

---

# FRUIT QUALITY AND DEFECT IMAGE CLASSIFICATION WITH CONDITIONAL GAN DATA AUGMENTATION

---

A PREPRINT

**Jordan J. Bird**

Department of Computer Science  
Nottingham Trent University  
Nottingham, United Kingdom  
jordan.bird@ntu.ac.uk

**Chloe M. Barnes**

College of Engineering and Physical Sciences  
Aston University  
Birmingham, United Kingdom  
barnecm1@aston.ac.uk

**Luis J. Manso**

ARVIS Lab  
Aston University  
Birmingham, United Kingdom  
l.manso@aston.ac.uk

**Anikó Ekárt**

College of Engineering and Physical Sciences  
Aston University  
Birmingham, United Kingdom  
a.ekart@aston.ac.uk

**Diego R. Faria**

ARVIS Lab  
Aston University  
Birmingham, United Kingdom  
d.faria@aston.ac.uk

January 27, 2023

## ABSTRACT

Contemporary Artificial Intelligence technologies allow for the employment of Computer Vision to discern good crops from bad, providing a step in the pipeline of selecting healthy fruit from undesirable fruit, such as those which are mouldy or damaged. State-of-the-art works in the field report high accuracy results on small datasets (<1000 images), which are not representative of the population regarding real-world usage. The goals of this study are to further enable real-world usage by improving generalisation with data augmentation as well as to reduce overfitting and energy usage through model pruning. In this work, we suggest a machine learning pipeline that combines the ideas of fine-tuning, transfer learning, and generative model-based training data augmentation towards improving fruit quality image classification. A linear network topology search is performed to tune a VGG16 lemon quality classification model using a publicly-available dataset of 2690 images. We find that appending a 4096 neuron fully connected layer to the convolutional layers leads to an image classification accuracy of 83.77%. We then train a Conditional Generative Adversarial Network on the training data for 2000 epochs, and it learns to generate relatively realistic images. Grad-CAM analysis of the model trained on real photographs shows that the synthetic images can exhibit classifiable characteristics such as shape, mould, and gangrene. A higher image classification accuracy of 88.75% is then attained by augmenting the training with synthetic images, arguing that Conditional Generative Adversarial Networks have the ability to produce new data to alleviate issues of data scarcity. Finally, model pruning is performed via polynomial decay, where we find that the Conditional GAN-augmented classification network can retain 81.16% classification accuracy when compressed to 50% of its original size.

**Keywords** Fruit Quality · Cultivar · Image Classification · Data Augmentation · Convolutional Neural Networks · Generative Adversarial Networks

## 1 Introduction

Recognition of fruit quality is important in smart agriculture to increase production efficiency. Contemporary Artificial Intelligence technologies allow for the employment of Computer Vision to discern good crops from bad, providing a step in the pipeline of selecting healthy fruit from undesirable fruit, such as those which are mouldy or damaged.

Even during modern times, collecting a dataset that generally represents a species of fruit poses difficulties. For example, the widely used Fruits 360 dataset [1] contains 2134 images of apples belonging to one of thirteen cultivars. According to the United Nations, 87.2 million tonnes of apples were farmed globally in 2019 alone [2]. In terms of smart agriculture, this highlights the issue of data scarcity. On the problem of crop quality recognition, much data is required to generalise to a population and thus become apt for real-world use. Considering the yield of fruit globally compared to practical data collection, bridging this gap manually i.e., *collecting more data*, is simply an impossible task.

Computer vision has long been noted as an important approach for fruit processing[3, 4, 5]. Further work also explores varying techniques for computer vision in smart agriculture such as transfer learning from pre-trained weights[6, 7, 8], and current research has started to explore the approach of data augmentation in improving such problems[9, 10]. In this work, we focus on exploring a solution to this problem for lemon harvesting. Recording data for both lemons and limes, the United Nations noted that 20 million tonnes were harvested in 2019 [2]. According to the CIA World Factbook [11], lemon and lime fruit exports comprised over \$3.3 billion USD of international exports in 2019. The largest exporters were Spain (\$828 million USD), Mexico (\$523 million USD), and the Netherlands (\$339 million USD). Although the top 15 countries exported 93.5% of all lemons and limes in 2019, many countries are expanding their efforts; for example, Belize increased lemon and lime exports by 11,100% from 2018 to 2019, followed by Timor-Leste and Georgia which increased exports by 4,200% and 2,115%, respectively. With this information in mind, the ability to autonomously select and reject lemon fruit based on health, e.g., sorting out those which have developed mould or gangrene, would allow for further increases in an already growing market by increasing production efficiency.

In this article, we explore a Conditional Generative Adversarial Network-based solution to data scarcity through training data augmentation. The main scientific contributions we present in chronological order are as follows:

- Exploration of Convolutional Neural Network (CNN) topologies for fruit quality recognition.
- Implementation of a Conditional Generative Adversarial Network (Conditional GAN) for the generation of synthetic healthy and unhealthy fruit images. The trained synthetic data generation model is made available for future work<sup>1</sup>.
- Augmentation of the original dataset with synthetic images to improve the CNN performance.
- Exploration of features within synthetic images shows that the Conditional GAN learns to generate healthy synthetic fruit images with no defects, as well as unhealthy synthetic fruit images with defects such as mould and gangrene.
- Model pruning shows that a Conditional GAN-augmented classification network can retain 81.16% classification accuracy when compressed to 50% of its original size.

## 2 Background and Related Work

### 2.1 Fruit Quality Recognition

Fruit Quality Recognition is a technique where a fruit can be scored or classified autonomously by an algorithm given input features such as photographs. As previously mentioned, the ability to perform this task autonomously (and non-invasively [12]) allows for an increase of production efficiency since sorting can be performed by machines endowed with such an algorithm.

Reduction in cost of such a system is of particular interest in the field given that several Lower Economically Developed Countries (LEDCs) are expanding the production and export of lemon fruit [11]. Solutions such as electronic noses can cost up to 100,000\$ USD [13], whereas a camera and computer are a fraction of the cost, arguing that image recognition is a more viable option when cost is an issue. It is worth noting though, that low-cost electronic noses are currently an expanding line of research within the Sensors and Internet of Things (IoT) fields [14]. Electronic noses are indeed strongly performing solutions to fruit quality recognition as shown by [15], [16], and [17].

Deep learning approaches to fruit classification are abundant, as shown by multiple literature reviews on the subject [18, 19, 20, 21, 22]. In comparison, the application of deep learning to the more fine-grained problem of quality classification are relatively rare. Bhargava and Bansal show that the field of automated fruit quality recognition is rapidly growing, in part due to the availability of newer technologies [23].

Given that one can discern the quality of a fruit based on observation, visual features are often noted as of importance when it comes to fruit quality classification. Yamamoto et al. [24] suggested that Linear Discriminant Analysis (LDA)

---

<sup>1</sup>Image generation model weights and code are available at:  
<https://github.com/jordan-bird/synthetic-fruit-image-generator>

of colour, shape, and size enabled the formation of a distance matrix which could be used to classify both the cultivar and quality of strawberries. Usage of a single LDA led to a classification accuracy of 42% whereas combining the three analyses caused accuracy to rise to 68%. Capizzi et al. followed a similar technique through texture and gray-features with a Radial Basis Probabilistic Neural Network, which scored around 97.25% on a limited set of images of orange fruits [25]. Azizah et al. [26] provided a solution to defect classification of the mangosteen fruit, attaining a mean 97.5% 4-fold classification accuracy, albeit with a limited dataset. This pattern of data scarcity continues as would be expected in the field; in 2020, Fan, et al. [27] found that CNNs could classify apple defects with around 96.5% accuracy after processing 300 fruit images (150 per class). This study also notes the efficiency of deep learning algorithms post-training, the algorithm was capable of processing 5 fruit images per second (0.2 seconds each).

In this work, we take inspiration from Osako et al.’s approaches to cultivar discrimination of lychee fruit [28]. The study showed the success of fruit image classification (albeit for a different task of cultivar recognition) when fine-tune transfer learning with the VGG16 CNN [29], and predictions were analysed super-imposed upon the images via Grad-CAM [30] in order to explain useful features for discrimination. We follow a similar approach in this work (applied to a new problem) in terms of fine-tuning of VGG16 and analysis with Grad-CAM, and go a step further in improving classification through data augmentation with a Conditional Generative Adversarial Network to self-regularise the network by creating new, synthetic fruit images.

## 2.2 Data Scarcity and Augmentation

Regarding the global fruit yield statistics versus dataset size examples within the introduction, data scarcity in machine learning is the reliance of models on exhaustive labelling, providing a limitation to their real-world use [31]. Given that the use of a model is to aim towards generalisation of a population, a lack of data can lead to a situation wherein training accuracy scores are high and yet deployment to industry would lead to failure. As noted in the introduction, gathering enough data of cultivated fruit is impractical. Without enough data to properly represent the population, models will be prone to overfitting. Given this, other methods are required to prevent overfitting and encourage generalisation towards the real-world use of a machine learning model outside the realm of simply collecting more data.

Data augmentation is the process of creating new training data by either slightly modifying the data at hand or generating new, synthetic data [32]. An augmented dataset thus provides more training examples for a given task.

Image recognition tasks for Convolutional Neural Network image classification are affected by data scarcity due to their data requirements [33, 34], where many generative models have been recommended to alleviate such issues [35, 36]. Generative models have also been noted to positively impact biological signal classification [37, 38], semantic Image-to-Image Translation [39], speech processing [40, 41], and Human Activity Recognition [42, 43] among many others. In this work, we use a Conditional GAN for data augmentation, which are described in the following section. This is based on the literature wherein GANs and Conditional GANs have been noted to perform particularly well in image classification [44, 45, 46, 47]. We note specific inspiration from Fu et al. [48], where Conditional GANs have been noted to perform well on fine-grained images such as classification of bird and dog breeds (rather than classification of a whole species). This bares similarity to our problem, where finer details on generally similar images dictate which class they belong to.

## 2.3 GAN and Conditional GAN

The Generative Adversarial Network (GAN) was first introduced in 2014 [49]. The idea behind the GAN is to have two neural networks compete in a zero-sum game ergo *adversarial*, i.e., the loss of one network is directly beneficial to the other and vice versa. To give an example of image generation, as this work performs, there are two networks; a generator network which creates images, and a discriminator network which classifies the inputs as either real or fake. As with most deep learning approaches, the gradients of each network are updated after each training batch with a stochastic gradient algorithm. The output of the generator network feeds directly into the discriminator network and thus training of the two networks is automated via their competition. In terms of categorical cross-entropy, a score can be calculated as follows:

$$E_x[\log(D(x))] + E_z[\log(1 - D(G(z)))], \quad (1)$$

where the first part of the equation ( $E_x[\log(D(x))]$ ) is the recognition of real images and the second part ( $E_z[\log(1 - D(G(z)))]$ ) is the recognition of fake images.  $E_x$  and  $E_z$  are the expected values over all real and fake data, respectively; e.g.,  $x$  is a real input from the dataset and  $z$  may begin as a random noise input to a generator. The function  $D(x)$  is the probability that a given data is real and is thus therefore being reversed to discern fake images. Note that  $D(x)$  is replaced by  $G(z)$  in the second part of the equation, this is due to input to the discriminator being Generator  $G$ ’s output when presented with random input vector  $z$ . This is known as a *minimax loss*, since the generator’s aim is to maximise Equation 1 while the discriminator aims to minimise it.

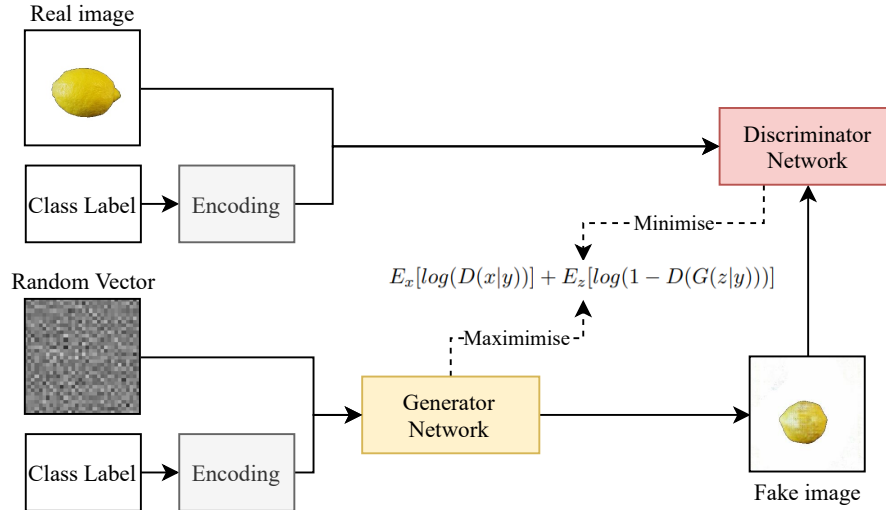


Figure 1: Generalisation of a Conditional GAN topology.

A Conditional Generative Adversarial Network (CGAN or Conditional GAN) [50] is an extension of the above technology, but with a given class label. That is, the generator now aims to learn to generate images belonging to one of  $n$  classes, in this work this is a binary label of “healthy” and “unhealthy”. Equation 1 can be extended as follows:

$$E_x[\log(D(x|y))] + E_z[\log(1 - D(G(z|y)))], \quad (2)$$

where data objects  $x$  and  $D(z)$  are given class label  $y$ . Therefore,  $D(x|y)$  is the discriminator’s probability that  $x$  is real given class label  $y$ , and  $G(z|y)$  is the output of the generator with random vector  $z$  given class label  $y$ . This minute difference in topology from a GAN, as can be observed in Figure 1, allows for the generation of objects belonging to multiple classes. If the dataset in this work was presented to a vanilla GAN, the network would learn to generate fake fruit images by learning from real fruit, thus two networks would then be needed for the generation of either class. Said networks would have to train independently of one another. By using a Conditional GAN, we can specify to the network whether we want it to generate healthy or unhealthy fruit by learning not only to generate them in the general sense, but also learning from the significance of a class label.

### 3 Method

A general overview of the proposed approach can be observed in Figure 2, where the training data is augmented with a Conditional GAN approach. In this section, we describe the method for each step of the experiments performed.

#### 3.1 Data Collection and Preprocessing

Initially, an open source dataset of lemon images were acquired from SoftwareMill [51]<sup>2</sup>. The dataset contains 2690 images of lemons at a resolution of  $1056 \times 1056$  pixels and are annotated in COCO format. Given that each COCO annotation describes one class, i.e., a fruit that exhibits both mould and gangrene will have two individual entries, we sort through the dataset to apply a single binary class label of “healthy” or “unhealthy” to each of the fruit images.

Given the computational complexity of the algorithms involved, the images are then resized to  $256 \times 256$  pixels; this resolution still allows for the visualisation of undesirable features while reducing the total number of RGB pixel values from 3,345,408 ( $1056 \times 1056 \times 3$ ) to 196,608 ( $256 \times 256 \times 3$ ). This reduction to 5% of the original model inputs reduces the amount of memory required for training of all models, since the use of full resolution images is not feasible for consumer-grade hardware. In terms of deployment and real-world usage, robots themselves will have energy restrictions due to the processing cost and profit tradeoff regarding automation of fruit sorting. Thus, this reduction in image size increases the practicality of the approach.

To better discern noise throughout the generative learning process, the black background is replaced with white. Although it would have no effect on the training process of the model, the background is replaced so visual glitches

<sup>2</sup>Note: None of the authors of this work are affiliated with SoftwareMill

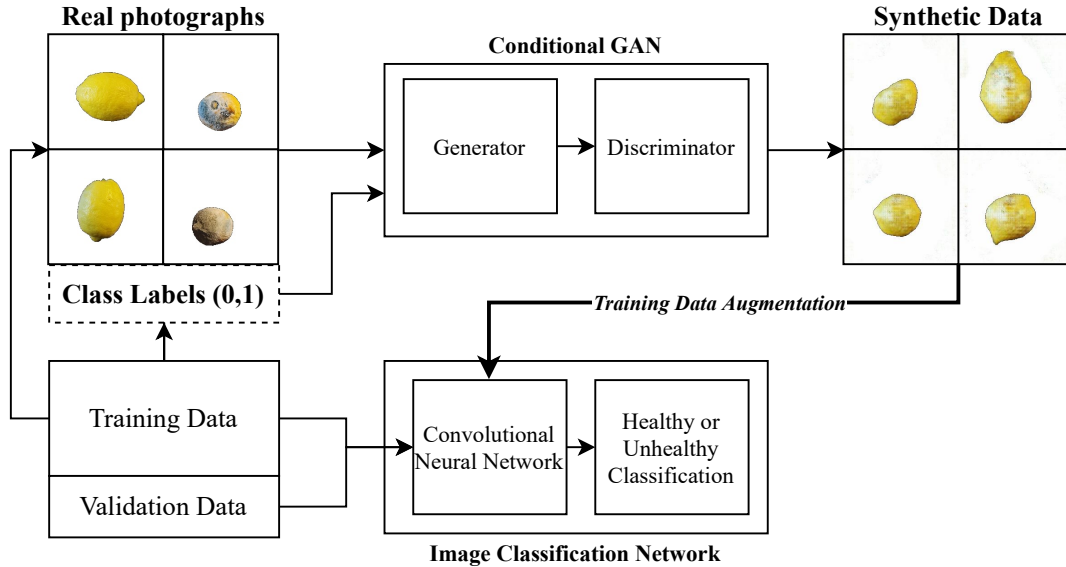


Figure 2: A general overview of our proposed approach for data augmentation towards fruit quality image classification.

can be better discerned through manual observation throughout training. For example, later in Figure 7, several small glitches occur in the eighth generation of outputs that would have been more difficult to observe in the presence of a dark background.

Examples of some images in the preprocessed dataset can be found in Figure 3. Healthy lemons are gathered for the healthy class, whereas mouldy, gangrenous, and those with a dark style remaining are gathered to form the unhealthy class.

### 3.2 Data Augmentation via Conditional GAN

For data augmentation, a Conditional GAN is utilised. The model is selected since it supports the concatenation of the generator and discriminator networks with a second input of the class label. That is, the context of a healthy or unhealthy fruit is specified, and so the model will learn to generate images as belonging to either one of the two classes.

The initial input to the generator is a vector representing a three-channel  $8 \times 8$  pixel image ( $8 \times 8 \times 3$ ). By using Convolutional Transpose layers, this is eventually upsampled to a  $256 \times 256$  RGB image. The discriminator network downsamples twice with convolutional layers of 128 neurons, a kernel size of (3, 3) and a stride of (2, 2). Each layer utilises LeakyReLU activation [52] whereas the output is set as a hyperbolic tangent for scaling, and the ADAM optimiser [53] is used to train. Latent space for class label interpretation is of size 100. Hyperparameter selections are based on the findings of the studies in [54].

The Conditional GAN was first initially trained for 500 epochs, manual exploration of the produced synthetic data showed promise, but several severe visual glitches still occurred. Due to this, the training was extended and performed for 2000 epochs in total with a batch size of 64. It was also observed that batch sizes below 64 caused the generator to cease training after around 10 epochs and failed to learn any further.

### 3.3 Classification, Model Analysis and Pruning

The image classification network itself utilises the concept of fine-tune transfer learning from a large ImageNet-trained model, VGG16 [29]. A diagram of the Convolutional Neural Network topology we use for the classification of fruit quality images can be observed in Figure 4. The final three ReLU layers and SoftMax predictions have been replaced by a single interpretation layer and a single output neuron with a sigmoid activation function for the optimisation of binary cross-entropy.

The number of neurons within the interpretation layer is optimised through a linear search of (8, 16, 32, 64, ..., 8192) neurons. The network is given a maximum of 100 epochs to train, but training is stopped early if no further learning

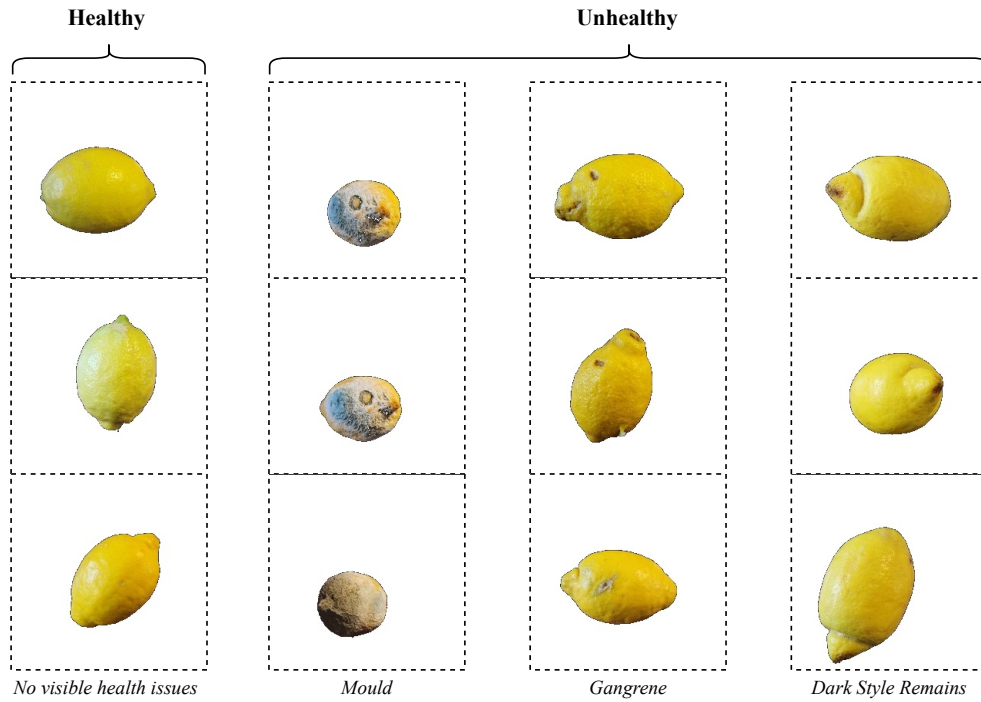


Figure 3: Visualisation of healthy and unhealthy lemons within the dataset. Mouldy, gangrenous, and those with a dark style remaining are all considered unhealthy.

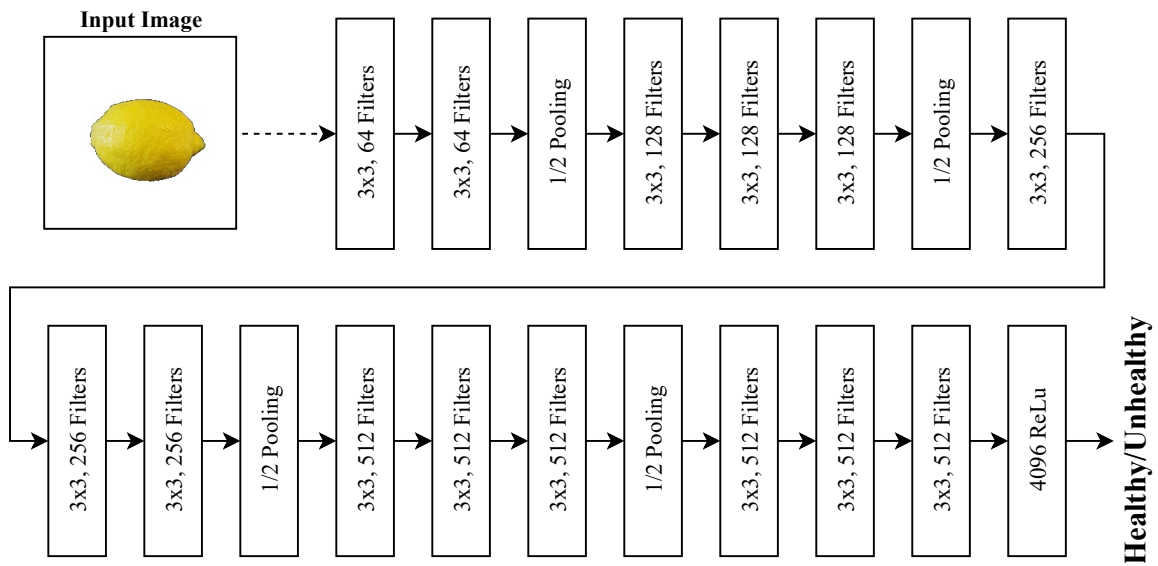


Figure 4: Overview of the VGG16 topology and custom interpretation and output layer which is used for binary image classification.

Table 1: A linear search for the best number of CNN interpretation neurons for the classification of the real dataset, repeated three times with different random seeds.

CNN Output Interpretation Neurons	Validation Accuracy (%)			
	Seed = 1	Seed = 2	Seed = 3	Mean
8	60.1	60.97	57.13	59.40
16	60.1	78.31	71.5	69.97
32	81.41	60.97	57.13	66.50
64	75.71	78.93	57.13	70.59
128	81.04	79.8	76.33	79.06
256	73.23	75.84	65.55	71.54
512	81.91	80.79	79.06	80.59
1024	73.85	83.02	80.17	79.01
2048	83.4	81.91	79.68	81.66
4096	82.28	<b>83.77</b>	79.93	<b>81.99</b>
8192	81.54	81.78	81.41	81.58

occurs within a time of 10 epochs. Hyperparameters are set to a batch size of 64, and the Adam optimiser has a learning rate of 0.001,  $\beta_1 = 0.9$ ,  $\beta_2 = 0.999$ , and  $\epsilon = 1e - 07$ .

Explainability in AI is important, especially when such algorithms are considered for real world usage. Algorithms tend to operate in a black-box like nature, for example, the algorithm in this work would take as input an image of a lemon fruit and produce a class label output, corresponding to whether the fruit is healthy or not. With this in mind, regardless of the training accuracy scores attained, further analysis is needed to explain why decisions are made and predictions are given. We analyse several synthetic images through Gradient-weighted Class Activation Mapping (Grad-CAM) [30, 55]. Class activation maps are produced by the convolutional neural network trained only on real images when given synthetic data as input. This allows us to confirm that undesirable characteristics are indeed both generated and classified as being important; since a GAN generator’s goal is simply to learn to outperform the discriminator, and this may be done through other means i.e. finding methods to *trick* the classifier.

Considering that processing is performed on large quantities of fruit and that there may be time and energy restrictions, model pruning is performed on the network to explore the possibility of smaller model sizes, more apt for real-world usage [56, 57], through polynomial decay. For each model, 9 pruning experiments were performed with weight sparsity ranging from 0.9 (10% of original size) to 0.1 (90% of original size). Pruning was performed on the whole model for 20 epochs, from a sparsity of 0 (full size) to the given value for the individual experiment (0, 0.1, 0.2, ..., 0.9).

### 3.4 Experimental Software and Hardware

The models in this work were implemented in the Keras library with a TensorFlow backend. Models were trained on an RTX 2080Ti GPU (4352 CUDA cores).

## 4 Results

### 4.1 Non-Augmentation Results

Table 1 shows the exploration of interpretation neurons following the VGG16 ImageNet Convolutional Neural Network. Following the three tests, the lowest observed accuracy was that from 8 interpretation neurons with a mean classification accuracy of 59.4%. The best model was the network with 4096 interpretation neurons, which scored 81.99% classification accuracy, and the best individual run was the second run of the 4096-neuron network which scored 83.77% classification accuracy.

### 4.2 Conditional GAN Training

Separated into two graphs for readability purposes, Figures 5 and 6 show the observed losses for the generator and discriminator networks of the Conditional GAN, respectively. Albeit with several anomalous spikes, it can be observed that the generator starts at a high loss of 5.5 which drops throughout the first few epochs. The discriminator, as can be expected, starts low given the quality of output by the generator. The generator can be seen to rise steadily for the first

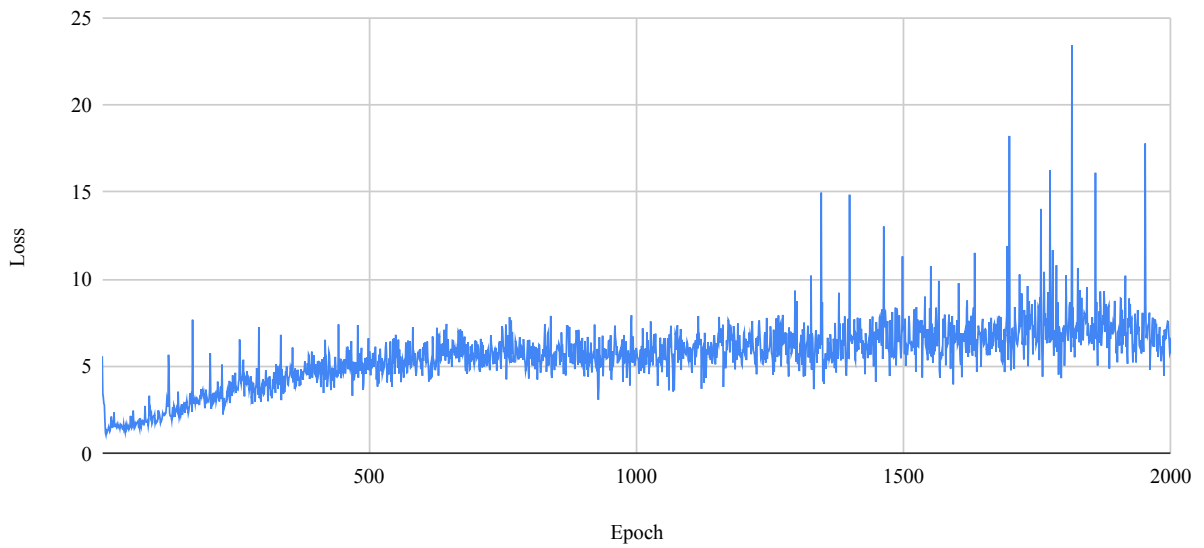


Figure 5: Observed loss (errors) for the generator network of the Conditional GAN during training.




Figure 6: Observed loss (errors) for the discriminator network of the Conditional GAN during training. Loss is calculated for recognition of real and fake images separately.



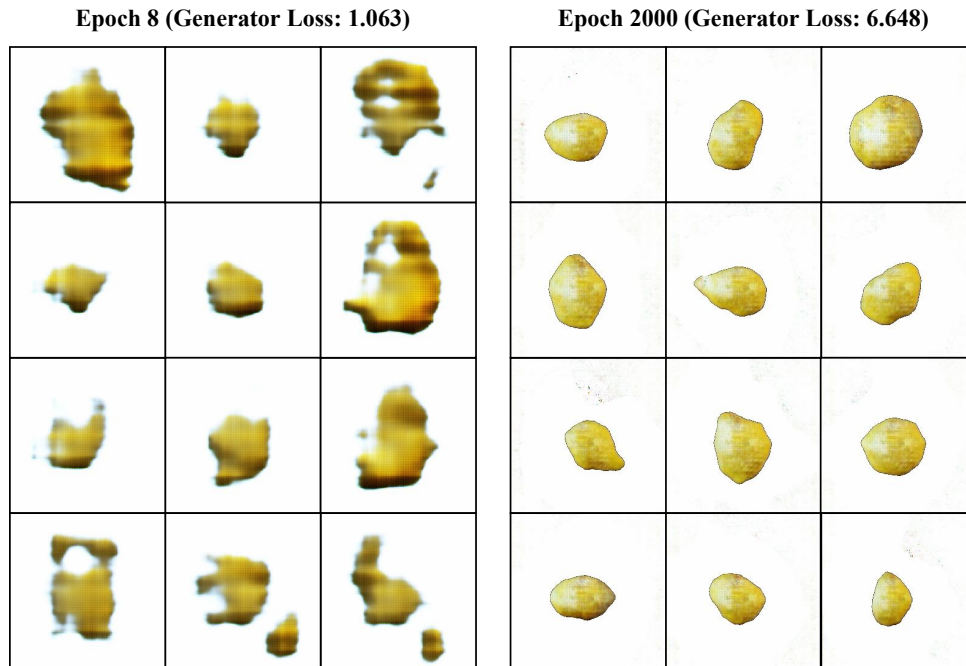


Figure 7: Examples of generator outputs at the lowest observed loss compared to the final epoch.

500 epochs before becoming relatively stabilised. The two discriminator network losses were lower throughout these first 500 epochs before showing an oscillatory nature during the remainder of the learning process.

At the end of the final epoch, the discriminator losses were 0.013 for the real images and 0.005 for the fake images produced by the generator. The generator loss was 6.648. Indeed, this value of 6.648 is by far not the lowest observed, but it is important to consider the nature of GANs; the losses of the two networks are relative to one another, i.e. it is an adversarial score. To provide an example of this, Figure 7 shows a comparison of the images produced by the generator with the lowest loss (epoch 8 at 1.063) and then the generator at the final epoch. Evidently, the images produced by the final generator are of much higher quality than those output by the generator when it experiences the lowest observed loss. For this reason, the final generator is selected as the synthetic data producing model in this work, although it is suggested in the future to explore the quality at multiple stages. Figure 8 shows 18 real photographs for healthy and unhealthy classed lemons, as well as 18 examples of Conditional GAN generator output for healthy and unhealthy classed lemons. Interestingly, many of the synthetic images seem to be more reminiscent of potatoes than lemons; given the nature of GANs, this is due to the model's generalisation of the dataset which contained lemons photographed at different angles - and so this generalisation of a shape is reflected in the produced images. A more uniform colour can be observed in healthy synthetic lemons, whereas unhealthy lemons are given mould and dark styles, as well as several instances of gangrene. Similarly to the potato-like shape of the outputs, this could be attributed to the generalisation nature of GANs, several patterns have been observed during training, and these patterns are then applied while generating new images. A higher resolution example of synthetic lemon images showing features of mould and gangrene can be observed in Figure 9. The generator has seemingly learnt to cast light and shadows on the fruit as well as undesirable characteristics; both fruits have a texture and colour similar to mould on the surface, and the second fruit seems to have a darker patch towards the top.

Figure 10 shows class activation maps on six synthetic images (three per class) from the convolutional neural network trained **only** on the real data. All six images were predicted to belong to their ground-truth classes. Note that on the bottom row the Convolutional Neural Network focuses on issues on the flesh of the fruit such as mould in the first two images and a dark patch which may indicate gangrene on the right-most bottom image. Additionally, the class activation maps for the healthy fruit exist more generally around the shape. These behaviours are seemingly reminiscent of how a human would analyse the images, either focusing on unhealthy characteristics if the fruit is bad or observing the general image when the fruit shows no undesirable features.

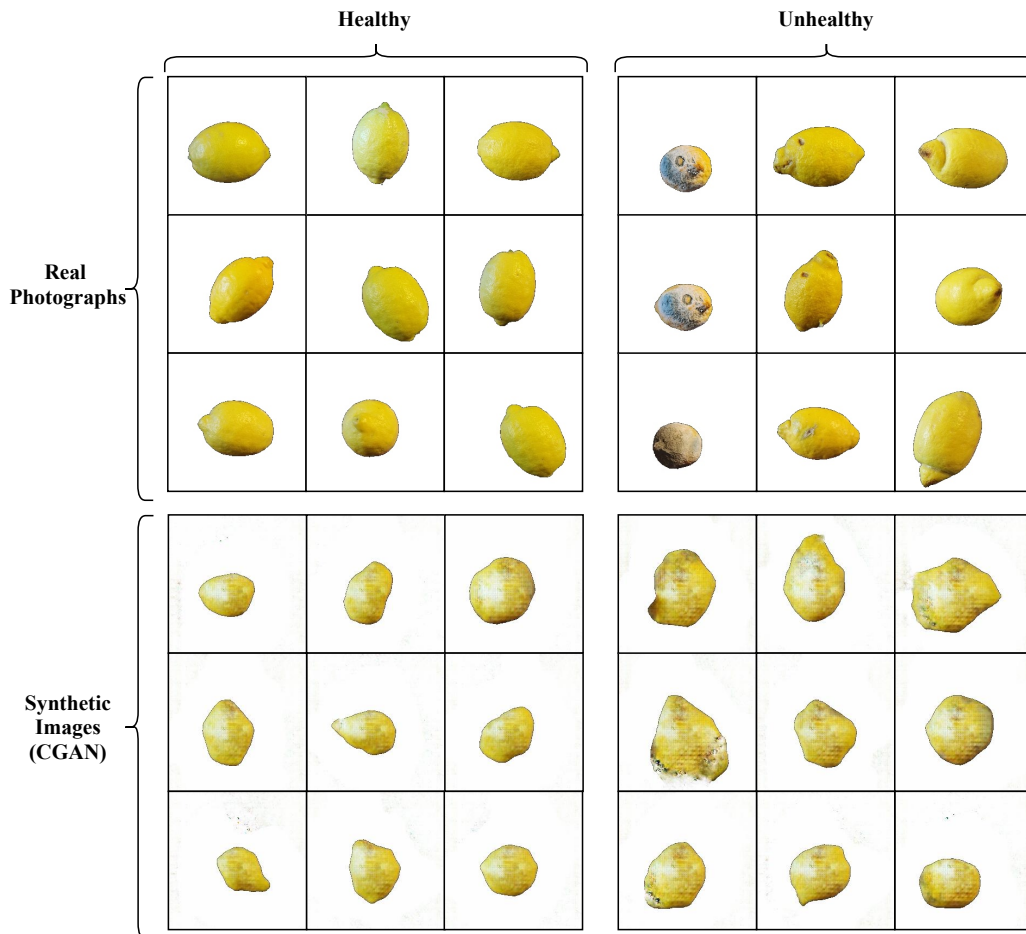


Figure 8: Examples of both real photographs and Conditional GAN outputs (Epoch 2000) for the two classes of healthy and unhealthy lemons.




Figure 9: Two synthetic images of lemons belonging to the “unhealthy” class generated by the Conditional GAN at full resolution. The lemons are seemingly presenting with symptoms of both mould and gangrene, showing that the generator model has learnt to generate fruit with undesirable characteristics. An unrealistic ‘checker-boarding’ effect has also been generated by the Conditional GAN on the surface of the fruit.




Figure 10: Grad-CAM analysis of six outputs from the Conditional GAN. Top row shows images belonging to the “healthy” class and bottom row shows images belonging to the “unhealthy” class. This Grad-CAM VGG16 CNN is trained only on real photographs, and no training has been performed on synthetic images.

Table 2: Comparison of non-augmented (first row) to training data augmentation with Conditional GAN-generated images.

Synthetic Images Augmented (per Class)	Synthetic Data (Total)	Classification Accuracy (%)
0	0	83.77
100	200	87.25
200	400	<b>88.75</b>
300	600	87.63
400	800	87.14
500	1000	87.51
600	1200	87.75
700	1400	85.64
800	1600	86.51
900	1800	87.01
1000	2000	87.63
1100	2200	85.02
1200	2400	86.88
1300	2600	85.89
1400	2800	85.52
1500	3000	87.63

This analysis further shows that both desirable and undesirable characteristics are generalised, and, to an extent reproduced by the generative model. Indeed, the synthetic images are not perfect (as can be observed when comparing them via Figure 8), but these activation maps provide insight into the useful synthetic knowledge existing within them.

### 4.3 Classification Comparison

The results from the best CNN experiment along with the augmentation approaches can be observed in Figure 11 and Table 2. All augmentation approaches scored higher than training only on the real images, showing that augmentation has had a positive effect on the learning process. The best set of results were achieved augmenting the dataset with 400 images (13.51% of the whole dataset, 200 per class), which scored a classification accuracy of 88.75% when classifying unseen images of both healthy and unhealthy lemons.

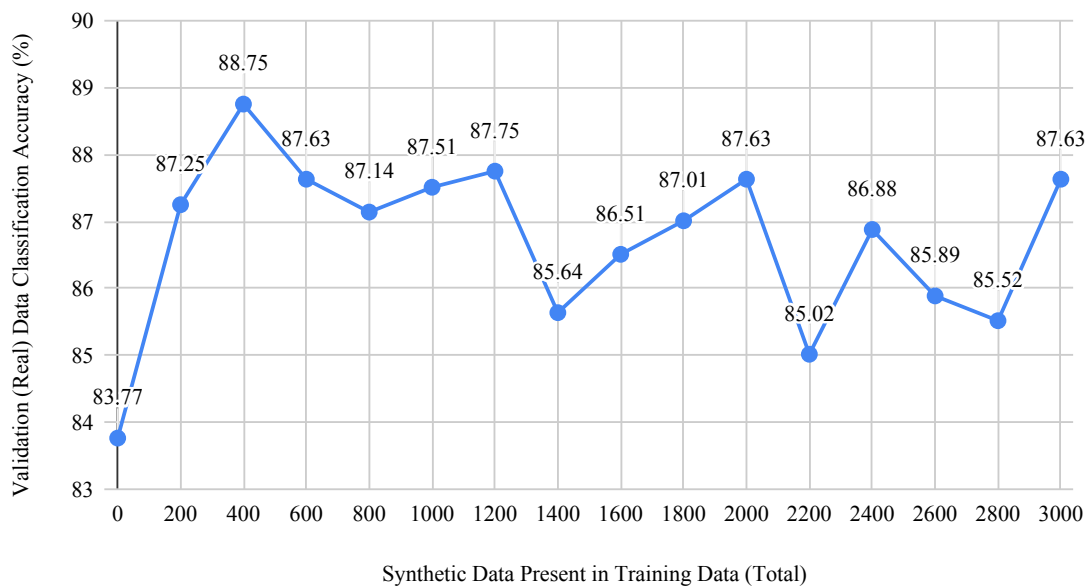


Figure 11: Comparison of how the vanilla CNN (no augmentation) performs against the models which have additional synthetic training data present.

Table 3: Comparison of models when pruning with polynomial decay (for smaller model sizes) for both the vanilla and data augmentation approaches.

Pruned model size (% of original)	Final Sparsity	Post-pruning Classification Accuracy (%)	
		Vanilla	Augmented (200 synthetic images)
10	0.9	60.1	60.1
20	0.8	73.85	73.98
30	0.7	75.71	79.18
40	0.6	77.45	78.31
50	0.5	80.79	81.16
60	0.4	81.54	82.16
70	0.3	82.16	82.65
80	0.2	83.64	84.76
90	0.1	82.9	83.51
100	0	83.77	88.75

Although there was a varying number of classification accuracies recorded, note that even the weakest augmentation approach (2200 images) caused the image recognition ability to rise from 83.77% to 85.02%. That is, of the 15 trials performed, all augmentation approaches outperformed the vanilla CNN. These results thus argue that Conditional GAN-based training data augmentation is a promising approach to improve fruit quality image classification. The generator model and weights are made publicly available for further exploration.

#### 4.4 Pruning

Table 3 and Figure 12 show a comparison of the pruning experiments from 90% to 10% compression (through sparsity *i.e.* zeroed weights). Although the final two models with all neurons present the best results, we note earlier that a faster model would be ideal for the problem at hand, since fruits are processed in large quantities and the models might need being ran under time and energy restrictions. It is interesting to note from these results that the data augmented network tends to outperform the vanilla network in all cases. It is also interesting to note that a data augmentation network compressed to half of its original size can still perform at an accuracy of 81.16%, effectively doubling production at a loss of 7.59% accuracy.

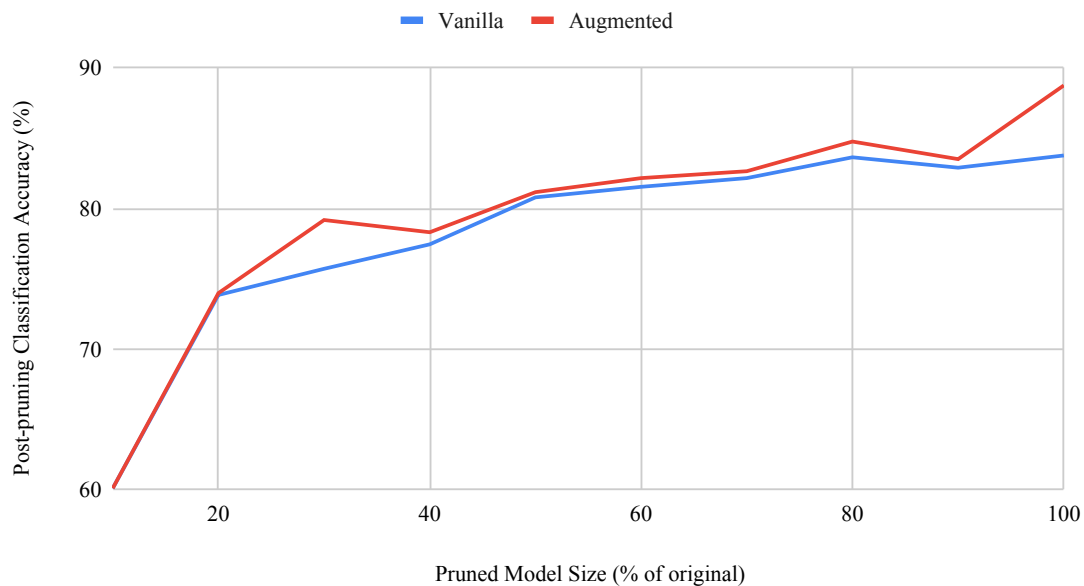


Figure 12: Comparison of the polynomial decay pruning accuracies for both vanilla and data augmentation models.

## 5 Conclusion and Future Work

Lemon fruit are of growing financial importance in LDC nations, in some cases the percentage annual increase of exports were measured in the thousands. The possibility of increasing productivity and reducing the operating costs of those systems are thus particularly interesting. This work has suggested methods to both improve autonomous fruit quality recognition, followed by linear pruning to reduce the complexity of the computer vision model.

To finally conclude, we first noted that autonomous sorting of healthy fruit from undesirable fruit is possible through contemporary computer vision technologies. We explored the concepts of fine-tuning, transfer learning, and Conditional GAN-based training data augmentation; our results showed that the recognition (and thus sorting) of fruit images was improved when augmentation was introduced. We found that introducing 400 synthetic data points had the largest impact, raising the recognition accuracy from 83.77% to 88.75% via a convolutional neural network. Finally, we performed Grad-CAM analysis on the model trained only on real photographs to show that the Conditional GAN was successful in imagining the undesirable characteristics of the flesh of the synthetic fruit generated. Thus, this work argues that Conditional Generative Adversarial Networks have the ability to produce new data to alleviate issues of data scarcity in the problem of fruit health classification.

Although we explored the lowest loss of the generator compared to the final loss, and a vast improvement was made albeit with the said higher loss, future further exploration of model weights throughout the training process may be useful to explore. Given the behaviour of losses spiking, several of the later models should be explored in order to ascertain whether the final epoch did produce the best model for this approach. Although the Conditional GAN-based approach showed promise to the classical CNN, one limitation of this work is that we compare the best train-test CNN to the Conditional GAN; in future, given resource availability to quickly train several Conditional GANs, a better metric for evaluation would be to train  $k$  Conditional GANs during  $k$ -fold cross validation. This would allow for better scientific accuracy. Note that the Conditional GAN was trained for 17 hours on a leading GPU, and thus such a solution would only likely be viable in the future with better technology.

Following the augmentation approach, the classification accuracy rose from 83.77% to 88.75%. In future work, further exploration into the classification model should be performed to further increase this level of recognition accuracy. In addition to further improvement of binary classification, future work could also discern the type of unhealthy attribute, for example, the specificity between damaged and rotten fruit, which directs the fruit to the correct destination; in this case, a juicer or disposal, respectively.

Since the lemon dataset is relatively new, there are no published works to directly compare our results to. In the future, once the dataset has been explored by other works, it would be useful to have a direct comparison of all results when attempting to classify the data.

Pruning showed that the models could be presented at a fraction of their original size and only lose a small amount of classification ability. With this in mind, other pruning techniques could be explored in the future to further maximise the model's ability in terms of real-world usage.

Some future work is also needed for the real-world use of this approach. The nature of the dataset that these models are trained with dictate that image segmentation is required in order to separate the fruit from the background (such as a surface or other fruit). If a segmentation network is trained and applied prior to the preprocessing described in these techniques, then the approach can be tested on more data in a much higher volume, i.e., fruit production.

## 6 Model and Code Availability

For replicability purposes as well as future research, the generator model and synthetic image generation Python code have been made available at: <https://github.com/jordan-bird/synthetic-fruit-image-generator>.

## References

- [1] H. Mureşan and M. Oltean, "Fruit recognition from images using deep learning," *Acta Universitatis Sapientiae, Informatica*, vol. 10, no. 1, pp. 26–42, 2018.
- [2] Food and Agriculture Organization of the United Nations, "FAOSTAT: Crops.." <http://www.fao.org/faostat/en/#data/QC>, February 2021. Accessed: 16th March 2021.
- [3] M. Khojastehnazhand, V. Mohammadi, and S. Minaei, "Maturity detection and volume estimation of apricot using image processing technique," *Scientia Horticulturae*, vol. 251, pp. 247–251, 2019.
- [4] A. Jahanbakhshi, M. Momeny, M. Mahmoudi, and Y.-D. Zhang, "Classification of sour lemons based on apparent defects using stochastic pooling mechanism in deep convolutional neural networks," *Scientia Horticulturae*, vol. 263, p. 109133, 2020.
- [5] A. Jahanbakhshi and K. Kheiralipour, "Evaluation of image processing technique and discriminant analysis methods in postharvest processing of carrot fruit," *Food Science & Nutrition*, vol. 8, no. 7, pp. 3346–3352, 2020.
- [6] Q. Xiang, X. Wang, R. Li, G. Zhang, J. Lai, and Q. Hu, "Fruit image classification based on mobilenetv2 with transfer learning technique," in *Proceedings of the 3rd International Conference on Computer Science and Application Engineering*, pp. 1–7, 2019.
- [7] R. Siddiqi, "Effectiveness of transfer learning and fine tuning in automated fruit image classification," in *Proceedings of the 2019 3rd International Conference on Deep Learning Technologies*, pp. 91–100, 2019.
- [8] S. K. Behera, A. K. Rath, and P. K. Sethy, "Maturity status classification of papaya fruits based on machine learning and transfer learning approach," *Information Processing in Agriculture*, 2020.
- [9] Y.-C. Chou, C.-J. Kuo, T.-T. Chen, G.-J. Horng, M.-Y. Pai, M.-E. Wu, Y.-C. Lin, M.-H. Hung, W.-T. Su, Y.-C. Chen, *et al.*, "Deep-learning-based defective bean inspection with gan-structured automated labeled data augmentation in coffee industry," *Applied Sciences*, vol. 9, no. 19, p. 4166, 2019.
- [10] A. Abbas, S. Jain, M. Gour, and S. Vankudothu, "Tomato plant disease detection using transfer learning with c-gan synthetic images," *Computers and Electronics in Agriculture*, vol. 187, p. 106279, 2021.
- [11] Central Intelligence Agency, "The world factbook field listing: Exports – commodities." <https://www.cia.gov/the-world-factbook/field/exports-commodities>, 2020. Accessed: 16th March 2021.
- [12] N. Vetrekar, R. Gad, I. Fernandes, J. Parab, A. Desai, J. Pawar, G. Naik, and S. Umopathy, "Non-invasive hyperspectral imaging approach for fruit quality control application and classification: case study of apple, chikoo, guava fruits," *Journal of Food Science and Technology*, vol. 52, no. 11, pp. 6978–6989, 2015.
- [13] J. B. Chang and V. Subramanian, "Electronic noses sniff success," *IEEE Spectrum*, vol. 45, no. 3, pp. 50–56, 2008.
- [14] C. J. García-Orellana, M. Macías-Macías, H. M. González-Velasco, A. García-Manso, and R. Gallardo-Caballero, "Low-power and low-cost environmental iot electronic nose using initial action period measurements," *Sensors*, vol. 19, no. 14, p. 3183, 2019.
- [15] J. Brezmes, E. Llobet, X. Vilanova, J. Orts, G. Saiz, and X. Correig, "Correlation between electronic nose signals and fruit quality indicators on shelf-life measurements with pink lady apples," *Sensors and Actuators B: Chemical*, vol. 80, no. 1, pp. 41–50, 2001.

- [16] C. Di Natale, A. Macagnano, E. Martinelli, R. Paolesse, E. Proietti, and A. D'Amico, "The evaluation of quality of post-harvest oranges and apples by means of an electronic nose," *Sensors and Actuators B: Chemical*, vol. 78, no. 1-3, pp. 26–31, 2001.
- [17] J. Brezmes, M. L. Fructuoso, E. Llobet, X. Vilanova, I. Recasens, J. Orts, G. Saiz, and X. Correig, "Evaluation of an electronic nose to assess fruit ripeness," *IEEE Sensors Journal*, vol. 5, no. 1, pp. 97–108, 2005.
- [18] H. M. Zawbaa, M. Hazman, M. Abbass, and A. E. Hassaniien, "Automatic fruit classification using random forest algorithm," in *2014 14th International Conference on Hybrid Intelligent Systems*, pp. 164–168, IEEE, 2014.
- [19] S. R. Dubey and A. S. Jalal, "Application of image processing in fruit and vegetable analysis: A review," *Journal of Intelligent Systems*, vol. 24, no. 4, pp. 405–424, 2015.
- [20] S. Naik and B. Patel, "Machine vision based fruit classification and grading-a review," *International Journal of Computer Applications*, vol. 170, no. 9, pp. 22–34, 2017.
- [21] K. Hameed, D. Chai, and A. Rassau, "A comprehensive review of fruit and vegetable classification techniques," *Image and Vision Computing*, vol. 80, pp. 24–44, 2018.
- [22] J. Naranjo-Torres, M. Mora, R. Hernández-García, R. J. Barrientos, C. Fredes, and A. Valenzuela, "A review of convolutional neural network applied to fruit image processing," *Applied Sciences*, vol. 10, no. 10, p. 3443, 2020.
- [23] A. Bhargava and A. Bansal, "Fruits and vegetables quality evaluation using computer vision: A review," *Journal of King Saud University-Computer and Information Sciences*, 2018.
- [24] K. Yamamoto, S. Ninomiya, Y. Kimura, A. Hashimoto, Y. Yoshioka, and T. Kameoka, "Strawberry cultivar identification and quality evaluation on the basis of multiple fruit appearance features," *Computers and Electronics in Agriculture*, vol. 110, pp. 233–240, 2015.
- [25] G. Capizzi, G. L. Sciuto, C. Napoli, E. Tramontana, and M. Woźniak, "Automatic classification of fruit defects based on co-occurrence matrix and neural networks," in *2015 Federated Conference on Computer Science and Information Systems (FedCSIS)*, pp. 861–867, IEEE, 2015.
- [26] L. M. Azizah, S. F. Umayah, S. Riyadi, C. Damarjati, and N. A. Utama, "Deep learning implementation using convolutional neural network in mangosteen surface defect detection," in *2017 7th IEEE International Conference on Control System, Computing and Engineering (ICCSCE)*, pp. 242–246, IEEE, 2017.
- [27] S. Fan, J. Li, Y. Zhang, X. Tian, Q. Wang, X. He, C. Zhang, and W. Huang, "On line detection of defective apples using computer vision system combined with deep learning methods," *Journal of Food Engineering*, vol. 286, p. 110102, 2020.
- [28] Y. Osako, H. Yamane, S.-Y. Lin, P.-A. Chen, and R. Tao, "Cultivar discrimination of litchi fruit images using deep learning," *Scientia Horticulturae*, vol. 269, p. 109360, 2020.
- [29] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *arXiv preprint arXiv:1409.1556*, 2014.
- [30] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, "Grad-cam: Visual explanations from deep networks via gradient-based localization," in *Proceedings of the IEEE international conference on computer vision*, pp. 618–626, 2017.
- [31] J. Zhang, *Towards Robust Machine Learning Models for Data Scarcity*. PhD thesis, Arizona State University, 2020.
- [32] C. Shorten and T. M. Khoshgoftaar, "A survey on image data augmentation for deep learning," *Journal of Big Data*, vol. 6, no. 1, pp. 1–48, 2019.
- [33] M. D. Bloice, C. Stocker, and A. Holzinger, "Augmentor: an image augmentation library for machine learning," *arXiv preprint arXiv:1708.04680*, 2017.
- [34] N. Andriyanov and D. Andriyanov, "The using of data augmentation in machine learning in image processing tasks in the face of data scarcity," in *Journal of Physics: Conference Series*, vol. 1661, p. 012018, IOP Publishing, 2020.
- [35] J. Nalepa, M. Marcinkiewicz, and M. Kawulok, "Data augmentation for brain-tumor segmentation: a review," *Frontiers in computational neuroscience*, vol. 13, p. 83, 2019.
- [36] N.-T. Tran, V.-H. Tran, N.-B. Nguyen, T.-K. Nguyen, and N.-M. Cheung, "On data augmentation for gan training," *IEEE Transactions on Image Processing*, vol. 30, pp. 1882–1897, 2021.
- [37] R. Anicet Zanini and E. Luna Colombini, "Parkinson's disease emg data augmentation and simulation with dcgans and style transfer," *Sensors*, vol. 20, no. 9, p. 2605, 2020.

- [38] J. J. Bird, M. G. Pritchard, A. Fratini, A. Ekart, and D. Faria, "Synthetic biological signals machine-generated by gpt-2 improve the classification of eeg and emg through data augmentation," *IEEE Robotics and Automation Letters*, 2021.
- [39] R. B. Arantes, G. Vogiatzis, and D. R. Faria, "Csc-gan: Cycle and semantic consistency for dataset augmentation," in *International Symposium on Visual Computing*, pp. 170–181, Springer, 2020.
- [40] Y. Qian, H. Hu, and T. Tan, "Data augmentation using generative adversarial networks for robust speech recognition," *Speech Communication*, vol. 114, pp. 1–9, 2019.
- [41] J. J. Bird, D. R. Faria, C. Premebida, A. Ekárt, and P. P. Ayrosa, "Overcoming data scarcity in speaker identification: Dataset augmentation with synthetic mfccs via character-level rnn," in *2020 IEEE International Conference on Autonomous Robot Systems and Competitions (ICARSC)*, pp. 146–151, IEEE, 2020.
- [42] I. Alnujaim, D. Oh, and Y. Kim, "Generative adversarial networks to augment micro-doppler signatures for the classification of human activity," in *IGARSS 2019-2019 IEEE International Geoscience and Remote Sensing Symposium*, pp. 9459–9461, IEEE, 2019.
- [43] B. Erol, S. Z. Gurbuz, and M. G. Amin, "Gan-based synthetic radar micro-doppler augmentations for improved human activity recognition," in *2019 IEEE Radar Conference (RadarConf)*, pp. 1–5, IEEE, 2019.
- [44] M. Frid-Adar, I. Diamant, E. Klang, M. Amitai, J. Goldberger, and H. Greenspan, "Gan-based synthetic medical image augmentation for increased cnn performance in liver lesion classification," *Neurocomputing*, vol. 321, pp. 321–331, 2018.
- [45] C. Han, L. Rundo, R. Araki, Y. Nagano, Y. Furukawa, G. Mauri, H. Nakayama, and H. Hayashi, "Combining noise-to-image and image-to-image gans: Brain mr image augmentation for tumor detection," *IEEE Access*, vol. 7, pp. 156966–156977, 2019.
- [46] M. B. Lee, Y. H. Kim, and K. R. Park, "Conditional generative adversarial network-based data augmentation for enhancement of iris recognition accuracy," *IEEE Access*, vol. 7, pp. 122134–122152, 2019.
- [47] M. Loey, G. Manogaran, and N. E. M. Khalifa, "A deep transfer learning model with classical data augmentation and cgan to detect covid-19 from chest ct radiography digital images," *Neural Computing and Applications*, pp. 1–13, 2020.
- [48] Y. Fu, X. Li, and Y. Ye, "A multi-task learning model with adversarial data augmentation for classification of fine-grained images," *Neurocomputing*, vol. 377, pp. 122–129, 2020.
- [49] I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 2, NIPS'14*, (Cambridge, MA, USA), p. 2672–2680, MIT Press, 2014.
- [50] M. Mirza and S. Osindero, "Conditional generative adversarial nets," *arXiv preprint arXiv:1411.1784*, 2014.
- [51] M. Adamiak, "Lemons quality control dataset." <https://github.com/softwaremill/lemon-dataset>, July 2020. Accessed: 16th March 2021.
- [52] A. L. Maas, A. Y. Hannun, and A. Y. Ng, "Rectifier nonlinearities improve neural network acoustic models," in *Proc. icml*, vol. 30, p. 3, Citeseer, 2013.
- [53] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [54] A. Radford, L. Metz, and S. Chintala, "Unsupervised representation learning with deep convolutional generative adversarial networks," *arXiv preprint arXiv:1511.06434*, 2015.
- [55] D. T. Huff, A. J. Weisman, and R. Jeraj, "Interpretation and visualization techniques for deep learning models in medical imaging," *Physics in Medicine & Biology*, vol. 66, no. 4, p. 04TR01, 2021.
- [56] P. Molchanov, A. Mallya, S. Tyree, I. Frosio, and J. Kautz, "Importance estimation for neural network pruning," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11264–11272, 2019.
- [57] A. N. Fountsop, J. L. Ebongue Kedieng Fendji, and M. Atemkeng, "Deep learning models compression for agricultural plants," *Applied Sciences*, vol. 10, no. 19, p. 6866, 2020.