

# Deep Learning for High Throughput Decision Making on Diamond Content of PET Activated Kimberlite Rocks

Nicholas Connell<sup>1</sup>, Timothy Brooks<sup>2</sup>, Tiego Mpai<sup>2</sup>, Thendo Nemakhavhani<sup>2</sup>,  
Simon Connell<sup>2</sup> and Charis Harley<sup>1</sup>

<sup>1</sup>Department of Electrical and Electronic Engineering Science, Faculty of Engineering and the Built Environment, University of Johannesburg, Johannesburg, South Africa

<sup>2</sup>Department of Mechanical Engineering Science, Faculty of Engineering and the Built Environment, University of Johannesburg, Johannesburg, South Africa

E-mail: nickpconnell@gmail.com

**Abstract.** The application of computer-vision and deep-learning for the sorting of crushed rock in diamond mines is investigated. To this end vast quantities of positron emission tomographic images are synthesised to train convolutional neural networks as classifiers.

## 1 Introduction

The goal of the MinPET project is to develop, test, and validate a machine that can sort crushed rocks from a diamond mine. The machine is expected to be capable of sorting at least 500 tonnes of rock per hour. The sheer volume of rock that needs to be processed necessitates the use of autonomous AI agents. This rest of this paper describes the combination of synthetic data generation and deep-learning needed to create such agents and highlights the achievements of the MinPET research team [1–5]. Figure 1 depicts a hypothetical run-of-mine set up for the MinPET system.

The scanning technique employed is positron emission tomography (PET). The rock is activated with 40 MeV gamma ray beam to induce the giant dipole resonance and transmute certain atoms into PET isotopes, namely carbon to carbon-11 so diamond content can be ascertained. This paper aims to show that the viability of training a convolutional neural network classify PET images of kimberlite rocks as containing diamond or barren.

Traditional diamond mining sorting methods employ staged crushing, between which any revealed diamonds are retrieved. This requires vast amounts of energy and water, which is expensive and wasteful because deposits contain on average only a few carats of diamond per 10 000 kg of rocks. Furthermore, it forces compromises to be made between profitability and acceptable diamond breakage; in mines known to produce diamonds larger than 100 carats these compromises often lead to the diamonds being broken into smaller less profitable pieces [6]. Mines are often situated in energy insecure and water scarce countries where they negatively impact the local environment and put great stress on a host country’s resources.

MinPET offers three key benefits over traditional rock processing techniques; reduced diamond breakage, lowered costs to extract diamonds, and lowered environmental impact by identifying and isolating diamond bearing rocks at an earlier stage in the crushing process. Immediate benefits include higher profits, lower energy costs, and lowered chance of damaging diamonds. The quantity of material requiring further processing will be greatly lowered because the majority of rocks are barren and will be discarded. The will make viable methods that are more environmentally friendly and are better at preserving diamond integrity [6]. Additionally, the MinPET process will help audit the quantity and size of diamonds contained in the rocks as the PET image contains all of this information.

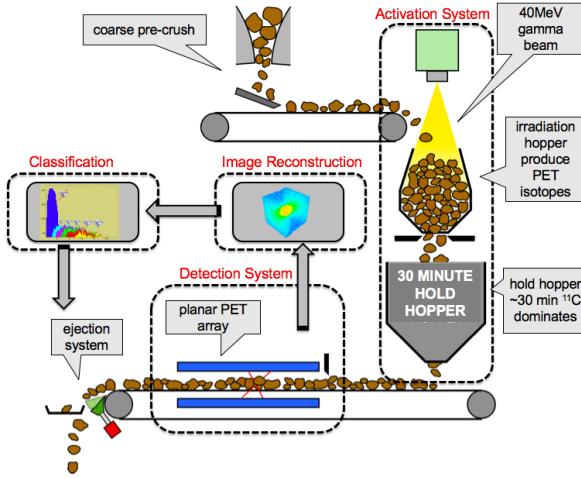


Figure 1: Proposed run-of-mine MinPET system showing all critical functions of the system: activation, holding, detection, reconstruction, classification, and ejection

## 2 Methodology

The study consisted of four stages. First, the generation of synthetic data by Monte Carlo simulations using the Geant4 [7] software. Second, the reconstruction of 3D images of the distribution of PET-isotopes using the Astra Toolbox Python library [8] and maximum likelihood expectation maximisation (MLEM). Third, training data was extracted from the images using custom Python code. Fourth, the AI was trained using an adapted version of a 3D ConvNet from the second place solution [9] to the 2017 Kaggle competition written in Keras a Python library for machine learning.

### 2.1 Simulation

In total 3000 simulations were conducted to produce the training data and 3010 simulation were conducted to produce the test data. Each simulation contained one spherical kimberlite rock which housed one spherical diamond. The rock always had a diameter of 100 mm. The composition by mass fraction of the compounds of kimberlite used in the simulations is detailed in Table 2.1, which were determined for southern Africa [10,11]. For each training data simulation the diameter of the diamond in mm was sampled uniformly from the range (0,15]. The testing data was split into 10 groups each with different constant diamond diameter: 1 mm, ..., 10 mm. The diamond's location in every simulation was randomised by sampling each coordinate uniformly and resampled if the position placed any amount of the diamond outside of the rock.

Compound	Si <sub>4</sub> O <sub>2</sub>	MgO	CaO	CO <sub>2</sub>	H <sub>2</sub> O	Fe <sub>2</sub> O <sub>3</sub>	FeO
Mass Fraction (%)	42.5562	28.6058	5.3233	5	5	4.6938	3.3138
Compound	Al <sub>2</sub> O <sub>3</sub>	TiO <sub>2</sub>	K <sub>2</sub> O	P <sub>2</sub> O <sub>5</sub>	MnO	Na <sub>2</sub> O	
Mass Fraction (%)	2.5158	1.7979	0.6799	0.1914	0.1889	0.1309	

Table 1: Kimberlite Composition by Mass Fraction.

The activity of the kimberlite and the diamond was determined by assuming a uniform electron flux of  $2.8086 \times 10^{13} \text{ cm}^{-2}$  incident on a copper bremsstrahlung target, and incorporating a delay of 23 minutes to simulate the time spent in the hold hopper after irradiation and before detection. Table 2.1 summarises the possible pathways to produce PET isotopes, their half life in seconds, and the effective reaction cross-sections for a photon beam produced by an electron beam of 40 MeV.

The rock and the diamond are simulated passing through the detectors at  $1 \text{ m s}^{-1}$ . There were two detector planes situated 150 mm above and 150 mm below the path of the rock. The detectors were rotated by  $1^\circ$  around the horizontal and vertical axes to take advantage of digital oversampling. The pixels of the detector were Bismuth Germanate ( $\text{Bi}_4\text{Ge}_3\text{O}_{12}$ ). Pixels are grouped into photo-multiplier tubes (PMT), and PMTs were grouped into MinPET Units (MPU). Table 2.1 summarises the geometry of the detectors' components by their side length, the

Reaction	$^{12}\text{C} \rightarrow ^{11}\text{C}$	$^{16}\text{O} \rightarrow ^{11}\text{C}$	$^{16}\text{O} \rightarrow ^{15}\text{O}$	$^{40}\text{Ca} \rightarrow ^{38}\text{K}$	$^{54}\text{Fe} \rightarrow ^{53}\text{Fe}$	$^{46}\text{Ti} \rightarrow ^{45}\text{Ti}$
Cross-Sec. (mb)	0.18791177	0.001010126	0.18038939	0.1121721	1.4198512	1.60765
Half life (s)	1220.04	1220.04	122.24	458.40	510.60	11088.00

Table 2: PET-isotopes and their parent atoms with reaction cross-section at 40 MeV and isotope half life.

gap between consecutive elements, and the grid they form. Each detector array was approximately 1.6 m long in the rock's direction of travel and 1.2 m wide.

Component	Length (mm)	Gap (mm)	Grid (XxY)
Pixel	2.84	0.2	16x16
PMT	48.44	2.0	4x4
MPU	199.76	5.0	8x6

Table 3: Detector component definitions.

The output of simulations was a collection of coincident hit pairs. For two gamma ray hits to be considered a coincident they needed to come from different detectors, the difference in their hit times needed to be 2 ns or less, and there needed to be no other gamma rays within 2 ns of either of their hit times. This meant the final output file consisted of nine data points, two energy deposits by each hit, two sets of 3D coordinates, and one hit time.

## 2.2 Image Reconstruction

Each coincident hits pair can be used to calculate a line of response that connects the two hit locations. It can be assumed with some confidence that a positron will have annihilated somewhere along this line. Given that most positrons annihilate with an electron close to their parent atom one can begin to reconstruct the most likely distribution of PET isotopes in the rock given many lines of response.

First, a global translation was applied to all hit locations such that the reference changed from the detector frame to the moving rock frame. Equation 1 show how this was done, relying on three quantities: the velocity of the rock/belt  $v_{belt}$ , the hit time  $t_{hit}$ , and the time at which the rock was centred between the detectors  $t_{centre}$ .

$$x_{new} = x_{old} + v_{belt}(t_{centre} - t_{hit}) \quad (1)$$

Second, the line of response for each coincident pair was determined. The intersection of this line with the mid plane was computed, where the mid plane is the horizontal plane mid-way between the detectors. The angle that the line made with the mid plane was also computed in spherical coordinates. This resulted in four data points for each pair: the polar (measured from the vertical) and azimuthal angles, and the x and y intersection coordinates. Any lines of response with a polar angle greater than  $\pi/4$  radians were discarded.

Third, the four dimensional data computed from the coincident pairs were binned to get a 4D rock sinogram. The bin size in the x and y axes was 2 mm. The bin size for the polar angle was  $\pi/90$  radians. For the azimuthal bins, the bin size depended on which polar bin it is associated with. The number of azimuthal bin at given polar bin  $n_\theta$  is governed by equation 2.

$$N_\phi = \begin{cases} 1 & , \text{ if } n_\theta = 0, \\ \text{round} \left( 180 \cdot \sin \left( \frac{n_\theta}{45} \cdot \frac{\pi}{2} \right) \right) & , \text{ otherwise.} \end{cases} \quad (2)$$

Finally, the 4D sinogram was given to an MLEM algorithm that computed a density map of PET isotopes. The code for the MLEM algorithm used the Astra Toolbox. The output was a 150x150x150 array where a given value represented the estimated density of PET-isotopes for the voxel. Each voxel had a side length of 2 mm. Figure 2 contains an example of a reconstructed PET image of a 10 mm diameter diamond in a rock.

## 2.3 Convolutional Neural Network Architecture

The convolutional neural network (CNN) architecture presented in Table 2.3 was adapted from a 2017 solution to a Kaggle competition on classifying lung tissue and predicting nodule malignancy [9]. The domains of diamond classification and lung tissue classification overlap nicely, given that both have a lot of background noise and the target can be of any size or shape. Furthermore, CNNs are a foundational aspect of deep learning for computer vision and have been shown to excel at image classification tasks.

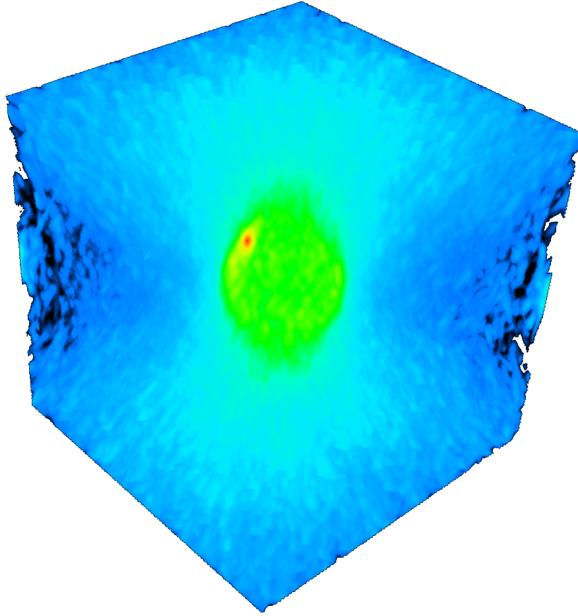


Figure 2: Example of a reconstructed image showing a 10 mm diameter diamond (red) in a 100 mm diameter kimberlite rock (green) surrounded by noise (blue).

The layers listed in Table 2.3 can be divided into seven stages. The first is the input stage which immediately downsamples the input via averaging so that the training time is faster and more manageable. The next four stages are variations on convolutions followed by a max pool. The first of these captures information about patterns in the data, the subsequent stages build upon the previous ones to create and monitor ever more complex patterns and features. The number of features doubles between the stages to capture the exponential growth of pattern complexity and combinations. The sixth stage is a bottleneck which imposes regularisation constraints on the network and forces it to discard useless features. The final stage is the prediction, which uses a sigmoid activation function because it outputs a value between 0 and 1, allowing us to interpret the output as the CNN’s confidence that a sample of rock contains a diamond.

#### 2.4 Training, Testing, and Data preparation

The reconstructed PET images were sorted into three groups according to the size of the diamond in the simulation. The diamond diameters in the first group (DG1) ranged from 0 mm to 5 mm, in the second group (DG2) they ranged from 5 mm to 10 mm, and in the third group (DG3) they ranged from 10 mm to 15 mm. There are seven possible ways to combine these groups, each of which will be used to train a CNN. The images were de-noised with Gaussian blurring using a 1 voxel wide kernel. Then the values in the images were rescaled to the interval  $[0, 1]$  using the global minimum and maximum values of both the training and testing datasets.

Second, each image was subsampled to get 20 diamond samples and 20 barren samples. Each sample is a cube with a side length of 21 voxels. The cubes were sampled by randomising the location of the central voxel. For the 20 diamond cubes the diamond was always fully contained within the cube. While for the 20 barren cubes their central voxel was always in the kimberlite sphere and they were as far away as possible from the diamond. This meant that the 3000 training simulations became 60,000 diamond cubes and 60,000 barren cubes. The test data underwent the same sub sampling treatment as the training data. Giving rise to 6020 diamond cubes and 6020 barren for each of the diamond sizes.

### 3 Results and Discussion

In total, seven CNNs were trained on the seven combinations of the three training dataset groups (DG1, DG2, and DG3). The results are recorded below in Figure 3. These results show that none of the combinations dominate in terms of performance in all four metrics. However, CNNs that were exposed to 0 mm-5 mm (DG1), generally have higher recall on diamonds less than 7 mm in diameter, but at the same time they let through many false positives. Conversely, the CNNs that weren’t exposed to DG1 perform worse on diamonds less than 7 mm in diameter, but almost never let through false positives. This discrepancy is likely down to the cubes containing diamonds less than

Layer	Window	Activation	Output shape	Features
Input			21x21x21	1
Average	2x2x2		11x11x11	1
Convolution	3x3x3	ReLU	11x11x11	64
Max pool	2x2x2		10x10x10	64
Convolution	3x3x3	ReLU	10x10x10	128
Max pool	2x2x2		9x9x9	128
Convolution	3x3x3	ReLU	9x9x9	256
Convolution	3x3x3	ReLU	9x9x9	256
Max pool	2x2x2		8x8x8	256
Convolution	3x3x3	ReLU	8x8x8	512
Convolution	3x3x3	ReLU	8x8x8	512
Max pool	2x2x2		7x7x7	512
Convolution	2x2x2	ReLU	6x6x6	64
Convolution	6x6x6	Sigmoid	1x1x1	1

Table 4: High-level summary of the CNN classifier architecture.

2 mm in diameter beginning to fade into the background. This would confuse the CNN training process giving rise to a low threshold for classifying a cube as containing a diamond. More experimentation is needed to verify which subset of diamond training samples allows a CNN to classify smaller diamonds without increasing the false positive rate.

Image quality may also be limiting the performance of the CNNs, the MLEM algorithm does not provide image quality improvements after only one iteration. Improving the image quality may lead to substantial gains in discriminatory power, which could be achieved by improving the MLEM algorithm used.

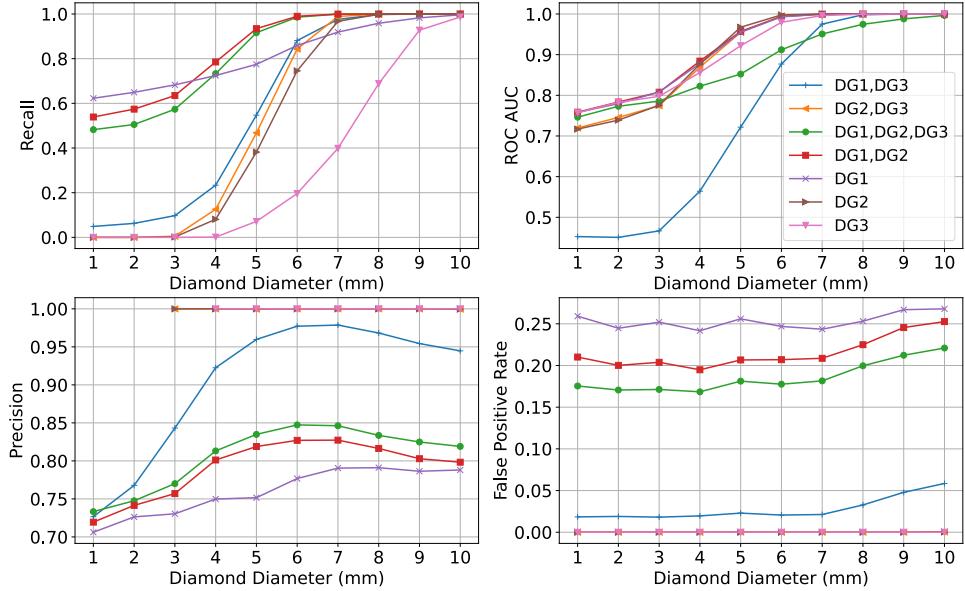


Figure 3: Plots showing test data performance of seven CNNs: Recall (upper left), ROC area under the curve (upper right), Precision (lower left), and False positive rate (lower right).

Out of the seven CNNs, the orange one, trained on DG2 and DG3, is the best performing model. It dominates the pink and brown models because it saw twice as many examples during training. It also maintains a very low false positive rate and a high precision. However, it is not the climax of this study. Importantly it captures only 50% of diamonds at the 5 mm size, which are a critical component of any mine's revenue source. Moreover, the

training data can be improved by increasing the complexity of rock and diamond shapes as well as the number of rocks and the size of cuboid the CNN is trained on. The aim is for a CNN look at 100 cmx30 cmx10 cm belt of rocks and decide whether there are any diamonds. Furthermore, the CNN does not tell where a rock is which is crucial information for the ejection system.

#### 4 Conclusion

The results provide compelling evidence that a CNN can be trained to tell the difference between diamond containing rocks and barren rocks. The limit of the models trained in this study appears to be related to the identification of diamonds of diameter between 5 mm and 6 mm. Further work needs to be carried out to investigate how to lower this limit without increasing the false positive rate. Steps in this direction include improving the image quality and identifying and removing confounding training examples. These can be achieved by improving the MLEM algorithm. Moreover, the structure and hyperparameters of the CNN were not fully investigated, and this should be addressed in future work, for example by using a large cube size.

More work also needs to be done to increase the realism of the simulations. This can be achieved via increasing the number and shape variations of the rocks and diamonds. Modelling a non-uniform activation field will also help increase the realism and thus reliability of the simulations and any results derived from them.

#### References

- [1] E. M. Chinaka, Z. Zibi, J. Van Rooyen, S. H. Connell, and M. M. Cook, “Radiation shielding analysis and optimisation for the minpet kimberlite sorting facility using the monte carlo calculation code, mcnp,” in *The Proceedings of SAIP2016*, 2016, pp. 386–391.
- [2] M. N. H. Cook and S. H. Connell, “Using geant4 to create 3d maps of dosage received within a minpet diamond sorting facility,” in *The Proceedings of SAIP2016*, 2016, pp. 392–397.
- [3] M. Cook and S. Connell, “A genetic algorithm approach to enhancing the performance of a pet detector array,” in *The Proceedings of SAIP2016*, 2016, pp. 398–403.
- [4] M. Cook, M. Tchonang, E. Chinaka, M. Bhamjee, F. Bornman, and S. Connell, “Overview of the mineral-pet run-of-mine diamond bearing rock sorter,” in *The Proceedings of SAIP2016*, 2016, pp. 479–484.
- [5] M. Cook, “Remote detection of light elements using positron emission tomography,” PhD thesis, University of Johannesburg, 2014.
- [6] U. Andres, “Development and prospects of mineral liberation by electrical pulses,” *International Journal of Mineral Processing*, vol. 97, pp. 31–38, 2010. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0301751610000864>
- [7] S. Agostinelli et al. (GEANT4 Collaboration), “GEANT4: a simulation toolkit,” *Nucl. Instrum. Meth.*, vol. A 506, pp. 250–303, 2003.
- [8] W. Van Aarle, W. J. Palenstijn, J. Cant, E. Janssens, F. Bleichrodt, A. Dabrowski, J. De Beenhouwer, K. J. Batenburg, and J. Sijbers, “Fast and Flexible X-ray Tomography Using the ASTRA Toolbox,” *Optics Express*, vol. 24, pp. 25 129–25 147, 2016.
- [9] J. Dewit, “2nd place solution for the 2017 national datascience bowl,” <https://juliandewit.github.io/kaggle-ndsb2017/>, 2017, accessed: 2024-07-29.
- [10] R. Hart, “Table of kimberlite compositions,” *Private Correspondence*, 2007.
- [11] S. Tappe, “Report on Global Kimberlite Major Element Compositions,” *Private UJ report*, 2016.