Fibrous Image Analysis and Counting

Mark Robson

Ajmal Mian
School of Computer Science and Software Engineering
University of Western Australia

Andrew Mehnert
Centre for Microscopy, Characterisation and Analysis
University of Western Australia

Ruairi Ward
CEED Client: OHMS Hygiene

Abstract

OHMS Hygiene requires a software tool to aid in the counting of asbestos fibres present on glass slide specimens imaged using phase contrast microscopy. The process of counting fibres is currently completed by human operators and is time-intensive and subjective. This project aims to create software that can identify and determine the dimensions of fibres in a microscope image. The software is intended to aid the operators to count faster and to more accurately measure the sizes of the fibres for the purpose of assessing fibre concentration levels. The project uses Mask R-CNN technology to find and isolate fibres from the background as well as measure their dimensions. The machine learning model was trained upon a data set of example fibres collected from the historical slides kept by OHMS Hygiene and hand annotated to mark the location of the fibres.

1. Introduction

OHMS works in many areas of occupational health management, providing its clients with a wide range of services to support their businesses. One service that OHMS provides is environmental air sampling and testing for the presence of harmful asbestos fibres. This is achieved by passing air through a filter that traps the fibres. The filter is carefully prepared to create a microscope slide that is then examined by an operator using phase contrast microscopy. The operator manually counts fibres in several regions or fields-of-view (FOVs) on the slide to estimate the number of fibres per millilitre of air. This information is used by OHMS to advise the client on appropriate actions and precautions to take.

The current manual process of examining a slide for the presence of asbestos is a time-intensive and skilled task. The development of computer software to aid the operator with this task is thus highly desirable. Such software would decrease the time taken to count fibres on a slide therefore increasing the number of air samples that could be processed in a day. Moreover, it can potentially reduce or eliminate operator subjectivity. Such software must be able to identify potential fibres from background particles, and for each candidate evaluate its length and width to determine if it meets specified criteria for a countable fibre.

2. Asbestos Monitoring Background

Asbestos fibres are closely monitored and controlled due to the long-term negative health effects that they can have. Individuals exposed to asbestos can develop a wide range of serious illnesses such as asbestosis, lung cancer, mesothelioma and pleural plaques (Doll and Peto 1985). The health effects of the fibres were not always well known, and the fibres were used widely in commercial applications. The strength, fire resistance and insulative properties of asbestos make it an ideal construction material, which was widely used until the negative health effects lead to reductions in its use during the 1980's and a complete ban in Australia by 2003 (Eller 1996; Leigh 1996). The historical use of asbestos in construction has resulted in building materials containing asbestos that can become hazardous if disturbed.

Asbestos also occurs naturally in rock formations (Doll and Peto 1985; Leigh 1996). Many valuable ores are contained in areas with asbestos-bearing rock. The mining process can disturb the fibres, and it is therefore important to monitor and control asbestos when mining. The mining and construction industries are often faced with the problem of managing and controlling the health risks associated with asbestos, and many standards and systems have been developed to aid in the safe management of asbestos.

To create effective plans and control strategies to manage asbestos it is important to estimate the concentration of airborne fibres. The membrane filter method is commonly used to estimate the concentration of fibres in an air sample (Guidance Note on the Membrane Filter Method for Estimating Airborne Asbestos Fibres 2005). The air sample is collected by an air pump that passes air through a filter trapping the fibres. This filter is turned into a microscope slide that is used to count the number of fibres in the sample. This information is combined with the time the pump was running and the flow rate of the pump to provide the estimated concentration of fibres from the air sample. The count of fibres on the filter is completed by an operator using phase contrast microscopy. The operator looks at one hundred different FOVs at 40x magnification. The FOVs are selected by the operator such that they are evenly distributed over the whole observable area of the slide. Typically, the number of fibres counted in 100 FOVs is determined, and this is taken to be representative of the whole slide. This fibre count along with the volume of air sampled is used to calculate the result in fibres per millilitre.

The concentration in fibres per millilitre is then compared to the Safe Work Australia Exposure Standard. This defines an upper limit on the airborne concentration that must not be exceeded. For asbestos fibres, the Safe Work Australia Hazardous Chemical Information System (HCIS) lists an exposure standard of 0.1 f/mL (Guidance on the Interpretation of Workplace Exposure Standards for Airborne Contaminants 2013).

3. Background Technology

This section describes the basic operation of the technologies used by this project to detect fibres. This short description only provides the fundamental concepts, and further reading is advised for a more comprehensive understanding (He et al. 2018; Liu et al. 2015; Ren et al. 2017).

3.1 Artefactual Neural Networks

Artefactual neural networks (ANN) are used in a wide range of machine learning systems. A network takes a set of inputs and produces a set of outputs. The system is built from a set of nodes which at their simplest take a set of inputs, perform an operation and produce an output. A node for example could take each of its inputs and output the weighted sum of the inputs. Nodes take inputs from either the original input of the network or from outputs of other nodes. A network of nodes is often arranged in layers or levels progressing from input to a final useful output (Yegnanarayana 2006).

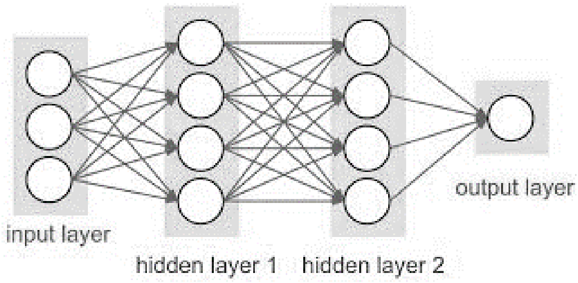

Figure 1 An example of a small network (Yegnanarayana 2006)

The nodes of the network begin with randomly assigned values and weights, and initialy give no meaningful results for a given input. The network needs to be trained to provide useful results. Training is achieved through a process of back propagation, where the network is given a set of inputs for which the correct output is known. At each level of the network moving succssively from the output layer, the nodes' weights are edited such that the output moves slightly closer to the correct output. This step is repeated over many iterations and leads to a network that can predict the output to a given input.

3.2 Convolutional Neural Networks

Convolutional neural networks are an extension of ANN where each node represents a new image built from the input layer or the previous layer. This is achieved through an operation called a convolution. Each convolution takes an image and a convolution kernel and returns a new image where each pixel's value is now a combination of its neighbours' values. The kernel is a matrix describing the weight each neighbouring pixel has towards the new value of that pixel. In the same way that the weights of each node in an ANN are trained, the weights in the kernel are trained in a CNN. Each time such a convolution operation is completed the resulting image is of a smaller size. Over many such operations the image is reduced to a set of individual values. In the classic categorisation problem each image is reduced to a single set of numbers which is typically a percentage representing the confidence that the object is of that category (for example, a fibre). This, however would only be used to tell if an image is or is not of a fibre, not where in the image the fibre is or how many fibres there are. One approach to finding the location of fibres would be to sample many smaller parts of the image and test whether that region is classified as an object. Such a solution is often called a sliding window approach. This approach does, however, require a lot of processing power, especially as it is often necessary to look at hundreds of sub-images to find smaller objects.

3.3 Regional Convolutional Neural Networks & Mask RCNN

The problems with sliding window approaches have lead to the development of several methods to provide not only more accurate detection but also to reduce the amount of processing power needed to find and classify objects in an image. One such approach is to use a region proposal network to find areas that are likely to include an object, This network however has fewer layers, and is therefore computationally less expenive. This process is explained in detail in (He et al. 2018; Ren et al. 2017). Further improvements were made in the Fast RCNN and Faster RCNN papers that allowed for images to be processed at faster rates while improving accuracy. This work developed systems which can reuse the calculations from the region proposal network in the classification network ,and therefore reduce the time taken to make accurate detections (He et al. 2018; Ren et al. 2017).

The development of Mask RCNN extends the Faster RCNN approach by adding a mask prediction sub-network to the detection system. This collection of layers provides a mask highlighting which pixels of an image belong to the object. The following examples show how the system can detect and accurately mask objects from the Common Objects in Context dataset (Lin et al. 2018). Note how the system performs detection at differing scales (He et al. 2018).

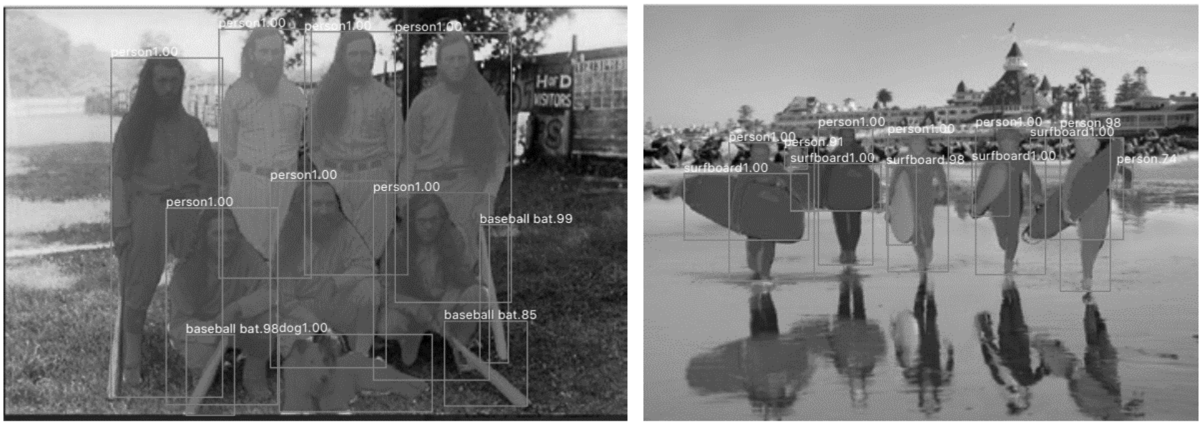

Figure 2 An Example of Mask RCNN detection (He et al. 2018)

4. Fibrous Image Analysis and Counting Project

4.1 Data Collection and Augmentation

The development of a large set of annotated data was critical to the training of a reliable classification model. The lack of any data necessitated the collection of many images from historical physical samples, whith hand annotation of the data used to highlight the locations of the fibres. A tool-assisted method was used to save and process the mask for each image.

The data annotation of a fibre included its location and a mask highlighting which areas of the image are part of the fibre. This allowed for the use of more advanced methods of classification such as Mask RCNN (He et al. 2018).

The quality of detection of the neural network is affected by the amount of training data that can be generated and annotated. Given that this is a time intensive task it was important to utilise the data that was generated as efficiently as possible. The context of the problem allows for the creation of many fibre examples from each annotated fibre in the training data. This is possible since a fibre is rotation invariant, meaning that the fibre can be placed at many different orientations without losing contextual information. Each fibre that is in the training image can be mirrored or flipped as well as rotated to generate many more fibre examples than are present in the data set. This augmentation can be used to increase the effective size of the data set to allow for more accurate results. However, care was needed to ensure that artefacts from interpolation did not result from such data augmentation.

4.2 Transfer Learning

The large amount of data that is needed to train a machine learning approach is often time consuming and difficult to create; however, many of the features that a network develops for a task are useful in many other tasks. This transfer learning can be used to reduce the amount of data needed to create an accurate network, as well as reducing the computational time taken to perform the necessary calculations. This approach was used to reduce the amount of training data needed by this project. The pretrained network ResNet-50 was used as a backbone to the developed software. The smaller network of the ResNet-50 was selected over the ResNet-101 as the context of the problem contained a relatively limited input domain, and the performance improvement will greatly improve the usability of the software.

4.3 Network Adaptions

The development of a network able to sufficiently detect fibres at different scales required many adaptations to the Mask RCNN repository. The original network was developed to detect objects as large as half the input image area at aspect ratios limited to 2:1. This is unsuitable for fibre detection, as fibres are often only very small components of an image, and the network needed to be reshaped to detect these smaller fibres with a wider range of aspect ratios (up to 4:1). The detection software was also trained to mask not only the fibre itself but also the surrounding area, This was done to include contextual information of the fibre, allowing the detection software to correctly mark fibres that have been partly covered by a particulate as a single fibre rather than two individual fibres.

5. Results

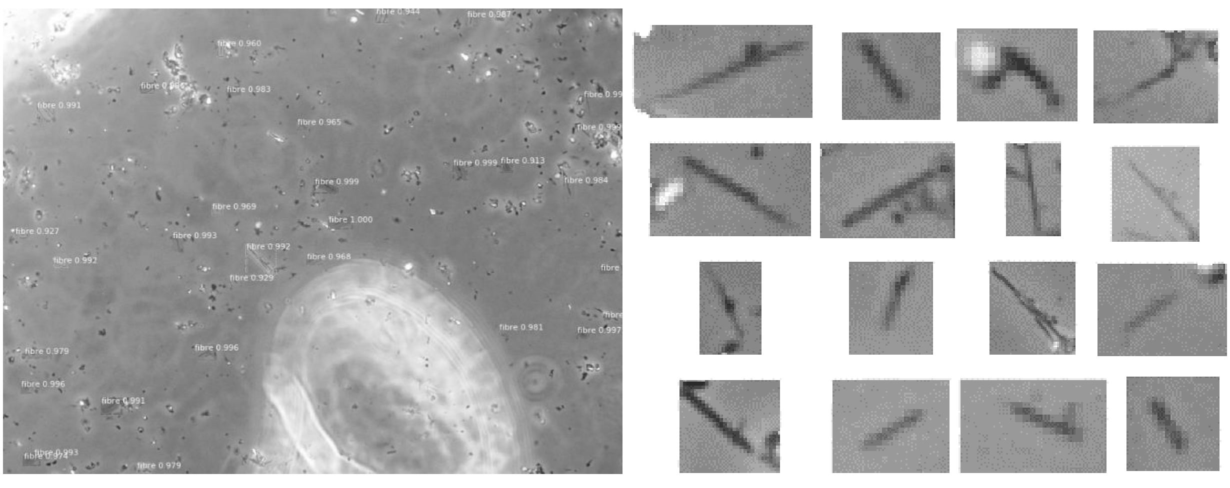

Figure 3 An example of fibre detections and enlarged fibre examples

The network that has been developed and trained takes a single microscope image and produces a set of fibre locations. The network can accurately detect the location and shape of fibres in an image. This information is further used to compute the size of the fibres. The network can accurately detect 70% of the fibres in the 50 image testing dataset images with a confidence of 90% or more. The results are also on a network that has not reached a fully trained state. The fibres which are detected are then sized to determine whether they are of an adequate size to be counted under the National Occupational Health and Safety Commission (NOHSC) Criteria and Department of Mines, Industry Regulation and Safety (DMIRS) Criteria.

6. Conclusions and Future Work

The software has reached a critical stage of development where it can accurately produce annotations for further training. The module developed has been trained on one thousand annotated images representing only one tenth of the recommended amount of data yet already shows great promise. Future work will focus on using the developed system to produce the remaining images to train and tests the network. In addition, the algorithm used to produce the dimensions of the fibres is inaccurate for curved fibres and requires additional work to provide more accurate information.

For OHMS to make the most use of the software it will additionally need to be embedded into a simplistic graphics user interface with proper documentation to highlight the software's use and internal workings. This software is designed to be integrated into a potential project aimed at automating the imaging of samples with a motorised microscope system.

7. References

Doll, R. and Peto, J. (1985). *Effects on health of exposure to asbestos*. 1st ed. London: Health and Safety Commission.

Eller, P. (1996). NIOSH manual of analytical methods, 4th ed. *Choice Reviews Online*, 33(05), pp.195-230.

Guidance note on the membrane filter method for estimating airborne asbestos fibers. (2005). 2nd ed. [Canberra, A.C.T.]: NOHSC. Safe Work Australia.

Guidance on the interpretation of workplace exposure standards for airborne contaminants. (2013). 1st ed. [Canberra, A.C.T.]: NOHSC. Safe Work Australia.

Ren, S., He, K., Girshick, R. and Sun, J. (2017). Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(6), pp.1137-1149.

He, K., Gkioxari, G., Dollar, P. and Girshick, R. (2018). Mask R-CNN. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pp.1-1.

Leigh, J. (1996). Malignant mesothelioma in Australia 1982–1995. Lung Cancer, 15(2), p.272.

Lin, T., Maire, M., Belongie, S., Bourdev, L., Girshick, R., Hays, J., Perona, P., Ramanan, D., Zitnick, C. and Doll, P. (2018). Microsoft COCO: Common Objects in Context. [online] Arxiv.org. Available at: http://arxiv.org/abs/1405.0312 [Accessed 3 Aug. 2018].

Liu, W., Anguelov, D., Erhan, D., Szegedy, C., Reed, S., Fu, C. and Berg, A. (2015). SSD: Single Shot MultiBox Detector. *CoRR*, [online] abs/1512.02325(9905), pp.21-37. Available at: https://arxiv.org/abs/1512.02325

Yegnanarayana, B. (2006). Artificial neural networks. New Delhi: Prentice-Hall of India.