Optimizing Crop Cut Collection for Determining Field-Scale Yields in an Insurance Context

Ritvik Sahajpal[†]
Geographical Sciences
University of Maryland
College Park, Maryland, USA
ritvik@umd.edu

Inbal-Becker Reshef
Geographical   Sciences
University of Maryland
College Park, Maryland, USA
ireshef@umd.edu

Sylvain Coutu
Innovation Special Lines Center
Swiss-Re
Switzerland
sylvain.coutu@nasaharvest.org

ABSTRACT

Accurately determining crop yields at field-scale can help farmers estimate their net profit, enable insurance companies to ascertain payouts, and when aggregated at regional and national scales, crop yield estimates are critical in ensuring food security. Over the last few decades, crop cuts have been widely used to estimate field-scale crop yields. Crop cuts, while cost prohibitive, are the most reliable way to estimate yields at field level. We present a novel machine-learning based method to optimize the number and location of fields selected for performing crop cuts to drive down costs while maintaining the capacity to accurately predict crop yields at field-scale. This method is applied to crop cut data collected through a partnership between NASA Harvest and Swiss Re (Public-Private Partnership) in Ukraine in 2018 and 2019 for multiple crops (including Winter Wheat, Maize and Soybeans). We demonstrate the utility of specific bands and extracted features in improving model performance and show that our machine-learning model can explain nearly 70% of the variation in yields while saving up to 20% of the costs incurred in obtaining these crop cuts. We explore the trade-off between the number of crop cuts performed and model performance and demonstrate that our method has generalizability across agro-ecological zones.

KEYWORDS

Crop Yield Modeling, Mixed Effect Models, Optimizing Crop Cut Collection, Food Security, Public-Private Partnership

1   Introduction

Ensuring food-security via timely and accurate estimation of crop yields from field to regional to national and global scales is a priority policy goal for the U.N., as enshrined in the U.N. sustainable development goals [1]. Apart from the food security goals, an accurate estimation can also assist farmers in maximizing their profits and enable insurance companies to determine and deliver payouts in an efficient and cost-effective manner in case of adverse conditions. Given the need to deliver the insurance payouts in a timely manner, traditionally insurance

[†]Corresponding author
[*]All authors are also members of NASA Harvest (https://nasaharvest.org/)

companies have relied on a large pool of trained manual labor, also called loss adjusters, to visit farmers who report yield shortfalls, assess whether their claims are accurate and determine the requisite payoff.

To perform the assessment of crop yields in a field, a loss adjuster typically performs multiple crop cuts. A crop cut involves marking a subplot within the field, and subsequent measurement of production and area to get a field-scale estimate. A typical subplot has the size of a square which is 1m across. By averaging the yields from multiple representative crop cuts within a field, the loss adjuster can generally get an accurate estimate of the field scale yield without having to perform a much more time consuming and expensive full field harvest [2].

However, there are logistical and scaling challenges involved in performing crop cuts at large spatial scales. These challenges pertain to both having a small enough team of loss adjusters to keep overhead costs low for the insurance company as well as having enough people to perform crop- cut based assessments in a timely manner, once a farmer reports a pre-harvest yield loss event. To develop a strategy towards addressing this challenge, NASA Harvest collaborated with Swiss Re - a global reinsurer, Green Triangle - a tech company specialized in agro data collection, and AgroRisk Ltd - a local loss adjustment company. Several agronomists were deployed in Ukraine in order to collect crop cut information for the six main crops grown in Ukraine (Wheat, Maize, Soybean, Barley, Sunflower and Canola) for 2018 and 2019. The fields were spread across Ukraine (figure 1), and, for each field, the agronomists collected georeferenced information on crop type, planting date, yield and crop condition. The agronomist estimated crop conditions correspond closely to the yield measures via crop cuts (figure 2).
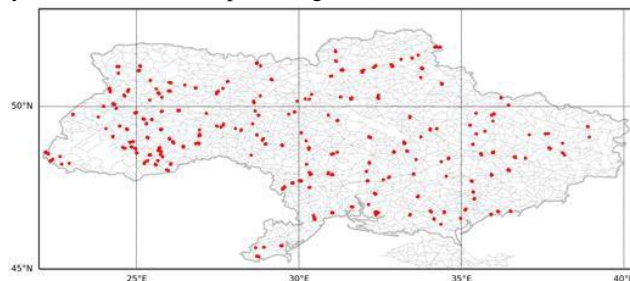
Figure 1. Fields where crop cuts were collected for Wheat in Ukraine in 2018 and 2019. The red dots indicate field locations and the gray lines are the road networks.
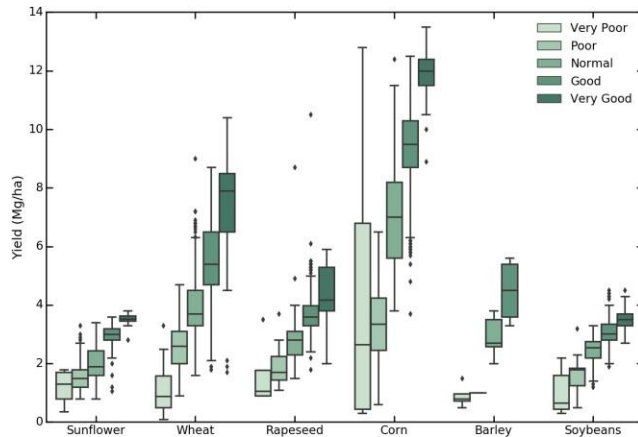

Figure 2. Crop cut yield varying by crop and estimated crop condition.

In the present work, we focus on the dominant crop in Ukrainian agriculture: Wheat. Our strategy focused on using satellite data that can provide a scalable, repeatable, low cost and crop agnostic approach to monitor crop yields and conditions on large acreages of agricultural land [3]. We attempt to answer the following science questions:

1. How can satellite derived vegetation vigor metrics be used to assess crop yields at field scale?
2. How can the number and location of crop cuts be optimized to drive down costs and time while maintaining model accuracy?

2 Methods

Our approach has two components. First, we extracted the Harmonized Landsat-8 Sentinel-2 (HLS) data [4] for nearly 1800 crop cuts in 487 Wheat fields for the years 2018 and 2019. HLS provides 30m resolution optical data from which we derived a measure of crop vigor called the Green Chlorophyll Vegetation Index (GCVI). GCVI has been demonstrated as an effective metric to assess crop yields [5] in a variety of growing conditions and for a variety of crops. The biophysical basis of this approach is that the factors impacting crop growth and yield first affect the photosynthetically active biomass in a plant, and that effect can be captured through indicators like NDVI. Since optical satellite data has gaps and often suffers from cloud contamination, we applied a smoothing algorithm (figure 3, [6]) to obtain an upper envelope time-series to the HLS data points for each field to fill gaps in the data.
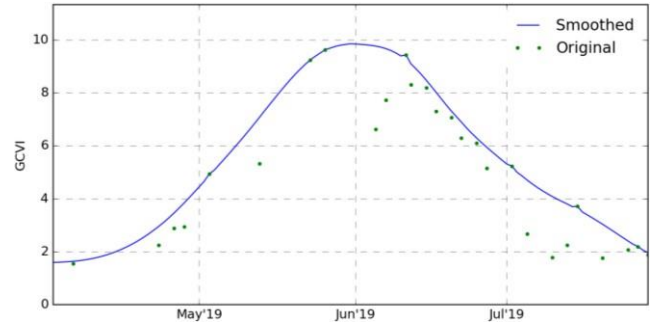

Figure 3. A smooth upper envelope curve fitted over the original satellite data.

For each of the time-series, we computed multiple time-series characteristics or features [7]. These features capture various properties of the vegetation index that can potentially impact crop growth and yields and include the following: value of peak GCVI, area under the curve till peak GCVI, time needed (number of days) to attain peak GCVI, count above mean, count below mean, longest strike above mean, longest strike below mean, mean change, ratio beyond 1 sigma, decline after senescence, absolute sum of changes and standard deviation of the time-series signal.

We averaged the crop cut yields for each field to obtain an estimate for the field-scale yield. Nested cross-validation was used to avoid high variance in performance estimate due to small size of test set, and exhaustive search (based on a hyper-parameter search across a grid of parameter values) was used to determine optimal model hyper-parameters for each ML model. We applied the following machine-learning algorithms on the 13 features computed for each field from the time-series to predict crop yields for Wheat at field scale.

1. Linear regression
2. Lasso regression: constrains the selection of features leading to a parsimonious model
3. Random Forests
4. Cubist: An extension to random forests where the terminal leaves contain linear regression models. This helps them extrapolate better than vanilla random forests [8].
5. Mixed Effect Random Forests (MERF, [9]): A modification to random forest models that are relevant if clustering is present in data (e.g. crop yields can show different variations w.r.t satellite indicators at regional scales)

To optimize the selection of crop cuts, we wanted to minimize the distance traveled between crop cuts while maximizing model performance, based on metrics like RMSE or R-squared. We measured the distance between crop cuts based on the length of the minimum spanning tree (MST) that connects all the crop cut sites with the shortest possible network of roads such that there is no cycle in the network [10]. We subsequently combined the machine learning models with a genetic algorithm [11] to optimize the number and location of crop cuts by minimizing the

length of the MST connecting the crop cuts while simultaneously maximizing the model R-squared.

3   Results

We evaluated model performance based on the R-squared estimate relating model predicted crop cut yields to the measured yields. MERF model performed the best, followed closely by Cubist, while the linear regression model was the worst. Our best R-squared values approach 0.67 for the MERF model. As we reduce the number of crop cuts used in training the model, there is a drop-off in the R-squared values. The drop-off is the steepest in the linear regression model and the slowest in the MERF and Cubist models (figure 4).
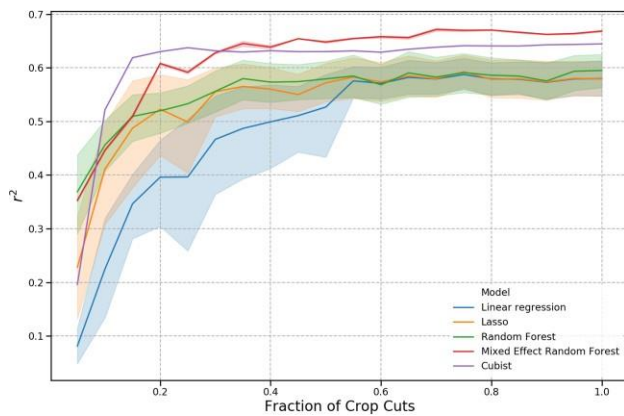

Figure 4. R-squared values of model performance as a function of the fraction of crop cuts used in training the model.

Generally, all models maintain their respective R-squared values as the fraction of crop cuts used in model training is reduced from 1 to nearly 0.6. This can be explained by the fact that multiple crop cuts are sampled from each fields, and reduction in the number of crop cuts per field will not impact model performance as long as the remaining crop cuts are representative of that field. This is true especially for larger fields with more homogeneous growth patterns and small variations n yield outcomes across the field.

A reduction in the fraction of crop cuts also implies that a smaller MST can span the remaining crop cuts. Indeed, our results show that we can reduce the MST distance by 22% (from 6,750 to 5,250 km) without significant reduction in model performance for the MERF model for both R-squared and RMSE (figure 5, 6). A visual inspection of the MST for three different values of crop cut fraction shows that the genetic algorithm based optimization model does tend to optimize the MST distance (figure 7).

We also computed crop condition classes from the observed and predicted yields by dividing them into 4 equal quartiles. Our

confusion matrix for Wheat (Figure 8) demonstrates that the best performing model captures the crop condition class very well, rarely mistaking a poor performing field for a field that is performing in the top tier of yields.
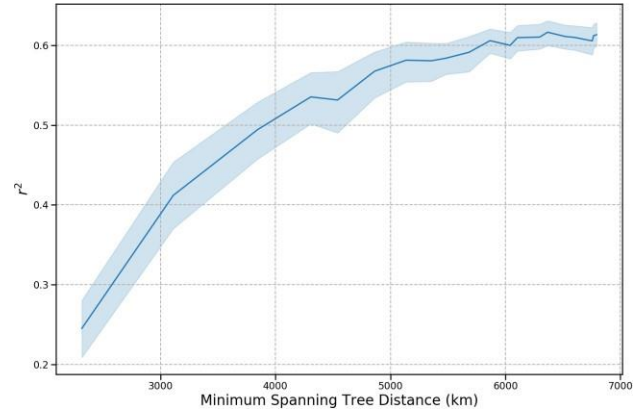

Figure 5. Reduction in R-squared value as a function of the MST distance between the fraction of crop cuts selected for training the model.
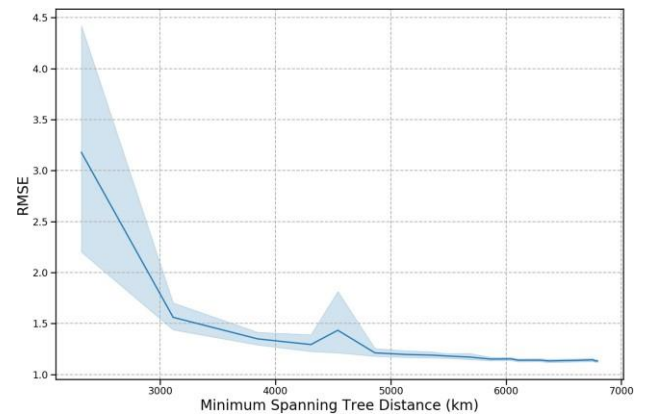

Figure 6. Reduction in RMSE value as a function of the MST distance between the fraction of crop cuts selected for training the model.
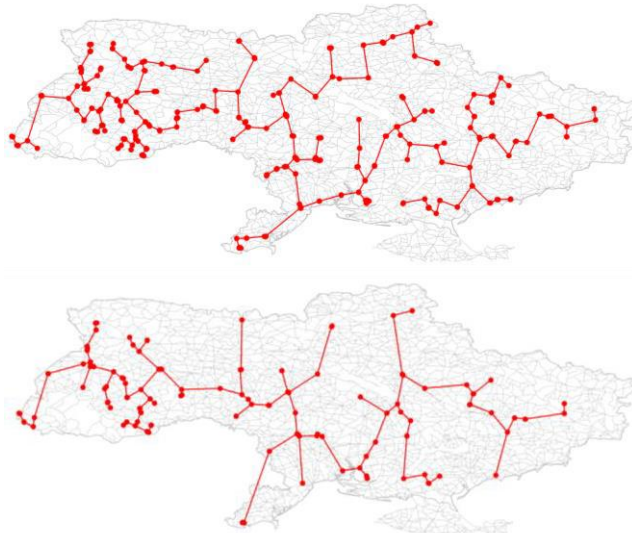
Figure 7. Minimum spanning tree in red based on genetic algorithm based optimized selection of crop cuts for three different fractions of crop cuts (from top to bottom: 1.0, 0.6, 0.2).

Figure 8. A confusion matrix comparing the crop condition classes from observations vs model predictions.

4   Conclusions

Our approach has several benefits over traditional crop yield monitoring techniques. First, Earth Observation (EO) data does not suffer from observational bias, allowing for more objective, quantitative, and scientific estimation of crop conditions. Second, the EO datasets capture different biophysical components that relate to various attributes of crop health including leaf temperature and water use efficiency. This will enable the tracking of a range of crop threats, each of which varies in impact on plant physiology and yield. Finally, our model can be updated in a Bayesian framework using field observations of crop condition and threats. In the present problem, we demonstrated

that we could reduce the distance traveled to collect crop cuts by 22%, and the number of fields visited by 40% without sacrificing model performance. Mixed effect based approaches provide an effective approach towards yield forecasting since they are able to capture the regional variability in crop yields. By combining these approaches with genetic algorithms and a clustering metric like MST, we can determine the optimal location and number of crop cuts in a variety of geographies and for diverse cropping systems.

Beyond the straightforward benefit of such approaches for reduction of costs and increasing timeliness of payouts in the insurance industry, optimized collection of agriculture data can help the development of national crop statistics, and provide the necessary underlying database for future machine learning based models [12].

REFERENCES

[1] Whitcraft, A.K., Becker-Reshef, I., Justice, C.O., Gifford, L., Kavvada, A. and Jarvis, I., 2019. No pixel left behind: Toward integrating Earth Observations for agriculture into the United Nations Sustainable Development Goals framework. Remote Sensing of Environment, 235, p.111470.

[2] Fermont, A. and Benson, T., 2011. Estimating yield of food crops grown by smallholder farmers. International Food Policy Research Institute, Washington DC, 1, p.68.

[3] Lobell, D.B., Thau, D., Seifert, C., Engle, E. and Little, B., 2015. A scalable satellite-based crop yield mapper. Remote Sensing of Environment, 164, pp.324-333.

[4] Claverie, M., Masek, J.G., Ju, J. and Dungan, J.L., 2017. Harmonized landsat-8 sentinel-2 (HLS) product user's guide. National Aeronautics and Space Administration (NASA): Washington, DC, USA.

[5] Jain, M., Singh, B., Srivastava, A.A.K., Malik, R.K., McDonald, A.J. and Lobell, D.B., 2017. Using satellite data to identify the causes of and potential solutions for yield gaps in India's Wheat Belt. Environmental Research Letters, 12(9), p.094011.

[6] Chen, J., Jönsson, P., Tamura, M., Gu, Z., Matsushita, B. and Eklundh, L., 2004. A simple method for reconstructing a high-quality NDVI time-series data set based on the Savitzky–Golay filter. Remote sensing of Environment, 91(3-4), pp.332-344.

[7] Christ, M., Braun, N., Neuffer, J. and Kempa-Liehr, A.W., 2018. Time series feature extraction on basis of scalable hypothesis tests (tsfresh–a python package). Neurocomputing, 307, pp.72-77.

[8] Quinlan, J.R., 1993, June. Combining instance-based and model-based learning. In Proceedings of the tenth international conference on machine learning (pp. 236-243).

[9] Hajjem, A., Bellavance, F. and Larocque, D., 2014. Mixed-effects random forest for clustered data. Journal of Statistical Computation and Simulation, 84(6), pp.1313-1328.

[10] Graham, R.L. and Hell, P., 1985. On the history of the minimum spanning tree problem. Annals of the History of Computing, 7(1), pp.43-57.

[11] Deb, K., Pratap, A., Agarwal, S. and Meyarivan, T.A.M.T., 2002. A fast and elitist multiobjective genetic algorithm: NSGA-II. IEEE transactions on evolutionary computation, 6(2), pp.182-197.

[12] Coutu, S., Becker-Reshef, I.,Whitcraft, A.K., Justice, C., 2020. Food security: underpin with public and private data sharing. Nature, 578 (515).