An Analysis of Satellite Imagery for Crop Production using Deep Learning Techniques

Thamizhamuthu R[1*], Revathi M[2], Vartika Jaiswal[3] and Parul Kohli[4]

[1, 2, 3, 4] Department of Computing, SRM Institute of Science and Technology, Tamil Nadu, India.
thamizhr@srmist.edu.in

Abstract. A research project aimed at ascertaining how the most advanced technologies are integrated into enhancing agricultural productivity and sustainability is conducted. The study utilized satellite imagery to analyze key factors determining crop production that can be condensed into three key areas, namely: Weed detection, seeding and planting scheme and water body detection. This improvement in the accuracy and efficiency of the identification of areas overgrown and infested by weeds within agricultural fields is done with the use of CNN. High-resolution satellite images are thus being processed for the proper delineation of distinct zones of vegetation such that interventions can be made much more specific through better crop management practices. Another major study that could be conducted is optimizing the planting and sowing strategy. Using machine learning algorithms applied over data from satellites combined with planting patterns and densities, there would maximized crop yields and optimized land use for better seed and plant selection towards more efficient and productive agriculture. It focuses on water body management wherein images obtained from satellites are applied for the control and monitoring purpose of water use. It identifies and classifies maintained and non-maintained water bodies, analyzing their nature of distribution and availability. This helps in finding a sustainable strategy for the management of water, which then forms an important cause for sufficient water supply during the growing season. All these technologies, and several more, uncover a lot of potential in elevating agricultural productivity by providing actionable insights to the farmers and farm managers. Bringing this conclusion to culmination, the study is going to make agriculture more precise, efficient, and even more profitable by using deep learning algorithms in analyzing satellite imagery.

Keywords: Precision Agriculture, Satellite Imagery, Machine Learning, Deep Learning, Weed Detection, Seeding Schemes, Water Body Management.

1 Introduction

Satellite imaging and computational technologies have made rapid advancements, and several options are now available for precision agriculture that provides powerful monitoring, analysis, and optimization tools in crop production activities.

In the first half of the project, a module of crop/weed detection based on CNNs is developed. CNNs are widely used in image recognition problems because they learn spatial hierarchies in images, which becomes a great advantage wherein the system can differentiate between crops and weeds.

© The Author(s) 2025
S. Kannadhasan et al. (eds.), Proceedings of the International Conference on Sustainability Innovation in Computing and Engineering (ICSICE 24), Advances in Computer Science Research 120,
https://doi.org/10.2991/978-94-6463-718-2_12

Module processing high-resolution satellite images enhance accuracy and efficiency in detecting weed infestation within agricultural fields. Precise weed identification would lead to targeted interventions rather than broad-spectrum herbicides, thus furthering the cause of sustainable agriculture.

The other area would be on the seeding and planting module where machine learning algorithms could optimize planting patterns or densities. All that satellite imagery offers provides tremendous data that might under process with such algorithms then reveal insights into the best strategies for planting. It makes use of variables in terms of soil quality and moisture levels and takes into account the crop's historic performance to maximize the seeding process-that is, crops are seeded in the best location with optimal density, which thus enhances crop yield and improves land use efficiency for attaining a better level of productivity without needing more resources.

It is in this direction that the module for the management of the water body goes towards monitoring and managing water resources in agriculture. The component of this module deep learning approach to identify and classify water bodies from satellite images. In this respect, the module helps in developing strategic plans on sustainable water resource management by looking into the aspect of distribution and availability, particularly in areas where water resources are scarce. The system could distinguish between maintained and non-maintained water bodies; hence, control of irrigation and water usage would be much more accurate, and this would enhance the long-term sustainability in agriculture.
The integration of these technologies upscales agricultural productivity with the project. In that case, study aims at making agriculture more specific, effective, and more rewarding by showing actionable insights derived from satellite imagery to farmers and farm managers. Adding in that use of CNNs, machine learning, and deep learning in this context hail the new computational technologies' transformative impact on traditional farming, paving the way to make agriculture sustainable in the future.

2 Proposed Architecture

The architecture for proposed system is designed to integrate various machine learning and deep learning models into a cohesive system (Fig. 1). The architecture consists of several interconnected components and modules, each responsible for specific tasks within the overall workflow.

2.1 Data Collection and Storage

Satellite Imagery Source. High-resolution satellite imagery is obtained from sources such as Sentinel, Landsat, and commercial providers.

Data Repository. A cloud-based storage solution (e.g., AWS S3, Google Cloud Storage) is used to store raw and processed satellite images, along with any additional metadata.

2.2 Data Preprocessing

Noise Reduction. Techniques such as Gaussian blurring or median filtering to reduce noise in the images.

Normalization. Adjusting image pixel values to a standard range to facilitate model training.

Segmentation. Separating different regions of interest (e.g., crop areas, water bodies) using methods like thresholding or edge detection.

2.3 Analysis and Modelling

Weed Detection Module. A CNN model tailored for image classification tasks. Uses frameworks like TensorFlow to train CNN on annotated dataset.

Seeding and Plating Scheme Module. ML Algorithms for optimization through feature extraction, model selection (random forest, neural network etc.) and training.

Water Body Management Module. Water body detection using unsupervised learning (e.g. K-means clustering) and water resource management using supervised learning models (e.g. SVM, decision trees)

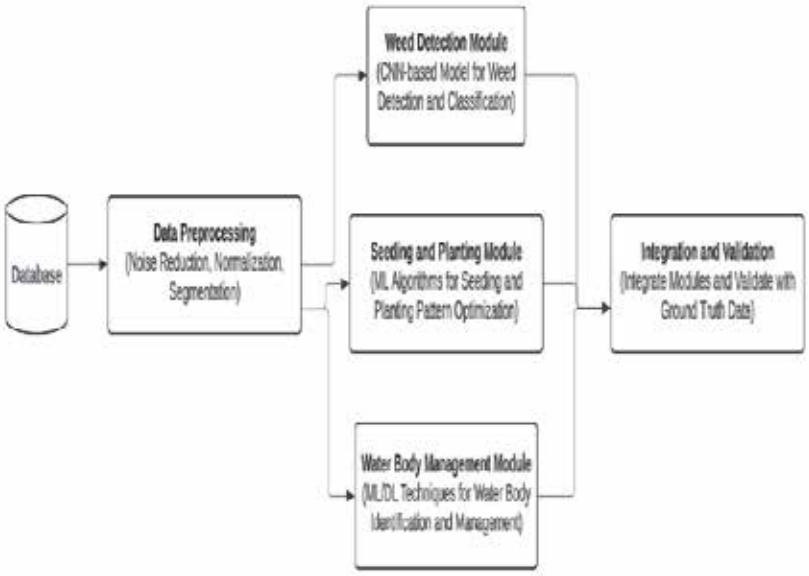

Fig. 1. System Architecture.

2.4 Integration and Validation

Each module (Weed Detection, Seeding and Planting, Water Body Management) is developed as a microservice, allowing independent deployment and scaling. Providing a web-based dashboard for visualizations of detection results, optimization recommendations, and water management insights.

2.5 Weed Detection Module

The Weed Detection Module is an advanced system which helps accurately classify and identify infestations of weeds in agricultural fields with further clarification made by using UAV high-resolution images. Compared to satellite imagery, UAV images benefit greatly by achieving a higher resolution that will enable capturing particular images at certain altitudes and times that are even more precisely identifiable even for the smallest weed patches that go otherwise unnoticed.

A massive dataset which is a very popular source hub in data science was provided for designing this module. It contains annotated images regarding agricultural fields, and it requires annotated images in the supervising learning process to determine and validate the model. The images were labeled with much care where crops and weeds were differentiated and become a great basis for the training of the model.

We make use of InceptionV3 architecture in this task. It is one of the most recent Convolutional Neural Networks (CNNs) through which the model has been significantly noted for a high precision while executing the recognition processes on the images. Pre-trained InceptionV3 is imported; here, base layers of the model become frozen. Thus, the general features on image recognition are saved, but specific characteristics regarding weeds can be learned.

Custom classification layers were put above the frozen InceptionV3 base layer, such that the model can be adapted to weed detection. These layers were designed so as to interpret the features this CNN had learnt and made predictions about the presence of weeds or not. Thus, the custom layers learnt on the Kaggle dataset such that the model learned how to distinguish between crops and weeds. The trained model was then saved and it was later used to create further applications for real-time weed detection tasks. The model monitored performance in the course of training by plotting graphs for accuracy and loss both on training and validation. This furnished critical insights into learning abilities and capabilities concerning generalization, thus allowing us to adjust parameters such that the final model may be correct yet robust.

Image preprocessing was one of the modules of the Weed Detection Module. UAV images are preprocessed to ensure that predictions made homogenous and highlight features that may be important for weed detection before prediction could take place. Some of these processes include resizing the images according to the model's input requirement for InceptionV3, normalizing the pixel values, and data augmentation by increasing its robustness to differences in conditions like lighting, scale, and orientation.

Fig. 2. Original UAV image capturing the presence of weed.

These images are then passed through the trained model for prediction, which is, in turn, to classify each image for the existence of weeds and where they stay. A process known as heatmap was used for better justification of the classification of areas that contain weeds. The heatmap is the method that provides visualization to the regions of the image most likely to be containing weeds a visualization of where the model puts focus and wherein it makes decisions. Fig. 2. Shows the Original UAV image capturing the presence of weed.

For better visualization, the heatmaps produced were overlayed on the original UAV images. Overlay appears to give an excellently clear view of the found locations of weeds within the field, hence aiding in easy validation of model predictions and therefore planning of targeted interventions. The heatmap technique has improved accuracy in weed detection but enabled understanding in the spatial distribution of weeds too.

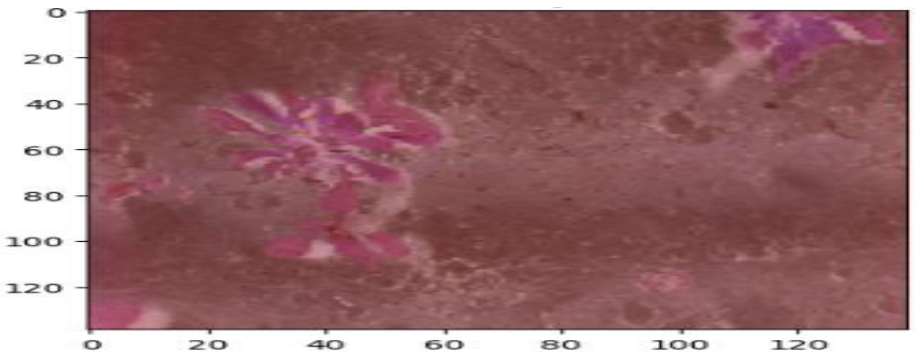

Fig. 3. Overlay image generated using heatmaps to highlight the presence of weed in specific regions.

This concludes by saying that the Weed Detection Module allows for high-accuracy identification of weeds by fusing UAV-acquired high-resolution imagery with deep learning techniques. The tool of precision agriculture which is here being proposed has been built using the InceptionV3 model, custom classification layers, and heatmap-based segmentation. The identification and localization of weeds on UAV images would basically support targeted interventions rather than application of herbicides throughout the whole area in something much more sustainable in the agricultural perspective. Fig. 3. Shows the Overlay image generated using heatmaps to highlight the presence of weed in specific regions.

2.6 Seeding and Planting Module

The seeding and planting module are essentially the heart module in our paper, which is a module to process satellite imagery efficiently and their associated metadata which can help in assisting in land classification and optimizing crop production. It starts by reading data from a CSV file which is a structured repository of information on the satellite images and their labels along with the paths to the images. This is the metadata that forms the starting point at which the imagery is processed, segmented, then analyzed for different types of land cover. Ensuring a balanced and unbiased dataset, the module randomly shuffles the data frame so that the distribution in neither training nor validation can be biased toward one class.

The shuffling step helps in the minimization of overfitting that the model might do very well on some specific areas of the data but miserably fails at unseen parts. This dataset has to be divided into a training set and validation set, using a 90:10 split ratio, so that 90% of the data is being used to train the model whereas 10% is held as validations to ensure that the generalizability of the model is properly tested.

Fig. 4. Satellite image capturing agricultural land.

The following module consists of extracting class information and the RGB values assigned to the different land covers. In this case, the class names differ as far as the set of kinds of lands taken are all encoded differently for easier interpretation and to visualize or distinguish between those classes. For example, it provides color [0, 255, 255] for urban land, agriculture land comes out to be yellow-colored and has [255, 255, 0], while forest land is marked as [0, 255, 0] having green color. These RGB values can map the satellite imagery into meaningful, color-coded segments that are easy to identify. Further, these RGB values are fed into the model. The model treats these as different categories pertaining to their application for segmentation. Fig. 4. Shows the Satellite image capturing agricultural land.

Actually, so vital a part of the module are the class labels to be translated into a machine-friendly format. The one-hot encoding is used for this purpose, which explains a process where class labels have already been transformed to a binary vector. Each of the classes has its unique position marked by a "1," while all others by "0. ". For example, if one pixel represents the agriculture land class, then its one-hot encoded version will be [0, 1, 0, 0, 0, 0, 0], supposing that agriculture is the second class in an order. It is actually true that binary encoding forms a basis during training on the deep learning model. This is because, for instance, the classification for each pixel contained in the image is made easier. Once this model predicts these classes, reverse one-hot encoding is applied to transform those predictions back into the same original RGB values for further interpretation visually. Fig. 5. shows the Ground truth mask. Fig. 6. Shows the Predicted mask.

One more very important function in the module, land cover performs extreme data processing. It not only loads data from a frame of data but also treats class values of RGB and applies some kinds of data augmentation.

Fig. 5. Ground truth mask.

Another key augmentation technique is the diversity artificially increase the training dataset so that it's much more robust. Such techniques include flipping and scaling, rotation, and random cropping simulating conditions that exist in real-world satellite images. That way the model does not rely too much on images having some specific orientation or under certain conditions, and it will generalize better to unseen data. Preprocessing involves the following: normalization refers to the process of getting the pixel intensity values in the same image to the same scale, and reshaping, where all images have the same dimensions, so that the model performs better.

Fig. 6. Predicted mask.

Another key augmentation technique is the diversity artificially increase the training dataset so that it's much more robust. Such techniques include flipping and scaling, rotation, and random cropping simulating conditions that exist in real-world satellite images. That way the model does not rely too much on images having some specific orientation or under certain conditions, and it will generalize better to unseen data. Preprocessing involves the following: normalization refers to the process of getting the pixel intensity values in the same image to the same scale, and reshaping, where all images have the same dimensions, so that the model performs better.

Then, we used DeepLabV3+, which is recently known for a property of strong spatial representations on image data, to apply the actual segmentation task. The convolution in DeepLabV3+ allows the network to detect objects of different sizes, which is an advantage in satellite image analysis. For instance, the patches considered here-farmland or water body-might range from hundreds to thousands of pixels in size. After all augmentations have been applied to the training data, the dataset is learned using the appropriate hyperparameters, such as learning rate, batch size, and epochs. Continuously testing and checking the model as training continues the validation set, the Mean IoU score of the model was reported on test data to be 0.6961-the common metric in the segmentation task measuring an overlap between the segmentation masks learned and the ground truth masks. Dice Loss, which measures the difference of the mask predicted with respect to the ground truth mask, was 0.3123, which further supported that the model was good.

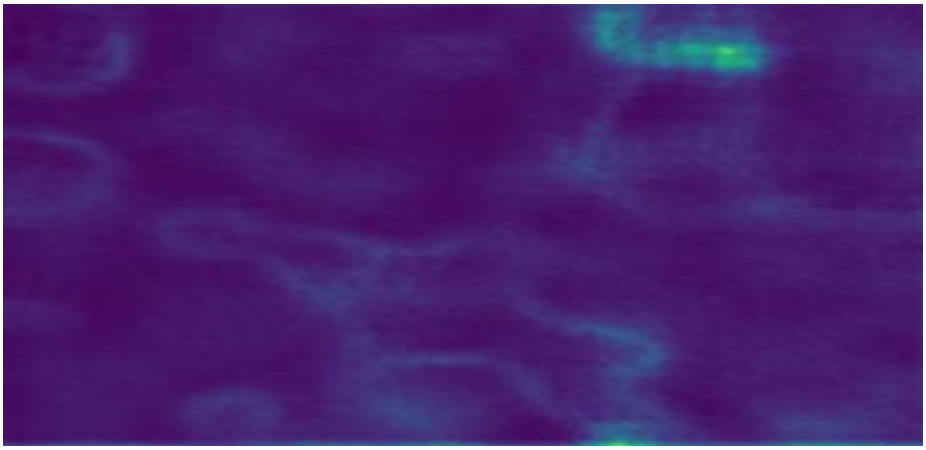

Fig. 7. Predicted urban land heatmap.

While training, besides monitoring IoU and Dice Loss, plots are created to view Dice Loss and IoU for both the training as well as validation set as an assistant in tracing the learning behavior of the model. Fig. 7. Shows the Predicted urban land heatmap.

This allows the possibility of detecting problems such as overfitting, which may imply that the model may overlearn the training data but fail to validate it. On training and validation, final predictions are developed on test data, and the segmented regions corresponding to different land cover types are computed in terms of area.

Such a calculation provides necessary information on land distribution in the studied area and allows agriculturists to make even more effective plans. With this module on seeding and planting, farmers and planners of agriculture would sift through various satellite images in their decisions for good management of the farmlands as well as crops. Such differentiation and quantification of distribution over different types of land cover allow for the optimization of seeding, which can be planned with better crop cycles, resource allocation, and yield prediction. In that sense, in substance, this module carries the crux of our deep learning approach of taking raw satellite data and converting it into actionable insights for sustainable agricultural practices.

2.7 Waterbody Detection Module

Land classification is highly critical in identifying what makes up water bodies from satellite imagery, always essential in agriculture planning. Water bodies are picked based on the unique RGB value of [0, 0, 255], meaning blue color. The same deep learning-based segmentation technique may be applied to differentiate water bodies from other classes of land such as urban land, forest regions, agricultural fields for accurate mapping of water bodies. This way, the DeepLabV3+ model with one-hot encoding and reverse one-hot encoding processes may efficiently differentiate between classes while representing water bodies appropriately and may be a great tool for the detection of water bodies.

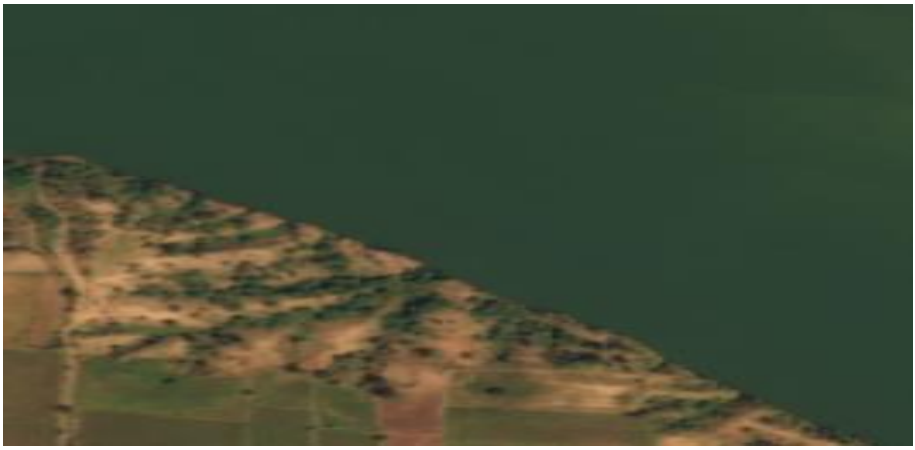

Fig. 8. Satellite image capturing waterbody.

Of course, the water body above the agricultural fields is so important that its detection cannot be overemphasized. In fact, most irrigation takes place from nearby sources; hence, its closeness to lakes, rivers, or reservoirs can significantly influence the plans on crops and yield potential. Detection from satellite images allows farmers and agriculturists to make the correct decision on irrigation systems to optimize water use, waste not, and especially plan more efficiently for the use of this scarce water resource. Fig. 8. Shows the Satellite image capturing waterbody.

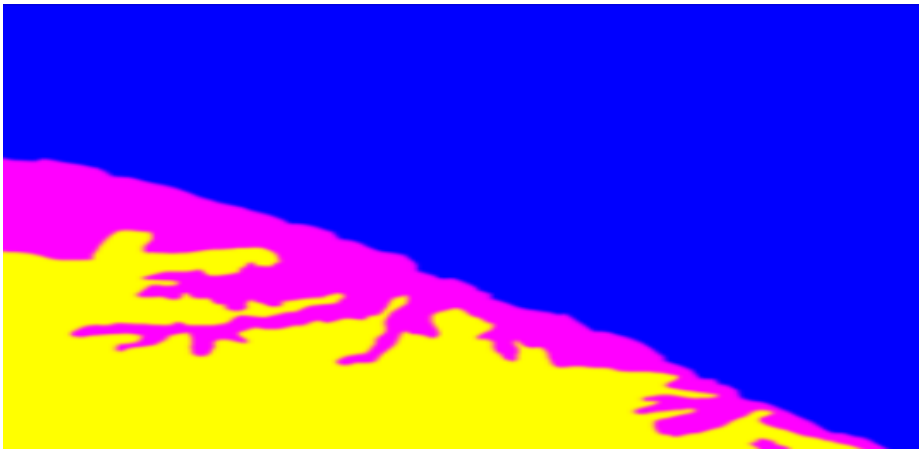

Fig. 9. Ground truth mask.

In addition to general water availability, identification of water bodies and their proximity and size may give an idea about the possibility of flooding within a season of rain. Water body identification also can be helpful for an idea about seasonal fluctuations in water levels and thus predict and plan crop cycles within a region. Water bodies can also be identified to help keep the health of the ecosystem with identification of possible water pollution or decline in water resources that are very important in sustainable agriculture. Fig. 9. Shows the Ground truth mask.

Fig. 10. Predicted urban land heatmap.

That approach ensures this model does not merely segment different kinds of land; it integrates the overall satellite image analysis framework to result in the detection of water bodies. It ends up being such that farmers can incorporate availability of water into their broader agricultural strategy. Fig. 10. shows the Predicted urban land heatmap.

Acknowledgments. We sincerely express our gratitude to our professors, for their invaluable support, insightful guidance, and encouragement at every stage of this research. Their expertise and constructive feedback have been instrumental in shaping this work.We also extend our heartfelt thanks to the organizers of the *ICSICE* for providing us with the opportunity to present our paper. This platform has allowed us to share our findings and gain valuable insights from the academic and research community.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Chenghai Yang, James H. Everitt, Qian Du, “Using High-Resolution Airborne and Satellite Imagery to Assess Crop Growth and Yield Variability for Precision Agriculture”, vol. 101, no. 3, pp. 582-592, March 2013
2. C. Yang, J. H. Everitt, and J. M. Bradford, “Airborne hyperspectral imagery and yield monitor data for estimating grain sorghum yield variability” Trans. ASAE, vol. 47, no. 3, pp. 915–924, 2004.
3. B. Zhang, Y. Wu, F. Zheng, and X. Zhang, “Deep learning for remote sensing image classification: A review,” WIREs Data Mining and Knowledge Discovery, vol. 9, no. 5, 2019.
4. X. Hu, Y. He, F. Xu, and Y. Huang, “A novel deep learning-based method for crop classification using remote sensing data,” IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing, vol. 13, pp. 1452–1463, 2020.
5. L. Ruiz, J. Ballesteros, and D. Escobar, “Object-based crop mapping using Sentinel-2 imagery and convolutional neural networks,” Remote Sensing, vol. 11, no. 6, p. 597, 2019.

6. M. Liu, Y. Liu, J. Shi, and W. Xu, “Crop yield prediction based on multi-temporal satellite imagery using a hybrid deep learning framework,” Remote Sensing, vol. 12, no. 5, p. 707, 2020.
7. G. Maimaitijiang, S. Sagan, B. Sidike, et al., “Soybean yield prediction from UAV using multimodal data fusion and deep learning,” Remote Sensing of Environment, vol. 237, p. 111599, 2020.
8. Z. Zhang, P. Xie, and D. Shi, “High-resolution mapping of cropland in China using multi-temporal Landsat-8 imagery and deep learning techniques,” ISPRS Journal of Photogrammetry and Remote Sensing, vol. 164, pp. 112–124, 2020.
9. N. Kamilaris and F. X. Prenafeta-Boldú, “Deep learning in agriculture: A survey,” Computers and Electronics in Agriculture, vol. 147, pp. 70–90, 2018.
10. Kussul, M. Lavreniuk, S. Skakun, et al., “Deep learning classification of land cover and crop types using remote sensing data,” IEEE Geoscience and Remote Sensing Letters, vol. 14, no. 5, pp. 778–782, 2017.
11. Z. Zhong, L. Li, J. Cui, Z. Jiang, and Y. Wang, “A large-scale land cover classification dataset based on georeferenced images,” IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing, vol. 14, pp. 1074-1086, 2021.
12. X. You, G. Tang, J. Ding, et al., “Crop identification and acreage estimation using UAV and deep learning,” Computers and Electronics in Agriculture, vol. 178, p. 105770, 2020.
13. J. Gutiérrez, J. V. Hinojosa, R. Rodríguez, and G. Quiroz, “Crop water stress detection using remote sensing and machine learning,” Remote Sensing, vol. 12, no. 3, p. 754, 2020.
14. H. Dong, L. Wang, Y. Sun, and X. Sun, “U-Net based segmentation and deep learning classification of agricultural land cover using high-resolution remote sensing images,” Sensors, vol. 21, no. 5, p. 1451, 2021.
15. R. Diop, J. Carreiras, and M. Shimabukuro, “Mapping deforestation in the Brazilian Amazon using deep learning and Sentinel-2 imagery,” Remote Sensing, vol. 12, no. 14, p. 2271, 2020.
16. M. Du, L. Sun, and F. Zhang, “Automatic crop classification using deep learning methods with high-resolution satellite images,” International Journal of Remote Sensing, vol. 41, no. 15, pp. 5898–5914, 2020.
17. J. Shao, X. Fu, and Z. He, “Crop classification using multiscale deep learning and multitemporal Sentinel-2 imagery,” Remote Sensing, vol. 13, no. 4, p. 690, 2021.
18. G. Campanile, P. Iacobellis, and F. Mastrorilli, “Multispectral satellite imagery for real-time monitoring of crop nitrogen status using deep learning algorithms,” Agricultural Water Management, vol. 227, p. 105798, 2020.
19. D. You, H. Lei, and Y. Wang, “Mapping agricultural drought using deep learning and Landsat-8 imagery,” Computers and Electronics in Agriculture, vol. 171, p. 105298, 2020.
20. X. Ma, W. Guo, X. Wang, and Y. Yang, “A crop water deficit estimation method based on deep learning using remote sensing data,” Agricultural Water Management, vol. 240, p. 106379, 2020.
21. L. Ouyang, Z. Zhu, Y. Zhang, and X. Lu, “Using deep learning and high-resolution satellite imagery to assess the impact of agricultural land use change on ecosystem services,” Ecological Indicators, vol. 117, p. 106594, 2020.
22. H. Xu, Y. Li, J. Wu, and F. Shi, “Crop classification with UAV images using deep learning and spectral-temporal data fusion,” Computers and Electronics in Agriculture, vol. 171, p. 105318, 2020.

Open Access This chapter is licensed under the terms of the Creative Commons Attribution-NonCommercial 4.0 International License (http://creativecommons.org/licenses/by-nc/4.0/), which permits any noncommercial use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.