HIGH-THROUGHPUT IMAGE-BASED PLANT STAND COUNT ESTIMATION USING CONVOLUTIONAL NEURAL NETWORKS

Saeed Khaki

Industrial Engineering Department
Iowa State University
Ames, Iowa 50011, USA
skhaki@iastate.edu

Hieu Pham

Syngenta Seeds
Slater, Iowa 50244, USA
hieu.pham@syngenta.com

Ye Han

Syngenta Seeds
Slater, Iowa 50244, USA

Wade Kent

Syngenta Seeds
Slater, Iowa 50244, USA

Lizhi Wang

Industrial Engineering Department
Iowa State University
Ames, Iowa 50011, USA

November 5, 2020

ABSTRACT

The future landscape of modern farming and plant breeding is rapidly changing due to the complex needs of our society. The explosion of collectable data has started a revolution in agriculture to the point where innovation must occur. To a commercial organization, the accurate and efficient collection of information is necessary to ensure that optimal decisions are made at key points of the breeding cycle. However, due to the sheer size of a breeding program and current resource limitations, the ability to collect precise data on individual plants is not possible. In particular, efficient phenotyping of crops to record its color, shape, chemical properties, disease susceptibility, etc. is severely limited due to labor requirements and, oftentimes, expert domain knowledge. In this paper, we propose a deep learning based approach, named DeepStand, for image-based corn stand counting at early phenological stages. The proposed method adopts a truncated VGG-16 network as a backbone feature extractor and merges multiple feature maps with different scales to make the network robust against scale variation. Our extensive computational experiments suggest that our proposed method can successfully count corn stands and out-perform other state-of-the-art methods. It is the goal of our work to be used by the larger agricultural community as a way to enable high-throughput phenotyping without the use of extensive time and labor requirements.

Keywords Corn stand counting, Convolutional neural networks, Plant phenotyping, Deep learning

1 Introduction

The “phenotyping bottleneck” refers to the phenomenon that the agricultural community is not able to accurately and efficiently collect data on physical properties of crops (Furbank and Tester, 2011). This high-throughput phenotyping (HTP) bottleneck is oftentimes due to resource limitations such as labor, time, and domain expertise in pathology and genetics, even for large scale farming operations and commercial breeding programs. Certain phenotypic traits such as early stage stand counting for corn (*Zea mays* L.) can only be accurately performed during its early growth stage. If this data collection time-window is missed, then the task is near impossible to complete (McWilliams et al., 1999).

Understanding the data collection challenges facing modern agriculture, agronomists and researchers have turned to modern solutions that combine modern machine learning and imaging analytics (Yang et al., 2020). Image-capturing devices such as unmanned aerial vehicles, high definition cameras, and, even, cell phone cameras are being used as devices to collect information to be analyzed either live or at a later time (Mogili and Deepak, 2018; Kulbacki

et al., 2018). With these new technologies organizations and farmers are no longer bounded by the data collection time-window and now have the ability to manually analyze images at a later time. However, with this new approach comes additional problems such as storing massive amounts of image-based data and a familiar but new challenge - analyzing massive reserves of images accurately and efficiently. To advance modern agriculture, tools must be made easily available to agronomists to enable real-time decision making. The information contained within these images allows for timely management decision to optimize yield against harmful attack vectors (pests, diseases, drought, etc.).

To analyze large reserves of images quickly modern deep learning tools have been invoked by various crops. Recent literature has seen the combination of planting phenotyping and traditional machine learning techniques to count crops, detect color and classify stress in various crops through images (Singh et al., 2016; Naik et al., 2017; Yuan et al., 2018; Guo et al., 2018). These recent works demonstrate the impact traditional machine learning has on the future of agriculture. However, these methods are not without limitations. Using traditional approaches oftentimes, requires high quality images, constant lighting conditions, and fixed camera distances. These limitations act as a barrier to true HTP. With advances in state-of-the-art deep learning techniques, these constraints are no longer binding. Robust models can be constructed to analyze crops in numerous variable conditions. This is seen in the current literature combining deep learning and HTP.

Image-based phenotyping and deep learning can broadly be labeled as an application area of computer vision. Traditional tasks include classifying single images, counting objects and detection objects. Common deep learning frameworks using AlexNet, LeNet, and ResNet-50 architectures have been applied to classify various fruits and vegetables from single images (Mohanty et al., 2016; Cruz et al., 2017; Wang et al., 2017). Other deep learning models using VGG-16 as a feature extractor and the “You Only Look Once” model has been used to count and detect leaves, sorghum heads, and corn kernels (Giuffrida et al., 2018; Ghosal et al., 2019; Mosley et al., 2020; Khaki et al., 2020c; Redmon et al., 2016). Using novel frameworks to count corn tassels, Lu et al. (2017) combined convolutional neural networks (CNNs) and local counts regression into a framework they call TasselNet. Additionally, open-access, high-quality, annotated datasets are being created and released to the public to engage researchers in applying their deep learning knowledge to agriculture (Zheng et al., 2019; Sudars et al., 2020; Haug and Ostermann, 2014). These recent works showcase the potential for combining modern deep learning and agriculture in hopes of mitigating the so called “phenotyping bottleneck”. For the curious reader who is interested in a clear, concise, and thorough review of image-based HTP, we point the reader towards a survey paper by Jiang et al. (2020).

Corn is known to be one of the world’s most essential crop due to the number of products that it can create (e.g. flour, bio-fuels) (Berardi et al., 2019). Additionally, a large percentage of corn is used in livestock farming to feed pigs, cattle, and cows. The world’s reliance on corn cannot be understated. Aside from the manufacturing aspect, corn has a large impact on the United States’ economy. In 2019, it is estimated that the U.S. corn market contributed approximately \$140 billion to the U.S. economy. It is evident that the agricultural community and the world must act to ensure the continued optimal production of corn. By 2050, the world’s population is estimated to be approximately 9 billion (Stephenson et al., 2010). With the increase in population combined with the non-increasing arable land, changes will need to occur so that we can continue to optimize corn yield while utilizing less resources. Previous studies invoke deep learning based approaches to predict corn yield based using genetics, environment, and satellite imagery, but these studies are not considered HTP on commercial corn and only act as a way to estimate yield during the growing season (Khaki et al., 2019a; Russello, 2018; Khaki et al., 2019b; Khaki and Wang, 2019; Khaki et al., 2020a).

The ultimate goal of this paper is to count the number of corn stands in an image of a specific area on the field taken during the early phenological stages (VE to V6) (Ciampitti et al.). Roughly these phenological stages refer to the visible leaves on the stem. For instance, VE (emergence) is the first phenological stage where the stem breaks through the soil. V1 is the appearance of the first leaf. V2 is appearance of the second leaf, and VN is the appearance of the n -th leaf. From a practical perspective, an estimated stand count value allows for the establishment of a planting rate and, ultimately, yield potential. If the proper rate is planted, then farmers/breeders can estimate yield based on product by population. However, if population is not there i.e. poor germination/ bad planter, then farmers can identify the issue and replant, or at the very least establish what the new yield will be. Knowing that the planting density is below its desired threshold enables farmers to decide how they want to best manage their corn to make up for the difference in planting rating (e.g. more fertilizer, more aggressive pesticide control, etc.). Traditionally, farmers perform stand count estimations manually. However, this process is time consuming, labor intensive and prone to human error. Because of this, there is a reluctance to perform a stand count, ultimately, leaving farmers at a net-loss for the corn yield. Utilizing an image-based approach to this problem will allow for the timely estimation of stand counts and as well as a consistent measure to the quality of data.

Due to the need of efficient and effective HTP, in this paper, we present a deep learning based approach, named DeepStand, to alleviate the concerns of manual, labor intensive stand counting. The proposed method adopts a truncated VGG-16 network as a backbone feature extractor and merges multiple feature maps with different scales to make the

network robust against scale variation. This approach is similar to common methods in crowd counting where models are used to detect individual people in large crowds. Due to the similarity of these problems, we utilize a point density based approach for detecting the corn stands.

2 Methodology

Image-based corn stand counting is challenging compared to the counting tasks in other fields due to multiple factors, including occlusions, scale variations, and small distance between corn stands. Figure 1 shows the corn stands at different growing stages. This paper proposes a deep learning based method, DeepStand, to count the number of corn stands based on using a 180-degree image taken at 4-6 feet above the ground. It is worth mentioning that as corn progresses through its phenological stages, accurately counting the planting density becomes a difficult task for a computer due to the amount of overlapping leaves. However, from a pragmatic perspective, stand counting should be performed before V4 to ensure agronomists can act in a timely manner to mitigate any crop issues.

Figure 1: The images show corn stands at vegetative stages V1 to V6. The images include high scale variations, occlusions, and small distance between corn stands especially at stages V4 to V6.

2.1 Network Architecture

Corn stand images usually include high scale variations and occluded corn stands by nearby corn leaves. As a result, the proposed counting method should be robust against these factors. Our proposed stand counting method is inspired by methods proposed for the counting task in other fields such as crowd counting (Gao et al., 2020a) and dense object counting (Gao et al., 2020b).

The architecture of the proposed network is outlined in Figure 2. The proposed network generates a density map given an image of corn stands, where integral over the density map gives the total number of corn stands. Our proposed method is a CNN-based density estimation method. We do not use other methods such as detection-based (Li et al., 2008; Zeng and Ma, 2010) or regression-based (Chan et al., 2008; Idrees et al., 2013; Wang et al., 2015) methods for the following main reasons. Detection-based approaches usually apply an object detection method such as faster R-CNN (Ren et al., 2015), SSD (Liu et al., 2016) or a detector via a sliding window (Khaki et al., 2020b) on an image to first detect the objects and then count them. However, these approaches may not work well when applied to the scenes with occlusion and dense objects. Moreover, training these methods requires a considerable amount of annotated images, which is not publicly available for the task of corn stand counting. Regression-based approaches directly map an image patch to the count. These approaches deal with the problems of occlusion and background clutter successfully, however, they ignore the spatial information in the input image. As such, these approaches do not know how much each region of image contribute the final count (Gao et al., 2020b).

Figure 2: The outline of the DeepStand architecture. The parameters of the convolutional layers and deconvolutional layers are denoted as ‘Conv-(kernel size)-(number of filters)’ and ‘Deconv-(kernel size)-(number of filters)’, respectively. All layers have the stride of 1 except for the layers with “S2” notation, which have stride of 2. The padding type is the ‘same’ except for the first deconvolutional layer for which we use ‘valid’ padding. © denotes matrix concatenation.

We use a truncated VGG-16 (Simonyan and Zisserman, 2014) as a backbone for feature extraction in our proposed network. The truncated VGG-16 is composed of convolutional layers with a fixed kernel size of 3×3 that extracts discriminative features from input image for further analysis of the network. We use the VGG-16 network backbone in our network mainly because of its good generalization ability to other vision tasks such as counting and object detection (Liu et al., 2016; Gao et al., 2020b; Liu et al., 2019; Li et al., 2018). The truncated VGG-16 network includes all layers of VGG-16 network except the last max-pooling layer and all fully connected layer. The truncated VGG-16 shrinks the input images’ resolution to the 1/8 of its original size. We increase input size of the truncated VGG-16 network from 224×224 to 300×300 in our proposed network to learn more fine-grained features and patterns from the input image to further improve accuracy of our proposed method (Tan and Le, 2019).

The proposed network merges feature maps from multiple scales of the network to make it robust against scale variations in images. Similar scale-adaptive architectures have been used in other vision studies (Zhang et al., 2018; Ronneberger

et al., 2015; Bai et al., 2019). To concatenate feature maps with different spatial resolutions, we use zero padding to enlarge the smaller feature maps to the size of the largest feature map. Finally, we use three deconvolutional layers (transposed convolution) (Dumoulin and Visin, 2016) with stride of 2 to up-sample the output of the network to the size of the original input image.

Finally, we do a post processing on the predicted density map to draw a bounding box around each corn stand. The post processing includes the following steps: (1) threshold the estimated density map to zero-out regions where the value of density map is insignificant, (2) find peak coordinates on the density map as the center location of corn stands, (3) draw a bounding box around each corn stand, and (4) apply non-maximum suppression to remove overlapping bounding boxes. The above-mentioned post processing has a very low computational cost and does not increase the inference time.

2.2 Loss Function

Let I_i , D_i , $F(I_i, \Theta)$, Θ , and N denote i th image, the i th ground truth density map, the predicted density map of the i th image, network parameters, and the number of input images, respectively. As such, the network loss can be defined as below:

$$L(\Theta) = \frac{1}{N} \sum_{i=1}^N \|F(I_i, \Theta) - D_i\|_2^2 \quad (1)$$

Euclidean loss measures estimation error at pixel level and has been used in other crowd counting studies (Boominathan et al., 2016; Gao et al., 2020b; Lian et al., 2019).

3 Experiments and Results

This section first introduces the dataset used in our study, data augmentation, evaluation metrics, and training procedure. Then, we report the results of our proposed method along with other competing methods. All experiments were conducted in Tensorflow framework (Abadi et al., 2016) on a NVIDIA Tesla V100 SXM2 GPU.

3.1 Data

3.1.1 Ground Truth Density Maps Generation

We generate ground truth density maps following the procedure of density map generation in Boominathan et al. (2016) to train the network parameters. A corn stand located at pixel x_i can be represented by a delta function $\delta(x - x_i)$. As such, ground truth output for an image with M annotated corn stands can be defined as follows:

$$H(x) = \sum_{i=1}^M \delta(x - x_i) \quad (2)$$

Then, the $H(x)$ is convoluted with a Gaussian kernel with a standard deviation σ to generate the density map $D(x)$, where the standard deviation can be defined based on the average distance of k-nearest neighboring annotations. The summation over the density map is equal to the number of corn stands presented in the image.

$$D(x) = \sum_{i=1}^M \delta(x - x_i) * G_\sigma(x) \quad (3)$$

The use of such density maps as ground truth can help CNNs learn from the spatial information in images.

3.1.2 Stand Count Data

This section presents the procedure to prepare sufficient data to train and evaluate our proposed method. Our original dataset includes 394 images of corn stands at growing stages V1 to V6 with a fixed size of 1024×768 taken at 4-6 feet

above the ground. This includes a total of 6154 total stands across all images. The following table shows the summary of statistics of the dataset.

Number of Images	Resolution	Min	Max	Avg	Total
394	1024 × 768	5	31	15.62	6154

Table 1: The summary statistics of the corn stand dataset. The Min, Max, Avg, and Total denote the the minimum, maximum, average, and total number of corn stands in the dataset, respectively

We randomly selected 20% of the images (80 images) as test data and used the rest of the images as training data (314 images). We augment the training dataset to generate sufficient data to train our proposed method. To make the the proposed method robust against scale variation, we construct a multi-scale pyramidal representation of each training image following the work Boominathan et al. (2016). The multi-scale pyramidal representation of images includes scales of 0.4 to 1.3, incremented in steps of 0.1, times the original image resolution. Then, patches with size of 300×300 are cropped at random locations, which are followed by randomly flipping and adding Gaussian noise.

3.2 Evaluation Metrics

We use standard evaluation metrics, Mean Absolute Error (MAE) and Root Mean Squared Error (RMSE), to measure the counting performance of the proposed method. Generally, MAE indicate the accuracy of the results and RMSE measures the robustness. These two metrics are defined as follows:

$$MAE = \frac{1}{N} \sum_{i=1}^N |C_i^{pred} - C_i^{GT}| \quad (4)$$

$$RMSE = \sqrt{\frac{1}{N} \sum_{i=1}^N |C_i^{pred} - C_i^{GT}|^2} \quad (5)$$

where N is the number of test images, C_i^{pred} and C_i^{GT} are the estimated count and the ground truth count corresponding to the i th test image.

3.3 Training Hyperparameter

DeepStand network is trained end-to-end from scratch. The network parameters are initialized with Xavier initialization (Glorot and Bengio, 2010). Adam optimizer (Kingma and Ba, 2014) with learning rate of $3e-4$ and a mini-batch size of 24 are used to minimize the loss function defined in Equation (1). The learning rate is gradually decayed to $25e-6$ during the training process. The network is trained 80,000 iterations on 93,258 image patches generated following the data augmentation procedure in 3.1.2.

3.4 Comparison with State-of-the-art

To evaluate the efficiency of our proposed method, we compare our method with five state-of-the-art models. These five models were originally proposed for crowd counting problem, but they are also applicable for other object counting problems, which are as follows:

CSRNet: proposed by Li et al. (2018), uses a fully convolutional architecture which includes truncated VGG-16 network backbone as the front-end feature extractor and a set of dilated convolutional layers as back-end to estimate the density map.

SaCNN: proposed by Zhang et al. (2018), employs a scale-adaptive CNN network to cope with scale and perspective change in images. The SaCNN uses a backbone similar to VGG-16 architecture for feature extraction and merges feature maps from different layers of network to make the model robust against scale variation.

MSCNN: proposed by Zeng et al. (2017), extracts scale-relevant features using a multi-scale CNN network. The MSCNN network consists of multiple Inception-like (Szegedy et al., 2015) modules for multi-scale feature extraction.

CrowdNet: proposed by Boominathan et al. (2016), uses a multi-column network architecture. The network consists of a deep CNN and a shallow CNN to predict density map. The deep CNN has a VGG-like network architecture and the shallow CNN includes three convolutional layers.

DeepCrowd: proposed by Wang et al. (2015), is a regression-based method which directly learns a mapping from the image patches to the count. DeepCrowd network includes five convolutional layers which are followed by two fully connected layers.

3.5 Results

This section reports the evaluation results and compares our proposed method with other state-of-the-art methods for the task of corn stand counting. After having trained all methods, we evaluated their performance on the hold-out test data which includes 80 images of corn stands from growing stages V1 to V6. Table 2 illustrate the stand counting performances of the proposed and comparison methods on the test data with respect to MAE and RMSE evaluation metrics.

Method	MAE	RMSE
CSRNet (Li et al., 2018)	2.23	2.84
SaCNN (Zhang et al., 2018)	4.38	5.25
MSCNN (Zeng et al., 2017)	2.26	2.65
CrowdNet (Boominathan et al., 2016)	4.49	5.53
DeepCrowd (Wang et al., 2015)	2.61	3.13
DeepStand (our proposed)	1.73	2.46

Table 2: The counting performances of the proposed and comparison methods on the test data.

Table 2 illustrates that our proposed stand counting method outperforms the other methods to varying extents. The MSCNN has a comparable performance with CSRNet and both perform better than other methods except our proposed method. The main reason for the good performance of CSRNet is because of utilizing dilated convolutional layers to aggregate the multi-scale contextual information. The good performance of MSCNN can also be attributed to the use of multi-scale features which make it robust against scale variation. The DeepCrowd method as a regression-based method outperformed the SaCNN and the CrowdNet methods. The proposed method performed better than other methods for a couple of reasons: (1) our proposed method merges feature maps from multiple scales of the network to cope with scale variation, and (2) the use of deconvolution layers for up-sampling the feature maps increases the quality of the predicted density maps.

Even though the counting performance of the CSRnet, MSCNN, DeepCrowd, and SaCNN are good, the performances of these methods are not satisfactory when the whole image is fed to these methods at once. As a result, an input image should be cropped into some non-overlapping patches and fed to these methods which increases their inference time. Our proposed method and CrowdNet take the whole image and predict the density map in a single forward path, which makes them considerably faster than other methods.

Figure 3 visualizes some stand counting results of our proposed method including original image, predicted density map, and the detected corn stands in an image.

Figure 3: Visual results of our proposed stand counting method. The first, second and third rows indicate, respectively, input images, estimated density maps, and detected corn stand in images. GT and Pred stand for the ground truth and the predicted number of corn stands, respectively. The figure is best viewed in color.

In order to see how the growing stages (VE to V6) affect the counting performance of the proposed method, we divide the test data based on the growing stages into three classes of VE-V1, V2-V4, and V5-V6 and report the counting performances of the proposed and comparison methods on these classes. As shown in Table 3, the proposed method has a consistently low error across all growing stages which indicates the robustness of our proposed method. The results also indicate that the highest counting error of all methods except MSCNN belong to the stage V5-V6 due to the occlusion and background clutter.

Method	Growing stage					
	VE-V1		V2-V4		V5-V6	
	MAE	RMSE	MAE	RMSE	MAE	RMSE
CSRNet (Li et al., 2018)	2.34	3.12	1.51	1.89	3.04	3.56
SaCNN (Zhang et al., 2018)	5.87	6.28	2.25	2.49	6.24	6.90
MSCNN (Zeng et al., 2017)	2.72	2.98	2.37	2.76	1.91	2.32
CrowdNet (Boominathan et al., 2016)	3.71	4.88	4.43	4.94	5.97	6.41
DeepCrowd (Wang et al., 2015)	3.23	3.67	2.57	2.91	2.39	3.11
DeepStand (our proposed)	1.0	1.31	1.39	1.94	2.47	3.30

Table 3: The counting performances of the proposed and comparison methods on the test data across different growing stages.

4 Discussion and Conclusion

In this paper, we presented a deep learning based approach, named DeepStand, for corn stand counting problem. The proposed method adopts a truncated VGG-16 network as a backbone feature extractor and merges multiple feature maps

with different scales to make the network robust against scale variation. Finally, DeepStand uses a set of deconvolutional layers as a back-end to up-sample the network output. Our extensive experimental results indicate that our proposed method can successfully count and subsequently detect corn stands regardless of the image scale and the lighting condition. Our proposed method also outperformed other state-of-the-art methods commonly used in similar tasks.

The results of our proposed model, DeepStand, accurately detects and counts corn stands. This work gives hope towards the future of modernizing agriculture to enable true high-throughput phenotyping and, ultimately, better enabling agronomists with information to enhance decision making. Indeed, DeepStand provides accurate stand estimates for phenological stages VE-V4, where farmers traditionally desire stand counts. Simply due to the difficulty of overlapping leaves and background clutter at this moment we are not able to construct a model which can handle stages beyond V4. However, from a practical perspective, this is not necessary. When the continued development of deep learning, we hope to one day revisit this problem to construct a robust model that will work no matter the growth stage.

Although this work focuses on counting corn stands, a similar approach may be generalized to similar yet different applications such as detecting sorghum heads, counting wheat heads, detecting weeds, identifying planting gaps, etc. The open-challenges facing high-throughput phenotyping is vast and requires the joint collaboration of machine learning researchers to be able to solve these problems.

Conflicts of Interest

The authors declare no conflict of interest.

References

- Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., Devin, M., Ghemawat, S., Irving, G., Isard, M., et al., 2016. Tensorflow: A system for large-scale machine learning, in: 12th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 16), pp. 265–283.
- Bai, H., Wen, S., Gary Chan, S.H., 2019. Crowd counting on images with scale variation and isolated clusters, in: Proceedings of the IEEE International Conference on Computer Vision Workshops, pp. 0–0.
- Berardi, D., Hartman, M.D., DeLucia, E.H., Hudiburg, T.W., 2019. Flooding in the us corn belt: Mitigating climate change and crop loss by converting to flood tolerant bioenergy crops. AGUFM 2019, B33E–04.
- Boominathan, L., Kruthiventi, S.S., Babu, R.V., 2016. Crowdnet: A deep convolutional network for dense crowd counting, in: Proceedings of the 24th ACM international conference on Multimedia, pp. 640–644.
- Chan, A.B., Liang, Z.S.J., Vasconcelos, N., 2008. Privacy preserving crowd monitoring: Counting people without people models or tracking, in: 2008 IEEE Conference on Computer Vision and Pattern Recognition, IEEE. pp. 1–7.
- Ciampitti, I.A., Elmore, R.W., Lauer, J., . Corn growth and development. Dent 5, 75.
- Cruz, A.C., Luvisi, A., De Bellis, L., Ampatzidis, Y., 2017. X-fido: An effective application for detecting olive quick decline syndrome with deep learning and data fusion. Frontiers in plant science 8, 1741.
- Dumoulin, V., Visin, F., 2016. A guide to convolution arithmetic for deep learning. arXiv:1603.07285.
- Furbank, R.T., Tester, M., 2011. Phenomics–technologies to relieve the phenotyping bottleneck. Trends in plant science 16, 635–644.
- Gao, G., Gao, J., Liu, Q., Wang, Q., Wang, Y., 2020a. Cnn-based density estimation and crowd counting: A survey. arXiv preprint arXiv:2003.12783 .
- Gao, G., Liu, Q., Wang, Y., 2020b. Counting dense objects in remote sensing images, in: ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), IEEE. pp. 4137–4141.
- Ghosal, S., Zheng, B., Chapman, S.C., Potgieter, A.B., Jordan, D.R., Wang, X., Singh, A.K., Singh, A., Hirafuji, M., Ninomiya, S., et al., 2019. A weakly supervised deep learning framework for sorghum head detection and counting. Plant Phenomics 2019, 1525874.
- Giuffrida, M.V., Doerner, P., Tsafaris, S.A., 2018. Pheno-deep counter: a unified and versatile deep learning architecture for leaf counting. The Plant Journal 96, 880–890.
- Glorot, X., Bengio, Y., 2010. Understanding the difficulty of training deep feedforward neural networks, in: Proceedings of the thirteenth international conference on artificial intelligence and statistics, pp. 249–256.
- Guo, W., Zheng, B., Potgieter, A.B., Diot, J., Watanabe, K., Noshita, K., Jordan, D.R., Wang, X., Watson, J., Ninomiya, S., et al., 2018. Aerial imagery analysis–quantifying appearance and number of sorghum heads for applications in breeding and agronomy. Frontiers in plant science 9, 1544.

- Haug, S., Ostermann, J., 2014. A crop/weed field image dataset for the evaluation of computer vision based precision agriculture tasks, in: European Conference on Computer Vision, Springer. pp. 105–116.
- Idrees, H., Saleemi, I., Seibert, C., Shah, M., 2013. Multi-source multi-scale counting in extremely dense crowd images, in: Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 2547–2554.
- Jiang, Y., Li, C., et al., 2020. Convolutional neural networks for image-based high-throughput plant phenotyping: A review. *Plant Phenomics* 2020, 4152816.
- Khaki, S., Khalilzadeh, Z., Wang, L., 2019a. Classification of crop tolerance to heat and drought—a deep convolutional neural networks approach. *Agronomy* 9, 833.
- Khaki, S., Khalilzadeh, Z., Wang, L., 2020a. Predicting yield performance of parents in plant breeding: A neural collaborative filtering approach. *Plos one* 15, e0233382.
- Khaki, S., Pham, H., Han, Y., Kuhl, A., Kent, W., Wang, L., 2020b. Convolutional neural networks for image-based corn kernel detection and counting. *Sensors* 20, 2721.
- Khaki, S., Pham, H., Han, Y., Kuhl, A., Kent, W., Wang, L., 2020c. Deepcorn: A semi-supervised deep learning method for high-throughput image-based corn kernel counting and yield estimation. *arXiv preprint arXiv:2007.10521*.
- Khaki, S., Wang, L., 2019. Crop yield prediction using deep neural networks. *Frontiers in Plant Science* 10, 621.
- Khaki, S., Wang, L., Archontoulis, S.V., 2019b. A cnn-rnn framework for crop yield prediction. *Frontiers in Plant Science* 10.
- Kingma, D.P., Ba, J., 2014. Adam: A method for stochastic optimization. *arXiv:1412.6980*.
- Kulbacki, M., Segen, J., Kniec, W., Klempous, R., Kluwak, K., Nikodem, J., Kulbacka, J., Serester, A., 2018. Survey of drones for agriculture automation from planting to harvest, in: 2018 IEEE 22nd International Conference on Intelligent Engineering Systems (INES), IEEE. pp. 000353–000358.
- Li, M., Zhang, Z., Huang, K., Tan, T., 2008. Estimating the number of people in crowded scenes by mid based foreground segmentation and head-shoulder detection, in: 2008 19th international conference on pattern recognition, IEEE. pp. 1–4.
- Li, Y., Zhang, X., Chen, D., 2018. Csrnet: Dilated convolutional neural networks for understanding the highly congested scenes, in: Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 1091–1100.
- Lian, D., Li, J., Zheng, J., Luo, W., Gao, S., 2019. Density map regression guided detection network for rgb-d crowd counting and localization, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 1821–1830.
- Liu, C., Weng, X., Mu, Y., 2019. Recurrent attentive zooming for joint crowd counting and precise localization, in: 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), IEEE. pp. 1217–1226.
- Liu, W., Anguelov, D., Erhan, D., Szegedy, C., Reed, S., Fu, C.Y., Berg, A.C., 2016. Ssd: Single shot multibox detector, in: European conference on computer vision, Springer. pp. 21–37.
- Lu, H., Cao, Z., Xiao, Y., Zhuang, B., Shen, C., 2017. Tasselnet: counting maize tassels in the wild via local counts regression network. *Plant methods* 13, 79.
- McWilliams, D.A., Berglund, D.R., Endres, G., 1999. Corn growth and management quick guide .
- Mogili, U.R., Deepak, B., 2018. Review on application of drone systems in precision agriculture. *Procedia computer science* 133, 502–509.
- Mohanty, S.P., Hughes, D.P., Salathé, M., 2016. Using deep learning for image-based plant disease detection. *Frontiers in plant science* 7, 1419.
- Mosley, L., Pham, H., Bansal, Y., Hare, E., 2020. Image-based sorghum head counting when you only look once. *arXiv preprint arXiv:2009.11929*.
- Naik, H.S., Zhang, J., Lofquist, A., Assefa, T., Sarkar, S., Ackerman, D., Singh, A., Singh, A.K., Ganapathysubramanian, B., 2017. A real-time phenotyping framework using machine learning for plant stress severity rating in soybean. *Plant methods* 13, 23.
- Redmon, J., Divvala, S., Girshick, R., Farhadi, A., 2016. You only look once: Unified, real-time object detection, in: Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 779–788.
- Ren, S., He, K., Girshick, R., Sun, J., 2015. Faster r-cnn: Towards real-time object detection with region proposal networks, in: Advances in neural information processing systems, pp. 91–99.
- Ronneberger, O., Fischer, P., Brox, T., 2015. U-net: Convolutional networks for biomedical image segmentation, in: International Conference on Medical image computing and computer-assisted intervention, Springer. pp. 234–241.

- Russello, H., 2018. Convolutional neural networks for crop yield prediction using satellite images. IBM Center for Advanced Studies .
- Simonyan, K., Zisserman, A., 2014. Very deep convolutional networks for large-scale image recognition. *arXiv:1409.1556*.
- Singh, A., Ganapathysubramanian, B., Singh, A.K., Sarkar, S., 2016. Machine learning for high-throughput stress phenotyping in plants. *Trends in plant science* 21, 110–124.
- Stephenson, J., Newman, K., Mayhew, S., 2010. Population dynamics and climate change: what are the links? *Journal of Public Health* 32, 150–156.
- Sudars, K., Jasko, J., Namatevs, I., Ozola, L., Badaukis, N., 2020. Dataset of annotated food crops and weed images for robotic computer vision control. *Data in Brief* , 105833.
- Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., Rabinovich, A., 2015. Going deeper with convolutions, in: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1–9.
- Tan, M., Le, Q.V., 2019. Efficientnet: Rethinking model scaling for convolutional neural networks. *arXiv preprint arXiv:1905.11946* .
- Wang, C., Zhang, H., Yang, L., Liu, S., Cao, X., 2015. Deep people counting in extremely dense crowds, in: *Proceedings of the 23rd ACM international conference on Multimedia*, pp. 1299–1302.
- Wang, G., Sun, Y., Wang, J., 2017. Automatic image-based plant disease severity estimation using deep learning. *Computational intelligence and neuroscience* 2017.
- Yang, W., Feng, H., Zhang, X., Zhang, J., Doonan, J.H., Batchelor, W.D., Xiong, L., Yan, J., 2020. Crop phenomics and high-throughput phenotyping: past decades, current challenges, and future perspectives. *Molecular Plant* 13, 187–214.
- Yuan, W., Li, J., Bhatta, M., Shi, Y., Baenziger, P.S., Ge, Y., 2018. Wheat height estimation using lidar in comparison to ultrasonic sensor and uas. *Sensors* 18, 3731.
- Zeng, C., Ma, H., 2010. Robust head-shoulder detection by pca-based multilevel hog-lbp detector for people counting, in: *2010 20th International Conference on Pattern Recognition, IEEE*. pp. 2069–2072.
- Zeng, L., Xu, X., Cai, B., Qiu, S., Zhang, T., 2017. Multi-scale convolutional neural networks for crowd counting, in: *2017 IEEE International Conference on Image Processing (ICIP), IEEE*. pp. 465–469.
- Zhang, L., Shi, M., Chen, Q., 2018. Crowd counting via scale-adaptive convolutional neural network, in: *2018 IEEE Winter Conference on Applications of Computer Vision (WACV), IEEE*. pp. 1113–1121.
- Zheng, Y.Y., Kong, J.L., Jin, X.B., Wang, X.Y., Su, T.L., Zuo, M., 2019. Cropdeep: the crop vision dataset for deep-learning-based classification and detection in precision agriculture. *Sensors* 19, 1058.