

Particle Identification in Smartphone Camera Images Using the Distributed Electronic Cosmic-ray Observatory

James Bourbeau^{1,2}, Felipe Campos^{2,3}, Blake Gallay^{2,4}, Matthew Meehan^{1,2}, Robert Morgan¹, Jeffrey Peacock⁵, Alex Pizzuto^{*1,2}, Cassidy Schneider^{1,2}, Ariel Levi Simons⁶, Justin Vandenbroucke^{1,2}, Miles Winter^{1,2}

¹ Department of Physics, University of Wisconsin, Madison, WI 53706, USA

² Wisconsin IceCube Particle Astrophysics Center, Madison, WI 53706, USA

³ University of California, Berkeley, Berkeley, CA, 94720, USA

⁴ Cornell University, Ithaca, NY, 14850, USA

⁵ Sensorcast, Boulder, CO, 80305, USA

⁶ University of Southern California, Los Angeles, CA, 90007, USA

E-mail: pizzuto@wisc.edu

The Distributed Electronic Cosmic-ray Observatory (DECO) is a global network of smartphones that searches images for evidence of cosmic rays and other ionizing charged particles. DECO was released to the public in 2014 and has citizen scientists from 80 countries on all seven continents participating in the project. We previously demonstrated that tracks seen in the DECO data set are caused by cosmic-ray muons by comparing the track length distribution of candidate events in the data set to the expected distribution from a cosmic-ray flux. However, robust particle identification necessary to separate cosmic rays from the radioactive background on an event-by-event basis had not previously been identified. We present a deep learning, computer vision algorithm for identifying and classifying charged particles in camera image sensors. The convolutional neural network was trained using images from the DECO data set and achieves classification performance comparable to human quality across four distinct DECO event classifications. We apply our model to the entire DECO data set and determine a selection that achieves a purity of 95% when applied to cosmic-ray muons and $\geq 90\%$ for all event types. The automated classification is run on the public DECO data set in real time in order to provide classified particle interaction images to users of the app and other interested members of the public. The model and techniques used to develop it are applicable to other smartphone-based cosmic-ray detectors and data sets consisting of images of charged particles from solid-state camera sensors.

*36th International Cosmic Ray Conference -ICRC2019-
July 24th - August 1st, 2019
Madison, WI, U.S.A.*

*Speaker.

1. Introduction

The Distributed Electronic Cosmic-ray Observatory (DECO) is a citizen science project that turns cell phones into cosmic-ray detectors. The project is made possible by remarkable parallels between modern smart-phone cameras and the technology found in astronomical telescopes and particle detectors. This has meant that cosmic rays and other forms of ionizing radiation can be readily observed from billions of devices around the world. Since September 2014, DECO has been publicly available as an Android application¹. During this time, data has been collected from 74 countries on all seven continents, as well as forty-seven US states, and all data taking locations are displayed in Figure 1. DECO's uniquely low barrier to entry has made it a useful medium through which to introduce students, and the public at large, to fundamental concepts in astrophysics, particle physics, computer science, and data analysis. The automated classification of DECO events is facilitated by a convolutional neural network. A public online database², allows users to examine the data they collect in near real-time. These data have been previously used to accurately measure the depletion thickness of cell phone image sensors, enabling the construction of simulations that confirm and refine our understanding of these particle interactions.

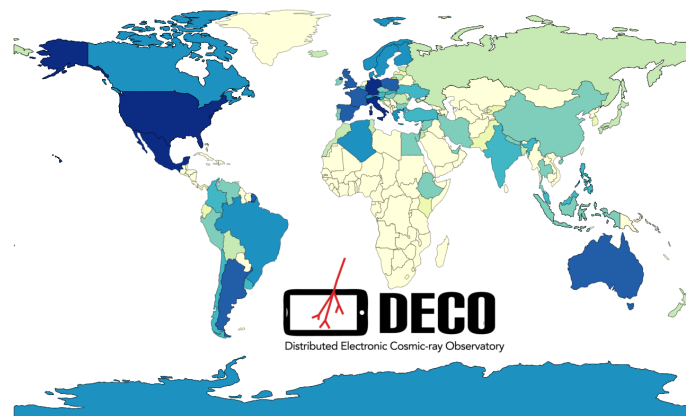


Figure 1: World map showing the global distribution of DECO users. Darker shades of green and blue indicate a higher number of images collected from the corresponding country, while countries colored beige have not yet contributed to DECO. The United States leads in data collection, with over 300,000 images. Map plotted using an equirectangular projection and is up to date as of July 2019.

2. DECO App

The DECO app seeks to detect ionizing radiation from secondary particles created by cosmic-ray primaries while cell-phone cameras are not being used to record optical photons. Therefore, the DECO mobile application is designed to be run with the camera lens face down or covered with opaque tape. While running, the app records camera images (~50 ms exposures) continuously and applies a filter algorithm to select particle interaction candidates. These data are synchronized to a browsable central server for analysis by scientists and members of the public. Associated metadata

¹App available at <https://wipac.wisc.edu/deco/app>

²<https://wipac.wisc.edu/deco/data>

for each event is also collected, including a unique event ID, phone model, Android version, time, geolocation, altitude, local magnetic field values, pressure, and temperature. Included in our processing of DECO data is the automatic event classification, and the creation of a JPEG image for each event. Users may query the database based on data, location, device ID, phone model, event category (standard or minimum bias) or event classification.

The events which pass all of the filters are then classified into morphological event classifications, in accordance with the convention in [1]. Tracks are long, straight clusters of pixels created by high-energy (GeV) ionizing particles, primarily muons at sea level, and cosmic rays above 20,000 ft. Worms are named for the curvy path of pixels, often composed of fragmented clusters, caused by the meandering path of electrons (produced by radioactive decays) undergoing multiple Coulomb scattering inside the sensor. Spots are small, approximately circular clusters of pixels that can be produced by lower energy electrons (e.g. produced by Compton scatters of gamma rays from radioactive decay) being quickly absorbed, or may also be indicative of alpha particles or cosmic rays normally incident to the sensor. Additionally, there are also events that result from artifacts in the image sensor: light exposure, hot spots, thermal noise fluctuations, and larger scale sensor artifacts such as rows of bright pixels.

Reliable classification of these events is non-trivial, as any DECO-specific classification algorithm must take into account the high levels of rotational and translational symmetry possessed by particle events, as well as an understanding of systematic inhomogeneities present in mobile hardware and software. In the section that follows, we describe the construction and optimization of a convolutional neural network that has increased the accuracy and efficiency of event classification.

3. Convolutional Neural Network

Convolutional neural networks (CNNs) [2] are a subclass of deep learning neural networks that have shown extraordinarily good performance learning features from datasets that are characterized by a grid-like topology, and so naturally lend themselves to the pixel images produced by DECO. The architecture begins with alternating layers to extract meaningful features that take into account the topological and morphological structure of the input data [3]. First, convolutional layers take a stack of inputs (e.g. color channels in an image) and convolve each with a set of learnable filters to produce a stack of output feature maps, where each feature map is a particular abstraction of the input image. These layers are followed by pooling layers, which reduce the dimensionality of a feature map by using an aggregation function to compute a summary statistic across a small, local region of the input. The features extracted from the pooling and convolutional layers are then used as input for a standard, fully connected, feedforward neural network, producing the desired output, here, a classification of our input image.

3.1 DECO CNN

In order to train our model, we require a set of images with human-determined categorical labels. While preliminary DECO CNN training worked off of a small data set of approximately 1000 images [4], further training would require a large number of training examples. This presented a challenge, as training data must be accommodated by a set of corresponding human-determined categorical labels that allow the error between the model's predictions and the ground truth to be

assessed. Assigning human labels to the entire DECO image database of ~45,000 events would be a very time consuming task. With this in mind, we opted for an iterative approach in which the number of labeled training examples was successively increased in parallel with the optimization of the model. That is, during the optimization process, images classified by the CNN were inspected by eye for incorrect classifications. These events were assigned a correct human label and added to the training data available to the next iteration of the model. The full details of the architecture, data augmentation, and training are presented in [5], but we summarize some notable results below.

3.2 Results

The trained model has exhibited excellent performance across all four categories (tracks, worms, spots, and noise) in the model. Overall, the model has obtained greater than 90 percent accuracy at correctly identifying the classification of events compared to human labeling. Figure 2 shows a category-by-category summary, known as a confusion matrix, quantifying the error between human and CNN classifications for each category. Each row has been normalized to the total count of human-labeled events, thus the matrix describes the conditional CNN probability distribution for each of the categories. The model was trained until improvements to the loss, a measurement of the model error, were negligible, and this is displayed in Figure 3. We find that this plateau occurred after 800 epochs (iterations through the data).

For the purpose of providing real-time classifications for the events listed in the public DECO database, we seek to maintain a high-purity set of events identified as tracks. After evaluating constant cut-off values of 0.7, 0.8 and 0.9 on the test set, we opted for a probability threshold of 0.8, which yields an event selection with a track efficiency³ of 80%, and, most importantly, a track purity of 90% [5].

4. Detector Characterization

The path shapes visible in DECO images are directly determined by the type of interactions the ionizing radiation undergoes within the detector. Although DECO has proven its capability to detect ionizing radiation and classify it into distinct morphological event classes, mapping between detected morphology and incident particle identity requires measurements made with known incident fluxes and comparison to simulation. DECO has recently detected events in the presence of understood radiative sources as well as through the use of GEANT4 based simulations.

4.1 Simulation

To simulate cell phone camera image sensors, we use Allpix-squared, a C++ open-source modular framework for the simulation of silicon pixel detectors [6] (code available at [7]). Allpix-squared interfaces GEANT4 [8], a simulation toolkit which handles the particle interactions in matter, and requires a version of ROOT [9]. Modules can be selected to choose the type of silicon pixel, the dimensions of the pixel array, depletion voltage and thickness, diffusion parameters, magnitude of electronics noise, and digitization. We simulate a pixel array similar to the one

³The proportion of events identified by humans as one of the four categories that were correctly classified by the CNN




Figure 2: The probability of the CNN correctly identifying each event type, along with the probability of mis-identifying each category, can be read directly off of this row-normalized confusion matrix, i.e. the model correctly identifies human-labeled tracks as tracks 92% of the time, while incorrectly identifying them as worms 9% of the time.

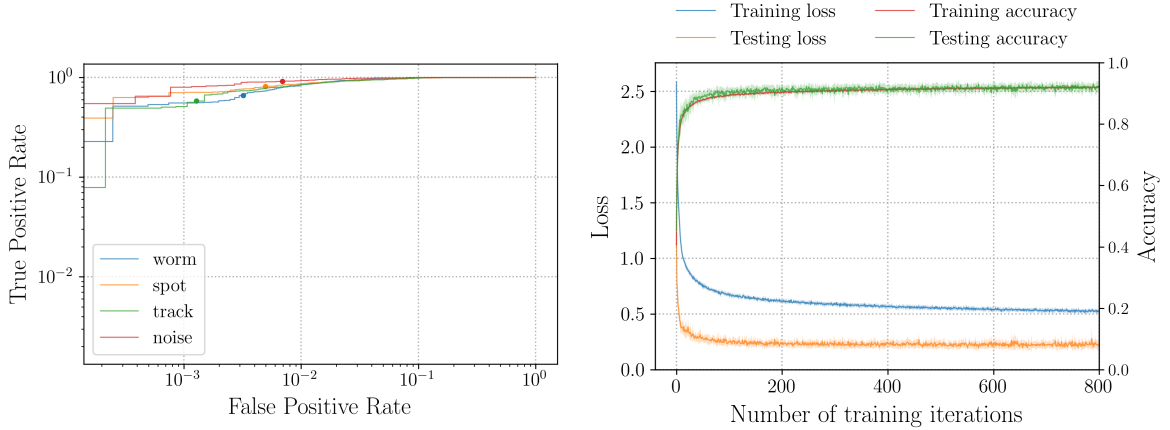


Figure 3: (Left) Receiver operating characteristic (ROC) curve displaying the true positive rate vs. false positive rate for a variety of threshold values. A threshold of 0.9 is indicated with a dot for each category. (Right) The training and testing loss as a function of training iteration (epoch).

measured in [10], and an image of the detector interacting with a beam of muons is displayed in Figure 4.

As we are mainly concerned with ionizing particles from cosmic-ray interactions, we restrict our analysis here to injected muons, though photon induced signatures could also be studied. We inject 1,000 muons at discrete energies every half decade from 10 keV to 10 GeV. Incident angles are every 15 degrees with respect to the pixel normal, and from every incident particle, we extract the number of pixels that registered a hit, the charge on each of these pixels, and the locations of the pixels hit. With this information and knowledge of the geometry of the detector, we can

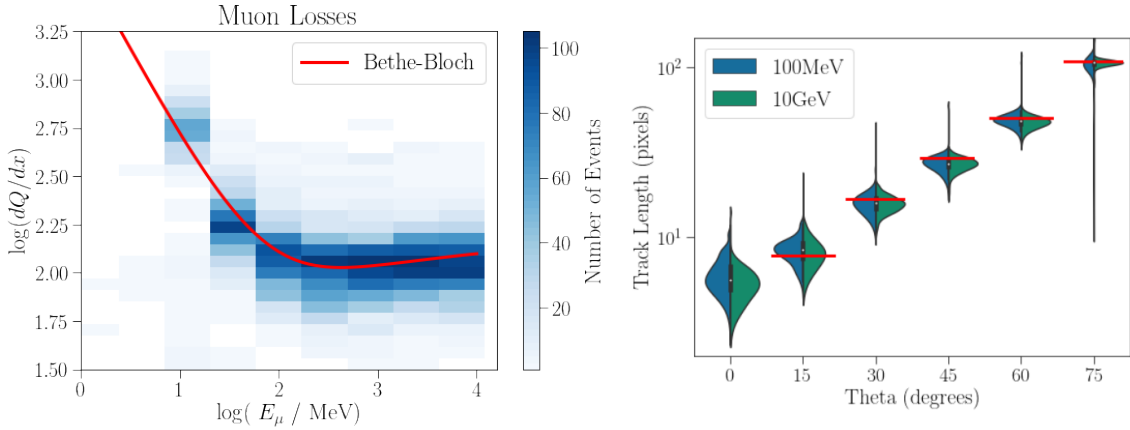


Figure 5: Observables from simulations of muons through the detector. (Left) Energy losses of incident muons compared to the expectation from Bethe-Bloch. Charge per length is in units of deposited electron charge per pixel width. The line is obtained by taking the Bethe-Bloch equation and transforming from dE/dx to dQ/dx by scaling by the energy required to liberate an electron charge in the sensitive region of the detector. (Right) violin plot of observed track length as a function of incident angle with respect to the normal of the pixel array for two different benchmark energies. Red shows the analytical approximation from the geometry of the detector. Observed track length calculation takes into account inter-pixel charge diffusion, assuming uniform diffusion around the muon’s path.

extract information about the energy deposited per unit length, and this is displayed in Figure 5.




Figure 4: Simulation of a beam of $4\text{GeV } \mu^+$ (blue) interacting with the camera image sensor. The lens and screen are modeled as GEANT4 plexiglass and the remaining parts of the phone are modeled with aluminum. Negatively charged particles that are created from interactions are displayed in red.

At energies between 0.1 - 10 GeV, muons are minimum ionizing particles (MIPs), losing approximately the same amount of energy per unit length.

To compare the charge deposited to energy lost per unit length, we assume the energy losses are dominated by ionization, where the loss is described by the Bethe-Bloch equation [11] We assume each ionization loss transfers exactly the amount of energy required to create an electron-pair hole in silicon, 3.62 eV.

In addition to energy losses, the track length of events in the detector is a crucial observable. With minimum ionizing muons, the track length holds the promise of providing a basis for per-event direction reconstruction.

Figure 5 displays the track length and deposited charge distributions for two benchmark energies covering the range of energies of cosmic-ray muons at sea level. For minimum ionizing muons, although the energy of the incident muon might vary by orders of magnitude, the total deposited charge does not vary much, as is expected.


4.2 Irradiation Measurements

To directly measure known sources of ionizing radiation, we exposed a Samsung Galaxy S2 to two different sources. For a gamma-ray source, $^{60}_{27}\text{Co}$ was used, which undergoes the decay



The two gamma rays are emitted from the transitioning of the doubly-excited Nickel to the ground state (this occurs with a branching ratio of 99.88%; the remaining decay is directly to a singly-excited state of Nickel). This process has a half life of 5.27 years, and the source we used had an activity of $1.0 \mu\text{Ci}$.

For a beta source, we used Strontium-90, which undergoes the decay



with a half life of 28.8 years. Here, the $0.546 \text{ MeV } \beta^-$ is the desired radiation, and our source had an initial activity of $0.1 \mu\text{Ci}$ and was approximately 5 years old.

The activity of the sources was measured using a Geiger counter at various distances, and was compared to the expectations from the quoted activities. Following source measurements, data collection took place over the course of 3 weeks, during which the DECO app was running on the phone placed at distances within 5 cm of the radioactive source.

By comparing the rate of detected events during exposures to radioactive sources to historic rates from the same device, we find a significant increase in the event rate, which is indicative of the device detecting events from the source. During sessions in which the device was not exposed to a local radioactive source, we find the rate to be approximately $5 \times 10^{-5} \text{ Hz}$, whereas during sessions with a source nearby, we find a rate of 2×10^{-4} over a livetime of roughly $1.1 \times 10^6 \text{ s}$.

Figure 6 displays detected events from both normal DECO events and irradiation measurements as well as simulated events. In the case of the irradiation measurements, the $\mathcal{O}(\text{MeV}) \beta^-$ emitter can be seen leaving curved signatures in the detector, consistent with the belief that low energy electrons can undergo multiple scatters in the sensitive region of the detector, and this is also evident in the 10 MeV simulated β^- . For simulated muons, nearly all signatures are elongated tracks, as is expected for MIPs.

5. Discussion and Conclusion

With the implementation of the online CNN, DECO has proven its capability to detect and classify ionizing radiation in realtime. Preliminary studies of simulated events in DECO show promise for using the morphological classifications of events to determine incident particle identification, as well as incident particle direction. These data are made available to the public, and provide a valuable dataset for cosmic-ray physicists and citizen scientists alike.

6. Acknowledgements

DECO is supported by the American Physical Society, the Knight Foundation, the Simon Strauss Foundation, QuarkNet, and by National Science Foundation Grant #1707945. We are grateful for DECO users across the world who contribute to the data set.




Figure 6: Comparison of DECO events from our database (top) to those which were recorded during irradiation measurements (middle) as well as to simulated events (bottom). All images have dimensions of 100×100 pixels. Colors indicates the pixel luminance for data events, defined as the RGB sum of the luminance of each subpixel. For simulation events, color indicates the deposited charge, and typically ranges from 0-200 electron depositions per pixel.

References

- [1] D. Groom, *Experimental Astronomy* **14** (2002) 45–55.
- [2] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, *Proceedings of the IEEE* **86** (1998) 2278–2323.
- [3] P. Y. Simard, D. Steinkraus, and J. C. Platt, *Best practices for convolutional neural networks applied to visual document analysis*, in *Proceedings of the Seventh International Conference on Document Analysis and Recognition - Volume 2, ICDAR '03*, (Washington, DC, USA), pp. 958–, IEEE Computer Society, 2003.
- [4] M. Meehan, S. Bravo, F. Campos, J. Peacock, T. Ruggles, C. Schneider, A. L. Simons, J. Vandenbroucke, and M. Winter, *POS (ICRC2017)* **375** (2018). [35,375(2017)].
- [5] M. Winter, J. Bourbeau, S. Bravo, F. Campos, M. Meehan, J. Peacock, T. Ruggles, C. Schneider, A. L. Simons, and J. Vandenbroucke, *Astropart. Phys.* **104** (2019) 42–53.
- [6] S. Spannagel, K. Wolters, D. Hynds, N. Alipour Tehrani, M. Benoit, D. Dannheim, N. Gauvin, A. NÄijrnberg, P. SchÄijtze, and M. Vicente Barreto Pinto, *Nucl. Instrum. Meth.* **A901** (2018) 164–172.
- [7] K. Wolters et al., “Allpix Squared: Generic Pixel Detector Simulation Framework.” <https://project-allpix-squared.web.cern.ch/project-allpix-squared/>.
- [8] **GEANT4** Collaboration, S. Agostinelli et al., *Nucl. Instrum. Meth.* **A506** (2003) 250–303.
- [9] R. Brun, *AIP Conf. Proc.* **583** (2001) 297.
- [10] J. Vandenbroucke et al., *JINST* **11** (2016) P04019.
- [11] S. Tavernier, *Experimental Techniques in Nuclear and Particle Physics*. Springer, 2010.