31st International Conference Nuclear Energy for New Europe
NENE 2022 September 12-15 PORTOROŽ SLOVENIA

Radioactive Environment Characterisation Using Multi Detector Arrays and a 3-Dimensional Scanning Lidar

Matthew Ryan Tucker)
University of Bristol
H H Wills Physics Laboratory
BS8 1TL, Bristol, United Kingdom
matthew.ryantucker@bristol.ac.uk

Thomas B Scott
University of Bristol
H H Wills Physics Laboratory
BS8 1TL, Bristol, United Kingdom
t.b.scott@bristol.ac.uk

ABSTRACT

Radiation data sets can be divided into 4 regimes, and environments should be classified to inform decision making. In this paper, the data collection process with a handheld laser scanner is outlined, and a neural network to classify the data sets is designed, trained and tested. The network was trained on simulated data, but was successfully able to classify several real world data sets.

1 INTRODUCTION

Radiation mapping is a key part of monitoring the continued status of all nuclear sites, both those currently active and those undergoing decommissioning. This mapping can be done by human operators, carrying the necessary sensors, or by robots, such as drones or unmanned ground vehicles. Using robots has many advantages[1], as they minimise risk to humans and can increase efficiency, but some areas of nuclear sites may still be inaccessible to conventional robots. This may be due to physical access constraints, or due to concern from stakeholders about untested technology. Preparing a robot for every conceivable scenario can be challenging, and if the robot fails then it may add to the hazards in the site. Operator carried radiation detectors can therefore be more practical, providing the dose levels are not significant enough to pose a risk to humans.

To position the measurements made in such sites, a laser scanning system can be used to make detailed 3 dimensional maps of the space, through a process known as Simultaneous Localisation and Mapping, or SLAM. Multiple radiation detectors can be mounted on the operator, which provides an immediate directionality to all measurements made: this directionality is enhanced by the attenuation provided by the human body[2], blocking a significant fraction of the radiation. The mapping system used for this work was the Frontier, developed by the Oxford Robotics Institute[3]. This produces high resolution point clouds which can be combined

with coloured cubes representing radiation data: an example of this can be seen in figure 1. This sensing system can also be mounted on robots, but the data gathered in this paper was collected by a human operator.

4 gamma spectrometers are attached to the operator, evenly spaced around them. The sensor system is carried in the user's hands, and a map is built up as they explore the space. The data is fed back in real time to a laptop, or can be streamed to a tablet that the operator can see. The operator can then assess what areas are of most interest by seeing the data as it comes in, and focus their efforts there. A good quality map of the space can quickly be built up, and then the area of the site covered by following a raster pattern, to provide complete coverage. A raster pattern may not always be possible, in which case exploration should continue as most appropriate.

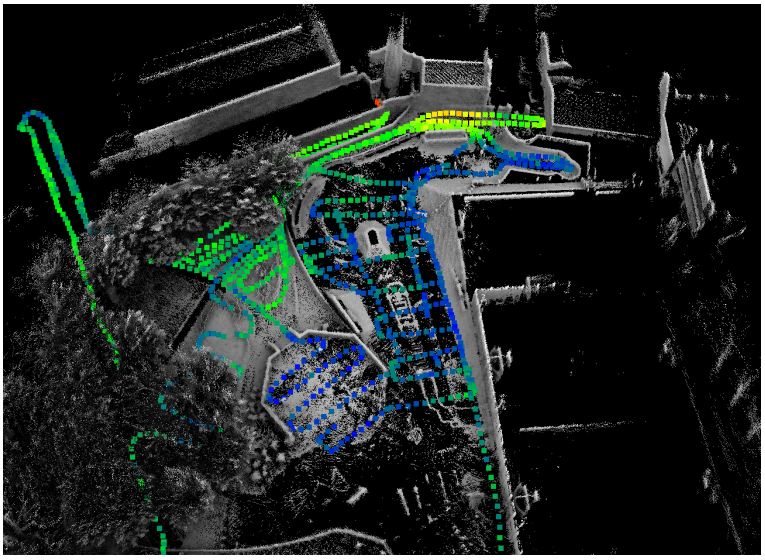

Figure 1: A radiation map on the site of a closed coal mine in Bristol. A hot spot can be seen towards the top of the image. This data set was correctly classified as a hot spot by the network

Once the data is collected, it must then be interpreted to characterise the environment, and influence decision making on the remediation strategy for a site. The environments generally observed can be separated into 4 regimes:

1. Uncontaminated areas
2. Evenly contaminated areas
3. Hot particles
4. Shine paths

Uncontaminated areas will return background readings throughout: the magnitude of this background reading will vary due to local conditions and geology. Evenly contaminated areas, such as sites displaying homogeneous deposition of radioactive dust, will have radiation readings substantially above background, but with no clear hot spots. Areas with hot particles will have areas of elevated radiation, and detected counts will rise roughly following the inverse square law as the source location is approached. A shine path occurs when a source is partially shielded from the survey area, but with some gaps: this can result in abrupt discontinuities in radiation levels in areas close to each other.

These regimes can be identified by human operators, but this can be prone to error, and information can be missed. Automating this may reduce error, and would be a key part of a

completely automated monitoring system. Developing the techniques to automatically categorise these environments is therefore crucial, even if the data used was gathered by a human. There are several approaches that could be taken. One approach would be to analyse the gradients between the measurements: if there are any large discontinuities, this would indicate a shine path: if the gradient is very low, this would indicate either even contamination or an uncontaminated area. This can be difficult to implement, and can be computationally expensive.

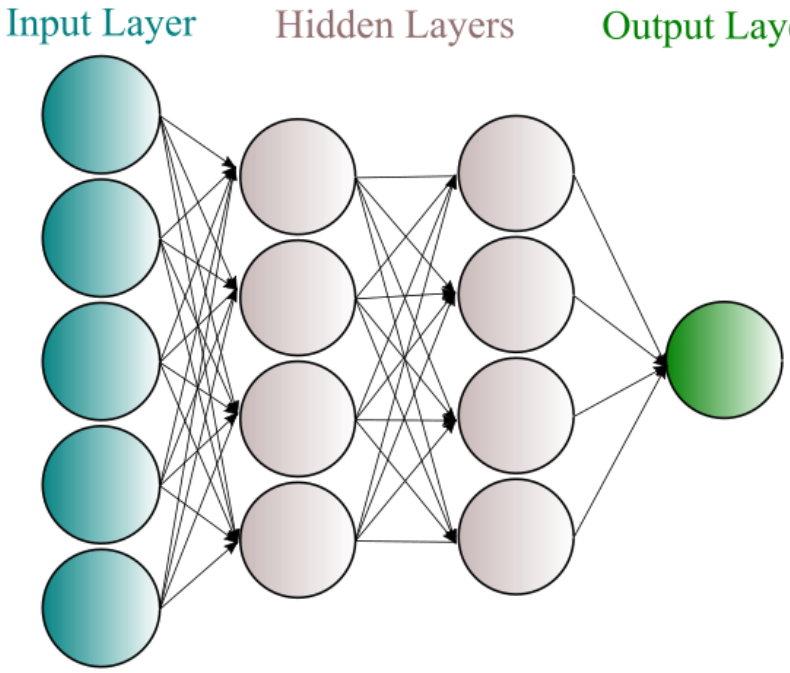

Figure 2: A simplified neural network, with 5 nodes in the input layer, two hidden layers with 4 nodes in each, and an output layer with 1 node. All nodes are connected to all nodes in the subsequent layers.

Another approach is to train a neural network to interpret this data and predict which category the data set falls into. A neural network is an example of machine learning, where a computer can be supplied with training data and learn how to perform a task such as classification. Neural networks were initially designed to emulate models of human thought[4], but the similarity between modern neural networks and the human brain is low. Neural networks have one input layer, which has a size equal to the data input: if the input is a 20x20 pixel picture, the input layer would be made up 400 nodes, one for each pixel. The network has a number of hidden layers, with each layer returning 'weights': variables which are applied to your input. The number of nodes used by each of these hidden layers is set by the designer of the network, with each node returning a weight. A simplified example is shown in figure 2.

In a simple neural network, every node in a layer is connected to every node in the successive layer (fully connected). The final layer of a neural network is the output later, which can be used to make a prediction. The output layer shape can vary: if the network is performing a binary classification task, the output layer may be one node, and return a value between 0 and 1. If the value is greater than 0.5, this could be interpreted as one of the two categories: if the value is below 0.5, the other prediction would be made. In a multiple classification problem, the output layer is normally the same size as the number of options: the classification can be made by choosing the output with the highest value.

Neural Networks are trained by inputting training data. The network will assess the 'loss', which is a measure of difference between the network's prediction and the true label of the data. It will then adjust the weights of each layer, and over many iterations attempt to converge on the weights which give the most accurate prediction. Customarily, training data is split into two groups, with roughly a third of the data held back and not used to train the data. This 'test' data can then be used to analyse the performance of the network, once training is complete. A danger with neural networks is that they can 'over fit', showing a high accuracy on training data but a much lower accuracy on test data. This can sometimes be mitigated through use of a regularisation term, which is set up to favour lower weights, and as such can make the model less specific.

One limitation to machine learning is that a large quantity of data is needed. For problems such as radioactive environment characterisation, it can be difficult to acquire this quantity of data. For this reason, the training data for this network has been simulated, and then evaluated against real data sets.

2 METHODS

The $10,000$ generated examples were randomly divided between each of the 4 regimes. First, a field was generated, with a number sources randomly distributed on a 100x100 grid. If the simulated regime was a shine path, obstacles were added around the source, which were assumed to provide total attenuation. The distance of each point on the grid to each of the sources was calculated, and whether or not the path from the point to the source intersected any attenuating obstacles. A successful approach should be able to differentiate between the different regimes using this data, but it is incredibly unlikely that complete coverage of the area will be achieved. Therefore, it is also important to sub-sample this data to mimic the process of data collection.

To do this, a path was generated, with a random starting location. Two path algorithms were used: the first generated a completely random path, with the direction changing at random intervals. The second generated raster paths, covering the grid systematically. For both paths, noise was added to make the paths less regular, and therefore closer to a real world example.

Once the path was generated, a reading was taken for each point in the path. This reading was a combination of the contributions from the source, adjusted for the distance, and a background reading. Both readings used a Poisson distribution to mimic the randomness of radioactive decay. The attenuation of radiation due to air and any scattering affects were not taken into account for the training data set.

To ensure that this classification problem could be solved by machine learning, the full data sets were worked with initially, rather than the sampled ones. Machine learning models often converge faster if inputs are lower in value, so all datasets were initially normalised by dividing them by the maximum counts possible in the simulation, $10,000$. The neural network architecture used for this was 2 hidden layers, the first with 25 nodes, and the second with 15, both using relu activation functions and a small regularisation term. The network was implemented through Tensor Flow, a popular Python package for machine learning. The model rapidly converged to an accuracy of over 99% on both the training data and the test data, indicating that the classification was effective.

The same architecture was then reused, and the network trained again, on the sample arrays, those generated using the path. An accuracy of 96% was found on the training data, but the test data accuracy was substantially lower, at 75%. The sample arrays will not display all of the features that are present, and will be less similar to one another than the full field

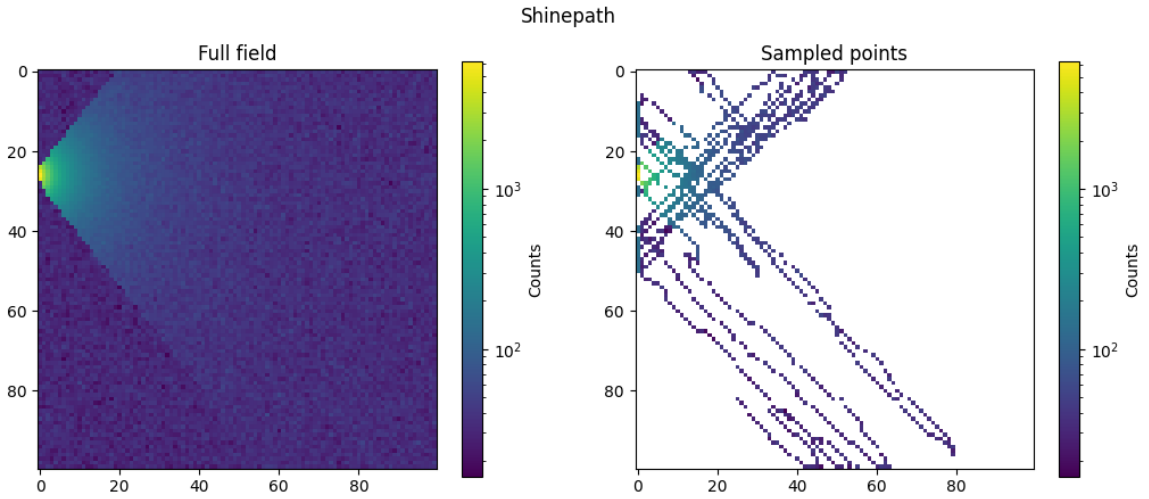

Figure 3: A example of the simulated shine path data, with the full field on the left and a simulated path on the right. The source is placed outside the bounds of the image, with a simulated border with a small gap in it.

arrays. As such, the neural network is more liable to over-fit: the network essentially learns the features of the training data set, instead of generalising to all data. Increasing the regularisation term decreased the accuracy on the training data to 75%, but did not result in an increase to the accuracy of the test data set, which remained at 75%.

Sometimes the simulated path may not have captured data which accurately predicts the field. It is important to find some point of comparison for the 75% accuracy, so humans were tested at the same task, using the sample data. Experienced humans could achieve an accuracy of approximately 85%, indicating that there is room for improvement, but the neural network outperformed inexperienced operators, with an average score of 70%.

The network was then tested on several real world data sets, to evaluate its performance. To successfully interpret data sets, they must be stretched to the same size as the training data, 100x100 pixels, and the data normalised by dividing by the maximum counts possible. The network outputs a vector containing 4 numbers, which can be processed using the softmax algorithm[5] to return the probability of the data set falling into each classification.

3 DISCUSSION

The network correctly categorised the Dean Lane dataset (Figure 1) as a hotspot, the Trevalour clay pit as background, and the Chernobyl dataset as Elevated (both in figure 4). The Trevalour dataset does not have any particular regions of interest: the left hand region is higher than the right, but it is not a substantial difference, and the radiation levels do not fall off in an inverse square manner: the categorisation of background therefore seems appropriate. The counts present are similar to the background level the model was trained with.

The Chernobyl data has background radiation counts which are several orders of magntiude higher than the value the dataset was trained on. A characterisation of Elevated seems accurate for the majority of the dataset, but there is a clear hot spot at the top of the image. The model returned a prediction of Elevated (72%) and Hot spot (28%), so the hot spot has not completely been disregarded. A shine path data set has not been collected with the equipment modelled: a data set should be recorded to provide a more comprehensive test.

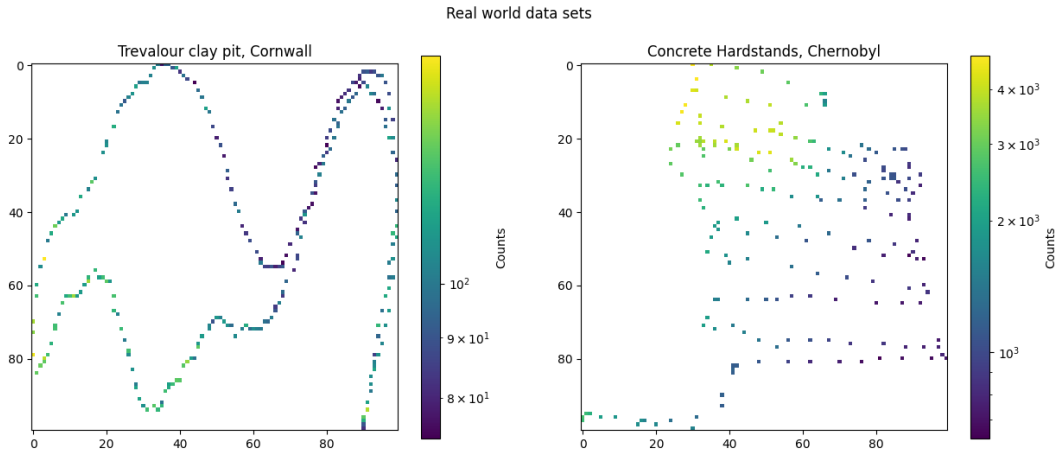

Figure 4: Datasets from Trevalour and Chernobyl. Trevalour prediction: Background (71%), Shinepath(19%), Hotspot(10%). Chernobyl prediction: Elevated (72%), Hotspot(28%).

Overall, the accuracy of network was acceptable, with a similar accuracy to human classification. In order to test it more extensively, it should be applied to more real world data sets, and its performance assessed. One issue with the network is that it was trained on individual examples: all of the hot spot and shine path data sets had a background level set to that present in Bristol. For examples like the Chernobyl data (figure 4), the background is substantially higher, but the hot spot is still the predominant feature. To fix this, the training data could be diversified, with various background rates. The network was also only trained and tested on data with one source or region of interest: this should be improved on for future work.

To increase the accuracy of the network, it may be necessary to alter the architecture. The current network works excellently when it has a full data set, which will not occur often in reality. When the data is subsampled along a path, much less data is collected. It may be that the only indication a shine path is present are two readings side by side, one which is at background and the next several orders of magnitude higher. To notice these small differences, the network could incorporate convolution steps, which can subdivide the data into smaller sections and therefore be more likely to notice this.

One issue with any neural network is that the reasons networks make decisions are unclear. For example, the Trevalour data set (figure 4) returns a classification of background, but with a 19% probability of a hotspot, which does not fit the observed data. For this reason, it is important that neural networks have a high accuracy and have been extensively tested before relying on them for critical responsibilities, and they should be deployed alongside other fail safe measures: if the counts detected are above a threshold, categorisations could be flagged for double checking by a human.

The ability to train a neural network to a good degree of accuracy with data generated by simulated python code is surprising, and opens up new areas for machine learning to be deployed in. Several assumptions were made for the sake of simplicity in the simulation, but it was still sufficient to train the network. The use of sophisticated particle modelling packages such as GEANT4 to properly account for attenuation and scattering would increase the quality of the training data, and may therefore produce better result. This data would take substantially longer to generate, and this is unlikely to be the limiting step for performance: time would be better spent improving the architecture. Substantially more sophisticated Python scripts for generating similar data do exist, such as those used by *Nicholson et al (2020)*[6], but these

require more computation time than was available. The data set demonstrated was also only suitable for straight line travel as if along a street, but the approach described could be adapted to future work.

Successful radioactive environment classification is a key component to the goal of increasing automation in the nuclear industry, and therefore making it safer and more affordable. With more data being gathered in nuclear power plants than at any point in the past, there is great potential for approaches such as machine learning to make use of this.

References

[1] Ioannis Tsitsimpelis et al. "A review of ground-based robotic systems for the characterization of nuclear environments". In: *Progress in Nuclear Energy* 111 (Mar. 2019), pp. 109–124. ISSN: 0149-1970.

[2] E Buchanan et al. "Operator related attenuation effects in radiometric surveys". In: *Radiation Measurements* 86 (), pp. 24–31.

[3] Milad Ramezani et al. "The newer college dataset: Handheld LiDAR, inertial and vision with ground truth". In: *IEEE International Conference on Intelligent Robots and Systems* (Oct. 2020), pp. 4353–4360. ISSN: 21530866.

[4] Warren S. McCulloch and Walter Pitts. "A logical calculus of the ideas immanent in nervous activity". In: *The bulletin of mathematical biophysics 1943 5:4* 5.4 (Dec. 1943), pp. 115–133. ISSN: 1522-9602.

[5] John S. Bridle. "Probabilistic Interpretation of Feedforward Classification Network Outputs, with Relationships to Statistical Pattern Recognition". In: *Neurocomputing* (1990), pp. 227–236.

[6] Andrew D. Nicholson et al. "Generation of Synthetic Data for a Radiation Detection Algorithm Competition". In: *IEEE Transactions on Nuclear Science* 67.8 (Aug. 2020), pp. 1968–1975. ISSN: 15581578.

ACKNOWLEDGMENTS

The author is grateful to their collaborators in Ukraine for access to sites of scientific interest, and to Andrew Ng's online machine learning courses for providing valuable instruction