

Crop Lodging Prediction from UAV-Acquired Images of Wheat and Canola using a DCNN Augmented with Handcrafted Texture Features

Sara Mardanisamani^{*1}, Farhad Maleki^{*1}, Sara Hosseinzadeh Kassani¹, Sajith Rajapaksa¹, Hema Duddu², Menglu Wang², Steve Shirliffe², Seungbum Ryu², Anique Josuttes², Ti Zhang², Sally Vail³, Curtis Pozniak², Isobel Parkin³, Ian Stavness¹ and Mark Eramian¹

¹Department of Computer Science
University of Saskatchewan
Saskatoon, SK, Canada

sara.samani@usask.ca, farhad.maleki@usask.ca

mark.eramian@usask.ca

²Department of Plant Sciences
University of Saskatchewan
Saskatoon, SK, Canada

³Agriculture and Agri-food Canada
Saskatoon Research and Development Center
Saskatoon, SK, Canada

Abstract

Lodging, the permanent bending over of food crops, leads to poor plant growth and development. Consequently, lodging results in reduced crop quality, lowers crop yield, and makes harvesting difficult. Plant breeders routinely evaluate several thousand breeding lines, and therefore, automatic lodging detection and prediction is of great value aid in selection. In this paper, we propose a deep convolutional neural network (DCNN) architecture for lodging classification using five spectral channel orthomosaic images from canola and wheat breeding trials. Also, using transfer learning, we trained 10 lodging detection models using well-established deep convolutional neural network architectures. Our proposed model outperforms the state-of-the-art lodging detection methods in the literature that use only handcrafted features. In comparison to 10 DCNN lodging detection models, our proposed model achieves comparable results while having a substantially lower number of parameters. This makes the proposed model suitable for applications such as real-time classification using inexpensive hardware for high-throughput phenotyping pipelines. The GitHub repository at <https://github.com/FarhadMaleki/LodgedNet> contains code and models.

1. Introduction

Lodging occurs when plant stems break or bend over so that plants are permanently displaced from their optimal upright position. It is a common problem for many crops, including wheat and canola, and can be caused by external forces, including wind, rain, or hail [23], and morphological factors, such as thin or weak stem structures.

Multiple studies on rice, wheat, and oats have shown that lodging can cause grain yield loss and deterioration in seed quality [10]. In addition, lodging can cause problems for harvest operations often resulting in increasing the demand for grain drying, which raises production costs. In most crops, severe lodging results in as much as a 50% yield reduction [2]. For plant breeders, it is important to identify lodging-resistant varieties from thousands of experimental plots. Therefore, automatic lodging detection methods from overhead images are valuable. In addition, crop insurance claims following wind/hail storms currently require manual assessment of crop damage. Lodging detection from aerial drone-acquired images could provide a faster and more accurate assessment of the area and severity of lodging within a field.

Researchers have developed some field methods to cope with lodging and determine the best way to obtain the maximum harvestable product. However, manual in-the-field assessment requires estimating the part of the field that is lodged and the degree of lodging plants relative to their vertical axis. These manual measurements are time-consuming

* Co-first authors.

and costly [1], as well as quite subjective. To monitor crop lodging within an entire field, unmanned aerial vehicles (UAVs) may automatically collect high-resolution aerial images to detect the lodging in a simple, flexible, cost-effective way [25]. Then image analysis techniques can be used for automatic lodging detection.

We propose a deep convolutional neural network architecture that couples handcrafted and learned features to detect lodging. Previously in the literature, only methods based on handcrafted features have been used for lodging detection from images [14, 22, 24, 25]. Deep convolutional neural networks (DCNNs) have been successfully applied to a wide range of image classification tasks [4, 6, 7, 9, 15, 19]. However, to the best of our knowledge, they have not been used for lodging detection. In general, methods that only rely on handcrafted features often achieve lower accuracy in comparison to deep CNN models. They are also sensitive to noise. DCNN-based models, on the other hand, often disregard research on problem-specific handcrafted features. Furthermore, they need a substantial amount of training data to achieve high accuracy.

To avoid the shortcomings of both approaches and to benefit from their strengths, in this paper, we propose LodgedNet, an architecture that uses a DCNN-based model together with two texture feature descriptors: local binary patterns (LBP) and gray-level co-occurrence matrix (GLCM) for crop lodging classification. We also developed 10 DCNN-based models using well-established architectures. LodgedNet is designed to offer rapid training and prediction time while achieving accuracy comparable to that of the 10 DCNN-based models. To the best of our knowledge, there is no DCNN-based model in the literature applied to lodging detection. This work offers a comprehensive study of CNN architectures (including LodgedNet) for lodging detection.

The rest of the paper is organized as follows. Section 2 presents the related work on lodging detection. In Section 3, we describe the lodging datasets for wheat and canola and present the proposed architecture. Section 4 presents experimental results. Section 5 discusses the results and the utility of the proposed architecture for similar applications in agriculture. Finally, Section 6 ends the paper with a short summary and conclusion.

2. Related Works

Rajapaksa et al. [14] used handcrafted features and a support vector machine (SVM) for lodging classification using data obtained from drone imagery. They extracted features using a gray level co-occurrence matrix (GLCM), local binary patterns (LBP), and Gabor filters. Then they used an SVM to classify the feature vectors extracted from images of lodged and non-lodged plots. Their method was designed for grayscale images and they used information only from a

single image channel to predict lodging.

Wang et al. [22] proposed a method for lodging detection using pixel information obtained from wheat plot images taken by drones. They calculated nine colour features based on pixel values. In addition, they obtained 13 features from the nine colour features using the ENvironment for Visualizing Images (ENVI) software. They then used a thresholding approach to discriminate lodged pixels. Their approach solely relies on thresholding and pixel values and disregards spatial information and high correlation among pixel neighbourhoods. This makes the result sensitive to noise.

Yang et al. [25] proposed a spectral and spatial hybrid image classification method to detect rice lodging using images taken by drones. They obtained spatial information, including height data, using the IBM 3D construction algorithm and texture features. In addition to spatial information of the field, they extracted spectral information of each pixel using single feature probability (SFP). Then the extracted features were used by a decision tree classifier and a maximum likelihood classifier to detect lodging. Using RGB images and texture features, they achieved an accuracy of 88.14%. Incorporating pixel-wise spatial information, they achieved the accuracy values of 90.76% and 96.17% using the maximum likelihood classifier and the decision tree classifier, respectively. However, extracting spatial features requires using extra equipment that makes this approach expensive and time-consuming.

Yang et al. [24] (a different group of researchers) used satellite data (RADARSAT-2) to detect lodging in wheat fields. They extracted a set of sensitive polarimetric features and backscattering intensity features from five consecutive RADARSAT-2 images throughout the entire growing season to detect lodging. Using this approach they were able to identify lodged fields. However, plant breeders require lodging detection on the much smaller scale of their breeding plots which typically rectangular plots measuring a few meters on each side. Satellite imagery cannot provide the required resolution for assessment of individual breeder plots of this size.

In this research, we design an architecture based on handcrafted and DCNN features. Combining these types of features have been proposed previously [5, 12, 21]. Using an image along with Gabor filters extracted from that image as input to the network, Hosseini et al. [5] achieved a higher accuracy compared to several traditional and CNN-based models. Wang et al. [21], using a cascaded approach based on combining a CNN model and handcrafted features, proposed a computationally efficient model for counting the number of cells undergoing mitosis. Nguyen et al. [12] used a CNN model and multi-level local binary pattern (MLBP) for presentation attack detection (PAD) in face recognition. They combined features extracted from a

CNN-based model and the multi-level local binary pattern (MLBP) method to build a support vector machine classifier. Their model achieved a higher accuracy compared to previous PAD methods.

To the best of our knowledge, all of the published work on crop lodging detection have been developed using handcrafted features tailored to one or a few specific types of crops. Although models based on handcrafted features are often computationally efficient and applicable even in situations where we do not have access to a large number of training examples, these models often have been designed for a specific crop type and might not achieve a comparable accuracy when applied to other crop types. Furthermore, adjusting a handcrafted feature to a different task often is not straightforward and requires further research. Deep convolutional neural networks (DCNN), on the other hand, have proven to be an effective approach in machine vision. However, DCNN-based models often require a large amount of training data. In this paper, we propose a model that benefits from the strengths of both handcrafted features and also DCNNs, and avoids the shortcomings of these approaches.

3. Materials and Methods

3.1. Data Set Description

The wheat and canola datasets used in this study were obtained from two breeding field trials. Plot images were taken with a Draganfly X4P quad-copter (Draganfly Innovations Inc., Saskatoon, SK, Canada) carrying a MicaSense RedEdge camera (Micasense Inc. Seattle, WA, USA) in the summer of 2016. This camera captures images with five spectral channels: red, blue, green, near infrared, and red-edge. Agisoft Photoscan (Agisoft LLC, St. Petersburg, Russia) was used to stitch images and obtain a high-resolution orthomosaic image of each field. The ground resolution of the obtained images is approximately 15 to 26 mm/pixel for each band. Images were taken at a height of 20 meters for canola and 30 meters for wheat. Figure 1 shows the orthomosaic image extracted from the red, green, and blue channels.

3.2. Data

Two orthomosaic images for wheat (9492×8340 pixels) and seven orthomosaic images for canola (9492×8340 pixels) are used in this study. Each field is organized into several columns, and each column is divided into several small plots (see Figure 1). Table 1 represents the number of extracted plots for wheat and canola categorized as lodged and non-lodged.

We extracted image samples with dimensions 60×100 and 118×348 pixels from wheat and canola plots, respectively (the different sample sizes are due to the wheat and



Figure 1: An RGB orthomosaic image for a wheat trial.

Samples	Non-lodged	Lodged	Total
Wheat	285	180	465
Canola	1170	468	1638

Table 1: Number of samples for wheat and canola datasets for each class.

canola plots being of different physical size). Ground truth labels of “lodged” or “not lodged” were provided by a crop agronomist and a plant scientist. Figure 2 illustrates lodged and non-lodged samples from wheat and canola plots.

We randomly split each dataset into training and test sets. We used about 80% of the images for training and the remaining 20% for the test set. Training data were split randomly with 80% of training samples used for a training set and the remaining 20% for a validation set. Table 2 shows the number of samples in the training, validation, and test sets.

Data	Wheat		Canola	
	NL	L	NL	L
Train	187	113	754	300
Validation	48	28	188	75
Test	50	39	228	93

Table 2: The number of non-lodged (NL) and lodged (L) samples in the training, validation, and test sets.

3.3. The proposed architecture

Convolutional neural networks are a well-established approach in computer vision contributing to the success of applications in image segmentation, object detection, and




Figure 2: RGB images for wheat and canola: a) a non-lodged wheat plot; b) a lodged wheat plot; c) a non-lodged canola plot; and d) a lodged canola plot.

image classification [11, 16, 18]. In this section, we propose LodgedNet, which is a DCNN-based architecture for lodging detection. LodgedNet, as illustrated in Figure 3, consists of three main components: a CNN backbone, a texture feature extraction module, and a classification module. The CNN backbone consists of seven convolutional layers, with 16, 16, 32, 32, 32, 32, and 64 filters respectively. We used filters of size three, paddings of size one, and Rectified Linear Unit (ReLU) activation functions. A *SpatialDropout* [20] layer with a dropout rate of 0.5 was used after each convolutional layer, except for the first and third layers, to prevent overfitting. Max-pooling layers are used before the third, fifth, and seventh convolutional layers, and after the seventh convolutional layer (See Figure 3). The texture extraction module extracts LBP and GLCM features from each channel of the input images. These features are image properties related to second-order statistics that account for the spatial inter-dependencies of pixels at specific relative positions. The extracted texture features and the features generated from the last layer of the CNN module are flattened and concatenated to be used as input to the classification module. The classification module consists of two fully connected layers, respectively of 128 and two neurons, followed by a final softmax layer. To further prevent overfitting one dropout layer [17] was applied before, and another after, the first fully connected layer. A dropout rate of 0.5 was used for these layers. This architecture has been designed to achieve high accuracy while having a small number of parameters to facilitate training

and deployment on low-cost hardware.

4. Experimental Results

4.1. Feature Extraction Module

In this section, we describe the model specification for LodgedNet as well as the implementation details. The LBP [13] and GLCM [3] texture extractor methods used in the proposed architecture were applied with the following parameters. They are the same as those selected by Rajapaksa et al. [14] in their work which uses only GLCM and LBP for lodging classification. We extracted 16 *contrast* texture features with GLCM using the normalized and asymmetric 2D co-occurrence matrices for each channel of an input image using four orientations (0° , 90° , 135° , 180°) and four distances (1px, 2px, 4px, 5px). For LBP, we used 8-bit rotationally-invariant uniform local binary patterns and constructed 10-bin LBP histograms to be used as features. In addition, we extracted 16-bin LBP variance histograms and concatenated them with the 10 LBP features as recommended by Ojala et al. [13]. For each image channel, 16 GLCM features and 26 LBP features were extracted. Then these features were concatenated to be used as an input to the classification module. The LBP and GLCM features obtained using the “greycomatrix”, “greycoprops”, and “local_binary_pattern” functions from “scikit-image” library.

4.2. Data Augmentation and Training

It is common practice to perform data augmentation to further increase the generalizability of trained models. We employed a number of transformations for data augmentation. We resized the original images while preserving their aspect ratios and then center-cropped the resized images to achieve fixed-size inputs of 64×128 for LodgedNet and fixed-size inputs of 224×224 for all well-established models except Inception-V3, where a fixed-size input of 299×299 is required. A cropped image undergoes a vertical and a horizontal flipping each independently with a probability of 0.5. Each image was normalized by subtracting the mean and dividing by the standard deviation, where mean and standard deviation were calculated as the per-channel mean and standard deviation of the training and validation data.

All DCNN architectures were implemented in Python using the PyTorch package Version 1.0 on a Intel Core i7-5930K 3.5 GHz processor and NVIDIA GTX 1080 Ti with 11 GB graphical processing unit (GPU) and 32 GB RAM.

LodgedNet was trained on both wheat and canola samples from scratch, while for the 10 well-established models, a version of the models pretrained with ImageNet data [9] were used. To make a fair comparison, for the 10 well-established models we redefined the classifier component of each model to match that of LodgedNet’s classification

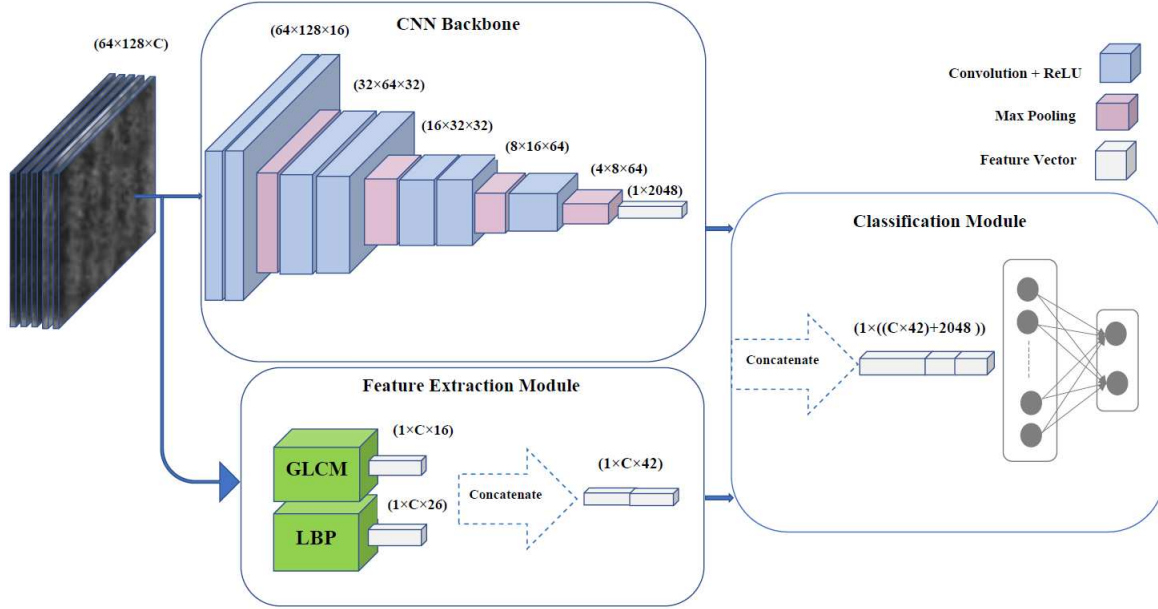


Figure 3: A schematic diagram of LodgedNet, the proposed architecture. All convolutional and fully connected layers are followed by a dropout layer with a ratio of 0.5 except convolutional layers one and three; C represents the number of image channels.

module. The pretrained models were further trained using two approaches. In the first approach, we let all parameters of the models be further trained using the lodging datasets. In the second approach, we froze all parameters other than the parameters of the classifier component, which were learned using the lodging datasets. In all experiments, an Adam optimizer with a learning rate of 0.001, and $\beta_1 = 0.9$, and $\beta_2 = 0.999$ were used. Also, we used a batch size of 16 and number of epochs equal to 50.

Table 3 shows the results of LodgedNet and the 10 well-established models for wheat and canola datasets, respectively. Also, Table 4 illustrates the results of two state-of-the-art methods using handcrafted features. These results represent the first applications of DCNNs to the problem of lodging detection.

5. Discussion

As can be observed from Table 3 and 4, LodgedNet achieved comparable results to the 10 DCNN-based classifiers and outperformed Rajapaksa et al.'s [14] state-of-the-art model for both crops.

Table 3 shows the number of parameters for the architectures used in this study. LodgedNet has a substantially lower number of parameters compared to all of the other DCNN-based classifiers. For example, the number of parameters in VGG19 is about 432 times more than that of LodgedNet. Among the DCCN-based models, SqueezeNet has the smallest number of parameters, which is almost 3.7

times more than that of LodgedNet. Despite the substantially smaller number of parameters, LodgedNet was among the top three methods with the highest test accuracy for the canola dataset and only misclassified two test samples from the wheat dataset.

Furthermore, the small number of parameters makes deploying LodgedNet on low-cost hardware possible, leading to near real-time inference. Using drone imagery along with mobile or low-cost portable computers, models such as LodgedNet that are computationally less demanding will be accessible to a wide range of agricultural applications. However, a sound comparison of the inference time for LodgedNet and the 10 DCNN-based models is difficult to achieve due to several factors: LodgedNet works with both three- and five-channel images, the input image dimensions of the tested networks differ, and the DCNN backbones of the 10 lodging detection models are highly-optimized implementations from the *torchvision* package of PyTorch. As a coarse comparison of these models, we used three-channel (RGB) images to estimate the inference time as the average of inference times for test samples. As depicted in Table 3, LodgedNet comes second to AlexNet in prediction time while achieving comparable accuracy for wheat and a substantially higher accuracy for canola (about 10% higher) when compared to AlexNet.

In this paper, we focused on proposing an architecture that is not computationally demanding and makes using handcrafted features possible. We used handcrafted texture

Architecture	Wheat	Wheat (FW)	Canola	Canola (FW)	Number of Parameters	Prediction Time Canola (ms) $\mu \pm \sigma$	Prediction Time Wheat (ms) $\mu \pm \sigma$
VGG19 [15]	98.68%	97.75%	98.75%	99.06%	143,667,240	5.75 \pm 0.92	9.99 \pm 0.00
VGG16 [15]	98.68%	96.62%	97.19%	98.44%	138,357,544	4.99 \pm 0.75	4.99 \pm 0.00
AlexNet [9]	98.85%	97.75%	89.09%	90.65%	61,100,840	2.95 \pm 0.65	2.83 \pm 0.37
ResNet101 [4]	100%	95.50%	98.44%	98.44%	44,549,160	34.94 \pm 4.67	41.66 \pm 6.87
Inception-V3 [19]	100%	97.75%	99.68%	100%	27,161,264	36.66 \pm 6.42	33.13 \pm 4.52
ResNet50 [4]	97.70%	100%	99.06%	99.06%	25,557,032	18.88 \pm 3.77	21.66 \pm 3.73
DensNet201 [6]	98.85%	100%	98.75%	99.06%	20,013,928	81.39 \pm 7.21	88.33 \pm 14.62
DensNet169 [6]	100%	100%	98.75%	99.37%	14,149,480	67.12 \pm 7.08	71.66 \pm 10.67
ResNet18 [4]	100%	100%	99.37%	98.44%	11,689,512	7.18 \pm 1.05	7.49 \pm 0.5
SqueezeNet [7]	97.70%	98.87%	98.75%	98.75%	1,235,496	10.46 \pm 1.73	9.99 \pm 5.77
LodgedNet	97.70%	-	99.06%	-	332,306	2.99 \pm 0.31	3.49 \pm 0.49

Table 3: Comparison between the proposed model and the 10 lodging detection models developed based on well-established architectures. Red, green, and blue channels were used as input images for wheat and canola; FW denotes the case of training with frozen weights; μ and σ denote average and standard deviation of the prediction time in milliseconds. The models are presented in descending order of number of parameters.

Case	Wheat			Canola			
	Methods	RGB	Five Channel	Rededge	RGB	Five Channel	Rededge
GLCM-based [14]	-	-	-	84.94%	-	-	72.86%
LBP-based [14]	-	-	-	96.77%	-	-	90.54%
LodgedNet	97.70%	97.33%	97.33%	-	99.06%	99.38%	-

Table 4: Comparison between our method and Rajapaksa’s method [14] (state-of-the-art in the literature for lodging classification) which used two handcrafted features for wheat and canola datasets on only one image channel.

feature extractors for lodging detection. However, the use of handcrafted features is not limited to lodging detection and such features are available for various application domains. Using these features can help to increase model accuracy, more specifically in domains where there is not a large amount of training data available.

In this study, we used five-channel images taken by drones for lodging detection. The results of our experiments showed that even in the absence of red edge and near-infrared channels, lodging detection can be performed with high accuracy. However, red-edge and near-infrared channels might contribute to achieving higher accuracy in other agricultural applications. We suggest using these channel data, which are supported by LodgedNet, for other agricultural applications.

5.1. Future Work and Limitations

LodgedNet used two handcrafted feature extraction approaches, namely GLCM and LBP. However, LodgedNet is not limited to using these two feature extractors. Extending LodgedNet to use other handcrafted features is suggested as future research. Because lodging tends to make the texture of the overhead plot views more directional, we expect

that using Gabor filters [8] alongside GLCM and LBP might lead to higher accuracy.

In this paper, we used LodgedNet for lodging classification. However, the proposed architecture can be used to tackle other image classification problems where there are handcrafted features available from previous research. Although models that only rely on handcrafted features often lead to lower accuracy and higher sensitivity to noise, in comparison to their DCNN-based counterparts, incorporating these handcrafted features in an architecture similar to LodgedNet could potentially help with improving accuracy. This approach could be of more value where there are relatively few training samples available.

A limitation of our assessment of prediction time is that it did not take into account the time to compute the LBP and GLCM features. This took about 30ms on average for a plot image input to LodgedNet because they were computed with a single-threaded CPU algorithm but this is not directly comparable to the DCNN forward pass prediction times since those computations were performed on a GPU. A parallelized implementation would greatly reduce the LBP and GLCM feature computation time. Regardless of this cost, the lower forward pass computation time of

LodgedNet should reduce training time because the LBP and GLCM features can be pre-computed and do not need to be extracted repeatedly for each training epoch or during the tuning of hyper-parameters.

One limitation of the transfer learning used for the 10 DCNN-based models is that they have been trained using RGB images and therefore the transfer learning using images that use extra channels such as red edge and near infrared cannot be used if we choose to use transfer learning. Considering the number of parameters for these models, training them from scratch requires a substantially large number of training samples; otherwise, considering their large capacity, they tend to overfit and the trained models are less likely to be generalizable to unseen samples. LodgedNet, however, can be used with a variable number of input image channels because the number of channels, C , is a hyper-parameter. Our extensive use of *SpatialDropout* [20] and regular dropout layers [17] provides resilience to overfitting.

6. Conclusion

In this paper, we used DCNNs to address crop lodging classification. We trained 10 models based on well-established DCNN architectures pre-trained on ImageNet data. We then proposed a new architecture, LodgedNet, that utilizes both DCNN and handcrafted feature extractors to build a lodging prediction network. Our comprehensive study of lodging prediction compared 11 DCNN architectures, including LodgedNet. LodgedNet, as well as all of the other DCNN architectures tested, outperformed current state-of-the-art lodging detection models using only handcrafted features. LodgedNet's prediction accuracy compares favourably with the 10 other architectures tested while having about one quarter of the number of trainable parameters compared to SqueezeNet, which is the next smallest network. The fewer number of parameters in LodgedNet accelerates training and inference time. It also facilitates deploying LodgedNet on low cost hardware.

LodgedNet can be used for lodging detection in high-throughput plant phenotyping scenarios. Such pipelines will be critical in the future search for higher yield crops varieties needed to feed a growing population.

7. Acknowledgements

This research was enabled thanks to funding from the Canada First Research Excellence Fund.

References

[1] P. Berry, J. Spink, R. Sylvester-Bradley, A. Pickett, M. Sterling, C. Baker, and N. Cameron. Lodging control through variety choice and management. In *Proceedings of the 8th*

Home-Grown Cereals Association R & D Conference on Cereals and Oilseeds, pages 7–1, 2002.

[2] P. Berry, M. Sterling, J. Spink, C. Baker, R. Sylvester-Bradley, S. Mooney, A. Tams, and A. Ennos. Understanding and reducing lodging in cereals. *Advances in Agronomy*, 84(04):215–269, 2004.

[3] R. M. Haralick, K. Shanmugam, et al. Textural features for image classification. *IEEE Transactions on Systems, Man, and Cybernetics*, (6):610–621, 1973.

[4] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016.

[5] S. Hosseini, S. H. Lee, and N. I. Cho. Feeding hand-crafted features for enhancing the performance of convolutional neural networks. *arXiv preprint arXiv:1801.07848*, 2018.

[6] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4700–4708, 2017.

[7] F. N. Iandola, S. Han, M. W. Moskewicz, K. Ashraf, W. J. Dally, and K. Keutzer. Squeezenet: Alexnet-level accuracy with 50x fewer parameters and 0.5 mb model size. *arXiv preprint arXiv:1602.07360*, 2016.

[8] A. K. Jain, S. Prabhakar, L. Hong, and S. Pankanti. Filterbank-based fingerprint matching. *IEEE Transactions on Image Processing*, 9(5):846–859, 2000.

[9] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems*, pages 1097–1105, 2012.

[10] B. Ma, D. Biswas, Q. Zhou, and C. Ren. Comparisons among cultivars of wheat, hulled and hullless oats: Effects of n fertilization on growth and yield. *Canadian Journal of Plant Science*, 92(6):1213–1222, 2012.

[11] S. P. Mohanty, D. P. Hughes, and M. Salathé. Using deep learning for image-based plant disease detection. *Frontiers in plant science*, 7:1419, 2016.

[12] D. T. Nguyen, T. D. Pham, N. R. Baek, and K. R. Park. Combining deep and handcrafted image features for presentation attack detection in face recognition systems using visible-light camera sensors. *Sensors*, 18(3):699, 2018.

[13] T. Ojala, M. Pietikäinen, and T. Mäenpää. Multiresolution gray-scale and rotation invariant texture classification with local binary patterns. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, (7):971–987, 2002.

[14] S. Rajapaksa, M. Eramian, H. Duddu, M. Wang, S. Shirliffe, S. Ryu, A. Josuttis, T. Zhang, S. Vail, C. Pozniak, and I. Parkin. Classification of crop lodging with gray level co-occurrence matrix. In *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 251–258, March 2018.

[15] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.

[16] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *CoRR*, abs/1409.1556, 2014.

- [17] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1):1929–1958, 2014.
- [18] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1–9, June 2015.
- [19] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2818–2826, 2016.
- [20] J. Tompson, R. Goroshin, A. Jain, Y. LeCun, and C. Bregler. Efficient object localization using convolutional networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 648–656, 2015.
- [21] H. Wang, A. Cruz-Roa, A. Basavanthally, H. Gilmore, N. Shih, M. Feldman, J. Tomaszewski, F. Gonzalez, and A. Madabhushi. Cascaded ensemble of convolutional neural networks and handcrafted features for mitosis detection. In *Medical Imaging 2014: Digital Pathology*, volume 9041, page 90410B. International Society for Optics and Photonics, 2014.
- [22] J.-J. Wang, H. Ge, Q. Dai, I. Ahmad, Q. Dai, G. Zhou, M. Qin, and C. Gu. Unsupervised discrimination between lodged and non-lodged winter wheat: a case study using a low-cost unmanned aerial vehicle. *International Journal of Remote Sensing*, 39(8):2079–2088, 2018.
- [23] W. Wu and B.-L. Ma. A new method for assessing plant lodging and the impact of management options on lodging in canola crop production. *Scientific reports*, 6:31890, 2016.
- [24] H. Yang, E. Chen, Z. Li, C. Zhao, G. Yang, S. Pignatti, R. Casa, and L. Zhao. Wheat lodging monitoring using polarimetric index from radarsat-2 data. *International Journal of Applied Earth Observation and Geoinformation*, 34:157–166, 2015.
- [25] M.-D. Yang, K.-S. Huang, Y.-H. Kuo, H. P. Tsai, and L.-M. Lin. Spatial and spectral hybrid image classification for rice lodging assessment through uav imagery. *Remote Sensing*, 9(6):583, 2017.