**Regular Article - Experimental Physics** 



# Background rejection in atmospheric Cherenkov telescopes using recurrent convolutional neural networks

R. D. Parsons<sup>1,3,a</sup>, S. Ohm<sup>2</sup>

<sup>1</sup> Max-Planck-Institut für Kernphysik, P.O. Box 103980, 69029 Heidelberg, Germany

<sup>2</sup> DESY, 15738 Zeuthen, Germany

<sup>3</sup> Institut für Physik, Humboldt-Universität zu Berlin, Newtonstr. 15, 12489 Berlin, Germany

Received: 23 October 2019 / Accepted: 20 April 2020 / Published online: 6 May 2020 © The Author(s) 2020

Abstract In this work, we present a new, high performance algorithm for background rejection in imaging atmospheric Cherenkov telescopes. We build on the already popular machine-learning techniques used in gamma-ray astronomy by the application of the latest techniques in machine learning, namely recurrent and convolutional neural networks, to the background rejection problem. Use of these machine-learning techniques addresses some of the key challenges encountered in the currently implemented algorithms and helps to significantly increase the background rejection performance between 100 GeV and 100 TeV energies. We apply these machine learning techniques to the H.E.S.S. telescope array, first testing their performance on simulated data and then applying the analysis to two well known gamma-ray sources. With real observational data we find significantly improved performance over the current standard methods, with a 20-25% reduction in the background rate when applying the recurrent neural network analysis. Importantly, we also find that the convolutional neural network results are strongly dependent on the sky brightness in the source region which has important implications for the future implementation of this method in Cherenkov telescope analyses.

# 1 Introduction

Historically, one of the largest challenges in ground-based gamma-ray astronomy with imaging atmospheric Cherenkov telescopes (IACTs) is the identification and rejection of hadron-initiated air showers based on shower images. This is due to the extreme outnumbering of gamma-ray induced air showers by those from cosmic-ray hadrons (a factor  $10^4$  in even the brightest fields of view). Therefore, in order to detect most sources, the difference in development of hadronic and

electromagnetic air showers, which causes a corresponding difference in the observed IACT camera image, must be used to discriminate gamma-ray candidates from the mostly hadronic background. Although cosmic-ray electrons are a much less abundant particle species in the cosmic-ray spectrum ( $\sim 10^3$  cosmic-ray hadrons per cosmic-ray electron at 1 TeV), they constitute an irreducible background for the classification because of the electromagnetic nature of the air shower the initiate.

Traditionally, this background rejection has been performed through the use of Hillas Parameters [1], which parameterise the cleaned camera images using their second moments. The so-called tail cuts cleaning requires a pixel to have an intensity exceeding a threshold  $t_1$  and a neighbour exceeding a threshold  $t_2$  [2]. Images from hadronic showers appear to be both, longer and wider, than those from gamma rays, due to the larger transverse momentum transfer within the hadronic interactions in the shower cascade. By placing cuts on the image width and length, the first generation of very-high-energy (VHE;  $0.1 \text{ TeV} \le E \le 50 \text{ TeV}$ ) gammaray sources were detected [3,4]. In the following generation of gamma-ray observatories the use of multiple telescopes to image air showers from different directions improved their characterisation and the classification capability. The information from more than one image of the same shower has been combined by using the Hillas parameters from multiple telescopes to construct mean scaled parameters [2], for example mean scaled width (MSCW) defined below:

$$MSCW = \sum_{i=1}^{n} \frac{w_i - \langle w \rangle}{\sigma_w} \cdot \frac{1}{n}$$
(1)

where w is the Hillas width,  $\langle w \rangle$  is the expected width determined from lookup tables derived from Monte Carlo air-shower simulations and  $\sigma_w$  is the expected RMS of the width. In this way the Hillas parameters can be combined into

<sup>&</sup>lt;sup>a</sup> e-mail: daniel.parsons@physik.hu-berlin.de (corresponding author)

a single parameter with a mean of zero and a standard deviation of one. Background rejection can then be achieved by placing cuts on both MSCW, and mean scaled length (MSCL) parameters. Note that lookup tables are typically produced for a broad phase-space range, covering a variety of observing conditions and telescope configurations.

The use of Hillas parameters for background rejection has proven extremely effective in the rejection of hadronic background, however, it is clear that these parameters do not effectively contain all the information from individual camera images like asymmetries or pixel-wise information. By construction, mean-scaled parameters also average over multiple telescopes employing different weightings. Necessarily, this leads to a loss of information on the separation power stored in individual images. Another limitation of the classical mean-scaled parameter based box cuts is that they do not take into account linear and non-linear correlations between input parameters. Machine-learning techniques such as random forests [5], boosted decision trees [6,7] or neural networks [8] have been developed and successfully applied to data taken with the third generation of IACTs based on telescope- and event-wise input. These algorithms do exploit correlations between variables, but cannot compensate the information lost in the construction of the input parameters. The same is true for the much more powerful state-of-the-art likelihood methods that base the classification on parameters from pixel-wise comparisons between the recorded shower images and the expected image from a semi-analytical [9] or template-base model [10]. Attempts to include the maximum possible information into the background rejection with IACTs are therefore clearly an important topic.

The vast progress made in recent years in the field of deep learning techniques allows to apply new data-driven feature extraction methods to scientific data. For instance, convolutional neural networks (CNNs, [11]) extract spatial features or identify objects in image data using complex, hierarchical artificial neural networks (e.g. [12]). Recurrent neural networks (RNNs), on the other hand, are designed to classify sequenced or temporally correlated data. Methods employing long short-term memory (LSTM) units [13] are particularly powerful in areas such as speech recognition [14] and are nowadays applied in astronomy in different variants (e.g. [15,16]). Lately the development of combined convolutional and recurrent neural networks (CRNNs) have been explored to classify transient objects in optical astronomy (e.g. [17]).

The application of a convolutional, recurrent network structure has also been demonstrated on IACT simulations [18] with encouraging results. However, the expected improvements were not seen when applied to H.E.S.S. data. This demonstrates the challenges regarding the stability and reputability when deploying those advanced analysis methods on experimental data. Tackling those aspects is one of the primary focus of this work. This paper is organised as following: in the first section we motivate the usage of deep neural networks to address the apparent information-loss problem in classical parameterbased IACT classifiers and introduce the deep neural network designs used for this study. The following section explains how the networks are constructed and trained, followed by a section addressing the performance of the network with Monte-Carlo events. Finally, we will test the performance and stability of the network against real H.E.S.S. data.

# 1.1 Convolutional input layers

If we wish to move beyond the paradigm of image parameterisation as used in state-of-the-art machine-learning techniques such as [19] and [8], using a multi-layer perceptron (MLP) is no longer sufficient. Although an MLP could be created using individual image pixels as input, such a network would require an extremely large training data set in order to properly classify data. This is due to two effects: firstly the maximally connected nature of the MLP (each neuron in a layer is connected to each neuron in the proceeding layer) would create a huge number of parameters to fix in the case of even a very coarsely pixelated image. Secondly all spatial information of the pixels relative to their neighbouring pixels would be lost, therefore the network would not be stable against the translation of a given classifying feature through the image.

Convolutional neural networks offer a way around this problem by instead extracting the information from the shower image in the Cherenkov camera itself, and by applying a series of convolutional kernels on it. The result of this application is a 2-dimensional feature map of the image. Typically in such networks the most important features are selected (and the dimensionality reduced) through the use of a *max pooling* layer, where the maximum value of the feature map in a given 2D window is selected. The results of this pooling can then be passed through further convolutional and pooling layers, allowing features on larger scales than the convolutional kernel to be extracted. Different CNN architectures have been successfully implemented and applied in particle and astroparticle physics (e.g. [20-22]).

# 1.2 Recurrent network layers

Often in machine learning problems the classification of a number of sequential correlated images (for example images of the same shower seen from different perspectives) is required. Again, in this case the construction of a traditional network structure with each image of the series as an input would introduce an unsatisfactory number of free parameters to the network. Additionally, as the same features are being searched for in all images, such separate input is counterproductive.

To counter this problem, recurrent layers are constructed in such a way that the correlated inputs can be fed through the same network in sequence with each input modifying the behaviour of the network for all subsequent inputs. In this way the network is able to process inputs while retaining knowledge of information, which has already been seen. Typically, recurrent network implementations such as the LSTM contain mechanisms to "forget" older inputs, such that the sum of potentially extreme former inputs does not lead to a runaway of the network weights to infinity. One example of a RNN in particle physics, is the work by [23], where identification of b quarks based on particle jet properties in ATLAS at the LHC is performed. In this work, different tracks associated to the same jet are sequentially input into the RNN, which learns about the correlations between tracks associated to the same vertex.

# 2 Neural network design

In order to quantify the performance of the recurrent neural network on both, simulations and IACT data, three networks were designed using different inputs and network topologies (see Figs. 1 and 2). All networks were created using the Keras [24] python-based machine-learning interface, using *TensorFlow* [25] as the back-end module.

## 2.1 Mean scaled input

Firstly a simple multi layer perceptron (MLP) was created using mean scaled parameters as input. This network was created to provide a baseline comparison for the recurrent networks. The input parameters for this network are the mean scaled width and length of the shower in comparison to both simulated gamma-ray (MSCL, MSCL) as well as back-



Fig. 1 Network topology of the mean scaled (left) and Hillas-based recurrent (right) networks, stacked boxes show regions of the network where telescope inputs are processed in parallel. More details given in Sect. 2



Fig. 2 Network topology of the convolutional network, stacked boxes show regions of the network where telescope inputs are processed in parallel

ground events (MSCWO, MSCLO), the reconstructed depth of the shower maximum in the atmosphere  $X_{max}$  and the  $\Delta E/E$  parameter, calculated as the averaged spread in energy reconstruction between. These input parameters are similar to those used in the BDT method of [19], and hence should perform similarly to the technique already implemented in the H.E.S.S. framework. However, a retraining is performed to ensure consistency of the MVA tools used and the training data set. In general this network performs similarly to that described in [19].

# 2.2 Parametric input

The second network created also used the Hillas parameters for the input information, however, rather than combining these parameters using the aforementioned mean scaled method instead the unscaled parameters are used as input and combined within the network using a recurrent layer. The inputs to this network are the Hillas width and length, sum of pixel amplitudes in the cleaned image, reconstructed impact parameter, the displacement of the image centroid from the reconstructed source position and the distance of the image centroid from the camera field of view.

## 2.3 Image input

Finally, in order to quantify the effects of adding more image information, a network was created, which also takes the camera images as input. However, as most convolutional algorithms are created to operate on a regular image of square pixels, rather than the hexagonal arrangement used in many IACT cameras (such as those of the H.E.S.S. array) some preprocessing must be performed. Firstly the camera images are cleaned using the standard split level tail cut scheme used in H.E.S.S. [2] and then four rows of additional pixels are added to the edge of the cleaned image (in the same procedure used by [10]). A linear interpolation is then performed between these pixels, using Delauney triangulation, allowing them to be mapped onto a square grid with pixel size of  $0.05^{\circ}$  and a total width of 5°. Although the image cleaning step is not strictly necessary for the convolutional analysis, the reduction in the number of image pixels greatly increases the speed of the interpolation step and reduced the amount of data stored. Additionally the removal of noisy pixels not in the vicinity of the shower image may help to produce a more stable result. Finally the image is rescaled such that the image intensity lies between 0 and 1 (with negative intensity pixels set to 0). This rescaling was found to greatly ease training and although it does remove some normalisation information from the network the amplitude information is added to the network as part of the parameteric input layer (Fig. 3).

Once this preprocessing is complete the data are passed to the convolutional neural network pictured in Fig. 2. This network takes the interpolated images as input, passing them through two steps of convolution and max pooling and then flattening the resulting feature map into one dimension. This flattened information was then passed through a densely connected neural network. To avoid significant over-training of the network, dropout layers were added to this Sect. [26]. During the training of the network these layers randomly remove a fraction of the network connection (in this case 50%) to ensure no individual connections can dominate the network. This convolutional section was purposefully designed to be rather simple (in comparison with cutting-edge image classification algorithms) to try to avoid the situation where classification power is dependent on subtle image features present only in simulated data. In this case we sacrifice some potential performance for stability. In general this convolutional section is rather similar to that presented in [18], however more information is then added to the network.

The result of this convolutional section is then concatenated with the densely connected layer of the parametric network described earlier and fed into a recurrent layer and ultimately to the output layer. In principle concatenating the results of the network in this fashion is not required to perform image classification, however, it is useful in this case for two main reasons. Firstly, it allows us to assess the rejection power of the information added to the network by the image data over the parametric. Secondly, and most crucially, it provides information to the network which cannot be easily extracted from the camera images, such as the distance of the telescopes from the shower core. Given a sufficiently large training data set such information could be included in the network implicitly, by learning the locations of the telescopes in addition to how to perform event reconstruction. However, this would significantly increase training time and could potentially introduce systematic effects to the results.

# 3 Network training

The three networks were trained using simulated data generated from the CORSIKA Monte-Carlo air shower simulation code [27] and the *sim\_telarray* telescope and camera simulation [28]. This simulation chain has been proven within the H.E.S.S. and CTA collaborations to provide an accurate representation of the telescope data. To train the network, a sample of simulated gamma rays and protons was created which simulates the performance on the phase 1 H.E.S.S. array (4 × 12 m telescopes) at 70% of their design optical efficiency. Gamma-ray and proton events were simulated in a diffuse cone of opening angle 2.5° with an energy spectrum of  $E^{-1.5}$  and an energy range covering from below 100 GeV to over 100 TeV (dependent of the simulated species).

The simulated events were then passed through the H.E.S.S. Analysis Program (HAP) and the standard event selection cuts applied, requiring at least two camera images in an event over 60 photoelectrons and with an image centroid less than 2° from the camera centre [19]. The remaining events were then reconstructed using the standard H.E.S.S. Hillas parameter based shower reconstruction and events



**Fig. 3** Example images of gamma-ray and proton events both before (upper panels) and after (lower panels) the image pre-processing step described. Hillas parameter ellipses are overlaid on top of the camera

images. Note, in the network input an additional random rotation is added to the processed images but this is omitted here for clarity

reconstructed as lying within the central  $1^{\circ}$  from the camera centre passed to the neural network (to ensure the training sample is not dominated by highly truncated large offset events). This event selection resulted in a total of around 100,000 gamma-ray and proton events. These events were then split into four energy bins (0.1–0.4, 0.4–1, 1–5 and

5–100 TeV), with energy ranges chosen as a compromise between keeping a small range to ensure similar events are compared and having sufficient event statistics to perform the training.

The network was then trained in these 4 energy bins using 80% of the events as the training sample and the remaining



Fig. 4 Distribution of the CRNN classifier for a sample of gammaray and proton events. Events are re-weighted to represent an energy spectrum of  $E^{-2}$  for both signal and background

20% as an independent validation sample used to modify the network learning rate during training. A single loss function was introduced by computing the cross entropy of the two event categories after passing through each network type. A maximum of 100 training iterations was allowed, however an early stopping criterion was implemented to terminate the training in the case that the loss function (evaluated on the validation sample) shows no improvement. Typically the network training was terminated after 40–50 iterations. Training was performed using an Nvidia P100 GPU and the CRNN



training took from 0.5 to a few hours to complete, dependent on the energy bin and the number of iterations completed.

#### 4 Monte Carlo performance

Once training was complete the performance of the network was tested using an independent set of Monte-Carlo data representing the four H.E.S.S. phase one telescopes at 70% of their nominal optical efficiency and a zenith angle of 20°. In order to represent the typical data taking mode of H.E.S.S., gamma-ray events were simulated as a point source with an offset from the telescope pointing direction of  $0.5^{\circ}$ , while protons were simulated as a diffuse source with an opening angle of 2.5°, however only events reconstructed in the central 1° were included in performance evaluations.

The output of the neural network when evaluated on this dataset is a classification value between 0 and 1, roughly representing the probability that the event is a gamma ray  $(P_{\gamma})$ . However, as most gamma-ray events lie so close to 1 this classifier was reformulated to make the distribution more easily visible.

$$\zeta = -\log_{10}(1 - P_{\gamma}) \tag{2}$$

The resultant classifier distribution of  $\zeta$  is shown in Fig. 4 and is strongly peaked at 0 for the tested protons and lies between 0 and 10 for gamma-ray events. However, it is typically useful when cutting on this parameter to select events based on an energy-dependent (as the classifier distribution is typically strongly energy-dependent) gamma-ray efficiency.



Fig. 5 Left: Comparison of background rejection performance (as a ratio to mean scaled network performance) vs energy-dependent signal cut for the two recurrent networks. Right: Energy dependence of background rejection performance (as a ratio to mean scaled network

performance) for the two recurrent networks at 80% (solid line) and 60% signal efficiency. Events are re-weighted to represent an energy spectrum of  $E^{-2}$  for both signal and background

#### 4.1 Background rejection performance

Figure 5 (left) shows the performance of the recurrent networks in comparison with the traditional mean scaled parameter based network at different gamma-ray efficiency cut levels. The performance improvement of the recurrent networks is clear, with improvements seen at all levels of signal efficiency. The Hillas RNN shows around a 20% reduction in background in comparison to the mean scaled network, while the CRNN shows almost a 60% improvement in rejection power.

Figure 5 (right) shows the energy dependent comparison of background rate to mean scaled network at performance at 80% and 60% gamma-ray efficiency. In the lower energy bins (< 5 TeV) a clear improvement is seen in the performance of the recurrent networks over the mean scaled network. A  $\sim 20-25\%$  reduction is seen in the proton rate in the Hillas RNN at both 80% and 60% signal efficiency . Such an improvement at low energies could be expected due to the relatively large fluctuations in the Cherenkov light distribution in this energy range, potentially resulting in significantly different images being seen in the different telescopes. In this case taking the mean of the shower parameters will result in a loss of information and performance, whereas the recurrent network can use the full information from all telescopes.

The CRNN shows an even larger improvement in the lowest energy bins, showing a reduction in proton rate of more than 60% at both 80% and 60% signal efficiency. At low energies the convolutional layers are able to provide additional image information to the background rejection, most likely using shower information from pixels that were eliminated from the Hillas parameter construction by the image cleaning.

Above 5 TeV, the Hillas RNN performance matches closely the mean scaled network at both signal efficiencies, as the more well defined air showers in this energy range reduce the observed differences in the different telescopes. The larger images available in this energy range, however, provide significant information to the CRNN maintaining the 60% improvement in background rejection seen at lower energies.

# 4.2 Sensitivity to NSB level

The performance of the neural networks presented so far were evaluated at the nominal, per pixel, night-sky background (NSB) of 100 MHz used within H.E.S.S. simulations. However, this simulated value is a compromise between that observed level in extragalactic regions of as low as 50 MHz and that seen in the Galactic plane, which can reach to 300 MHz or above in some bright regions (e.g. the Carinae region). In order to test the robustness of the networks against



Fig. 6 Degradation in signal (gamma-ray) efficiency for the three tested networks as a function of NSB level

differing levels of noise we created gamma-ray simulations at 5 different NSB levels (100–300 MHz) and tested the fraction of events passing a background rejection cut defined using the 100 MHz simulations.

Figure 6 shows the acceptance of gamma-ray events at the different NSB levels, when defining the background rejection cut level at 80% gamma-ray acceptance based on the simulations at 100 MHz NSB. It is clear from these results that the Hillas-based networks are rather robust, falling in acceptance by only around 10% from 100 to 300 MHz. The robustness of the Hillas parameters can be understood from the two tail cut cleaning levels (typically 5 and 10 p.e.) applied to the image being significantly higher than the expected pixel to pixel fluctuations resulting from NSB noise (around 1 p.e. at 100 MHz in H.E.S.S.), resulting in the noise level in the included pixels being low. The CRNN however is strongly affected by the NSB, with a 50% reduction in gamma-ray event acceptance from 100 to 300 MHz. This reduction in acceptance is due to increased noise in the convolutional portion of the network, with the larger noise level introducing larger fluctuations in the image as well as potentially larger NSB fluctuations being mistaken for hadronic sub-showers in the uncleaned image section. Therefore increasing the noise level in the non-signal distribution clearly affects the CRNN classifier distribution, producing lower  $\zeta$  values for equivalent events.

This strong sensitivity to NSB is clearly a concern when evaluating the performance of the CRNN on data and it must be ensured that the results are compared with simulations of an appropriate NSB when extracting results.

## 5 Performance on H.E.S.S. data

Tests on an independent Monte Carlo have shown a significant increase in performance of both recurrent neural networks over the mean scaled network. However, these simulations are based on an idealised representation of the instrument behaviour. In reality camera images may contain a number of issues that affect the quality of the data, for example some camera pixel may be broken or the level of NSB may vary across the field of view. The network was therefore tested on H.E.S.S. phase one data. The outcome of this analysis was then used to test the stability of the results and check the performance in comparisons to the predictions of Monte Carlo simulations.

In order to ensure the different classifiers are compared in a fair way, energy dependent cut sets were created for the three classifiers that maintain a fixed efficiency of gamma rays passing cuts. In this case values of 80% and 60% gamma-ray efficiency were chosen as typical values for soft and hard cuts respectively. Table 1 shows the statistics for the cuts tested using the three different neural network configurations.

# 5.1 PKS 2155-304

The first source tested was the well known BL Lac object PKS 2155-304 observed throughout the operation of the H.E.S.S. instrument (e.g. [29,30]). A sample of around 15 h of observation with zenith angle of close to 20° and optical efficiency similar to that in the MC simulations was selected from the non-flaring periods of PKS 2155-304. This dataset contains a similar number of more than 1300 excess events and takes place over a relatively diverse set of observing conditions. Figure 7 (left) shows the resultant significance map (created using the *gammapy* software package [31]) of this region, clearly showing a strong source at a position consistent with the catalogue position of PKS 2155-304. Figure 7 (right) shows the distribution of significance from non-source pixels, which is well fit by a Gaussian with a mean of 0.14 and a width of 1.07, quite consistent with the expectation for

well normalised background in signal free regions (mean of 0, width of 1).

Table 1 shows the detection statistics for the dataset using the reflected background [32] to estimate the residual background contamination in the source region. For all network configurations tested a similar number of excess gammaray events are detected due to the cut being made on the expected gamma-ray efficiency. As expected the RNN-based networks show a reduction in the estimated level of background contamination. However as in this case as there is some small variation in the number of excess events it is fairer to make comparisons of the signal to background ratio (S/B) i.e. the number of excess events divided by estimated background contamination. In both the 60% and 80% cut set the Hillas-RNN shows an improvement in S/B of around 5-10%, while the CRNN shows an improvement of around 20% over the mean scaled network. This improvement in background rejection does not translate into large increases in source significance due to the extremely bright source being investigated.

Figure 8 (left) shows the distribution of the CRNN classification parameter obtained from this dataset in comparison with the results of MC simulations (at 100 MHz NSB rate) re-weighted to a spectral index of -3.4. In this case the data distribution provides an excellent match to the Monte Carlo expectation, demonstrating a stable behaviour of the classifier on strong, steep spectrum sources.

## 5.2 HESS J1745-290

The second case studied was the Galactic Centre point source HESS J1745-290 [33] commonly associated with the supermassive black hole Sagittarius A\*. In this case a selection of data was made from 2004 to 2008 datasets, resulting in a total of around 30 hours of observations. This field of view represents a rather different analysis to PKS 2155-304. Firstly the spectrum of this source is comparatively hard, with a spectral index of -2.1 and a cut-off at around 14 TeV. In addition to this the level of NSB in this region is significantly

Table 1Detection statistics for the two run lists tested with background cuts tuned to retain 80% and 60% of the gamma-ray events at energiesbetween 100 GeV and 100 TeV

	Network	80% Gamma-ray efficiency					60% Gamma-ray efficiency				
		Non	$\alpha N_{\rm OFF}$	Excess	S/B	σ	Non	$\alpha N_{\rm OFF}$	Excess	S/B	σ
	MSC NN	2602	560.4	2041.6	3.64	57.7	1841	288.7	1552.3	5.38	55.7
PKS 2155-304	Hillas RNN	2590	529.4	2060.6	3.89	59.1	1825	268.7	1556.3	5.79	56.7
(quiescent)	CRNN	2634	477.1	2156.9	4.52	62.8	1904	248.5	1655.5	6.66	60.4
	MSC NN	3071	1553.4	1517.6	0.98	31.5	2068	844.2	1223.8	1.45	32.8
HESS J1745-290	Hillas RNN	2813	1327.1	1485.9	1.12	32.9	1906	716.5	1189.5	1.66	33.9
	CRNN	2968	1320.2	1647.8	1.25	36.0	2030	693.0	1337.0	1.93	37.8



**Fig. 7** Left: Significance map of the PKS 2155-304 region created using an oversampling radius of 0.12° at 80% gamma-ray efficiency, the position of the source is marked with the dotted circle. Right: 1D



Fig. 8 Comparison of the CRNN classifier distribution with three energy thresholds obtained in the analysis of PKS 2155-304 (left) and HESS J1745-290 (right), data is compared with Monte Carlo simu-



distribution of significance from signal-free pixels (solid histogram), shown in comparison with the best-fit Gaussian



lations re-weighted to match the source spectrum index with 100 and 200 MHz NSB level respectively, and scaled to number of excess events in each energy bin

higher at around 200 MHz in comparison with the approximately 60 MHz in the PKS 2155-304 observations. For this dataset an improvement in S/B of around 15% is seen for the Hillas-RNN over the mean scaled NN and around 25% in the CRNN.

The CRNN classifier distribution shown in Fig. 8 (right) again shows an excellent match to the MC simulations (with

an NSB level of 200 MHz) re-weighted to the source spectrum. Again this demonstrated the stable behaviour of the network even with diverse observation conditions and higher NSB levels, although clearly care must be taken to choose the correct NSB level in the simulations.

# 6 Discussion

In this paper we have demonstrated the potential sensitivity gains available to imaging atmospheric Cherenkov telescopes by using the latest generation of machine-learning tools for background rejection and for the first time demonstrated a successful application of this scheme to data from the H.E.S.S. gamma-ray observatory. Applications of the convolutional-recurrent neural network to Monte Carlo air shower simulations and real data show an improvement in background rejection power of around 20–25% over the use of mean-scaled parameters typically used in previous background rejection implementations.

Although this does not match the even stronger performance gains predicted from simulated events (20% and 60% for the Hillas RNN and CRNN respectively) this mismatch could be caused by several factors. First the presence of cosmic-ray electrons which are present in the data (e.g. [34]) is not accounted for in the simulation predictions. These electron-induced air showers develop almost identically to gamma-ray induced air showers and are often considered to represent an irreducible background in IACT data which becomes more and more important as the hadron rejection power improves.

Secondly the network training is performed using simulated protons as the background events, however significant systematic uncertainties exist in the modelling of hadronic interactions in this energy range [35]. This behavioural uncertainty could result in a reduced performance when applying the trained networks to data due to incorrectly reproducing features within the air shower. However, due to the "black box" nature of the network behaviour it is difficult to identify any features that do not match between data and simulations. In principle this issue could be negated by using data from gamma-ray source free regions as the background training sample. However, to do this one must address the issues associated with the network efficiently learning the difference between simulated events and real data (discussed in detail in [18]).

This improvement is in line with the performance of that of goodness of fit cuts from image template based event reconstruction (e.g. [9]). The reproduction of the sensitivity of goodness of fit based cuts is to be expected in the case of gamma-rays where the air showers behave in relatively predictable way and the images seen in the individual telescopes are strongly correlated. It is important to take note of the sensitivity of the network performance to different observing conditions and that care must be taken to ensure that particularly the NSB level of the simulations matches that of the data to which it is being compared. This strong sensitivity to NSB level could potentially be lowered by careful preprocessing and denoising of the image, however it is possible that a runwise-simulation scheme (e.g. [36]) may be required to ensure

the lowest possible systematic uncertainties if such a scheme is deployed.

# 7 Conclusion

Although no significant performance gains are seen in background rejection power over the current state of the art goodness of fit based background rejection, use of this machine learning scheme does add some benefits. Firstly, the systematic uncertainties of this method, while likely as large, are different from the goodness of fit based approach. Thus allowing evaluation of the systematic uncertainties of analysis at the limits of the instrumental threshold. Secondly, the goodness of fit based approach relies on comparing shower images to a mean expected image template, limiting it is usefulness in the classification of particle species which produce large shower-to-shower fluctuations (such as protons or heavier nuclei). However, the training step of the RNNs naturally includes these fluctuations, meaning the RNNs may also be extremely useful in measuring the mass composition of hadrons in IACT data.

This work represents an early step in the use of such CRNNs for the rejection of hadronic background in IACTs. However, in order for such a system to be deployed for regular use in this generation of IACTs (or the next) significant further study is required. Particularly the network must be generalised for use with a much wider range of observing conditions. This work shows that the level of NSB of the observations is particularly important and some method of accounting for this must be found, but also one must consider the zenith angle, azimuth angle and also the optical efficiency of the telescopes observing to ensure a stable classifier performance.

Acknowledgements The authors would like to thank the H.E.S.S. Collaboration for allowing the use of H.E.S.S. data and simulations in this publication, as well as providing useful discussions and input to the paper.

**Data Availability Statement** This manuscript has no associated data or the data will not be deposited. [Authors' comment: The H.E.S.S. raw data and simulations used in this study are not public but belong to the H.E.S.S. collaboration. The use of the data in this study was kindly granted by the H.E.S.S. collaboration.]

**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 copy-

right holder. To view a copy of this licence, visit http://creativecomm ons.org/licenses/by/4.0/. Funded by SCOAP<sup>3</sup>.

# References

- 1. A.M. Hillas, Cerenkov light images of EAS produced by primary gamma. Int. Cosmic Ray Conf. **3** (1985)
- F. Aharonian, A.G. Akhperjanian, A.R. Bazer-Bachi et al., Observations of the Crab nebula with HESS. Astron. Astrophys. 457, 899 (2006). https://doi.org/10.1051/0004-6361:20065351
- T.C. Weekes, M.F. Cawley, D.J. Fegan et al., Observation of TeV gamma rays from the Crab nebula using the atmospheric Cerenkov imaging technique. ApJ 342, 379 (1989). https://doi.org/10.1086/ 167599
- M. Punch, C.W. Akerlof, M.F. Cawley et al., Detection of TeV photons from the active galaxy Markarian 421. Nature 358, 477 (1992). https://doi.org/10.1038/358477a0
- J. Albert, E. Aliu, H. Anderhub et al., Implementation of the Random Forest method for the Imaging Atmospheric Cherenkov Telescope MAGIC. Nucl. Instrum. Methods Phys. Res. A 588, 424 (2008). https://doi.org/10.1016/j.nima.2007.11.068
- S. Ohm, C. van Eldik, K. Egberts, /hadron separation in veryhigh-energy -ray astronomy using a multivariate analysis method. Astropart. Phys. **31**(5), 383 (2009). https://doi.org/10.1016/j. astropartphys.2009.04.001. ISSN 0927-6505
- Y. Becherini, A. Djannati-Ataï, V. Marandon et al., A new analysis strategy for detection of faint γ-ray sources with Imaging Atmospheric Cherenkov Telescopes. Astropart. Phys. 34, 858 (2011). https://doi.org/10.1016/j.astropartphys.2011.03.005
- 8. T. Murach, M. Gajdus, R.D. Parsons. A neural network-based monoscopic reconstruction Algorithm for H.E.S.S. II. (2015) (arXiv e-prints)
- M. de Naurois, L. Rolland, A high performance likelihood reconstruction of γ-rays for imaging atmospheric Cherenkov telescopes. Astropart. Phys. 32, 231 (2009). https://doi.org/10.1016/j.astropartphys.2009.09.001
- R.D. Parsons, J.A. Hinton, A Monte Carlo template based analysis for air-Cherenkov arrays. Astropart. Phys. 56, 26 (2014). https:// doi.org/10.1016/j.astropartphys.2014.03.002
- K. Fukushima, Neocognitron: a self-organizing neural network model for a mechanism of pattern recognition unaffected by shift in position. Biol. Cybern. 36, 193 (1980)
- A. Krizhevsky, I. Sutskever, G.E. Hinton. ImageNet classification with deep convolutional neural networks. In: F. Pereira, C.J.C. Burges, L. Bottou, et al. (eds.) Advances in Neural Information Processing Systems, vol. 25, pp. 1097–1105. Curran Associates, Inc. (2012)
- S. Hochreiter, J. Schmidhuber, Long short-term memory. Neural Comput. 9(8), 1735 (1997)
- A. Graves, A. Rahman Mohamed, G.E. Hinton, Speech recognition with deep recurrent neural networks (2013). arXiv:1303.5778
- D. Muthukrishna, G. Narayan, K.S. Mandel et al., RAPID: early classification of explosive transients using deep learning. Publ. Astron. Soc. Pac. 131(1005), 118002 (2019). https://doi.org/10. 1088/1538-3873/ab1609
- A. Möller, T. de Boissière, SuperNNova: an open-source framework for Bayesian, neural network-based supernova classification. Mon. Not. R. Astron. Soc. 491(3), 4277 (2019). https://doi.org/10. 1093/mnras/stz3312 (ISSN 0035-8711)
- R. Carrasco-Davis, G. Cabrera-Vives, F. Förster et al., Deep learning for image sequence classification of astronomical events. Publ. Astron. Soc. Pac. 131(1004), 108006 (2019). https://doi.org/10. 1088/1538-3873/aaef12

- I. Shilon, M. Kraus, M. Büchele, et al., Application of deep learning methods to analysis of imaging atmospheric Cherenkov telescopes data. Astropart. Phys. **105**, 44 (2019). https://doi.org/10.1016/j. astropartphys.2018.10.003
- S. Ohm, C. van Eldik, K. Egberts, γ/hadron separation in very-high-energy γ-ray astronomy using a multivariate analysis method. Astropart. Phys. **31**, 383 (2009). https://doi.org/10.1016/ j.astropartphys.2009.04.001
- A. Radovic, M. Williams, D. Rousseau et al., Machine learning at the energy and intensity frontiers of particle physics. Nature 560(7716), 41 (2018). https://doi.org/10.1038/ s41586-018-0361-2
- IceCube Collaboration, M.G. Aartsen, M. Ackermann, et al., The IceCube Neutrino Observatory—Contributions to ICRC 2017 Part V: Solar flares, Supernovae, Event reconstruction, Education, Outreach (2017) (arXiv e-prints)
- A. Aurisano, A. Radovic, D. Rocco et al., A convolutional neural network neutrino event classifier. JINST 11(09), P09001 (2016). https://doi.org/10.1088/1748-0221/11/09/P09001
- Identification of Jets Containing *b*-Hadrons with Recurrent Neural Networks at the ATLAS Experiment. Technical Report ATL-PHYS-PUB-2017-003. CERN, Geneva (2017)
- 24. F. Chollet, et al., Keras (2015). https://github.com/fchollet/keras
- M. Abadi, A. Agarwal, P. Barham, et al., TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems (2015) (Software available from tensorflow.org)
- N. Srivastava, G. Hinton, A. Krizhevsky, et al., Dropout: A simple way to prevent neural networks from overfitting. J. Mach. Learn. Res. 15(1), 1929 (2014) (ISSN 1532-4435)
- D. Heck, J. Knapp, J.N. Capdevielle, et al., CORSIKA: a Monte Carlo code to simulate extensive air showers (1998)
- K. Bernlöhr, Simulation of imaging atmospheric Cherenkov telescopes with CORSIKA and sim\_ telarray. Astropart. Phys. 30, 149 (2008). https://doi.org/10.1016/j.astropartphys.2008.07.009
- F. Aharonian, A.G. Akhperjanian, A.R. Bazer-Bachi et al., Multiwavelength observations of PKS 2155–304 with HESS. A&A 442, 895 (2005). https://doi.org/10.1051/0004-6361:20053353
- H.E.S.S. Collaboration, H. Abdalla, A. Abramowski, et al., Gamma-ray blazar spectra with H.E.S.S. II mono analysis: the case of PKS 2155-304 and PG 1553+113. A&A 600, A89 (2017). https://doi.org/10.1051/0004-6361/201629427
- A. Donath, C. Deil, M.P. Arribas, et al., Gammapy: An open-source Python package for gamma-ray astronomy. In: 34th International Cosmic Ray Conference (ICRC2015), vol. 34 of International Cosmic Ray Conference, p. 789 (2015)
- D. Berge, S. Funk, J. Hinton, Background modelling in very-highenergy γ-ray stronomy. A&A 466, 1219 (2007). https://doi.org/ 10.1051/0004-6361:20066674
- F. Aharonian, A.G. Akhperjanian, K.-M. Aye et al., Very high energy gamma rays from the direction of Sagittarius A\*. A&A 425, L13 (2004). https://doi.org/10.1051/0004-6361:200400055
- 34. F. Aharonian, A.G. Akhperjanian, U. Barres de Almeida et al., Energy spectrum of cosmic-ray electrons at TeV energies. Phys. Rev. Lett. **101**(26), 261104 (2008). https://doi.org/10.1103/ PhysRevLett.101.261104
- R.D. Parsons, H. Schoorlemmer, Systematic differences due to high energy hadronic interaction models in air shower simulations in the 100 GeV–100 TeV range. Phys. Rev. D 100(2), 023010 (2019). https://doi.org/10.1103/PhysRevD.100.023010
- M. Holler, J. Chevalier, J.-P. Lenain, et al., Run-wise simulations for imaging atmospheric cherenkov telescope arrays. In: Proceedings of the 35th International Cosmic Ray Conference. Busan (2017)