

# Standardizing and Centralizing Datasets to Enable Efficient Training of Agricultural Deep Learning Models

Amogh Joshi<sup>1,2,3</sup>, Dario Guevara<sup>1,2,3</sup>, and Mason Earles<sup>1,2,3</sup>

<sup>1</sup> Department of Viticulture and Enology, University of California, Davis

<sup>2</sup> Department of Biological and Agricultural Engineering, University of California, Davis

<sup>3</sup> AI Institute for Next-Generation Food Systems (AIFS)


{amnjoshi,dguevara,jmearles}@ucdavis.edu

**Abstract.** In recent years, deep learning models have become the standard for agricultural computer vision. Such models are typically fine-tuned to agricultural tasks using model weights that were originally fit to more general, non-agricultural datasets. This lack of agriculture-specific fine-tuning potentially increases training time and resource use, and decreases model performance, leading an overall decrease in data efficiency. To overcome this limitation, we collect a wide range of existing public datasets for three distinct tasks, standardize them, and construct standard training and evaluation pipelines, providing us with a set of benchmarks and pretrained models. We then conduct a number of experiments using methods which are commonly used in deep learning tasks, but unexplored in their domain-specific applications for agriculture. Our experiments guide us in developing a number of approaches to improve data efficiency when training agricultural deep learning models, without large-scale modifications to existing pipelines. Our results demonstrate that even slight training modifications, such as using agricultural pretrained model weights, or adopting specific spatial augmentations into data processing pipelines, can significantly boost model performance and result in shorter convergence time, saving training resources. Furthermore, we find that even models trained on low-quality annotations can produce comparable levels of performance to their high-quality equivalents, suggesting that datasets with poor annotations can still be used for training, expanding the pool of currently available datasets. Our methods are broadly applicable throughout agricultural deep learning, and present high potential for significant data efficiency improvements.

## 1 Introduction

Deep learning models have become standard for modern computer-vision-based agricultural tasks. Examples of common standard tasks now largely automated by deep learning include fruit detection (Sa et al., 2016; Bargoti and Underwood, 2016), crop and weed segmentation (Jeon et al., 2011; Milioto et al., 2018), and plant disease classification (Mohanty et al., 2016; Ümit Atila et al., 2021). Certain deep learning models have even been used for tasks beyond standard single-RGB-image predictions, involving hyperspectral and thermal imagery for various analyses (Elsherbiny et al., 2021; Wang et al., 2021) or tasks involving multiple scales of data, such as spatiotemporal crop yield prediction (Nevavuori et al., 2020). Existing approaches to such tasks often involve direct applications of state-of-the-art models developed for other tasks in the field of deep learning, which have been thoroughly evaluated by works such as Li et al. (2020).

A largely prevalent problem in the agricultural domain is a significant deficiency of task-specific data. Transfer learning (Pan and Yang, 2010) is an approach undertaken when training large deep learning models which attempts to offset this data deficiency, therein transferring knowledge from a source task to the new, reduced size target dataset. In practice, this generally consists of replacing random weight initialization for model parameters with existing pretrained weights from a prior task, converting a complete training task into a finetuning task using existing pretrained models – for instance, many image classification and object detection tasks start off with pretrained weights from the ImageNet (Deng et al., 2009) and COCO (Lin et al., 2014) datasets, respectively. However, these datasets are mostly generalized to common objects in environments ranging from inside a home to the streets of a city. Such environments consist of a broad range of objects and conditions, allowing models to locate specific objects with relatively greater ease. In contrast, agricultural environments are highly domain-specific, and these existing pretrained model approaches may not offer as significant of a knowledge transfer for agricultural domains. An example of the contrast between the aforementioned conditions is shown in Figure 1. This, in turn, suggests the potential need for alternative pretrained model approaches in order to improve data efficiency when training agricultural deep learning models. Previously, Nowakowski et al. (2021), Moon and Eek Son (2021), and Shrivastava et al. (2021) applied transfer learning to agricultural image



**Fig. 1.** Sample of images from the COCO dataset (row 1) and from assorted agricultural datasets (row 2), showcasing the contrast between the general environment of COCO imagery and the specific environment of agricultural imagery.

classification tasks, demonstrating improved performance using pretrained weights on the ImageNet dataset, citing its large magnitude of data and certain specific classes relevant to agriculture. However, little to no prior work has been conducted assessing the viability of transfer learning in other, more relevant agricultural tasks, like semantic segmentation and object detection. Furthermore, there does not exist a single centralized repository for agriculture-specific datasets, preventing a large-scale ImageNet-style dataset for agriculture from being developed. While work done by [Zheng et al. \(2019\)](#) and [Sahili and Awad \(2022\)](#) have proposed new, larger-scale datasets for agriculture, increasing progress towards such a goal, such datasets are usually specific to certain conditions and designed for image classification, leaving other tasks without as much data. As a result, the potential for transfer learning as a viable improvement for data efficiency, at least in agriculture, is diminished.

Another common approach in deep learning which attempts to offset a deficiency of data is the use of data augmentations. For computer vision tasks, this refers to the transformation of an input image along with its corresponding annotations, with the stated intent of increasing diversity of data. Standard augmentations include rotation or flipping of an input image, or adjustment of visual parameters like saturation or brightness. The usefulness of augmentations have been studied in a general case in works such as [Shorten and Khoshgoftaar \(2019\)](#); however these applications have largely been focused either on general image classification tasks or for a specific domain which is not similar with agriculture. Agricultural environments are largely different than general environments and are often more complex, involving less quantities of obvious features and more irregular and inconsistent shapes ([Wang et al., 2022](#)). Certain studies of augmentations have been conducted in works such as [Moreira et al. \(2022\)](#), where custom colorspace and vegetation indexes were used when processing input data in order to boost performance on agricultural datasets, but there has not been done any large-scale analysis of augmentation effectiveness for agricultural environments. In fact, a potential benefit of the domain constraints of agricultural data is the prospect for assessing features specific to these environments. This can be used to both develop data processing pipelines specific to agriculture, using such approaches, and adapt existing model architectures to agricultural tasks using methods like transfer learning. Standing in the way of this development is a lack of a large-scale centralized database for agricultural data.

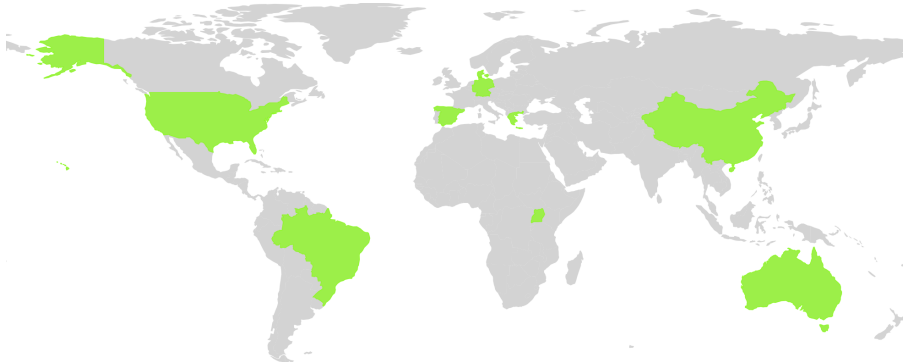
Bringing together datasets from previous studies, we present a novel set of centralized and standardized public agricultural datasets and benchmarks using state-of-the-art deep learning models. This collection of datasets is composed of real data collected from agricultural datasets across three different tasks: image classification, semantic segmentation, and object detection. For each task, we develop a standard data format and training pipeline, which we use to generate benchmarks and pretrained models for each dataset. These pipelines and models enable us to conduct a set of experiments in which we adapt existing methods for improving data efficiency, such as transfer learning and image augmentation, as described above, for agricultural environments. We first conduct an assessment of agricultural models for transfer learning, by (1) assessing data efficiency and

performance when using agricultural pretrained weights for fruit detection models. We apply this study further by (2) assessing data efficiency and performance when using pretrained weights for backbones of agricultural semantic segmentation models for fruit and plant segmentation. These two experiments provide us with a quantifiable way to improve data efficiency for agricultural models. Following our analysis of methods to improve models, we continue our study by evaluating methods to improve data, by (3) assessing the effectiveness of standard spatial and visual image augmentations for improving model performance on fruit detection and fruit segmentation tasks. Finally, we note that while our methods suggest methods to improve data, there exists a pool of low-quality data which may potentially be disregarded, we (4) assess the effect of annotation quality, primarily accuracy and quantity of bounding boxes, on fruit detection tasks. Through these experiments, we find a strong potential for improving data efficiency for agricultural deep learning models without significant modifications to existing pipelines, with potential implications for a broad range of tasks. Furthermore, we open-source our standardized datasets and pretrained models through our framework AgML<sup>4</sup>, to facilitate the adoption of our described methods.

## 2 Methods and Experiments

To provide a baseline for our experiments, we have collected a set of agricultural datasets and developed a novel set of standard benchmarks and pretrained models on these datasets. These datasets have been standardized by task and are included in the dataset catalogue in the open-source agricultural machine learning framework called AgML. We discuss our methodology in developing a standard task-based data processing pipeline for training agricultural machine learning models, and our models’ results on these datasets. These pipelines are designed to be adaptable for the experiments conducted later in this research, and serve as our primary tool for them.

### 2.1 Collecting Datasets



**Fig. 2.** Countries from which agricultural data in the AgML public dataset collection was sourced.

For the purposes of this work, we restrict datasets to three tasks: image classification, semantic segmentation, and object detection. Furthermore, we focus on datasets with only a single 3-channel RGB image input, in order to facilitate the creation of benchmarks for a variety of datasets using standard deep learning models without any architectural modifications, allowing easier accessibility for a broader range of future applications. We compile datasets from a number of sources, with the majority coming from [Lu and Young \(2020\)](#). We provide a full listing of each of the datasets used in this dataset as well as key details and references in Table 1, which also contains benchmark performance on metrics as detailed later in Section 2.2.

Figure 2 displays, shaded in green, the countries from which data in AgML’s collection of public data sources was gathered from. This data is sourced from each of the six human-inhabited continents. Furthermore, this data covers a wide variety of conditions, including lightning, camera angle, and amount of obfuscation. We illustrate

<sup>4</sup> <https://github.com/Project-AgML/AgML>

**Table 1.** A listing of publicly available agricultural datasets used as part of this research (as described in *Dataset Name*), along with their number of images (*Images*), agricultural task (*Ag. Task*), and corresponding benchmark performance for the models described above on a testing set (*Benchmark*)


Dataset Name	Images	Classes	Ag. Task	Benchmark
<b>Image Classification</b>				<b>Accuracy</b>
bean_disease_uganda	1295	3	<i>B</i>	<b>96.90%</b>
plant_seedlings_aarhus (Giselsson et al., 2017)	5539	12	<i>A</i>	<b>94.39%</b>
soybean_weed_uav_brazil (dos Santos Ferreira et al., 2017)	15336	4	<i>A</i>	<b>100.0%</b>
sugarcane_damage_usa (Alencastre-Miranda et al., 2018)	153	6	<i>B</i>	<b>100.0%</b>
crop_weeds_greece (Espejo-Garcia et al., 2020)	508	4	<i>A</i>	<b>100.0%</b>
rangeland_weeds_australia (Olsen et al., 2019)	17509	10	<i>A</i>	<b>97.94%</b>
leaf_counting_denmark (Teimouri et al., 2018)	9372	9	<i>C</i>	<b>88.90%</b>
plant_village_classification (Hughes and Salath'e , 2015)	55448	39	<i>B</i>	<b>98.91%</b>
plant_doc_classification (Singh et al., 2020)	2336	28	<i>B</i>	<b>89.27%</b>
<b>Semantic Segmentation</b>				<b>mIoU</b>
carrot_weeds_germany (Haug and Ostermann, 2015)	60	2	<i>D</i>	<b>52.18%</b>
sugarbeet_weed_segmentation (Sa et al., 2018)	125	2	<i>D</i>	<b>53.59%</b>
apple_flower_segmentation (Dias et al., 2018)	148	1	<i>E</i>	<b>68.38%</b>
apple_segmentation_minnesota (Häni, Roy, and Isler, Häni et al.)	670	1	<i>F</i>	<b>79.08%</b>
rice_seedling_segmentation (Khan et al., 2020)	224	2	<i>D</i>	<b>52.20%</b>
peachpear_flower_segmentation (Dias et al., 2018)	42	1	<i>E</i>	<b>72.58%</b>
red_grapes_and_leaves_segmentation (Kalampokas et al., 2020)	258	2	<i>F</i>	<b>49.18%</b>
white_grapes_and_leaves_segmentation (Kalampokas et al., 2020)	273	2	<i>F</i>	<b>51.93%</b>
<b>Object Detection</b>				<b>mAP@0.5</b>
fruit_detection_worldwide (Sa et al., 2016)	565	7	<i>G</i>	<b>70.35%</b>
apple_detection_usa (Karkee et al., 2019)	2290	1	<i>G</i>	<b>94.16%</b>
apple_detection_spain (Gené-Mola et al., 2019)	967	1	<i>G</i>	<b>86.65%</b>
apple_detection_drone_brazil (Santos and Gebler, 2021)	689	1	<i>G</i>	<b>79.62%</b>
mango_detection_australia (Koirala et al., 2019)	1242	1	<i>G</i>	<b>95.32%</b>
grape_detection_californiaday (AI and Lab, 2019)	126	1	<i>G</i>	<b>69.01%</b>
grape_detection_californianight (AI and Lab, 2019)	150	1	<i>G</i>	<b>63.99%</b>
<b>Task Legend</b>				$A,B$ {Weed, Disease/Damage} Classification, $C$ Leaf Counting $D,E,F$ {Weed, Flower, Fruit} Segmentation, $G$ Fruit Detection

this diversity in two examples. Figure 4 contains an image from each of the object detection datasets used in our research. Each of the datasets contains a varying set of conditions, ranging from environmental differences, such as day versus night, to camera angle, including ground-level versus drone captures, to the type of fruit (the images shown include apple, grape, mango, and cantaloupe). Similarly, Figure 3 displays examples of imagery from three semantic segmentation datasets used in our research. We note the variation of agricultural tasks within these datasets, with one dedicated to solely weed segmentation, one being a combination of plant and weed segmentation, and another being flower segmentation. Furthermore, the difference in camera angle and environmental conditions are also present in these datasets.

## 2.2 Models and Methods

We select our model architectures with two primary focuses: state-of-the-art performance, and ease of access using existing deep learning frameworks. This approach allows us to construct standard pretrained models with both high-scoring benchmarks as well as ease of reproducibility and application to future tasks. We assess state-of-the-art performance using traditional performance benchmarks, namely ImageNet (Deng et al., 2009) for image classification, CityScapes (Cordts et al., 2016) for semantic segmentation, and COCO test-dev (Lin et al., 2014) for object detection. Similarly, we define the best-case scenario for ease of access to be the scenario where a model is contained entirely within the PyTorch (Paszke et al., 2019) library family, to reduce external dependencies. We now elaborate on our training procedure for each category of deep learning tasks.





**Fig. 3.** Sample images from a collection of the semantic segmentation datasets used in this research, containing the original image with annotated segmentation masks. The dataset from which each image is sourced from is annotated on top of it.

**Image Classification** In selecting an image classification model, we restrict our search to those available in the `torchvision` library, in order to facilitate our goal of ease-of-use and reduction of external dependences. The EfficientNet family of models has the highest performance out of all models in the library. So, we select **EfficientNetB4** as our model for image classification, considering its high ImageNet top-1 accuracy of 82.9% with a similar number of parameters to traditional state-of-the-art models like ResNet50. We select categorical cross-entropy loss as our criterion, defined as

$$L_{CE} = \sum_{i=1}^n y_i \cdot \log \hat{y}$$

where  $n$  is the number of classes,  $y_i$  is the ground truth, and  $\hat{y}$  is the prediction. We evaluate model performance using categorical accuracy.

**Semantic Segmentation** We again restrict our search to the `torchvision` library, and in turn we select the **DeepLabV3** model with a ResNet50 backbone, as described in the original implementation (and as is available in the library). We select a loss function based on the number of classes in the dataset being trained upon. For binary segmentation tasks, involving only a single class, we use binary cross-entropy loss with logits, defined as

$$L_{BCE} = -(y \log \hat{y} + (1 - y) \log(1 - \hat{y}))$$

over each pixel, where  $y$  represents the ground truth value, and  $\hat{y}$  represents the predicted value. For multi-class segmentation tasks, we use dice loss (Sudre et al., 2017):

$$L_D = 1 - \frac{2y\bar{y} + 1}{y + \bar{y} + 1}$$


where  $y$  again represents the ground truth and  $\hat{y}$  represents the predicted value. We evaluate model performance using the mean Intersection-over-Union (mIoU) metric:

$$\text{mIoU} = \frac{1}{n + 1} \sum_{i=0}^n \frac{p_{ii}}{\sum_{j=0}^n p_{ij} + \sum_{j=0}^n (p_{ji} - p_{ii})}$$

where  $n$  is the number of classes and  $p_{ij}$  is the number of pixels of class  $i$  predicted to belong to class  $j$ .

**Object Detection** Here, we expand our reach beyond the `torchvision` library, due to the reduced amount of object detection models available in it. We select the EfficientDet family of models –keeping in line with our prior selection of EfficientNet for image classification – choosing the **EfficientDetD4** model, due to its high performance of 49.3% mAP on the COCO evaluation dataset, with a comparable number of parameters to traditional state-of-the-art models like YOLOv3. We use the open source `effdet` library<sup>5</sup> as our implementation

<sup>5</sup> See: <https://github.com/rwightman/efficientdet-pytorch>



**Fig. 4.** Sample images from a collection of the object detection datasets used in this research, containing the original image with annotated fruit bounding boxes. The dataset from which each image is sourced from is annotated on top of it.

of the EfficientDetD4 model. The loss function used is focal loss (Lin et al., 2017) with  $\alpha = 0.25$  and  $\gamma = 1.5$ , as described in the original implementation, and given by:

$$L_F = -(1 - p_t)^\gamma \log(p_t)$$

where  $p_t = p$  if  $y = 1$ , otherwise  $p_t = 1 - p$  (and  $p$  is the estimated probability for the class  $y$ ). We evaluate model performance using the mean average precision metric, at an IOU threshold of 0.5 (which we denote as  $\text{mAP@0.5}$ ).

### 2.3 Model Training Parameters

For each of the three deep learning tasks, we develop training pipelines which are adaptable to each of the different agricultural datasets. These training pipelines are used both for our benchmark and pretrained model development as well as for our experiments, proving their versatility in adapting to a number of different agricultural deep learning tasks. We summarize the main training parameters for each task in Table 2. Data preprocessing is minimal; our only step involves resizing images to the image size in Table 2 and model-specific preprocessing: namely normalization for image classification and semantic segmentation and bounding box format conversion for object detection. In addition, for semantic segmentation and object detection, we decay the initial learning rate by  $\gamma = 0.75$  every 8 epochs. All models are trained using the PyTorch Lightning library (Falcon et al., 2019), on an NVIDIA Titan RTX GPU.

### 2.4 Agricultural Model Weights for Object Detection

Object detection models are traditionally trained using domain-generic weights: either random initialization or a common benchmark like the COCO dataset. For highly specified environments like agriculture, however,

**Table 2.** Summary of core training parameters for each task.

Task	Image Size	Optimizer	Learning Rate	Epochs
Image Classification	224	Adam	0.001	100
Semantic Segmentation	512	NAdam (Dozat, 2016)	0.005	50
Object Detection	512	AdamW (Loshchilov and Hutter, 2017)	0.0002	100

datasets are often of a much smaller magnitude. In many cases, as explored in Lu and Young (2020), datasets may only consist of a few hundred images, common in the case of object detection where data annotation takes significant time. In turn, using domain-specific agricultural weights as a starting point for deep learning models can significantly boost data efficiency, leading to models attaining high performance with less data and less training time.

To assess the performance of agricultural pretrained weights, we finetune an EfficientDetD4 model on the `fruit_detection_worldwide` dataset, and evaluate a set of five distinct pretrained model approaches: *COCO*, *GRAPE*, *GRAPENIGHT*, *APPLEDRONE* and *NONE*. The parameters of these different models are summarized in Table 3. *NONE* and *COCO* serve as our baselines for traditional pretrained weight approaches, while *GRAPE*, *GRAPENIGHT*, and *APPLEDRONE* represent agricultural pretrained models trained for different types of environments – *GRAPE* for daytime, *GRAPENIGHT* for nighttime, and *APPLEDRONE* for drone as opposed to ground imagery.

**Table 3.** Summary of pretrained model approaches for single-class detection.

Model Name	Datasets Pretrained On
<i>NONE</i>	None
<i>COCO</i>	COCO (Lin et al., 2014)
<i>GRAPE</i>	<code>grape_detection_californiaday</code>
<i>GRAPENIGHT</i>	<code>grape_california_night</code>
<i>APPLEDRONE</i>	<code>apple_detection_drone_brazil</code>

We then use these pretrained models for two experiments involving the `fruit_detection_worldwide` dataset. We first finetune a model on each of its 7 classes (specifically, 7 one-class models), and then one on the entire 7-class dataset (one 7-class model) – enabling an assessment of the results of pretrained weights on not just fruit localization, but also classification. Our experimental method and pipeline is carried from Section 2.2, though we restrict our evaluation to the first 50 epochs to obtain a better picture of data efficiency.

## 2.5 Agricultural Backbone Weights for Semantic Segmentation

In contrast to the standard available pretrained weights for object detection, tasks like semantic segmentation often do not have even a domain-general set of pretrained weights available for training – especially due to the fact that a network with a certain number of classes cannot properly utilize the pretrained weights from one with a different number of classes, unlike in object detection. In such cases, pretrained weights are only used for the *backbone*, a feature extraction model which is used as the input stem into a semantic segmentation network. For instance, in our DeepLabV3 model, the backbone was ResNet50, a model otherwise used for image classification tasks. Traditionally, such classifiers are trained on the ImageNet dataset (Deng et al., 2009). However, as we intend to demonstrate, using agricultural domain-specific weights for even the backbone can performance gain for the larger network as a whole.

Image classification datasets are largely different from semantic segmentation datasets – while most of our semantic segmentation datasets in Table 1 are for locating fruits or weeds, many of our image classification datasets are for *distinguishing* between leaves and weeds. For a thorough examination of the viability of agricultural weights, we select three different agricultural image classification datasets to use as pretrained models, summarized in Table 4. *VILLAGE* is our largest image classification dataset, with over 55,000 images for species and

**Table 4.** Summary of pretrained backbones for semantic segmentation. Note that in contrast to object detection, the number of classes of the image classification models are not relevant, as they are being used as feature extractors.

Model Name	Datasets Pretrained On
<i>NONE</i>	None
<i>IMAGENET</i>	ImageNet (Deng et al., 2009)
<i>VILLAGE</i>	plant_village_classification
<i>COUNTING</i>	leaf_counting_denmark

disease classification, while *COUNTING* is a much smaller dataset focused on a largely different application – leaf counting. We assess these four pretrained backbones on three different datasets: **apple\_flower\_segmentation**, **apple\_segmentation\_minnesota**, and **rice\_seedling\_segmentation**. Our experimental method and pipeline is the same as in Section 2.2, though we restrict our evaluation to only the first 20 epochs of training.

## 2.6 Effectiveness of Augmentations for Generating Diverse Data

A standard procedure for improving the performance, and in particular, generalizability of deep learning models, is to apply visual and spatial augmentations to the input data, with the intended goal of increasing the diversity of the input data. Standard spatial augmentations include rotation, distortion (such as affine transforms), and cropping, while standard visual augmentations range from saturation and brightness contrast to Gaussian noise addition. For agricultural deep learning models, generalizability is a crucial goal, as environmental conditions vary significantly in different scenarios.


In turn, we conduct a robust analysis of various augmentations across a number of datasets in order to collect insight into the most viable augmentations for improving generalizability for agriculture. Performance is assessed of a collection of different augmentations on a set of different datasets across our three different tasks – image classification, semantic segmentation, and object detection – to provide a greater insight into each augmentation’s effectiveness. Our analysis only involves the usage of a single augmentation per trial, without combinations, in order to provide an insight as to the capability of each augmentation independently, relevant to our goal of assessing augmentations in a agricultural context. Additionally, we collect final metric scores early in training, namely at only 20 epochs. As noted from our prior observations, most models begin approaching the relative peak of their performance at this threshold. So, a better analysis of the effectiveness of augmentations is done by analyzing their performance improvements earlier in training. We use the albumentations (Buslaev et al., 2020) library for our augmentations, noting its seamless integration into our existing PyTorch training pipeline. The augmentations we use, alongside model performance corresponding to the datasets we select, are summarized in Table 5 to eliminate redundancy. Furthermore, we provide a visual example of each of the augmentations we use in Figure 5.

## 2.7 Effects of Annotation Quality on Learning

Data annotation is a heavily time-consuming task within deep learning. In the field of agriculture, this is an especially prevalent issue, due to the present nuance and even simply the quantity of objects such as fruits which need to be annotated in fruit detection tasks. Oftentimes, the amount of effort needed to annotate a certain dataset can result in that dataset being generated with imperfect and even obviously inaccurate annotations, potentially leading to them being unusable for deep learning tasks. However, with the already present scarcity of agricultural data, any loss of existing data has larger impacts on the open-source data community. Thus, we now assess how the quality of annotations on agricultural data affects models’ overall performance. In particular, we aim to observe whether models are capable of achieving high performance without perfect annotations, potentially leading to less strict annotation requirements for deep learning models in the future.

We use three fruit detection datasets: **grape\_detection\_californiaday**, **apple\_detection\_drone\_brazil**, and **fruit\_detection\_worldwide**. For each dataset, we set aside a common test set. For the remaining data, we create a set of datasets with reduced quality annotations, removing a certain percentage of bounding boxes for each image. This procedure is conducted for retention of only 30% of annotations, to 90% of annotations, with 10% increments. Using our object detection pipeline from Section 2.2, we train for 25 epochs on each dataset and record their **mAP@0.5** on the aforementioned test sets.





**Fig. 5.** An example image from `grape_detection_californiaday` with each of the augmentations used in this research applied to it.

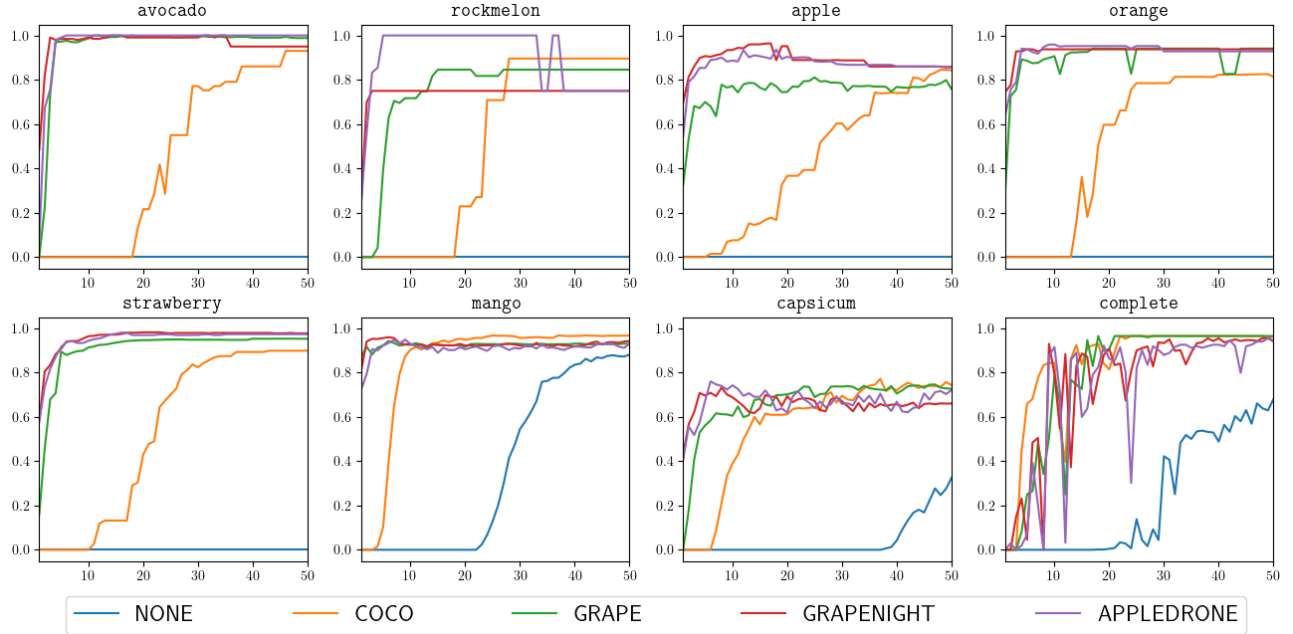
### 3 Results and Discussion

#### 3.1 Standard Benchmarks for Agricultural Datasets

The performance of each of our task-based models on the datasets we have collected is summarized in Table 1. Our standard pipelines enable our models to achieve comparable performance with existing benchmarks on our collected datasets, in certain cases even exceeding them. As mentioned in Section 2.2, these datasets, along with pretrained models which achieved these benchmarks, are available through the open-source framework AgML, enabling further research on developing even more efficient data and model pipelines.

#### 3.2 Performance of Agricultural Pretrained Weights for Object Detection

A summary of the results of our experiments using agricultural pretrained weights for object detection are shown in Figure 6, with each of the fruits representing the 7 one-class models, and **complete** representing the one 7-class model. For all of the 7-class models, the pretrained agricultural models significantly outperform the *COCO* and *NONE* baselines. We observe that the agricultural models tended to plateau at their maximum mean average precision before 10 epochs, for most fruits – some, like strawberry, avocado, and mango, see this plateau as early as 5 epochs. On the other hand, *COCO*, the standard baseline, usually takes between 25 to 40 epochs to reach its own maximum value, while the *NONE* model fails to even break zero mean average precision for five out of seven fruits, and only reaches a comparable score to *COCO* for one. For the 7-class model, we see




**Fig. 6.** Comparison of  $\text{mAP@0.5}$  values on seven different individual fruits from the `fruit_detection_worldwide` dataset, alongside the complete dataset (listed as `complete`, the bottom right plot).

that the agricultural pretrained models follow a similar trajectory to *COCO*, taking around or over 30 epochs to plateau at their maximum mean average precision.

A number of observations can be drawn from these results. Pretrained weights are clearly a better approach when training object detection models, as they converge significantly faster than random weights, saving time and resources in training. The poor performance of the *NONE* model demonstrates the value of transferring prior knowledge when training on new tasks. Furthermore, for agricultural domain-specific detection tasks, using agricultural pretrained weights can provide, at minimum, comparable performance to *COCO*, or other standard weights, but at best, even better performance and much faster convergence. Our results are consistent with previous observations for other tasks, such as [Too et al. \(2019\)](#), where agricultural pretrained weights were found to increase performance. Our results are further consistent with previous observations specifically for object detection such as in [Bresilla et al. \(2019\)](#), where using pretrained fruit detection weights from similar environments resulted in a significant performance increase. A key insight from our work, however, is the fact that agricultural environment played little role in performance improvement – all three models, despite their distinct environments, achieved similar performance on a dataset in yet another noticeably different environment. Furthermore, each agricultural pretrained model followed a similar performance *trajectory*, corroborating our observations that using agricultural pretrained weights of any form result in not only higher performance, but faster convergence. In turn, this approach, which only requires a slight modification to existing pipelines, can significantly improve training results for future agricultural deep learning work.

### 3.3 Performance of Agricultural Backbone Weights for Semantic Segmentation

Similar to the previous section, we summarize the results of our semantic segmentation models in Figure 7. We find that our models with agricultural pretrained backbones, *VILLAGE* and *COUNTING*, tend to outperform models with the existing *NONE* and *IMAGENET* backbones to a measureable degree. For the datasets `apple_segmentation_minnesota` and `rice_seedling_segmentation`, our agricultural pretrained models reach their maximum  $\text{mIOU}$  almost within the first couple of epochs, while the other two models either take between 10-15 epochs to reach a similar degree of performance, or do not reach it at all. For `apple_flower_segmentation`, our models reach similar levels of performance as *COCO*, but maintain a more consistent trajectory of performance increase, while *COCO* and *NONE* see a spike in performance followed by a quick drop, suggesting relative inconsistency.



**Fig. 7.** Comparison of mIOU values for four different pretrained backbones on three different evaluation datasets. The  $y$ -axis is on a consistent scale, although each dataset has its models converge to a different range of values.

Semantic segmentation is a unique task due to the required level of data annotation required for it – each pixel in an image must be assigned to a certain class. This is reflected in the reduced size of segmentation datasets, an example of which can be seen in the image counts of Table 1. As a result, there is potentially not enough available data to generate a large-scale pretrained network for semantic segmentation, especially when noting the potential difference in classes between a pretrained network and an applied, task-specific network – resulting in loss of knowledge in weights which cannot be transferred. This is noted in works like [Sodjinou et al. \(2022\)](#) and [Peng et al. \(2021\)](#), where standard domain-general weights are used in lieu of their agricultural equivalents. However, image classification datasets are in relative surplus in comparison to semantic segmentation, and most segmentation models use classification models as their feature extractors, representing the encoder module for a typical encoder-decoder architecture. For most tasks, a standard approach is to freeze weights for this feature extraction module, the encoder, and only update parameters for the segmentation head, the decoder. This is inspired by the fact that the feature extraction model obtains the most relevant information regarding the location of the objects being segmented – however, for agricultural networks, this approach seems largely counter-intuitive, as the feature extractor is usually trained on domain-general images and this unable to recognize the most relevant features for plants. Using agricultural pretrained weights, on the other hand, can improve the localization performance of the feature extractor, in turn improving the network’s performance as a whole. Our results show that using agricultural weights for feature extractors serves as a potential workaround for the lack of a standard semantic segmentation pretrained model, still enabling improved performance and reduced training time.

### 3.4 Augmentation Effectiveness

We summarize the performance of each of the augmentations used in this work, on six different datasets ( $A, B, C$  object detection, and  $D, E, F$  semantic segmentation), in Table 5. When using a single augmentation, as in our experiments, observe that spatial augmentations – namely augmentations (1) through (5) as denoted in the table, when used independently – generally outperform visual augmentations – (6) through (11) – for all of the datasets used in this experiment. The augmentations which provide the largest performance increase are all spatial, while the visual augmentations, in many cases, fail to even provide *any* performance increase whatsoever. For our subsequent analysis, we refer to each dataset using its key in the table ( $A$  refers to `grape_detection_californiaday`, and so on).

Our next key observation regards the effect of different augmentations on different environments and conditions. For instance, datasets  $A$  and  $B$  consist of the same fruit in the same environment, simply in day conditions versus night conditions. While certain visual augmentations, such as rain and fog, provide a performance boost for  $A$ , they actually significantly reduce performance for  $B$ , indicating that certain factors, such as rain and fog, are less valuable information in nighttime environments as opposed to daytime environments. Another example of the effects of different augmentations, this time for spatial augmentations with reference to camera positioning, is found when observing the results for  $C$ . While  $A$  and  $B$ , the other object detection datasets, consist of ground cameras capturing images of plants,  $C$  consists of aerial imagery captured by a drone. In turn, augmentations like rotation, translation provide a better opportunity for the model to generalize to different aerial camera positions, as opposed to less distortive augmentations like horizontal and vertical flips. For se-




**Table 5.** Summary of performance of different augmentations on the aforementioned semantic segmentation and object detection datasets. Bold indicates the highest performance for a certain dataset.

Augmentation	<i>A</i>	<i>B</i>	<i>C</i>	<i>D</i>	<i>E</i>	<i>F</i>
	mAP@0.5			mIOU		
original (0)	64.08%	53.62%	75.11%	<b>69.78%</b>	78.52%	51.04%
horizontal-flip (1)	62.76%	58.32%	74.46%	66.97%	79.19%	<b>51.35%</b>
vertical-flip (2)	<b>68.45%</b>	56.12%	74.48%	62.48%	75.35%	49.19%
shear (3)	63.31%	58.28%	80.96%	65.61%	77.54%	51.04%
rotate (4)	60.77%	<b>61.06%</b>	84.42%	64.51%	73.96%	49.75%
translate (5)	65.53%	53.05%	<b>86.60%</b>	66.26%	<b>79.22%</b>	48.57%
brightness (6)	59.01%	52.69%	78.80%	66.45%	77.72%	44.50%
hsv-shift (7)	56.51%	49.32%	82.39%	51.34%	73.12%	42.95%
gaussian-blur (8)	65.17%	57.09%	84.60%	69.25%	78.61%	50.96%
rain (9)	60.44%	42.25%	81.20%	69.49%	73.09%	50.11%
fog (10)	65.35%	48.66%	78.81%	62.73%	70.81%	44.31%
sun-flare (11)	65.92%	55.87%	80.72%	49.86%	74.70%	43.63%
Dataset Legend	<sup>A</sup> grape_detection_californiaday					
	<sup>B</sup> grape_detection_californianight					
	<sup>C</sup> apple_detection_drone_brazil					
	<sup>D</sup> apple_segmentation_minnesota					
	<sup>E</sup> apple_flower_segmentation					
	<sup>F</sup> rice_seedling_segmentation					

semantic segmentation tasks, we observe a less obvious boost in performance from augmentations. In fact, for *D*, no augmentations provide the best result, while for *E* and *F* the performance boost is within one percent of no augmentations. This potentially stems from the different conditions of semantic segmentation tasks as opposed to object detection tasks – semantic segmentation tasks, especially for fruit or leaf segmentation, like *D* and *E*, can come down to individual pixels, and such precision may likely be distorted by spatial augmentations. In all tasks, while visual augmentations appear to provide less of an obvious benefit as opposed to spatial augmentations, our results still provide some insight into their potential viability. In particular, augmentations which affect the hue of an image, such as HSV shift and brightness or contrast shifts, tend to actually decrease performance, as they result in a model adapting to those input colors. On the contrary, Gaussian blurring provides a consistent, though sometimes minute, performance increase relative to the input. This suggests that a reduction of features, as done by adding a slight blur to the input, can in fact increase performance by potentially having the model learn more general features in the input environment.

A major benefit of augmentations is their ability to generate data potentially modified for a large variety of conditions, enabling greater generalizability for agricultural models. Existing state-of-the-art work done for domain transfer often involves the usage of generative adversarial networks (GANs), such as [Fei et al. \(2021\)](#), who developed a GAN for transferring sample imagery between day and night domains, and [Zhang et al. \(2021\)](#), who used a CycleGAN network to edit the fruits present in imagery while maintaining the environmental conditions. While augmentations may not necessarily be able to provide an entire domain transfer as in the prior methods, they can still provide a generalization of conditions – for instance, rain and fog augmentations, as used in our research, can generalize data to a broader range of common conditions across the world. Other augmentations, like sun flare, can reduce the impact of edge cases where images are obscured. We note that certain augmentations expand beyond a single data points, potentially involving multiple images being transformed together. This is the case in common augmentations such as mosaic and patch, as explored in their applications for agricultural robotics in [Su et al. \(2021\)](#). While we do not include these augmentations in our research, as they involve larger modifications to our pipelines, they do present potential further performance improvements for agricultural models. If used properly based on the input environment, augmentations provide an efficient way to produce a larger set of environments for agricultural data, improving generalizability of models.





**Fig. 8.** Performance of models on three different datasets dependent on the percentage of bounding box annotations retained per image.


### 3.5 Annotation Quality

The performance of each of our annotation quality models are recorded in Figure 8, which displays the performance of models relative to the quality of annotations on the datasets they are trained on. For the datasets `grape_detection_californiaday` and `apple_detection_drone_brazil`, we find a consistent increase in performance as the quality of annotations increase. Notably, we do not observe a one-to-one correspondence between quality of annotations and level of performance: e.g., 30% of annotations actually corresponds to over 50% of peak performance, suggesting that models trained on lower quality annotations are, at least to some extent, able to localize similar objects to a slightly higher than expected degree. We provide a sample of predicted bounding boxes for each `grape_detection_californiaday` iteration in Figure 9. This observance, however, does not hold for `fruit_detection_worldwide`, which has a distinction of being a multi-class dataset in contrast to the prior two. Surprisingly, for this dataset, we find that our models achieve maximum performance with only 50% of annotations. This example, in turn, provides an insight into the hazardous potential of using lower-quality annotations. At 50% of annotations, our model may have only learned general features, thus allowing it to still perform much better on other sample images. However, higher percentages of still low-quality annotations may result in a model learning to distinguish between different instances of the same fruit by extremely nuanced features, overfitting on the samples on which it is trained.

A number of methods have been proposed in recent work to try and automate the data annotation process, which is one of the major bottlenecks for agricultural deep learning. Some approaches, such as [Beck et al. \(2020\)](#), use robotic systems to capture data from a wider range of angles and automatically annotate bounding boxes knowing the location of plants in controlled environments. Other approaches ([Fatima and Mahmood, 2021](#); [Ciarfuglia et al., 2022](#)) involve semi-supervised learning, taking advantage of highly precise data to train high-performance models which in turn improve predictions on unlabeled images. Such approaches show potential for not only making new, high-quality agricultural datasets, but even potentially improving upon existing datasets, annotating missed fruits or other objects which may have been missed out on. Nevertheless, our work demonstrates that fruit detection models can still obtain relatively high performance with some fruits unannotated. In turn, potential new approaches may involve using lower-quality datasets in coordination with better datasets, expanding the pool of available data for agriculture and further boosting model performance.

## 4 Conclusion

In our work, we have developed a novel set of standardized and centralized agricultural datasets, alongside benchmarks and pretrained models using state-of-the-art models. Our custom pipelines achieve comparable



**Fig. 9.** Sample predictions by models trained on different levels of annotation quality. The percentages above each image reflect the amount of bounding boxes retained for the data used in training that specific model.

performance with existing benchmarks using no extensive data or architectural modifications, making them widely applicable to a variety of agricultural deep learning tasks. We have also assessed a number of existing methods for improving model performance domain-specific to agriculture, including using agricultural pretrained model weights and image augmentations. Furthermore, we have even explored the viability of traditionally overlooked lower-quality data, potentially expanding the data pool for agriculture. Our results demonstrate that slight training modifications can significantly boost model performance and result in shorter convergence time. We have open-sourced our standardized and centralized versions of the datasets used in our work, alongside our pretrained models and benchmarks, to guide easier adoption of our described methods.

## 5 Acknowledgements

This project was partly supported by the USDA AI Institute for Next Generation Food Systems (AIFS), USDA award number 2020-67021-32855.

## References

- AI, Plant and Biophysics Lab. 2019 “Grape Detection 2019 Day.”.
- Alencastre-Miranda, Moises, Joseph R. Davidson, Richard M. Johnson, Herman Waguespack, and Hermano Igo Krebs. “Robotics for Sugarcane Cultivation: Analysis of Billet Quality using Computer Vision.” *IEEE Robotics and Automation Letters* 3 (2018): 3828–3835.
- Bargoti, Suchet and James Underwood. 2016 “Deep Fruit Detection in Orchards.”.
- Beck, Michael A., Chen-Yi Liu, Christopher P. Bidinosti, Christopher J. Henry, Cara M. Godee, and Manisha Ajmani. “An embedded system for the automated generation of labeled plant images to enable machine learning applications in agriculture.” *PLOS ONE* 15 (Dec 2020): e0243923.
- Bresilla, Kushtrim, Giulio Demetrio Perulli, Alexandra Boini, Brunella Morandi, Luca Corelli Grappadelli, and Luigi Manfrini. “Single-Shot Convolution Neural Networks for Real-Time Fruit Detection Within the Tree.” *Frontiers in Plant Science* 10 (2019).
- Buslaev, Alexander, Vladimir I. Iglovikov, Eugene Khvedchenya, Alex Parinov, Mikhail Druzhinin, and Alexandr A. Kalinin. “Albumations: Fast and Flexible Image Augmentations.” *Information* 11 (2020).
- Ciarfuglia, Thomas A., Ionut Marian Motoi, Leonardo Saraceni, and Daniele Nardi. “Pseudo-Label Generation for Agricultural Robotics Applications.” *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*. June 2022, 1686–1694.
- Cordts, Marius, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. “The Cityscapes Dataset for Semantic Urban Scene Understanding.” *CoRR* abs/1604.01685 (2016).
- Deng, Jia, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. “ImageNet: A large-scale hierarchical image database.” *2009 IEEE Conference on Computer Vision and Pattern Recognition*. 2009, 248–255.
- Dias, Philippe A., Amy Tabb, and Henry Medeiros. “Multispecies Fruit Flower Detection Using a Refined Semantic Segmentation Network.” *IEEE Robotics and Automation Letters* 3 (2018): 3003–3010.
- Santos Ferreira, Alessandro, Daniel Matte Freitas, Gercina Gonalves da Silva, Hemerson Pistori, and Marcelo Theophilo Folhes. “Weed Detection in Soybean Crops Using ConvNets.” *Comput. Electron. Agric.* 143 (dec 2017): 314–324.
- Dozat, Timothy. “Incorporating Nesterov Momentum into Adam.” (Feb 2016).
- Elsherbiny, Osama, Lei Zhou, Lei Feng, and Zhengjun Qiu. “Integration of Visible and Thermal Imagery with an Artificial Neural Network Approach for Robust Forecasting of Canopy Water Content in Rice.” *Remote Sensing* 13 (2021).
- Espejo-Garcia, Borja, Nikos Mylonas, Loukas Athanasakos, Spyros Fountas, and Ioannis Vasilakoglou. “Towards weeds identification assistance through transfer learning.” *Computers and Electronics in Agriculture* 171 (2020): 105306.
- Falcon, William et al. “Pytorch lightning.” *GitHub*. Note: <https://github.com/PyTorchLightning/pytorch-lightning> 3 (2019).
- Fatima, Tazeen and Tariq Mahmood. “Semi-Supervised Learning in Smart Agriculture: A Systematic Literature Review.” *2021 6th International Multi-Topic ICT Conference (IMTIC)*. Nov 2021, 1–8.
- Fei, Zhenghao, Alex Olenskyj, Brian N. Bailey, and Mason Earles. “Enlisting 3D Crop Models and GANs for More Data Efficient and Generalizable Fruit Detection.” *2021 IEEE/CVF International Conference on Computer Vision Workshops (ICCVW)*. 2021, 1269–1277.
- Gené-Mola, Jordi, Verónica Vilaplana, Joan R. Rosell-Polo, Josep-Ramon Morros, Javier Ruiz-Hidalgo, and Eduard Gregorio. “KFuji RGB-DS database: Fuji apple multi-modal images for fruit detection with color, depth and range-corrected IR data.” *Data in Brief* 25 (2019): 104289.
- Giselsson, Thomas Mosgaard, Mads Dyrmann, Rasmus Nyholm Jørgensen, Peter Kryger Jensen, and Henrik Skov Midtby. “A Public Image Database for Benchmark of Plant Seedling Classification Algorithms.” *arXiv preprint* (2017).
- Haug, Sebastian and Jörn Ostermann. “A Crop/Weed Field Image Dataset for the Evaluation of Computer Vision Based Precision Agriculture Tasks.” *Computer Vision - ECCV 2014 Workshops*. Ed. Lourdes Agapito, Michael M. Bronstein, and Carsten Rother Cham: Springer International Publishing, 2015, 105–116.
- Hughes, David P. and Marcel Salath’e. “An open access repository of images on plant health to enable the development of mobile disease diagnostics through machine learning and crowdsourcing.” *CoRR* abs/1511.08060 (2015).

- Häni, Nicolai, Pravakar Roy, and Volkan Isler. “MinneApple: A Benchmark Dataset for Apple Detection and Segmentation.”
- Jeon, Hong Y., Lei F. Tian, and Heping Zhu. “Robust crop and weed segmentation under uncontrolled outdoor illumination.” *Sensors (Basel, Switzerland)* 11 (2011): 6270–6283.
- Kalampokas, Theofanis, Konstantinos Tziridis, Alexandros Nikolaou, Eleni Vrochidou, George A. Papakostas, Theodore Pachidis, and Vassilis G. Kaburlasos. “Semantic Segmentation of Vineyard Images Using Convolutional Neural Networks.” *Proceedings of the 21st EANN (Engineering Applications of Neural Networks) 2020 Conference*. Ed. Lazaros Iliadis, Plamen Parvanov Angelov, Chrisina Jayne, and Elias Pimenidis Cham: Springer International Publishing, 2020, 292–303.
- Karkee, Manoj, Santosh Bhusal, and Qin Zhang. “Apple Dataset Benchmark from Orchard Environment in Modern Fruiting Wall.” (2019).
- Khan, Abbas, Talha Ilyas, Muhammad Umraiz, Zubaer Ibna Mannan, and Hyongsuk Kim. “CED-Net: Crops and Weeds Segmentation for Smart Farming Using a Small Cascaded Encoder-Decoder Architecture.” *Electronics* 9 (2020).
- Koirala, A., K. B. Walsh, Z. Wang, and C. McCarthy. “Deep learning for real-time fruit detection and orchard fruit load estimation: benchmarking of ‘MangoYOLO’.” *Precision Agriculture* 20 (Dec 2019): 1107–1135.
- Li, Min, Zhijie Zhang, Liping Lei, Xiaofan Wang, and Xudong Guo. “Agricultural Greenhouses Detection in High-Resolution Satellite Images Based on Convolutional Neural Networks: Comparison of Faster R-CNN, YOLO v3 and SSD.” *Sensors (Basel, Switzerland)* 20 (Aug 2020): 4938.
- Lin, Tsung-Yi, Priya Goyal, Ross B. Girshick, Kaiming He, and Piotr Dollár. “Focal Loss for Dense Object Detection.” *CoRR* abs/1708.02002 (2017).
- Lin, Tsung-Yi, Michael Maire, Serge J. Belongie, Lubomir D. Bourdev, Ross B. Girshick, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. “Microsoft COCO: Common Objects in Context.” *CoRR* abs/1405.0312 (2014).
- Loshchilov, Ilya and Frank Hutter. “Fixing Weight Decay Regularization in Adam.” *CoRR* abs/1711.05101 (2017).
- Lu, Yuzhen and Sierra Young. “A survey of public datasets for computer vision tasks in precision agriculture.” *Computers and Electronics in Agriculture* 178 (2020): 105760.
- Milioto, Andres, Philipp Lottes, and Cyrill Stachniss. “Real-Time Semantic Segmentation of Crop and Weed for Precision Agriculture Robots Leveraging Background Knowledge in CNNs.” *2018 IEEE International Conference on Robotics and Automation (ICRA)*. 2018, 2229–2235.
- Mohanty, Sharada P., David P. Hughes, and Marcel Salathé. “Using Deep Learning for Image-Based Plant Disease Detection.” *Frontiers in Plant Science* 7 (2016).
- Moon, Taewon and Jung Eek Son. “Knowledge transfer for adapting pre-trained deep neural models to predict different greenhouse environments based on a low quantity of data.” *Computers and Electronics in Agriculture* 185 (2021): 106136.
- Moreira, Germano, Sandro Augusto Magalhães, Tatiana Pinho, Filipe Neves dos Santos, and Mário Cunha. “Benchmark of Deep Learning and a Proposed HSV Colour Space Models for the Detection and Classification of Greenhouse Tomato.” *Agronomy* 12 (2022).
- Nevavuori, Petteri, Nathaniel Narra, Petri Linna, and Tarmo Lipping. “Crop Yield Prediction Using Multitemporal UAV Data and Spatio-Temporal Deep Learning Models.” *Remote Sensing* 12 (2020).
- Nowakowski, Artur, John Mrziglod, Dario Spiller, Rogerio Bonifacio, Irene Ferrari, Pierre Philippe Mathieu, Manuel Garcia-Herranz, and Do-Hyung Kim. “Crop type mapping by using transfer learning.” *International Journal of Applied Earth Observation and Geoinformation* 98 (2021): 102313.
- Olsen, Alex, Dmitry A. Konovalov, Bronson Philippa, Peter Ridd, Jake C. Wood, Jamie Johns, Wesley Banks, Benjamin Girgenti, Owen Kenny, James Whinney, Brendan Calvert, Mostafa Rahimi Azghadi, and Ronald D. White. “DeepWeeds: A Multiclass Weed Species Image Dataset for Deep Learning.” *Scientific Reports* 9 (Feb 2019): 2058.
- Pan, Sinno Jialin and Qiang Yang. “A Survey on Transfer Learning.” *IEEE Transactions on Knowledge and Data Engineering* 22 (2010): 1345–1359.
- Paszke, Adam, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zach DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019 “PyTorch: An Imperative Style, High-Performance Deep Learning Library.”
- Peng, Yun, Aichen Wang, Jizhan Liu, and Muhammad Faheem. “A Comparative Study of Semantic Segmentation Models for Identification of Grape with Different Varieties.” *Agriculture* 11 (Oct 2021): 997.



- Sa, I., Z. Chen, M. Popović, R. Khanna, F. Liebisch, J. Nieto, and R. Siegwart. “weedNet: Dense Semantic Weed Classification Using Multispectral Images and MAV for Smart Farming.” *IEEE Robotics and Automation Letters* 3 (Jan 2018): 588–595.
- Sa, Inkyu, Zongyuan Ge, Feras Dayoub, Ben Upcroft, Tristan Perez, and Chris McCool. “DeepFruits: A Fruit Detection System Using Deep Neural Networks.” *Sensors* 16 (2016).
- Sahili, Zahraa Al and Mariette Awad. 2022 “The Power of Transfer Learning in Agricultural Applications: AgriNet.”
- Santos, Thiago T. and Luciano Gebler. “A methodology for detection and localization of fruits in apples orchards from aerial images.” *CoRR* abs/2110.12331 (2021).
- Shorten, Connor and Taghi M. Khoshgoftaar. “A survey on Image Data Augmentation for Deep Learning.” *Journal of Big Data* 6 (Jul 2019): 60.
- Shrivastava, Vimal K., Monoj K. Pradhan, and Mahesh P. Thakur. “Application of Pre-Trained Deep Convolutional Neural Networks for Rice Plant Disease Classification.” *2021 International Conference on Artificial Intelligence and Smart Systems (ICAIS)*. 2021, 1023–1030.
- Singh, Davinder, Naman Jain, Pranjali Jain, Pratik Kayal, Sudhakar Kumawat, and Nipun Batra. “PlantDoc: A Dataset for Visual Plant Disease Detection.” *Proceedings of the 7th ACM IKDD CoDS and 25th COMAD*. New York, NY, USA: Association for Computing Machinery, 2020, 249–253.
- Sodjinou, Sovi Guillaume, Vahid Mohammadi, Amadou Tidjani Sanda Mahama, and Pierre Gouton. “A deep semantic segmentation-based algorithm to segment crops and weeds in agronomic color images.” *Information Processing in Agriculture* 9 (Sep 2022): 355–364.
- Su, Daobilige, He Kong, Yongliang Qiao, and Salah Sukkarieh. “Data augmentation for deep learning based semantic segmentation and crop-weed classification in agricultural robotics.” *Computers and Electronics in Agriculture* 190 (Nov 2021): 106418.
- Sudre, Carole H., Wenqi Li, Tom Vercauteren, Sebastien Ourselin, and M. Jorge Cardoso. “Generalised Dice Overlap as a Deep Learning Loss Function for Highly Unbalanced Segmentations.” *Lecture Notes in Computer Science* (2017): 240–248.
- Teimouri, Nima, Mads Dyrmann, Per Rydahl Nielsen, Solvejg Kopp Mathiassen, Gayle J. Somerville, and Rasmus Nyholm Jørgensen. “Weed Growth Stage Estimator Using Deep Convolutional Neural Networks.” *Sensors* 18 (2018).
- Too, Edna Chebet, Li Yujian, Sam Njuki, and Liu Yingchun. “A comparative study of fine-tuning deep learning models for plant disease identification.” *Computers and Electronics in Agriculture* 161 (Jun 2019): 272–279.
- Wang, Chunying, Baohua Liu, Lipeng Liu, Yanjun Zhu, Jialin Hou, Ping Liu, and Xiang Li. “A review of deep learning used in the hyperspectral image analysis for agriculture.” *Artificial Intelligence Review* 54 (Oct 2021): 5205–5253.
- Wang, Lijun, Jiayao Wang, Zhenzhen Liu, Jun Zhu, and Fen Qin. “Evaluation of a deep-learning model for multispectral remote sensing of land use and crop classification.” *The Crop Journal* (Feb 2022).
- Zhang, Wenli, Kaizhen Chen, Jiaqi Wang, Yun Shi, and Wei Guo. “Easy domain adaptation method for filling the species gap in deep learning-based fruit detection.” *Horticulture Research* 8 (Jun 2021): 1–13.
- Zheng, Yang-Yang, Jian-Lei Kong, Xue-Bo Jin, Xiao-Yi Wang, Ting-Li Su, and Min Zuo. “CropDeep: The Crop Vision Dataset for Deep-Learning-Based Classification and Detection in Precision Agriculture.” *Sensors (Basel, Switzerland)* 19 (Mar 2019): 1058.
- AtilaÜmit , Murat Uçar, Kemal Akyol, and Emine Uçar. “Plant leaf disease classification using EfficientNet deep learning model.” *Ecological Informatics* 61 (2021): 101182.