Dynamic Poverty Prediction with Vegetation Index

Binh Tang
Department of Statistical Science
Cornell University
Ithaca, NY 14850
bvt5@cornell.edu

Ying Sun
School of Integrative Plant Science
Cornell University
Ithaca, NY 14850
ys776@cornell.edu

Yanyan Liu
Markets, Trade and Institutions Division
International Food Policy Research
Washington, DC 20005
y.liu@cgiar.org

David S. Matteson
Department of Statistical Science
Cornell University
Ithaca, NY 14850
matteson@cornell.edu

Abstract

Accurate and timely estimates of economic status are critical to policy-makers in the world's poorest countries. Previous work has applied convolutional neural networks (CNNs) to high-resolution satellite imagery to perform poverty prediction. Although promising, such imagery has limited access and lacks a temporal signal. We show that publicly available, moderate-resolution vegetation index can be used with CNNs to produce equally accurate poverty estimates for developing countries that are heavily dependent on agriculture. In contrast to previous work, the continuous streaming of well known vegetation indices also allows us to update our estimates in light of weather shocks, opening up the possibility of making dynamic poverty mapping at minimal cost.

1 Introduction

Despite global economic growth, 330 million people are still living in extreme poverty in Africa [2]. The United Nations has acknowledged poverty as one of the greatest challenges facing humanity and aims to end extreme poverty in all forms by 2030 [11]. To achieve the goal, policy makers often rely on complex household surveys to measure poverty and allocate resources accordingly [1]. Since the data collection process is costly and time consuming, there is a lack of good-quality data to assess poverty regularly [8]. Inexpensive and scalable approaches to poverty prediction are therefore needed to complement household surveys.

Recent advances in remote sensing and machine learning have opened up a new path for poverty prediction. High resolution satellite images are rich in content and available globally, providing an objective view on the economic conditions of developing countries [4, 7, 10]. The increasing abundance of these images lends themselves to convolutional neural networks (CNNs), a deep learning approach that has recently seen tremendous success in many computer vision tasks [15, 16, 18]. In a critically acclaimed paper [8], Jean *et al.* applied CNNs on Google's daytime satellite images to measure regional poverty in Africa, yielding results comparable to estimation based on past surveys.

Although promising, most current studies with satellite images are limited to providing fixed estimates of poverty maps. The poverty predictions in Jean et al. [8], for example, were constant scalars regardless of the time of prediction, simply due to the lack of continuous access to Google's proprietary

Workshop on Modeling and Decision-Making in the Spatiotemporal Domain, 32nd Conference on Neural Information Processing Systems (NIPS 2018), Montréal, Canada.

data. Given the unprecedented climate changes anticipated in sub-Saharan countries [13], dynamic poverty prediction is critical to timely interventions and policy evaluation.

In this paper, we leverage the continuous streaming of the normalized difference vegetation index (NDVI), one of the widely known satellite measurements of Earth's vegetation greenness, to estimate poverty indicators more frequently. NDVI measures the difference in red and near-infrared light reflectance resulted from photosynthesis; areas of barren rock have very low NDVI while dense croplands often have high NDVI values (see Figure 1). As ultra-poor regions heavily depend on agriculture [5], NDVI provides an uninterrupted signal for crop heath and poverty tracking in general.

Our contribution is twofold: (1) we demonstrate that publicly-available, moderate-resolution NDVI can help predict poverty in Malawi, Nigeria, Rwanda, Tanzania, and Uganda as well as state-of-art methods, and (2) we perform poverty prediction for an out-of-sample period and capture changes in poverty measures for ultra-poor regions in Uganda. Based on the availability of survey data, these countries are chosen for direct comparison with [8].

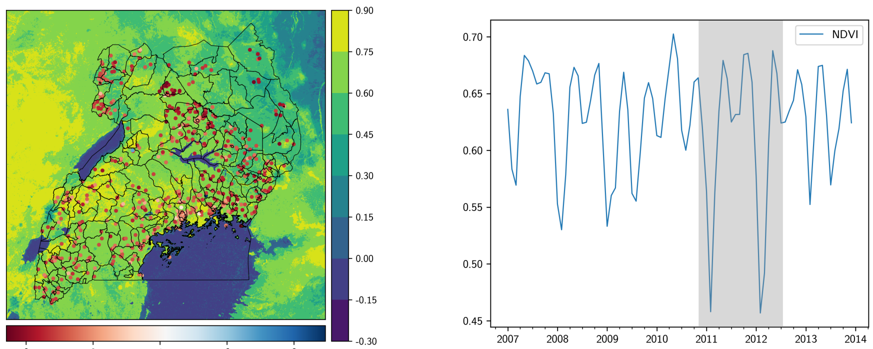

Figure 1: NDVI measurements for Uganda in 2011. On the *left*, the background image shows annual average NDVI with a vertical colorbar while the foreground scatters depict log consumption expenditures with a horizontal colorbar. On the *right*, the annual NDVI, spatially averaged over all survey locations, with notable drops during the 2011-2012 East Africa drought highlighted in gray.

2 Related Work

Recent studies on poverty prediction rely on passively collected data and statistical methods to circumvent the scarcity of household surveys. Researchers have fit linear models to nighttime light luminosities and found they strongly correlate with the gross domestic product of various countries [4, 7, 10]. Proprietary cell phone records of millions of subscribers in Rwanda and Senegal have also been used with tree-based classifiers and Gaussian process regressors to provide asset wealth index estimates [3, 14]. In another line of research, CNNs help extract predictive features from Google's high-resolution daytime images, providing accurate estimates of both consumption expenditure and asset wealth in multiple African countries [8, 21].

Although existing technologies can only measure NDVI at a lower resolution, recent work on poverty and health in Africa has proven its significant predictive power. Through increased crop yields, NDVI has been found to be positively correlated with child survival, nutrition, and anthropometric variables such as wasting [9]. Using spatial statistics techniques, Sedda *et al.* also [17] showed that the intensity of poverty varies inversely with NDVI in West Africa.

3 Data and Methods

Inspired by previous work, we apply CNNs to publicly available NDVI images to learn features useful for poverty prediction. Following Jean *et al.* [8], we use transfer learning and a two-step procedure to bypass the lack of labeled responses: (1) fine-tune a VGG-16 network [19] on NDVI images to predict nighttime light intensities, and (2) fit random forest regression models using NDVI features to predict poverty indicators. The combination of NDVI images and nighttime lights allows vegetation features indicative of economic activity to be learned and generalized to the poverty prediction task.

Country	Consumption Prediction (LSMS)			Asset Index Prediction (DHS)		
	Year	Jean et al. [8]	NDVI	Year	Jean et al. [8]	NDVI
Malawi	2013	0.37	0.341 ± 0.038	2010	0.55	0.498 ± 0.020
Nigeria	2013	0.42	0.387 ± 0.013	2013	0.68	0.738 ± 0.005
Rwanda	-	-	-	2010	0.75	0.725 ± 0.022
Tanzania	2012	0.55	0.603 ± 0.019	2010	0.57	0.638 ± 0.012
Uganda	2011	0.41	0.490 ± 0.012	2011	0.69	0.751 ± 0.007

Table 1: Spatially cross-validated r^2 values of the predictions of NDVI models relative to Jean *et al.* [8]. Separate models are fine-tuned and evaluated for different countries and surveys. For NDVI models, the means and standard deviations of r^2 values are reported using 5 independent trials.

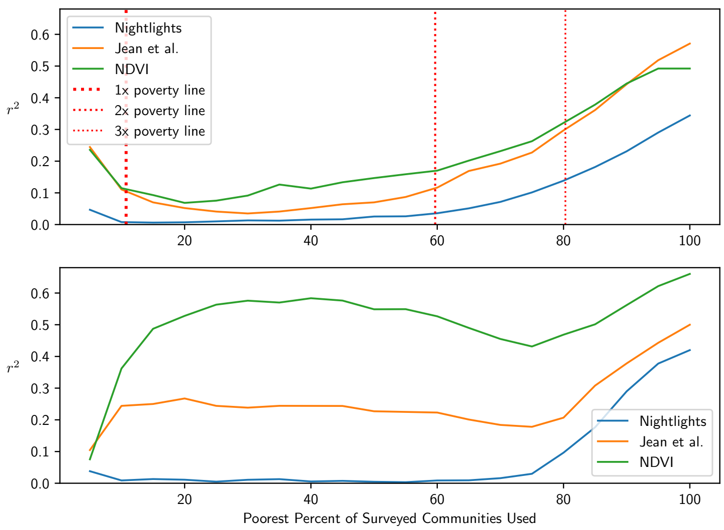

Figure 2: Spatially cross-validated results of NDVI models relative to nightlights and Jean *et al.* [8]. Nightlight-based models are random forests trained on scalar nighttime light intensities. The *top* figure shows r^2 values for estimating **consumption** using pooled observations across the four LSMS countries. We run separate trials for increasing percentages of the pooled dataset (e.g., the x-axis value of 60 indicates all surveyed communities below the 60th percentile of consumption are included. The *bottom* figure show similar r^2 values for estimating **asset index**.)

In the first step of our procedure, we start with a VGG-16 network pre-trained on ImageNet and adapt its fully connected layers to fit our input image sizes. Our inputs are annual average NDVI images, each 64 × 64 pixels in size and at a spatial resolution of 250 square meters per pixel. The images are sampled from a dataset produced by NASA's Terra satellite [6] and represent areas that are evenly spaced at 0.025 degree intervals. The network learns to map each NDVI image to the average value of annual nighttime light intensities that describe the same geographical region, as provided by the National Oceanic and Atmospheric Administration [12]. In contrast to [8, 21], we take a log transformation, but we do not discretize nightime light intensities.

In the second step, we extract features from NDVI images and fit regression models to predict two survey variables - logarithm of consumption expenditure and asset index. For direct comparison, we select the same surveys and follow the same preprocessing steps as in [8] (see Table 1 for the list of surveys). The conv5-2 layer of the fine-tuned network outputs a feature map of size 512 × 1 for each NDVI image, and we average feature maps of images whose centers are within five kilometers of a surveyed community. For each survey, we then train random forests on the 512-dimensional feature maps, using nested 5-fold spatial cross-validation to select hyperparameters, and output predictions.

In response to weather shocks, NDVI often changes over time (see Figure 1), and its feature maps can potentially capture and reflect these events in poverty predictions. Hence, we also fine-tune the network in the first step with updated NDVI images and nighttime lights. Predictions for out-of-sample periods are obtained by first training a random forest on the previous NDVI feature maps and testing it on the updated ones.

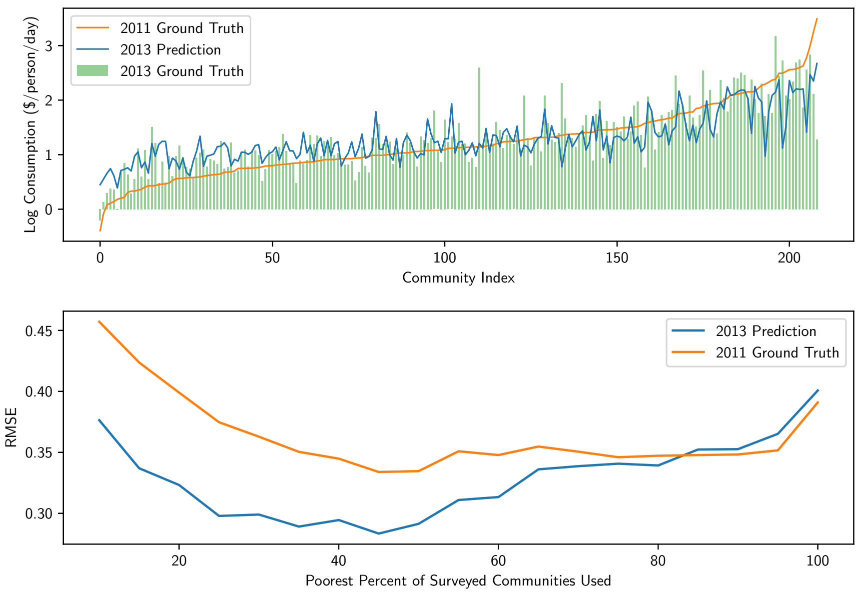

Figure 3: Consumption predictions for LSMS communities in Uganda made by a random forest model trained on 2011 data and tested on 2013 data. The *top* figure shows the ground-truth consumption along with predictions for LSMS communities ordered by 2011 data. The *bottom* figure shows RMSE values of the predictions for increasing percentages of the LSMS communities (e.g., the x-axis value of 60 indicates all communities below the 60th percentile in 2011 consumption are included).

4 Results

Our first experiment is to evaluate the predictive power of NDVI for poverty estimation. Following [8], we use expenditure data from the World Bank's Living Standards Measurement Study (LSMS) and asset index data from the Demographic and Health Surveys (DHS). Table 1 shows that our NDVI models are highly predictive of both average household consumption and average asset wealth. Spatially cross-validated predictions explain 34 to 60% of the variation in average consumption and 50 to 75% of the variation in average asset wealth across surveyed countries. In general, our models perform comparably to Jean *et al.* [8] when fit using data from individual countries.

When trained on pooled consumption or asset observations across all countries, our models perform significantly and consistently better than the state-of-art method (see Figure 2). We see an improvement of more than 100% in r^2 for asset index predictions for regions below the 2x poverty line, which is set at $1.9 per person per day by the World Bank. This observation agrees with our intuition that extremely poor communities depend most heavily on crop production. The modest improvements in consumption prediction can be partly explained by the fact that consumption data is noisier [8, 20].

In the second experiment, we study whether temporal changes in NDVI are indicative of poverty changes. Because surveys from different years are generally conducted at different locations, we limit this experiment to 209 communities in Uganda that are part of both the 2011-2012 and 2013-2014 LSMS surveys. As the 2011-2012 East Africa drought affected a large area of Uganda, the consumption distribution for these communities changes quite significantly between the surveys. Figure 3 shows that our random forest model can translate the increase in annual NDVI from 2011 to 2013 to reflect increased consumption in the poorest communities following the drought. In contrast, models that rely on static inputs such as *et al.* [8] can only perform as well as the 2011 ground truth when tested on the 2013 data.

5 Conclusion

In this paper, we have leveraged CNNs to extract features from NDVI images that are highly predictive of poverty. We demonstrate that publicly available, moderate-resolution NDVI can predict poverty as well as high-resolution images constrained by Google's licensing terms. Our model based on NDVI can also produce dynamic poverty estimates, potentially helping policy-makers make more informed and timely decisions.

6 Acknowledgement

This work was supported by the Cornell University Atkinson Center for a Sustainable Future (AVF-2017).

References

[1] Abhijit Banerjee, Esther Duflo, Nathanael Goldberg, Dean Karlan, Robert Osei, William Parienté, Jeremy Shapiro, Bram Thuysbaert, and Christopher Udry. A multifaceted program causes lasting progress for the very poor: Evidence from six countries. *Science*, 348(6236):1260799, 2015.

[2] Kathleen Beegle, Luc Christiaensen, Andrew Dabalen, and Isis Gaddis. *Poverty in a rising Africa*. The World Bank, 2016.

[3] Joshua Blumenstock, Gabriel Cadamuro, and Robert On. Predicting poverty and wealth from mobile phone metadata. *Science*, 350(6264):1073–1076, 2015.

[4] Xi Chen and William D Nordhaus. Using luminosity data as a proxy for economic statistics. *Proceedings of the National Academy of Sciences*, 108(21):8589–8594, 2011.

[5] Xinshen Diao, Peter Hazell, and James Thurlow. The role of agriculture in african development. *World development*, 38(10):1375–1383, 2010.

[6] Kamel Didan. MOD13Q1 MODIS/Terra Vegetation Indices 16-Day L3 Global 250m SIN Grid V006. *NASA EOSDIS LP DAAC*, 2015.

[7] J Vernon Henderson, Adam Storeygard, and David N Weil. Measuring economic growth from outer space. *American economic review*, 102(2):994–1028, 2012.

[8] Neal Jean, Marshall Burke, Michael Xie, W Matthew Davis, David B Lobell, and Stefano Ermon. Combining satellite imagery and machine learning to predict poverty. *Science*, 353(6301):790–794, 2016.

[9] Kiersten Johnson and Molly E Brown. Environmental risk factors and child nutritional status and survival in a context of climate variability and change. *Applied Geography*, 54:209–221, 2014.

[10] Yong Suk Lee. International isolation and regional inequality: Evidence from sanctions on north korea. *Journal of Urban Economics*, 103:34–51, 2018.

[11] United Nations. Transforming our world: The 2030 agenda for sustainable development. *Resolution adopted by the General Assembly*, 2015.

[12] NOAA National Centers for Environmental Information. Nighttime Lights Time Series. 2010.

[13] Martin Parry, Martin L Parry, Osvaldo Canziani, Jean Palutikof, Paul Van der Linden, and Clair Hanson. *Climate change 2007-impacts, adaptation and vulnerability: Working group II contribution to the fourth assessment report of the IPCC*, volume 4. Cambridge University Press, 2007.

[14] Neeti Pokhriyal and Damien Christophe Jacques. Combining disparate data sources for improved poverty prediction and mapping. *Proceedings of the National Academy of Sciences*, 114(46):E9783–E9792, 2017.

[15] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 779–788, 2016.

[16] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision*, 115(3):211–252, 2015.

[17] Luigi Sedda, Andrew J Tatem, David W Morley, Peter M Atkinson, Nicola A Wardrop, Carla Pezzulo, Alessandro Sorichetta, Joanna Kuleszo, and David J Rogers. Poverty, health and satellite-derived vegetation indices: their inter-spatial relationship in west africa. *International health*, 7(2):99–106, 2015.

[18] Pierre Sermanet, David Eigen, Xiang Zhang, Michaël Mathieu, Rob Fergus, and Yann LeCun. Overfeat: Integrated recognition, localization and detection using convolutional networks. *arXiv preprint arXiv:1312.6229*, 2013.

[19] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.

[20] Jessica E Steele, Pål Roe Sundsøy, Carla Pezzulo, Victor A Alegana, Tomas J Bird, Joshua Blumenstock, Johannes Bjelland, Kenth Engø-Monsen, Yves-Alexandre de Montjoye, Asif M Iqbal, et al. Mapping poverty using mobile phone and satellite data. *Journal of The Royal Society Interface*, 14(127):20160690, 2017.

[21] Michael Xie, Neal Jean, Marshall Burke, David Lobell, and Stefano Ermon. Transfer learning from deep features for remote sensing and poverty mapping. *arXiv preprint arXiv:1510.00098*, 2015.