

## Research Article

**Cite this article:** Sharpe SM, Schumann AW, and Boyd NS (2019) Detection of Carolina Geranium (*Geranium carolinianum*) Growing in Competition with Strawberry Using Convolutional Neural Networks. *Weed Sci* 67:239–245. doi: 10.1017/wsc.2018.66

Received: 28 June 2018  
Revised: 7 August 2018  
Accepted: 10 August 2018

**Associate Editor:**

Ramon G. Leon, North Carolina State University

**Key words:**

Deep-learning GPU training system; *Fscore*; image classification; machine vision; object detection

**Author for correspondence:**

Nathan S. Boyd, University of Florida, Gulf Coast Research and Education Center, 14625 Count Road 672, Wimauma, FL 33598. (Email: nsboyd@ufl.edu)

# Detection of Carolina Geranium (*Geranium carolinianum*) Growing in Competition with Strawberry Using Convolutional Neural Networks

Shaun M. Sharpe<sup>1</sup>, Arnold W. Schumann<sup>2</sup> and Nathan S. Boyd<sup>3</sup>

<sup>1</sup>Postdoctoral Associate, University of Florida, Gulf Coast Research and Education Center, Wimauma, FL, USA,

<sup>2</sup>Professor, University of Florida, Citrus Research and Education Center, Lake Alfred, FL, USA and <sup>3</sup>Associate Professor, University of Florida, Gulf Coast Research and Education Center, Wimauma, FL, USA

**Abstract**

Weed interference during crop establishment is a serious concern for Florida strawberry [*Fragaria* × *ananassa* (Weston) Duchesne ex Rozier (pro sp.) [*chiloensis* × *virginiana*]] producers. In situ remote detection for precision herbicide application reduces both the risk of crop injury and herbicide inputs. Carolina geranium (*Geranium carolinianum* L.) is a widespread broadleaf weed within Florida strawberry production with sensitivity to clopyralid, the only available POST broadleaf herbicide. *Geranium carolinianum* leaf structure is distinct from that of the strawberry plant, which makes it an ideal candidate for pattern recognition in digital images via convolutional neural networks (CNNs). The study objective was to assess the precision of three CNNs in detecting *G. carolinianum*. Images of *G. carolinianum* growing in competition with strawberry were gathered at four sites in Hillsborough County, FL. Three CNNs were compared, including object detection–based DetectNet, image classification–based VGGNet, and GoogLeNet. Two DetectNet networks were trained to detect either leaves or canopies of *G. carolinianum*. Image classification using GoogLeNet and VGGNet was largely unsuccessful during validation with whole images (*Fscore* < 0.02). CNN training using cropped images increased *G. carolinianum* detection during validation for VGGNet (*Fscore* = 0.77) and GoogLeNet (*Fscore* = 0.62). The *G. carolinianum* leaf–trained DetectNet achieved the highest *Fscore* (0.94) for plant detection during validation. Leaf-based detection led to more consistent detection of *G. carolinianum* within the strawberry canopy and reduced recall-related errors encountered in canopy-based training. The smaller target of leaf-based DetectNet did increase false positives, but such errors can be overcome with additional training images for network desensitization training. DetectNet was the most viable CNN tested for image-based remote sensing of *G. carolinianum* in competition with strawberry. Future research will identify the optimal approach for in situ detection and integrate the detection technology with a precision sprayer.

**Introduction**

Carolina geranium (*Geranium carolinianum* L.) is a problematic broadleaf weed in Florida strawberry [*Fragaria* × *ananassa* (Weston) Duchesne ex Rozier (pro sp.) [*chiloensis* × *virginiana*]] plasticulture production (Webster 2014). *Geranium carolinianum* emerges from the planting hole in the plastic mulch during crop establishment and escapes current control practices. While the exact cause for escape is unknown, it is likely a consequence of poor residual PRE herbicide control during the long growing season. Clopyralid is a registered POST herbicide that can be safely applied over the top of the strawberry crop (Boyd and Dittmar 2015; Figueroa and Doohan 2006; McMurray et al. 1996; Sharpe et al. 2018a). Safety concerns exist for application under higher temperatures that applicators may experience earlier in the production cycle (Sharpe et al. 2018b). These concerns can be alleviated by reducing input volumes with custom precision application technology. Specifically in the case of strawberry production, this would involve targeted application of clopyralid to *G. carolinianum* using machine vision–based applicators. Weed establishment often occupies low percentages of the field and demonstrates a patchy distribution. Site-specific precision application of POST herbicides may reduce inputs by 90% compared with broadcast applications (Zijlstra et al. 2011).

The concept of targeted precision sprayers, particularly autonomous ones, has been well studied, and the general framework has been reviewed elsewhere (Slaughter et al. 2008a).

Those authors note the necessity of several subsystems that are important for development in strawberry production: machine vision, weed control, and displacement-sensing subsystems. The machine vision component is the most complex subsystem, with a variety of sensor options for species discrimination. These include cameras to capture multispectral reflectance (Slaughter et al. 2008b), hyperspectral reflectance (Sharpe 2008; Vrindts et al. 2002), and color images (Tian et al. 1997; Tillett et al. 2008). Use of hyperspectral technology has led to highly accurate (96%) discrimination of weeds in tomatoes (*Lycopersicon esculentum* Mill.) (Zhang et al. 2012a, 2012b). Unfortunately, the cost of hyperspectral cameras is high compared with digital color cameras, which limits implementation (Fennimore et al. 2016).

One of the most promising prospects for species discrimination is digital color image-based deep convolutional neural networks (CNNs) (Schmidhuber 2015). In 2011, a graphics processing unit (GPU)-implemented neural network using a combination of convolutional and max-pooling layers (GPU-MPCNNs) achieved recognition rates greater than 98.73% on classifying traffic signs (Ciresan et al. 2011). Forward-fed GPU-MPCNNs became one of the most competition-winning deep-learning techniques employed thereafter (Schmidhuber 2015). The success of neural networks is due to their architecture, derived from the visual perception of animals (Gu et al. 2017). Increased computational speed through GPU implementation and altering connectivity and operation of layers has allowed for successful development of deep, multilayered CNNs capable of classifying images through pattern and color recognition (Schmidhuber 2015).

For GPU-enhanced CNNs, successful classification and detection is aided by filter-based pattern recognition at the convolutional layers (Gu et al. 2017). Convolutional layers use small, pixel-based filters to both increase the network depth and reduce the dimension of the input images (Szegedy et al. 2014). Pooling layers are an essential element for CNNs, as they reduce the number of connections between convolutional layers and, in turn, reduce computational load (Gu et al. 2017). Finally, the output of a layer may contain an activation function, which creates non-linearity in the network (Teimouri et al. 2018). The most common type is the rectified linear unit, which equates negative outputs to zero, reducing overall computation, and increases identification of sparse targets (Gu et al. 2017).

The recent developments for deep CNNs make them a feasible option for using digital color cameras as a cost-effective sensor for machine vision subsystems of precision applicators. A shallow image-classification CNN (10 layers) was able to discriminate soybean [*Glycine max* (L.) Merr.] plants from three Chinese weeds (*Convolvulus* sp., *Digitaria* sp., and *Cirsium* sp.) with 93% accuracy (Tang et al. 2017). Several weed species (18 in total) were identified with class accuracy between 46% and 78%, with a 70% accuracy of estimating leaf number with the Inception-V3 architecture (Szegedy et al. 2016; Teimouri et al. 2018). Monocot and dicot weeds were separated from the cereal fields with an average precision of 0.76 using the object-detecting CNN SSD512 (Dyrmann et al. 2018; Liu et al. 2016). Uncategorized weeds could be detected in winter wheat (*Triticum aestivum* L.) during occluded conditions using the object detector CNN DetectNet with moderate success (precision = 87%), but many weeds were missed (*recall* = 46%) (Dyrmann et al. 2017; Tao et al. 2016). Which CNNs may be best adapted for use within in situ smart sprayers in Florida strawberry production is currently unknown. Therefore, the objective of the study was to determine the feasibility

of use and accuracy of three CNNs in detecting *G. carolinianum* growing in competition with strawberry.

## Materials and Methods

Images of *G. carolinianum* in competition with strawberry plants were taken using a digital camera (DSC-HX1, Sony® Cyber-shot Digital Still Camera, Sony, Minato, Tokyo, Japan) at a resolution of 1,920 × 1,080 pixels. Training images were taken at two locations: the Gulf Coast Research and Education Center in Balm, FL (27.76°N, 82.22°W), and the Florida Strawberry Growers Association field site in Dover, FL (28.02°N, 82.23°W). Validation images were taken at two commercial strawberry farms (27.93°N, 82.10°W and 27.98°N, 82.10°W) on February 23, 2018. Strawberry plants were transplanted on October 10, 2017, in Balm, October 16, 2017, in Dover, and October 6, 2017, and October 10, 2017, for the validation sites. Images were taken at Balm at 63, 70, and 120 d after transplanting (DATr). Images were taken in Dover at 58, 120, and 130 DATr. Validation images were taken at two commercial strawberry farms at 134 and 136 DATr, respectively. Strawberry plants were at early- to mid-anthesis between 58 and 63 DATr, early fruiting by 70 DATr, and full fruiting by 120 DATr. *Geranium carolinianum* was growing vegetatively between 63 and 120 DATr and had developed reproductive structures after 120 DATr, primarily at validation sites.

The training data set contained a total of 705 images with *G. carolinianum* (positive) and 570 images without *G. carolinianum* (negative). The validation data set contained 88 positive and 109 negative images. Negative images were not restricted to only strawberry plants and contained images of other species growing in competition with strawberry within the bed. These species included black medic (*Medicago lupulina* L.), goosegrass [*Eleusine indica* (L.) Gaertn.], crabgrass species (*Digitaria* spp.), livid amaranth (*Amaranthus blitum* L.), crowfootgrass [*Dactyloctenium aegyptium* (L.) Willd.], and common ragweed (*Ambrosia artemisiifolia* L.). Images were taken centered on a strawberry bed and contained approximately four planting holes. Camera height was 130 cm from the soil surface. Bed height was 30.5 cm. Natural sunlight was taken as the light source, with approximately 10% to 20% cloud cover, and images were taken within 2 h of solar noon.

Images were resized to 1,280 × 720 pixels (720p) using Irfan-view (v. 4.50, Irfan Skijan, Jajce, Bosnia). The 720p image resolution was chosen to develop networks for future in situ video input on developed smart sprayers. This resulted in a ground-sampling distance of 0.2 cm pixel<sup>-1</sup> at bed height and 0.3 cm pixel<sup>-1</sup> at ground height. Image classification images were divided into four equal subimages through cropping, then resorted according to their classification. The resulting data set contained 3,542 negative images and 1,447 positive images.

Data were imported into the NVIDIA Deep Learning GPU Training System (DIGITS) (v. 6.0.0, NVIDIA, Santa Clara, CA) for data set compilation. Three CNNs within the DIGITS library were used for species discrimination. The DIGITS interface was of interest due to its advanced user-friendly interface, graphical layout, and ease of implementation compared with other deep-learning systems. The CNNs can be classified as either image-classification or object-detection based. Image classification uses predetermined groups of images (classes) to train the CNN to select a single classification for each unknown image. Object detection uses input images and label files containing the location



**Figure 1.** Canopy-trained DetectNet-generated bounding boxes generated on validation images for *Geranium carolinianum* growing in competition with the strawberry crop in a plasticulture setting at Plant City, FL, in 2018.

of the target (*G. carolinianum*) in each image. These files were generated by drawing bounding boxes onto positive images using custom software compiled with Lazarus (<https://www.lazarus-ide.org>). Two DetectNet networks for object detection were trained by either labeling whole *G. carolinianum* canopies (Figure 1) or labeling individual *G. carolinianum* leaves (Figure 2).

The CNN training was conducted in DIGITS using the Convolutional Architecture for Fast Feature Embedding (Caffe) (v. 0.15.14, UC Berkeley EECS, Berkeley, CA) framework to train the image-classification and object-detection neural networks (Tao et al. 2016). The three CNNs used for comparison were the visual geometry groups: 16-layer CNN (VGGNet) (Simonyan and Zisserman 2015), the 22-layer CNN deep network GoogLeNet (Szegedy et al. 2014), and the DIGITS network DetectNet (Tao et al. 2016). DetectNet contains five parts: (1) data ingest and augmentation layers, (2) a fully convolutional neural network (GoogLeNet), (3) loss functions to measure training error, (4) a

clustering function during validation for bounding box prediction, and (5) an average precision (*mAP*) metric to determine network performance against the validation set (Jia et al. 2014; Szegedy et al. 2014; Tao et al. 2016). All three CNNs were available through the DIGITS model store and were pretrained with the ImageNet data set (Jia et al. 2009); DetectNet was also pretrained with the KITTI data set (Geiger et al. 2013). Network pretraining facilitates transfer learning for deep CNNs to increase performance (Ge et al. 2015). These CNNs were selected due to their availability and ease of training through the DIGITS interface. Data augmentation was achieved using the augmentation layer in DetectNet and the mirroring option for VGGNet and GoogLeNet. The DetectNet augmentation layer included parameters for random cropping, flipping, rotating, hue adjustments, and desaturation of input images.

The ADADELTA solver type was selected for backpropagation of all CNNs (Zeiler 2012). This was to overcome the fundamental



**Figure 2.** Leaf-trained DetectNet-generated bounding boxes generated on validation images for *Geranium carolinianum* growing in competition with the strawberry crop in a plasticulture setting at Plant City, FL, in 2018.

problem with deep learning, wherein the activation function signal error during cumulative backpropagation grows or shrinks rapidly (Hochreiter et al. 2001; Schmidhuber 2015). During object-detection training, an intersection over union (IoU) method was used to evaluate whether the object detected was a true positive (*tp*), with a threshold of 0.7 (Tao et al. 2016). IoU is the ratio of overlap between the actual and predicted bounding boxes. Training continued, with the training rate progressively reduced, until the loss function output no longer decreased or the desired parameters (*maP*, *precision*, and *recall*) ceased to increase over 100 epochs.

Object-detection and image-classification results for *G. carolinianum* were arranged in a binary-classification confusion matrix under four conditions, a true positive classification (*tp*), a false negative (*fn*), a false positive (*fp*), or a true negative (*tn*) (Table 1). For object-detection networks, validation images were considered on the networks effectiveness at detecting both individual leaves and whole plants. Two scales of effectiveness are considered: leaves and whole plants. This technology is under development for precision-sprayer implementation to target herbicide application across the strawberry plant to ensure proper coverage of *G. carolinianum*, so precision on a plant-level scale was considered.

Three measures of neural network effectiveness were used: *precision*, *recall*, and *Fscore* (Sokolova and Lapalme 2009). *Precision* is a measure of how accurate the neural network was at positive identification and was calculated by:

$$Precision = \frac{tp}{tp+fp} \quad [1]$$

(Hoiem et al. 2012; Sokolova and Lapalme 2009; Tao et al. 2016). *Recall* is a measure of how effectively the neural network identified the target and was calculated by:

$$Recall = \frac{tp}{tp+fn} \quad [2]$$

(Hoiem et al. 2012; Sokolova and Lapalme 2009; Tao et al. 2016). The *Fscore* is the harmonic mean of the *recall* and *precision*, gives an overall impression of the network's positive labels, and was calculated by:

$$Fscore = \frac{2 * Precision * Recall}{Precision + Recall} \quad [3]$$

(Sokolova and Lapalme 2009). For validation, if an image would not converge to produce bounding boxes on visible *G. carolinianum*, all plants present were marked as a miss (0% *precision* and *recall*).

**Table 1.** Confusion matrix demonstrating the designation terminology for network output in detecting *Geranium carolinianum* in competition with strawberry in Hillsborough County, FL, in 2018.

Images actually contain <i>G. carolinianum</i>	Images predicted to contain <i>G. carolinianum</i>	
	Predicted: Yes	Predicted: No
Actual: Yes	True positive	False negative
Actual: No	False positive	True negative

## Results and Discussion

### Image Classification

GoogLeNet performed better than VGGNet for whole-image training (Table 2). GoogLeNet did not increase either *precision* or *recall* with cropped-image training. This contrasted with VGGNet, which demonstrated substantial gains in fit for cropped-image training. Validation for both GoogLeNet and VGGNet developed on whole images were not successful, and produced low values for *recall* and *Fscores*, making them unsuitable for field applications (Table 3). These networks were unable to identify most to all of the *G. carolinianum* present in the images, again making them unsuitable for field applications. Cropped-image training increased the validation results for both GoogLeNet and VGGNet (Table 3). VGGNet made substantial gains in detecting *G. carolinianum*, particularly in rejecting false positives, though the network did still struggle with missing the target (false negatives).

*Recall* was primarily the limiting factor for image classification of *G. carolinianum* while in competition with strawberry. This was likely due to the frame of the input image and the amount of negative space within (Figure 1). The camera field of view contained between three to four planting holes along the bed, and *G. carolinianum* occupied one to two planting holes. The number of *G. carolinianum* leaves that grew through the strawberry canopy and were visible from above were variable, ranging from 4 to more than 30.

Cropping images into four subimages substantially increased the validation *Fscore* of the networks for classifying images containing *G. carolinianum*. Improvements in the *Fscore* were likely due to three factors: (1) the cropped images increased the number of available training images to parameterize the CNN; (2) the overall image contained fewer pixels to analyze; and (3) due to cropping, the overall “negative” area outside the target (*G. carolinianum*) was reduced. The latter two points are likely important, as they increase the percentage of the image that has a positive pattern recognition potential to be amplified by the convolutional filters. The additional layers to facilitate object detection further benefited this by reducing the negative space in which classification occurs.

### Object Detection

Both the leaf-trained and canopy-trained DetectNet CNNs outperformed GoogLeNet and VGGNet CNNs in both training (Table 2) and validation (Table 3). Both the canopy-trained and leaf-trained DetectNet networks fit the data very well. Canopy-based

**Table 2.** Detection-training results using various convolutional neural networks (CNNs) for *Geranium carolinianum* growing in competition with strawberry in Balm, FL, in 2018.

Network	Training specifications	Network type <sup>b</sup>	<i>Precision</i>	<i>Recall</i>	<i>Fscore</i> <sup>a</sup>
DetectNet	Canopy	OD	0.98	1.00	0.99
DetectNet	Leaf	OD	0.99	0.99	0.99
GoogLeNet	Whole image	IC	0.96	0.21	0.35
GoogLeNet	Cropped	IC	0.68	0.22	0.34
VGGNet	Whole image	IC	0.94	0.09	0.16
VGGNet	Cropped	IC	0.73	0.61	0.66

<sup>a</sup>*Fscore* = the harmonic mean of the *precision* and *recall*.

<sup>b</sup>IC, image classification-based CNN; OD, object detection-based CNN.

**Table 3.** Detection validation using various convolutional neural networks (CNNs) for *Geranium carolinianum* growing in competition with strawberry in Balm, FL, in 2018.

Network	Specifications	Detection level	Network type <sup>a</sup>	Precision	Recall	Fscore <sup>b</sup>
DetectNet	Canopy as target	Plants	OD	0.99	0.78	0.87
DetectNet	Canopy as target	Leaves	OD	1.00	0.53	0.69
DetectNet	Leaf as target	Plants	OD	0.88	1.00	0.94
DetectNet	Leaf as target	Leaves	OD	0.99	0.95	0.97
GoogLeNet	Whole	Image	IC	N/C <sup>c</sup>	0	N/C <sup>c</sup>
GoogLeNet	Cropped	Quarter-image	IC	0.81	0.50	0.62
VGGNet	Whole	Image	IC	1.00	0.03	0.05
VGGNet	Cropped	Quarter-image	IC	0.90	0.67	0.77

<sup>a</sup>IC, image classification-based CNN; OD, object detection-based CNN.

<sup>b</sup>Fscore = the harmonic mean of the precision and recall.

<sup>c</sup>N/C, noncalculable value, no positive or false positives were generated by the neural network.

training produced a network that was precise at detecting plants at the cost of recall. As expected, canopy-based training missed more plants and many more leaves than the leaf-based training, but given the overall larger training target, made fewer false identifications. The canopy-trained CNN did produce larger bounding boxes that encompassed the strawberry leaves. Due to how precision was considered, this did not have a negative effect on the calculated accuracy, but depending on the application, this output may be undesirable.

The leaf-trained DetectNet was the best evaluated network for potential precision-sprayer implementation. The network identified all target plants, 142 in total. The smaller target had led to increased identification error (false positives), 63% of which occurred on weed species not prominent in the training data set, and particularly *M. lupulina* leaf clusters (Figure 3). Other false positives included strawberry flowers, immature fruit, and leaf edges. False-positive network desensitization is possible by adding false positive-rich training images for further training. To demonstrate this, an additional 309 images containing *M. lupulina* were taken during the same time periods as the training set and used to supplement the training data. The network was

retrained and successfully desensitized to *M. lupulina* such that false positives no longer occurred.

The strawberry field used to collect validation images contained a *G. carolinianum* morphological variant that differed from plants in the training fields. Plants of this variant contained a more open palmate leaf compared with the more closed palmate leaf predominant in the training fields. Observed morphological variation was unexpected, but the CNN detection performance remained high, making results promising for across-site applications. Missed leaves were primarily at the edges of larger plants that were oriented toward the row middle. A combination of unusual leaf orientation and background noise of the row middle likely led to missed identification; this may be remedied with selective further training. Even so, these larger plants are not necessarily the optimal target for clopyralid application, and sprays should occur on smaller plants, so this limitation may have minimal consequence for implementation. Using a CNN trained to detect individual leaves within the strawberry canopy would increase detection of smaller plants protruding through the canopy.

The leaf-trained DetectNet outperformed the canopy-trained DetectNet in detecting leaves during validation (Table 3). Leaf



**Figure 3.** Leaf-trained DetectNet-generated bounding boxes generated on validation images for *Geranium carolinianum* growing in competition with the strawberry crop at Plant City, FL in 2018. This figure demonstrates false-positive identifications.

detection produced the greatest *Fscore* for any CNN with nearly perfect *precision*. Similar results, namely limited *recall*, were observed for DetectNet trained to detect weeds in a wheat field (Dyrmann et al. 2017). The authors had not differentiated any particular weed species or growth stage, which increased target variability and reduced network *recall*, along with target occlusion by the crop. A similar study using SSD512 resulted in a similarly trained network with limited *recall* (Dyrmann et al. 2018), a consequence of target selection and the necessity of a larger training data set to compensate. For weed science applications, especially scenarios where the available image quantity may be low, target selection is an important consideration. It is advisable to select a target with widespread application and reduced variability, such as a leaf.

Other object-detection networks such as SSD (Liu et al. 2016) and You Only Look Once (Redmon et al. 2016) are designed to identify multiple classes per image. The ability to identify several classes is an important network quality for weed management strategies including scouting or concurrently spraying multiple herbicides. Similar target-dependent *recall* shortcomings have been identified for SSD (Dyrmann et al. 2018). Success may further be limited due to the use of an IoU and rectangular bounding boxes for object-detection CNNs. *Geranium carolinianum* was an excellent test species, because leaves were fairly symmetrical and round. The variability in the target was minimal across varying angles and orientations, which led to a minimum of background noise within the bounding box. Application to Cyperaceae or Poaceae species may be difficult to implement using bounding-box techniques, due to leaf shape and potential orientations. Alternatively, image-segmentation CNNs may be a better approach for such circumstances (Milioto et al. 2017).

In conclusion, results demonstrate a promising avenue for detection of a broadleaf species in a broadleaf crop. A leaf-trained DetectNet CNN achieved a high *Fscore* (0.94) in detecting *G. carolinianum* plants within the strawberry crop. Use of DetectNet as the decision system for a digital camera-based machine vision subsystem appears a viable option for precision control of *G. carolinianum* in Florida strawberry production. Future research is required to test DetectNet and other object-detection networks in situ coordinated with the proper *G. carolinianum* growth stage for ideal control with clopyralid. This would reduce chemical inputs and exposure, thus alleviating concerns regarding temperature-induced damage for earlier applications (Sharpe et al. 2018a, 2018b).

**Acknowledgments.** No conflicts of interest have been declared. This research received no specific grant from any funding agency or the commercial or not-for-profit sectors.

## References

- Boyd NS, Dittmar P (2015) Impact of application time and clopyralid rate on strawberry growth and yield. *Weed Technol* 29:821–826
- Cireşan D, Meier U, Schmidhuber J (2011) A committee of neural networks for traffic sign classification. Pages 1918–1921 in *The International Joint Conference on Neural Networks 2011 Conference Proceedings*. San Jose, CA: IEEE Computational Intelligence Society and the International Neural Network Society
- Dyrmann M, Jørgensen RN, Midtby HS (2017) RoboWeedSupport—detection of weed locations in leaf occluded cereal crops using a fully convolutional neural network. *Adv Anim Biosci* 8:842–847
- Dyrmann M, Skovsen S, Laursen MS, Jørgensen RN (2018) Using a fully convolutional neural network for detecting locations of weeds in images from cereal fields. Pages 1–7 in *The 14th International Conference on Precision Agriculture*. Montreal, QC: International Society of Precision Agriculture
- Fennimore SA, Slaughter DC, Siemens MC, Leon RG, Saber MN (2016) Technology for automation of weed control in specialty crops. *Weed Technol* 30:823–837
- Figuroa RA, Doohan DJ (2006) Selectivity and efficacy of clopyralid on strawberry (*Fragaria X ananassa*). *Weed Technol* 20:101–103
- Ge Z, McCool C, Sanderson C, Corke P (2015) Subset feature learning for fine-grained classification. Pages 46–52 in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshop*. Boston: Institute of Electrical and Electronics Engineers
- Geiger A, Lenz P, Stiller C, Urtasun R (2013) Vision meets robotics: the KITTI dataset. *Int J Rob Res* 32:1231–1237
- Gu J, Wang Z, Kuen J, Ma L, Shahroudy A, Shuai B, Liu T, Wang X, Wang G, Cai J, Chen T (2017) Recent advances in convolutional neural networks. *Pattern Recognit* 77:354–377
- Hochreiter S, Bengio Y, Frasconi P, Schmidhuber J (2001) Gradient flow in recurrent nets: the difficulty of learning long-term dependencies. Pages 1–15 in *Kremer S, Kolen J, eds. A Field Guide to Dynamical Recurrent Neural Networks*. Piscataway, NJ: Wiley-IEEE Press
- Hoiem D, Chodpathumwan Y, Dai Q (2012) Diagnosing error in object detectors. Pages 340–353 in *Fitzgibbon A, Lazebnik S, Perona P, Sato Y, Schmid C, eds. Computer Vision—ECCV 2012*. Berlin: Springer
- Jia D, Wei D, Socher R, Li-Jia L, Kai L, Li F-F (2009) ImageNet: a large-scale hierarchical image database. Pages 248–255 in *2009 IEEE Conference on Computer Vision and Pattern Recognition*. Miami Beach, FL: Institute of Electrical and Electronics Engineers
- Jia Y, Shelhamer E, Donahue J, Karayev S, Long J, Girshick R, Guadarrama S, Darrell T (2014) Caffe: Convolutional Architecture for Fast Feature Embedding. Pages 675–678 in *Proceedings of the 22nd ACM International Conference on Multimedia*. Orlando, FL: Association for Computing Machinery
- Liu W, Anguelov D, Erhan D, Szegedy C, Reed S, Fu C, Berg AC (2016) SSD: single shot multibox detector. <https://arxiv.org/abs/1512.02325>. Accessed: March 12, 2018
- McMurray GL, Monks DW, Leidy RB (1996) Clopyralid use in strawberries (*Fragaria x ananassa* Duch.) grown on plastic mulch. *Weed Sci* 44:350–354
- Milioto A, Lottes P, Stachniss C (2017) Real-time blob-wise sugar beet vs weeds classification for monitoring fields using convolutional neural networks. Pages 41–48 in *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*. Bonn, Germany: International Society for Photogrammetry and Remote Sensing
- Redmon J, Divvala S, Girshick R, Farhadi A (2016) You only look once: unified, real-time object detection. Pages 779–788 in *Proceedings of the 29th IEEE Conference on Computer Vision and Pattern Recognition*. Las Vegas, NV: IEEE Computer Society
- Schmidhuber J (2015) Deep learning in neural networks: an overview. *Neural Networks* 61:85–117
- Sharpe SM (2008) Potential for Hyperspectral Technology in Wild Blueberry (*Vaccinium angustifolium* Ait.) Production. MS thesis. Halifax, NS, Canada: Dalhousie University. 129 p
- Sharpe SM, Boyd NS, Dittmar PJ, MacDonald GE, Darnell RL (2018a) Clopyralid tolerance in strawberry and feasibility of early applications in Florida. *Weed Sci* 66:508–515
- Sharpe SM, Boyd NS, Dittmar PJ, MacDonald GE, Darnell RL (2018b) Effect of temperature on clopyralid safety in strawberry. *Weed Technol* 32:347–351
- Simonyan K, Zisserman A (2015) Very deep convolutional networks for large-scale image recognition. Pages 1–14 in *Proceedings of the 2015 International Conference on Learning Representations*, San Diego, CA: Computational and Biological Learning Society.
- Slaughter DC, Giles DK, Downey D (2008a) Autonomous robotic weed control systems: a review. *Comput Electron Agric* 61:63–78
- Slaughter DC, Giles DK, Fennimore SA, Smith RF (2008b) Multispectral machine vision identification of lettuce and weed seedlings for automated weed control. *Weed Technol* 22:378–384

- Sokolova M, Lapalme G (2009) A systematic analysis of performance measures for classification tasks. *Inf Process Manag* 45:427–437
- Szegedy C, Liu W, Jia Y, Sermanet P, Reed S, Anguelov D, Erhan D, Vanhoucke V, Rabinovich A (2014) Going Deeper with Convolutions. <http://arxiv.org/abs/1409.4842>. Accessed: May 11, 2018
- Szegedy C, Vanhoucke V, Ioffe S, Shlens J, Wojna Z (2016) Rethinking the inception architecture for computer vision. Pages 27–30 *in* Proceedings of the 2016 IEEE Conference on Computer Vision and Pattern Recognition. Las Vegas, NV: Institute of Electrical and Electronics Engineers
- Tang JL, Wang D, Zhang ZG, He LJ, Xin J, Xu Y (2017) Weed identification based on K-means feature learning combined with convolutional neural network. *Comput Electron Agric* 135:63–70
- Tao A, Barker J, Sarathy S (2016) DetectNet: Deep Neural Network for Object Detection in DIGITS. <https://devblogs.nvidia.com/detectnet-deep-neural-network-object-detection-digits>. Accessed: May 11, 2018
- Teimouri N, Dyrmann M, Nielsen P, Mathiassen S, Somerville G, Jørgensen R (2018) Weed growth stage estimator using deep convolutional neural networks. *Sensors* 18:1580
- Tian L, Slaughter DC, Norris RF (1997) Outdoor field machine vision identification of tomato seedlings for automated weed control. *Trans ASAE* 40:1761–1768
- Tillett ND, Hague T, Grundy AC, Dedousis AP (2008) Mechanical within-row weed control for transplanted crops using computer vision. *Biosyst Eng* 99:171–178
- Vrindts E, Baerdemaeker JDE, Ramon H (2002) Weed detection using canopy reflection. *Precis Agric* 3:63–80
- Webster TM (2014) Weed survey—southern states 2014. Vegetable, fruit and nut crop subsection. Page 288 *in* Proceedings of the Southern Weed Science Society 67th Annual Meeting. Birmingham, AL: Southern Weed Science Society
- Zeiler MD (2012) ADADELTA: An Adaptive Learning Rate Method. <http://arxiv.org/abs/1212.5701>. Accessed: May 14, 2018
- Zhang Y, Slaughter DC, Staab ES (2012a) Robust hyperspectral vision-based classification for multi-season weed mapping. *ISPRS J Photogramm Remote Sens* 69:65–73
- Zhang Y, Staab ES, Slaughter DC, Giles DK, Downey D (2012b) Automated weed control in organic row crops using hyperspectral species identification and thermal micro-dosing. *Crop Prot* 41:96–105
- Zijlstra C, Lund I, Justesen AF, Nicolaisen M, Jensen PK, Bianciotto V, Posta K, Balestrini R, Przetakiewicz A, Czembor E, de Zande J van (2011) Combining novel monitoring tools and precision application technologies for integrated high-tech crop protection in the future (a discussion document). *Pest Manag Sci* 67:616–625