Inference of cosmic-ray source properties by conditional invertible neural networks

Teresa Bister¹ , Martin Erdmann^{1,a} , Ullrich Köthe² , Josina Schulte¹

¹ Physikalisches Institut 3A, RWTH Aachen University, Otto-Blumenthal-Str., 52074 Aachen, Germany

² Interdisciplinary Center for Scientific Computing, University of Heidelberg, Im Neuenheimer Feld 205, 69120 Heidelberg, Germany

Received: 19 October 2021 / Accepted: 15 February 2022 / Published online: 24 February 2022

© The Author(s) 2022

Abstract The inference of physical parameters from measured distributions constitutes a core task in physics data analyses. Among recent deep learning methods, so-called *conditional invertible neural networks* provide an elegant approach owing to their probability-preserving bijective mapping properties. They enable training the parameter-observation correspondence in one mapping direction and evaluating the parameter posterior distributions in the reverse direction. Here, we study the inference of cosmic-ray source properties from cosmic-ray observations on Earth using extensive astrophysical simulations. We compare the performance of conditional invertible neural networks (cINNs) with the frequently used Markov Chain Monte Carlo (MCMC) method. While cINNs are trained to directly predict the parameters' posterior distributions, the MCMC method extracts the posterior distributions through a likelihood function that matches simulations with observations. Overall, we find good agreement between the physics parameters derived by the two different methods. As a result of its computational efficiency, the cINN method allows for a swift assessment of inference quality.

1 Introduction

Information about our universe is obtained from four different messengers: electromagnetic waves, cosmic rays, neutrinos, and gravitational waves. Large observatories exist worldwide for all these messengers. Each messenger carries information about its origin and, in addition, the history of its propagation to Earth as well as the distortion due to observational technology. To obtain dedicated information about properties of the universe, e.g., the source of the messenger, the influences of the other contributions (propagation, detector effects) have to be reversed by mathematical methods. This process is usually referred to as inference.

Among the four messengers, ultra-high-energy cosmic rays are characterized by two complicating effects in propagation and detection. First, the ionized nuclei can interact with cosmic background fields and may undergo directional deflections, lose energy, or decay sequentially on their way towards Earth. Second, as the nuclei enter the Earth's atmosphere, they cause showers of more than a billion secondary particles that allow only indirect determination of the primary's properties. Observatories enable an investigation of cosmic-ray frequency, the direction of arrival, energy, and cross-section, all of which can be used to unveil characteristics about the cosmic ray's origin.

In this work, we exemplarily investigate two methods for characterizing cosmic-ray sources from observations on Earth. Our focus is on correcting the measured energy distribution together with the distribution of the shower depths in the atmosphere for propagation effects. The measurements of today's observatories are so precise that bias and smearing of the detectors can be corrected comparatively easily [1,2]. Instead of the complete true energy distribution at the source, we aim to determine a set of characteristic quantities at the source, describing the set of different atomic nuclei (composition), the power of the energy spectrum (spectral index), and the maximum accelerator energy. This astrophysical scenario is kept very simple, but already shows the sensitivity of the measurements to source properties of cosmic rays [3]. This sensitivity is expressed in terms of posterior distributions of the characteristic source parameters.

Propagation of cosmic rays is associated with the above-mentioned interactions and nuclear decays, which follow each other sequentially in a random order. Propagation from source to observation is simulated by software such as CRPropa3 [4]. Because of the many random processes, the simulation can only be performed in a forward direction and inversion did not seem possible at first.

Therefore, the first determinations of the source properties were usually performed using forward simulations:

^a e-mail: erdmann@physik.rwth-aachen.de (corresponding author)

the characteristic quantities of the sources (source parameters) were modified until the measured distributions on Earth were reproduced by the forward simulations after cosmic-ray propagation. To avoid simulating each parameter setting individually, databases for the astrophysical scenario with different parameter settings were created. These databases contain weight factors for each measured spectrum and allow interpolation between the simulated parameter settings. Finally, to efficiently search for associated source parameters and their posterior distributions with the measured distributions, one could use Bayesian methods like the Markov Chain Monte Carlo (MCMC).

With regard to new developments in the context of neural networks, the inversion of the above-described propagation of cosmic rays is now possible after all. So-called normalizing flow networks consist of invertible blocks which enable network training in the forward direction and an evaluation in the backward direction, while preserving probability in both directions [5–7]. During the training process, a database of forward simulations in the one direction of computation is used, similarly as in Bayesian methods. The evaluation of a measured distribution by the trained network however happens in the backward computational direction, where the output consists of the posterior distributions of the source parameters. Recently, a number of inference methods based on deep learning methods have been developed and investigated (refer to the collection in [8]), including the unfolding of particle distributions [9] and the characterization of spatially correlated γ -ray maps [10].

In this paper, we present a normalizing flow network for the determination of cosmic-ray source parameters from measured distributions. The quality of this so-called *conditional invertible neural network* (cINN) is investigated in a comparative study with the traditional MCMC method.

This work is structured as follows: First, we introduce the astrophysical scenario and describe how we generated the database for the mapping between source parameters and observed distributions. Then, we briefly recall the MCMC procedure before going into detail about the functionalities of the cINN. In the subsequent section we then compare the results of the two methods and the computing resources used. Finally, we present our conclusions.

2 Astrophysical scenario and database

The astrophysical scenario used to investigate the performance of the inference methods is based on the findings of [3]. It consists of homogeneous sources which isotropically emit cosmic rays with different mass numbers A_{inj} and corresponding charge numbers Z_{inj} . Acceleration is possible up to a maximum rigidity R_{cut} where rigidity denotes energy divided by charge: $R = E/Z$. The cosmic-ray emission with

energy E_{inj} is described by a power law with spectral index γ and a broken exponential cutoff:

$$J_{\text{inj}}(E_{\text{inj}}, A_{\text{inj}}) = J_0 \cdot a(A_{\text{inj}}) \left(\frac{E_{\text{inj}}}{10^{18} \text{ eV}} \right)^{-\gamma} \cdot \begin{cases} 1 & Z_{\text{inj}} R_{\text{cut}} < E_{\text{inj}} \\ \exp\left(1 - \frac{E_{\text{inj}}}{Z_{\text{inj}} R_{\text{cut}}}\right) & Z_{\text{inj}} R_{\text{cut}} \geq E_{\text{inj}} \end{cases} \quad (1)$$

Here, J_0 is a normalization constant of the cosmic-ray flux, $a(A_{\text{inj}})$ denotes the injected fraction of the respective element with mass A_{inj} , defined below the cutoff.

Figure 1 shows an example of an injected spectrum at the source. Here, the source parameters have been adjusted to best-fit values identified by [3]. They are given by the spectral index $\gamma = 0.87$, rigidity cutoff $R_{\text{cut}} = 10^{18.62} \text{ V}$, fractions of nitrogen $a(\text{N}) = 88\%$, and silicon $a(\text{Si}) = 12\%$ following Tables 8 and 9 in [3] for a CRPropa3-based model, abbreviated with CTG. This set of parameters will hereafter be referred to as the *benchmark parameters*.

After the cosmic rays have been emitted at the source following the injected spectrum, they propagate through the universe, undergoing interactions before being detected at Earth. The mapping between the observables (the properties of the detected cosmic rays at Earth) and the source parameters at injection (γ , R_{cut} , $a(\text{H})$, $a(\text{He})$, $a(\text{N})$, $a(\text{Si})$ and $a(\text{Fe})$) is learned by the cINN. Hence, we first create a database to describe this mapping, which is then used for the network training. The same database is also used for the MCMC evaluation to find the set of source parameters that leads to the best agreement between the simulated observables and the measured data.

This simulation database is created in a modular way using one-dimensional CRPropa3 [4] simulations and a reweighting analogously to [3]. The database contains 10^6 simulated cosmic rays for each injected energy between 10^{18} eV and 10^{21} eV in bins of width $10^{0.02} \text{ eV}$ and each source distance, binned logarithmically between 1 Mpc and 5,700 Mpc in 118 bins. The simulations are performed for each representative element (hydrogen, helium, nitrogen, silicon, and iron) at source injection, and all secondary particles produced on the way to Earth are stored. For a homogeneous, isotropic, three-dimensional source population, a uniform distribution of the comoving source distances is expected before propagation effects, which we achieve by reweighting the simulated distances appropriately.

Upon detection at Earth, the cosmic rays are binned into energy bins e of width $10^{0.1} \text{ eV}$ between 10^{18} eV and 10^{21} eV and mass bins $A_{\text{det}} \in \{1, [2, 4], [5, 22], [23, 38] \text{ and } [39, 56]\}$. Neither the mass nor the charge of the arriving cosmic ray can be directly measured by today's cosmic-ray observatories. Therefore, the depth of the shower maximum X_{max} is used as an observable instead, as it relates inversely

Fig. 1 Injected spectrum using the best-fit source parameters from [3], following a power law with a broken exponential cutoff above a maximum rigidity as given by Eq. (1)

Fig. 2 Observed energy spectrum at Earth after injection following Fig. 1. As symbols with Poissonian errors, the benchmark simulation spectrum with $\mathcal{O}(70,000)$ events is shown, with different colors for the different detected element groups. The curves depict the prediction by the propagation database scaled to the same number of cosmic rays. The gray area marks the part of the energy spectrum below the threshold at $10^{18.7}$ eV, which is not part of the fit

to the cross-section in air which is connected to the primary cosmic-ray mass. From the detected energies and masses in the simulation database, the expected values for $X_{\max}$ can be calculated using Gumbel distributions G [11] and the EPOS-LHC hadronic interaction model [12], as in [3].

The mapping of the example source spectrum, depicted in Fig. 1, to Earth is shown in Figs. 2 and 3, where the detected energy spectrum and the detected $X_{\max}$ distributions can be seen. One can see that the $X_{\max}$ histogram is binned

Fig. 3 Depth of the shower maximum $X_{\max}$ distributions in energy bins. The binning is the same as for the energy spectrum with a combined bin above $10^{19.6}$ eV due to the smaller statistics. The benchmark simulation with $\mathcal{O}(2,700)$ events containing an $X_{\max}$ value is shown as symbols with Poissonian errors. The curves refer to the reweighted distributions from the propagation database, scaled to the same number of cosmic rays. The contributions by the different element groups are color-coded as in Fig. 2

two-dimensionally into $X_{\max}$ bins x between 550 g/cm^2 and 1050 g/cm^2 of width 20 g/cm^2 and similar energy bins $\tilde{e}$ as the observed spectrum with a combined bin above $10^{19.6}$ eV due to the smaller event statistic. We show not only the shape of the observables predicted by the database for the benchmark parameters as curves, but also one specific simulation from the database with the same number of events as in the data of the Pierre Auger Observatory. For this simulation, we additionally include bin-wise Poisson fluctuations. The spectrum in Fig. 2 contains $\mathcal{O}(70,000)$ events [13] and the shower depths histogram in Fig. 3 contains $\mathcal{O}(2,700)$ events [2], both above $10^{18.7}$ eV. This specific simulation will hereafter be referred to as the *benchmark simulation*.

In both the detected energy spectrum and the shower depth $X_{\max}$ distributions, the rapid decay of the flux as a function of the energy is visible. It is evident that the propagation has a substantial impact on the cosmic-ray energies, and other elements apart from the injected ones have emerged after interactions and decays. With increasing energy, the composition becomes heavier as expected from the rigidity-dependent acceleration at the source. The shape and location of the $X_{\max}$ distributions contain information on the composition: lighter cosmic rays can penetrate deeper into the atmosphere as the cross-section for air interactions is smaller, and the shower-to-shower fluctuations are larger than for heavy particles due to the superposition principle [14].

Altogether, we produced a database for the mapping from the injection at the source to the detected observables at Earth for different source parameters:

$$J_{\text{inj}}(E_{\text{inj}}, A_{\text{inj}}) \xrightarrow{\gamma, R_{\text{cut}}, a(A_{\text{inj}})} (E_{\text{det}}, X_{\max}^{\text{det}}) \quad (2)$$

In the following sections, we will describe how the MCMC and the cINN methods use this database for determining the source parameters. Afterward, both methods will be applied to the *benchmark simulation* in Sect. 5.

3 MCMC method for inference

Markov Chain Monte Carlo methods can be used to determine an unknown posterior probability density function (pdf) by sampling from it. The basis of parameter inference with MCMC methods is Bayes theorem, which connects the unknown posterior pdf $p(\theta|y)$ of the fit parameters θ given the data y to the likelihood of the data given the fit parameters $p(y|\theta)$ multiplied by the prior pdf $p(\theta)$:

$$p(\theta|y) = \frac{p(y|\theta) p(\theta)}{p(y)} \propto p(y|\theta) p(\theta) \quad (3)$$

Here, $p(y)$, which is often called the *Bayes integral*, is generally hard to calculate. For the inference of parameters, $p(y)$ does not have to be known, as it does not depend on the parameters θ and thus the shape of the posteriors can be determined without this normalization. The likelihood $p(y|\theta)$ corresponds to the *forward* direction described in Sect. 1, so it is usually known and can be calculated. The MCMC method does not require any derivatives or integrals to be calculated as is the case for example for minimizers [15].

In our case, the likelihood can be calculated using the propagation database described in Sect. 2, which predicts the energy spectrum and $X_{\max}$ distributions (corresponding to y) from the source parameters (corresponding to θ). Specifically, we use the same likelihood function $\mathcal{L} = \mathcal{L}_E \cdot \mathcal{L}_{X_{\max}}$ as in [3]. It contains a Poissonian likelihood $\mathcal{L}_E$ for the energy

spectrum, which compares the predicted spectrum calculated from the simulation database (event counts p in the energy bin e) to the benchmark simulation (corresponding event counts k):

$$\mathcal{L}_E = \prod_e \frac{(p^e)^{k^e}}{k^e!} \exp(-p^e) \quad (4)$$

The information on the energy spectrum is already used in the energy likelihood, so a multinomial likelihood $\mathcal{L}_{X_{\max}}$ is used for the $X_{\max}$ distributions:

$$\mathcal{L}_{X_{\max}} = \prod_{\tilde{e}} k^{\tilde{e}}! \prod_x \frac{(G^{\tilde{e},x})^{k^{\tilde{e},x}}}{k^{\tilde{e},x}!} \quad (5)$$

Here, $k^{\tilde{e},x}$ again describes the measured number of events in each energy bin $\tilde{e}$ and $X_{\max}$ bin x , and $G^{\tilde{e},x}$ represents the Gumbel distributions for the respective bin as in [3].

For the fit of the benchmark simulation, we use a sequential Monte Carlo based sampling algorithm within the PyMC3 framework [17, 18]. We chose the sequential algorithm over the Metropolis-Hastings algorithm because the latter required tuning of the proposal distribution but produced similar results. We let the sampler run for 5,000 steps in three different chains. Convergence is ensured by calculating the Gelman–Rubin coefficient $\hat{R}$ [19]. Additionally, the effective sample size is monitored and required to be $\gg 200$. As fit parameters θ , we use the two spectral parameters γ and $\log_{10}(R_{\text{cut}}/V)$ along with four representative values for the five elemental fractions utilizing the side condition of sum to unity with a simplex transformation [16]. We use the same flat bounded prior distributions for the source parameters as in [3]. Each chain runs for around twelve hours on a CPU.

4 Conditional invertible neural networks (cINN)

An alternative to MCMC sampling is a new method that uses deep learning techniques, introduced in [20] as an *invertible neural network* (INN) and extended to the *conditional* setup in [21] and [22]. The idea is based on the concept of normalizing flows, by which an invertible mapping is created between the physics parameters of interest θ , the source parameters in our case, and internal network parameters, referred to as latent variables z . The remarkable property of this bijective mapping is that it preserves probability.

4.1 Architecture of a cINN

To create the invertible mapping $z = f(\theta)$ between the internal network parameters z and the parameters of interest θ , *reversible blocks* [23] are used in our case. Figure 4 shows a

schematic sketch of one reversible block. It is based on the architecture introduced in [23] and [24] and can be evaluated in both the forward as well as the backward direction. In the forward direction, the parameters of interest vector θ is first split into two halves. The latent vector $z = [z_1, z_2]$, where z_1, z_2 correspond to the first and second half of the latent vector, is determined as follows:

$$\begin{aligned} z_1 &= \theta_1 \odot \exp(s_2(\theta_2)) + t_2(\theta_2) \\ z_2 &= \theta_2 \odot \exp(s_1(z_1)) + t_1(z_1) \end{aligned} \quad (6)$$

where $\odot$ refers to element-wise multiplication and the mappings $s_i()$ and $t_i()$ can be arbitrarily complicated and do not have to be invertible themselves. The mappings s_i, t_i are in general represented by additional neural networks. In the GLOW [24] setup, the mappings s_i, t_i are computed by a single subnetwork for each i . The inverse of the affine transformation can be easily obtained since the exponential function prevents division by zero and the subnetworks are always evaluated in the same direction:

$$\begin{aligned} \theta_2 &= (z_2 - t_1(z_1)) \odot \exp(-s_1(z_1)) \\ \theta_1 &= (z_1 - t_2(\theta_2)) \odot \exp(-s_2(\theta_2)) \end{aligned} \quad (7)$$

This network structure can be extended to the *conditional* case, in which the mapping between parameters of interest θ and latents z is learned under a specific condition. In our case, this condition are the observables y , the energy spectrum and the depth of shower maximum distributions. In the conditional network architecture, the condition is concatenated to the input of the subnetworks $s_i(\dots)$ and $t_i(\dots)$, which then become $s_i(\dots, y)$ and $t_i(\dots, y)$. This means that the output of the reversible blocks, the latent z , is now not only influenced by the parameters of interest θ , but also by the condition $y(\theta)$.

4.2 Training and loss function

During the training, the network learns to map each source parameter vector θ onto the corresponding latent vector z , taking into account the respective observables $y(\theta)$, which are calculated using the simulation database described in Sect. 2. For that, it is presented with multiple values for θ and the corresponding conditions $y(\theta)$. The combination of all these values then represents a *distribution* $p(\theta)$ of possible inputs and their respective conditions $p(y(\theta))$. Here, $p(\dots)$ refers to the collection of all elements of the respective quantity. This distribution of parameters of interest and conditions is then mapped onto a distribution of internal network parameters, the latent distribution $p(z)$, for which a specific form can be enforced by the loss function as described below.

For the inference of the source parameters, the trained network is evaluated in the backward direction. As the condition,

Fig. 4 Structure of the reversible block used for the conditional invertible neural network. It can be evaluated in two directions. The upper part shows the training mode or forward direction, the lower part displays the evaluation mode or backward direction

it is now presented with the measured data which represent one specific observation $\tilde{y}$. The full posterior distribution $p(\theta|\tilde{y})$ is then obtained by inserting the enforced latent distribution $p(z)$ into the trained network using the backwards direction, see Fig. 4. Thus, not only discrete values for the source parameters θ can be reconstructed with the network, but the whole posterior distribution is obtained.

A suitable loss function for the training of a cINN is introduced in [21]. The goal is to train a network that represents a mapping of a distribution in the latent space $p(z)$ to the true posterior space $p(\theta|y)$ (backward direction). Thus, we want to minimize the difference between the cINN posterior $p_\phi(\theta|y)$, where ϕ denote the network parameters, and the true posterior $p(\theta|y)$. The Kullback–Leibler divergence $\mathbb{KL}$ provides a measure on the difference of two probability distributions and is used as the basis of the loss L :

$$\begin{aligned} L &= \mathbb{KL}(p(\theta|y) \parallel p_\phi(\theta|y)) \\ &= \mathbb{E}_{\theta \sim p(\theta|y)} (\log p(\theta|y) - \log p_\phi(\theta|y)) \\ &= \text{const.} + \mathbb{E}_{\theta \sim p(\theta|y)} (-\log p_\phi(\theta|y)) \end{aligned} \quad (8)$$

Here, $\mathbb{E}$ denotes the expectation value, with parameter values θ sampled from the distribution $p(\theta|y)$. In the last step, the true posterior distribution is constant with respect to the network parameters and can thus be omitted in the loss function. Next, we apply the concept of probability conservation $p_\phi(\theta|y)d\theta = p(z)dz$ to transform the network posterior to the latent space:

$$\begin{aligned} L &= \mathbb{E}_{\theta \sim p(\theta|y)} (-\log p_\phi(\theta|y)) \\ &= \mathbb{E}_{\theta \sim p(\theta|y)} \left(-\log \left(p(z) \cdot \left| \det \left(\frac{\partial z}{\partial \theta} \right) \right| \right) \right) \\ &= \mathbb{E}_{\theta \sim p(\theta|y)} \left(-\log(p(z)) - \log \left(\left| \det \left(\frac{\partial z}{\partial \theta} \right) \right| \right) \right) \end{aligned} \quad (9)$$

The Jacobian $\partial z / \partial \theta$ of the reversible blocks (Fig. 4), which map from the physics parameter space θ to the latent space z via $z = f(\theta)$, turns out to be a triangular matrix. This simplifies the calculation of the determinant substantially. To see the argument, we decompose the transformation of the reversible block in Eq. (6) into two functions f_1 and f_2 . Using as an example f_1

$$f_1(\theta) = \begin{cases} z_1 = \theta_1 \odot \exp(s_2(\theta_2)) + t_2(\theta_2) \\ \theta_2 = \theta_2, \end{cases} \quad (10)$$

its Jacobian is calculated as follows:

$$\det \frac{\partial f_1(\theta)}{\partial \theta} = \det \begin{pmatrix} \frac{\partial z_1}{\partial \theta_1} & \frac{\partial z_1}{\partial \theta_2} \\ \frac{\partial \theta_2}{\partial \theta_1} & \frac{\partial \theta_2}{\partial \theta_2} \end{pmatrix} \quad (11)$$

$$= \det \begin{pmatrix} \text{diag}(\exp(s_2(\theta_2))) & \frac{\partial z_1}{\partial \theta_2} \\ 0 & \mathbb{I} \end{pmatrix} \quad (12)$$

$$= \prod_j \exp(s_{2,j}(\theta_2)) \quad (13)$$

Equivalently, the Jacobian of $f_2(\theta)$ is calculated, resulting in the total determinant:

$$\begin{aligned} \left| \det \left(\frac{\partial z}{\partial \theta} \right) \right| &= \left| \frac{\partial f_1(\theta)}{\partial \theta} \frac{\partial f_2(\theta)}{\partial \theta} \right| \\ &= \prod_j \exp(s_{2,j}(\theta_2)) \cdot \exp(s_{1,j}(z_1)) \\ &= \exp \left(\sum_j s_{2,j}(\theta_2) + s_{1,j}(z_1) \right) \end{aligned} \quad (14)$$

Note that the sum runs over the components j of the output of the mappings s_1 and s_2 , respectively. Now one can decide on the form of the distribution $p(z)$ that is enforced on the latent variables. To simplify the loss function as in [21], we choose a unit Gaussian distribution, denoted in one dimension by $p(z) = p(f(\theta)) = \exp(-f(\theta)^2/2)$. With the logarithmic functions in Eq. (9), this results in the following loss function for all parameter dimensions and the two subnetworks, averaged over m training datasets:

$$L = \frac{1}{m} \sum_{i=1}^m \left(\frac{1}{2} \|f(\theta_i)\|^2 - \sum_{l=1}^2 \sum_j s_{l,j} \right) \quad (15)$$

4.3 cINN for inference

The framework we use for the cINN for the inference of source parameters using the energy spectrum and the depth

of shower maximum distribution as observables is called Framework for Easily Invertible Architectures (FrEIA) [22] and is based on the PyTorch library [25]. The network consists of six reversible blocks with a GLOW [24] subnetwork structure. The mappings $s_{1,2}, t_{1,2}$ are represented by three fully connected layer transformations with an internal width of 256 with ReLU activation functions. Like in [22] and [26], prior to the exponential transformation of s_i , a non-linear transformation according to $\tilde{s} = 0.636 \alpha \arctan(s/\alpha)$ is applied to support stable training, here using $\alpha = 1.9$. After each reversible block, a permutation layer is used to enhance mixing between the different latent variables. The conditions y are the binned energy and shower maximum values (see Fig. 4).

The training data are created with the aforementioned database for mapping the source parameters θ , namely the spectral index γ , the maximum rigidity R_{cut} , and the five elemental fractions $a(\text{H})$, $a(\text{He})$, $a(\text{N})$, $a(\text{Si})$ and $a(\text{Fe})$, to the detected observables on Earth. The spectral index and the maximum rigidity have already been constrained by [3], which we use to limit our training data to reasonable pairs of (γ, R_{cut}) around the found minimum. The elemental fractions can be sampled uniformly using $(5 - 1) = 4$ representative sorted variables [16] to satisfy the condition that the sum equals one.

1,000,000 training samples and 100,000 validation samples with their corresponding energy spectrum and depth of shower maximum are generated. An interesting note is that a factor of 10 fewer training examples compromised the results. Before entering these into the network, they have to be preprocessed. The spectral parameters γ and $\log_{10}(R_{\text{cut}}/\text{V})$ are transformed to values between 0 and 1. For the energy spectrum, we use the aforementioned 17 energy bins e (Fig. 2). During the training, each of the energy bin contents is modified according to a Poisson distribution with the number of events corresponding to the typical event statistics measured by the Pierre Auger Observatory (see Sect. 2). This is important for the network to learn to evaluate different scenarios with the underlying statistical fluctuations. Afterward, the bin content of each energy bin is multiplied by E^3 . This helps flatten the steeply decreasing spectrum measured at Earth (Fig. 2), and its effectiveness in improving the reconstruction quality of the cosmic-ray source parameters was checked. The network is given this modified bin content of the 17 energy bins as conditional input, where the sum over all bins is normalized to one.

The depth of shower maximum distributions are binned into the bins $\tilde{z}, x$ as described in Sect. 2. Again, as for the energy spectrum, each bin content is altered using a Poisson distribution with reduced statistics, as expected by the measurements at the Pierre Auger Observatory. Here, by normalizing the X_{max} distribution in each energy bin to unity we remove the energy spectrum information that is already

used as a separate observable. To feed this (10 x 24) matrix into the network, we use flattening which results in a one-dimensional array with 240 entries.

The 17 energy bins are entered into the first 3 layers of the network as conditional input, and the $X_{\max}$ distributions into the last three layers. We verified that the information of both observables is indeed used by the network. The training of the network takes thirteen hours on a GPU and the evaluation of a single scenario can be completed in seconds.

5 Determination of source parameters with the cINN compared with the MCMC method

In the following, we evaluate the *benchmark simulation* presented in Sect. 2 using the MCMC and the cINN methods. Both methods yield posterior distributions of the fit parameters, which can be used to determine the most probable value and the uncertainty on the parameters by the 68 % interval as well as unveil correlations between the parameters.

Even though both methods can be used to characterize the posteriors, they use inherently different mathematical bases for it. The MCMC uses a likelihood, which is engineered according to the experimental statistics of the observables measured at Earth, as presented in Sect. 3. The difference between the predicted and measured observables is minimized when maximizing the likelihood, and the MCMC algorithm ensures asymptotic convergence of the samples to the posteriors [15].

The cINN, on the other hand, uses a likelihood-free inference where a loss function minimizes the distance between the true posterior distributions of the source parameters and the posterior distribution as predicted by the network. It does not ensure agreement of the network prediction with the measured observables at Earth, which is only implicitly achieved by the agreement of the source parameter posteriors.

Figure 5 shows the posterior distributions for the spectral index γ and the rigidity cutoff R_{cut} from the MCMC in the upper part and from the cINN in the lower part, respectively. The posterior distributions of γ and $\log_{10}(R_{\text{cut}}/V)$ are generally similar for both methods: the one-dimensional histograms show a symmetric distribution with the *true* benchmark parameters within one standard deviation of the posterior mean. Both methods find a positive correlation between the parameters, shown in the two-dimensional lower histograms. One can see that both methods slightly underestimate the two parameters. It was checked that this finding depends on the specific scenario chosen; hence, using different Poissonian variations of the observables leads to slightly shifted posteriors for both methods.

One can see that the widths of the posteriors are slightly larger for the cINN than for the MCMC. For the cINN the widths of the posteriors depend on the size of the training

Fig. 5 Posterior distributions obtained with the MCMC (*upper*) and the cINN (*lower*) methods of the spectral index γ and the rigidity cutoff $\log_{10}(R_{\text{cut}}/V)$. The mean of the distributions is shown by the black solid curve and the *true* underlying value of the *benchmark simulation* is marked by the red curve. In the lower left of the plots of both methods, the mean and standard deviation of both parameters are shown, and for the MCMC also the Gelman–Rubin index $\hat{R}$

dataset and the training time of the network, therefore it must be ensured that the network is trained with a sufficiently large training dataset for a long enough time. The same applies for the MCMC, where a sufficient number of chains have to be run with enough sampling steps. This can be ensured by keeping the Gelman–Rubin coefficient $\hat{R}$ [19] close to one, which is shown in the figures. Also, we trained multiple networks with different initializations and compared the posteriors, which appear to be quite similar. The example shown here is obtained from the cINN with the lowest validation loss value.

Figure 6 shows the posterior distributions for the composition fractions of the five representative elements. One can see that both methods again lead to similar results in gen-

Fig. 6 Posterior distributions of the composition fractions obtained with the MCMC (upper) and the cINN (lower) methods. The mean of the distributions is shown by the black solid curve and the *true* underlying value of the *benchmark simulation* is marked by the red curve. In the lower left of the plots of both methods, the mean and standard deviation of both parameters are shown and for the MCMC also the Gelman–Rubin index $\hat{R}$

eral. Both are able to identify that the composition at the source is dominated by nitrogen and silicon and that the iron fraction is tiny. The posteriors of the lighter elements are very broad, ranging down to zero contribution. This indicates that those parameters are more difficult to determine, which is attributable to the fact that the cutoff energy (cf. Fig. 1) for hydrogen is at $E_{\text{cut}} = Q_e \cdot R_{\text{cut}} = 10^{18.62}$ eV (Q_e denotes the elementary charge) and the one for helium is at $2Q_e \cdot R_{\text{cut}} = 10^{18.92}$ eV, which means that almost no light primaries are expected to survive above the energy threshold of the observables at $10^{18.7}$ eV.

Fig. 7 Modeled energy spectra for the cINN (dashed) and the MCMC (dotted) using the predicted source parameters as given in Figs. 5 and 6. Both methods are able to find a set of source parameters that describe the benchmark simulation energy spectrum, depicted as black symbols including Poissonian error bars. Also, the individual element contributions predicted at Earth, shown in different colors for different mass groups, agree with the ones of the benchmark simulation in Fig. 2

The correlations between the fractions look similar using both methods. The histograms for the MCMC look less smooth than for the cINN; this is due to the fact that we have several chains combined into one posterior. The Gelman–Rubin coefficient $\hat{R}$ is close to one for all parameters, indicating a convergence of the chains, but the posteriors could still become slightly smoother with more sampling steps or more chains.

In general, we observe very similar posterior distributions generated by the two different methods. One can additionally compare the reconstructed observables as shown in Fig. 7 exemplarily for the energy spectrum. In comparison with Fig. 2, one can see that both methods yield good agreement between the modeled spectrum and the benchmark simulation. The same applies to the X_{max} histograms (not shown here). The level of agreement can be quantified by calculating the deviance D [3], which is two times the negative log-likelihood ratio of the model, and the *saturated* model that would describe the data perfectly. For the two observables, the energy spectrum and the X_{max} distributions, we use the likelihood functions which are used directly for the MCMC sampling, as given in Sect. 3. For the MCMC we achieve a deviance of $D^{\text{MCMC}} = D_E + D_{X_{\text{max}}} = 14.7 + 123.7 = 138.4$ and for the cINN $D^{\text{cINN}} = 14.3 + 124.7 = 139.0$. The value of the deviance should be in the range of the number of degrees of freedom (number of non-zero bins = 124). Thus, we obtain reasonable, quite similar values for both methods, indicating a good description of the observables.

5.1 Stability of the cINN results

To evaluate the performance of the new cINN method in more detail, we evaluated not only this single simulation, but also a test dataset of 10,000 simulations. This extensive test was performed only with the cINN and not with the MCMC, as the computing time for the MCMC exceeds reasonable times within our computational resources. Figure 8 shows two-dimensional histograms of the source parameters. The true value of each simulation in the test set is shown on the x -axis, the mean of the cINN posterior is indicated on the y -axis.

For the spectral index γ and the cutoff rigidity R_{cut} we see good agreement between the mean estimate and the true simulation value which can be confirmed by the small normalized root mean square error

$$\text{NRMSE} = \frac{\sqrt{\frac{\sum_{i=1}^N (\langle \theta_i \rangle - \theta_i)^2}{N}}}{\max(\theta) - \min(\theta)}, \quad (16)$$

which is 0.014 for the spectral index γ and 0.018 for the rigidity cutoff. No far outliers are found and the small widening of the distribution for larger values of R_{cut} is due to the degeneracy of high rigidity values for larger spectral indices γ as revealed by the previous data analysis presented in [3].

For the composition fractions, the lighter element fractions cannot be reconstructed, as described above. In this case, the cINN mostly just predicts the average value for five elements 1/5. The NRMSE value is larger than 0.15 for the light elements. For the heavier elements, the reconstruction ability improves and the NRMSE decreases to 0.082 for iron.

Additionally, the widths of the posterior distributions can be examined. For this, we calculate the median calibration error e_{cal} following [26]. The calibration error is defined as the difference between a confidence level q and the actual fraction of observations $q_{\text{inliers}} = N_{\text{inliers}}/N$ of the whole test dataset of size N within that q -confidence interval. We calculate the calibration error for a range of confidence intervals (0.01, 0.99) in 0.01 steps and take the median over the absolute values. Appropriate posterior distributions would result in values close to zero. We reach a median calibration error of 0.001 to 0.006 for all parameters as given in Fig. 8, confirming suitable widths of the posterior distributions from the cINN for the whole test dataset. This also applies to the light element fractions with often too large posterior means, indicating that the cINN predicts a suitably large uncertainty for these unrecoverable parameters, as was also seen in Fig. 6.

6 Conclusion

We presented the application of a new method using deep learning techniques, the so-called conditional Invertible Neu-

Fig. 8 2D histograms of the 7 source parameters γ , R_{cut} and $a(\text{H,He,N,Si,Fe})$ using 10^4 test datasets. The true value of the simulation is shown on the horizontal axis, the mean of the cINN posterior on the vertical axis. The gray straight line represents the position of perfect agreement

ral Network (cINN), to a scenario from astroparticle physics constraining characteristic cosmic-ray source parameters. Using the energy spectrum and shower depth distributions on Earth as observables, the network is able to assess posterior distributions of the source parameter space. These allow not only a best-fit value to be estimated, but also uncertainties, possible degeneracies, and correlations between the parameters to be unveiled. The accuracy of the approach has been tested and verified to provide promising results for a large phase space of the source parameters. Given the speed of the method, it is easily possible to extend the scenario to more observables and more characterizing parameters of cosmic-ray sources. This allows for potential future applications of the technique.

Additionally, we compared the method with the conventional MCMC method on a specifically simulated scenario similar to the measurements of the Pierre Auger Observatory. The two inference methods use rather different techniques. While the cINN method aims at matching the true and the predicted distributions of the source parameters, the MCMC method is based on a likelihood analysis where the simulations are adapted to the observed data distributions. Nevertheless, we found good reconstruction of the source parameters within one standard deviation for both methods and an overall agreement of the posterior distributions.

Training of the cINN takes approximately thirteen hours while the evaluation of several test scenarios can be done instantaneously. For the MCMC however, each chain runs for around twelve hours, and several chains are needed to ensure convergence. Each new test scenario has to be evaluated single-handedly with the MCMC, making the cINN significantly more computationally effective overall.

Acknowledgements This work is supported by the Ministry of Innovation, Science, and Research of the State of North Rhine-Westphalia, and by the Federal Ministry of Education and Research (BMBF).

Data Availability Statement This manuscript has no associated data or the data will not be deposited. [Authors' comment: The simulated data used here are for demonstration purposes of the methods only.]

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>. Funded by SCOAP³.

References

1. A. Aab et al. [Pierre Auger Collaboration], *Phys. Rev. D* **102**, 062005 (2020). <https://doi.org/10.1103/PhysRevD.102.062005>. [arXiv:2008.06486](https://arxiv.org/abs/2008.06486)
2. A. Yushkov [for the Pierre Auger Collaboration], *Proc. 36th ICRC, Madison*, PoS (ICRC2019), 482 (2019). [arXiv:1909.09073](https://arxiv.org/abs/1909.09073)
3. A. Aab et al. [Pierre Auger Collaboration], *J. Cosmol. Astropart. Phys.* **04**, 038 (2017). <https://doi.org/10.1088/1475-7516/2017/04/038>. [arXiv:1612.07155](https://arxiv.org/abs/1612.07155)
4. R. Alves Batista et al., *J. Cosmol. Astropart. Phys.* **1605**, 038 (2016). <https://doi.org/10.1088/1475-7516/2016/05/038>. [arXiv:1603.07142](https://arxiv.org/abs/1603.07142)
5. E.G. Tabak, E. Vanden-Eijnden, *Commun. Math. Sci.* **8**, 217–233 (2010). <https://doi.org/10.4310/CMS.2010.v8.n1.a11>
6. I. Kobyzev, S.J.D. Prince, M.A. Brubaker, *IEEE Trans. Pattern Anal. Mach. Intell.* **43**, 3964–3979 (2020). <https://doi.org/10.1109/TPAMI.2020.2992934>. [arXiv:1908.09257](https://arxiv.org/abs/1908.09257)
7. G. Papamakarios, E. Nalisnick, D. Jimenez Rezende, S. Mohamed, B. Lakshminarayanan, *J. Mach. Learn. Res.* **22**, 57 (2021). <http://jmlr.org/papers/v22/19-1028.html>. [arXiv:1912.02762](https://arxiv.org/abs/1912.02762)
8. M. Feickert, B. Nachman, [arXiv:2102.02770](https://arxiv.org/abs/2102.02770)
9. M. Bellagente, A. Butter, G. Kasieczka, T. Plehn, A. Rousset, R. Winterhalder, L. Ardizzone, U. Köthe, *SciPost Phys.* **9**, 074 (2020). <https://doi.org/10.21468/SciPostPhys.9.5.074>. [arXiv:2006.06685](https://arxiv.org/abs/2006.06685)
10. S. Mishra-Sharma, K. Cranmer, [arXiv:2110.06931](https://arxiv.org/abs/2110.06931)
11. M. De Domenico et al., *J. Cosmol. Astropart. Phys.* **1307**, 050 (2013). <https://doi.org/10.1088/1475-7516/2013/07/050>. [arXiv:1305.2331](https://arxiv.org/abs/1305.2331)
12. T. Pierog et al., *Phys. Rev. C* **92**, 034906 (2015). <https://doi.org/10.1103/PhysRevC.92.034906>. [arXiv:1306.0121](https://arxiv.org/abs/1306.0121)
13. V. Verzi [for the Pierre Auger Collaboration], *Proc. 36th ICRC, Madison*, PoS (ICRC2019), 450 (2019). [arXiv:1909.09073](https://arxiv.org/abs/1909.09073)
14. K.-H. Kampert, M. Unger, *Astropart. Phys.* **35**, 10 (2012). <https://doi.org/10.1016/j.astropartphys.2012.02.004>. [arXiv:1201.0018](https://arxiv.org/abs/1201.0018)
15. D. Hogg, D. Foreman-Mackey, *ApJS* **236**, 11 (2018). <https://doi.org/10.3847/1538-4365/aab76e>. [arXiv:1710.06068](https://arxiv.org/abs/1710.06068)
16. S. Onn, I. Weissman, *Ann. Oper. Res.* **189**, 331–342 (2011). <https://doi.org/10.1007/s10479-009-0567-7>
17. S.E. Minson, M. Simons, J.L. Beck, *Geophys. J. Int.* **194**(3), 1701–1726 (2013). <https://doi.org/10.1093/gji/ggt180>
18. J. Salvatier, T.V. Wiecki, C. Fonnesbeck, *PeerJ Comput. Sci.* **2**, e55 (2016). <https://doi.org/10.7717/peerj-cs.55>
19. A. Gelman, D.B. Rubin et al., *Stat. Sci.* **7**, 457 (1992). <https://doi.org/10.1214/ss/1177011136>
20. L. Ardizzone et al., [arXiv:1808.04730](https://arxiv.org/abs/1808.04730)
21. S.T. Radev et al., [arXiv:2003.06281](https://arxiv.org/abs/2003.06281)
22. L. Ardizzone et al., [arXiv:1907.02392](https://arxiv.org/abs/1907.02392)
23. L. Dinh et al., [arXiv:1605.08803](https://arxiv.org/abs/1605.08803)
24. D.P. Kingma, P. Dhariwal, [arXiv:1807.03039](https://arxiv.org/abs/1807.03039)
25. A. Paszke et al., *Adv. Neural Inf. Process. Syst.* **32**, 8024 (2019). <https://proceedings.neurips.cc/paper/2019/hash/bdbca288fee7f92f2bfa9f7012727740-Abstract.html>. [arXiv:1912.01703](https://arxiv.org/abs/1912.01703)
26. V.F. Koll et al., *MNRAS* **499**(4), 202 (2020). <https://doi.org/10.1093/mnras/staa2931>. [arXiv:2007.08391](https://arxiv.org/abs/2007.08391)