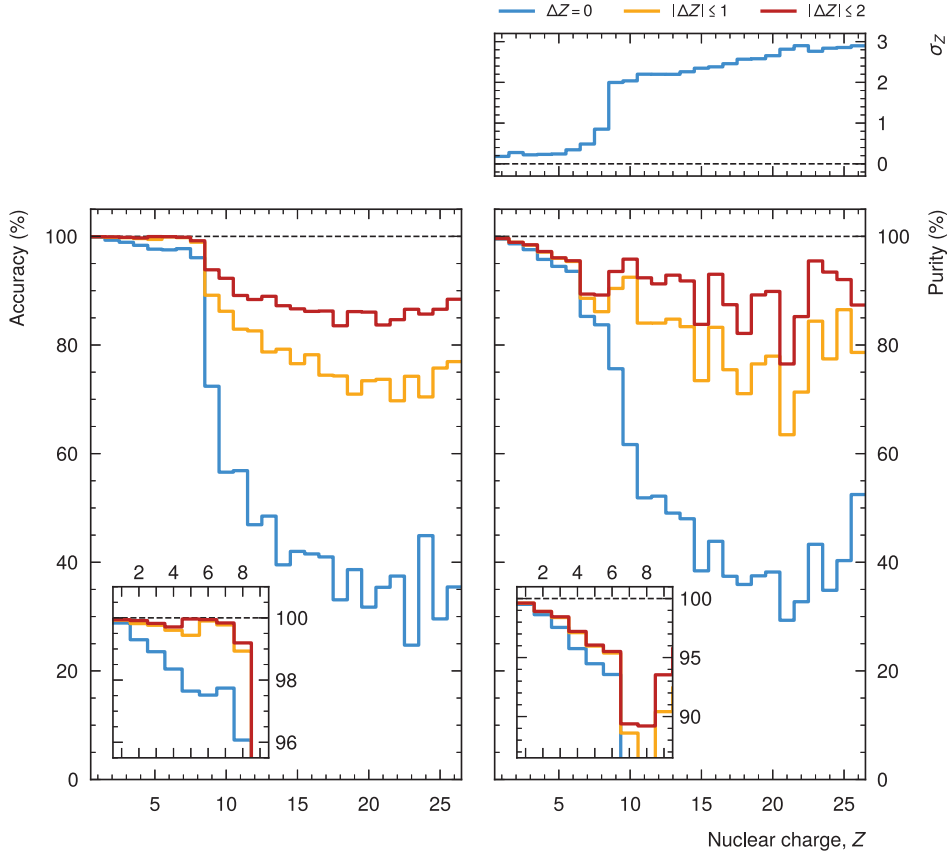


Figure 8. The accuracy (left panel) and purity (bottom right panel) of the charge determination as a function of Z . We define the accuracy as the fraction of correctly identified events for a given Z_{mc} and the purity as the fraction of events in a Z_{rec} class for which $Z_{rec} = Z_{mc}$. The top right panel shows the standard deviation, σ_Z , of the reconstructed-charge distributions (the columns in Fig. 7). The blue curves show the performance for an exact identification of the charge ($\Delta Z = Z_{rec} - Z_{mc} = 0$). The yellow and red curves illustrate the improvement when we allow $|\Delta Z| \leq 1$ and $|\Delta Z| \leq 2$, respectively. The insets in the lower panels provide zoomed-in views for $Z \leq 8$.

the track reconstruction network (with nine million training events and one million testing events) but included cobalt ($Z = 27$) in the data set to avoid boundary effects for iron. The nuclei had energies between 20 MeV and 5 TeV with a log-uniform distribution.

5.1 Results

Figure 7 summarizes the performance of the charge determination in the form of a confusion matrix. Each column shows the distribution of reconstructed nuclear charges, Z_{rec} , for all events with true nuclear charge Z_{mc} . The framework finds the true charge, or a very close one, for the majority of events. This suggests that the networks learn that nuclei with similar energy-deposition profiles have similar charges. Averaging over the full charge range, we find that 59% of events are assigned to their true Z class.

As expected, the accuracy, i.e., the fraction of correctly identified events for a given Z_{mc} , decreases with increasing Z . Hydrogen (protons) and helium events are correctly identified in 99.8% and 99.3% of cases, respectively. For light nuclei up to oxygen, we achieve values well over 95%. It is only for heavier nuclei that the exact charge determination evidently becomes

more difficult. However, the network still learns to roughly recognize the energy-deposition profiles of nuclei with larger Z and in most cases assigns the biggest fraction of events to the correct class. Incorrectly reconstructed events are assigned to neighboring classes, with Z_{rec} close to Z_{mc} .

This is illustrated in Figure 8, which shows the accuracy and the standard deviation, σ_Z , (here calculated as the square root of the variance) of the reconstructed-charge distributions as a function of Z . The blue curves show the performance for an exact identification of the charge ($\Delta Z = Z_{rec} - Z_{mc} = 0$). The fraction of correctly classified events slowly decreases from nearly 100% (hydrogen) to a little over 96% (oxygen), then drops to values below 75% (fluorine) when the second network takes over. It ultimately reaches values between 30% and 40% for the highest Z . The same behavior is reflected in the width of the reconstructed- Z distributions: For light elements, σ_Z is close to 0.2; it then jumps to two for fluorine and steadily rises to almost three for iron. If we allow $|\Delta Z| \leq 1$ (i.e., we count every event with $Z_{mc} - 1 \leq Z_{rec} \leq Z_{mc} + 1$ as correctly identified), the values for light nuclei increase to over 99% all the way to oxygen (yellow curve in Fig. 8). Though we still observe a substantial break at the boundary between the two networks, the performance for heavy nuclei improves significantly and stays

above 70% for all elements. If we allow $|\Delta Z| \leq 2$ (red curve), we obtain a further substantial improvement for large Z , with values consistently above 83% across the whole range. In these cases, 84% and 91% of all events are assigned to their true Z class, respectively. This behavior strongly suggests that our framework is able to reliably determine that a particle is, for example, iron-like even if it cannot exactly determine its charge.

We also investigated the purity of the charge determination, which we here define as the fraction of events in a Z_{rec} class for which $Z_{\text{rec}} = Z_{\text{mc}}$. It is 99.6% and 98.8% for hydrogen and helium, respectively, and exceeds 84% for elements through oxygen. For heavier nuclei, it decreases to values around 35%. For $\Delta Z = 0$, the mean purity over the entire charge range is 58%. If we allow $|\Delta Z| \leq 1$ or $|\Delta Z| \leq 2$, the effect is analogous to that discussed above for the reconstruction accuracy (see bottom right panel of Fig. 8), and the overall mean purity increases to 83% and 91%, respectively.

5.2 Discussion

The accuracy and purity of the charge determination for hydrogen (99.8% and 99.6%, respectively) and helium (99.3% and 98.8%) represent a significant improvement over the performance of currently operational instruments. Together, these elements make up 99% of all cosmic-ray nuclei, and the accuracy and purity of their identification is therefore crucial to dosimetry. For light nuclei ($Z \leq 8$), the performance for $\Delta Z = 0$ (accuracy and purity over 95% and 84%, respectively) is also adequate. The fact that accuracies better than 83% for heavier elements can only be obtained if we allow $|\Delta Z| \leq 2$ is not critical in the context of our work. Though the biological effectiveness of nuclei with similar charge differs somewhat due to the quadratic dependence of the energy deposition on Z , the uncertainty introduced by allowing $|\Delta Z| \leq 2$ is still much smaller than for a determination of Z from the LET alone. Despite the reduced performance for heavy nuclei, the ADU can therefore measure the biologically relevant dose with higher accuracy than the majority of sensors used today.

We observe a sharp drop in performance at the interface between the low- Z and high- Z networks (see Fig. 8). This suggests that the network architecture struggles most with learning to reconstruct the highest Z , which negatively affects the performance over the whole charge range the network is trained for. This hypothesis was corroborated by our attempts to shift the hand-over point between the networks to higher Z , which increased the accuracy of the low- Z network for medium-light elements at the expense of a significantly reduced performance for hydrogen and helium. The difficulty in determining the charge of heavy nuclei reveals the inherently limited separation power of the detector: Ionization quenching in the plastic scintillators leads to ever smaller differences between the mean energy deposition of nuclei with larger Z (Losekamm, 2025). At the same time, stronger energy-loss straggling and a higher fragmentation probability lead to larger variations in the energy-deposition profiles. Combined, these effects result in an increasingly large overlap in the profiles for particles with similar charge, restricting our ability to accurately determine Z on an event-by-event basis.

The impact of the network's difficulty to separate heavier nuclei on its performance for lighter ones also reveals one of the challenges of using machine-learning algorithms. During

training, the network learns that the energy-deposition profiles of nuclei with higher Z can vary widely. Since it attempts to find features that are common to all events, it (incorrectly) applies this knowledge to lighter elements. In the extreme – if the charge range of a network is chosen too wide – the bias introduced by large Z can completely degrade the excellent performance that can be achieved for the clearly separable light elements. Therefore, a promising approach to improve the charge determination for medium- Z elements might be to increase the number of networks, thereby reducing the charge range covered by each of them. A framework with more branches does, however, have a higher complexity (and thus longer execution time) and more hand-over points at which boundary effects can occur. In addition, a network architecture with more parameters may also help to increase the accuracy.

6 Energy measurement

Determining the kinetic energy of nuclei traversing the ADU requires knowledge of their mass, which we determine from Z by assuming a corresponding average A . We then use 26 element-specific neural networks of identical architecture (see Fig. 4) to determine their energy. For low-energy nuclei that stop in the detector, the energy is given by the sum of their energy depositions. For penetrating particles, the networks must learn to relate the energy-deposition profile to their velocity. They perform a regression analysis and return the energy per nucleon as a real number in the range of 20 MeV/n to 1 GeV/n in their single output node (except for the heaviest elements, for which the upper limit is 10 GeV/n). We trained each network on an individual data set of 1.8 million simulated events and evaluated the combined framework using a separate data set of eight million events.

The purity of the subsets of events passed on to the energy-reconstruction networks ranges from 99.6% for hydrogen to 52.3% for iron (see previous section). At the same time, the error of reconstructing a nucleus' energy assuming a charge that is slightly off becomes smaller for heavier nuclei because of the decreasing relative difference in mass of adjacent elements. Therefore, to reduce the impact of an incorrectly reconstructed Z for elements with $Z \geq 2$, the individual training data sets comprised events with charges in the range $[Z - 1, Z + 1]$ or $[Z - 2, Z + 2]$ (which is equivalent to allowing $|\Delta Z| \leq 1$ or $|\Delta Z| \leq 2$, see previous section), with all three (five) possible values equally likely.

6.1 Results

Measuring the energy of stopping nuclei in principle does not require a minimum number of fibers with a signal. To find the ideal lower limit of our detector's sensitivity range, we investigated two scenarios: one in which we applied no selection cuts to our test data set and one where we selected only events with $N_{\text{sig}}^{\text{Dx}} \geq 3$ and $N_{\text{sig}}^{\text{Dz}} \geq 3$ (like for charge determination). The results for hydrogen are illustrated in the left and right panels of Figure 9, respectively. We show the bin-wise energy resolution, σ_E/E_{kin} , where σ_E is the Gaussian width of the reconstructed energies for nuclei within a given energy bin and E_{kin} is the central value of that bin. For this purpose, we divide the respective energy ranges into 100 log-uniform bins.

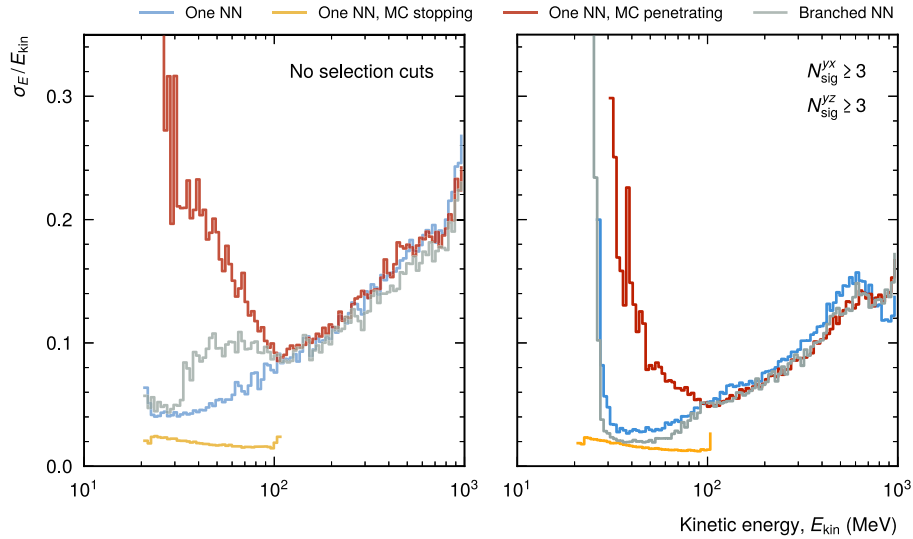


Figure 9. Bin-wise energy resolution for protons (hydrogen) with energies from 20 MeV to 1 GeV. The left panel shows the results if we apply no selection cuts to our test data, the right panel those for events with $N_{\text{sig}}^{\text{yx}} \geq 3$ and $N_{\text{sig}}^{\text{yz}} \geq 3$. The blue histograms give the performance for one neural network trained over the full data set of (stopping and penetrating) hydrogen nuclei; the yellow and red histograms that of separate networks for stopping and penetrating particles. The gray histograms show the performance of a branched framework, where a filter network separates stopping from penetrating particles and passes them to the separate energy-reconstruction networks.

The blue histograms show the resolution that a single network trained over the full data set of both stopping and penetrating hydrogen nuclei achieves. If we apply no selection cuts (see left panel), we obtain $\sigma_E/E_{\text{kin}} \leq 8\%$ for energies up to 100 MeV, for which most protons stop and deposit their full kinetic energy in the detector. The resolution monotonically decreases with increasing energies and reaches about 25% at 1 GeV. For events with $N_{\text{sig}}^{\text{yx}} \geq 3$ and $N_{\text{sig}}^{\text{yz}} \geq 3$ (see right panel), we achieve a resolution of better than 5% for $E_{\text{kin}} \leq 100$ MeV and better than 16% for higher energies. Despite this significant overall improvement, we observe a worsening of the resolution at the lower end of the reconstruction range and a non-monotonic behavior towards the upper boundary. While the former effect is simply due to the small range of protons of such low energies and the correspondingly small number of events making it through the selection cuts, the reason for the latter one is not immediately clear.

We separately trained and tested our network on penetrating and non-penetrating particles to assess how it copes with the different information content of the respective events. We manually created two data sets based on knowledge from the MC simulation: one with particles stopping in the detector (*MC stopping*) and one with particles penetrating it (*MC penetrating*). The results are shown in Figure 9 as yellow and red histograms and labeled accordingly. For stopping particles, we achieve $\sigma_E/E_{\text{kin}} \leq 2\%$, with essentially no discernible difference between the two scenarios up to 100 MeV, above which the range of protons is larger than the size of the detector. For penetrating particles, the resolution closely resembles that of the single-network reconstruction above 100 MeV in the full data set and is slightly better if we apply selection cuts. At lower energies, σ_E/E_{kin} substantially worsens with decreasing energy in either scenario, reflecting both the limited number of events in that range and the fact that penetrating particles with such low energies traverse only few fibers. The energy limit of about

100 MeV at which protons begin to penetrate the full detector translates to a break in the distribution.

We also investigated a more realistic scenario in which a neural network separates stopping from penetrating particles. Based on the decision of this filter network, events are passed on to appropriate energy-reconstruction networks specifically trained with stopping or penetrating particles (*branched NN*, see gray histograms in Fig. 9). If we apply no selection cuts, this branched framework performs marginally better than the previously discussed networks above 100 MeV but worse than the single-network reconstruction at lower energies. This suggests that the negative impact of incorrect decisions by the filter network, i.e. stopping or penetrating protons being passed to the wrong energy-reconstruction network, outweighs the advantage of using a dedicated network for stopping particles. The performance of the branched framework significantly improves if we require $N_{\text{sig}}^{\text{yx}} \geq 3$ and $N_{\text{sig}}^{\text{yz}} \geq 3$. At energies above 100 MeV, the resolution is essentially identical to the case where we manually select penetrating particles using knowledge from the simulation. Below 100 MeV, it is noticeably better than that of the single-network reconstruction and approaches the lower limit of MC-selected stopping protons for $E_{\text{kin}} \leq 50$ MeV. This indicates that our filter network can more effectively separate stopping from penetrating particles due to the (on average) larger number of fiber signals per event. The improvement comes at the expense of sensitivity at the lowest energies, where σ_E/E_{kin} drastically worsens for $E_{\text{kin}} \leq 30$ MeV because only few events pass the selection criteria.

Finally, we evaluated the performance of the single-network approach for heavier nuclei. We used a data set comprising all elements up to cobalt for this analysis. For illustration, Figure 10 shows the examples of hydrogen (H), helium (He), carbon (C), and iron (Fe) nuclei. We trained the networks for He and C with nuclei in the range $[Z - 1, Z + 1]$ of the target charge – i.e., $[1, 3]$ for He and $[5, 7]$ for C – and the one for Fe ($Z = 26$) with

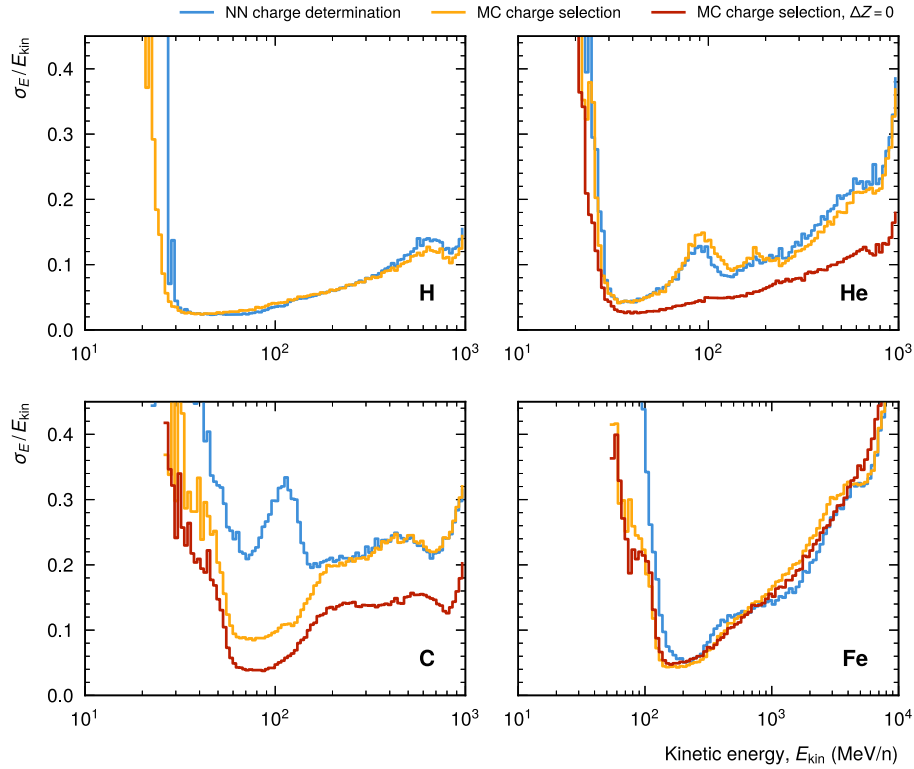


Figure 10. Energy resolution for hydrogen (H), helium (He), carbon (C), and iron (Fe) nuclei with energies from 20 MeV/n to 1 GeV/n (10 GeV/n for iron). The yellow histograms show the performance for testing events of the target charge which we selected based on knowledge from the simulation ($Z_{\text{mc}} = Z_{\text{target}}$), the blue ones the result for particles that were selected by our charge-determination network from a test data set containing all elements up to cobalt ($Z_{\text{rec}} = Z_{\text{target}}$). For helium and carbon, we trained the networks with nuclei in the range $[Z - 1, Z + 1]$ of the target charge and the one for iron in the range $Z \in [24, 28]$. The network for hydrogen was always trained for the target charge only. The red histograms show the performance for networks trained only with the target charge determined from the simulation ($\Delta Z_{\text{mc}} = 0$).

nuclei in the range $Z \in [24, 28]$. The network for hydrogen was trained for the target charge only. We tested each of the four networks with two different data sets. The first contained events of the target charge ($Z_{\text{mc}} = Z_{\text{target}}$), which we selected based on knowledge from the MC simulation (*MC charge selection*, yellow histograms). For the second data set, we used the charge-determination framework described in Section 5 to select events with $Z_{\text{rec}} = Z_{\text{target}}$ (*NN charge determination*, blue histograms). For comparison, the red histograms show the ideal performance, for which we trained all networks on the target charge only (i.e., $\Delta Z = 0$) and selected training and testing events based on knowledge from the MC simulation.

For hydrogen, the blue and yellow histograms allow to investigate the impact of charge confusion (i.e., data sets with sub-100% purities being passed on to the energy-measurement networks). We achieve the previously discussed good resolution in the case of manually selected events, and we see that wrongly assigned nuclei (primarily helium) affect it below 30 MeV and above about 500 MeV. The large decrease in resolution at the low-energy bound is intuitively understandable because $\Delta Z \geq 1$ has a huge impact on the energy loss of stopping light nuclei. The effect can be reduced by training the network for charges in the range $Z \in [1, 2]$. This, however, comes at the expense of a lower resolution across all energies, which has a larger impact on the overall performance than improving the lower sensitivity limit and is therefore not desirable. For penetrating

particles, the maximum difference is about two percentage points and occurs at around 600 MeV. In either case, we achieve $\sigma_E/E_{\text{kin}} \leq 5\%$ for $E_{\text{kin}} < 100$ MeV, $\sigma_E/E_{\text{kin}} \leq 10\%$ for $E_{\text{kin}} < 300$ MeV, and $\sigma_E/E_{\text{kin}} \leq 16\%$ for $E_{\text{kin}} < 1$ GeV.

For helium, the difference between the blue and yellow histograms is less pronounced at low energies. This illustrates that the effect of charge confusion can largely be suppressed if the network sees neighboring elements during training. The red histogram (training with $\Delta Z = 0$ and evaluation with manually selected helium events), shows that in the unrealistic case of perfect charge determination, we get a performance comparable to that for hydrogen. The largest differences are a ten-percentage-point bump around 90 MeV/n and an even larger divergence toward the upper energy boundary. Overall, we achieve $\sigma_E/E_{\text{kin}} \leq 14\%$ for $E_{\text{kin}} < 300$ MeV/n and $\sigma_E/E_{\text{kin}} \leq 24\%$ for $E_{\text{kin}} < 800$ MeV/n.

The case of carbon is a good example for medium- Z nuclei and shows that charge confusion can affect the energy resolution much more significantly than observed for helium. The impact is most pronounced at around 100 MeV/n, where the blue histogram for neural-network-based charge determination peaks at almost 34%. The yellow curve for MC-selected events, on the other hand, is at around 10%, only about four percentage points above the ideal limit (red histogram). At higher energies, the blue and yellow curves agree almost perfectly, indicating

that the network-based charge selection works well enough to be compensated for by the extended charge range of the energy network. However, charge confusion significantly affects the performance of the energy network on the NN-selected testing data for energies around 100 MeV/n, an effect that we need to further look into. Realistically achievable σ_E/E_{kin} for carbon appear to be in the 10%-to-25% range.

To cover a comparable section of the Bragg curve for all elements, we trained the networks for the heaviest elements for energies up to 10 GeV/n. The results for iron are overall better than for carbon, and the curves for the three evaluations largely agree. The fully neural-network-based framework (blue histogram) outperforms the supposedly ideal case (red histogram) for some energies above 1 GeV/n. Here, it seems that the combined effects of charge confusion and the broader training set ($|\Delta Z| \leq 2$) more than compensate for the uncertainty introduced by the low purity of the data set selected by the charge-determination network. The exact cause, however, is not yet clear and requires further investigation. Overall, we achieve $\sigma_E/E_{\text{kin}} \leq 7\%$ for $E_{\text{kin}} \approx 150$ MeV/n, $\sigma_E/E_{\text{kin}} \leq 10\%$ for $E_{\text{kin}} < 400$ MeV/n, and $\sigma_E/E_{\text{kin}} \leq 20\%$ for $E_{\text{kin}} < 2$ GeV/n.

6.2 Discussion

Our results show that energy resolutions on the order of 10% to 20% can be achieved, with even better values for hydrogen and helium over a wide range of the most relevant particle energies. However, the cases of helium, carbon, and iron also demonstrate that our charge measurement can still be improved. There are many features whose cause we do not yet fully understand but which suggest that both the charge-determination and energy-measurement networks face similar challenges. The present results nonetheless demonstrate that an energy measurement based on a neural-network-based reconstruction is possible and can deliver acceptable resolutions. An important aspect of our future work is to find a better approach for dealing with charge confusion and minimizing its impact on the energy measurement. We also need to introduce a mechanism that allows networks to deal with particle energies outside their target range.

7 Conclusion

In summary, our findings show that a detailed characterization of cosmic-ray nuclei with energies in the MeV-to-GeV range can be performed with a simple and compact instrument like the ADU of the RadMap Telescope. They also reveal the inherent physical limits of the detector and highlight some of the challenges of using a reconstruction framework based on neural networks. The results presented here are based on a simulation that is largely realistic but represents an ideal-world scenario because it does not take into account some of the aspects of a real detector that are harder to model – for example optical and electrical crosstalk, the misalignment of individual scintillating fibers, as well as detector efficiencies and digitization uncertainties. The study of how well our network architecture can cope with such effects remains open work.

We already discussed the performance and limits of the networks for the individual reconstruction tasks and how they

might be improved. Now, we also want to draw attention to potential improvements that are common to them.

For example, enhancements in all three tasks can surely be achieved by refining the training of our networks. For example, we deliberately used training data in which all elements are equally abundant to avoid biases towards those with a naturally higher flux. We also did so because using realistic abundances is impractical, as the fluxes of the most and least abundant elements differ by almost six orders of magnitude. The knowledge that the chances of encountering certain elements in the real cosmic-ray environment may differ significantly will very likely improve the overall performance of our framework. A similar argument can be made for the energy spectrum, which we likewise did not model realistically to reduce the computational effort of training. We therefore need to develop an approach for teaching our networks the real elemental and spectral abundances without using fully realistic training data. The challenge of doing so lies in ensuring that the networks are optimally trained across their whole sensitivity range.

Some of the framework's limitations may be caused by our networks not being able to capture all information contained in a detector event. We chose rather simple convolutional architectures that make use of a representation of event signatures as two images. A multitude of more complex approaches has since been applied to track and characterize particles in detectors. It was shown that transformer architectures can perform a variety of tasks (Vaswani et al., 2017); graph neural networks are now widely used to reconstruct particle tracks (Scarselli et al., 2009); and physics-informed networks may be useful for refining our charge and energy determination (Raissi et al., 2019). Part of our future work is to investigate the impact of a more complex architecture on our reconstruction performance and computational requirements.

Another limiting factor common to all reconstruction tasks is the little information contained in events with only few fiber signals, for example of particles that have very low energies or traverse one of the corners of the detector. Tracks or energy-deposition profiles with fewer data points suffer more from statistical fluctuations caused by straggling and path-length differences than longer ones. Though we have not yet studied this effect systematically, the case of the track reconstruction shows that such events can have substantially higher uncertainties. Excluding them from analysis, potentially at the expense of sensitivity at the lowest energies, may improve the overall performance of the networks, though care must be taken to avoid introducing biases during event selection.

Finally, all our results are based on simulations for a detector in open space, i.e., with almost no material around it. This is, of course, not a realistic application scenario. In the RadMap Telescope, the ADU is contained in a housing made from aluminum and surrounded by layers of electronics with a non-uniform material budget. The instrument is deployed inside the ISS, which provides yet more substantial shielding. Even in a (hypothetical) scenario where the detector was mounted to the outside of a spacecraft with little need for protection against the space environment, significant shielding across parts of the solid angle would be provided by the spacecraft itself. All this shielding significantly alters both the composition and the energy distribution of the radiation field, creating yet another challenge for optimally training the networks. However, neglecting the effects

from additional material allowed us to evaluate the capabilities of the instrument and our reconstruction framework under near-optimal conditions. In the context of our work, the positive effect of shielding is that the relative contribution of hydrogen and helium to the absorbed dose becomes ever more dominant for increasing material thickness (Slaba et al., 2017). Since our charge determination and energy measurement work best for these light nuclei, this leads to a smaller uncertainty of the biologically relevant dose calculated from our measurements.

In conclusion, our study demonstrates that the RadMap Telescope can perform spectroscopic measurements that significantly exceed the capabilities of most instruments used for radiation monitoring aboard spacecraft today. A distinguishing feature is the nearly omnidirectional acceptance of the detector, i.e., its field of view covering the full solid angle. Our work also shows that using a neural network-based reconstruction is feasible and can produce acceptable results. Even though the numerical values cited here may shift slightly (for better or worse) once we implemented a more realistic detector model and changes to our framework, we very much expect this qualitative conclusion to hold.

Acknowledgments

The editor and authors thank Robert Wimmer-Schweingruber and an anonymous reviewer for their assistance in evaluating this paper, and for their constructive and helpful feedback.

Funding

This work was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) via grant 414049180 and via Germany's Excellence Initiative – EXC-2094 – 390783311.

Conflicts of interest

The authors declare no conflict of interest.

Data availability statement

The models and settings used to generate the GEANT4 data that this work is based on are available from the authors upon reasonable request.

References

- Agostinelli S, Allison J, Amako K, Apostolakis J, Araujo H, et al. 2003. GEANT4 – a simulation toolkit. *Nucl Instr Meth Phys Res Section A* **506(3)**: 250–303. [https://doi.org/10.1016/s0168-9002\(03\)01368-8](https://doi.org/10.1016/s0168-9002(03)01368-8).
- Aguilar M, Cavazonza LA, Ambrosi G, Arruda L, Attig N, et al. 2021. Periodicities in the Daily Proton Fluxes from 2011 to 2019 measured by the alpha magnetic spectrometer on the international space station from 1 to 100 GV. *Phys Rev Lett* **127(27)**: 271102. <https://doi.org/10.1103/PhysRevLett.127.271102>.
- Alaghband Y, Klein PM, Kramar EA, Cranston MN, Perry BC, et al. 2023. Galactic cosmic radiation exposure causes multifaceted neurocognitive impairments. *Cell Mol Life Sci* **80(1)**: 29. <https://doi.org/10.1007/s00018-022-04666-8>.
- Allison J, Amako K, Apostolakis J, Arce P, Asai M, et al. 2016. Recent developments in Geant4. *Nucl Instr Meth Phys Res Sect A* **835**: 186–225. <https://doi.org/10.1016/j.nima.2016.06.125>.
- Barthel J, Sarigul-Klijn N. 2019. A review of radiation shielding needs and concepts for space voyages beyond Earth's magnetic influence. *Prog Aerospace Sci* **110**. <https://doi.org/10.1016/j.paerosci.2019.100553>.
- Birks JB. 1951. Scintillations from organic crystals: specific fluorescence and relative response to different radiations. *Proc Phys Soc Sec A* **64(10)**: 874–877. <https://doi.org/10.1088/0370-1298/64/10/303>.
- Caldwell A, Kollár D, Kröninger K. 2009. BAT – The Bayesian analysis toolkit. *Comp Phys Commun* **180(11)**: 2197–2209. <https://doi.org/10.1016/j.cpc.2009.06.026>.
- Chancellor JC, Blue RS, Cengel KA, Auñón-Chancellor SM, Rubins KH, Katzgraber HG, Kennedy AR. 2018. Limitations in predicting the space radiation health risk for exploration astronauts. *NPJ Micrograv* **4**: 8. <https://doi.org/10.1038/s41526-018-0043-2>.
- Chancellor JC, Scott G, Sutton J. 2014. Space radiation: the number one risk to astronaut health beyond low earth orbit. *Life* **4**: 491–510. <https://doi.org/10.3390/life4030491>.
- Cucinotta FA. 2015. Review of NASA approach to space radiation risk assessments for mars exploration. *Health Phys* **108**: 131–142. <https://doi.org/10.1097/HP.0000000000000255>.
- Cucinotta FA, Alp M, Rowedder B, Kim M-HY. 2015. Safe days in space with acceptable uncertainty from space radiation exposure. *Life Sci Space Res* **5**: 31–38. <https://doi.org/10.1016/j.lssr.2015.04.002>.
- Cucinotta FA, To K, Cacao E. 2017. Predictions of space radiation fatality risk for exploration missions. *Life Sci Space Res* **13**: 1–11. <https://doi.org/10.1016/j.lssr.2017.01.005>.
- Delp MD, Charvat JM, Limoli CL, Globus RK, Ghosh P. 2016. Apollo lunar astronauts show higher cardiovascular disease mortality: possible deep space radiation effects on the vascular endothelium. *Sci Rep* **6**: 29901. <https://doi.org/10.1038/srep29901>.
- Di Fino L, Romoli G, Santi Amantini G, Boretti V, Lunati L, et al. 2023. Radiation measurements in the International Space Station, Columbus module, in 2020–2022 with the LIDAL detector. *Life Sci Space Res (Amst)* **39**: 26–42. <https://doi.org/10.1016/j.lssr.2023.03.007>.
- Dietze G, Bartlett DT, Cool DA, Cucinotta FA, Jia X, McAulay IR, Pelliccioni M, Petrov V, Reitz G, Sato T. 2013. Assessment of radiation exposure of astronauts in space. ICRP Publication 123. *Annals of the ICRP* **42(4)**: 1–339. <https://doi.org/10.1016/j.icrp.2013.05.004>.
- Durante M, Cucinotta FA. 2011. Physical basis of radiation protection in space travel. *Rev Modern Phys* **83**: 1245–1281. <https://doi.org/10.1103/RevModPhys.83.1245>.
- Fogtman A, Baatout S, Baselet B, Berger T, Hellweg CE, et al. 2023. Towards sustainable human space exploration – priorities for radiation research to quantify and mitigate radiation risks. *npj Micrograv* **9(1)**. <https://doi.org/10.1038/s41526-023-00262-7>.
- George SP, Gaza R, Matthia D, Laramore D, Lehti J, et al. 2024. Space radiation measurements during the Artemis I lunar mission. *Nature* **634(8032)**: 48–52. <https://doi.org/10.1038/s41586-024-07927-7>.
- Gruhn C, Binimi M, Legrain R, Loveman R, Pang W, et al. 1982. Bragg curve spectroscopy. *Nucl Instr Meth Phys Res* **196**: 33–40. [https://doi.org/10.1016/0029-554X\(82\)90612-7](https://doi.org/10.1016/0029-554X(82)90612-7).
- Guo Z, Zhou G, Hu W. 2022. Carcinogenesis induced by space radiation: A systematic review. *Neoplasia* **32**: 100828. <https://doi.org/10.1016/j.neo.2022.100828>.
- Hassler DM, Zeitlin C, Wimmer-Schweingruber RF, Böttcher S, Martin C, et al. 2012. The radiation assessment detector (RAD) Investigation. *Space Sci Rev* **170**: 503–558. <https://doi.org/10.1007/s11214-012-9913-1>.

- Hassler DM, Zeitlin C, Wimmer-Schweingruber RF, Ehresmann B, Rafkin S, et al. 2014. Mars' surface radiation environment measured with the mars science laboratory's curiosity rover. *Science* **343**: 1244797–1244797. <https://doi.org/10.1126/science.1244797>.
- Hellweg CE, Berger T, Matthiä D, Baumstark-Khan C. 2020. *Radiation in space: relevance and risk for human missions*, Springer Briefs in Space Life Sciences. Springer, Cham. ISBN 978-3-030-46743-2. <https://doi.org/10.1007/978-3-030-46744-9>.
- Hollender L. 2019. *A Bayesian particle filter for particle identification*, Technical University of Munich.
- Klein PM, Parihar VK, Szabo GG, Zoldi M, Angulo MC, et al. 2021. Detrimental impacts of mixed ion radiation on nervous system function. *Neurobiology of Disease* **151**: 105252. <https://doi.org/10.1016/j.nbd.2021.105252>.
- Koontz SL, Boeder PA, Pankop C, Reddell B. 2005. The ionizing radiation environment on the International Space Station: performance vs. expectations for avionics and materials. *IEEE Radiation Effects Data Workshop*, IEEE, Seattle, WA, pp. 110–116. <https://doi.org/10.1109/redw.2005.1532675>.
- Lecun Y, Bottou L, Bengio Y, Haffner P. 1998. Gradient-based learning applied to document recognition. *Proc IEEE* **86(11)**: 2278–2324. <https://doi.org/10.1109/5.726791>.
- Longair MS. 2012. *High energy astrophysics*, Cambridge University Press, Cambridge, 3rd edn. ISBN 9780521756181.
- Looper MD, Mazur JE, Blake JB, Spence HE, Schwadron NA, et al. 2020. Long-term observations of galactic cosmic ray LET spectra in lunar orbit by LRO/CRA TER. *Space Weather* **18(12)**: e2020SW002543. <https://doi.org/10.1029/2020sw002543>.
- Losekamm MJ. 2025. *Scintillator-based particle detectors for radiation measurements on the international space station and for the exploration of the moon*. TUM.University Press. ISBN 9783958840935.
- Losekamm MJ, Berger T, Hinderberger P, Kaseman M, Kendelbacher T, et al. 2025. First results from the RadMap Telescope. In: *Proceedings of 39th International Cosmic Ray Conference – PoS (ICRC2025)*. <https://doi.org/10.22323/1.501.0073>.
- Losekamm MJ, Berger T, Hinderberger P, Kasemann M, Kendelbacher T, et al. 2023. Measuring Cosmic Rays with the RadMap Telescope on the International Space Station. In: *Proceedings of 38th International Cosmic Ray Conference – PoS(ICRC2023)*. <https://doi.org/10.22323/1.444.0099>.
- Losekamm MJ, Milde M, Pöschl T, Greenwald D, Paul S. 2017. A new analysis method using Bragg curve spectroscopy for a multi-purpose active-target particle telescope for radiation monitoring. *Nucl Instr Meth Phys Res Sec A* **845**: 520–523. <https://doi.org/10.1016/j.nima.2016.05.029>.
- Losekamm MJ, Paul S, Pöschl T, Zachrau HJ. . 2021. The RadMap Telescope on the international space station. In: *2021 IEEE Aerospace Conference (50100)*. IEEE, Big Sky, MT. <https://doi.org/10.1109/AERO50100.2021.9438435>.
- Losekamm MJ, Paul S, Pöschl T. 2024. Position-dependent light yield in short, coated SCSF-78 scin tillating fibers. *Radiation Measurements* **174**: 107116. <https://doi.org/10.1016/j.radmeas.2024.107116>.
- Loshchilov I, Hutter F. 2019. *Decoupled weight decay regularization*. <https://arxiv.org/abs/1711.05101>, 1711.05101.
- Mao XW, Pecaüt MJ, Gridley DS. 2021. Acute risks of space radiation. In: Young LR, Sutton JP (eds.), *Handbook of Bioastronautics*. Springer, Cham, pp. 263–276. ISBN 978-3-319-12190-1, https://doi.org/10.1007/978-3-319-12191-8_27.
- Milde M. 2016. *Development of a data-analysis framework for the multi-purpose active-target particle telescope*. Master's thesis, Technical University of Munich.
- Montesinos CA, Khalid R, Cristea O, Greenberger JS, Epperly MW, et al. 2021. Space radiation protection countermeasures in microgravity and planetary exploration. *Life* **11(8)**: 829. <https://doi.org/10.3390/life11080829>.
- Naito M, Kodaira S. 2022. Considerations for practical dose equivalent assessment of space radiation and exposure risk reduction in deep space. *Sci Rep* **12(1)**: 13617. <https://doi.org/10.1038/s41598-022-17079-1>.
- National Academies of Sciences, Engineering, and Medicine. 2021. *Space radiation and astronaut health: Managing and communicating cancer risks*. The National Academies Press, Washington, DC. ISBN 978-0-309-47966-0. <https://doi.org/10.17226/26155>.
- Navas S, Amsler C, Gutsche T, Hanhart C, Hernández-Rey JJ, et al. 2024. Review of particle physics. *Phys Rev D* **110(3)**: 030001. <https://doi.org/10.1103/PhysRevD.110.030001>.
- Norbury JW, Slaba TC, Sobolevsky N, Reddell B. 2017. Comparing HZETRN, SHIELD, FLUKA and GEANT transport codes. *Life Sci Space Res* **14**. <https://doi.org/10.1016/j.lssr.2017.04.001>.
- Parihar VK, Allen B, Tran KK, Macaraeg TG, Chu EM, et al. 2015. What happens to your brain on the way to Mars. *Sci Adv*. **1**: e1400256. <https://doi.org/10.1126/sciadv.1400256>.
- Prechelt L. 1998. Early stopping- but when?. In: *Neural networks: tricks of the trade* Orr GB, Müller K-R (eds.), Springer, Berlin, Heidelberg, pp. 55–69. ISBN 978-3-540-65311-0. https://doi.org/10.1007/3-540-49430-8_3.
- Pöschl T. 2022. *Modeling of the Galactic Cosmic-Ray antiproton flux and development of a multi purpose active-target particle telescope for cosmic rays*. Ph.D. thesis, Technical University of Munich. <https://nbn-resolving.org/urn:nbn:de:bvb:91-diss-20220706-1659625-1-4>.
- Pöschl T, Greenwald D, Losekamm MJ, Paul S. 2021. Measurement of ionization quenching in plastic scintillators. *Nucl Instr Meth Phys Res Sec A* **988(164865)**: <https://doi.org/10.1016/j.nima.2020.164865>.
- Raissi M, Perdikaris P, Karniadakis G. 2019. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J Comput Phys* **378**: 686–707. <https://doi.org/10.1016/j.jcp.2018.10.045>.
- Rankin JS, Bindi V, Bykov AM, Cummings AC, Della Torre S, Florinski V, Heber B, Potgieter MS, Stone ECZhang M. 2022. Galactic cosmic rays throughout the heliosphere and in the very local interstellar medium. *Space Sci Rev* **218(5)**: 42. <https://doi.org/10.1007/s11214-022-00912-4>.
- Reames DV. 2021. Solar energetic particles. *Lecture Notes in Physics*. Springer, Cham, 2nd edn. ISBN 978-3-030-66401-5. <https://doi.org/10.1007/978-3-030-66402-2>.
- Reames DV. 2022. Energy spectra vs. element abundances in solar energetic particles and the roles of magnetic reconnection and shock acceleration. *Solar Phys* **297(3)**: 32. <https://doi.org/10.1007/s11207-022-01961-2>.
- Rikhi R, Samra G, Arustamyan M, Patel J, Zhou L, Bungo B, Moudgil R. 2020. Radiation induced cardiovascular disease: An odyssey of bedside-bench-approach. *Life Sci Space Res*. **27**: 49–55. <https://doi.org/10.1016/j.lssr.2020.07.005>.
- Rodríguez-Pacheco J, Wimmer-Schweingruber RF, Mason GM, Ho GC, S'anchez-Prieto, S, et al. 2020. The energetic particle detector. *Astron Astrophys*. **642**. <https://doi.org/10.1051/0004-6361/201935287>.

- Scarselli F, Gori M, Tsoi AC, Hagenbuchner M, Monfardini G. 2009. The graph neural network model. *IEEE Transact Neural Netw* **20** (1): 61–80. <https://doi.org/10.1109/TNN.2008.2005605>.
- Slaba TC, Bahadori AA, Reddell BD, Singleterry RC, Cloudsley MS, Blattinig SR. 2017. Optimal shielding thickness for galactic cosmic ray environments. *Life Sci Space Res.* **12**: 1–15. <https://doi.org/10.1016/j.lssr.2016.12.003>.
- Stoffle N, Pinsky L, Kroupa M, Hoang S, Idarraga J, et al. 2015. Timepix-based radiation environment monitor measurements aboard the International Space Station. *Nucl Instr Meth Phys Res Sect A* **782**: 143–148. <https://doi.org/10.1016/j.nima.2015.02.016>.
- Szegedy CLiu W, Jia Y, Sermanet P, Reed S, et al. 2015. Going deeper with convolutions. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 1–9, Boston, MA, USA. <https://doi.org/10.1109/CVPR.2015.7298594>.
- van Laarhoven PJM, Aarts EHL. 1987. Simulated annealing: theory and applications *Mathematics and its applications*. Springer, Dordrecht. ISBN 978-90-481-8438-5. <https://doi.org/10.1007/978-94-015-7744-1>.
- Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, Kaiser L, Polosukhin I2017. Attention is all you need. In: *Proceedings of the 31st International Conference on Neural Information Processing Systems*, vol. 30 of *NIPS'17*, 6000–6010. Curran Associates Inc., Red Hook, NY.
- Vozenin MC, Alaghband Y, Drayson OGG, Piaget F, Leavitt R, et al.2024. More may not be better: enhanced spacecraft shielding may exacerbate cognitive decrements by increasing Pion exposures during deep space exploration. *Radiation Res.* **201**(2): 93–103. <https://doi.org/10.1667/RADE-23-00241.1.S1>.
- Walsh L, Schneider U, Fogtman A, Kausch C, McKenna-Lawlor S, et al. 2019. Research plans in Europe for radiation health hazard assessment in exploratory space missions *Life Sci Space Res.* **21**: 73–82. <https://doi.org/10.1016/j.lssr.2019.04.002>.
- Wimmer-Schweingruber RF, Yu J, Böttcher SI, Zhang S, Burmeister S, et al. 2020. The Lunar Lander Neutron and Dosimetry (LND) experiment on Chang'E 4. *Space Sci Rev* **216**(6). <https://doi.org/10.1007/s11214-020-00725-3>.
- Zeitlin C, Castro AJ, Beard KB, Abdelmelek M, Hayes BM, Johnson AS, Stoffle N, Rios RR. 2023. Results from the radiation assessment detector on the international space station: Part 1, the charged particle detector. *Life Sci Space Res* **39**: 67–75. <https://doi.org/10.1016/j.lssr.2023.01.003>.
- Zeitlin C, La Tessa C. 2016. The role of nuclear fragmentation in particle therapy and space radiation protection. *Front Oncol* **6**: 1–13. <https://doi.org/10.3389/fonc.2016>.

Cite this article as: Meyer-Hetling L, Losekamm MJ, Paul S & Pöschl T. 2026. A neural-network framework for tracking and identifying cosmic-ray nuclei in the RadMap Telescope. *J. Space Weather Space Clim.* **16**, 8. <https://doi.org/10.1051/swsc/2026001>.