Project title:	Detection of Strawberry Disease with Deep Learning and Computer Vision
Project number:	SF TF 170
Project leader:	Dr B. Al-Diri, University of Lincoln, Supervisor, Dr A Hunter, University of Lincoln, Supervisor, Dr X. Xu, NIAB EMR, Supervisor
Report:	Annual Report, Year 1
Previous report:	None
Key staff:	Wayne Andrews
Location of project:	University of Lincoln, Riseholme & Brayford Campus, NIAB EMR
Industry Representative:	Richard Harnden, Berry Garden Growers Ltd
Date project commenced:	07/11/2018

DISCLAIMER

While the Agriculture and Horticulture Development Board seeks to ensure that the information contained within this document is accurate at the time of printing, no warranty is given in respect thereof and, to the maximum extent permitted by law the Agriculture and Horticulture Development Board accepts no liability for loss, damage or injury howsoever caused (including that caused by negligence) or suffered directly or indirectly in relation to information and opinions contained in or omitted from this document.

The results and conclusions in this report are based on an investigation conducted over a one-year period. The conditions under which the experiments were carried out and the results have been reported in detail and with accuracy. However, because of the biological nature of the work it must be borne in mind that different circumstances and conditions could produce different results. Therefore, care must be taken with interpretation of the results, especially if they are used as the basis for commercial product recommendations.

AUTHENTICATION

We declare that this work was done under our supervision according to the procedures described herein and that the report represents a true and accurate record of the results obtained.

Wayne Andrews

PhD Student

University of Lincoln

Signature .... Date ....

Dr B Al-Diri

Associate Professor

University of Lincoln

Signature .... Date ....

CONTENTS

GROWER SUMMARY

Headline

Improve strawberry crop yield by embracing automated disease detection powered by Computer Vision and Deep Learning.

Background

The soft fruit industry in the UK is a multi-million-pound industry with demand doubling over the past 20 years. With the average price of a punnet of strawberries (400g) costing roughly £2 for the same period, representing a real term fall in price. Growers are having to be ever more efficient with their crop if this trend is to be maintained. For most of the peak-season (April – October) the UK is almost self-sufficient with ~90% of the crop grown under cover to extend the season from what is once was (June-July). With the UK consuming 131,000 tonnes of strawberries between 2018-19 and spending more than £653 million (FarmingUK, 2019), there is a great need to ensure the crop is healthy in order to meet the demand. Should the UK need to rely on imported strawberries it is estimated the price would increase 50% with farmers arguing negative impacts on the environment also with transporting the crop from mainland Europe and beyond to the UK (Wheeler, 2018)

With the ever-increasing surge towards automation in the Agri-Tech industry coupled with the looming potential shortfall in labour (Doward, 2019), there is an opening for autonomous systems for the soft fruit industry. Growers control disease in strawberry crops and extend shelf life of picked fruit using a range of crop protection products, however there are ever increasing restrictions placed upon what can be used, some herbicides and pesticides have been completely phased out due to new regulations. This has left an industry looking to new ways to maintain a healthy crop and produce a profitable yield.

Current crop disease management is accomplished in a very analogue manner with skilled agronomists having to painstakingly inspect the crop at a grower's site, it is hoped that by using Deep Learning this work can be made easier by gathering data from the entire crop traversed and highlighting areas that may require further attention or intervention.

Deep learning is a subset of the larger field of machine learning, vast neural networks comprised of many layers inspired by the way the human brain processes information, if given enough data to learn from a deep learning system can allow a machine to solve complex problems. There are deep learning systems everywhere around us, from virtual assistants such as Amazon’s Alexa, Apple’s Siri or Microsoft’s Cortana to vision systems for pilotless drones or autonomous vehicles. (Marr, 2018)

Over the past few years there has been a rapid surge in deep learning coupled with higher resolution imaging sensors becoming more readily available at more affordable prices. Solutions are becoming available to move towards a more automated management strategy to better utilise highly trained staff and deploy them where needed.

Botrytis alone is thought to be responsible for as much as €10 billion in lost crop globally each year, Phytophthora causes crown and root rots in strawberries and can swiftly progress to plant death, in turn reducing profitable yield (FungiAlert, 2019)

Powdery Mildew attacks the leaves, flowers and fruit of the strawberry, and can result in yield losses from 20% to 70% of crop potential (Avice M Hall, 2017). It is almost impossible to have the plants under continuous surveillance and with the current political climate combined with the dropping value of the pound there are less seasonal workers from the EU available to monitor, harvest and maintain healthy crops. Estimates of labour shortages as high as 30% have been reported on some farms with the Home Office estimating ~80,000 positions needed filling in the 2019 season, mostly by workers from other EU countries (Doward, 2019).

Summary

Using fruit that has been inoculated by three different commonly occurring fruit rot pathogens (Rhizopus, Mucor and Botrytis), this project has so far demonstrated the ability to use existing deep learning models to detect disease present on post-harvest fruit. This type of detection would be useful in a packing environment, checking punnets of fruit as they pass by on a conveyor and rejecting those which are potentially unsuitable for sale. Using a state-of-the-art ‘Mask R-CNN’ model it was possible to achieve 78.54% accuracy for instance detection (instance detection not only detects the presence of a class within an image, it also detects the individual instances of each class), based on training accomplished with a very small amount of data. Classification of disease present in an image of a strawberry was giving accuracies of upto 92.05% using an increased amount of data. A dataset of powdery mildew has been collected covering the period from inoculation to visual symptoms becoming present, this dataset is currently in the labelling stage.

Financial Benefits

Once completed this project will enable growers to monitor their crops more effectively and address potential problems before issues spread to neighbouring crops. It may also be possible to visually screen fruit as it is packed and help identify fruit that is showing signs of infection that may drastically reduce shelf life.

Action Points

There are no grower action points at this early stage of the project.

SCIENCE SECTION

Introduction

Over the years, attempts have been made to use computer vision techniques to detect various diseases in strawberry plants and fruit. One such study, by Ayesha and Jani, 2017, used images segmented using K-Means clustering before extracting features using Grey Level Co-Occurrence Matrix, these features were then classified using a Support Vector Machine (SVM) and an Artificial Neural Network (ANN). However, no further information is given to indicate which classification method performed better. The authors of the paper simply state one of the advantages of the SVM is a high accuracy with less training data as opposed to a large amount of data that is required for the ANN with the increased complexity compared to using the SVM. The authors also state that using a Random Forest Classifier also requires a large number of training images compared to the SVM, however it is not clear whether they used Random Forests or were simply outlining a difference between the two approaches (Ayesha and Jani, 2017).

Pattern recognition has also been used by Changqi Ouyang et al. (2013) where a segmentation algorithm was created utilising heavy pre-processing, mean shift and pattern recognition to extract features to use in an SVM. The system was dealing with three factors commonly leading to strawberries being out of specification; powdery mildew, shrinkage and uneven ripening. The authors used a combination of 12 colour features, 8 texture features and 9 shape features. The authors used a three-layer neural network for classification as well as a SVM in testing and found that the SVM outperformed the ANN in their experiments which led them to favour the SVM for their classification of an ‘out of specification’ strawberry. There are no published results for either of the classification methods used however (Ouyang *et al.*, 2013).

There is not a publicly available dataset that contains real-world images of strawberry diseases and pests in real world environments. Although there are resources such as Plant Village, containing 54k+ datasets of plant leaves this dataset is not exclusive to strawberry leaves and the images were taken under lab conditions with a homogenous background and the leaves are presented face up, all of which is not representative of real world examples. Besides this, the current project will be looking at much more than just the leaves, with the stems, flowers, fruit and crowns all forming part of the proposed solution.

To this end the focus of work carried out in the first year of the current study has been compiling a dataset of images for disease found on the crop itself in the field due to the lack of comprehensive publicly available datasets for strawberry disease and also detection of disease found in post-harvest fruit.

Data collection, and ground truthing (labelling the data), is an essential step that must be completed. Any system proposed must be robust enough to handle the differences in cultivars used and the vast difference in growing conditions, one such way is to ensure there is access to an adequate number of training data to begin with.

Methods and Results

Post-Harvest Fruit Data Collection & Inoculation – Trial 1

Location - NIAB EMR (Kent)

Aim/Purpose - Understanding the visual differences between pathogens that can affect the shelf-life of strawberries.

Materials used

- Store bought strawberries with no visible defects
- 4 trays
- Paper roll/blue roll to line the trays
- Clear plastic bags large enough to cover the trays with headroom
- Tape to seal the bag
- Camera to capture the images

Procedure

Strawberries were purchased from a local supermarket before being placed onto 4 trays at random, these strawberries were then inoculated with one of three commonly occurring fruit rot pathogens (Rhizopus, Mucor and Botrytis) taken from petri dishes where the pathogen had been isolated and incubated on Potato Dextrose Agar (PDA), with a fourth tray being left uninoculated (control). The trays were then covered with a large plastic bag and sealed, before being stored in ambient temperature (18-20 degrees Celsius).

After 6 days of storage the trays were then placed into a “fume hood”, and images were captured of the trays and the strawberries.

Images were captured using the rear facing camera of a Huawei Honor 9 with the camera settings as described in table 1.

Table 1 Camera settings for the rear facing camera of a Huawei Honor 9.

Camera model	STF-L09
F-stop	f/2.2
ISO Speed	ISO-125
Exposure time	1/50 sec
Focal length	4mm
35mm focal length	27
Resolution	3264x2448
Colour representation	sRGB

Results

The images collected were labelled and used to train a Mask R-CNN network, the network was trained to detect 4 classes, Rhizopus, Botrytis, Mucor and Healthy. Four models were trained using a combination of different backbones and weights (as shown in Table 2).

On average the best performing combination for model selection was the ResNet 50 backbone with the weights trained initially on the COCO dataset, this combination achieved 78.54% on a very limited amount of data.

Table 2. Results from using the Mask R-CNN network on a small dataset of post-harvest fruit.

Label	Instances	ResNet 50 (coco weights)	ResNet 101 (coco weights)	ResNet 50 (imagenet weights)	ResNet 101 (imagenet weights)
Rhizopus	43	92.80%	76.10%	62.52%	56.29%
Botrytis	52	93.59%	90.12%	76.21%	81.48%
Mucor	32	94.44%	58.97%	94.44%	60.40%
Healthy	4	33.33%	0.00%	50.00%	50.00%
Total	**131**	**78.54%**	**56.30%**	**70.79%**	**62.04%**

Post-Harvest Fruit Data Collection & Inoculation – Trial 2

Location - Refrigeration Unit, Riseholme Campus, University of Lincoln

Aim/Purpose - Obtain data from early stages of infection and determine which architecture to use as a feature extractor for post-harvest fruit rot/disease object detection.

Materials used

- Store bought strawberries with no visible defects
- 4 trays
- Paper roll to line the trays
- Clear plastic bags large enough to cover the trays with headroom
- 0.5ml pipets
- Tape to seal the bag
- Camera to capture the images
- Iso-propanol alcohol (95%) to reduce cross contamination risk
- Lightbox to capture images while limiting exposure to ambient air

Procedure

The inoculum was taken from petri dishes using pathogens collected from NIAB EMR in Kent, under the guidance of Dr Tom Passey of NIAB EMR. The inoculum was incubated on PDA, before being transported whilst kept in a sealed bag inside an insulated "cool/hot" box to Riseholme Campus.

Strawberries were purchased from a local supermarket before being placed onto 4 trays at random, each strawberry on each tray was then individually inoculated with one of three pathogens, with one tray (control) being mock inoculated with distilled water.

For each of the three pathogens the inoculum was scraped from the Petri dish and mixed into 50ml of distilled water, before each strawberry was inoculated by pipette at three random locations on each fruit.

The trays were then covered with a large clear plastic bag and sealed, before being stored in the refrigerator (between 6-8 degrees Celsius). Images were captured on day one, then every day from day 3 until day 10, then once more on day 12 when the fruit was disposed of.

The temperature that the trays were stored in was raised from 7 degrees C to 12 degrees C on day 7 as the fruit was not showing obvious visible signs of infection and the temperature was the most readily available variable to change in order to facilitate the growth of the pathogens.

By far the most time-consuming task was the ground truthing/labelling of the images collected, a selection of 800 images were chosen and uploaded to servers at www.labelbox.com for labelling (LabelBox, 2019).

The images selected were, 200 from the uninoculated tray day 3 to day 5, 200 from the Mucor tray day 9 to day 12, 200 from the Botrytis tray day 9 to day 12, and 200 from the Rhizopus tray day 9 to day 12.

Images were labelled with LabelBox as outlined in the Data Labelling section. Once the strawberries had been labelled there was a slight imbalance in the amount of healthy strawberries verses each of the other classes (see **Figure 1**).

When labelling the data, an annotator creates a mask that highlights the affected strawberry and associates that mask with a class, a machine/computer doesn't know the difference between what is a strawberry or what is a tray or anything else, so it's this spatial information that is vital in order to train any machine learning system to distinguish which part of an image is of interest.by masking out the highlighted area of an image. The masks created during the labelling process will be used for object/instance detection in future work however, for this experiment only the class the image belongs to will be used. So instead of the system highlighting the individual strawberry and its class, it will instead give a detected class. It is anticipated that this will result in lower accuracies considering each image may have other strawberries present, however at this stage it is only necessary to determine which model has the more robust feature extraction layers and robust performance given these challenging circumstances.

The resulting labelled dataset yielded 1743 instances of which the classes were distributed as per the figure 1.

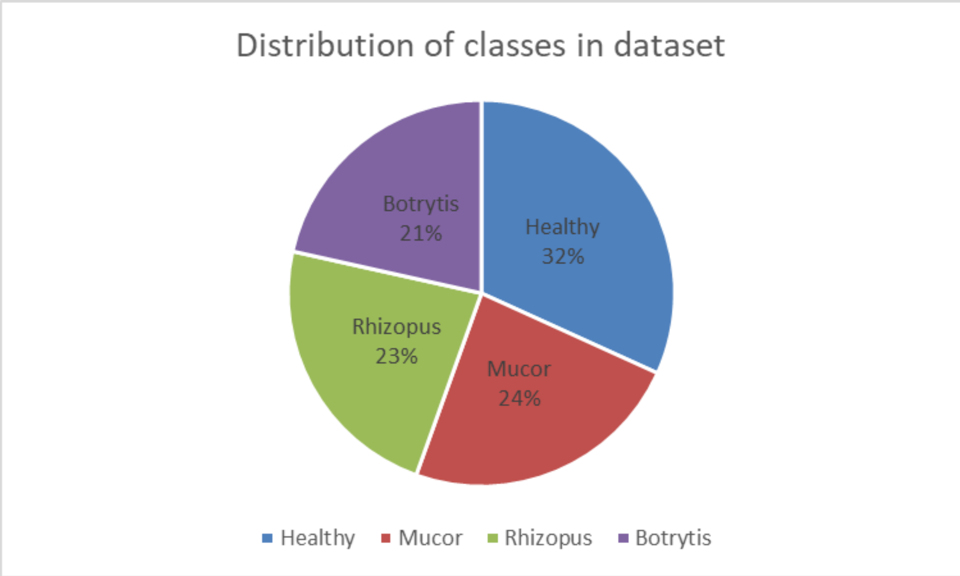

Figure 1. Distribution of classes amongst the dataset.

The dataset was shuffled and arranged using a 75/20/5 split before using 6 common architectures to classify the images. Each model was loaded from Keras and were initialised with the Imagenet weights as the networks were originally trained on more than 14 million images (the number of images with bounding box annotations in the dataset 1,034,908), whereas the COCO dataset is considerably smaller with 330,000 images. Future work will analyse the differences between the two sets of weights on the same dataset. (Lin, Patterson and Ronchi, 2019; Stanford Vision Lab, 2019).

With the pretrained loaded models obtained without the final classification layers, it was necessary to add layers to suit the four-class output required in this instance.

Results

Each model was then trained for 60 epochs and their performance evaluated over the testing set (new data not previously seen, but similar in appearance). Initially the data was passed into the model with no augmentation or manipulation other than resizing down to 224x224px. However, after training three of the models it was clear that the models were overfitting to the training data and the validation accuracy showed this to be the case. Following this result new models were trained, again for 60 epochs and this time each image was augmented when passed into the model to reduce this overfit.

From the six original models (Table 3) MobileNet was chosen due to its smaller size and performance based upon the initial 60 epochs training. This model was then trained using K-Fold cross validation (K was chosen to be 7) for 100 epochs, where the input images and labels were split using SK-Learn's method which resulted in a split of ~83/17 for training and testing for each fold.

Table 3. Summary of the 6 models trained to determine which model to use for feature selection.

Model	Validation Accuracy (%)	Model size (MB)
DenseNet201	92.05	80.01
ResNet 50	92.05	98.2
MobileNet V1	87.5	18.4
VGG 19	81.82	78.4
VGG 16	79.55	58.1
Inception V3	54.55	99.6

Powdery Mildew Cultivation and Data Collection

Location - NIAB EMR (Kent)

Aim/Purpose - To collect images that cover the period from healthy to infected plant

Materials used

- Malling Centenary cultivar Strawberry plants x60
- 5 trays
- Camera to capture the images
- Paint brush to transfer the inoculum
- Paper tags to label the leaves that were inoculated

Procedure

Healthy plants were transferred into greenhouse for inoculation. Inoculum was collected from plants that were inoculated 7 days prior with the use of a paint brush. 48 plants were chosen at random and two sets of leaves were tagged on each plant. Inoculum was transferred to the middle leaflet of the tagged leaves.

Table 4 Camera settings for the Sony DSLR.

Camera model	ILCE-6000
ISO Speed	ISO-100
Focal length	50mm
35mm focal length	75

Resolution	6000x4000
Colour representation	sRGB

Images (captured using the Sony ILCE-6000) were taken each day of each set of tagged leaves for each plant and of the whole plant to capture any infection that may be present upon untagged leaves. Ground truth data was collected by recording which plants/sets of leaves had visual powdery mildew colonies present.

Results

The purpose of this experiment was to gather temporal data for strawberry plants both before symptoms are visible and after, with a view to training a deep learning system to detect the presence of powdery mildew before obvious visual colonies appear.

Due to the lack of available inoculum present on the previously inoculated plants 7 days prior there was only enough to transfer using the brush method to 34 plants which was on a total of 68 leaves (both T1 and T2).

The results obtained and presented in table 5 show there was a ~57% infection rate from the plants that were inoculated. There were also a few plants that had naturally become infected with early visible symptoms that were not noticed until the plants had already been transferred into the greenhouse. Rather than tag the affected leaves and use these towards the inoculated plants, the decision was taken to observe the inoculated plants during the early stages of infection to then use this data in the future for early detection.

Table 5. Resulting observations from the Powdery Mildew Data collection at NIAB EMR.

Day	**0**	**1**	**2**	**3**	**4**	**5**	**6**	**7**
Number of plants	34	60	60	60	60	60	60	60
Innoculated	0	24	36	48	48	48	48	48
T1	3	6	6	12	17	18	25	32
T2	0	2	3	8	12	16	24	30
Other	0	8	8	8	9	9	18	21
Infection rate from innoculations								
T1 rate		25.00%	16.67%	25.00%	35.42%	37.50%	52.08%	66.67%
T2 rate		8.33%	8.33%	16.67%	25.00%	33.33%	50.00%	62.50%
Other rate		33.33%	22.22%	16.67%	18.75%	18.75%	37.50%	43.75%
Overall infection								
T1 infection %	8.82%	10.00%	10.00%	20.00%	28.33%	30.00%	41.67%	53.33%
T2 infection %	0.00%	3.33%	5.00%	13.33%	20.00%	26.67%	40.00%	50.00%
Other infection %	0.00%	13.33%	13.33%	13.33%	15.00%	15.00%	30.00%	35.00%

Discussion

It seems quite clear from the results that there is great potential for the detection of the three pathogens (Rhizopus, Mucor and Botrytis) on the post-harvest fruit that has been inoculated.

The Mask R-CNN model gave 78.54% accuracy in the detection of the instances and the segmentation was encouraging given the small number of images used for the training.

The selection of a model to use for the feature extractors was a good exercise in using the currently available toolsets for labelling data and ultimately using this data in a meaningful way. The labelling process itself was cumbersome and could be improved, and it is hypothesised that by using an additional pre-processing step before classification would increase the accuracies further as there would be the opportunity to remove ambiguity with multiple classes being present within the same image. The loss function used was "categorical cross-entropy", so the model was assuming there was only a single correct answer for each given image however this was not the case in images that had more than one strawberry in it. Further experimentation using different loss functions is needed.

The image augmentation used is also helping to improve the overall accuracy in the initial training runs, given the possibility of the system encountering strawberries that could be in any given orientation, it was only natural to allow the inclusion of images at various orientations and scales, and it is shown in the training curves that it made a noticeable improvement and reduced the overfitting of the models. There is a lot of variability between epochs which suggest that a lower learning rate would be advantageous in future training.

The K-Fold cross validation did not provide the expected results, the cause was found to be the inclusion of a dropout layer instead of a fully connected layer, the results showed a lower overall accuracy for both the training sets and the validation sets.

In the powdery mildew experiment, there were a few naturally occurring infections on day zero from the stock plants, by the end of the experiment the infection rate for the inoculated leaves was between 62-66% with 43% of other plants having naturally transferred colonies of powdery mildew present (see table 5).

Transference of the inoculum via the brushing method is effective as is placing test plants in proximity of infected spreader plants. It was noted that for the test cultivar (Malling Centenary) at least, most infections not only were visible on the top of the leaf, there was also mycelium present on the underside of the leaves which was able to transfer to neighbouring plants when the trays were placed close to one another.

1400+ images were gathered from this experiment covering a range of healthy plants to mildly infected plants with some leaves showing signs of infection, this data will be useful in determining the early stages of infection. The next stage would be to label all the data to enable use for training detection networks with the possibility of generating more data using Generative Adversarial Networks to vastly increase the available data and possibly increase both robustness and accuracy.

Conclusions

- Using a state-of-the-art Mask R-CNN model it was possible to achieve **78.54%** accuracy for instance detection based on training accomplished with a very small amount of data (experiment 1).
- Classification of disease in an image of a post-harvest strawberries in laboratory conditions (experiment 2) achieved up **92.05%** for the top performing model, and **87.5%** for the chosen model with the more efficient model in terms of model size.
- Data labelling is a process that requires optimisation to enable more efficient use of collect images.
- Future work will investigate the estimation of the shape of the leaves to aid the detection of powdery mildew where there are not yet visible mycelium present.

Knowledge and Technology Transfer

To aid the collection of data and to help leverage the collective expertise of Berry Garden's expert agronomists a mobile data collection app is being developed.

Attended Deep Learn 2019 in Warsaw, Poland and discussed a range of model architectures that may be suitable with academics from institutions from around Europe and further afield, whilst also learning the current state of the art with regards to model optimisation and size reduction.

References

Avice M Hall, X. J. and J. D. (2017) 'Control of strawberry powdery mildew under protection'. Available at: https://projectblue.blob.core.windows.net/media/Default/Imported Publication Docs/Control of strawberry powdery mildew under protection.pdf.

Ayesha, Lt. and Jani, P. D. G. (2017) 'Disease Detection in Leaves using Image Processing Techniques', *International Research Journal of Engineering and Technology(IRJET)*, 4(4), pp. 1639–1642. Available at: https://www.irjet.net/archives/V4/i4/IRJET-V4I4336.pdf.

Doward, J. (2019) *Bumper British strawberry crop may rot as EU pickers stay away | Business | The Guardian*, *The Guardian*. Available at: https://www.theguardian.com/business/2019/may/04/shortage-fruit-pickers-britain-compete-germany (Accessed: 29 October 2019).

FarmingUK (2019) *British strawberries hit shelves early due to warmer weather - FarmingUK News*, *farminguk.com*. Available at: https://www.farminguk.com/news/british-strawberries-hit-shelves-early-due-to-warmer-weather_52005.html (Accessed: 7 April 2020).

FungiAlert (2019) *The future looks rosy for the UK's berry market - FungiAlert*. Available at: https://fungialert.com/the-future-looks-rosy-for-the-uks-berry-market/ (Accessed: 28 October 2019).

LabelBox (2019) *LabelBox*. Available at: www.labelbox.com.

Lin, T.-Y., Patterson, G. and Ronchi, M. R. (2019) *COCO - Common objects in Context*, *cocodataset.org*. Available at: http://cocodataset.org/#home.

Marr, B. (2018) *What Is Deep Learning AI? A Simple Guide With 8 Practical Examples*, *forbes.com*. Available at: https://www.forbes.com/sites/bernardmarr/2018/10/01/what-is-deep-learning-ai-a-simple-guide-with-8-practical-examples/#68218f4a8d4b (Accessed: 7 April 2020).

Ouyang, C. *et al.* (2013) 'The Research of the Strawberry Disease Identification Based on Image Processing and Pattern Recognition', in. Springer, pp. 69–77. doi: 10.1007/978-3-642-36124-1_9.

Stanford Vision Lab (2019) *ImageNet: Summary & Statistics*, *image-net.org*. Available at: http://image-net.org/about-stats.

