Pest Detection on UAV Imagery using a Deep Convolutional Neural Network

***Yacine Bouroubi[1], Pierre Bugnet[2], Thuy Nguyen-Xuan[2], Claire Gosselin[2], Carl Bélec[3], Louis Longchamps[3] and Philippe Vigneault[3]**

[1] University of Sherbrooke, 2500 boul. de l'Université, Sherbrooke, Quebec, Canada, J1K 2R1.

[2] Effigis Geo-Solutions, 4101 Molson Street, Montreal, Quebec, Canada, H1Y 3L1

[3] Agriculture and Agri-Food Canada, 430 boul. Gouin, Saint-Jean-sur-Richelieu, Quebec, Canada, J3B 3E6

* Corresponding Author: yacine.bouroubi@usherbrooke.ca

A paper from the Proceedings of the
14th International Conference on Precision Agriculture
June 24 – June 27, 2018
Montreal, Quebec, Canada

Abstract. *Presently, precision agriculture uses remote sensing for the mapping of crop biophysical parameters with vegetation indices in order to detect problematic areas, and then send a human specialist for a targeted field investigation. The same principle is applied for the use of UAVs in precision agriculture, but with finer spatial resolutions. Vegetation mapping with UAVs requires the mosaicking of several images, which results in significant geometric and radiometric problems. Furthermore, even at such resolutions, it is still not possible to precisely identify the nature of the detected stresses. The concept proposed here aims to use UAVs for precise and automated pest detection and identification with images acquired a few meters above the crop canopy, at millimetric resolution.*

The image processing is based on artificial intelligence (deep learning) computer vision methods. These methods are trained with images collected for different crops and symptoms. The UAV image acquisitions calendar is optimized using a bioclimatic model that evaluates disease risk. The spatial acquisition plan prioritizes areas with persistent moisture, where the probability of pest presence is higher. These areas could be determined using optical or SAR satellite imagery.

This approach was applied to detect diseases in a vineyard (mildew), potato beetles and weeds (in lettuce, carrot and onion fields). All experimental fields were located in Quebec, Canada. Results show that the application of the deep learning technique to crop canopy UAV images can reach a success rate above 90%, which demonstrates the potential of this approach. Thus, the proposed concept is a major innovation in the application of UAVs in agriculture. It will allow the effective control of pests by optimizing pesticide use while reducing the waste of resources and the harmful effects of chemical products.

Keywords.
Pest detection, weed detection, UAV imagery, deep learning

Introduction

Precision agriculture is based on plot management according to intra-field variability by modulating interventions and treatments as needed in order to reconcile the principle of sustainable agriculture with the objective of economic profitability. In order to adapt farming practices to the spatial variability of plots, satellite remote sensing is an excellent way to map soils and crop growth on a scale suitable for examining plot conditions, detecting stresses, or applying inputs in "management zones" or "variable rates" modes. At this scale (resolution of 5 m to 2 m) it is possible to detect stresses, but not to identify them. Indeed, a weak growth detected on a satellite image can be due to a nutrient deficiency, a pathology, insects, problems of compaction or drainage, an extreme soil texture, etc. Products derived from satellite images, although very useful, do not allow to know the exact nature of the detected stresses.

Drone images have already been used to detect phytosanitary problems using traditional approaches such as classification. However, these methods have not been very successful due to the high confusion rates in the detection of biophysical and phytosanitary problems, especially during different phenological stages (Albetis et al., 2017, Malhein et al., 2012). In recent years, the use of deep learning has led to incredible advances in image recognition that can be used to identify objects of all kinds, even if they feature different colours, sizes or shapes (Russakovsky et al. al., 2015). These methods have recently been applied to the fields of agriculture. Although the number of studies conducted thus far on the subject are still small, their results are very promising. They include those on the identification of diseases (Sladojevic et al., 2017, Durmus et al., 2017, Hanson et al., 2017). Deep learning reveals resistance to the conditions of acquisition, but also to the variations that the same category of objects can exhibit (shape, colour, etc.). This performance fills the gaps of the traditional approaches developed in remote sensing.

The concept presented in this paper is to "send drones to take images with sufficient resolution to make accurate identification automatically". The application of this concept is based on the following principles: (1) drone images must be acquired at sufficient resolutions (of the order of one millimeter) to "see" the nature of the problem; (2) the image processing system, based on a deep convolutional neural network CNN, must be powerful enough to "see" the problem and sufficiently well-trained to "recognize" its nature; (3) the system must be optimized to avoid unnecessary operations (acquisition and processing) by evaluating the spatio-temporal distribution of the risk of developing disease and pests by combining a bioclimatic model (CIPRA, Beaudry et al., 2013) and surface moisture maps derived from optical or SAR satellite images.

This automated screening system has been tested for three different applications: detection of mildew and signs of grazing in the vineyard; detection of potato beetles; and detection of weeds in lettuce, carrot and onion plots.

Study Area and Data Acquisition

Applying automated screening for diseases and damages in the vineyards was conducted in the summer of 2017 on the Experimental Farm of Agriculture and Agri-Food Canada (AAFC) in Frelighsburg, Quebec, Canada. The vineyard at the Frelighsburg Experimental Farm consists of two plots of the same variety (Chancellor), the first of which is kept healthy using a combination of preventive treatments, and the second is heavily contaminated by the inoculum of mildew.

The tests presented here were performed on more than 600 images acquired by a DJI Phantom-3 drone on September 22, 2017, at three height levels above the canopy (approximately 2 m, 4 m and 6 m) and at three speeds (about 1 m/s and 2.5 m/s).

The application for the potato beetle was conducted at the AAFC Experimental Farm in Ste-Clotilde, Quebec, Canada. The potato plot (Hilite Russet) was established at the end of May 2017 and consisted of 14 rows, 17 metres in length. Weekly screening was conducted to track potato beetle populations in both time and space. No insecticide treatment was applied for experimental purposes. The images used in this paper were acquired by a DJI Phantom 3 on July 6, 2017.

The images used to test for weed screening are photographs taken in lettuce, carrot, and onion plots at the beginning of the growing season at the time of weed emergence and herbicide application. These images come from the database of images acquired in the region of Napierville, Quebec, and described in the article by Panneton and Brouillard (2009).

Surface moisture estimation using Sentinel-1 and Landsat-8 imagery was performed on relatively larger areas including some commercial vineyards near the Frelighsburg, Quebec area. Figure 1 shows the location of all sites of interest in this study.

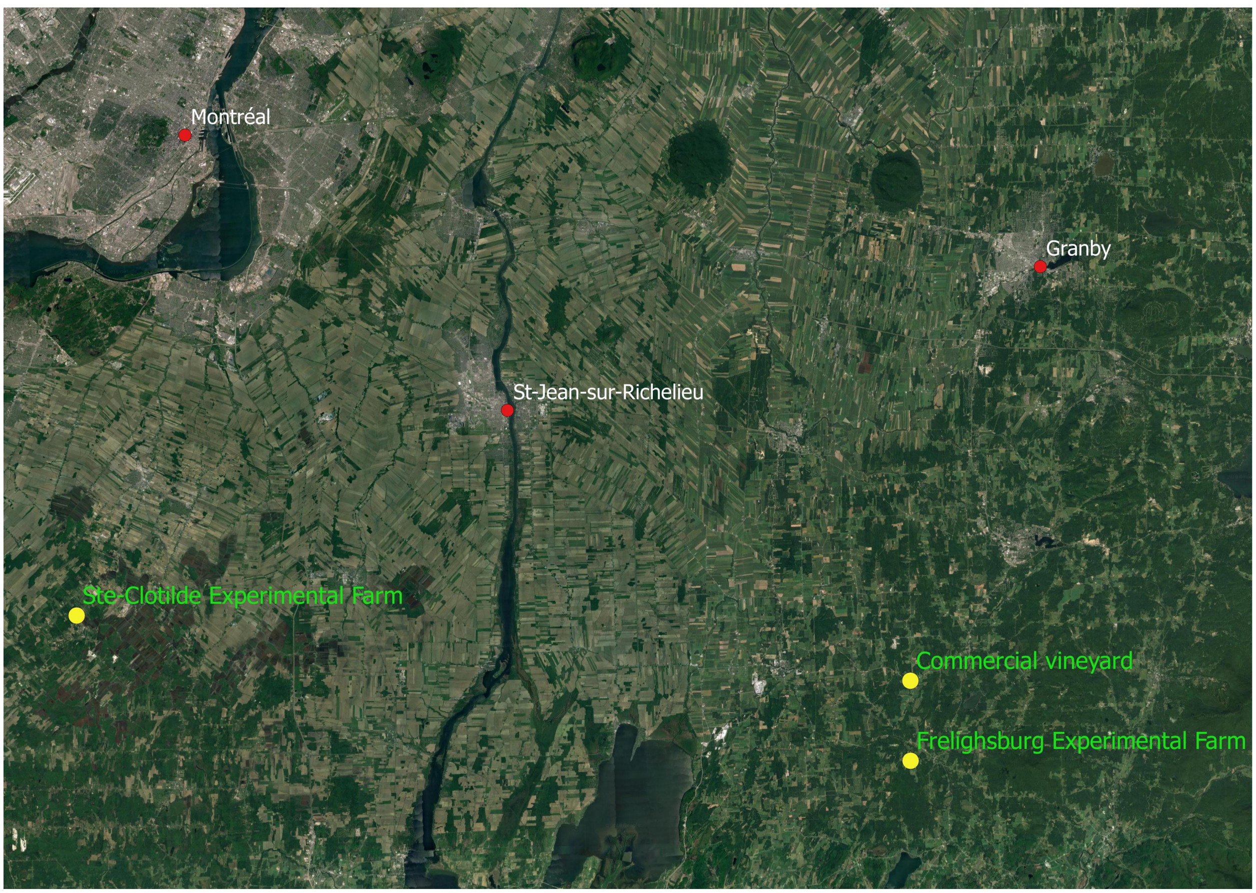

Figure 1. Location of study areas: vineyards at Freligshburg, potato and other vegetable fields at Ste-Clotilde. All study areas are in Quebec, Canada.

Methods

Convolutional Neural Networks for Pest Detection on UAV Images

In recent years, deep learning has revolutionized the field of image recognition (LeCun et al., 2015, Russakovsky et al., 2015). The deep learning algorithms most used in image processing are convolutional neural networks (CNN) consisting of three processing blocks: (i) a first block extracts, through multi-scale convolutions, information at several levels of resolution; (ii) the second block is a deep neural network that "learns", through training on a set of annotated images, to characterize the images and to identify their content; (iii) the third block is a classifier. There are currently several pre-trained CNNs that can be used according to this principle.

In our applications, we have used CNN CaffeNet (Jia et al., 2014) as illustrated in Figure 2. The UAV image is scanned systematically (via convolution) or on a targeted basis (following a segmentation) in order to extract the windows that will be analyzed by the CNN. The "deep signatures" (4096 deep features) are compared with those of the windows containing the searched for symptom classes (diseases, insects, grazing, weeds, etc.) by an SVM classifier. Note that the training windows underwent a number of rotations and translations to take into account the possible variation in the angle of view of the symptoms. The next section presents different image acquisition and analysis options that have been tested to identify the optimal conditions for applying CNN to our problem.

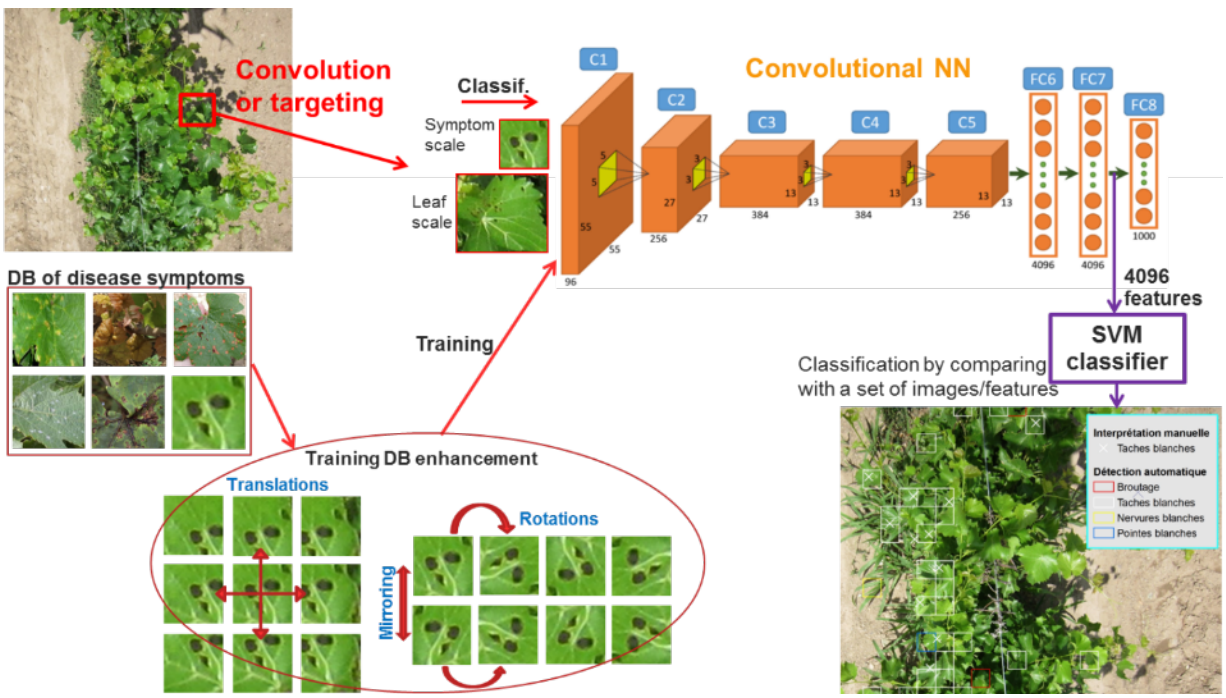

Figure 2. The use of CaffeNet CNN for disease identification on UAV images

Optimizing Imagery Acquisition and Analysis Parameters

UAV Image Resolution

Images acquired by a RGB camera with a resolution of 0.2 mm were degraded (by spatial rescaling) to resolutions of 0.4 mm, 0.6 mm, ..., up to 2.4 mm in order to identify the lowest resolution (and therefore the greater spatial coverage) that still allows for the detection of the looked for symptoms. This exercise (carried out only for vines) allowed evaluating the influence of the resolution of the image on the success of detection of the various symptoms.

Extraction of the Analysis Windows

As mentioned before, we have tried two techniques to extract the windows of analysis used as inputs of the CNN for the detection of the symptoms of disease or devastation. The first technique uses a regular grid defined by convolution, while applying a threshold on the "green/red" index in order to keep only the vegetated windows; and the second technique was based on the segmentation of the image and the application of the CNN to windows centered on the centroids of the objects resulting from the segmentation.

Size of the Analysis Windows

Symptom detection was tested by choosing different sizes of the analysis window, which can be focused on the symptom itself, on an affected leaf, or on a group of leaves. We therefore analyzed the same images with windows of different sizes: 24 x 24 pixels, 50 x 50 pixels, 100 x 100 pixels, 150 x 150 pixels, 200 x 200 pixels, 250 x 250 pixels and 500 x 500 pixels.

Size of the Training Datasets

In order to evaluate the influence of the number of reference signatures used on the results of the detection analyses, we carried out analyses with 100, 500, 1,000 and 2,000 reference signatures per symptom class. The reference signature bank also contains signatures for healthy plants and the background. This "other" class is usually much more common and represents a greater variety of situations than symptom classes. We therefore performed detection tests with different ratios (1:1, 1:10, 1:25 and 1:50) of "symptoms" class size to "other" class size in the training bank.

Spatio-Temporal Optimization of Imagery Acquisition

The systematic detection of diseases and infestations (at regular intervals in space and time) requires significant investments in terms of work and time. The targeting of time periods and screening areas optimizes UAV screening by selecting the right place (the wettest areas) at the right time (the weather conditions most conducive to the development of diseases and pests).

Weather conditions have a considerable influence on the risks of the development and spread of pests. The CIPRA (Beaudry et al., 2013) software developed by AAFC brings together several models for predicting the development of crop pests. These models rely primarily on weather observations and predictions. We compared the forecasts provided by CIPRA for vines mildew to field observations at the Frelighsburg experimental farm.

Mapping of soil moisture and vegetation would allow us to identify areas with a higher risk of developing this infection. The tests presented in this paper concern a wine-production sector near Frelighsburg (commercial plots). We compared SAR Sentinel-1 (VV intensity) images acquired during the 2017 season, on the one hand, and the NDMI index = (NIR - SWIR1)/(NIR + SWIR1) from Landsat-8 on the other hand, with the prevailing weather conditions a few days before the acquisition of the images. This exercise allowed to pattern out the moisture variations in the vineyards during the growing season.

Results

Optimal Parameters for UAV Image Acquisition and Analysis

Regarding the resolution of the images, the results of the same analysis performed on resampled images show that the initial resolution of 0.2 mm has the best success rate. In general, the success rate decreases as the resolution of the image decreases. However, it is relatively stable at resolutions of 0.6 mm to 1 mm. It is therefore recommended to aim for UAV acquisitions with a resolution of about 1 mm at the leaf level.

For the location of the analysis windows, the use of a regular grid (a convolution) associated with a selection of vegetated windows (green/red > threshold) provides results that are often better compared with the use of segmentation. Moreover, the addition of a segmentation phase significantly increases the processing time (3 minutes with segmentation versus 30 seconds without segmentation on an AWS cloud GPU machine). In some cases, when the colour of the object to be detected constitutes crucial information for its detection (e.g., the potato beetle), a segmentation followed by a colour testing is useful to apply before using CNN CaffeNet.

The results of detection obtained with different sizes of analysis windows show that the best success rates for the detection of downy mildew are achieved with window sizes of 100 x 100 to 250 x 250 pixels. The best grazing detection success rates in the vineyard are observed with windows of 50 x 50 to 150 x 150 pixels of 1 mm. Potato beetle detection is optimal for windows of 50 x 50 pixels of 1 mm. Weed detection at the beginning of the growth stage is also optimal for windows of 50 x 50 pixels of 1 mm.

Tests on the number of reference signatures (windows used for CNN training) by symptom class showed that success rates stabilize when there are at least 500 reference signatures for each symptom class. The ratio between the reference signatures for the symptom classes and those for the “other” class that gives the best results is “1:1”, i.e., when the size of the “other” class is roughly equal to the size of the classes of symptoms. Weighting techniques exist to overcome the imbalance between the class size of the objects searched and the size of the classes of the background (Wang et al., 2016).

These optimal parameters were used in the application of CNN CaffeNet for the three applications of this study: the detection of vine mildew, potato beetle and weeds in fields of vegetable crops (lettuce, carrots and onion). Note that the use of a single-GPU machine (Intel Xeon E5-2686 v4 with 61 Go RAM) allowed processing an entire Phantom-3 UAV image (12 megapixels) in about one minute, 115 times faster than a CPU (Intel i7-2600 CPU, 3.40 GHz and 16 Go RAM).

Pest Detection in Vineyards

Figure 3 shows an example of detection of mildew symptoms on one of the Phantom-3 images acquired on September 22, 2017, in which signs of advanced mildew (dried leaves) can be observed. The detection carried out with windows of 100 x 100 pixels shows that the majority of the traces of mildew were correctly detected. We can note an omission and some false alarms. Note that the dry leaves on the ground, which could have been false alarms, were not detected.

Figure 3. Example of mildew detection using deep learning

The detection of mildew on all UAV images associated with the GPS coordinates recorded during the flight allows producing a map that shows the spatial distribution of mildew infection rates (Figure 4). This rate represents the ratio between the number of windows classified as "mildew" and the total number of (vegetated) windows analyzed in the UAV image. The mapping of the detection results thus makes it possible to evaluate the extent and severity of the infection and infestation rates of a plot and to establish a phytosanitary intervention plan.

This will allow a variable dose application during the phytosanitary treatment.

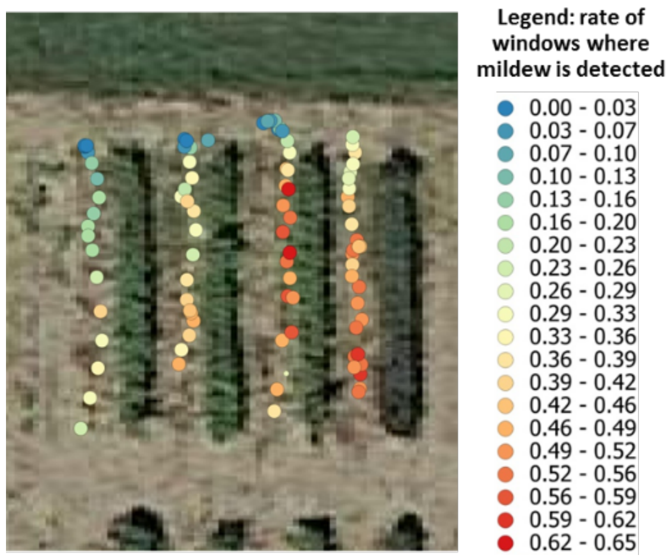

Figure 4. Severity of mildew infection map produced from UAV images analysed by deep learning

Insects Detected in Potatoes

The detection of potato beetles on a Phantom-3 image is illustrated in the example of Figure 5. All insects on the scene have been detected. Only one false alarm is observed in this example.

Figure 5. Example of detection of potato beetles on UAV images using deep learning

Weed Detection

Figure 6 shows examples of weed detection in lettuce, carrot and onion plots. The distinction between weed leaves and crop leaves, which have different shapes, is very satisfactory, according to the visual appreciation of the images of Figure 6.

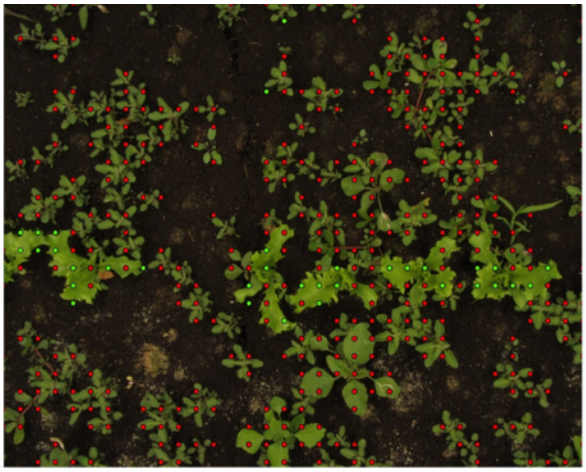

Weed (red dots) detection in a lettuce (green dots) parcel

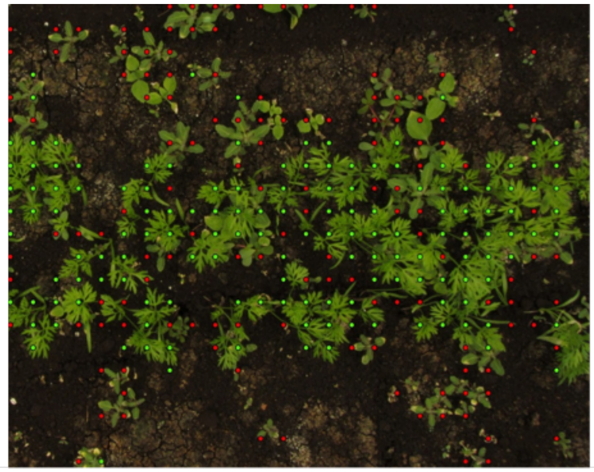

Weed (red dots) detection in a carrot (green dots) parcel

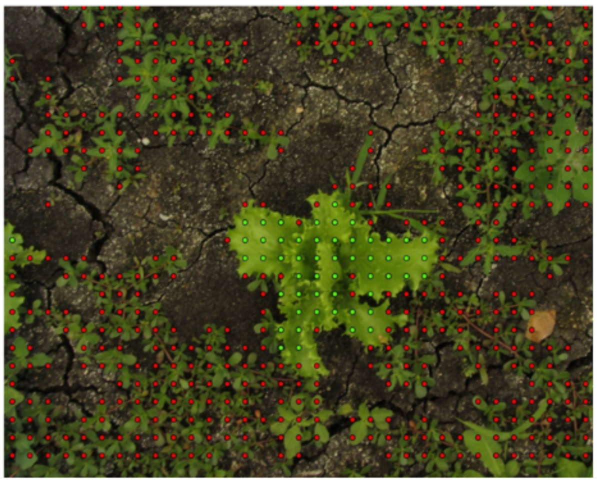

Weed (red dots) detection in a lettuce (green dots) parcel

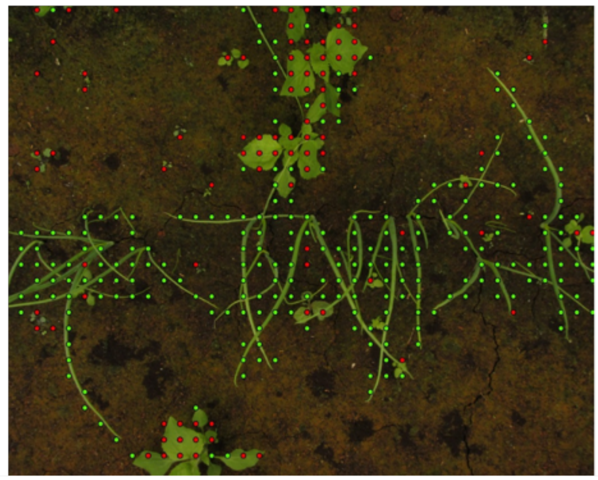

Weed (red dots) detection in a onion (green dots) parcel

Figure 6. Examples of weed detection in lettuce, carrot and onion fields using deep learning

Spatio-Temporal Optimization of UAV Acquisitions

The CIPRA model was run for the Frelighsburg farm. The software automatically retrieves the weather data and displays temperature and precipitation history. The CIPRA model also calculates, from these meteorological data, leaf wetting, which is used to predict the risk of developing diseases, such as mildew and pests. Field observations of mildew and grazing (presence of Japanese beetle) are consistent with the risk dates given by CIPRA. Mildew appeared early in the season, with a dramatic increase as early as mid-July. The signs of grazing became apparent later, around September. This agreement shows the usefulness of CIPRA in achieving an acquisition schedule that takes into account the risk of pests for a given growing season.

In order to assess the ability of Sentinel-1 satellite imagery to estimate surface moisture, we examined whether a relationship existed between VV and VH backscatter observed at different dates and the amount of rain that was measured during the three days before these dates. This analysis was limited to backscatter values from 17 vineyard plots with no vegetation on the ground (between rows). Changes in VV backscattering closely follow moisture level fluctuations. As with Sentinel-1 images, the NDMI moisture index calculated from Landsat-8 imagery is strongly affected by rainfall recorded three days prior to the date of acquisition of the image. Figure 7 shows an example of VV backscattering of Sentinel-1 and Landsat-8 NDMI for acquisition dates before and after precipitation.

This variability in moisture levels between vineyard plots can be an indicator of differences in the risk of developing diseases and insects between these plots. For example, the level of moisture often appears to be higher in Plot D than in Plot A. Thus, one might think that the risk of mildew is higher in Plot D than in Plot A. Plot D should therefore be prioritized in the planning of screening activities. Regarding intra-plot variations, it would be difficult to validate the patterns observed within the plots given the low resolution of the Sentinel-1 and Landsat-8 images and the lack of field measurements of moisture that could be used for validation.

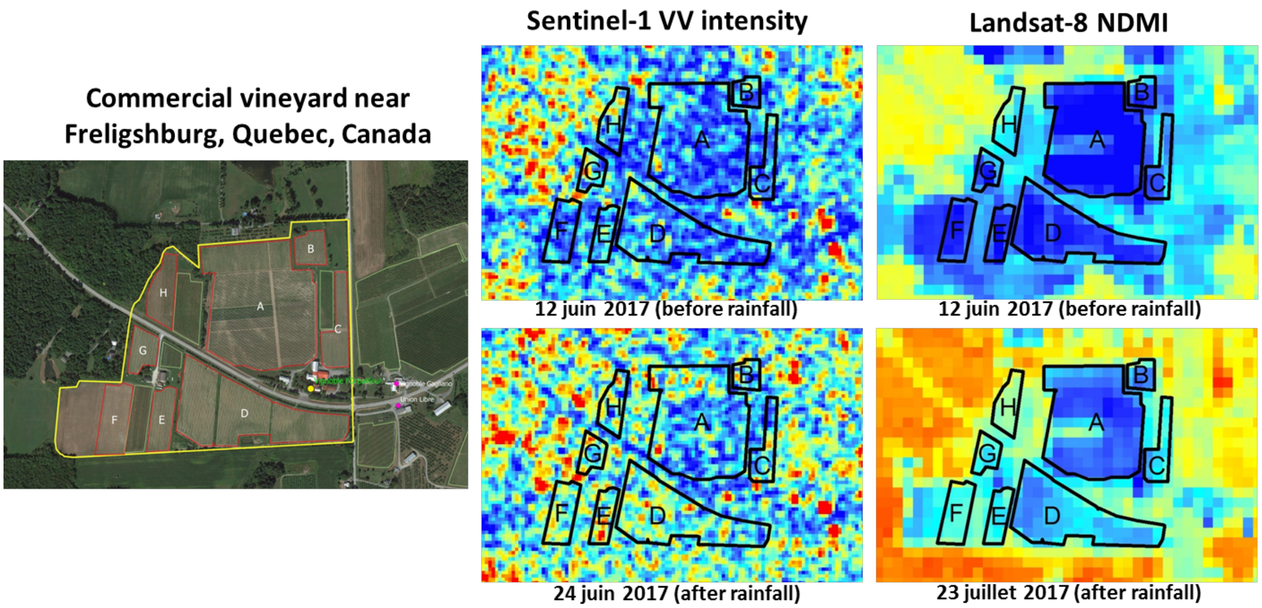

Figure 7. Estimation of the spatial distribution of surface moisture used to assess the risk of pest development

Conclusion and Perspectives

The concept presented in this paper proposes to combine the major developments in compact UAV systems, on the one hand, and in advanced computer vision, on the other. In the last five years, deep learning techniques allowed developing new image-based applications that are very difficult to implement with conventional image processing methods. Thanks to GPU computing power, the processing time is very reasonable for operational use.

The approach developed in this paper will help improving precision farming from zone-based to plant-based management, especially for agricultural input application. We can also assume that efficient and automatic pest detection and identification will help promote the smart use of chemical products thus reducing waste, pollution and human and animal health problems.

This system could be improved or adapted with respect to several components. For example, more recent CNN architectures could be tested, including fine-tuning the CNN through a retraining process, to make the network more sensitive to color. Also, the imagery could be captured with cameras installed on tractors during regular operations and used for a variety of crop monitoring issues.

Acknowledgements

The authors are grateful to the Canadian Space Agency who supported this research.

References

Albetis, J., Duthoit, S., Guttler, F., Jacquin, A., Goulard, M., Poilvé, H., Dedieu, G. (2017). Detection of Flavescence dorée Grapevine Disease Using Unmanned Aerial Vehicle (UAV) Multispectral Imagery. Remote Sensing, 9, 308.

Beaudry, N., G. Bourgeois, D. Choquette, G. Chouinard, D. Plouffe. 2013. CIPRA - Centre informatique de prévision des ravageurs en agriculture: Guide des cultures. Publication d'AAC.

Durmus, H., Günes, E. O., & Kirci, M. (2017). Disease Detection on the Leaves of the Tomato Plants by Using Deep Learning. 6th International Conference on Agro-Geoinformatics. Fairfax.

Hanson, A. M., Joy, A., & Francis, J. (2017). Plant Leaf Disease Detection using Deep Learning and Convolutional Neural Network. International Journal of Engineering Science and Computing, 7, 3, 5324-5326.

Jia, Y., E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell, “Caffe: Convolutional architecture for fast feature embedding,” arXiv:1408.5093, 2014.

LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. Nature, 521, 436-444.

Mahlein, A.-K., Rumpf, T., P. Welke b, H.-W. D., Plümer, L., Steiner, U., & Oerke, E.-C. (2012). Development of spectral indices for detecting and identifying plant diseases. Remote Sensing of Environment, 128, 21-30.

Panneton B., Brouillard M. (2009) Colour representation methods for segmentation of vegetation in photographs. Biosyst Eng 102:365–378

Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Fei-Fei, L. (2015). ImageNet Large Scale Visual Recognition Challenge. International Journal of Computer Vision, 115, 3, 211-252.

Sladojevic, S., Arsenovic, M., Anderla, A., Culibrk, D., & Stefanovic, D. (2016). Deep Neural Network Based Recognition of Plant Disease by Leaf Image Classification. Computational Intelligence and Neuroscience, 2016.

Wang S., W. Liu, J. Wu, L. Cao, Q. Meng, and P.J. Kennedy. 2016. Training Deep Neural Networks on Imbalanced Data Sets. 2016 International Joint Conference on Neural Networks (IJCNN) - IEEE. Vancouver, BC, Canada.