A deep learning based approach for monitoring sustainable farming practices at a parcel level

Arjun Verma[1] , **Vikram Sarbajna**[1]
[1]Credible, Mumbai, India
{arjun, vikram}@credible-india.com

Abstract

Monitoring the effectiveness of policy interventions that promote sustainable farming practices has always been a costly affair. It requires an extensive ground presence which is not always available or reliable. In this paper we present our work so far in the application of deep learning techniques to automate the identification of individual parcels (farms). Our study area is located in the central state of Madhya Pradesh in India, where the average landholding size is around 0.6 hectares per farmer. We created a methodology that uses CNN models for segmentation and Canny Edge detector for generating contours. Our future work concentrates on improving the quality of the reference data and applying additional post-processing methods. Overall, we demonstrate how deep learning could be used for providing specific agronomic advice to individual farmers across large areas and the monitoring thereof, something which is essential in mitigating the effects of climate change.

1 Introduction

Considering the majority of the global food production (estimated to be 70 percent) is done by small holder farmers, even a small change in farming practices can have a significant impact on sustainability parameters. There is a definite surge in the number of initiatives both by local and global organizations to advice and improve key aspects such as soil health and efficient water-usage. Measuring the results of these initiatives at an individual farmer (or parcel) level is a cumbersome process and involves ground truthing by human operators. This is where insights from remote sensing can provide much-needed scale in terms of time and cost. Digital cadastral maps provide an essential piece of infrastructure that facilitate remote sensing-based monitoring at a parcel-level. However, formal cadastral maps cover only one third of all the parcels in the world [Nyandwi *et al.*, 2019]. The uncovered areas are typically found in low and middle-income countries which also happen to have the highest proportion of small holder farmers. According to [Ministry of Agriculture GoI, 2019] the small and marginal land holdings in India (less than 2.00 hectares in area) constituted 86.21% of the total agricultural land holdings in 2015–16. In this paper we present the progress we have made in applying deep learning techniques for parcel-identification in an Indian context. These techniques can be used to automate the process of creating and maintaining cadastral maps, which will ultimately enable the monitoring of sustainable farming practices at a parcel level.

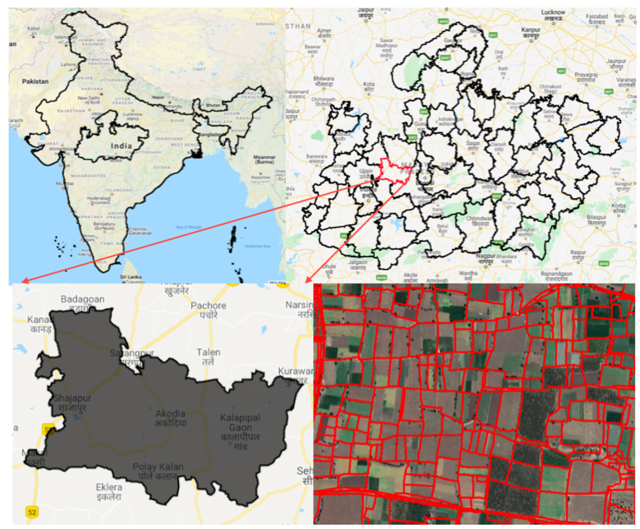

Figure 1: Study Area of Interest - Shajapur, Madhya Pradesh, India

2 Related Work

Image processing approaches that combine edge detection and morphological operations have been the norm for identifying parcels, [Usman and Beiji, 2012] provides such an approach. In recent years however, there has been a steady rise in deep learning based research and applications in remote sensing. [Zhang *et al.*, 2016] provides an overview of deep learning approaches and a general framework for it. The approaches can be grouped into two general categories [Zhang *et al.*, 2016], Pixel based approach (PBA) and Object based approaches (OBA). Our approach would belong to PBA along with the use of Convolution Neural Network (CNN). The majority of the work so far has been concentrated in areas which

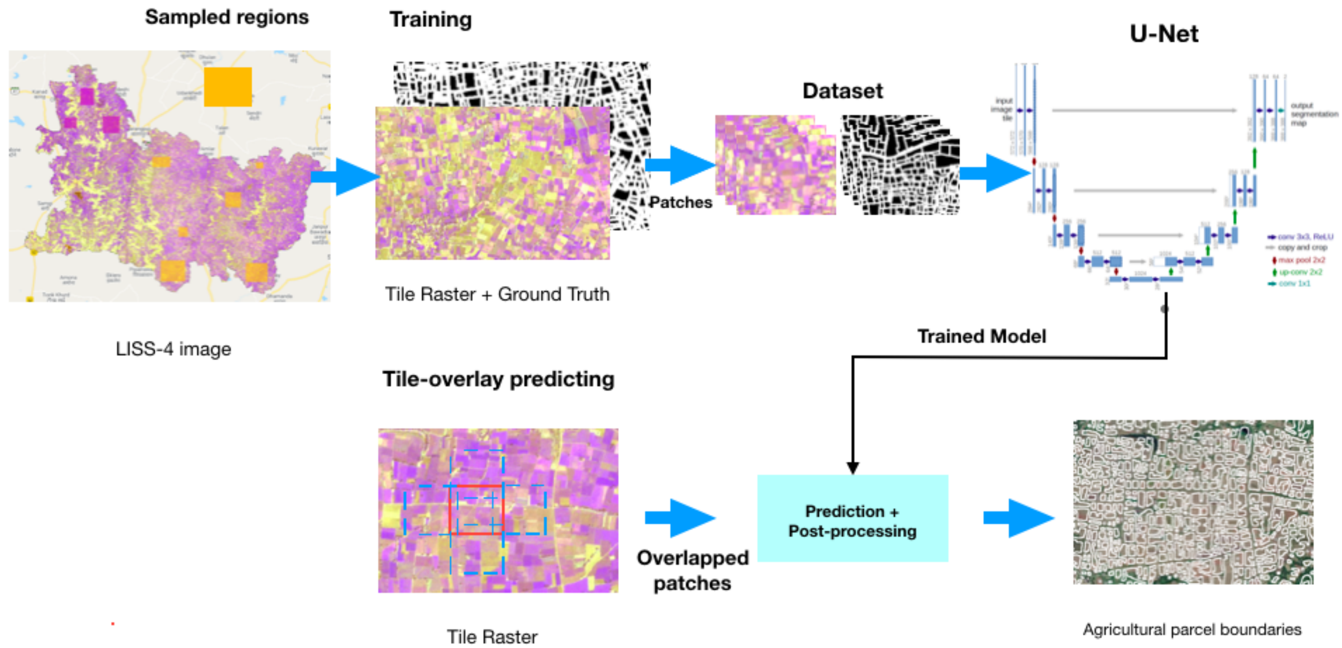

Figure 2: Workflow of the methodology

have predominantly large parcels and using very high resolution (VHR) satellite imagery [García-Pedrero *et al.*, 2017] [Garcia-Pedrero *et al.*, 2019] [Rieke, 2017] [Xia *et al.*, 2019] [Masoud *et al.*, 2020] [Persello *et al.*, 2019]. Usage of VHR imagery becomes a limitation when applied to a country like India due to its associated costs while scaling the approach. Besides satellite imagery costs, the parcels tend to be non-uniform geometric shapes (as opposed to uniformly shaped parcels in high income countries). This paper differs from previous work in its usage of relatively lower resolution imagery and smaller parcels.

3 Study area and Data description

Our study is set in the district of Shajapur in the central state of Madhya Pradesh (MP) in India. MP is one of the major agricultural (producer) states in India. Its major crops include Soybean, Wheat and Gram. The district of Shajapur was selected due to the availability of the cadastral map for this research and the large number of small parcels located within it. The cadastral map was available in an ESRI shapefile format. There are around 0.55 million parcels in the district with an average size of 0.6 hectares. The study area is shown in Figure 1. Satellite images from the LISS-4 camera sensor of ISRO's ResourceSat-2 program were used for this study. It has a spatial resolution of 5.8 m with three spectral bands Red, Green, and Near-Infrared. The spatial coverage of the data matches the entire district of Shajapur.

4 Methodology

The methodology consists of two main parts.

1. Segmentation
2. Prediction and Post-Processing

The first part involves creating a training data set and building a segmentation model. The LISS-4 image was clipped into patches of 256x256 pixels. The reference data was created by transforming the cadastral map into a binary mask. Each training data instance is thus represented by a pair of LISS-4 patch and the corresponding binary mask. The model training is further detailed in Section 4.1. In the second part, we generate the parcels using the trained segmentation model and apply post-processing steps. The predictions were generated in the form of a segmentation map. The output is a continuous [0,1] value in each pixel. The post-processing involves finding the edges and vectorizing the edge map to get individual parcels. The overall workflow of the methodology is shown in Figure 2.

4.1 Segmentation

Our proposed model is inspired from U-Net [Ronneberger *et al.*, 2015]. U-Net is an FCNN model originally developed for medical image segmentation and has an encoder-decoder based architecture. Few adaptions were made in the network: (1) The encoder part was replaced with resnet-34 [He *et al.*, 2015] architecture by extracting the same number of layers from the resnet-34 as the number of encoder layers in the model. (2) The encoder weights were initialized using the pretrained weights from ImageNet dataset [Deng *et al.*, 2009]. The network was trained in following settings: (i) Adam optimizer [Kingma and Ba, 2014] with a learning rate of 1e-4. (ii) Binary cross entropy as an objective function.

(iii) No. of epochs = 12 and batch size = 16.

Figure 3: L To R: Predicted Mask, Extracted parcels

Figure 4: Manually annotated parcels

4.2 Prediction and Post-Processing

While U-Net learns to identify strictly parcels, it does not necessarily generate closed boundaries, which is required to extract individual parcels. In order to obtain the agricultural parcels (in vector data), the following post-processing strategy is followed. The predicted mask shown in Figure 3 was firstly imported back (raster data) in Google Earth Engine (GEE). A Gaussian blur filter followed by the Canny Edge Detector [Canny, 1986] was applied to generate the contours. These contours represent individual parcels. In the vectorization process, we set the same resolution as LISS-4. Finally, we used GEE built-in reduceToVectors [Gorelick *et al.*, 2017] functionality which creates polygons at the boundary of homogeneous groups of connected pixels as shown in Figure 3.

5 Accuracy Assessment

The accuracy assessment was done at a parcel level. Manually annotated parcels were used as reference data for testing as these were more accurate than the cadastral map data. 500 digitised parcels were used as the reference data as shown in Figure 4. The digitised parcels were drawn on Google Earth Engine (GEE) at a 1m resolution. Dice coefficient was used as the metric for evaluation. It is a metric to quantify the percentage overlap between the target mask and our prediction output. It has been widely used to assess image segmentation tasks [Zou *et al.*, 2004] [Zhang *et al.*, 2019]. The Dice coefficient is defined as

$$Dice = \frac{2TP}{2TP + FP + FN} \tag{1}$$

where TP = True Positive, FP = False Positive and FN = False Negative. The reference data and the final output was rasterized back from its original vector format to calculate the dice coefficient.

	Dice Score
Plot-1	0.754
Plot-2	0.313
Plot-3	0.311
Plot-4	0.398
Plot-5	0.413
Plot-6	0.43
Plot-7	0.612
Plot-8	0.33
Plot-9	0.21
Plot-10	0.22

Table 1: Sample Plot level results in terms of Dice score

Mean	Median	SD
0.37	0.377	0.132

Table 2: Aggregated Dice Score of all the parcels

6 Results

Table 1 and 2 shows the dice coefficient scores (predicted parcels compared to reference data) for a random set of ten parcels and aggregated statistics for all the predictions respectively. The results show that 68% of the predictions had a dice score within the range of 0.25-0.50. While 16% fall in 0.5-0.75 range. Figure 5 shows output of predicted parcels overlaid on top of the reference parcels. Figure 6 provides an overview of the predictions for a set of parcels. A certain level of under-segmentation can be observed from the results. The predicted parcels are seen to contain multiple reference parcels within them. We elaborate on these points in the discussion section.

7 Discussion

The under-segmentation of the predicted parcels was expected and can be attributed to factors mentioned earlier such as non-uniform geometries and small size of the parcels. The predictions exhibited dangling features in the boundaries, similar results were observed during the study by [Nyandwi *et al.*, 2019]. A major limitation was the precision of the cadastral map which was used to create the segmentation model

Figure 5: Actual parcels(in red) overlayed with the predicted parcels(in white).

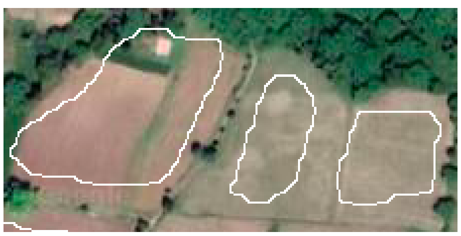

Figure 6: Sample Prediction

during the first part of our methodology. Incorrect or obsolete demarcation of the boundaries led to the inferior quality of the training data. The low dice scores were thus a result of the mismatch between the quality of the training data and that of the manually annotated reference data used for testing. Given this background, we feel that the number of parcels with a dice score above 0.75 can be considerably increased by using better quality training data and applying additional post-processing methods.

8 Future work

Our future work will be concentrated in two areas: Improving the training set (quality and quantity) and exploring additional post-processing methods. We shall replace the existing cadastral map as reference data by manually annotated parcels. This will improve the quality aspect of the training set. There is no benchmark for what constitutes an ideal number of parcels for the training set [Zhang *et al.*, 2016]. We aim to have atleast 5000 parcels.

As far as post-processing step is concerned we shall be exploring atleast four new methods. The first method is based on gPb [Arbeláez *et al.*, 2011] contour detector which combines colour and texture information of an image as opposed to our current method which only uses the gradient of an image. The second method utilizes Structured edge detector(SE) [Dollár and Zitnick, 2013] which represents a supervised learning approach for edge detection. Finally, we will be experimenting with active contours [Kass *et al.*, 1988]. Active contour models(or “snakes” algorithm) are widely used for systematically refining object contours. Active contours have been earlier applied to (radar) satellite images [Horritt *et al.*, 2001] to delineate flood damage extent.

9 Conclusion

Our ongoing work demonstrates how remote sensing can be used to automate the identification of individual parcels. It has the potential to provide a cost-effective way of creating and maintaining cadastral maps for large areas such as entire districts or states. It provides policy makers and change-agents with a powerful tool to monitor and assess the effectiveness of sustainable farming practices without being dependent on ground presence. It also opens up the possibility of providing agronomic advice to individual farmers based on insights (soil moisture, vegetation health, pest or disease attack). Overall, it shows deep learning can be used to speed up much needed (digital) infrastructural projects for largely agrarian low to middle-income countries.

References

[Arbeláez *et al.*, 2011] P. Arbeláez, M. Maire, C. Fowlkes, and J. Malik. Contour detection and hierarchical image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(5):898–916, 2011.

[Canny, 1986] John Canny. A Computational Approach to Edge Detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-8(6):679–698, November 1986.

[Deng *et al.*, 2009] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR09*, 2009.

[Dollár and Zitnick, 2013] Piotr Dollár and C. Lawrence Zitnick. Structured forests for fast edge detection. In *ICCV*, 2013.

[Garcia-Pedrero *et al.*, 2019] Angel Garcia-Pedrero, Mario Lillo, Dionisio Rodríguez-Esparragón, and Consuelo Gonzalo-Martin. Deep learning for automatic outlining agricultural parcels: Exploiting the land parcel identification system. *IEEE Access*, 7:1–1, 10 2019.

[García-Pedrero *et al.*, 2017] A. García-Pedrero, Consuelo Gonzalo-Martin, and Mario Lillo. A machine learning approach for agricultural parcel delineation through agglomerative segmentation. *International Journal of Remote Sensing*, 38:1–11, 01 2017.

[Gorelick *et al.*, 2017] Noel Gorelick, Matt Hancher, Mike Dixon, Simon Ilyushchenko, David Thau, and Rebecca Moore. Google earth engine: Planetary-scale geospatial analysis for everyone. *Remote Sensing of Environment*, 2017.

[He *et al.*, 2015] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition, 2015.

[Horritt *et al.*, 2001] M. S. Horritt, D. C. Mason, and A. J. Luckman. Flood boundary delineation from Synthetic Aperture Radar imagery using a statistical active contour model. *International Journal of Remote Sensing*, 22(13):2489–2507, January 2001.

[Kass *et al.*, 1988] Michael Kass, Andrew Witkin, and Demetri Terzopoulos. Snakes: Active contour models. *International Journal of Computer Vision*, 1(4):321–331, January 1988.

[Kingma and Ba, 2014] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2014.

[Masoud *et al.*, 2020] Khairiya Mudrik Masoud, Claudio Persello, and Valentyn A. Tolpekin. Delineation of Agricultural Field Boundaries from Sentinel-2 Images Using a Novel Super-Resolution Contour Detector Based on Fully Convolutional Networks. *Remote Sensing*, 12(1):59, January 2020.

[Ministry of Agriculture GoI, 2019] Govenment of India Ministry of Agriculture GoI. Agriculture census. Technical report, 2019.

[Nyandwi *et al.*, 2019] Emmanuel Nyandwi, Mila Koeva, Divyani Kohli, and Rohan Bennett. Comparing Human Versus Machine-Driven Cadastral Boundary Feature Extraction. *Remote Sensing*, 11(14):1662, July 2019.

[Persello *et al.*, 2019] Claudio Persello, Valentyn Tolpekin, John Ray Bergado, and Rolf de By. Towards Automated Delineation of Smallholder Farm Fields From VHR Images Using Convolutional Networks. In *IGARSS 2019 - 2019 IEEE International Geoscience and Remote Sensing Symposium*, pages 3836–3839, Yokohama, Japan, July 2019. IEEE.

[Rieke, 2017] Christoph Rieke. Deep learning for instance segmentation of agricultural fields, thesis, 2017.

[Ronneberger *et al.*, 2015] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation, 2015.

[Usman and Beiji, 2012] Babawuro Usman and Zou Beiji. Satellite imagery cadastral features extractions using image processing algorithms: A viable option for cadastral science. *International Journal of Computer Science Issues IJCSI*, Vol 9, 08 2012.

[Xia *et al.*, 2019] Xue Xia, Claudio Persello, and Mila Koeva. Deep fully convolutional networks for cadastral boundary detection from uav images. *Remote Sensing*, 11:1725, 07 2019.

[Zhang *et al.*, 2016] L. Zhang, L. Zhang, and B. Du. Deep learning for remote sensing data: A technical tutorial on the state of the art. *IEEE Geoscience and Remote Sensing Magazine*, 4(2):22–40, 2016.

[Zhang *et al.*, 2019] J. Zhang, A. Saha, Z. Zhu, and M. A. Mazurowski. Hierarchical convolutional neural networks for segmentation of breast tumors in mri with application to radiogenomics. *IEEE Transactions on Medical Imaging*, 38(2):435–447, 2019.

[Zou *et al.*, 2004] Kelly Zou, Simon Warfield, Aditya Bharatha, Clare Tempany, Michael Kaus, Steven Haker, William Wells, Ferenc Jolesz, and Ron Kikinis. Statistical validation of image segmentation quality based on a spatial overlap index. *Academic radiology*, 11:178–89, 02 2004.