

Received June 21, 2020, accepted July 9, 2020, date of publication July 28, 2020, date of current version August 18, 2020.

Digital Object Identifier 10.1109/ACCESS.2020.3012417

Deep Learning-Based Automated Forest Health Diagnosis From Aerial Images

CHIA-YEN CHIANG^{1,2}, CHLOE BARNES³, PLAMEN ANGELOV¹, (Fellow, IEEE),
AND RICHARD JIANG¹

¹Lancaster Intelligent, Robotic and Autonomous Systems (LIRA) Research Centre, Lancaster University, Lancaster LA1 4WY, U.K.

²School of Electronic Engineering and Computer Science, Queen Mary University of London, London E1 4FZ, U.K.

³Excel Geo, Northampton NN6 0BN, U.K.

Corresponding author: Richard Jiang (r.jiang2@lancaster.ac.uk)

This work was supported in part by the Engineering and Physical Sciences Research Council (EPSRC) under Grant EP/P009727/1.

ABSTRACT Global climate change has had a drastic impact on our environment. Previous study showed that pest disaster occurred from global climate change may cause a tremendous number of trees died and they inevitably became a factor of forest fire. An important portent of the forest fire is the condition of forests. Aerial image-based forest analysis can give an early detection of dead trees and living trees. In this paper, we applied a synthetic method to enlarge imagery dataset and present a new framework for automated dead tree detection from aerial images using a re-trained Mask RCNN (Mask Region-based Convolutional Neural Network) approach, with a transfer learning scheme. We apply our framework to our aerial imagery datasets, and compare eight fine-tuned models. The mean average precision score (mAP) for the best of these models reaches 54%. Following the automated detection, we are able to automatically produce and calculate number of dead tree masks to label the dead trees in an image, as an indicator of forest health that could be linked to the causal analysis of environmental changes and the predictive likelihood of forest fire.

INDEX TERMS Deep learning, tree detection, tree segmentation, forest health diagnosis, climate change, forest fire.

I. INTRODUCTION

More and more evidence indicates that global climate change has led to deforestation and extinction of species over the past decade. The incremental changes in temperature have not only caused extreme events such as floods, drought and gales, but inevitably affect landscapes and whole ecosystems. In such a situation, species may struggle to adapt to changes [1]. Moreover, it is reported [2] that 110,000 hectares of forest have died in 2018, which is around the size as 150,000 football pitches or one nineteenth of Wales, due to extreme heat and storms.

Certain species whose habitats were living woods have lost their habitats due to deforestation and dying trees. On the other hand, organism or creatures that take dead wood as habitats (for example: fungus, bacteria, predatory beetles, and parasitic wasps) increase their population. Pests such as bark beetles and cardinal beetle can nest in dead tree trunks, and then may move to healthy trees and kill them [2].

The associate editor coordinating the review of this manuscript and approving it for publication was Jiachen Yang.

In addition, the dry, inflammable nature of dead wood makes fire disasters happen more frequently [3]. In California, poor management of dead wood in forests has been identified as the major cause of wildfire. Both brought huge costs since the fire destroyed neighboring healthy trees and spread fast to villages nearby. Forest fires also trigger severe smog; for example, in 2015, wildfire smoke in Indonesia caused an estimated 100,000 premature deaths [4].

Because of the impacts on ecosystems, inhabitants and households mentioned above, technology to monitor forest areas has been developed rapidly in the past decades. Technologies such as remote airborne sensor or satellite imagery made forest monitoring less costly, and allows monitoring without human involvement. By using these technologies, forest experts and disaster specialists not only can monitor forests in real-time, but also adopt preventive measures for natural disasters.

On the other hand, pest disaster, one of the natural disasters caused by climate change, has recently drawn experts' attention. The Canopy Health Monitoring (CanHeMon) [5] project aimed to eliminate the danger to forests from pine wood nematode (PWN), and so previous research has been

done to detect dead trees within forests, to prevent them from becoming vectors for PWN. A MaxEnt-based iterative image analysis algorithm, applied to remote sensing data, as able to detect individual declining tree crowns with high accuracy [5]. The algorithm searches for the distribution with the maximum entropy for a given input set of occurrences and a set of constraining variables. As higher entropy indicates high disorder and abundance of information, the algorithm can filter out irrelevant variables, leaving the most useful variables highly correlated to dead trees. However, the results revealed that the rate of false positives increases as the sample size decreases. Moreover, eliminating small size samples (smaller than 2.5m^2) harmed the true positive detection rate, dropping from 80% to 65%. This may indicate that the algorithm is not suitable to detect small objects or that there were not enough samples of small trees in the training data.

In addition, the CanHeMon project suggested that using deep learning algorithms may provide a more generalised image analysis model and allow us to identify pixel-wise shapes of trees in an image. To investigate this, the CanHeMon team also performed tests using a neural network model. The results showed that rate of false positives was lower than for the MaxEnt-based algorithm, but the behaviour of the loss function indicated more epochs and iterations were needed in training the model. Hence, in this research, we extend the deep learning methods applied by the previous team.

In this study, we used a Mask Region with Convolutional Neural Networks (Mask-RCNN) algorithm, accompanied by transfer learning [6] to deal with the problem of pixel-wise object identification and to enhance the learning pattern, i.e. through additional training sessions, the overall score performance can be improved by using pre-trained weights produced in the previous session.

Moreover, to cope with small datasets, and problems with unlabelled data, we use synthetic tools: randomly assigning instances into backgrounds, both extracted from raw data. At the same time, we output COCO format annotations [7] to be used in inference. In general, manually annotating these masks is quite time-consuming, frustrating and commercially expensive [8]. However, we were inspired by Adam Kelly, the founder of Immersive Limit, and were able to automatically create thousands of synthetic images and their annotations with minimal manual effort. [9] The study makes four major contributions:

(1) Use synthetic data to do Mask RCNN modeling and reach about 54% on mean average precision score (mAP). The results show that it is possible to use synthetic methods to mitigate data scarcity.

(2) Mask prediction is limited by selected ROIs and can be found from high loss of RPN. Therefore, applying semantic segmentation may be a more practical way than instance segmentation to detect dead tree region. Figure VI-B

(3) Using two optimizers in training process caused loss converged faster than using one optimizer. Figure VI-B

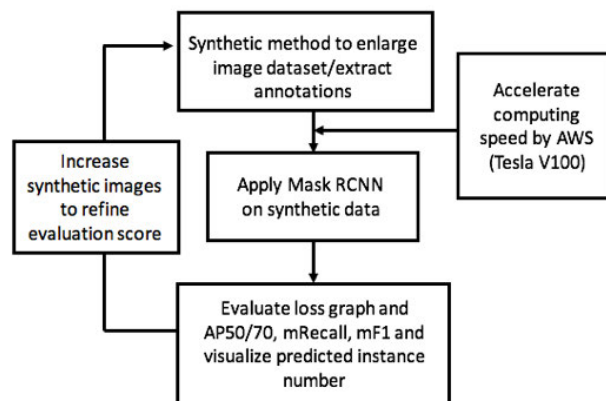


FIGURE 1. Structure of the Mask-RCNN model.

In addition, Figure 1 provides our framework for forest analysis. More details will be in the following sections.

II. RELATED WORK

A. SATELLITE IMAGING FOR FOREST ANALYSIS

Although we did not use satellite images in our research, they have been used to monitor forest cover for a long time. Analysing the change of forest cover generally relies on visual interpretation of a dataset. Accompanied with ground measurements, scientists are able to calculate the loss of forest cover over time and doing assess properties from the satellite imagery, such as biophysical and biochemical properties of forest. [10]

In addition, in various studies [11]–[13], Landsat imagery [14], a popular open database of satellite imagery, was sampled or trained to monitor forest cover changes. Landsat imagery resolution is between 15-60m with multiple bands. “Landsat is a joint effort of the U.S. Geological Survey (USGS) and the National Aeronautics and Space Administration (NASA)... The data are useful to a number of applications including forestry, agriculture, geology, regional planning, and education.” [15]

B. AI TECHNIQUES IN OBJECT DETECTION

With the remarkable advancement in artificial intelligence techniques in the past decades, there are several object detection algorithms which are now generally used, for example R-CNN [16], Fast-RCNN [17], Faster-RCNN [18], YOLO [19], SSD [20], and NAS-Net [21]. When performing object detection, we are interested in finding two things: where our target objects are located, and which categories they belong to [22]. Our goal is to draw bounding boxes around target objects in an image, and output a percentage indicating how certain we are that they belong to certain classes. Moreover, state of the art techniques allow real-time prediction of object localisation and classification, for example You Only Look Once (YOLO) and Single-Shot Detector (SSD).

However, the cost of the fast inference time of such methods may be lower accuracy and low-resolution feature maps, due to a single network evaluation [23]. They use a single stage to complete both object localisation and classification, but in contrast, region based algorithms (RCNN, Fast-RCNN, Faster-RCNN and Mask-RCNN) separate the process into two stages: first using a region proposal network (RPN) to propose a bounding box for a candidate object; then predicting the class and box offset, and outputting a binary mask for each region of interest's ROI [24]. That means we calculate both the classification scores and bounding box regression in second stage to decide which ROIs have the highest probability to contain target objects, and simultaneously perform instance segmentation on selected ROIs. Here, instance segmentation [24] means to differentiate each target object and cut them apart even when they are labeled in the same class. It is often compared to semantic segmentation, labeling on each pixel with the class of its enclosing object or region (see Figure 2). Note that we only took segmentation into account because it was more precise to cut off dead tree regions in pixel-wise manner than simply output bounding boxes; for example, some boxes may contain a big proportion of living trees.



FIGURE 2. Semantic segmentation (left) versus instance segmentation (right). Image from [24].

C. CNN AND FCN IN PIXEL-WISE DEAD FOREST MAPPING

Recent studies [25]–[28] showed CNN and FCN were used in semantic segmentation on forest mapping. In [25] a VGG neural network (an architecture of CNN) was proposed to classify pixels with a sliding window technique. One notable approach called ensemble learning was introduced to aggregate different predictions generated by the same model. This produced deterministic estimates and reduced uncertainties caused by weight initialization settings. The resulting global accuracy predicting health status of two tree species was around 80% to 85%. However, dead trees were misclassified (defined as “omission error”) at a rate of 35% to 40%.

Another study suggested sliding windows [29] cause high computational cost because neighbour pixel information, which contains the same information as centre pixels, is used to classify centre/target pixel. An algorithm such as FCN is developed to prevent duplicate computation by predicting many pixels at once. In addition, FCN can take various input sizes and output a dense prediction.

A DenseNet (an architecture of FCN) proposed in [26] has outstanding results for classifying dead trees and standing dead trees. The resulting recalls and precisions are mainly within 90 to 100% although a recall of a subset was 67.8%.

The author also encouraged that instance level prediction can be done to improve future research.

D. TRANSFER LEARNING

In the field of image processing, there have been many studies using transfer learning. The technique uses an existing trained network to improve the training accuracy for a new model.

Since an outstanding trained network such as COCO or ImageNet has been trained using many images with a big number of labels, it seems sensible to reuse fundamental parts of the pre-trained net as low-level feature extractors, i.e. to recognize edges, colours and textures - whether detecting complex and subtle objects such as a human face or simple objects such as rectangles, early layers in neural networks always detect the same basic shapes and edges.

To be more specific, the COCO dataset [7] is a large-scale object detection, segmentation, and captioning dataset. It has over 200 000 labeled images with over 80 labels, while the ImageNet dataset has tens of millions of images with thousands of labels. [30]

[6] So to train a model on a custom dataset, all we need to do is replace the classification and bounding box predictors with customised ones which are able to predict our labels. First, we take a pre-trained net without its FC part, since Fully Connected layers are used only for the old model's classification. Next, a new FC containing customized label classifiers replaces the old. Then, a convolution base which is initially close to the old classifier needs to be frozen since it contains high level features (for example cat ears, dog eyes, or a more representative feature) related to classification in the old model which may not be useful for our own customized classification problem.

E. TRANSFER LEARNING IN FOREST AERIAL IMAGERY

We used a COCO or Imagenet pre-trained net and retrained the “head” parts for our model in the first session, i.e. we kept the COCO or Imagenet backbone but retrained the RPN, classifier and mask heads of the network. [31]

In the first training session, we only trained the head since we want to leverage the advantage from COCO and ImageNet, that is, use the weights that have been trained from many images. However, we found that merely training the head was not sufficient to get higher accuracy. Because of that, in the second and third sessions, we train all layers twice. Therefore, we denote the deeper model as “head + all + all”.

Then in the second training session, we used the last weight produced in the first session to continue training. This gave us faster training times, and allowed us to use a well-trained model even when we had scarce data. Related studies on detecting objects in satellite imagery of urban landscapes [32] show that by using an ImageNet pre-trained net, accuracy was increased by around 5% for a fixed length of training time. However, the study also showed a less satisfying performance even after using transfer learning because the model is trained to predict objects from another dataset which has different domain features, which may not be similar to our own

training data. Therefore, problems such as how transferable features in deep neural networks are, and how to effectively reuse pre-trained nets when most features of the target objects are different from those of the training data, have been discussed in [33]. A Selective Learning Algorithm (SLA) has been proposed in this case, which learns useful features from unlabelled data and introduces an intermediate domain to close the gap between target objects and training data.

III. PRELIMINARY ON CNN AND MASK RCNN

A. CNN

1) WHAT IS A CONVOLUTIONAL NEURAL NETWORK?

A convolutional neuron network (convnet, or CNN) is a class of deep neural network that has proved very powerful for image classification. In 2012, Alex Krizhevsky used CNN to win the ImageNet competition which made neural nets come to prominence in the field of computer vision. [30]

2) WHY USE A CONVOLUTIONAL OPERATION?

We can view an image as having 3D inputs with (height, weight, channel) so the input features have dimension $H \times W \times C$. For deeper layers, the channel value is usually increased in order to learn more features. A channel value is also the number of filters which contain special features we want to detect. In addition, given pre-defined filters, we can obtain corresponding feature maps by applying convolutional operation on input images. The operation not only shrinks the input dimension but preserve the spatial relationship between pixels; therefore a series of features are as well preserved. Figure 3 shows different feature maps produced by applying filters to the same raw images. Usually, a typical Convnet looks like Figure 4.

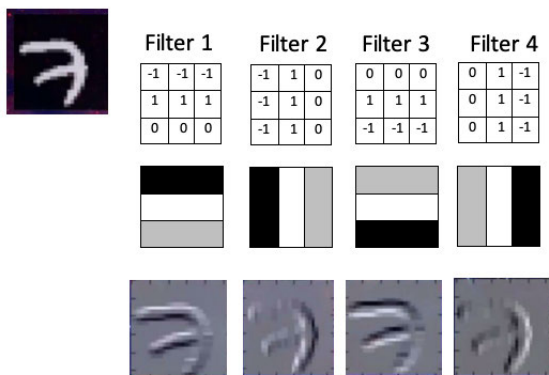


FIGURE 3. Different feature maps produced by applying various filters to the same raw image. From [34].

3) WHY CAN'T WE USE A SIMPLE CNN TO PREDICT OBJECT LOCATION?

- The number of target objects and their coordinates in the raw picture is unknown before detection, so output nodes can't be specified. In CNN, however, a fixed number of nodes is predefined for the output layer. Normally,

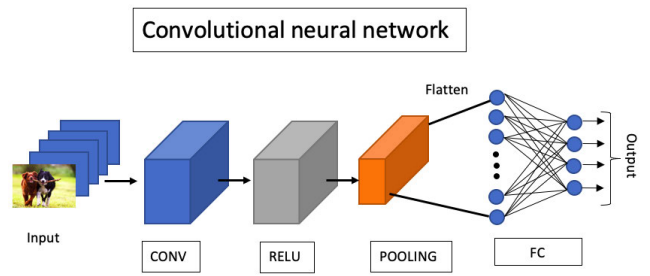


FIGURE 4. Structure of a convolutional neural network. From [35].

outputs in CNN are fixed, predefined classes. Therefore, additional layers are needed to locate their coordinates.

- To classify if an object is a target or not, a number of fixed-sized input images are fed to CNN. However, in the object detection problem, input objects may have different sizes. In this case, a ROI pooling layer in RCNN is added to scale all objects in fixed size before passing them to the classification process.

B. MASK RCNN STRUCTURE

In this section, we will first introduce Faster RCNN, the predecessor to Mask RCNN; then we will discuss Mask RCNN. This will give a broad picture of how Mask RCNN developed and the structure of the algorithm (see Figure 5).

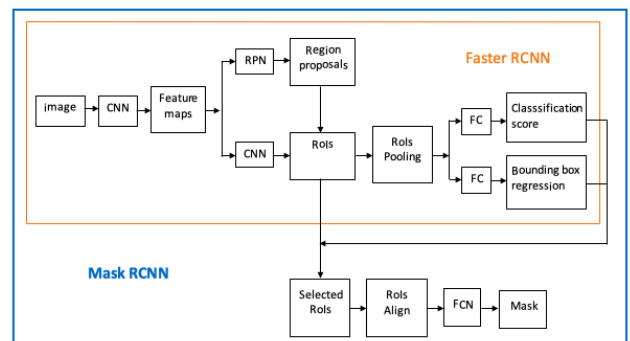


FIGURE 5. Structure of a Mask RCNN model. From [24].

1) FASTER RCNN

Faster RCNN begins with a standard convolution neural network (CNN) for feature extraction, which produces feature maps. A feature map is the output activation for a given filter in each layer. When a filter matches certain features on an image, the convolution operation [36] returns high activation values and reveal those features in the feature maps. A feature map is also a mapping of where a certain kind of feature is found in the image [34].

Then, it uses a Region Proposal Network (RPN) to extract Regions of Interest (ROIs) and resize the ROIs to a fixed size using ROI pooling. (Note: In section subsection IV-B, we will talk about how RPN generates ROIs.)

A pooling layer then reduces the dimension of the ROIs and the number of parameters in order to save training time and combat overfitting [35]. It divides an ROI into a number of sections by pre-defined output size (Note: it is not necessary that ROI is perfectly divisible by the section number.) Next, a maximum value is found for each section and we replace all other pixel values in the section with this value. [18] See Figure 6.

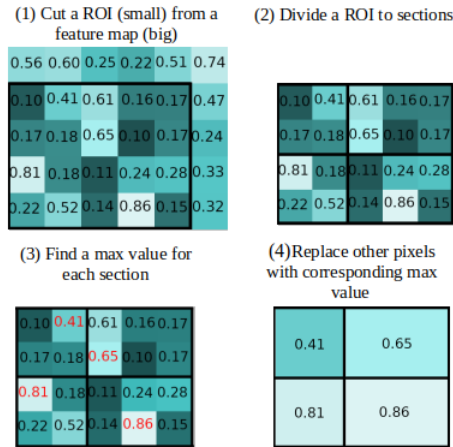


FIGURE 6. How a max pooling layer works.

Finally, the ROIs after max-pooling are flattened and fed to two fully connected layers (FCs). The flatten process makes pooled ROIs turn into vectors which store non-spatial features for later calculate the multilinear relationship between input neurons (in each ROI) and classes, and between neurons and bounding box size.

One FC layer calculates a bounding box regressor to refine the bounding box size, i.e. make each box covers the regions over its target object. Coordinates information will be sent to bounding box loss calculation. In next iteration, to minimize lost, bounding box coordinates will be updated therefore produce more accurate bounding box; and the other FC layer predicts a vector of probabilities over output classes using a softmax function. The probabilities will be used in calculation of cross-entropy loss which we adjusts probability information to minimize the cost in every iteration.

2) MASK-RCNN

Mask-RCNN can be thought of as a combination of a Faster RCNN that does object detection (class + bounding box) and a Fully Convolutional Network (FCN), a semantic segmentation method producing a pixel-wise mask. (See Figure 7)

In addition, the author of Mask RCNN [24] did some fine-tuning of Faster RCNN and FCN. First, since the ROI pooling layer in Faster RCNN can lead to the location of target objects in the feature map being misaligned from their location in the original image, this is dealt with by refining the mask with an ROI Align layer. ROI Align is a quantization-free layer, preserving exact spatial locations by replacing quantization with bilinear interpolation [24].

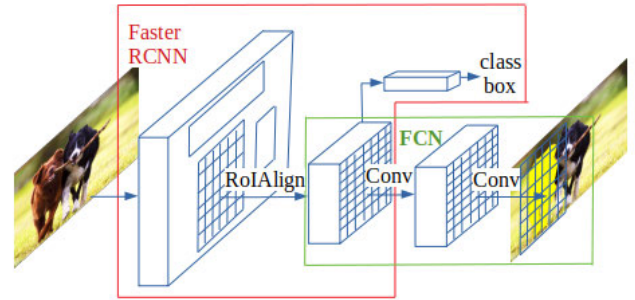


FIGURE 7. Structure of the Mask-RCNN model. Image from [link].

Bilinear interpolation is used for computing the floating-point location values in the input [33].

Then FCN, a semantic segmentation method, is applied to each ROI in order to predict a pixel-wise, binary mask (see Figure 8). Unlike the original setting of FCN, here the activation function in FCN was changed from softmax to sigmoid which was shown to significantly increase Average Precision (AP) [24]. The reason is that FCN only calculates per-pixel sigmoid values and an average binary cross-entropy loss for a single, wholly classified object each time, rather than calculating per-pixel softmax and a multinomial loss for all pixels. Because of this, label prediction is decoupled with mask generation which avoids competition among classes.

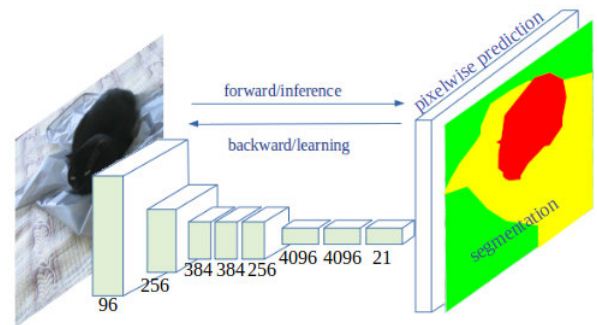


FIGURE 8. FCN predicts a pixel-wise binary mask. Image from [link].

3) BACKBONE OF MASK-RCNN

Mask-RCNN uses Resnet-FPN as its default backbone and is a type of CNN which produces feature maps; the structure of Feature pyramid networks (FPN) allows us to build high-level semantic feature maps at different scales. When images are fed into FPN, first they go into the bottom-up process: the process functions the same as classic convnet extracts features, four feature maps (C2, C3, C4, C5) are produced and get reduced size (1/2 the size of the previous map) as the level goes up. At the same time, more features are extracted on the upper level but have lower resolution because of the convolution operation process.

In addition to the loss of resolution introduced by this, high resolution feature maps from lower level are not used for

detection because of the lack of sufficient information (for example, it is possible that only some edges are detected). These two disadvantages make detection of small objects very difficult.

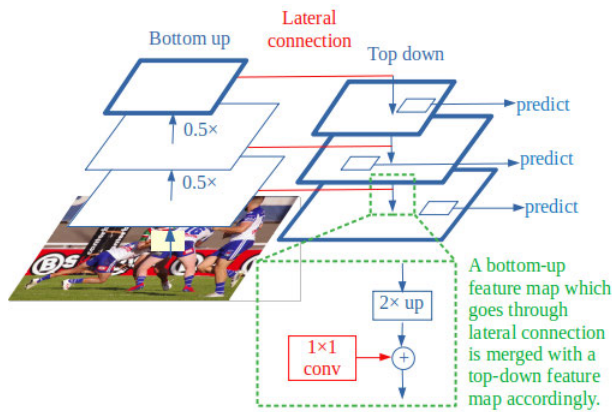


FIGURE 9. Structure of Resnet-FPN, with the bottom-up pathway on the left, the top-down pathway on the right. [37]. Image from [link].

In the FPN paper [37], however, the author added a top-down pathway (see Figure 9) to reconstruct a high resolution from a semantic rich layer by an upsampling method. The method doubles the map size (twice the size of the previous map) as the level goes down, with the resulting feature maps denoted as (P2, P3, P4, P5). A bottom-up feature map which goes through a lateral connection is then merged with a top-down feature map accordingly. The connection simply applies a 1×1 convolution filter on each bottom-up feature map; therefore, it retains map sizes but reduces channel depths. In addition, to deal with spatial aliasing (the resulting location of objects may be different from their locations in the raw image because of repeatedly downsampling and upsampling) in the merged maps, a 3×3 convolution filter is used to generate feature maps with the highest resolution and accurate spatial information; therefore we can use them to do final predictions.

IV. OUR FRAMEWORK FOR FOREST ANALYSIS

In this section, we explain how we applied Mask-RCNN to predict the location of dead trees in forests. The framework structure can be referred to Figure 1. All experiments are conducted using Tensorflow-GPU docker on a Mac Air with eight NVIDIA Tesla V100 GPUs in Amazon Web Services (AWS).

A. DATASET COLLECTION AND PRE-PROCESSING

The dataset was collected by aerial photography on May 15th 2019, from the Wood of Cree in Scotland, and was stored in a 10 GB merged tiff file. As the file was too large for much of our software to process, we divided it into 40×10 patches, each of 800×800 pixels, using the Gridsplitter plugin for the QGIS software. We then created our synthetic dataset from the raw data: we extracted over 300 images

of dead trees as foreground objects and randomly resized, rotated and changed the lightness before they were placed into 63 candidate backgrounds, also extracted from the raw data. Next we produced a mask by filling random colours over the dead trees and black for the background. We then retrieved x and y coordinates of the target objects and stored them in a COCO format annotation file. Eventually, the dataset was extended from 225 patches to 5000 patches and produced 5000 annotations in only a few hours.

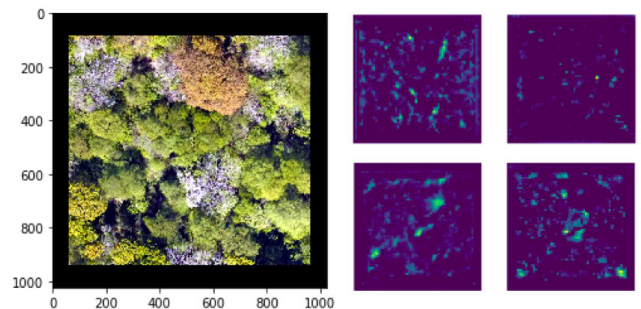


FIGURE 10. A raw image, and the four corresponding feature maps. The lighter-coloured “hot zones” indicate the locations of the features we’re interested in: shapes and colours of dead trees.

B. HOW THE BACKBONE GENERATES FEATURE MAPS

Once the dataset is fed to our convolutional backbone, it produces four feature maps for each picture in the dataset (see Figure 10).

The feature maps are then scanned by RPN in a sliding-window fashion, finding anchors (pre-defined bounding boxes with fixed height and width) of high objectness score, i.e. the probability that an anchor contains an object or a non-object. Positive anchors have high objectness scores when they contain an object.

Next, they are measured by value of 0.7 or more for the “Intersection over Union” (IoU) score. IoU is calculated as the ratio of the area of the overlap (intersection) between the pre-defined bounding box and the ground truthbox with the total area covered by both (union).

Note that because positive anchors do not always include the whole boundary of their target object, they are examined by a regressor to refine the anchor location. However, if several positive anchors overlap too much, we keep the one with the highest objectness score and discard the rest. This method is referred to as Non-max Suppression. Figure 11 shows the steps to merge the anchor boxes.

Before the region proposals are sent to the second stage, the algorithm scans and cuts out regions of interest (ROIs) from feature maps, as suggested by region proposals. ROIs are then scaled to a pre-defined, reduced dimension through the ROI pooling layer. (see Figure 6)

Simultaneously, the selected ROIs are moved to the spatial locations corresponding to the original picture in the ROI alignment layer, and then are individually sent for pixel-wise classification in the FCN layer, resulting in a prediction result as in Figure 12.

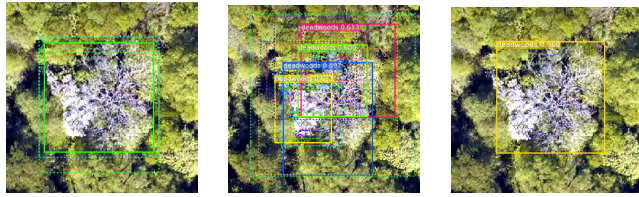


FIGURE 11. Step by step detection from finding positive anchor boxes to trim down low confidence objects.

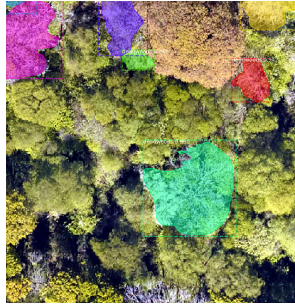


FIGURE 12. Mask-RCNN prediction result.

C. MODEL FINE-TUNING

To make our training sessions faster and to more precisely predict target objects, we made some modifications to the Mask-RCNN model.

First, we adopted a deeper model, i.e. training a ‘head + all + all’ model instead of a single-head model, and we redefined the learning rate for each training session. We set different learning rates based on optimizers, as in [38]. Our choices of parameters are summarised in Table 1. Column with bold type shows the a new, fine-tune value is assigned, compared to its last model. Models are seperated to 4 groups with more similarity within group and are allowed to be compared.

Secondly, to avoid overfitting, we used data augmentation, i.e. the algorithm accepts the original dataset, randomly transforms it and returns only the new, transformed data [39], [40]. We applied rotation by 0, 90, 180, and 270 degrees, tuned brightness and contrast between 0.9 to 1.1, applied horizontal flipping and clipped corners.

Thirdly, to make the loss function converge faster, we changed our activation function from Relu to Leaky Relu. The latter avoids the problem with “dying neurons” which Relu activation functions suffer from: when a neuron “dies”, whichever input it receives, its output will be equal to 0; in that case a neuron loses prediction power and stops learning. Leaky Relu, however, has a small but non-zero gradient for inputs below 0.

Fourthly, we applied a dropout function in the FPN layer to randomly drop out input nodes before doing classification and predicting bounding boxes. This decreases random correlations within neurons. It serves a similar purpose as data augmentation, i.e. to prevent overfitting [41].

Fifthly, we changed the optimizer from stochastic gradient descent (SGD) to Adaptive Moment Estimation (Adam) [42]. An optimizer aims to update parameters θ in the negative

TABLE 1. Comparison of settings for our fine-tuning of Mask-RCNN. Column with bold type shows the a new, fine-tune value is assigned, compared to its last model. Models are seperated to four groups with more similarity within group.

Model	Pre-trained net	Acti- vation function	Learn- ing rate	Opti- mizer	Drop- out	Data- set
Fine-tune-1	Image- Net	Relu	10^{-4} , 10^{-5} , 10^{-6}	Adam	Yes	2300
Fine-tune-2	Image- Net	Leaky Relu	10^{-4} , 10^{-5} , 10^{-6}	Adam	Yes	2300
Fine-tune-3	Image- Net	Leaky Relu	10^{-4} , 10^{-5} , 10^{-6}	Adam	Yes	3600
Fine-tune-4	Image- Net	Relu	10^{-4} , 10^{-5} , 10^{-6}	Adam	Yes	3600
Fine-tune-5	COCO	Relu	10^{-4} , 10^{-5} , 10^{-6}	SDG	Yes	3600
Fine-tune-6	COCO	Relu	10^{-4} , 10^{-5} , 10^{-6}	SDG	No	3600
Fine-tune-7	COCO	Relu	10^{-3} , 10^{-3} , 10^{-3}	SDG	No	3600
Fine-tune-8	Image- Net	Relu	10^{-4} , 10^{-5} , 10^{-6}	SDG	No	4000

gradient direction to minimize the loss. Studies show that Adam has less generalization error compared to SGD as well as being able to tolerate a noisy or sparse gradient. The algorithm is denoted

$$m_t = \beta_1 m_{t-1} + (1 - \beta_1)g_t$$

$$v_t = \beta_2 v_{t-1} + (1 - \beta_2)g_t^2$$

where m_t and v_t are the decaying averages of the past gradient and past square gradient respectively, and the default value of each of the two β s is close to 1. We then update the estimates of both to deal with the bias caused by the initial averages $v = 0$ and $m = 0$:

$$\hat{m}_t = \frac{m_t}{1 - \beta_1^t}$$

$$\hat{v}_t = \frac{v_t}{1 - \beta_2^t}$$

Because we initialize averages with zeros, the estimators are biased towards zero. However, the author of Adam optimizer used the above formulas to correct the estimators. [42]

Then, using the updated parameters, the Adam update rule is:

$$\theta_{t+1} = \theta_t - \frac{\eta}{\sqrt{\hat{v}_t} + \epsilon} \hat{m}_t,$$

where η is the learning rate, and we choose (as part of our tuning of the model) $\epsilon = 10^{-8}$.

Finally, we compare prediction results from fine-tune experiments (ex, activation function, dropout function, optimizer) to model without fine tune by visualizing loss history against epoch and a table reveals mean average precision (mAP) [43].

D. LOSS FUNCTION OF MASK-RCNN

Loss functions are used to evaluate how well a model performs [18] - in our case we aim to minimise a loss function by optimising weights for three different tasks: classification, bounding box prediction and mask prediction. We use a loss function which is the sum of a loss function for each of these tasks:

$$L = L_{\text{cls}} + L_{\text{box}} + L_{\text{mask}}$$

where L_{cls} is the same as in Faster RCNN:

$$L_{\text{cls}}(P_i, P_i^*) = -\frac{1}{N_{\text{cls}}} [P_i^* \log(P_i) + (1 - P_i^*) \log(1 - P_i)]$$

where P_i is the *predicted* probability that the i th anchor is a target object. In Matterport Mask RCNN, the total number of anchors is 256 (positive and negative). The i th anchor above “ i th” means a random number from 1 to 256.

If an anchor contains target objects, the *ground truth label* P_i^* is equal to 1, and otherwise it is equal to 0. In addition, the function is normalised by N_{cls} , i.e. the mini-batch size. A batch is the number of training samples in a given iteration. Overall, the classification loss L_{cls} is a log-loss over two classes (i.e. target object and not target object) which can be thought of as a binary classification loss (cross-entropy loss). In addition, L_{box} is also as in Faster RCNN:

$$L_{\text{box}} = \frac{\lambda}{N_{\text{cls}}} P_i R((t_i - t_i^*)_{i \in \{x, y, w, h\}})$$

where R is a robust L_1 loss:

$$R(\mathbf{t}^u - \mathbf{v}) = \sum_{i \in \{x, y, w, h\}} \text{smooth}_{L_1}(t_i^u - v_i)$$

and t_i^u for $i \in \{x, y, w, h\}$ are the position and dimension parameters of the predicted bounding box, v_i the parameters of the ground truth box, λ is a balancing parameter which is set to be around 10 in the paper (so that both losses L_{cls} and L_{box} are roughly equally weighted), and

$$\text{smooth}_{L_1}(s) = \begin{cases} 0.5s^2 & \text{if } |s| < 1, \\ |s| - 0.5 & \text{otherwise.} \end{cases}$$

We also have parameterised coordinates t_i and t_i^* ,

$$\begin{aligned} (t_i) &= (t_x, t_y, t_w, t_h) \\ (t_i^*) &= (t_x^*, t_y^*, t_w^*, t_h^*) \end{aligned}$$

where each individual parameterised coordinate is calculated given the predicted bounding box, ground truth box and anchor box coordinates as

$$\begin{aligned} t_x &= (x - x_a)/w_a, & t_y &= (y - y_a)/h_a, \\ t_w &= \log(w/w_a), & t_h &= \log(h/h_a) \end{aligned}$$

$$\begin{aligned} t_x^* &= (x^* - x_a)/w_a, & t_y^* &= (y^* - y_a)/h_a, \\ t_w^* &= \log(w^*/w_a), & t_h^* &= \log(h^*/h_a), \end{aligned}$$

where (x_a, y_a, w_a, h_a) are the coordinates of the anchor box.

Finally, L_{mask} is the average binary cross-entropy loss which is calculated when the k th mask is the associated ground truth mask:

$$L_{\text{mask}} = -\frac{1}{m^2} \sum_{1 \leq i, j \leq m} [y_{ij} \log y_{ij}^k + (1 - y_{ij}) \log(1 - y_{ij}^k)]$$

where $m \times m$ is the mask dimension for each ROI in each class, y_{ij} is the ground truth label (0 or 1) in cell (i, j) , while y_{ij}^k is the predicted probability that the same cell belongs to the k th class [18].

V. EXPERIMENTS

A. IMPLEMENT ON CLOUD SERVER WITH GPUS

A cloud server provides a virtual environment to perform computation which can be accessed remotely via internet. The major benefit is to overcome hardware limitation, for example, lack of disk space to store data or computing power to train models. In this study, in order to deal with the enormous amount of data and complicated algorithms, a cloud server(AWS) is chosen to provide the environment needed for the AI program (Mask RCNN).

Once an AWS computing instance has been registered, a private key file is given access to cloud server. Then, a Tensorflow-GPU docker file is used to activate eight GPUs as well as giving the user a temporary environment with numerous pre-installed python libraries. After the training session is done, the post-trained files are transferred back to cloud server before closing docker environment and retrieve files to local computer.

The training process takes about 4-6 hours for a dataset of 2000 to 4000 images.

B. EVALUATION OF PREDICTION PERFORMANCE

As well as the loss function, we use COCO mean average precision (COCO mAP) to evaluate prediction performance [43]. This indicator is widely used in evaluating models for object detection. Firstly, after IoU decides the ROIs, we calculate @mAP based on the prediction performance of those ROIs (ignoring objects with an IoU score lower than 50%). Then we calculate the mean precision, recall and F1. Then the AP is calculated as the area under the precision-recall curve: the so-called Area Under Curve (AUC). The curve shows the trade-off relationship between precision and recall. When the area is larger, the predicted results are more accurate (numbers of false positives and false negatives are lower) and more positive cases are detected (number of false negative is lower).

In addition, in COCO AP, the AP averaged over IoU thresholds is calculated, i.e. AP50, AP75 and AP95. In COCO mAP, an average for a 101-point interpolated AP is calculated. This means recalls are divided into 101 points and their maximized

precisions are averaged:

$$AP = \frac{1}{101} \sum_{r \in \{0,0,\dots,1,0\}} AP_r$$

$$= \frac{1}{101} \sum_{r \in \{0,0,\dots,1,0\}} P_{\text{interp}}(r).$$

To make an even more direct indicator, we also calculate the ROI number of two types of target objects (Figure 13) which will be revealed in the experimental results.

C. EXPERIMENTAL RESULTS AND ANALYSIS

Because CNN can easily suffer from overfitting, we started training our model with a dataset of 2000 images and a validation set of 300 images, to which we applied many transformations. We then looked to increase the size of the dataset gradually to find the training set which gives us the most general model. We trained each model with 100 epochs with different hyperparameter settings.

For the first pair of model, Fine-tune-1 and Fine-tune-2, the results showed that for the two types of targets (types A and B, see Figure 13) present in the data, ROI detection did not perform well for predicting type B objects (round shape) in the model “Fine-tune-1”. However, the results also showed the opposite trend in model “Fine-tune-2”, i.e. type B objects were detected reliably but few type A objects were. To resolve this, we decided to include more data in the next trial.



FIGURE 13. The model “Fine-tune-1” can reveal the type A pattern (left) but not type B (right), and “Fine-tune-2” can reveal the type B pattern but not type A.

For the next models, we reveal the evaluation statistics in Table 2. Note that we make an abbreviation of “fine-tune” to “FT”, for example, “FT3” is the model “Fine-tune-3”. In addition, we indicate the best three performances for each measurement in bold:

The table shows that results of FT6 are the best among all of the fine-tuned models. In addition, in Figure 14 we plotted the loss results to see if FT6 had the lowest training and validation loss.

Because FT7 has the lowest training loss and FT6 has lowest validation loss in Figure 14, we compared FT6 with FT7.

FT7 had a constant higher learning rate of 10^{-3} which is the closest of all our models to Matterport’s Mask RCNN default settings [31], while FT6 followed the fast-medium-slow learning rate of 10^{-4} - 10^{-5} - 10^{-6} . The results

TABLE 2. Comparison of the performance of the models. The best three performances for each measurement are in bold. The final two rows show bounding box numbers for the two types of target objects.

Model	FT3	FT4	FT5	FT6	FT7	FT8
AP50	0.39	0.57	0.43	0.54	0.60	0.38
AP75	0.14	0.53	0.13	0.48	0.56	0.10
mPrecision	0.62	0.78	0.62	0.76	0.82	0.74
mRecall	0.35	0.39	0.35	0.40	0.41	0.34
mF1	0.29	0.32	0.28	0.33	0.34	0.28
Type A (star)	40	59	33	69	32	53
Type B (round)	71	24	64	74	24	64

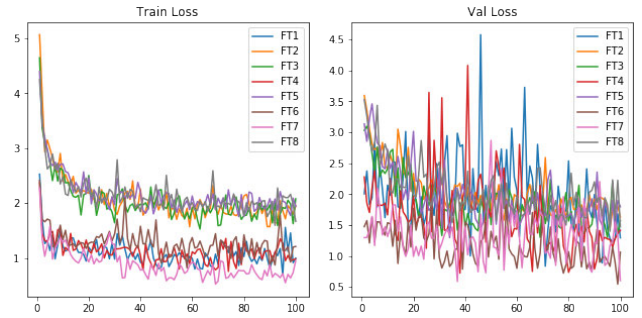


FIGURE 14. Training and validation loss for all models.

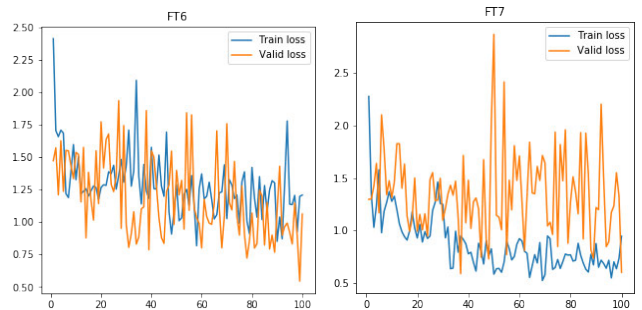


FIGURE 15. Total loss for the models FT6 and FT7.

in Figure 15 show that FT7 suffered from both fluctuating loss and overfitting. Overfitting happens when the training loss is less than the validation loss which means the algorithm learns local features appearing in the training set but not the validation set. This indicates that a fast and fixed learning rate in FT7 has no advantage over a learning rate which follows a convex pattern.

In addition, there were two major trends in Figure 14: a big shift in the training loss, and two peaks in the validation loss. Because all models had different settings, especially in the different groups, we assumed that the enlarged data set helped decrease the training loss. In addition, for the validation loss, two peaks were made by FT1 and FT4, and so we believe that this may be because of their model setting (Relu + Adam + Dropout).

The distributions of precision, recall and F1 from the model FT6 are shown below in Figure 16. The diagrams show that FT6 has very high precision but a rather low recall. This means number of false negatives is higher than the number of false positive.

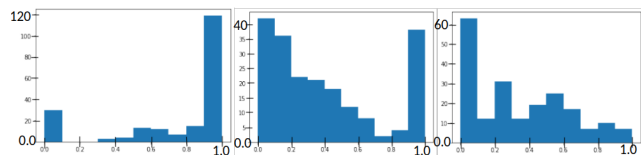


FIGURE 16. Distributions of precision, recall and F1 score for FT6 (respectively). The x-axis counts the number of predicted bounding boxes while the y-axis records the precision, recall and F1 scores respectively.

VI. DISCUSSION

A. ADVANTAGES AND DISADVANTAGES

Advantages of our method

1) Efficiently produces images and annotations:

The synthetic method was designed by Adam Kelly to avoid a time-consuming image collection and annotation process. In Adam's experiment [9], he created 2000 images containing cigarette butts and their annotations in 20 minutes from only 25 samples. Because we had small numbers of data and no longer collected dead tree images when we started to train our algorithm, we chose the synthetic method rather than trying to find more remote sensor data.

2) Less complex to apply compared to MaxEnt algorithm in CanHeMon project:

According to the image pre-processing section of the CanHeMon project, there are 14 potential predictor variables related to texture layer, containing significant image information. Then, after a series of background and AUC setting, the MaxEnt algorithm and principal component analysis are applied to reduce the number of superfluous variables. However, when applying CNN based algorithm, the variables used as inputs and outputs are rather simple: raw images and their labels. There is no manual elimination process for the variables/features but the algorithm automatically adjusts feature weights after backpropagation in each iteration.

3) False positive cases are rare:

The CanHeMon project suffered from a high rate of false positives, which does not occur in our Mask RCNN-based model. In the CanHeMon experiment, researchers manually removed false positive cases and fed the data back to the training session. However, the mask results from our model rarely showed false positive cases, although there were still many areas of dead trees which were not detected.

We found that a decrease in the IoU threshold can help detect more dead trees, and therefore decrease the number of false negatives.

Changing a hyperparameter in this way requires less manual effort than manually dropping FP cases.

Disadvantages of our method

1) Data scarcity:

We gathered all of the images of dead trees we could from both online resources and our partner company 2Excel. However, we did not generate enough data to overcome

the problems caused by data scarcity, as dead trees are an uncommon object, whereas most image datasets collect everyday objects (cars, people, etc.). The main problem was that there were several shapes of dead tree in our raw data (see the two types in Figure 13) but we could not generate a sufficient number of images to deal with this variety.

2) The process of producing COCO format annotations did not give clear boundaries in synthetic images:

The synthetic method we used sometimes led to target objects being divided into more than one piece by another object overlapping them. For example, see the red region in Figure 17. We had to manually drop the label when this happened, which is time consuming, reducing the benefit in time savings we gained from using the synthetic method. The overlapping type B dead trees give us inaccurate bounding box predictions.

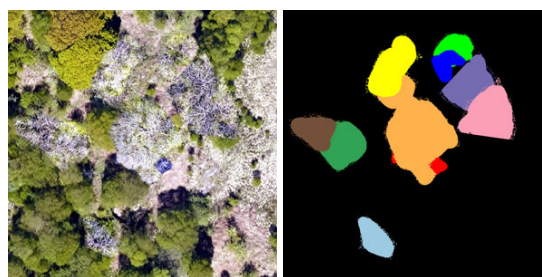


FIGURE 17. A target object in this synthetic image is divided in two by another object placed on top of it.

3) Predicted bounding boxes failing to accurately cover a whole region of dead wood

Our synthetic method can make target objects overlap, and if they belong to type A (star shaped) then the algorithm will fail to distinguish the objects from each other, as in Figure 18. At best, they are bound to a larger group and given a single label.

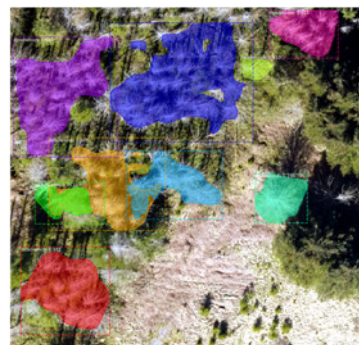


FIGURE 18. Several target objects of type A being assigned to a single bounding box.

4) Reliability of mAP and other statistics

The mAP scores in our results were still not as reliable as the ones produced and evaluated using a raw (not synthetic) validation set. The synthetic validation set was

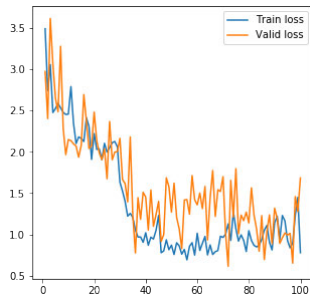


FIGURE 19. The loss functions obtained by using an SGD optimiser until epoch 30, and then using ADAM.

necessary because of the scarce data, but it would be ideal to use only raw images for validation if possible.

B. SUGGESTIONS FOR FUTURE RESEARCH

In this research, we only trained our model with the company’s aerial dataset instead of the Landsat dataset; however, we welcome to reseachers apply Mask RCNN to the Landset dataset to make comparison with our results. In addition, we did not take into account other variables such as the forest phenology in our model, because the aerial dataset was collected in summer time (May 15th 2019) when living trees are easily distinguished from dead trees. However, we cannot exclude the possibility that Mask RCNN detects and labels as “dead” rare species of trees who naturally have white leaves. Because of that, we encourage researchers to apply other possible features in their models rather than only relying on an image processing algorithm.

Secondly, to mitigate the disadvantage of data scarcity, we enlarged our dataset and produced labels automatically using a synthetic method. However, we pointed out some issues with this method in ??, although recent updates to the code we were using have dealt with some of the issues we saw in item 2 of the disadvantages in subsection VI-A. As Adam Kelly suggested [9], it is still necessary to evaluate the advantages and disadvantages of using synthetic data to train a model.

The issue with bounding box inaccuracy is the major contributor to the high loss in RPN. A possible solution is using a single FCN to classify each pixel if the instance-level segmentation is not necessary. Then, we could change the loss function from bbox + mask to FCN loss, decreasing the probability of high loss and enabling a more flexible and accurate prediction because mask prediction is not limited by the selected ROIs’ areas.

One other notable issue we experienced was how to increase the reliability of statistics (accuracy, mAP, mRecall and mF1) which were produced using the synthetic validations set to be consistent with the model’s predictive power on raw images. However, if we were able to collect a sufficiently large, manually annotated dataset, the statistics would be much more reliable. In addition, if the synthetic training set is replaced by manually-annotated raw data, it could increase

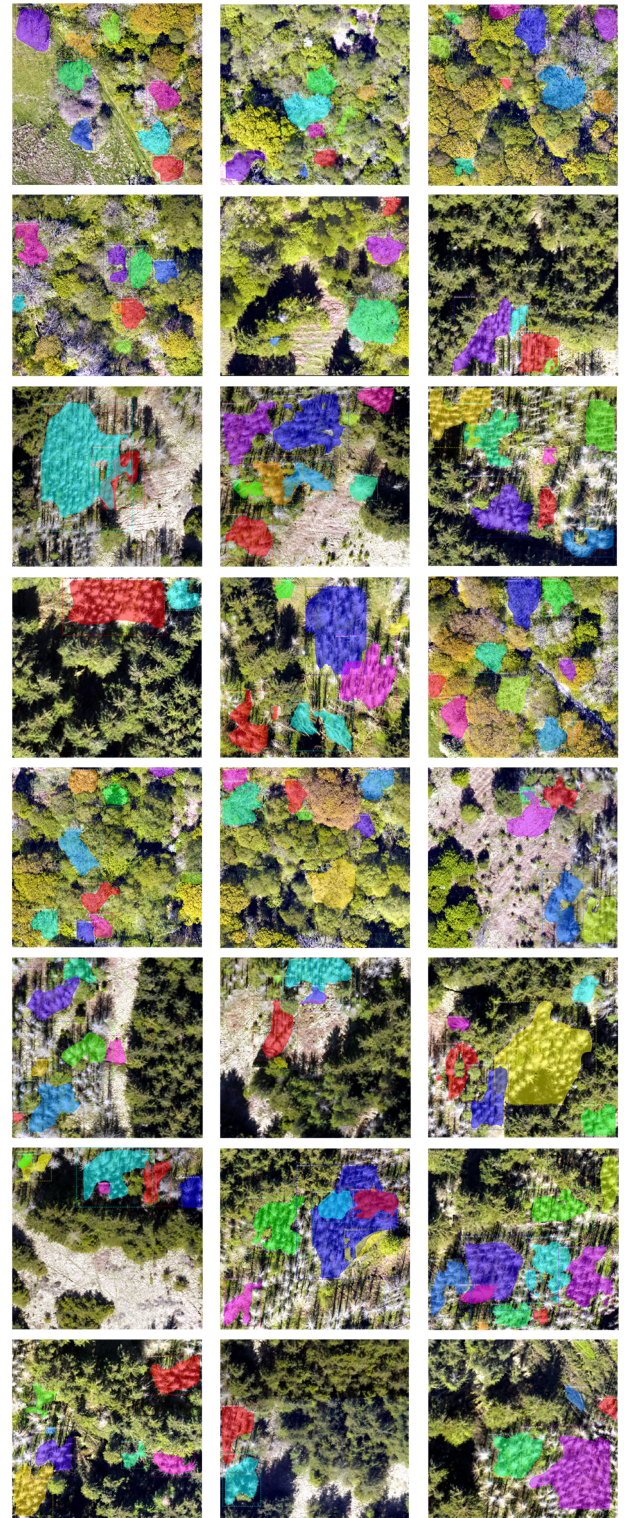


FIGURE 20. More test results on aerial images.

the predictive power since the synthetic data may not include features which are present in the raw data.

In the conventional training process, a single optimizer is used in a model. However, to take advantage of the merits of different optimizers, the author in [44] used ADAM first, and when the loss stopped decreasing, replaced ADAM with

SGD. The results obtained by doing this were better than those obtained using either optimizer on its own. After reading this, we performed the same experiment with our model. As Figure 19 shows, replacing SGD with ADAM at the 30th epoch led to a sudden drop in the loss function. As it is hard to compare directly with the other models we tested, we did not include this model in Table 1, but the results of our experiment are encouraging, so we would suggest that future research should look into using multiple optimizers for training.

Figure 20 shows more of our test results. Mostly the detection framework works well on these images while the training process was carried out on the same dataset. However, there is yet a gap between the expected detection accuracy and the performance of our deep learning framework, and further improvement is needed to extend our framework to practical applications.

The limits can be ascribed to two aspects. First, deep neural networks has been reported not robust to noises and other similar perturbations (adversary attacks) [45], while DNNs are somehow lack of explicit explainability unlike typical statistic approaches that model the data space explicitly [46]–[49]. Secondly, DNNs are data-driven and heavily rely on training datasets [50]–[53]. In our experiments, we have a limited amount of data samples. A typical remedy to this issue can be based on the neural network models of Generative Adversarial Networks (GANs) [54], which can generate more similar images similar to real ones to improve the training process.

Targeting at dead wood detection, our work is based on the assumption that the dead tree dataset collected in the spring-summer period have different colors in leaves. However, this may vary according to tree species and some rare species of trees may naturally have white leaves during four seasons. A future work can include forest phenology and classify the species of trees before the dead wood detection.

VII. CONCLUSION

In conclusion, we present a new framework for automated dead tree detection from aerial images using a re-trained Mask RCNN approach, with a transfer learning scheme. We apply our framework to our aerial imagery datasets, and compare eight fine-tuned models. The mean average precision score (mAP) for the best of these models reaches 54%. We also are able to automatically produce mask visualizations to label the dead trees in an image, so that the number of dead trees in a region can be automatically counted. Such aerial image-based forest analysis can give an early diagnosis of the conditions of forests and identify potential risk factors, and provide an potential solution to avoid of climate change caused disasters such as the recent forest fires in the California and Australia.

REFERENCES

- [1] *Ancient Tree Guide 5: Trees and Climate Change*, Woodland Trust, Grantham, U.K., 2008.
- [2] S. Florence. *Half a Billion Euros to Flow Into Germany's Forests*. euractiv.com. Accessed: Sep. 18, 2019. [Online]. Available: <https://www.euractiv.com/section/climate-environment/news/half-a-billion-euros-to-flow-into-germanys-forests/>
- [3] A. Mike. *Self-Induced Disaster: California Fires the Direct Result of Short Sighted Environmental Policies That Prohibit Forest Management*. www.ecology.news. Accessed: Sep. 18, 2019. [Online]. Available: <https://www.ecology.news/2018-11-10-self-induced-disaster-california-fires-the-direct-result-of-short-sighted-environmentalist-policies.html>
- [4] S. Joe and M. Bruno. (Aug. 3, 2019). *This Summer Uncontrolled Fires Have Raged Across the UK and Europe as Never Before. What, or Who, Sparked Them? And Might We Discover the Danger is Far From Over When the Flames are Doused?* www.telegraph.co.uk. Accessed: Aug. 3, 2019. [Online]. Available: <https://www.telegraph.co.uk/news/fire-storm-the-wildfires-sweeping-europe-and-britain/>
- [5] P. S. Beck et al., "The canopy health monitoring (canhemon) project," Ispra, Rome, Italy, Tech. Rep. JRC115400, 2019.
- [6] L. Torrey and J. Shavlik, "Transfer learning," in *Handbook of Research on Machine Learning Applications and Trends: Algorithms, Methods, and Techniques*. Hershey, PA, USA: IGI Global, 2010, pp. 242–264.
- [7] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, "Microsoft COCO: Common objects in context," in *Proc. Eur. Conf. Comput. Vis.* New York, NY, USA: Springer, 2014, pp. 740–755.
- [8] Y. Guo, Y. Liu, T. Georgiou, and M. S. Lew, "A review of semantic segmentation using deep neural networks," *Int. J. Multimedia Inf. Retr.*, vol. 7, no. 2, pp. 87–93, Jun. 2018.
- [9] *Complete Guide to Creating COCO Datasets*. Accessed: Jun. 5, 2019. [Online]. Available: <https://github.com/ak12twelve/cocosynth>
- [10] D. S. Boyd and F. M. Danson, "Satellite remote sensing of forest resources: Three decades of research development," *Prog. Phys. Geography, Earth Environ.*, vol. 29, no. 1, pp. 1–26, Mar. 2005.
- [11] K. S. Yankovich, E. P. Yankovich, and N. V. Baranovskiy, "Classification of vegetation to estimate forest fire danger using Landsat 8 images: Case study," *Math. Problems Eng.*, vol. 2019, pp. 1–14, Mar. 2019.
- [12] J. Miettinen, H.-J. Stibig, and F. Achard, "Remote sensing of forest degradation in Southeast Asia—Aiming for a regional view through 5–30 m Satellite data," *Global Ecol. Conservation*, vol. 2, pp. 24–36, Dec. 2014.
- [13] P. Borrelli, L. A. S. Rondón, and B. Schütt, "The use of Landsat imagery to assess large-scale forest cover changes in space and time, minimizing false-positive changes," *Appl. Geography*, vol. 41, pp. 147–157, Jul. 2013.
- [14] *Landsat Data Access*. Accessed: Aug. 18, 2019. [Online]. Available: <https://www.usgs.gov/land-resources/nli/landsat>
- [15] U.S. Geological Survey (USGS). (2016). *Landsat—Earth Observation Satellites*. Accessed: Aug. 18, 2019. [Online]. Available: <https://pubs.er.usgs.gov/publication/fs20153081>
- [16] R. Girshick, J. Donahue, T. Darrell, and J. Malik, "Rich feature hierarchies for accurate object detection and semantic segmentation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2014, pp. 580–587.
- [17] R. Girshick, "Fast R-CNN," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*. Washington, DC, USA: IEEE Computer Society, Dec. 2015, pp. 1440–1448, doi: 10.1109/ICCV.2015.169.
- [18] S. Ren, K. He, R. Girshick, and J. Sun, "Faster R-CNN: Towards real-time object detection with region proposal networks," in *Proc. Adv. Neural Inf. Process. Syst.*, 2015, pp. 91–99.
- [19] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You only look once: Unified, real-time object detection," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2016, pp. 779–788.
- [20] W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. Reed, C.-Y. Fu, and A. C. Berg, "SSD: Single shot multibox detector," in *Proc. Eur. Conf. Comput. Vis.*, 2016, pp. 21–37.
- [21] B. Zoph, V. Vasudevan, J. Shlens, and Q. Le, "Learning transferable architectures for scalable image recognition," Tech. Rep., Jun. 2018, pp. 8697–8710.
- [22] Z.-Q. Zhao, P. Zheng, S.-T. Xu, and X. Wu, "Object detection with deep learning: A review," *IEEE Trans. Neural Netw. Learn. Syst.*, pp. 1–21, Jan. 2019.
- [23] J. Huang, V. Rathod, C. Sun, M. Zhu, A. Korattikara, A. Fathi, I. Fischer, Z. Wojna, Y. Song, S. Guadarrama, and K. Murphy, "Speed/accuracy trade-offs for modern convolutional object detectors," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2017, pp. 3296–3297.
- [24] K. He, G. Gkioxari, P. Dollár, and R. Girshick, "Mask R-CNN," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, 2017, pp. 2980–2988.
- [25] J.-D. Sylvain, G. Drolet, and N. Brown, "Mapping dead forest cover using a deep convolutional neural network and digital aerial photography," *ISPRS J. Photogramm. Remote Sens.*, vol. 156, pp. 14–26, Oct. 2019.
- [26] S. Jiang, W. Yao, and M. Heurich, "Dead wood detection based on semantic segmentation of VHR aerial CIR imagery using optimized fcncdenset," *ISPRS-Int. Arch. Photogramm., Remote Sens. Spatial Inf. Sci.*, vols. XLII–2/W16, pp. 127–133, Sep. 2019.

- [27] T. Kattenborn, J. Eichel, and F. E. Fassnacht, "Convolutional neural networks enable efficient, accurate and fine-grained segmentation of plant species and communities from high-resolution UAV imagery," *Sci. Rep.*, vol. 9, no. 1, pp. 1–9, Dec. 2019.
- [28] A. Mazza, F. Sica, P. Rizzoli, and G. Scarpa, "TanDEM-X forest mapping using convolutional neural networks," *Remote Sens.*, vol. 11, no. 24, p. 2980, Dec. 2019.
- [29] Y. Wei and L. Tao, "Efficient histogram-based sliding window," in *Proc. IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit.*, 2010, pp. 3003–3010.
- [30] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Proc. Adv. Neural Inf. Process. Syst.*, 2012, pp. 1097–1105.
- [31] *Matterport Mask R-CNN*. Accessed: May 5, 2019. [Online]. Available: https://github.com/matterport/Mask_RCNN
- [32] A. Albert, J. Kaur, and M. C. Gonzalez, "Using convolutional networks and satellite imagery to identify patterns in urban environments at a large scale," in *Proc. 23rd ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, Aug. 2017, pp. 1357–1366.
- [33] B. Tan, Y. Zhang, S. J. Pan, and Q. Yang, "Distant domain transfer learning," in *Proc. 31st AAAI Conf. Artif. Intell.*, 2017, pp. 1–7.
- [34] R. Yamashita, M. Nishio, R. K. G. Do, and K. Togashi, "Convolutional neural networks: An overview and application in radiology," *Insights Imag.*, vol. 9, no. 4, pp. 611–629, Aug. 2018.
- [35] S. Dey, *Hands-on Image Processing With Python: Expert Techniques for Advanced Image Analysis and Effective Interpretation of Image Data*. Birmingham, U.K.: Packt Publishing, 2018.
- [36] V. Dumoulin and F. Visin, "A guide to convolution arithmetic for deep learning," 2016, *arXiv:1603.07285*. [Online]. Available: <http://arxiv.org/abs/1603.07285>
- [37] T.-Y. Lin, P. Dollar, R. Girshick, K. He, B. Hariharan, and S. Belongie, "Feature pyramid networks for object detection," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 2117–2125.
- [38] Y. Wu, L. Liu, J. Bae, K.-H. Chow, A. Iyengar, C. Pu, W. Wei, L. Yu, and Q. Zhang, "Demystifying learning rate policies for high accuracy training of deep neural networks," 2019, *arXiv:1908.06477*. [Online]. Available: <http://arxiv.org/abs/1908.06477>
- [39] *Image Preprocessing: Keras Image Data Generator Class*. Accessed: May 3, 2019. [Online]. Available: <https://keras.io/preprocessing/image/>
- [40] K. Li, G. Cheng, S. Bu, and X. You, "Rotation-insensitive and context-augmented object detection in remote sensing images," *IEEE Trans. Geosci. Remote Sens.*, vol. 56, no. 4, pp. 2337–2348, Apr. 2018.
- [41] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: A simple way to prevent neural networks from overfitting," *J. Mach. Learn. Res.*, vol. 15, no. 1, pp. 1929–1958, 2014.
- [42] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," 2014, *arXiv:1412.6980*. [Online]. Available: <http://arxiv.org/abs/1412.6980>
- [43] K. Oksuz, B. C. Cam, E. Akbas, and S. Kalkan, "Localization recall precision (LRP): A new performance metric for object detection," in *Proc. Eur. Conf. Comput. Vis. (ECCV)*, Sep. 2018, pp. 504–519.
- [44] N. S. Keskar and R. Socher, "Improving generalization performance by switching from Adam to SGD," 2017, *arXiv:1712.07628*. [Online]. Available: <http://arxiv.org/abs/1712.07628>
- [45] R. Jiang and D. Crookes, "Shallow unorganized neural networks using smart neuron model for visual perception," *IEEE Access*, vol. 7, pp. 152701–152714, 2019.
- [46] R. Jiang, A. T. S. Ho, I. Cheheb, N. Al-Maadeed, S. Al-Maadeed, and A. Bouridane, "Emotion recognition from scrambled facial images via many graph embedding," *Pattern Recognit.*, vol. 67, pp. 245–251, Jul. 2017.
- [47] R. M. Jiang, D. Crookes, and N. Luo, "Face recognition in global harmonic subspace," *IEEE Trans. Inf. Forensics Security*, vol. 5, no. 3, pp. 416–424, Sep. 2010.
- [48] R. M. Jiang, D. Crookes, N. Luo, and M. W. Davidson, "Live-cell tracking using SIFT features in DIC microscopic videos," *IEEE Trans. Biomed. Eng.*, vol. 57, no. 9, pp. 2219–2228, Sep. 2010.
- [49] R. M. Jiang, A. H. Sadka, and D. Crookes, "Multimodal biometric human recognition for perceptual human-computer interaction," *IEEE Trans. Syst., Man, Cybern. C, Appl. Rev.*, vol. 40, no. 6, pp. 676–681, Nov. 2010.
- [50] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, pp. 436–444, May 2015.
- [51] G. Storey, R. Jiang, and A. Bouridane, "Role for 2D image generated 3D face models in the rehabilitation of facial palsy," *Healthcare Technol. Lett.*, vol. 4, no. 4, pp. 145–148, Aug. 2017.
- [52] Z. Jiang, P. L. Chazot, M. E. Celebi, D. Crookes, and R. Jiang, "Social behavioral phenotyping of drosophila with a 2D–3D hybrid CNN framework," *IEEE Access*, vol. 7, pp. 67972–67982, 2019.
- [53] G. Storey, R. Jiang, S. Keogh, A. Bouridane, and C.-T. Li, "3DPalsyNet: A facial palsy grading and motion recognition framework using fully 3D convolutional neural networks," *IEEE Access*, vol. 7, pp. 121655–121664, 2019.
- [54] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *Proc. Adv. Neural Inf. Process. Syst.*, 2014, pp. 2672–2680.



CHIA-YEN CHIANG received the bachelor's degree in finance and international business from Fu Jen Catholic University, Taiwan, in 2016, and the M.Sc. degree in data science program. She is currently pursuing the Ph.D. degree with the Queen Mary University of London.



She is also an Intern Researcher with 2Excel Geo for her Master thesis project. She has taken part in a number of projects, including forest analysis, social media spam detection, sentimental analysis on post terrorist attacks, and deep learning on cloud. Her current research interests include statistical methods and modeling, statistical inference, data mining, machine learning, big data systems, and neural networks.

CHLOE BARNES received the Ph.D. degree in remote sensing from the University of Leicester.

She joined the 2Excel Geo team, in 2017. She is currently the Head of Remote Sensing, 2Excel-Geo. She is also a Domain Expert in tree disease detection using spectral imaging and LiDAR techniques from airborne platforms and is an Experienced Data Analyst. Since joining 2Excel, she has expanded her skills and experience and progressed into the Head of Remote Sensing role, in 2019. Her research interest includes hyperspectral data accessible to end-users, while machine learning can be exploited for the analysis of spectral data acquired by airborne hyperspectral remote sensing.



PLAMEN ANGELOV (Fellow, IEEE) received the Ph.D. and D.Sc. degrees. He is currently a Professor of intelligent systems, and the Director of Lancaster Intelligent, Robotic and Autonomous Systems (LIRA) Research Centre, Lancaster University. He has authored or coauthored 350 peer-reviewed journal and conference publications, including three research monographs, three granted USA patents, and so on. He is Fellow of IET and ELLIS, and the Vice

President of the International Neural Networks Society. He chaired multiple IEEE conferences, Technical Committees, and received awards by IEEE and INNS. He is also the Editor-in-Chief of *Evolving Systems*—a Nature Springer journal and an Associate Editor of a number of top AI journals, including the IEEE TRANSACTIONS ON CYBERNETICS and the IEEE TRANSACTIONS ON FUZZY SYSTEMS.



RICHARD JIANG is currently a Senior Lecturer (Associate Professor) with the School of Computing and Communications, Lancaster University, U.K. He has supervised and co-supervised over ten Ph.D. students. He has authored over 80 publications and was the Lead Editor of two Springer books. His research interests include artificial intelligence, XAI, neural computation, biomedical image analysis, intelligent systems, and biometrics and privacy. His recent research has been supported by grants from EPSRC, Leverhulme Trust, Qatar Science Foundation, and other industry funders. He is a Fellow of HEA, an Associate Member of the EPSRC College, and an EPSRC RISE Connector. He also holds a Leverhulme Research Fellowship.