Mine Closure 2024 – AB Fourie, M Tibbett & G Boggs (eds)

doi:10.36487/ACG_repo/2415_44

Mapping tailings storage facilities associated with abandoned mine sites

Mbalenhle Mpanza [a,*], Jillian Wistockk [b], Steven Rupprecht [a]

[a] *Department of Mining Engineering and Mine Surveying, University of Johannesburg, South Africa*

[b] *Esri, South Africa*

Abstract

Besides tailings storage facility (TSF) structural failure, there are various other threats that result from poor management of TSFs. For example, poor rehabilitation of TSFs can result in air pollution by wind-generated dust affecting human health and wellbeing. Furthermore, water pollution owing to acid rock drainage can result from a lack of TSF lining. Recently some TSFs have commenced posing a safety challenge for communities residing in proximity as they become infested with illegal mining activities. The 2020 Global Industry Standard on Tailings Management (GISTM) hopes to achieve zero harm to people and the environment, with zero tolerance for human fatality. To conform with the GISTM, companies are required to be accountable and responsible, and all stakeholders must prioritise TSF management. Two recent TSF-related cases (the Jagersfontein tailings dam failure and the crime activity by illegal miners at the North Sands Dump, both in South Africa) are significant motivations for effective tailings management. It appears that the significant challenges with TSF management in South Africa result largely from mine operation abandonment and legacy mine sites. This study suggests the mapping of TSFs associated with abandoned mine sites as a starting point for effective TSF management in South Africa. The study uses Faster Region Convolutional Neural Network (RCNN) and the Mask Region Convolutional Neural Network deep learning models to identify TSFs in Gauteng Province. Satellite (Sentinel-2) imagery data was used to train the RCNN models for TSF identification focusing on the red, green and blue band combinations. Sentinel-2 satellite imagery was obtained and processed using ESRI's Arc GIS Pro software and was used to train and test deep neural networks on the same platform using the Pro Notebook with inbuilt Python scripting. The Mask RCNN was able to accurately identify TSFs (which the model was not trained on) with minimal false positives. The efficacy of the approach is demonstrated by the discovery of 46 TSFs that were not part of the training data. This study highlights the potential of deep learning to assist in TSF management by identifying TSFs which might be ownerless and abandoned.

Keywords: *faster RCNN, Sentinel-2, object detection, TSFs, Mask RCNN*

1 Introduction

The conundrum of having large-scale mining companies divesting their mining operations and selling them to junior miners is part of the reasons mine abandonment occurs in South Africa (Liefferink 2019). Junior mining companies acquire mining assets that are generally marginal and have high rehabilitation costs, thus leading to them insolvency. It is common for junior mining companies to lack the necessary funding to finance expensive rehabilitation projects. Unfortunately, the large-scale companies never come to the rescue of the junior miners during the period of insolvency. Once a mining company becomes insolvent it ceases its operation, which results in sudden mine closure. The sudden closure results in mine abandonment and delays in mine rehabilitation projects (Mpanza 2022). Tailings storage facilities (TSFs) remain poorly managed and unrehabilitated when sudden mine closure occurs as companies can circumvent their management during

* Corresponding author. Email address: mmpanza@uj.ac.za

the liquidation period (Humby 2014). Poorly managed TSFs can result in dam breaks, economic losses and irreparable environmental pollution (Kossouf et al. 2014). Therefore, the mapping of TSFs associated with abandoned mine sites resulting from sudden mine closure is of great significance to regional effective resource utilisation, surrounding community health and safety monitoring. Traditional ground surveys of TSFs are costly and inefficient (Zabcic et al. 2014), with TSFs often difficult to detect from the ground and much effort being required. However, satellite or airborne sensing coupled with geographic information systems (GIS) can be used over a wide area to assess where hazardous natural or man-made phenomena have occurred or are likely to occur. This information can feed into risk assessments or illegal activities detection, along with information on natural resources, population and infrastructure. It can be used by governments, e.g. the Department of Mineral Resources and Energy in South Africa, and communities to prevent illegal activities and to design mitigation strategies that reduce environmental vulnerabilities to acceptable levels. The use of remote sensing and data science in mapping large-scale surface mines and TSFs has proven to be cost-effective (Balaniuk et al. 2020).

The mapping of tailings mine waste is a critical aspect of mining operations. Traditional methods of identifying tailings mine waste often rely on manual observation, surveys and analysis, which can be time-consuming, costly and subject to human error. The use of artificial intelligence such as convolutional neural networks (CNN) has emerged as a promising solution for the automatic mapping of TSFs. By analysing satellite images of TSFs using CNN, it is possible to develop a reliable and efficient system for accurately identifying and classifying different types of TSFs. Deep learning can obtain higher dimensional and abstract features of the images than traditional methods and it has been proven to be powerful technology for remote sensing image processing (Kussul et al. 2017; Zhang et al. 2016; Ball et al. 2017). Satellite imagery alone is inadequate to map TSFs owing to excessive variation in tone, shape and dimension between TSFs, and because the methods are difficult to apply TO a large area. Moreover, TSFs are more sparsely distributed than other objects, making their identification in a large area more difficult and time-consuming (Mazzia et al. 2020).

Davis et al. (2021) designed an onsite waste classification system using a deep learning method that can classify different categories of waste based on digital photographs taken from construction sites. This system achieved an accuracy rate of 95% for waste classification, demonstrating the potential of CNN in effectively identifying tailings mine waste and other types of waste materials. Furthermore, the implementation of CNN for identifying tailings mine waste can streamline the waste management process and improve overall operational efficiency.

In South Africa there is not yet a well-documented extensive inventory of TSFs classified according to their operational and ownership status, i.e. active-owned, dormant-owned, active-ownerless, dormant-ownerless and abandoned. The lack of comprehensive TSF inventories prevents a significant analysis of technical failures, crime action and illegal mining activities which could serve as a learning tool and help prevent future incidents. Existing records are often incomplete in crucial data elements: design height of the dam; the design footprint; construction type (e.g. upstream, downstream and centre line); age; design life; construction status; ownership status; capacity; release volume; location; etc. This study is, in part, driven by the lack of databases containing relevant information about TSFs. The specific objectives of this study were to: (1) map TSFs associated with abandoned mines to monitor activities around these sites; (2) use a low-cost method for detecting TSFs and abandoned mines by using deep learning and freely available satellite imagery; and (3) identify South African TSFs and mines of very different shapes, dimensions and hues in low-resolution imagery.

2 Materials and methods

This section discusses the research methodology and the procedure followed to conduct this study.

2.1 Study site

This study started by considering TSFs from the Gauteng Province alone. However, as machine learning requires significant amounts of data to train models, the area was expanded. TSFs in the Mpumalanga and Limpopo Provinces were also included as part of the study site (Figure 1).

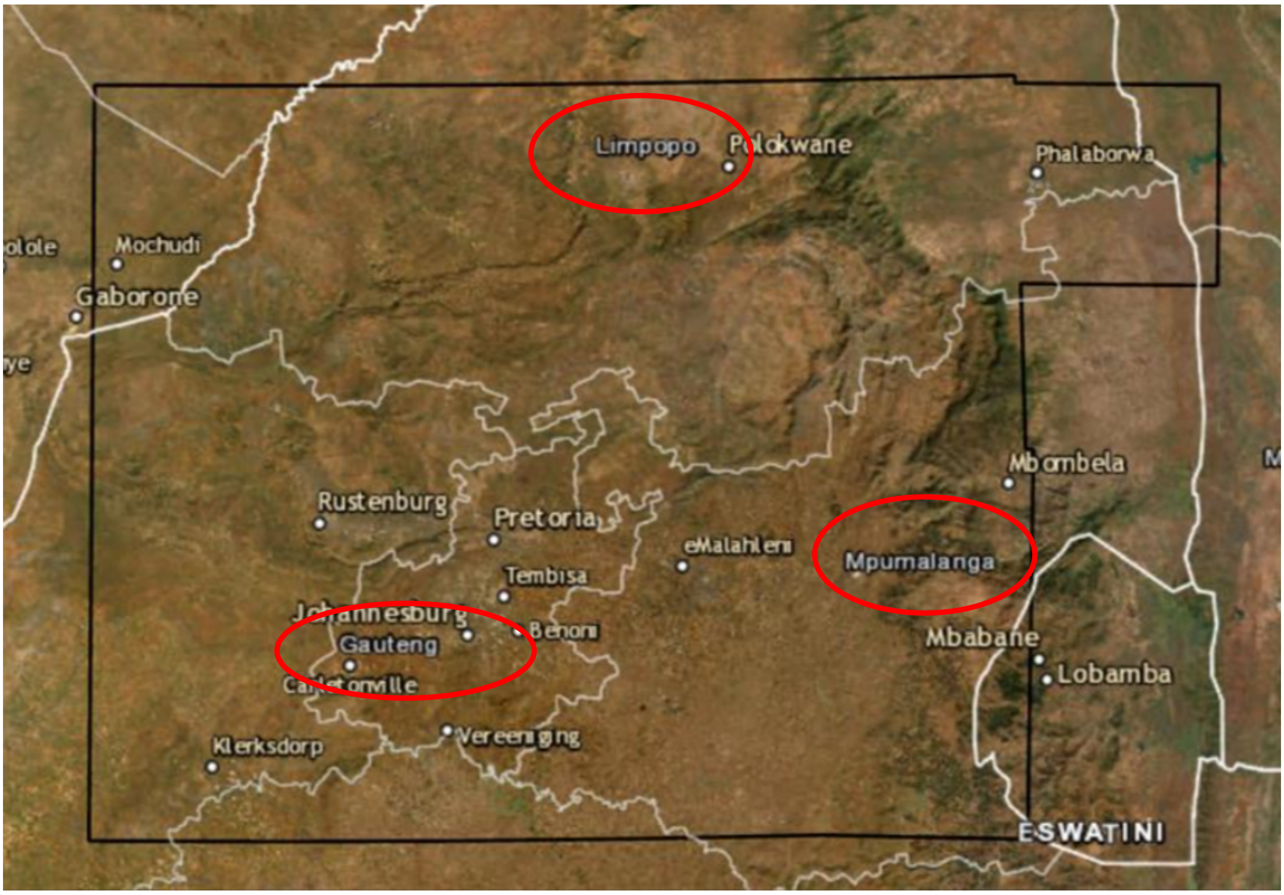

Figure 1 Study site map showing the Gauteng, Mpumalanga and Limpopo Provinces circled in red

2.2 Study approach

This study uses CNN algorithms to identify TSFs associated with abandoned mines from satellite imagery data. CNN are deep learning algorithms for computer vision tasks such as object detection, image segmentation, facial recognition and video analysis. In this study, TSF objects from satellite images were trained using Faster RCNN and Mask RCNN. Instead of working on a massive number of regions, this algorithm proposes a number of boxes in the image and checks to see if any of them contain an object. RCNN uses a selective search to extract these boxes (called regions) from an image. Four variables were used in this study to define the box of an object, and they were varying scales, colours, textures and enclosure.

The procedure for RCNN is that an image is input and then segmented to obtain multiple regions. The algorithm combines similar regions to form a larger region using the four variables. A pre-trained CNN is retrained using the last layer of the network, based on the number of classes that need to be detected. The region of interest (ROI) is defined for each image and the regions are reshaped to match the CNN input size. Once regions have been obtained, the support vector machine is trained to classify objects and background. Lastly, a linear regression model is trained to generate tighter bounding boxes for each identified object in the image.

In this study both Faster and Mask RCNN models were used. Faster RCNN uses a single, end-to-end, unified network. The network is said to accurately and quickly predict the locations of different objects, and the model returns bounding boxes for the locations of detected objects. It replaces the selective search method with a region proposal network (RPN) which makes the algorithm much faster, taking approximately 0.2 seconds. A limitation of Faster RCNN is that object proposal takes time and, as there are different systems working one after the other, the performance of systems depends on how the previous system has performed. Figure 2 summarises the Faster RCNN framework (Ren et al. 2015).

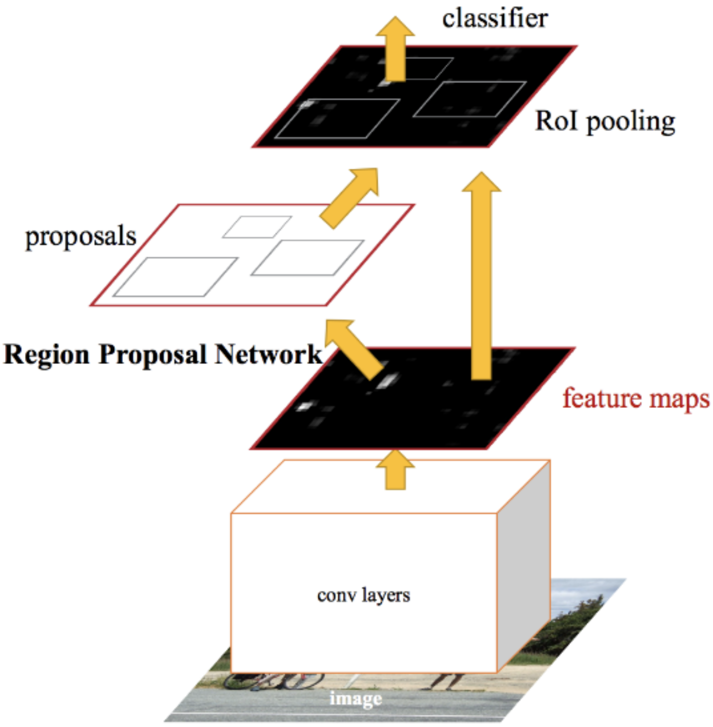

Figure 2 Framework for Faster RCNN: a single, unified network for object detection (Ren et al. 2015)

The steps followed by a Faster RCNN algorithm to detect objects in an image are as follows:

- take an input image and pass it to the convolutional network (ConvNet), which returns feature maps for the image
- apply an RPN to these feature maps and obtain object proposals
- apply an region of interest (RoI) pooling layer to make all the proposals the same size
- pass these proposals to a fully connected layer in order to predict and classify the bounding boxes.

The second model used in this study, Mask RCNN, is also an object detection model based on deep CNNs. The main benefit of the model is that it can return both the bounding box and a mask for each detected object in an image. The model combines elements from the classical computer vision tasks of object detection, where the goal is to classify individual objects and localise each using a bounding box, and semantic segmentation, where the goal is to classify each pixel into a fixed set of categories without differentiating object instances (He et al. 2018). Mask RCNN is applied to each RoI, predicting a segmentation mask in a pixel-to-pixel manner.

2.2.1 Data acquisition

A Department of Mineral Resources and Energy (DMRE) database contains information about each TSF and abandoned mine site in South Africa including the owner's name, the operational activity status (active or dormant) and the spatial coordinates. The data accessed for this paper covers Gauteng Province, which has predominantly gold mining activities. The database has 329 active and dormant-owned TSFs for Gauteng Province.

The database has data quality issues: a significant number of coordinates point to locations near the mines or TSFs but not exactly to them, or they are completely off-target. A visual inspection of each reference using satellite images was required to address this problem and correct the target coordinates, and only 90 TSFs in the database could be accurately used. The abandoned mine sites' dataset could not be used for validations as most of the coordinates pointed to existing buildings and incorrect areas.

Reliable ground truth labelled imagery is required for supervised machine learning. This study used known starting points to georeference data from the DMRE database in order to guide the RoI selection. For remote sensing applications it is imperative to obtain a representative collection of satellite imagery, resolve technical issues such as cloud cover, select appropriate spectral bands and undertake image registration.

The study used Sentinel-2 imagery owing to its free access. The images were extracted for 31 May 2023 because at this time rain had started to subside, there were less clouds to interfere with the images and the vegetation was still largely green. The date was an important consideration as there is a clear contrast between bare and vegetated land, and TSFs generally have low to no vegetation. The Sentinel imagery was obtained from the European Space Agency Copernicus Sentinel program datasets. Sentinel-2 was chosen for its resolution and revisit frequency: the former being 10 m and the latter being 10 days. This study used red, green and blue Sentinel band combinations, and the resolutions and wavelengths used are summarised in Table 1.

Table 1 Sentinel spectral bands used in this study (S2 MSI ESL Team 2022)

Bands	Wavelength (nm)	Bandwidth (nm)	Resolution (m)
Blue	492.4	66	10 m
Green	559.8	36	10 m
Red	664.6	31	10 m

The blue band reflects spectral signatures of soil, vegetation and man-made features, while the red band is commonly used in urban area mapping. The green band reflects spectral signatures of water, oil, vegetation and buildings. Other Sentinel bands were resulting in high false positives while the RGB band combinations had less false positives and were thus used in this study.

142 TSF samples and 131 non-TSF samples were used to train the deep learning models. The non-TSF class consisted of objects such as mine pits, mine stockpiles and other mining structures that a deep learning model could mistake for a TSF if not trained adequately in the difference. Image chips covering the areas of training samples were extracted from the Sentinel-2 imagery and were used for the deep learning models.

2.2.2 Satellite data and image processing

To process the Sentinel-2 COPERNICUS/S2 image collection, the ArcGIS Pro platform was used. Using the ArcGIS Pro raster toolboxes, the Sentinel images were mosaicked to form an image of the entire study area. The platform has an interface for the development of algorithms using Python in the Pro Notebook, and applications of image pre-processing features that allow the solving of some common problems in geoprocessing, such as cloud cover, were used. The steps involved were:

- selection of the study area and downloading of the Sentinel-2 imagery for that area
- image processing in order to obtain a single raster image
- application of a pixel-based median operation to the image stack to define a single raster image. The median value at a location was the most common valid (i.e. no cloud) pixel value from all images at that location.

2.2.3 Model training

The deep learning classifiers require samples of all target classes to train them. This includes the background class, which are non-TSFs: for example, dams, mine pits, stockpiles and other mining structures that a deep learning model may mistake for a TSF. It was also necessary to create a database with the coordinates of locations where no mines or TSFs were present. A careful choice of representative locations was necessary in order to capture the diversity of the background class, such as houses, towns, construction sites and lakes. From this dataset, bounding boxes centred on these (latitude, longitude) coordinates defined the sub-imagery used for training the classifier. A small square image was extracted around each coordinate, centred on the coordinate. Each of these sub-images contained enough information to determine the presence, or not, of a TSF. Images were extracted on a 448 × 448 pixel grid, with each pixel representing a 10 × 10 m 'on the ground' measurement. Each 448 × 448 pixel grid represented data in the spectral bands, being blue, green and red. The image extracts were used to train and validate the deep learning models on the presence of TSFs.

For input to the classification algorithm the coordinates of the points of interest must be in tabular format and made available as comma-separated values or Esri shapefiles. Shapefiles were imported into Arc GIS Pro as feature collections. A small square image was extracted around each coordinate, centred on the coordinate. Each sub-image contained enough information to determine the presence, or not, of a TSF. Figure 3 shows an example of a TSF and Figure 4 shows the same spot on a cropped image from the Sentinel-2 collection.

Figure 3 Sentinel-2 image of a tailings storage facility

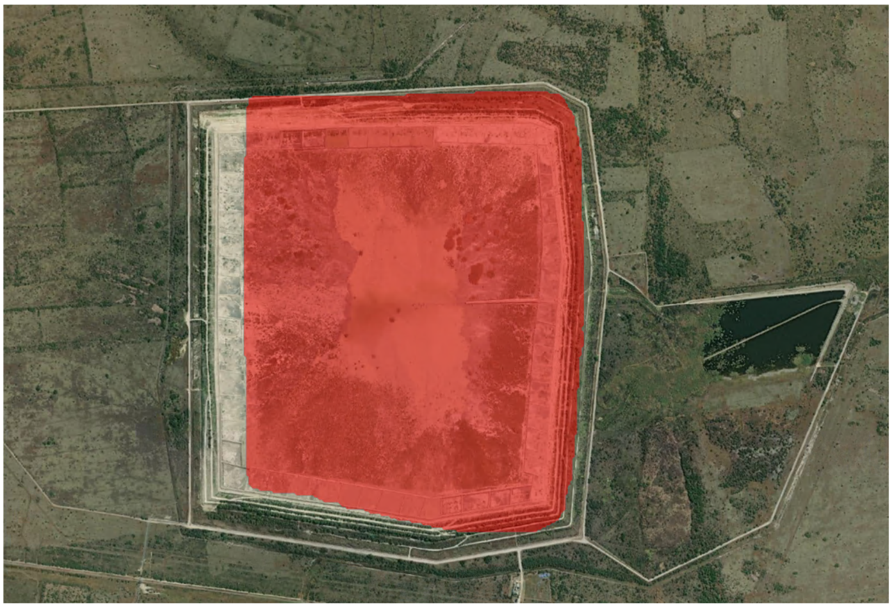

Figure 4 Sentinel-2 image with a tailings storage facility identified by Mask RCNN

3 Results

3.1 Object detection

Figure 5 shows a sample of results from the Faster RCNN model. Figure 5a images are the training samples provided to the model while Figure 5b images are the model predictions. The model returns bounding boxes of the identified objects. The Faster RCNN was able, without input, to identify the TSF shown in the bottom left corner of the top image in Figure 5b and the model had not been trained on this particular TSF.

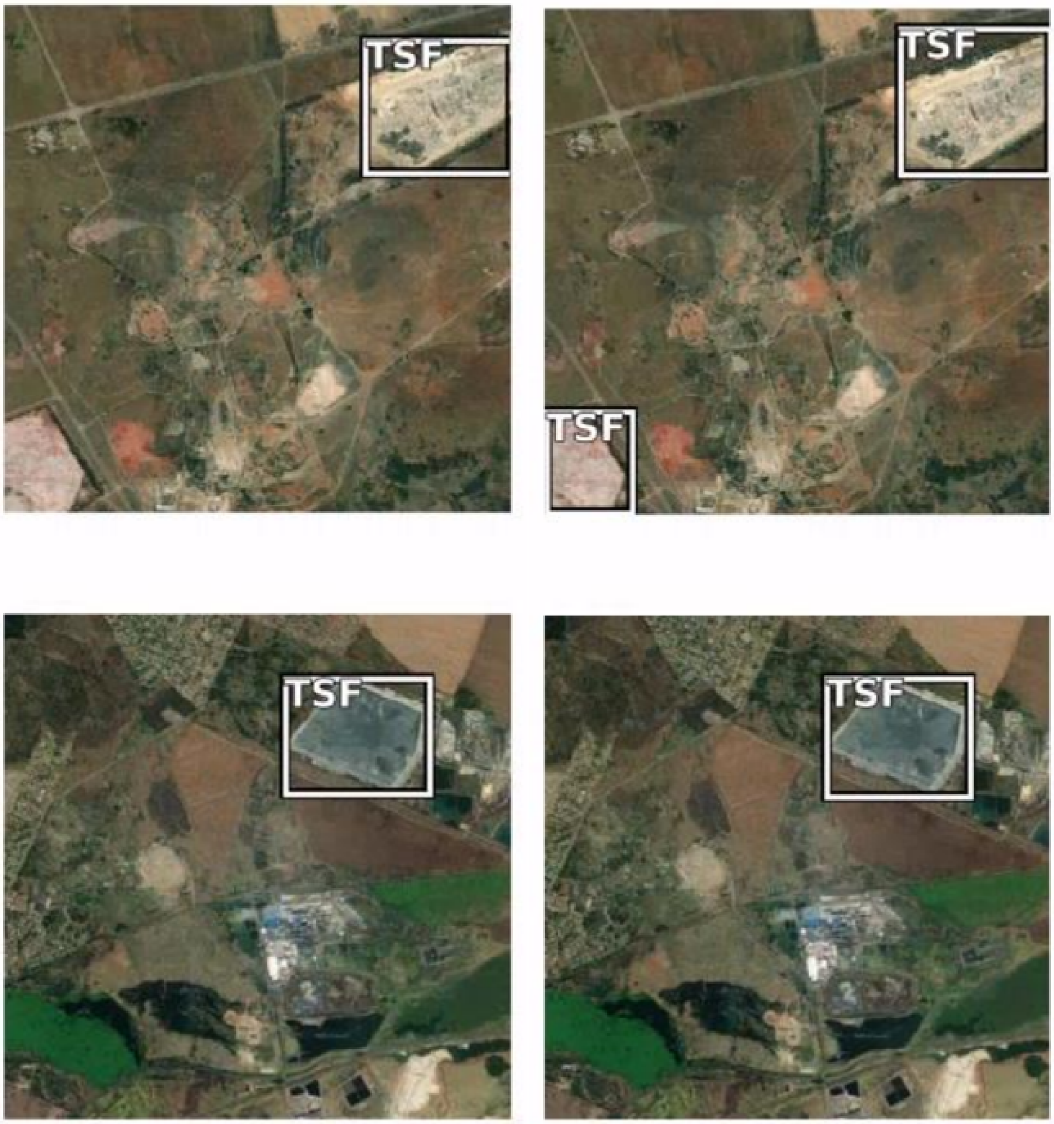

(a) (b)

Figure 5 Sample of results from the Faster RCNN model

The images in Figure 6a are the training samples provided to the model while the images in Figure 6b are the model predictions. The model returns the bounding boxes as well as the mask of the objects.

It was decided to only run further analysis for the Mask RCNN model as it provided the best results in this study.

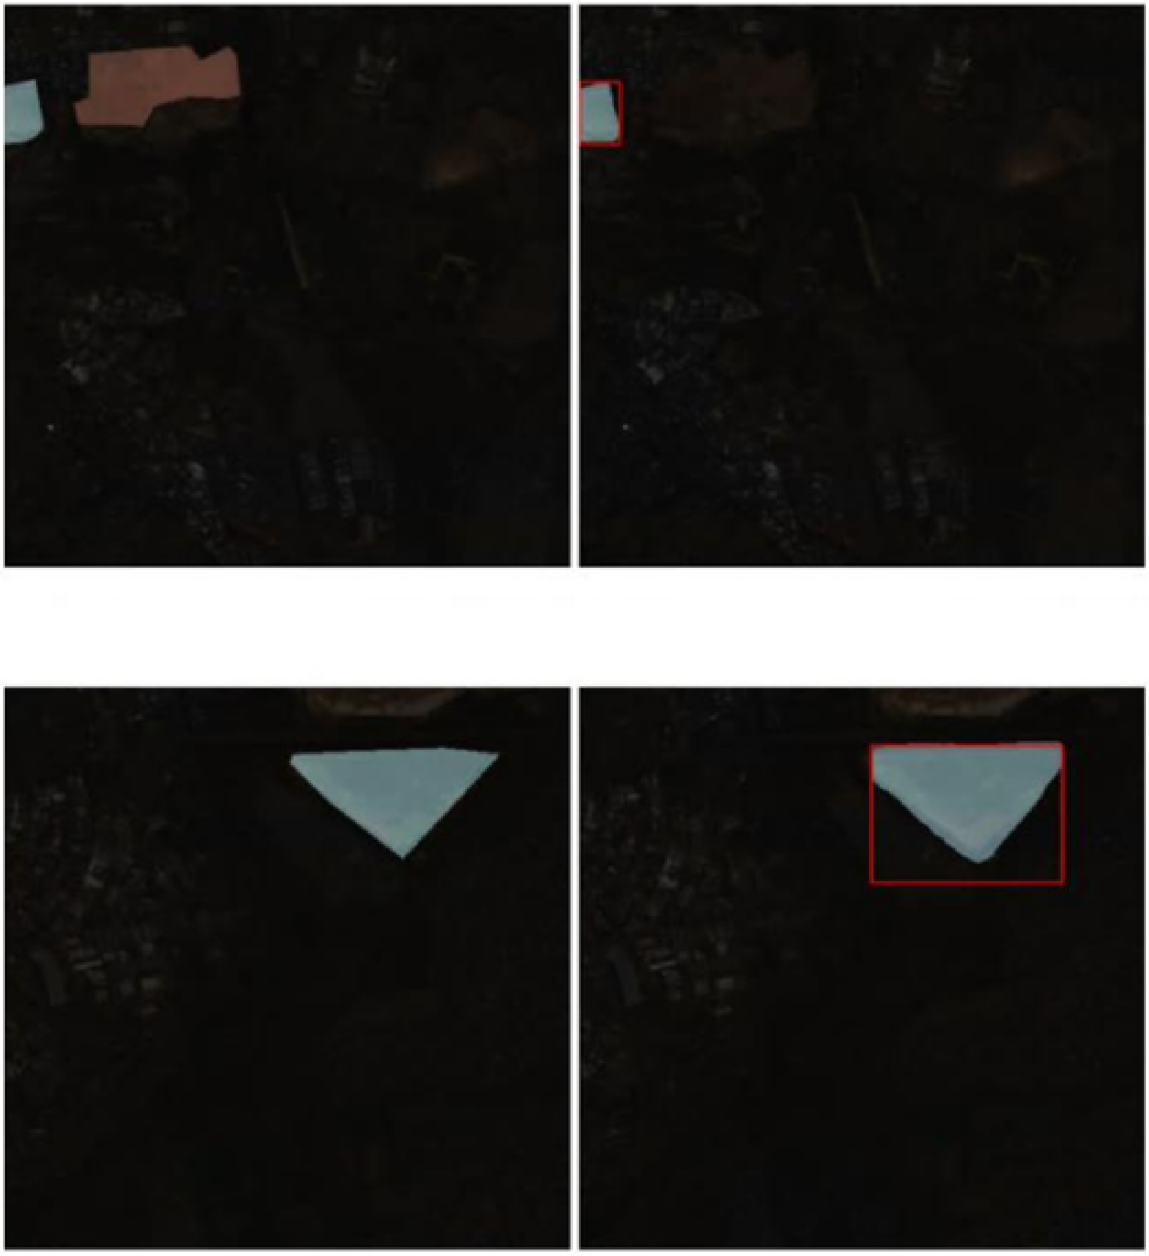

(a) (b)

Figure 6 Sample of results from the Mask RCNN model

When running the Detect Objects Using Deep Learning tools in Arc GIS Pro, a feature class containing a total of 2,353 objects was created. When analysing these results, 459 objects were classified as TSF and 1,894 were classified as other by the Mask RCNN model.

A process of manual verification (by visual inspection) was carried out on the 459 TSF objects. Each detected TSF was manually compared to the satellite image to verify if the object was indeed a TSF. The number of valid TSF detections was 183 objects and the invalid (false positives) TSF detections numbered 276.

Further analysis showed that of the 183 accurate positive detections, 137 intersected with the original training samples provided to the model while 46 of the TSF objects were detected by the model on its own and not provided in the training samples.

3.2 Tailings storage facility mapping

Figure 7 shows some of the TSFs that were accurately identified from the training data and those that were detected by the model. Figure 7a mostly shows TSF objects that were detected because of model training, indicated by a black bounding box around the TSF object.

(a)

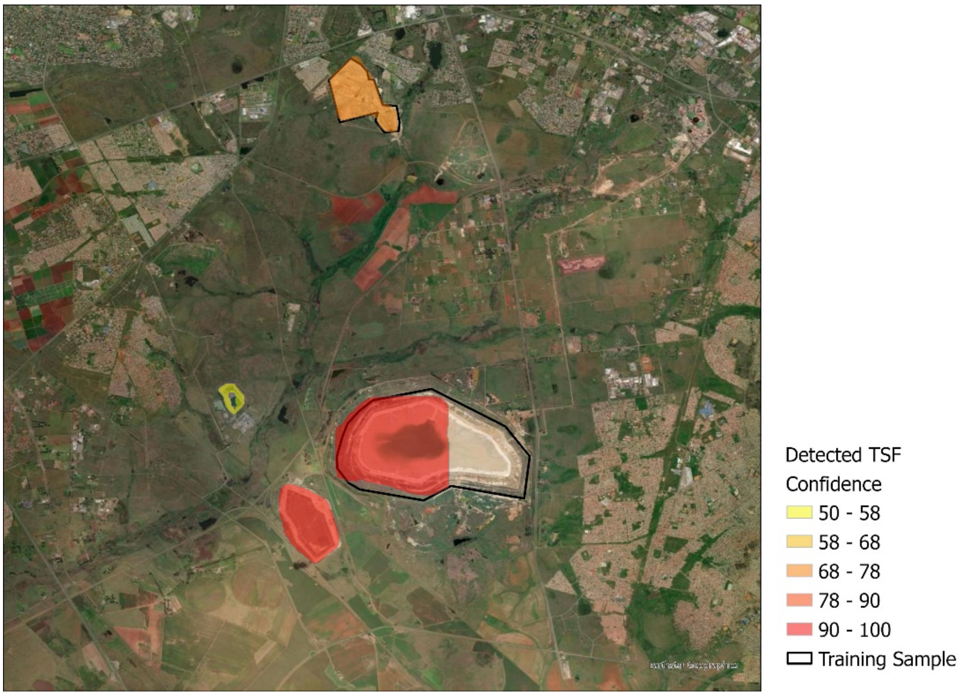

(b)

Figure 7 (a) Tailings storage facility (TSF) objects detected by the Mask RCNN model using training data; (b) TSF objects detected by the Mask RCNN model

Figure 7b shows some of the TSF objects that were detected by the Mask RCNN model 'on its own', i.e. without relying on the training data. The objects detected are shown with no black bounding box around. There is a level of confidence attached to the model detection: for example, the strong red-coloured TSF objects show high confidence (90–100%) while the yellow-to-orange TSF objects were detected with lower confidence (50–68%).

The results in Figure 8 shows some of the detected TSF objects. When analysing some of the results, the benefits and the downsides of using deep learning techniques to detect the TSF objects are clearly seen. The results show that TSF objects can indeed be detected by deep learning techniques but that the masks around these TSF objects are not always accurate. TSF objects are partially detected as not all the pixels are classified as part of the TSF Mask. The partial detection could be due to the model not recognising all the pixels of the TSF due to the model being trained using 448 x 448-pixel sizes. To improve this result, one could adjust the model to train on various-sized pixel grids.

Figure 8 Sample results of the tailings storage facility object detection by Mask RCNN

3.3 Model accuracy

Figure 9 shows the iterative evolution of the model accuracy on 10 re-sampling, training and validation procedures (k-fold cross validation). For the Mask RCNN, the highest precision on the training data of 67% was achieved after 85 epochs. The precision value is obtained when the model compares its predictions against a subset of the training samples, unseen by the model, and then awards a point per correctly identified object using the unseen samples. An epoch is a cycle of training, detecting and testing for the model. In each additional epoch the model learns from its errors in the previous epoch in order to improve its precision score. In each re-sampling step the Mask RCNN was validated on a separate dataset containing images.

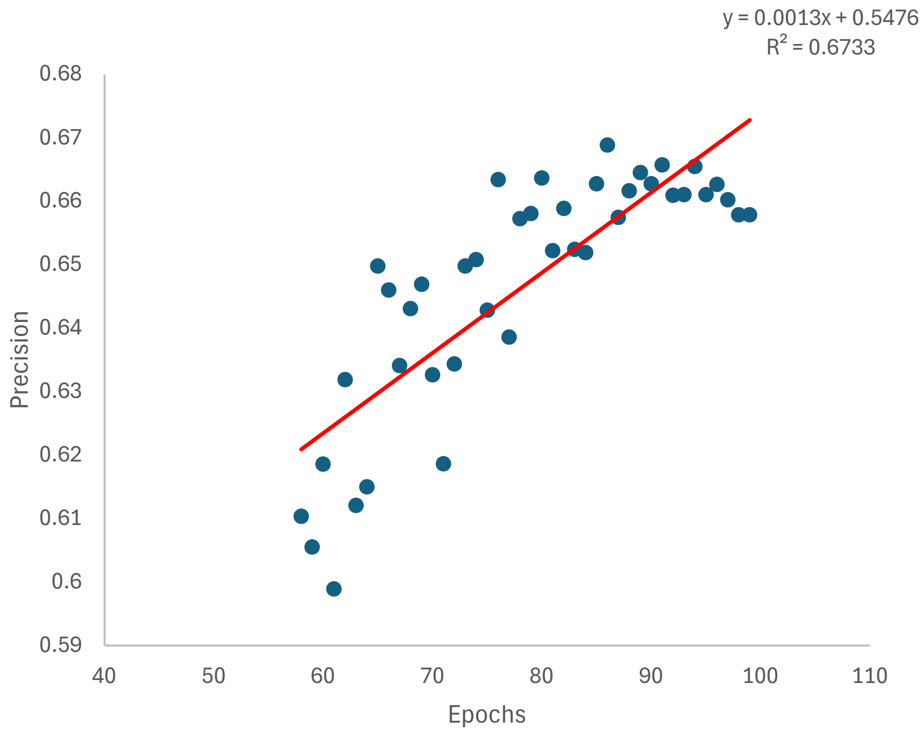

Figure 9 Evolution of the precision for tailings storage facility detection over 85 training epochs

With more epochs the model precision of the Mask RCNN increased. At about 85 epochs the precision was at 67% and this is the highest precision observed from this model. A direct relationship exists between epochs and model precision until beyond 85 epochs where a slight decline in precision is observed. The coefficient of determination R^2 =0.6733 means that 67% of the variation in y (the precision) is explained by the variation in x (the epochs).

4 Conclusion

From this study it is evident that deep learning models can be used to identify TSF sites. Both Faster and Mask RCNN could accurately identify TSFs at different accuracies. The capability of Sentinel-2 and its integration with deep learning was demonstrated in this study.

TSFs can be detected over a wide area using free low-resolution satellite imagery and automation. Although the majority of mines and TSFs are visible in images with good resolution and can be identified by a specialist without the aid of an algorithm, it is unlikely that a specialist will be able to do it in a cost-effective way in very large or unfamiliar areas. Furthermore, they will need high-resolution imagery, which usually has to be paid for or is out of date when it becomes freely available. Some mines and TSFs cannot be seen on Google Earth’s outdated imagery, but with the use of RCNN models they can be detected.

The three RGB Sentinel bands were most suitable to identify TSFs because visual patterns associated with mines and TSFs are very specific and not consistent with those found in the conventional image sets used to train the RCNNs available for transfer learning.

This study demonstrated the use of remote sensing images and deep learning in the mapping of TSFs. Although the identified TSFs were not necessarily associated with abandoned mine sites, this study has shown that TSFs can be identified by trained deep learning models. The Mask RCNN model was able to identify 46 TSFs that were not part of the training samples.

Despite the limited number of available training examples and the enormous diversity of visual patterns that TSFs present, the study demonstrated that TSFs can be identified automatically with an accuracy of over 67% using deep learning and free moderate-resolution satellite images.

Acknowledgement

We would like to thank the National Research Fund through the Black Academic Advancement Programme for funding the project. Special thanks go to the Esri South Africa team for all the training, deep learning modelling and support provided to us.

References

Balaniuk, R, Isupova, O & Reece, S, 2020, 'Mining and tailings dam detection in satellite imagery using deep learning', *Sensors*, vol. 20, no. 23, https://doi.org/10.3390/s20236936

Ball, JE, Anderson, DT & Chan, CS 2017, 'Comprehensive survey of deep learning in remote sensing: Theories, tools and challenges for the community', *Journal of Applied Remote Sensing*, vol. 11, no. 4, https://doi.org/10.48550/arXiv.1709.00308

S2 MSI ESL Team 2022, *Sentinel-2 Data Quality Report*, Report Issue 55, European Space Agency, Paris, https://sentinel.esa.int/documents/247904/4766895/OMPC.CS.DQR.002.09-2022+-+i54r0+-+MSI+L2A+DQR+October+2022.pdf

Davies, P, Aziz, FB, Newaz, T & Sher, W 2021, 'The classification of construction waste material using a deep convolutional neural network', *Automation in Construction*, vol. 122, https://doi.org/10.1016/j.autcon.2020.103481

He, K, Gkioxari, G, Dollar, P & Girshick, R 2018, *Mask RCNN*, Cornell University, Ithaca, https://doi.org/10.48550/arXiv.1703.06870

Humby, TL 2014, 'Facilitating dereliction? How the South African legal regulatory framework enables mining companies to circumvent closure duties', *Proceedings of the 9th International Conference on Mine Closure*, University of the Witwatersrand, Johannesburg, http://www.academia.edu/8621292/Facilitating_dereliction_How_the_South_African_regulatory_framework_enables_mining_companies_to_circumvent_closure_duties

Kossoff, D, Dbbin, WE, Alfredsson, M, Edwards, SJ, Macklin, MG & Hudson-Edwards, KA 2014, 'Mine tailings dams: characteristics, failure, environmental impacts and remediation', *Applied Geochemistry*, vol. 51, pp. 229–245.

Kussul, N, Lavreniuk, M, Skakun, S & Shelestov, A 2017, 'Deep learning classification of land cover and crop types using remote sensing data', *IEEE Geoscience and Remote Sensing Letters*, vol. 14, pp.778–782.

Liefferink, M 2019, *The Efficacy of South Africa's Environmental Impact Assessment (EIA) Regime: A Call for Responses to Government's Legislative and Policy Strengthen Environmental Governance and the Sustainability of our Developmental Growth Path*, presentation for Federation for a Sustainable Environment, Sandton.

Mazzia, V, Comba, L, Khaliq, A, Chiaberge, M & Gay, P 2020, 'UAV and machine learning based refinement of a satellite-driven vegetation index for precision agriculture', *Sensors*, vol. 20, no. 9, 2530.

Mpanza, M 2022, *Assessment of the Environmental and Socio-Economic Impacts of Unrehabilitated Tailings Storage Facilities: a Study of Blyvooruitzicht Gold Mine Liquidation Using Dispersal and Remote Sensing Models*, PhD thesis, University of the Witwatersrand, Johannesburg.

Ren, S, He, K Girshick, R, Zhang, X & Sun, J 2015, 'Faster R-CNN: towards real-time object detection with region proposal networks', *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 39.

Zabcic, N, Rivard, B, Ong, C & Muller, A 2014, 'Using airborne hyperspectral data to characterise the surface Ph and mineralogy of pyrite mine tailings', *International Journal of Applied Earth Observation and Geoinformation*, vol. 32, pp. 152–162.

Zhang, L, Xia G, Wu, T, Lin, L & Tai, XC 2016, 'Deep learning for remote sensing image understanding', *Journal of Sensors*, vol. 501, http://dx.doi.org/10.1155/2016/7954154