Exploring Deep Learning to Improve Compton Camera Based Prompt Gamma Image Reconstruction for Proton Radiotherapy

Gerson C. Kroiz[1], Carlos A. Barajas[1] Matthias K. Gobbert[1], Jerimy C. Polf[2]

[1] Department of Mathematics and Statistics,
University of Maryland, Baltimore County, Baltimore MD, USA
[2] Department of Radiation Oncology,
University of Maryland, School of Medicine Baltimore, MD, USA

Abstract. Proton beam radiotherapy is a cancer treatment method that uses proton beams to irradiate cancerous tissue while simultaneously sparing doses to healthy tissue. In order to optimize radiational doses to the tumor and ensure that healthy tissue is spared, many researchers have suggested verifying the treatment delivery through real-time imaging. One promising method of real-time imaging is through a Compton camera, which can image prompt gamma rays emitted along the beam's path through the patient. However, the images reconstructed with modern reconstruction algorithms are often noisy and unusable for verifying proton treatment delivery due to limitations with the camera. This paper demonstrates the ability of deep learning for removing false prompt gamma couplings and correcting the improperly ordered gamma interactions within the data for the case of Double events.

1 Introduction

Proton beams' primary advantage in cancer treatment as compared to other forms of radiation therapy, such as x-rays, is their finite range. The radiation delivered by the beam reaches its maximum, known as the Bragg peak, at the very end of the beam's range. Little to no radiation is delivered beyond this point. By exploiting the properties of the Bragg peak, it is possible to only irradiate cancerous tissues, avoiding any damage to the healthy surrounding tissues [1]. However, without some way to image proton beams in real time, limitations exist in our ability to take full advantage of the dose delivery properties of the proton Bragg peak. This is due to uncertainties in the beam's position in the body relative to important organs that should not be irradiated.

The Compton camera is one method for real-time imaging, which works by detecting prompt gamma rays emitted along the path of the beam. By analyzing how prompt gamma rays scatter through the camera, it is possible to reconstruct their origin. However, the raw data that the Compton camera outputs does not explicitly record the sequential order of the interaction data, which represents scatterings of a single prompt gamma ray. In addition, it often records false events, which mislabel scatterings of distinct prompt gamma rays as originating from a single ray. These problems make reconstructions based on Compton camera data noisy and unusable for practical purposes [1].

We approach these problems by leveraging deep learning techniques. We use neural networks, which, in general, represent data transformations. The network is trained by passing data through it, then updating it systematically so as to reduce the loss of its output compared with some desired output. Doing this properly can create a model that exploits subtleties in the data which traditional models are unable to use [2]. We show how this can be done in the following sections. Additional discussion about the impact of the approach on the application area can be found in [3].

The remainder of the paper is organized as follows. Section 2 discusses the biomedical application for this research. Section 3 describes a Compton camera and how we use it to collect data for our neural network. The network is then described in Section 4, where we discuss parts of the network's architecture. The model is tested on data in Section 5, where the resulting classifications are used for reconstruction. Section 6 makes general conclusions and proposes future directions for the research.

2 Proton Beam Therapy

To a first order approximation, the radiation dosage emitted by a proton beam is inversely proportional to the kinetic energy of the particles within the beam. Because the beam's particles lose kinetic energy as they traverse the patient, the amount of radiation delivered by the beam is low at its entry point, gradually rising until the beam nears the end of its range, at which point the delivered dosage rapidly reaches its maximum. This point of maximum dosage is called the Bragg peak. Little to no radiation is delivered beyond the Bragg peak. These characteristics of proton beam therapy give it a distinct advantage over x-rays. Exploiting its finite range, medical practitioners can confine the radiation of the beam to solely areas affected by cancerous tumors. Vital organs beyond the tumor can be spared [1].

While the characteristics of proton beam therapy explained above would in principle greatly reduce the negative effects of radiation therapy, there are still practical limitations. In current practice, the patient's body is imaged before undergoing treatment in order to map the position of the tumor. Because proton beam therapy consists of multiple sessions over a period of one to five weeks, the relative size and position of the tumor within the patient's body may change as surrounding tissues swell, shrink, and shift as a response to radiation. Therefore, whenever using proton beams, a safety margin must be added to the position of the Bragg peak in order to fully irradiate the tumor. This rules out certain beam trajectories that would otherwise minimize damage to healthy tissue [1].

Figure 1 compares two possible beam trajectories through a cross-section of the chest [1]. In this case, the heart, outlined in purple is positioned at the top-center of the figure and a tumor, outlined in green is located next to it.

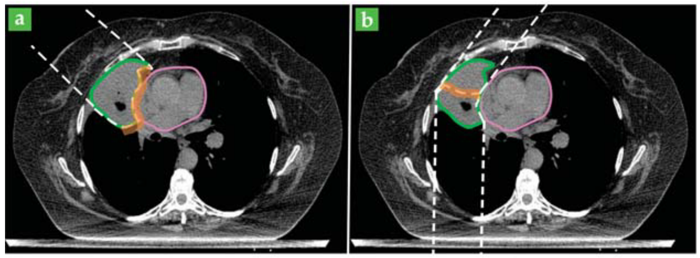

Fig. 1: (a) Optimal proton beam trajectory. (b) Suboptimal trajectory necessary to protect heart.

The optimal trajectory, shown in the left image, uses a single beam, which is represented as the space between the dashed white lines, to fully irradiate the tumor while stopping before reaching the heart. However, due to uncertainty in the exact location that the Bragg peak occurs (and the beam stops), a safety margin is added to the end of the optimal beam path to ensure the tumor always receives the prescribed dose even in the presence of day-to-day changes in patient setup and patient internal anatomy. This safety margin is represented in the figure as an orange strip at the end of the beam. This margin partially overlaps with the heart, which would mean a portion of the heart would receive the full treatment radiation dose. Since the heart is highly sensitive to radiation damage, it is very important that the radiation dose received by the heart is kept to a minimum. Therefore, in practice, the optimal trajectory is not used for treatment and instead the trajectory in the right image with two separate beams is used. Using two beams reduces the dose to the heart, but is considered suboptimal as these beams result in more radiation being delivered to a larger amount of healthy lung tissue. However, although suboptimal, it is considered preferable for treatment, since the beams avoid the heart [1].

If we were able to provide real-time information on the proton beam as it passes through the patient during live treatment, then we could ensure it is covering the tumor as intended, and safety margins (used to ensure tumor coverage) could be smaller and, thus, the more optimal treatments could be used. For instance, with proper real-time monitoring of the proton beam delivery, the optimal single beam treatment shown in Figure 1 (a) could be used without delivering high radiation doses to the heart while minimizing radiation to the lungs.

3 Compton Camera Imaging

3.1 Introduction to the Compton Camera

In order to exploit the full advantages of proton therapy, many researchers are investigating methods to image the beam in real time as it passes through the patient's body [1]. One proposed method for real-time imaging is by detecting prompt gamma rays that are emitted along the path of the beam using a Compton camera.

As the proton beam enters the body, protons in the beam interact with atoms in the body, emitting prompt gamma rays. These prompt gamma rays exit the body, some of which enter the Compton camera. Modules within the Compton camera record interactions with energy levels above some trigger-threshold. These modules have a non-zero time-resolution during which all interactions are recorded as occurring simultaneously. For each interaction (that is, Compton scatter) an (x, y, z) location and the energy deposited are recorded. The collection of all interaction data that a camera module collects during a single readout cycle is referred to as an event [4].

In principle, it is possible to use the data that the Compton camera outputs (paired with a suitable reconstruction algorithm) to image the proton beam, however, this has been shown to only be feasible at low energy levels. At the higher energy levels more typical of proton beam therapy, reconstructions of the beam are far too noisy to be helpful. This is a result of two main limitations in how the Compton camera records events.

At the higher energy levels typically used in treatment, proton beams emit a larger number of prompt gamma rays per unit time, increasing the likelihood of false events. Also, prompt gamma rays are more likely to scatter at higher energy levels, leading to more multi-scatter events which will be unordered. These two effects greatly diminish the accuracy of Compton camera reconstructions at high energy levels, making them unusable as explained further in [4].

There are several prompt gamma image reconstructions which can be used in conjunction with Compton camera data, such as the shifted histogram method seen in [5], but they produce bad results and thus cannot be effectively used with raw Compton camera data. The algorithms have a base assumption that the data has no false events and no misorderings. The order of interactions is directly connected to the path the prompt gamma took from the origin point. If the order of the interactions is shuffled, then the origin point of the prompt gamma will change and no longer be representative of the proton beam. These algorithms assume that the data they are working on is mostly perfect and all events within are suitable for use. If the Compton camera data could be denoised in an accurate, fast, and systematic way, then the Compton camera would be a viable method for prompt gamma image reconstruction.

3.2 The Representation of Events

Multi-scatter events can be classified into five categories: False Triples, False Doubles, Double to Triple, True Triples, and True Doubles. A False Triple event consists of three interactions which all originate from separate prompt gamma rays that happened to enter the same module of the camera at the same time. These should be removed from the data before reconstruction. Similarly, False Double events contain two interactions originating from separate prompt gamma rays. These too should be removed. A Double to Triple event contains two interactions corresponding to the same prompt gamma ray, and one interaction from a different prompt gamma ray. The non-corresponding interaction should be removed before reconstruction. The two remaining categories of events are True Double and True Triple events, which, once properly ordered, can be used for reconstruction. Detecting the correct ordering for the interactions of true events is, by itself, a non-trivial task. The studies in this paper focus solely on True and False Double events. However, the general neural network architecture used for the Double events in this paper is also used on the other multi-scatter event categories [6].

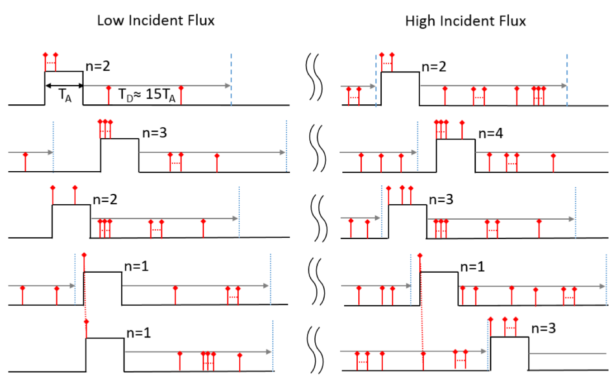

Fig. 2: An illustration of events.

Figure 2 shows a schematic of the Compton camera as it records events. The left side shows events produced at low energy levels and the right shows higher energy levels. Each row represents an independent module of the camera. The red arrows represent scatters, with those originating from the same prompt gamma ray being connected by a dotted line. A raised pulse represents a single readout cycle within a module of length T_A. The value n is how many interactions occur during the readout cycle. Looking at just the left side, the first two rows show a True Double and True Triple event, respectively. The third row shows a False Double event consisting of two scatters originating from different prompt gamma rays. The fourth and fifth rows show two True Single events that consist of separate scatters by the same prompt gamma ray. The right side representing higher energy levels shows a far greater proportion of false events.

The raw data output but the Compton camera for each interaction is of the form (e_i, x_i, y_i, z_i), where for $i = 1, 2, 3$, e_i is the energy level, and x_i, y_i, z_i are the x-, y-, and z-coordinates respectively. It is important to note that for any event i could be 1, 2, or 3 as each event can have up to 3 interactions included. For events with less than 3 interactions, for the values of i that are not included in the event, the corresponding (e_i, x_i, y_i, z_i) would be left empty or as NaN.

To improve the performance of our networks, we find it useful to append two extra values to the raw Compton camera data. In this version, we add the distances $\delta r_{i,j}$ between the ith and jth interactions, where $i, j = 1, 2, 3$. Since these values have physical significance with regards to the ordering of interactions, explicitly including them in the data makes it easier for the networks to learn.

4 Deep Learning

There is often the question whether or not deep learning is needed. If simpler, faster, easier to use methods exist, then one should just opt for those methods instead. In the context of our problem, there is currently no adequate physical model which will quickly and accurately determine the ordering of the Compton camera interactions. Since there is no physical model which we can use, then the next question is whether simpler and more accurate machine learning models can be used. In [6] we did a hyperparameter study using random forests with a maximum depth of 30 and 1400 estimators. Random forests near maximum depth yielded little benefit over faster to train, shallower, random forests with a depth of 6 or 7. The accuracy of all of our trained forests did not surpass 35% for a true event validation set. Shallow fully connected neural networks were used in limited capacity in [7] to classify a single interaction for a true triple. The work showed promising results but expanding beyond one interaction resulted in many issues in regards to output interpretation, loss calculations, and underfitting. When we incorporated false data and double to triples, the accuracy of prediction fell sharply for true events. This sharp drop in true event ordering detection coupled with heavy underfitting, when using the shallower networks, pushed us toward shallow but more complex fully connected neural networks. Additional studies with the shallow, complex, many neurons per layer, fully connected networks continued to produce poor validation scores for true scattering data which pushed us to deeper fully connected networks which maintained the same complexities, discussed in Section 4.1, but with fewer neurons per layer. Only in using deeper complex fully connected neural networks did we see the rampant underfitting eliminated leading to better accuracy for triple ordering detection noted in [6].

In this work, we train a neural network to eliminate false events while also correctly ordering true double events present in the Compton camera data. This cleaner dataset would allow prompt gamma image reconstruction algorithms to have more accurate prompt gamma travel paths with which to do their reconstructing. This more accurate image reconstruction provides the ability to better visualize the path of the proton beam within the patient, with the ultimate goal to make real-time adjustments to patient treatment plans.

The network contains three main components: an input layer which accepts the data, hidden layers which each perform some transformation on the data, and an output layer which returns the transformed data in some prescribed format [2]. We would like to train the neural network to transform the provided data in some useful way. In the case of the data output by the Compton camera, we would like the neural network to transform each multi-scatter event so that it contains only interactions originating from the same prompt gamma ray, and so that these interactions are in the correct order.

To improve the network's performance, it is typical to train the network on all available data multiple times. One pass through all the *training data*, or data isolated specifically for training the model, is referred to as an epoch. Often, the network will be trained for hundreds or thousands of epochs. It is standard practice to set aside some data with which to evaluate the network after each epoch. These data are called the *validation data*. By evaluating the network at the end of each epoch, it is possible to plot how the network's performance improves over the training process, giving insight into whether or not the network has been fully trained. After the network has finished training, a final data set separate from the training data and validation data is used to test the network. This data set is referred to as the *test data*.

4.1 Fully Connected Residual Blocks

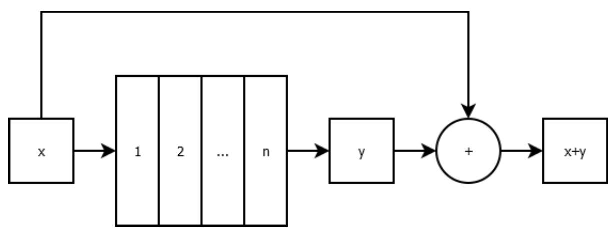

Fig. 3: A fully connected residual block takes an input and passes it through n layers eventually adding it to the output of the n layers.

The network used in Section 5 is a deep fully connected neural network. Neural networks, especially fully connected ones, break down once they start becoming notably deep and complex. One of the first problems is that the values start to become very small during the forward propagation process. This leads to zeros and like-zero values becoming more prominent as you go deeper and deeper. A fix to this forward propagation issue is to opt to use Leaky ReLU over the traditional ReLU. However, the second problem occurs during back propagation. During back propagation we start to see the gradient becoming like-zero causing little to no update to existing weights which causes learning stagnation. This phenomenon is discussed more intimately in [8] where they detail these effects. The major breakthrough to this problem inspired by ideas in [8] where they create ResNet, a network built from "residual blocks". A visual representation of a residual block can be seen in Figure 3. Consider some record x. We pass it as an input to a small group of n layers with their own activators. The result of the layer digestion we can call y. Finally, we concatenate x and y.

The concatenation in our case, and the case of the original ResNet, is addition. This addition operation helps push non-zero values through the forward propagation process which helps keep input data to each block fresh and non-zero. This also helps prevent vanishing gradients during the back propagation process. Residual blocks were originally used with convolutional layers whereas we are using fully connected layers. We also create a fully connected network which has considerably more layers than neurons per layers which yields a a super deep, yet thin, neural network.

4.2 Network Design

We used Tensorflow 2.5 with the bundled Keras module for our neural network backbone. The network design is built on the fully connected residual blocks described in Section 4.1. The architecture starts with an input layer with 10 neurons. The hidden layers are comprised of 128 fully connected residual blocks. Each block is made up of 2 layers, each consisting of 128 neurons with Leaky ReLU serving as the inter-layer activation function for the hidden layers. This gives us 256 hidden layers in total. Lastly, the output layer to the model consists of 3 neurons with a softmax activation function. The network training process used Keras' provided Adam optimizer with binary crossentropy loss. The training was done using Keras' `Model.fit` method on two NVIDIA K20 GPUs.

We conduct hyperparameter studies in [7] and [6] to optimize the architecture of the network and hyperparameters for training the model. Section 5 is based on these neural networks.

5 Results

After training the neural network, we want to test how accurate the model is with classifying the different forms of Doubles events and how these classifications can influence the reconstruction process of the proton beam.

5.1 Data and Hardware Configuration

The training and testing of the neural network alongside the reconstruction process use the distributed-memory cluster taki in the UMBC High Performance Computing Facility (`hpcf.umbc.edu`). In particular, the networks were trained and tested on a hybrid CPU/GPU node with two NVIDIA K20 GPUs (2496 computational cores over 13 SMs, 4 GB onboard memory), two 8-core Intel E5-2650v2 Ivy Bridge CPUs (2.6 GHz clock speed, 20 MB L3 cache, 4 memory channels), and 64 GB of memory.

For our studies, we use 150MeV (Mega electron Volt) beams with three different dosage rates: 0kMU (kilo Monitor Unit), 100kMU, and 180kMU. The larger kMU values correspond to more intense dosage rates. For the dosage rates 0kMU, 100kMU, and 180kMU, our test data set consisted of 165,000, 50,000, and 28,000 samples, respectively.

5.2 Neural Network Results

0kMU				100kMU				180kMU			
	12	21	44		12	21	44		12	21	44
12	84.3	8.8	6.9	12	84.2	9.0	6.8	12	84.7	8.3	7.0
21	8.8	84.7	6.5	21	9.2	84.2	6.6	21	8.3	84.8	6.8
44	5.9	5.9	88.2	44	5.6	5.9	88.5	44	5.4	5.9	88.7

Fig. 4: Confusion matrices for three different dosage rates. The fully connected network was trained on 150MeV beams with True and False Doubles data using a step-scheduler over 5000 epochs.

The results for how the neural network performed on the Doubles data is shown in Figure 4. Figure 4 displays three confusion matrices, one for each of the different dosage rates. The leftmost column is the correct input class and the number in each following column represents the percentage of data put into the class at the top of the column. Each cell within the matrix is shaded proportionally to the maximum matrix value.

The shading is based on the accuracy of the cell, where higher accuracies result in darker shades of green. The darkest entry in each row is the dominant classification of the input class. Within the matrices, there are three different labels: 12, 21, and 44. Note that since the data only includes Doubles, each label has two digits, one for each interaction. Label 12 represents the case where both interactions are correctly ordered and no adjustments need to be made. Label 21 represents the scenario where the second interaction should be first and the first interaction should be second. The 44 label is for False events indicating that the events should be thrown away. This is because the interactions originate from separate prompt gamma rays, and as such we cannot derive any meaningful conclusions from the two single interactions with regards to either gamma origin.

For each of the confusion matrices, the dominant classification for each input class is the class itself, where the accuracies range from 84.2% to 88.7%. This indicates that the model was trained successfully and is able to correctly classify most samples. In other words, given an unlabeled set of Doubles data, the model would be able to classify around 85% of the samples correctly. The matrix values between the dosage rates do not appear to change by more than 1%, where higher dosage rates have slightly higher accuracies. We see that the neural network classifies false events with around a 4% higher accuracy for all three dosage rates compared to the two true Double events 12 and 21.

5.3 Reconstruction Results

Based on the classifications of the data from the neural network, we created a new dataset of cleaned data, where events classified as False events were removed. In other words, the cleaned dataset consists of the classifications of the samples

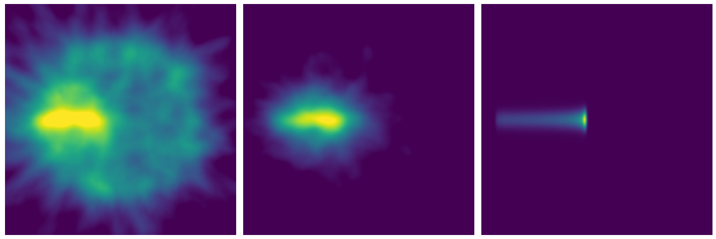

(a) Uncleaned (b) Cleaned (c) Dose

Fig. 5: Comparison of gamma ray images reconstructed with (a) uncleaned data and (b) cleaned data to the (c) dose delivered by the proton beam.

from the neural network. We reconstructed the proton beam using three version of the data, as shown in Figure 5. Figure 5 (a) shows the reconstruction image from the uncleaned data. Figure 5 (b) shows the reconstruction image from the cleaned data. Ultimately, we want the reconstruction images to resemble the Bragg Peak as shown in Figure 5 (c). When comparing the Figure 5 (a) and (b), we see that the cleaned data has removed a significant portion of the noise seen in the uncleaned reconstruction. However, relative to Figure 5 (c), there still is visual noise in the reconstruction using the cleaned data.

This is due to a combination of things. First the accuracy of the network is between 80% and 90% which means that between 10% and 20% of the reconstruction data consists of bad events. These incorrectly classified events will result in incorrect reconstruction for those prompt gamma rays. Additionally, the nature of reconstructing a proton beam based on scattering data is inherently difficult and not without its own challenges. The nature of reconstruction itself is beyond the scope of this work and can be seen in [1].

6 Conclusions

The results of the neural network when trained and tested on the Doubles data were quite promising, with testing accuracies in the mid-80s. This tells us that the neural network is capable of learning the structure and patterns that exist within the data at a reasonable level. When put to use, we can visually see changes in image quality when comparing the quality of the uncleaned versus cleaned reconstruction. The unclean data shows large amounts of noise and a rough shape whereas the clean data gives a thinner and more accurate profile when compared to the original dose. This demonstrates that we have significantly improved reconstruction quality by reducing the noise present in the data using our neural network.

There are several directions to take this research. First and foremost, the results in this paper are decent but can be improved. In-depth hyperparameter studies can be conducted to further improve network accuracy, which directly can reduce noise in the image reconstructions. Additionally, the work here only looked into fully connected neural networks. However, there are many deep learning models which may produce higher accuracies and better image reconstructions. We can explore other models such as convolutional or recurrent neural networks. Another approach for future research is to look into using deep learning techniques for the reconstruction algorithm. Using similar deep learning techniques like the ones shown in this work, we can study regression models that can accurately predict initial energy values for individual gamma rays. Improving the estimated initial energy values with the regression model will result in more accurate image reconstruction, which in practice can reduce the amount of uncertainty of the region influenced via the proton beam.

Acknowledgment

This work is supported by the National Institutes of Health National Cancer Institute under award number R01CA187416. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health. The hardware used in the computational studies is part of the UMBC High Performance Computing Facility (HPCF). The facility is supported by the U.S. National Science Foundation through the MRI program (grant nos. CNS–0821258, CNS–1228778, and OAC–1726023) and the SCREMS program (grant no. DMS–0821311), with additional substantial support from the University of Maryland, Baltimore County (UMBC). See https://hpcf.umbc.edu for more information on HPCF and the projects using its resources.

References

1. Polf, J.C., Parodi, K.: Imaging particle beams for cancer treatment. Phys. Today **68**(10), 28–33 (2015)
2. Chollet, F.: Deep Learning with Python. Manning Publications Co. (2018)
3. Polf, J.C., Barajas, C.A., Kroiz, G.C., Peterson, S.W., Maggi, P., Mackin, D.S., Beddar, S., Gobbert, M.K.: A study of the clinical viability of a prototype Compton camera for prompt gamma imaging based proton beam range verification. In: AAPM Virtual 63rd Annual Meeting (submitted (2021))
4. Maggi, P., Peterson, S., Panthi, R., Mackin, D., Yang, H., He, Z., Beddar, S., Polf, J.: Computational model for detector timing effects in Compton-camera based prompt-gamma imaging for proton radiotherapy. Phys. Med. Biol. **65**(12), 125,004 (2020). DOI 10.1088/1361-6560/ab8bf0
5. Mackin, D., Peterson, S., Beddar, S., Polf, J.: Evaluation of a stochastic reconstruction algorithm for use in Compton camera imaging and beam range verification from secondary gamma emission during proton therapy. Phys. Med. Biol. **57**(11), 3537–3553 (2012)
6. Barajas, C.A., Kroiz, G.C., Gobbert, M.K., Polf, J.C.: Deep learning based classification methods of Compton camera based prompt gamma imaging for proton radiotherapy. Tech. Rep. HPCF–2021–1, UMBC High Performance Computing Facility, University of Maryland, Baltimore County (2021). URL http://hpcf.umbc.edu

7. Basalyga, J.N., Barajas, C.A., Kroiz, G.C., Gobbert, M.K., Maggi, P., Polf, J.: Improvements to the deep learning classification of Compton camera based prompt gamma imaging for proton radiotherapy. Tech. Rep. HPCF–2020–29, UMBC High Performance Computing Facility, University of Maryland, Baltimore County (2020). URL http://hpcf.umbc.edu
8. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 770–778 (2016). DOI 10.1109/CVPR.2016.90