CARPATHIAN JOURNAL OF FOOD SCIENCE AND TECHNOLOGY

journal homepage: **http://chimie-biologie.ubm.ro/carpathian_journal/index.html**

ADVERSARIAL AUTOENCODERS FOR AGRICULTURE YIELD FORECASTING

Yueyang Symus Say[1], Mark Wong Kei-Fong[1], Eddie Ng Yin-Kwee [1]✉

School of Mechanical and Aerospace Engineering, Nanyang Technological University, Singapore 639798, Singapore

✉*MYKNG@ntu.edu.sg*

https://doi.org/10.34302/crpjfst/2022.14.3.9

Article history:
Received:
6 July 2022
Accepted:
12 August 2022
Published
September 2022

Keywords:
Agriculture;
Adversarial Autoencoder;
Crop yield forecasting;
Machine Learning;
Deep learning;
Machine Vision.

ABSTRACT

For sustainable food production. In agriculture, crop yields are increasingly affected by warmer temperatures, and pest infestations caused by climate change have increased agricultural losses. Increasing local production is important to reduce our dependence on imported food and provide a buffer in case of supply disruptions such as those caused by the COVID-19 pandemic. To increase food security, it is important to optimize agricultural yields, despite the high costs associated with factors such as supplemental feeding, pest control measures, or operating costs.

We present a Machine Vision method (MV) with Adversarial Autoencoder (AAE) as an approach to crop yield optimization. Predicted leaf area is projected from initial germination to early vegetative stages. Generative machine learning models are analyzed to determine a suitable architecture for crop yield prediction. Images of romaine lettuce grown over time under different conditions (e.g., light intensity) are used as the data set. Preliminary results show that the model created is able to predict an image with sufficient accuracy based on a single condition. With our method, corrective actions can be taken early, and yields recover from initial below-average values. Further work can be done to extend the model to other conditions such as moisture, strength of available sunlight, or soil nutrient content.

1.Introduction

Climate change has adversely affected crop yields and land suitable for agriculture (Zhang & Cai, 2011). The direct effects of global warming alone are projected to lead to a 2-13% decline in yields of major crops (Wang et al., 2020) Jägermeyr et al. (2021), and yield losses could be more severe than previously thought, putting the old agricultural model under production pressure to meet the demands of the future.

This concern is driven by unsustainable agricultural practices, lower crop yields due to climate change (Schmidhuber & Tubiello, 2007), and increasing scarcity of water and arable land (Downing, 2013; Porter et al., 2017; Rosegrant & Cline, 2003). By 2050, the world's population is projected to grow to 9 billion people, leading most countries to worry about numerous food security challenges, including quality, environmental and climate impacts, and reliable access to food sources sufficient to meet growing global demand (Diouf, 2009; Porter et al., 2017; Wise, 2013). Although demand is steadily increasing, 70% more food needs to be produced by 2050, however at the same time, agriculture's share of global GDP has shrunk to just 3 percent, a third of what it was just a several decades ago.

Food security is an important global issue and will remain so for the foreseeable future, as it is more relevant than ever given the immense global impact of the COVID-19 pandemic (Zurayk, 2020). The zoonotic virus was not only introduced by a food supplier, but now reveals

significant deficiencies in the current food supply system. As a result of border and market closures, the impact of quarantines, and the disruption of trade routes, the pandemic has created an unprecedented threat to food logistics and socioeconomic systems around the world, as well as to the livelihoods of billions of people (Galanakis, 2020; Laborde et al., 2020; Zurayk, 2020). Considering the impact of climate change on pests (Skendžić et al., 2021), maintaining current crop yields could be a colossal task in the future.

Predicting crop yield is a multifaceted problem because yield is determined by a variety of factors that are not limited to factors such as genotypic variation within crops, nutrient levels in their growing media, and weather. Despite the variability in crop yield predictions, it is important that such predictions be made because all crops require time to grow and, consequently, financial and spatial resources must be allocated for their growth. Optimizing yields relative to costs is a major concern for growers.

Early crop growth models such as the Decision Support System for Agrotechnology Transfer (DSSAT) (Jones et al., 2003) and CropSyst (Stockle et al., 1994) are still used for simulating crop growth, especially for simulating the effects of climate change and/or the interactions between genetics (G), environment (E), and management (M). However, these crop growth models require extensive calibration to produce accurate results. In addition, the significant cost of runtime and maintenance limits their use to farmers who have the appropriate knowledge, equipment, and financial resources. As a result, these systems cannot be used by low-income farmers who lack the necessary resources, and the scalability of the models is also limited.

Deep Learning (DL) methods such as Convoluted Neural Networks (CNN) have been proposed as an improved method for crop prediction (Khaki et al., 2020; Sakurai et al., 2019; Sun et al., 2019) and disease and pest detection (de Ocampo & Dadios, 2018; Ferentinos, 2018; Fuentes et al., 2017; Sladojevic et al., 2016; Wallelign et al., 2018). By learning and characterizing the performance of the plants, continuous monitoring will provide new insights to optimize the growth rate of the plants (Al-Shakarji et al., 2017). The proposed solution allows finding correlations between leaf area and biomass, helping to predict plant metrics, including growth rate and leaf area. This knowledge is of particular importance for controlling plant growth parameters in response to context and feedback (Shadrin et al., 2019; Shadrin et al., 2018). In addition, DL has been shown to be able to abstract nonlinear relationships that may not be detected by traditional statistical methods.

Recent developments in DL and its application to crop yield prediction aim to improve the accuracy of current models. Numerous research papers have applied DL to crop yield prediction using remote sensing techniques for data collection and various vegetation indices for yield quantification. Chlingaryan et al. (2018) extensively analyzed the use of Machine Learning (ML) for crop yield prediction and nitrogen estimation. They identified trends in the use of vegetation indices from satellite imagery with backpropagated neural networks for more accurate crop yield prediction and predicted that future applications of ML will be more optimized and focused on specific precision agriculture applications. Kulkarni et al. (2018), attempted time series prediction using recurrent neural networks to predict yields given soil and rainfall conditions.

A neuroevolutionary algorithm based on ML was developed to predict crop yields by providing access to information about trait importance (Kanimozhi & Akila, 2020). Similarly, a genetic algorithm-based approach for crop yield prediction has been identified and shown to outperform traditional neural networks and classical statistical methods in crop yield prediction (Bi & Hu, 2021). A comprehensive literature review of crop yield prediction using DL and remote sensing was also conducted in 2022 (Muruganantham et al., 2022) noting commonalities such as the fact that the most common remote sensing technology is the Moderate-Resolution Imaging

Spectroradiometer (MODIS), common indices used to quantify performance, and ML models used in previous research.

However, this newer research focuses on predicting crop yields at the macro level using overhead imagery and data to analyze yields across different land areas. In land-poor locations, conventional farming methods are costly and inefficient due to limited land, and many local farms have switched to high-yield urban/indoor agriculture, where a greater number of crops can be grown in the same space, using supplemental light to enhance plant growth (Jones, 2018).

In urban agriculture, crop yields can still be affected by variability, such as light conditions, but conventional yield forecasts based on satellite imagery cannot be used to accurately predict growth. Identifying these regions of suboptimal yield and correcting growth rates can improve overall crop yield and optimize the return on current operations. Current methods also do not account for dynamic changes in conditions; they cannot predict a change in conditions during growth. Yields are obtained by harvesting, which would result in an interruption of the plant's growth cycle.

Research Objectives

Hence, our work seeks to:

Utilize generative ML/DL models to visualize plant growth using images captured with simple, low-cost cameras.

Identify if the generative model can forecast a visual change in yield with different growing conditions.

This paper will first introduce past work pertaining to the use of ML in crop yield prediction, before elaborating on the generation of data and introducing the architecture of the generative model used in the work. The results will then be evaluated, and the project summarized thereafter.

2. Materials and methods

2.1. Data Generation

As mentioned in the introduction, most publicly available crop yield prediction datasets were selected for analysis. For this project, these datasets were unsuitable for analysis, so a suitable image dataset had to be created. We chose to create our own imagery using a simple, commonly available camera module so that we did not need expensive hyperspectral or high-resolution camera systems that may not be available to the average farmer. However, network issues limited the establishment of affordable microcontroller-based camera modules. The images were hence instead captured using a common smartphone.

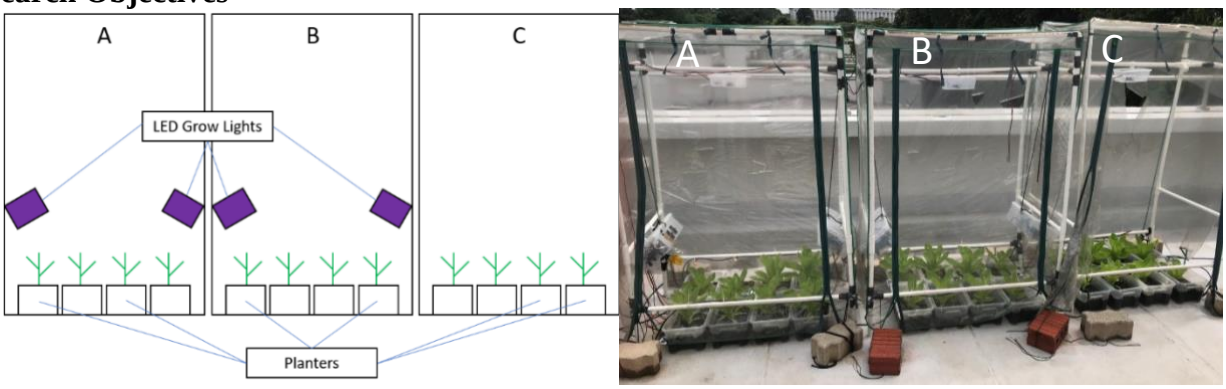

Figure 1. Schematic of experimental setup, each plant was grown in a single rectangular pot for ease of organization, 8 pots were placed in each camera perspective.

Romaine lettuce was selected as the plant for analysis because it has a short harvest period (approximately 1 month) and a large leaf area suitable for this analysis. Three experimental setups A, B, and C, as shown in Fig. 1, were created and placed outdoors under a tarpaulin to create a semi-controlled environment. Experimental setups A and B had LED grow lights attached to them, each of which was turned on 2 hours before sunrise and after sunset, providing a total of 4 hours of daylight extension. For setup C, the front of the tarpaulin was opened to create a semi-protected environment, as a fully protected tarpaulin could be too hot (in tropical regions) and humid for optimal plant growth.

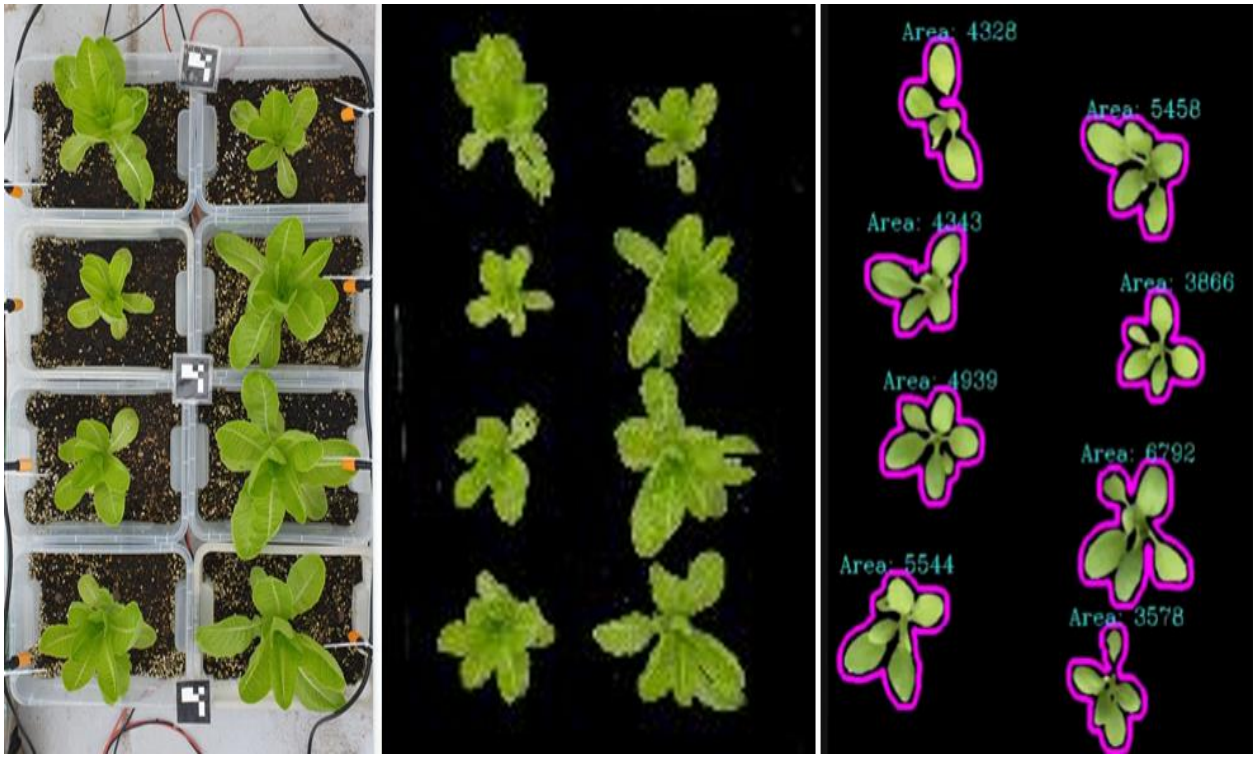

Figure 2. Comparison of raw and processed images

The images captured were then processed using the OpenCV python library (Bradski & Kaehler, 2008) to generate an image mask and resize the image. Furthermore, ArUco markers (Garrido-Jurado et al., 2014) were attempted to be utilized in this step. Due to variations in image capture, the angle, perspective, and size of the captured superstructure may change from image to image. The ArUco markers serve as image anchors, with known dimensions and orientations assigned to each marker. Accordingly, the images can be scaled and repositioned according to the known dimensions and orientations, standardizing the images and improving the accuracy of the dataset. In this step, only the information of the image relevant to the analysis is retained, and the image is converted to an appropriate size so that it can be run through the model in the next step with reasonable computational effort. Figure 2 illustrates the results of the image processing. Three batches of romaine lettuce were created over a harvest cycle of 30-40 days. For testing purposes, images of the lettuce after 34 days were used as a test data set to quantify the performance of the model.

2.2. Architecture

Generative Adversarial Networks (GANs) (Goodfellow et al., 2014) were originally considered, but the limited amount of available data and inability to account for conditions limited the application of the model in this project. Instead, the Adversarial Autoencoder

(AAE) (Makhzani et al., 2015) was identified as an appropriate architecture for the goals of this project. The AAE consists of the encoder, decoder, and discriminator and finds application in the supervised generation of data to be used for this analysis.

The combination of encoder and decoder aims to reconstruct the image and thus obtain a compressed representation of the image in terms of the hidden layer (Z), while the combination of encoder and discriminator aims to map Z to the known distribution.

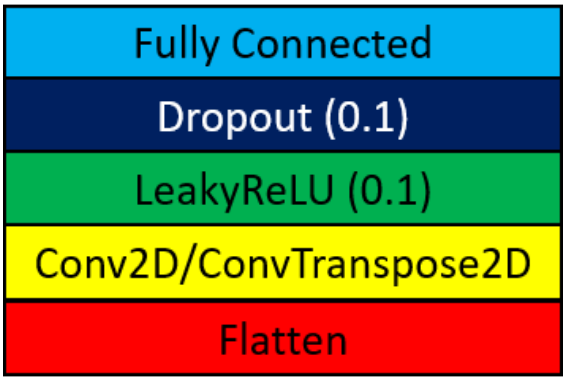

Figure 3. Legend for an Adversarial Autoencoder (AAE)

The PyTorch framework was used to code the model. For the Figures described in the following Sections 2.2.1 to 2.2.3, the layers will be identified as per the legend in Fig. 3, with the values in brackets stating the probability values as required for their respective layers.

2.2.1. Encoder

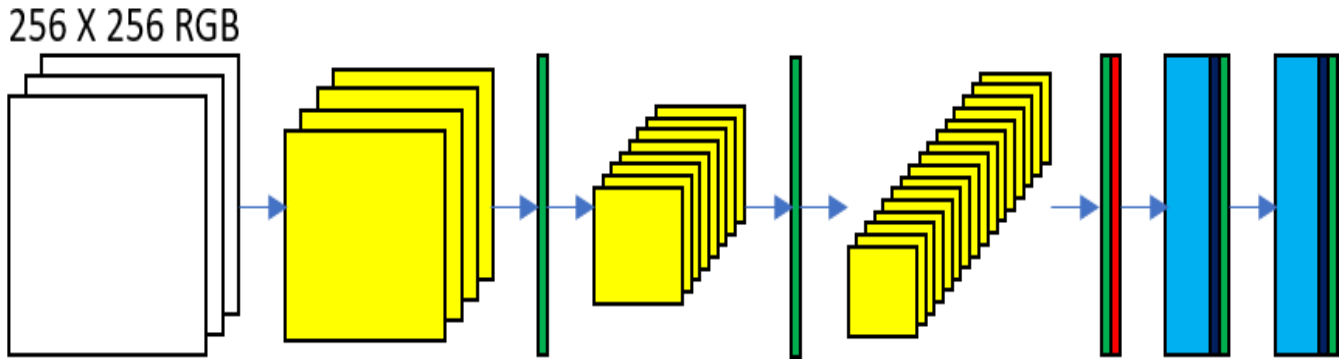

Figure 4. Encoder Architecture

Fig. 4 describes how the layers in the encoder are connected. The original encoder had only Fully Linked (FC) layers with the Rectified Linear Unit (ReLU) activation function. However, for the application in this project, convolutional layers (Lecun et al., 1998) were also used because convolutional layers act as filters to highlight areas of interest in an image. Three Conv2D layers were therefore added to the original encoder architecture before they underwent LeakyReLU activation with a probability of 0.1 and were flattened into a 1-dimensional (1D) array.

The LeakyReLU function was used because the resulting reconstruction loss was less than using ReLU and the fidelity of the generated image was slightly increased. This 1D array was then passed through 2 FC layers, with the output of the second FC layer being Z. Dropout layers with a probability of 0.1 were also added to reduce overfitting of the data set.

2.2.2. Decoder

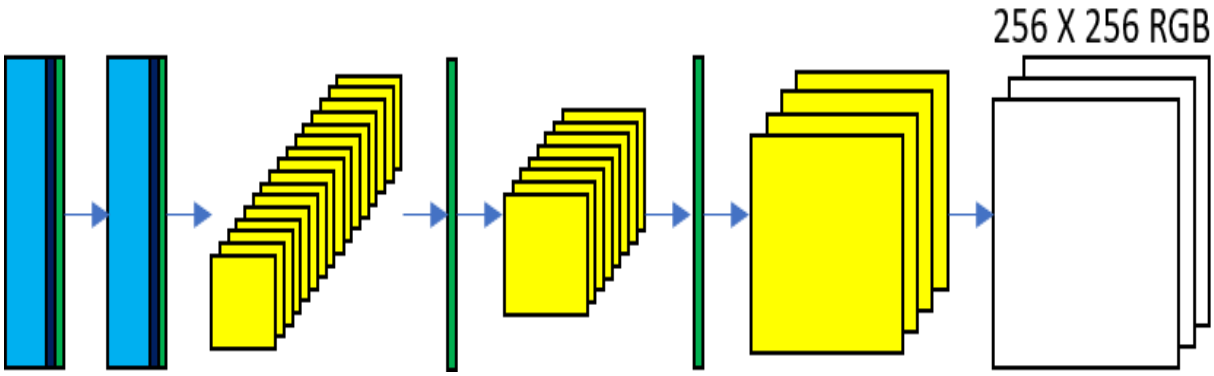

Figure 5. Decoder Architecture

Similar to 2.2.1, the original Decoder only consisted of FC layers and activation functions. The model was modified to include ConvTranspose2D layers, with inspiration from Radford et al. (2016), dropout layers, and the use of the LeakyReLU activation function for the same reasons as described in 2.2.1. The decoder takes Z and an array (Y) containing constraints such as time. Figure 5 describes how the layers are connected in the decoder. The reconstruction loss between the encoder/decoder is then calculated and tracked, and the weights of the encoder and decoder are updated.

2.2.3. Discriminator

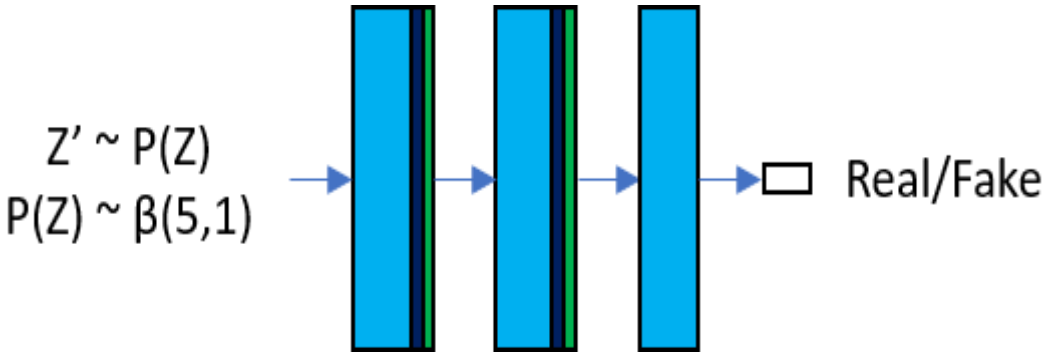

Figure 6. Discriminator Architecture

The discriminator uses 2 FC layers with the LeakyReLU activation function along with dropout layers to determine if a Z-layer sample is from the known distribution. The second FC layer is passed through another FC layer to produce a binary output indicating whether the sample is from the known distribution. The min-max GAN loss introduced by Goodfellow et al. (2014) is then calculated and then backpropagated to update the encoder and discriminator weights. For this project, a beta distribution with parameters (5,1) was used as it was found to give visually similar results to the actual image. Fig. 6 describes part of the architecture of the discriminator.

2.2.4. Combined

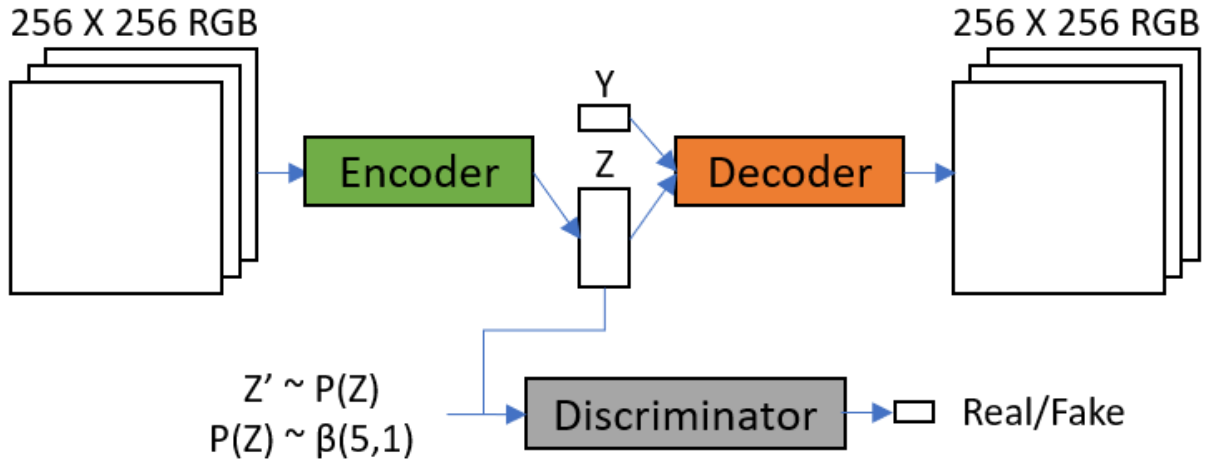

Figure 7. Combined Architecture

The encoder, decoder and discriminator are combined as shown in Fig. 7. The original image is passed through the encoder to obtain Z, an array of size 512, and Y, the array of conditions. The conditions in Y include: the time, a coded array for the degree of protection in the experimental setups, and whether additional light was present in the experimental setups. Y and Z were then input to the decoder, and an image was regenerated as described in 2.2.2.
A sample is taken from Z and Z', where Z' follows a known distribution. The two samples are passed through the discriminator in 2.2.3 and the loss is calculated. The model was run for 100 epochs using the Adam optimizer with a learning rate of $6e^{-4}$ for the generator (encoder/decoder) and a learning rate of $8e^{-4}$ for the regulator (encoder/discriminator).

2.3. Evaluation

The Mean Structural Similarity Index Measure (MSSIM) (Zhou et al., 2004) was used to quantify the fidelity of the generated image.

$$\mu_x = \frac{1}{N}\sum_{i=1}^{N} x_i \quad (1)$$

$$\sigma_x = \left(\frac{1}{N-1}\sum_{i=1}^{N}(x_i - \mu_x)^2\right)^{0.5} \quad (2)$$

$$l(\boldsymbol{x},\boldsymbol{y}) = \frac{2\mu_x\mu_y + C_1}{\mu_x^2 + \mu_y^2 + C_1} \quad (3)$$

$$c(\boldsymbol{x},\boldsymbol{y}) = \frac{2\sigma_x\sigma_y + C_2}{\sigma_x^2 + \sigma_y^2 + C_2} \quad (4)$$

$$s(\boldsymbol{x},\boldsymbol{y}) = \frac{\sigma_{xy} + C_3}{\sigma_x\sigma_y + C_3} \quad (5)$$

$$SSIM(\boldsymbol{x},\boldsymbol{y}) = [l(\boldsymbol{x},\boldsymbol{y})]^{\alpha} * [c(\boldsymbol{x},\boldsymbol{y})]^{\beta} * [s(\boldsymbol{x},\boldsymbol{y})]^{\gamma} \quad (6)$$

The indices x and y refer to the two images to be compared. μ refers to the arithmetic mean and σ to the standard deviation of the analyzed pixels. The luminance, contrast and structure scores are denoted to as *l*, *c* and *s* respectively.

This metric involves the extraction of three key features that contribute to the structure of the image – luminance (1), contrast (2), and structure, as described in Fig. 8. Luminance compares the average intensity of the pixel values, contrast compares the standard deviation of the signal, while structure compares the normalized signal so that the signal has a uniform standard deviation.

A series of comparison functions (3-5) are then created for luminance, contrast and structure to compare the actual and expected signals. The constants C_1, C_2, and C_3 provide numerical stability as the denominator approaches 0. The SSIM value (6) is then calculated based on the three comparison functions, with α, β and γ indicating the relative importance of the luminance, contrast and structure values. A value of 0 indicates complete dissimilarity (the images are completely different) and a value of 1 indicates complete similarity (the images are exactly the same).

A Gaussian weighting function of size 11x11 is introduced to compute the SSIM over localized regions of the image. The local SSIM values are then averaged to obtain the MSSIM values. The PyTorch Image Quality Assessment package (Rozet, 2022) was used to calculate this metric.

3. Results and discussions

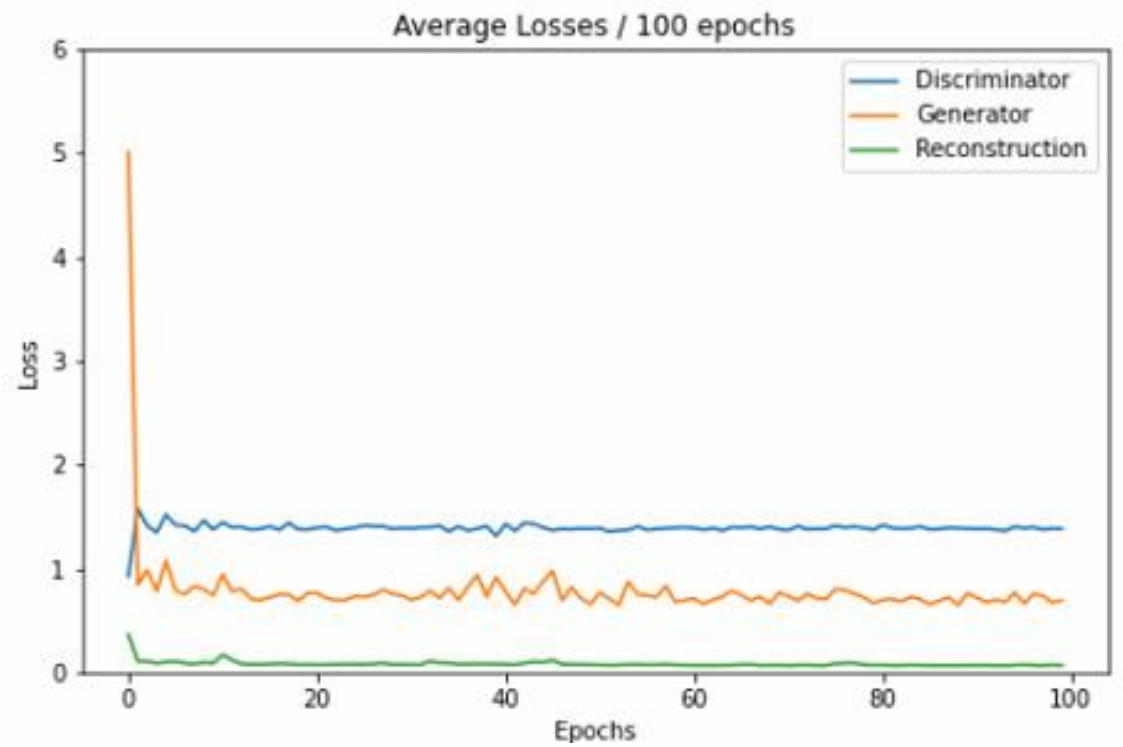

Figure 8. Loss curves for AAE

The loss curves for training the AAE model are shown in Fig. 8. Fig. 9 shows the transition from the real to the generated image for Setup A, Batch 4. It can be seen that the model is able to generate images that follow the trend of the real images.

From Fig. 10, The images produced vary over time and under different application conditions. Comparing the images from day 34, it is noticeable that the image from setting C has a larger green area than the image from setting A.

This difference becomes even more apparent when the same images are compared again on day 38. A more intense green is also observed when the Day 38 and Day 34 images are compared for both settings. The MSSIM scores for the 3 settings were then calculated using the data provided for the test. These scores are tabulated in Table 1.

Figure 9. Transition from actual to generated images, Setup A

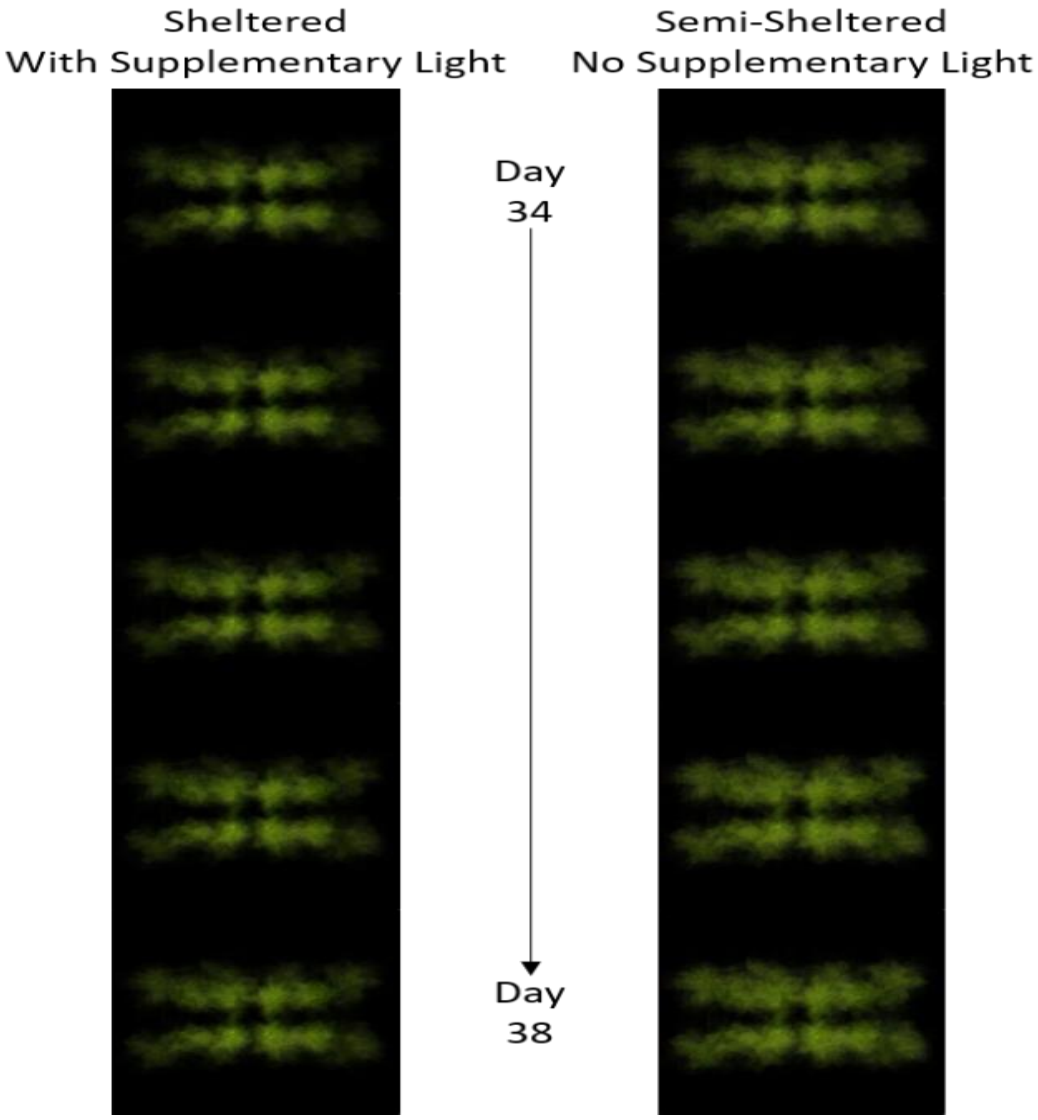

Figure 10. Comparison of generated images between Setup A (left) and C (right)

Table 1. MSSIM Scores, Setups A-C

Setup A						
Day	34	35	36	37	37	38
Score (x10)	5.76	5.76	5.81	5.82	5.73	5.70
Setup B						
Day	34	35	36	37	37	38
Score (x10)	5.83	5.76	5.76	5.67	5.70	5.69
Setup C						
Day	34	35	36	37	37	38
Score (x10)	5.26	5.09	5.24	5.01	5.12	5.16

The results from day 34 to the first instance of day 37 are from the 2nd batch, and the results from the second instance of day 37 to day 38 are from the 3rd batch. Experimental setups A and B are protected experimental setups with supplemental light, and experimental setup C is a semi-protected experimental setup without supplemental light, as described in 2.1.

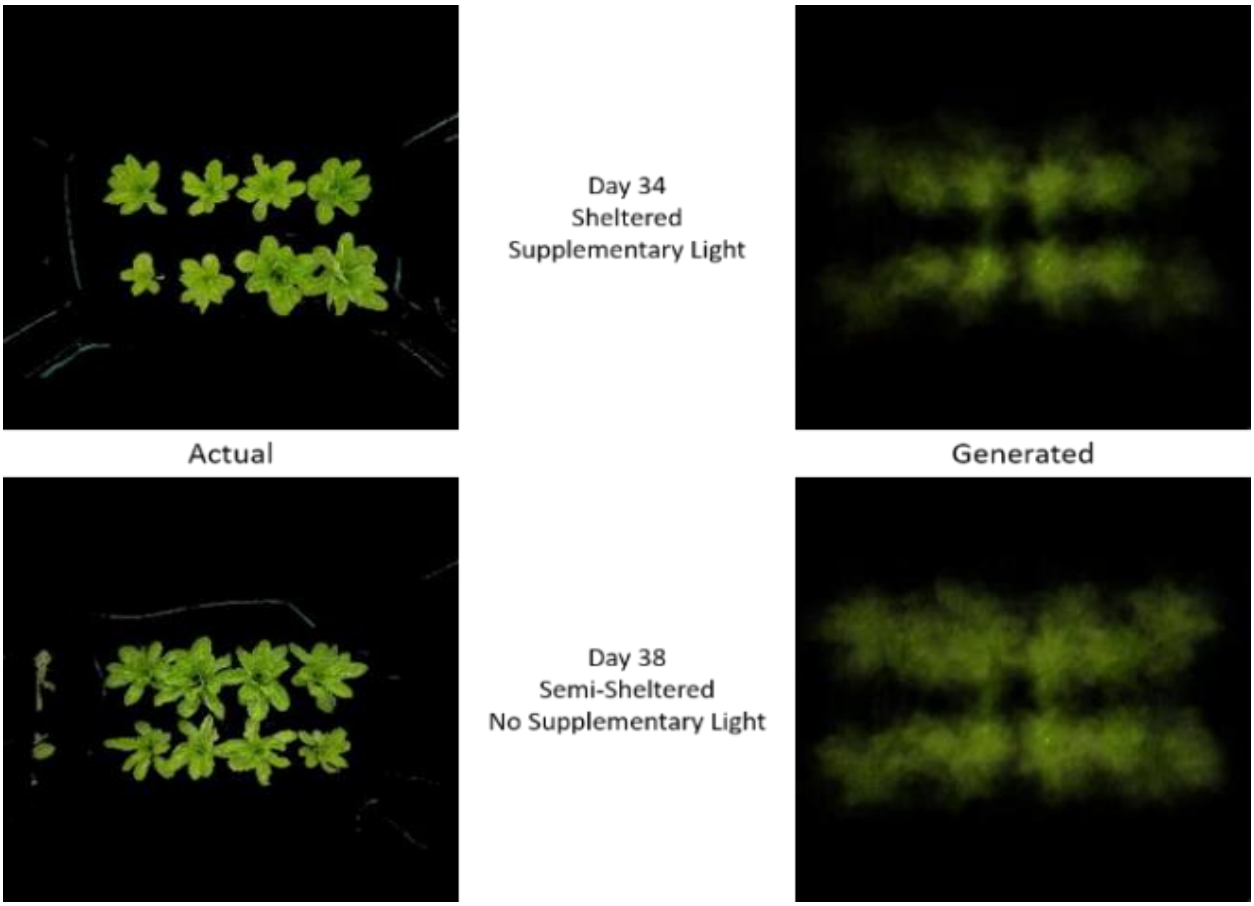

Figure 11. Sample of generated images against actual images

As can be seen from Table 1, the model did not perform poorly, but neither did it excel, with MSSIM values between 0.5 and 0.6, which is confirmed by the generated images, as shown in Fig. 11. Nevertheless, the models are able to capture the differences between configurations despite the limited amount of data. If more and more diverse data were available, the effectiveness of the model could be analyzed in more detail.

4. Conclusions

The AAE model has shown the ability to regenerate harvest images and generate images with different variables (e.g., time). This can help reduce harvest losses because the architecture can be modified to visually predict what the yield will be under changing conditions. With camera smartphones becoming more accessible, farmers can use any smartphone on hand for image capturing – avoiding the need for custom hardware deployment in a commercial setting. Regardless of the source of images (low-cost cameras or smartphones), the architecture can still be applied for crop yield prediction.

However, further work is needed. The data generated is insufficient (only 3 batches could be generated), and there are a number of external variables (e.g., outdoor temperature, solar radiation) that could not be accounted for in this project. Reference markers in the form of ArUco

markers were also used, but the OpenCV library was unable to consistently capture the markers throughout the dataset, possibly due to the reflection of sunlight on the laminated markers affecting the library's ability to read the markers. Therefore, there may be differences in the perspective and positions of the captured images, even though the authors did their best to keep them constant.

Despite the limitations mentioned above, the generation of such images is a testament to the potential of the model. If the model is developed further, the application of such an architecture in the distant future could change the definition of crop yield prediction.

In land-poor locations with the transition to high-tech, high-yield agriculture, the model can be extended with low-cost cameras to create a cost-effective image capture system, and the collected data can be applied to the model to enable a modern application of crop yield prediction given the current state of agriculture in land-poor locations. In this way, crop yields and costs could be optimized.

5. References

Al-Shakarji, N. M., Kassim, Y. M., & Palaniappan, K. (2017). Unsupervised learning method for plant and leaf segmentation. 2017 IEEE Applied Imagery Pattern Recognition Workshop (AIPR),

Bi, L., & Hu, G. (2021). A genetic algorithm-assisted deep learning approach for crop yield prediction. *Soft Computing*, *25*(16), 10617-10628. https://doi.org/10.1007/s00500-021-05995-9

Bradski, G., & Kaehler, A. (2008). *Learning OpenCV: Computer Vision with the OpenCV Library*. O'Reilly Media. https://books.google.com.sg/books?id=seAgiOfu2EIC

Chlingaryan, A., Sukkarieh, S., & Whelan, B. (2018). Machine learning approaches for crop yield prediction and nitrogen status estimation in precision agriculture: A review. *Computers and Electronics in Agriculture*, *151*, 61-69. https://doi.org/https://doi.org/10.1016/j.compag.2018.05.012

de Ocampo, A. L. P., & Dadios, E. P. (2018). Mobile Platform Implementation of Lightweight Neural Network Model for Plant Disease Detection and Recognition. 2018 IEEE 10th International Conference on Humanoid, Nanotechnology, Information Technology, Communication and Control, Environment and Management (HNICEM),

Diouf, J. (2009). How to feed the world in 2050. *Pop. Dev. Rev*, *35*, 837-839.

Downing, T. E. (2013). *Climate change and world food security* (Vol. 37). Springer Science & Business Media.

Ferentinos, K. P. (2018). Deep learning models for plant disease detection and diagnosis. *Computers and Electronics in Agriculture*, *145*, 311-318.

Fuentes, A., Yoon, S., Kim, S. C., & Park, D. S. (2017). A robust deep-learning-based detector for real-time tomato plant diseases and pests recognition. *Sensors*, *17*(9), 2022.

Galanakis, C. M. (2020). The Food Systems in the Era of the Coronavirus (COVID-19) Pandemic Crisis. *Foods*, *9*(4), 523.

Garrido-Jurado, S., Muñoz-Salinas, R., Madrid-Cuevas, F. J., & Marín-Jiménez, M. J. (2014). Automatic generation and detection of highly reliable fiducial markers under occlusion. *Pattern Recognition*, *47*(6), 2280-2292. https://doi.org/https://doi.org/10.1016/j.patcog.2014.01.005

Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., & Bengio, Y. (2014). Generative adversarial nets. *Advances in neural information processing systems*, *27*.

Jägermeyr, J., Müller, C., Ruane, A. C., Elliott, J., Balkovic, J., Castillo, O., Faye, B., Foster, I., Folberth, C., Franke, J. A., Fuchs, K., Guarin, J. R., Heinke, J., Hoogenboom, G., Iizumi, T., Jain, A. K., Kelly, D., Khabarov, N., Lange, S., Rosenzweig, C. (2021). Climate impacts on global agriculture emerge earlier in new generation of climate and crop models. *Nature Food*, *2*(11), 873-

885. https://doi.org/10.1038/s43016-021-00400-y

Jones, J. W., Hoogenboom, G., Porter, C. H., Boote, K. J., Batchelor, W. D., Hunt, L. A., Wilkens, P. W., Singh, U., Gijsman, A. J., & Ritchie, J. T. (2003). The DSSAT cropping system model. *European Journal of Agronomy*, *18*(3), 235-265. https://doi.org/https://doi.org/10.1016/S1161-0301(02)00107-7

Jones, M. A. (2018). Using light to improve commercial value. *Horticulture Research*, *5*(1), 47. https://doi.org/10.1038/s41438-018-0049-7

Kanimozhi, E., & Akila, D. (2020, 2020//). An Empirical Study on Neuroevolutional Algorithm Based on Machine Learning for Crop Yield Prediction. Intelligent Computing and Innovation on Data Science, Singapore.

Khaki, S., Wang, L., & Archontoulis, S. V. (2020). A cnn-rnn framework for crop yield prediction. *Frontiers in Plant Science*, *10*, 1750. https://www.ncbi.nlm.nih.gov/pmc/articles/PMC6993602/pdf/fpls-10-01750.pdf

Kulkarni, S., Mandal, S. N., Sharma, G. S., Mundada, M. R., & Meeradevi. (2018, 19-22 Sept. 2018). Predictive Analysis to Improve Crop Yield using a Neural Network Model. 2018 International Conference on Advances in Computing, Communications and Informatics (ICACCI),

Laborde, D., Martin, W., & Vos, R. (2020). Poverty and food insecurity could grow dramatically as COVID-19 spreads [Blog]. Retrieved 25th April 2020, from https://www.ifpri.org/blog/poverty-and-food-insecurity-could-grow-dramatically-covid-19-spreads

Lecun, Y., Bottou, L., Bengio, Y., & Haffner, P. (1998). Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, *86*(11), 2278-2324. https://doi.org/10.1109/5.726791

Makhzani, A., Shlens, J., Jaitly, N., & Goodfellow, I. J. (2015). Adversarial Autoencoders. *ArXiv, abs/1511.05644*.

Muruganantham, P., Wibowo, S., Grandhi, S., Samrat, N. H., & Islam, N. (2022). A Systematic Literature Review on Crop Yield Prediction with Deep Learning and Remote Sensing. *Remote Sensing*, *14*(9). https://doi.org/10.3390/rs14091990

Porter, J. R., Xie, L., Challinor, A. J., Cochrane, K., Howden, S. M., Iqbal, M. M., Travasso, M., Barros, V., Field, C., & Dokken, D. (2017). Food security and food production systems.

Radford, A., Metz, L., & Chintala, S. (2016). Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks. *CoRR, abs/1511.06434*.

Rosegrant, M. W., & Cline, S. A. (2003). Global food security: challenges and policies. *Science*, *302*(5652), 1917-1919. https://science.sciencemag.org/content/302/5652/1917.long

Rozet, F. (2022). *PyTorch Image Quality Assessement package*. In GitHub repository. https://github.com/francois-rozet/piqa

Sakurai, S., Uchiyama, H., Shimada, A., & Taniguchi, R.-i. (2019). Plant growth prediction using convolutional LSTM. 14th International Conference on Computer Vision Theory and Applications, VISAPP 2019-Part of the 14th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications, VISIGRAPP 2019,

Schmidhuber, J., & Tubiello, F. N. (2007). Global food security under climate change. *Proceedings of the National Academy of Sciences*, *104*(50), 19703-19708. https://www.ncbi.nlm.nih.gov/pmc/articles/PMC2148361/pdf/zpq19703.pdf

Shadrin, D., Menshchikov, A., Somov, A., Bornemann, G., Hauslage, J., & Fedorov, M. (2019). Enabling Precision Agriculture through Embedded Sensing with Artificial Intelligence. *IEEE Transactions on Instrumentation and Measurement*.

Shadrin, D., Somov, A., Podladchikova, T., & Gerzer, R. (2018). Pervasive agriculture: Measuring and predicting plant growth using statistics and 2D/3D imaging. 2018

IEEE International Instrumentation and Measurement Technology Conference (I2MTC),

Skendžić, S., Zovko, M., Živković, I. P., Lešić, V., & Lemić, D. (2021). The Impact of Climate Change on Agricultural Insect Pests. *Insects*, *12*(5), 440. https://doi.org/10.3390/insects12050440

Sladojevic, S., Arsenovic, M., Anderla, A., Culibrk, D., & Stefanovic, D. (2016). Deep neural networks based recognition of plant diseases by leaf image classification. *Computational intelligence and neuroscience, 2016*.

Stockle, C. O., Martin, S. A., & Campbell, G. S. (1994). CropSyst, a cropping systems simulation model: Water/nitrogen budgets and crop yield. *Agricultural Systems*, *46*(3), 335-359. https://doi.org/https://doi.org/10.1016/0308-521X(94)90006-2

Sun, J., Di, L., Sun, Z., Shen, Y., & Lai, Z. (2019). County-Level Soybean Yield Prediction Using Deep CNN-LSTM Model. *Sensors*, *19*(20), 4363.

Wallelign, S., Polceanu, M., & Buche, C. (2018). Soybean plant disease identification using convolutional neural network. The Thirty-First International Flairs Conference,

Wang, X., Zhao, C., Müller, C., Wang, C., Ciais, P., Janssens, I., Peñuelas, J., Asseng, S., Li, T., Elliott, J., Huang, Y., Li, L., & Piao, S. (2020). Emergent constraint on crop yield response to warmer temperature from field experiments. *Nature Sustainability*, *3*(11), 908-916. https://doi.org/10.1038/s41893-020-0569-7

Wise, T. A. (2013). Can we feed the world in 2050. *A scoping paper to assess the evidence. Global Development and Environment Institute Working Paper*(13-04).

Zhang, X., & Cai, X. (2011). Climate change impacts on global agricultural land availability. *Environmental Research Letters*, *6*(1). https://doi.org/10.1088/1748-9326/6/1/014014

Zhou, W., Bovik, A. C., Sheikh, H. R., & Simoncelli, E. P. (2004). Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, *13*(4), 600-612. https://doi.org/10.1109/TIP.2003.819861

Zurayk, R. (2020). Pandemic and Food Security. *Journal of Agriculture, Food Systems, and Community Development*, *9*(3), 1-5.

Acknowledgment and Conflict of Interest

Work presented in this paper is supported by institutional resources provided by Nanyang Technological University of Singapore. We wish to acknowledge the funding support for this project from Nanyang Technological University under the URECA Undergraduate Research Programme. The authors declare there is no conflict of interest.