

Applying image classification to develop artificial intelligence for tailings storage facility hazard monitoring using site-based cameras

J Engels *Tailpro Consulting, Chile*

H Gonzalez *Sierra Gorda SCM, Chile*

G Aedo *Sierra Gorda SCM, Chile*

Abstract

Image classification is a process whereby the spectral information of an image, based on its digital numbers, attempts to classify individual pixels to a theme or specific object (e.g. vegetation, water, vehicles, people, etc.). The output is generally an image map or mosaic of pixels, each of which belong to a particular theme or identification to produce an independent overlay of the original image. This overlay can be used to provide a post analysis regarding changes that are occurring in a sequence of images or, for example, identify a potential hazard that can trigger an action for human intervention.

The accuracy of image classification is based on having enough information to train a model to identify the theme or object of interest. This paper presents the results of a supervised machine learning technique whereby target objects were identified and models run to train the classification algorithm to identify changes in supernatant pond size, rates of rise, detection of inflows of water to an area and presence of mobile equipment. Training images were acquired from site-based static time-lapse cameras that have been taking images since early 2017 of different areas of a tailings storage facility in the north of Chile.

Keywords: *cameras, monitoring, machine learning, artificial intelligence*

1 Introduction to image classification

Computerised image classification is a technique that, over the last decade, has advanced significantly, principally due to the increased power of graphics processors designed to run matrix multiplication at high speeds. The increased power allows machine learning and deep learning techniques to be implemented to generate large neural networks for training purposes with the goal of improving the output classification accuracies. A neural network is a computational learning algorithm, inspired by biological neural networks, that processes information using a connectionist approach, modelling complex relationships between inputs and outputs to decipher patterns in data.

For a computer processing an image, it sees each pixel as a number and thus a complete image is a matrix of code that can be processed computationally. The following greyscale image (Figure 1) shows the pixel number identification based on the pixel depth (0–255 scale).

The general goal of image classification is to be able to generate a new dataset or output that is then used to trigger a post process and provide further analysis. For the purpose of this paper, the classified outputs are related to water and mobile equipment with the intention to develop artificial intelligence to monitor autonomously areas of a tailings dam using site-based cameras.

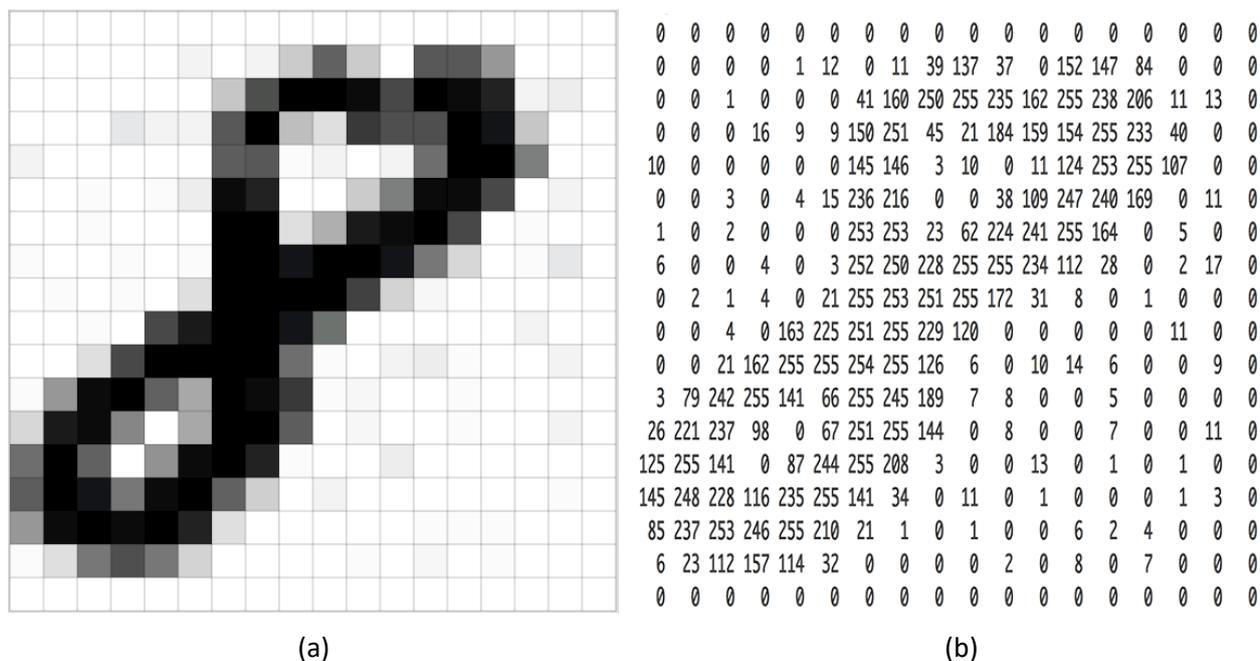


Figure 1 (a) Pixelated image of a number 8; (b) The corresponding greyscale pixel depth (Geitgey 2016)

2 Camera setup

The principal author has been developing time-lapse cameras for several years as a method of documenting tailings storage facility operation. This is predominantly to develop understanding and control of the following operational scenarios (Engels & Vega 2016):

- Pond size and migration changes as well as freeboard variations.
- Spigot discharge, channel meandering and beaching behaviour (drying and wetting, beach slope formation, flushing times, etc.).
- Document construction and daily operational activities.

The latest version of this camera has been in operation at the Sierra Gorda mine in the north of Chile since early 2017 and has been taking high resolution images of four distinct areas of the tailings storage facility every 20 minutes, 24 hours a day. The images are uploaded to a server for continual storage and a web portal provides tailings operators and senior management access to the latest three images to see an almost live picture of the areas of interest. Using this web portal, scripts were developed as part of this paper to subsequently mirror these images for processing and classification to generate additional data relating to the pond size and movement of both personnel and equipment.

Figure 2 presents the camera used for this paper and an example image of the entire view as generated by the camera prior to water inflow in the area of interest. This camera is monitoring the main decant area of a tailings storage facility where in late 2016, no further tailings were discharged into the area to allow construction of the main downstream embankment and the subsequent modifications to convert the water collection system from a decant tower to a floating barge array. The camera successfully monitored the embankment construction activities and provided a visual reference to verify that no water was accumulating against the main embankment. Prior to tailings discharge into the area in late 2018, the camera was prepared with the various scripts and server image recognition software to allow alarms to be established and tracking of the gradual increase in pond area that would be used to calibrate a pixel versus true area calculation in the future.

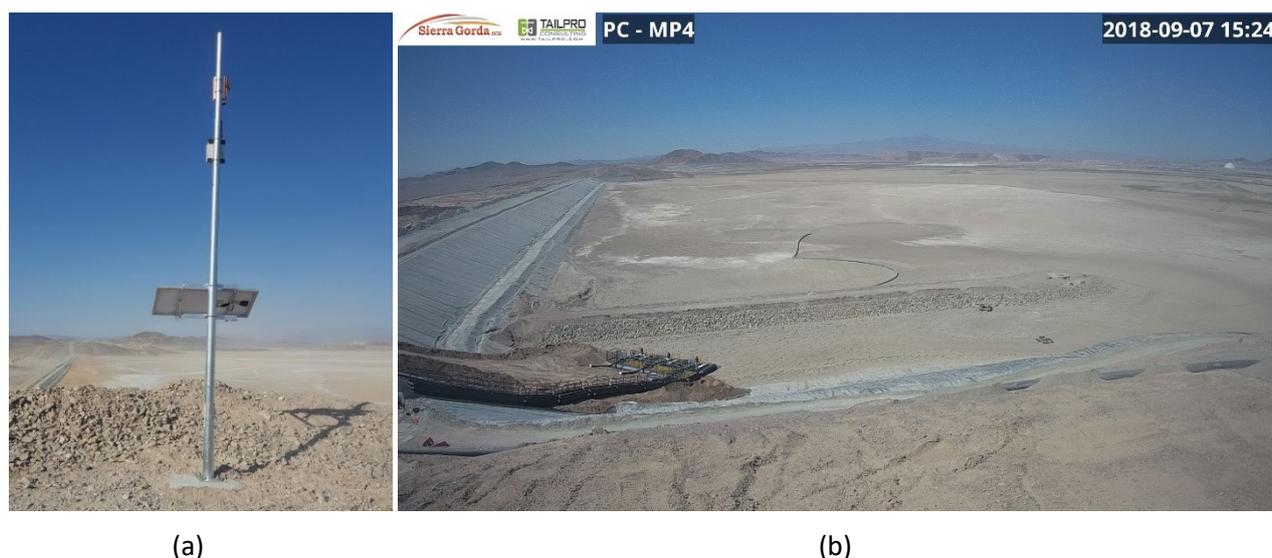


Figure 2 (a) Camera equipment installed since early 2017; (b) Camera view prior to water inflow

3 Concept, workflow and model training

3.1 Concept

The concept of image classification for these cameras was to develop a remote server back-end separate to the client front-end server. The client server would mirror newly acquired images to the remote server that would run the various algorithms and scripts to classify the image and update a percentage area graph of water present in the image. The graph is updated and hosted on the client server below the newly acquired images. The separation of the servers was essential so as not to slow down the client server and impact the operation of all cameras operated by the system. The graph was originally updated every 20 minutes during daylight hours but was then changed to only update four times between 8:00 am to 8:00 pm each day. The reasons being were to reduce server processing time and costs as an elastic server is used to process the images due to the high processing power demanded.

Since the camera is static (i.e. does not pan and tilt), each image is taken at the same position each time. An image mask was developed to draw a region of interest (ROI) in the area where water would accumulate in the images. This would increase the accuracy of the image classification and speed up the process of detection by discarding areas of the image that were not required to be analysed.

An alarm system was also integrated into the analysis to detect the initial inflow of water to the area to provide an alert via email showing the last three images together with the masked water overlay for easy comparison.

An additional graph was generated below the water area graph to identify if any earthmoving equipment and personnel were observed in the images. This was a model that was trained in parallel for an alternative embankment construction application in the future.

3.2 Workflow and model training

A relatively complex workflow for this camera was devised, as presented in Figure 3. The main intelligence behind the system is housed on an Amazon Web Server (AWS), which is the remote back-end system. This server hosts the deep learning model and mask system which was originally developed using Keras, a neural network written in Python, and then in Google TensorFlow (Google Inc. 2011) which allows supervised and unsupervised learning models to be trained and evaluated for a known dataset. The trained dataset and associated written codes were then modified to run on the Amazon servers under automated and unsupervised conditions.

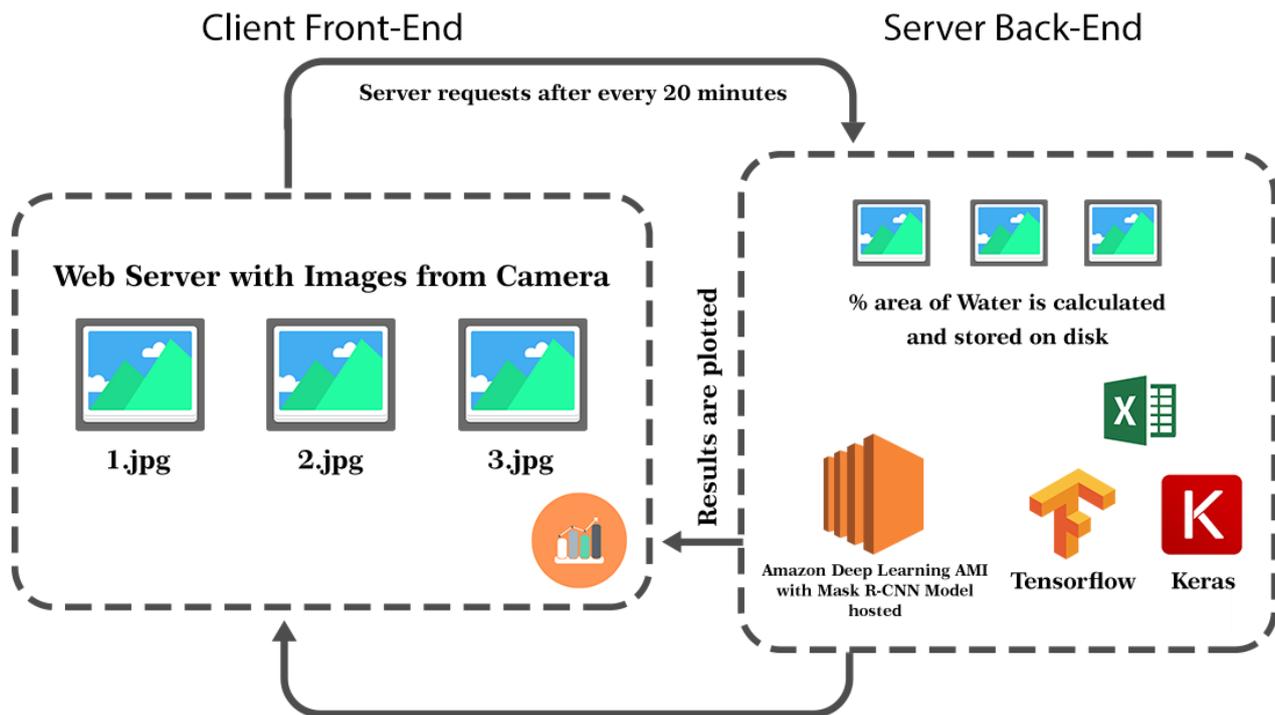


Figure 3 Workflow of image acquisition, post-processing and reporting

Convolutional neural networks (CNN) are the heart behind the image classification process to be able to identify the subject matter of interest or to be able to allow an image to be tagged, such as facial recognition of a celebrity for example. The details of CNN and the process of training go beyond the scope of this paper, however, information can easily be found via a simple web search.

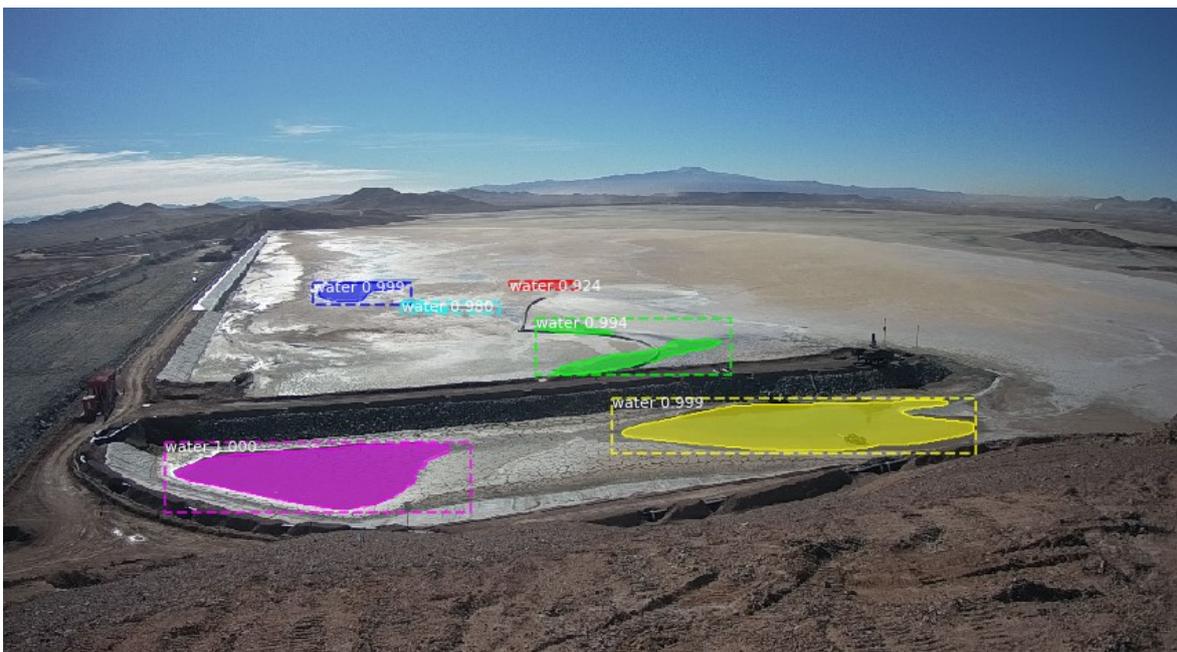
For the work presented in this paper, a masked CNN approach was considered for segmentation to allow a ROI to be applied to the area of where water is likely to occur within the camera image. This process was developed by Kaiming He, a researcher at Facebook AI and is branded Region-CNN or R-CNN (He et al. 2017). The model was trained using one month of images where a rain event occurred over the course of a few days allowing runoff to accumulate in the area of the decant pond and slowly evaporate with time. This provided a database of multiple pond locations and their sizes that would be ideal to train the model prior to the inundation of water associated with tailings deposition that would generate predominantly a single pond.

The detected water regions were used as a supervised technique together with Common Objects in Context (COCO) which is a large-scale image database of segmented objects whereby water is a known detected object that could be applied to the camera images to label the water and provide a higher degree of probability of correct detection as a starting point for the training.

Figure 4 presents a raw camera image together with the post processed water regions detected following a few iterations of training, whereby confidence scores of over 92% were determined for multiple ponds formed following the rainfall event.



(a)



(b)

Figure 4 (a) Raw camera image; (b) Initial water detection training identifying multiple ponds having over 92% probability of truth

For the equipment and personnel detection, the same workflow was considered as the water only imagery, but this time using excavators and people stances to train an independent model. Real images from the camera were then used relating to the construction of the floating barge array (Figure 2(b)) which involved excavators and personnel during mid to late 2018 which was subsequently documented in a separate graph on the same client front-end. Figure 5 presents the masked comparison showing detection of excavators (blue bounding boxes) and people (green bounding boxes).

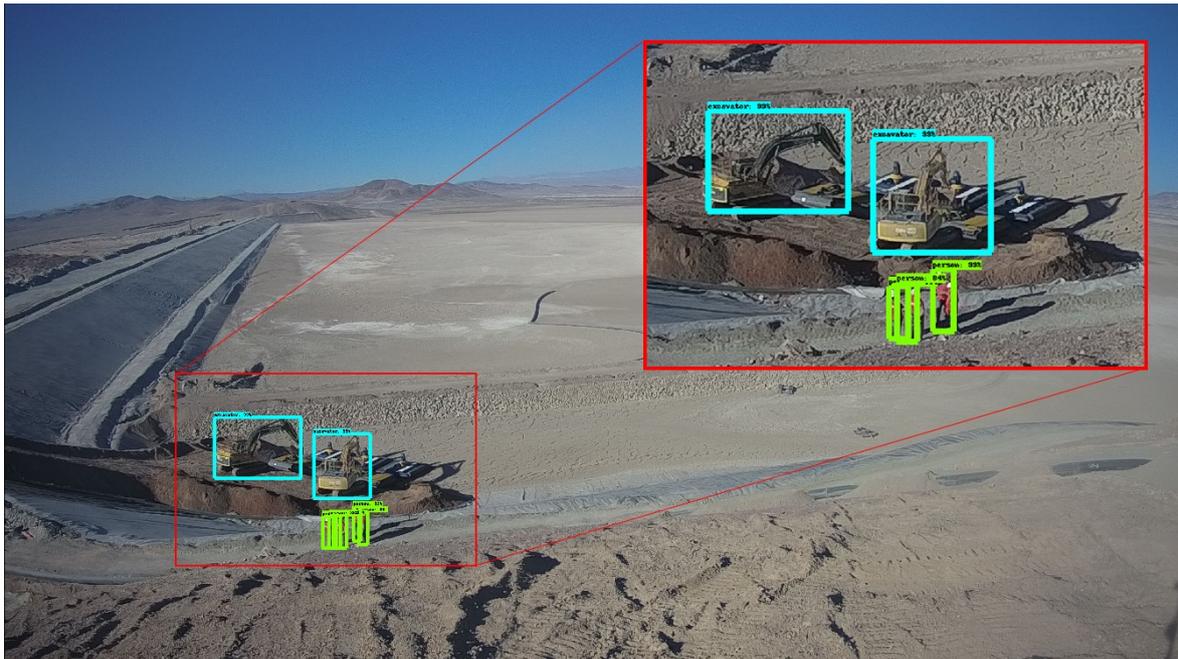


Figure 5 Equipment/personnel detection training having over 84% probability of truth

4 Classified outputs and data presentation

Following training of the model and the probabilities calculated, the algorithms were adapted and placed on the AWS system ready to receive new images from the front-end server as and when uploaded by the camera.

The following step was to calculate the percentage area of water defined by the ratio of number of pixels of detected water within the masked ROI. As new images are uploaded to the server, the output area is sequentially stored in an Excel file on the server and a graph of the outputs showing the last 30 days as a visual reference is pushed back to the client front-end server. Figure 6 presents the graph showing the detected rain events and reduction in pond area with time. The value of the y-axis is the total percentage area of water within the ROI area.

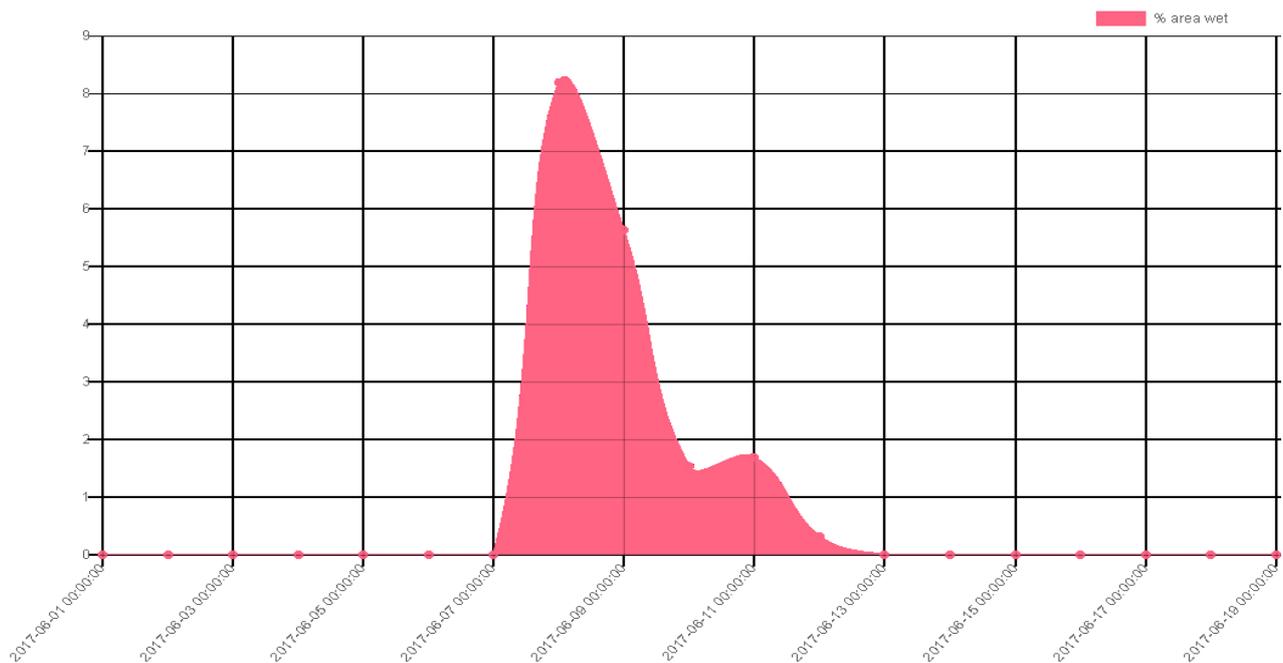


Figure 6 Detection of water ponding in the processed images

As part of the trial, an alarm trigger was used to send an email if the pond increased in size by 2% every four hours. The email contains the last three images taken together with the masked water areas for visual confirmation of detected changes. This trigger can also be set to a fixed value, for example, send an email once the pond reaches a set percentage area. As part of the trial, the alarm worked successfully, sending an email to a Gmail account which can easily be programmed to resend to multiple mailboxes as required (e.g. control room operator, mine management, operators onsite or even regulators).

The same email trigger can be applied to the equipment and personnel detection algorithms, whereby an alarm can be set for any specific condition or combination of conditions requiring tracking (e.g. two personnel and three excavators in any one image).

5 Conclusion and future work

The application of computerised image classification for tailings storage hazard monitoring has been presented in this paper. The techniques used have demonstrated the success in detecting water within an image as well as its accumulation/reduction in area with time. The training model used provided reliable predictions of water achieving over 92% probability of correctness, providing a reliable tool for future analysis and data collection from these types of cameras.

The alarm system developed also provided an automated means of increasing awareness to potential hazard occurrence, in this case when the pond area reached a certain size. This can provide an operator/owner with an external review of a known risk that does not rely on inspections by onsite personnel who may overlook minor changes which can be detected by image classification.

Following the development of this work a few limitations to this technology were determined as follows:

- Only daylight images can be processed using these techniques. Night images can be processed if illumination or high-powered infrared (IR) lamps are installed but may be impractical depending on the ROI requiring supervision.
- Rain on the camera lens and heavy dust accumulation can distort the interpretations and modifications would be required to protect the camera lens, particularly in environments of frequent rain.
- The camera position is required to be fixed, and no offsets should occur between images (e.g. windy conditions causing the camera pole to sway slightly between image capture). Although this is not critical, it will improve the accuracy slightly.

The authors' intentions are to continue to develop this technology and apply it to other cameras with the goal of providing a third review automated alarm system once defined operational parameters are surpassed. For this reason, the equipment and personnel recognition techniques were developed to assess how they could be applied to construction and maintenance operations within a tailings storage facility.

The following provides a summary of the potential improvements and new areas of development considered to take this technology to the next level:

- Improvement in the detection and classification of elements of interest under low level light conditions. The use of IR will also be investigated as a means of processing night images with the goal of providing a 24-hour monitoring solution using this technology.
- Optimisation of the workflow to reduce processing times that will ultimately increase efficiency and reduce the costs associated with operating the front-end and back-end server infrastructures.
- Calibrate the percentage area of water detected within an ROI of an image to output a real area coverage value in m². Thus, to be more practical for operators.
- Improve the alarm system to have upper and lower bound alarm triggers (e.g. upper and lower pond areas).

- Installation of stakes in the image frame to be able to try and detect rates of rise as a separate graphical output.
- Develop multiple regions of interest in a single image to detect other alarm triggers. For example, critical distances of water to the crest of an embankment whereby if water moves into this region an alarm is triggered.

Acknowledgement

The authors thank Sierra Gorda SCM for permission to publish this work together with their support relating to the implementation of these cameras onsite.

References

- Engels, J & Vega, F 2016, 'Remote sensing techniques to characterize tailings deposition behaviour to improve daily planning and provide early warning systems', in S Barrera & RJ Jewell (eds), *Proceedings of the 19th International Seminar on Paste and Thickened Tailings*, Gecamin, Santiago.
- Geitgey, A 2016, *Machine Learning is Fun! Part 3: Deep Learning and Convolutional Neural Networks*, Medium, viewed 18 August 2018, <https://medium.com/@ageitgey/machine-learning-is-fun-part-3-deep-learning-and-convolutional-neural-networks-f40359318721>
- Google Inc. 2011, *TensorFlow*, computer software, Google Inc., Menlo Park, <https://www.tensorflow.org/>
- He, K, Gkioxari, G, Dollár, P & Girshick, R 2017, 'Mask R-CONN', *Proceedings of the 2017 IEEE International Conference on Computer Vision*, IEEE Computer Society, Piscataway, pp. 2980-2988.