



L3HARRIS™
FAST. FORWARD.

Deep Learning in Agricultural Remote Sensing Applications

Jason Wolfe

L3Harris Geospatial

December 2019

TABLE OF CONTENTS

INTRODUCTION	3
DEEP LEARNING AND MACHINE LEARNING IN REMOTE SENSING	3
APPLICATIONS IN AGRICULTURE	4
Classifying land-use categories and crop types	4
Identifying and Mapping Weeds	6
Identifying Irregular Field Patterns	6
CONCLUSIONS	9
REFERENCES	9

LIST OF FIGURES

1. <i>Diagram of a deep neural network.....</i>	3
2. <i>Hierarchical representation of how deep learning algorithms learn features.....</i>	4
3. <i>SVM classification map of sugarcane development stages.....</i>	5
4. <i>UAV image of sugar beet plants and weeds.....</i>	6
5. <i>Classification image of circular irrigation fields.....</i>	7
6. <i>Google Earth image of smallholder farms in northern Nigeria.....</i>	7

Deep Learning in Agricultural Remote Sensing Applications

INTRODUCTION

Precise and frequent monitoring of agricultural health and productivity is critical for food security and economic sustainability. When remote sensing is used as a tool to monitor agriculture, the analytics must be reliable and accurate. Deep learning technology has provided highly accurate solutions to geospatial problems for many years. Its use in agricultural applications, however, is relatively new and continues to evolve as more research is conducted.

This paper addresses the question, “What types of agricultural remote sensing problems would be best solved using deep learning?” It describes how deep learning is superior to traditional machine learning methods for finding spatial patterns in land-use and agricultural regions, at the expense of significantly more training data. The paper explores common applications where deep learning is used in agriculture. The focus here is on remote observations from airborne and satellite images as opposed to leaf- or fruit-level observations.

DEEP LEARNING AND MACHINE LEARNING IN REMOTE SENSING

Remote sensing imagery is a valuable resource for monitoring agricultural practices throughout the world. In the era of “big data,” Earth observation satellites are being launched at a record pace, with consumers having access to imagery spanning the globe with more frequent revisit rates. For land-use classification and feature extraction, deep learning has received a lot of attention in recent years for its potential in producing highly accurate results. However, it is important to understand the nature of deep learning to determine what types of problems it can solve and where it can be the most helpful.

Deep learning is a subfield of machine learning, in which computer systems learn tasks and identify patterns on their own with only limited training data. What differentiates deep learning from machine learning is its ability to continually improve a prediction on its own without external guidance or intervention. Deep learning algorithms learn patterns by progressing through a series of layers in a neural network in order to draw conclusions (Figure 1).




Figure 1: Diagram of a deep neural network.

The ENVI® Deep Learning module is powered by TensorFlow, which is based on a convolutional neural network (CNN). A CNN uses a “deep” neural network to extract complex features from data in a hierarchical manner. It begins by learning spatial patterns (such as edges) at a low level, then extracts more complex features as the model progresses through deeper layers of data.

Here is a simple example of how machine learning and deep learning differ: A machine-learning algorithm can identify different colors of vehicles in high-resolution imagery based on the pixel values. However, a deep learning algorithm takes it a step further and not only identifies different colors but also different shapes and sizes of vehicles. It does this because the training data provides contextual information, which means the algorithm gets spatial information not only from each pixel, but also from its neighboring pixels. Thus, deep learning is ideal for solving *spatial* problems.

Figure 2 shows an example of how deep learning algorithms are initially presented with labeled images made up of simple pixels (Step 1). Next, they discover simple regularities that are present across many or all images (Step 2). They discover how the regularities are related to form higher-level concepts (Step 3). Finally, the system gains a high-level understanding of the target feature (Step 4).




Figure 2: Hierarchical representation of how deep learning algorithms learn features, using cars as an example.

Deep learning works best with remote sensing images that have a high spatial resolution because they provide more detail and richer spatial features. Deep learning algorithms do not tend to work as well with medium-resolution (10- to 30-meter) images because of the lack of fine details (Ma et al., 2019). Unmanned aerial vehicle (UAV) sensors are ideal for deep learning applications because of their high spatial resolution (Oghaz et al., 2019).

At the same time, too high of a resolution can result in longer processing times. For example, 7.5-centimeter data takes four times longer to process than 15-centimeter data, and it has four times the data volume. Determining the best resolution involves a trade-off between processing speed and having enough contextual information to locate features.

The next section explores agricultural use cases that can benefit from deep learning.

APPLICATIONS IN AGRICULTURE

Although deep learning technology has been around for decades, its use among novice remote sensing users is still relatively new. Ma et al. (2019) completed an extensive survey of deep learning studies in remote sensing and noted that the community has increasingly moved toward deep learning applications since 2014. Kamilaris and Prenafeta-Boldú (2018) completed a similar survey of deep learning applications in agriculture. Of the many documented uses of deep learning in agriculture, a few seem to be of particular interest to consumers of remote sensing images. These applications are described in the sections that follow.

Classifying Land-Use Categories and Crop Types

Land-use and crop-type classification can benefit from machine learning and deep learning methods. Because deep learning algorithms can learn complex spatial patterns in remote sensing imagery—while accounting for changes in solar illumination, rotation, and size—they can potentially yield much higher classification results than traditional machine learning methods. Their success depends on the size and quality of the training dataset (Kamilaris and Prenafeta-Boldú, 2018). The accuracy of deep learning land-use classifiers depends highly on the number of classes as well as the spectral and spatial distinction between the classes (Ma et al., 2018).

Mahdianpari et al. (2018) reported classification accuracies of 93.57% to 96.17% for three CNN-based algorithms, and Ndikumana et al. (2019) reported an F-score of 0.96 for a deep recurrent neural network (RNN). However, both studies used a large amount of training data. This is the key difference between machine learning and deep learning classifiers: Supervised deep learning classifiers require dozens, or even hundreds of training images for best results.

In contrast, machine learning classifiers such as K Nearest Neighbor (KNN), Support Vector Machine (SVM), and Random Forest (RF) usually rely on one or a few input images and only a handful of regions of interest (ROIs) for training data. Thus, they are faster to train than deep learning classifiers, although SVM is known to consume lots of system memory with large images (Kussel et al., 2017). While machine-learning classifiers

often yield accurate classification results, they rely on spectral information alone and ignore spatial and temporal information (Castro et al., 2017; Ndikumana et al., 2018). They cannot effectively identify spatial patterns or work with time-series data.

Figure 3 shows an example of an SVM classification image of different stages of sugarcane development. Using deep learning for this research would require dozens of training images over a growing season.

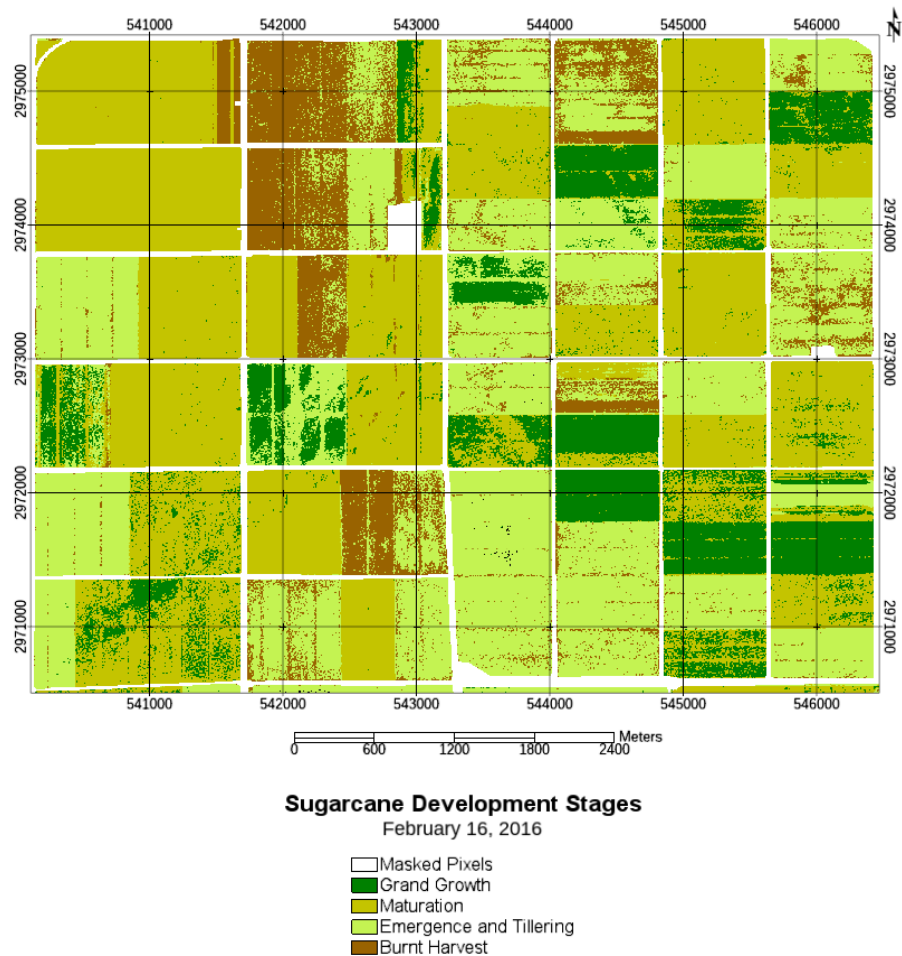


Figure 3. SVM classification map of sugarcane development stages (Wolfe, 2017).

For discriminating vegetation species or crop types, the general consensus is that training data is needed throughout the growing season (Castro et al., 2017; Kussel et al., 2017). Having a consistent and frequent time series of images can provide information about the phenological cycles of plants, which is essential to discriminating among them and monitoring their growth rates. Collecting data from multiple sources such as optical and synthetic aperture radar (SAR) can further improve classification accuracy.

In summary, deep learning has proven to achieve highly accurate land-use classification results in some studies such as Mahdianpari et al. (2018) and Ndikumana et al. (2019). However, those studies required much more training data when compared to traditional supervised classifiers. With crop types in particular, the training images should encompass multiple dates throughout the growing season.

Finally, high-resolution images will often yield the most accurate results with deep learning classifiers. When using medium (10- to 30-meter) resolution training images, the process of identifying and labeling features can be challenging as the features may be difficult to discern. This introduces the risk of providing bad training data, which can result in misclassified pixels and inaccurate predictions. Using high-resolution imagery for training will more effectively reveal spatial features, which are essential to accurate crop classification (Peña-Barragán et al., 2011).

Identifying and Mapping Weeds

The presence of weeds can contribute to decreased crop yields, and regular monitoring is needed to control their growth. Applying the same rate of herbicide to an entire field can create waste and cost farmers more money, in addition to creating environmental pollution (Bah et al., 2018). Combining image analytics from UAV imagery with precision agriculture can help agronomists advise farmers on where to target the herbicides to specific locations in the field. Precision agriculture is a practice that relies on Global Positioning Systems (GPS) and *in-situ* or remote sensors to provide detailed information about crop and soil health so that growers can ensure optimum productivity within fields.

Deep learning provides a practical solution to identifying weeds in fields from UAV imagery. Bah et al. (2018), Huang et al. (2018), and Sa et al. (2018) are examples of studies that followed this approach. Other studies such as Dyrmann et al. (2016) and Olsen et al. (2018) used deep learning to classify weeds at the leaf or ground level. This can be useful for recognizing and preventing mold growth, leaf rot, and insect damage.

Crops and weeds are spectrally similar, and they often look identical in color when viewing true-color or false-color imagery (Figure 4). Deep learning algorithms can detect subtle differences in their shapes and patterns.




Figure 4: UAV image of sugar beet plants and weeds, from Sa et al. (2018). Dataset available from <https://projects.asl.ethz.ch/datasets/doku.php?id=weedmap:remotesensing2018weedmap>.

Identifying Irregular Field Patterns

Since CNNs learn to identify different spatial patterns from imagery, they provide an ideal solution for locating agricultural fields with irregular shapes such as center-pivot irrigation fields, smallholder farms with indistinct borders, and fields containing illicit drug crops.

Circular Fields

Circular fields that are irrigated by center-pivot systems are popular throughout the United States but are especially prevalent in eastern Colorado, Kansas, and Nebraska. Although they incur lower labor costs than ground irrigation methods, they draw upon groundwater in underground aquifers. Researchers have been monitoring irrigated cropland from satellite imagery for decades, in part to predict levels of water consumption over time. While irrigated cropland is easy to identify in satellite imagery containing a near-infrared band, using a vegetation index such as NDVI, distinct shapes such as circular fields are more challenging to identify using automated tools.

Figure 5 shows an example of using ENVI Deep Learning to identify circular fields in a multispectral image.




Figure 5: Classification image of circular irrigation fields from a Sentinel-2 multispectral image of western Kansas, created with ENVI Deep Learning.

Smallholder Farms

Deep learning can help to delineate agricultural fields from surrounding land-use types in satellite images, particularly in subtropical countries whose agricultural practices are less regulated by the government. Having accurate maps of crop types and acreage is an important part of agricultural monitoring systems. They are essential to ensure food security in some regions and countries (Aguilar et al., 2018; Persello et al., 2019). In the United States and some European Union countries, agricultural activities are highly regulated by the government. This includes documentation of crop types and gross production. The U.S. Department of Agriculture (USDA), for example, publishes maps of crop types throughout the U.S.

In parts of Africa, however, much of the spatial information about agricultural fields is incomplete or not available. This is primarily due to the proliferation of smallholder farms, which are defined as having plot sizes of less than 2 hectares. They dominate the agricultural landscape in Africa and contribute about 75% of agricultural production there (Nyambo, Luhanga, and Yonah, 2019).

Mapping smallholder farms is challenging because they have indistinct boundaries, they are often mixed with other land-use types, and their sizes and crop types can widely vary within a given season (Figure 6).



Figure 6: Smallholder farms in northern Nigeria.

Persello et al. (2019) used deep learning segmentation methods to estimate the boundaries of smallholder farms in two study areas in Africa. They noted that field boundaries did not have clearly visible edges in high-resolution satellite imagery and that the boundaries had to be extracted by detecting subtle changes in texture and spectral patterns.

Finally, deep learning can be used to locate fields where illicit crops, such as *Cannabis Sativa*, are being grown and harvested. An example study is Ferreira et al. (2019). Often, these crops are mixed with regular crops in order to disguise them.

CONCLUSIONS

While deep learning is still evolving in remote sensing applications for agriculture, some consistent findings have emerged from recent studies. With regard to land-use and crop classification, accurate results can be achieved by using a large amount of high-resolution training images, preferably spanning multiple dates throughout the growing season. This is because deep learning algorithms consider the temporal element of the data as well as spatial and spectral.

If the agricultural application considers only spectral information from pixels, then machine learning algorithms such as Random Forest (RF) and Support Vector Machine (SVM) will continue to provide a quick and efficient solution for feature extraction or classification. Deep learning will not offer any added benefit. However, when the application involves a spatial component—such as locating fields with irregular patterns or mapping weeds that are spectrally similar to crops—then deep learning undoubtedly provides an accurate solution.

Finally, this paper only discussed the use of optical imagery; however, imagery from multiple types of sensors (such as LiDAR and SAR) can make deep learning algorithms more robust and suitable for a wider range of agricultural applications.

REFERENCES

- Aguilar, R., R. Zurita-Milla, E. Izquierdo-Verdiguier, and R. A. de By. "A Cloud-Based Multi-Temporal Ensemble Classifier to Map Smallholder Farming Systems." *Remote Sensing* 10, No. 729 (2018).
- Bah, M. D., A. Hafiane, and R. Canals. "Deep Learning with Unsupervised Data Labeling for Weed Detection in Line Crops in UAV Images." *Remote Sensing* 10, No. 11 (2018).
- Castro, J. B., R. Q. Feitosa, L. C. La Rosa, P. A. Diaz, and I. Sanches. "A Comparative Analysis of Deep Learning Techniques for Sub-tropical Crop Types Recognition from Multitemporal Optical/SAR Image Sequences." *Proceedings of the 30th SIBGRAPI Conference on Graphics, Patterns, and Images* (2017): 382-389.
- Dyrmann, M., A. K. Mortensen, H. S. Midtby, and R. N. Jørgensen. "Pixel-wise Classification of Weeds and Crop in Images by Using a Fully Convolutional Neural Network." *Proceedings of the CIGR-AgEng Conference* (2016).
- Ferreira, A., S. Felipussi, R. Pires, S. Avila, G. Santos, J. Lambert, J. Huang, and A. Rocha. "Eyes in the Skies: A Data-Driven Fusion Approach to Identifying Drug Crops from Remote Sensing Images." *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* (2019), in press.
- Huang, H., J. Deng, Y. Lan, A. Yan, X. Deng, and L. Zhang. "A Fully Convolutional Network for Weed Mapping of Unmanned Aerial Vehicle (UAV) Imagery." *PLoS ONE* 13, No. 4 (2018): 1-19.
- Kamilaris, A., and F. Prenafeta-Boldú. "A Review of the Use of Convolutional Neural Networks in Agriculture." *The Journal of Agricultural Science* (2018): 1-11.
- Kussul, N., M. Lavreniuk, S. Skakun, and A. Shelestov. "Deep Learning Classification of Land Cover and Crop Types Using Remote Sensing Data." *IEEE Geoscience and Remote Sensing Letters* 14, No. 5 (2017): 778-782.
- Ma, L., Y. Liu, X. Zhang, Y. Ye, G. Yin, and B. Johnson. "Deep Learning in Remote Sensing Applications: A Meta-analysis and Review." *ISPRS Journal of Photogrammetry and Remote Sensing* 152 (2019): 166-177.
- Mahdianpari, M., B. Salehi, M. Rezaee, F. Mohammadimanesh, and Y. Zhang. "Very Deep Convolutional Neural Networks for Complex Land Cover Mapping Using Multispectral Remote Sensing Imagery." *Remote Sensing* 10, No. 7 (2018).

- Ndikumana, E., D. H. T. Minh, N. Baghdadi, D. Courault, and L. Hossard. "Deep Recurrent Neural Network for Agricultural Classification Using Multitemporal SAR Sentinel-1 for Camargue, France." *Remote Sensing* 10, No. 8 (2018).
- Nyambo, D., E. Luhanga, and Z. Yonah. "A Review of Characterization Approaches for Smallholder Farmers: Towards Predictive Farm Typologies." *Hindawi: The Scientific World Journal*, Article ID 6121467 (2019).
- Oghaz, M. M., M. Razaak, H. Kerdegari, V. Argyriou, and P. Remagnino. "Scene and Environment Monitoring Using Aerial Imagery and Deep Learning." *Proceedings of the 15th International Conference on Distributed Computing in Sensor Systems* (2019).
- Olsen, A., et al. "DeepWeeds: A Multiclass Weed Species Image Dataset for Deep Learning." *Scientific Reports* 9, No. 2058 (2019): 1-12.
- Peña-Barragán, J., M. Ngugi, R. Plant, and J. Six. "Object-based Crop Identification Using Multiple Vegetation Indices, Textural Features and Crop Phenology." *Remote Sensing of Environment* 115 (2011).
- Persello, C., V. Tolpekin, J. Bergado, and R. de By. "Delineation of Agricultural Fields in Smallholder Farms from Satellite Images Using Fully Convolutional Networks and Combinatorial Grouping." *Remote Sensing of Environment* 231 (2019): 111253.
- Sa, I., M. Popivić, K. Raghav, Z. Chen, P. Lottes, F. Liebisch, J. Nieto, C. Stachniss, A. Walter, and R. Siegwart. "WeedMap: A Large-Scale Semantic Weed Mapping Framework Using Aerial Multispectral Imaging and Deep Neural Network for Precision Farming." *Remote Sensing* 10, No. 9 (2018).
- Wolfe, J. "Using ENVI for Agricultural Research. Case Study: Optical Remote Sensing of Sugarcane Development." Harris Geospatial Solutions (2017). Available from <https://www.harrisgeospatial.com/Industry-Solutions/Precision-Agriculture/Sugarcane-Whitepaper>.

