

MAPPING POVERTY IN THE PHILIPPINES USING MACHINE LEARNING, SATELLITE IMAGERY, AND CROWD-SOURCED GEOSPATIAL INFORMATION

I. Tingzon^{1*}, A. Orden¹, K.T. Go¹, S. Sy¹, V. Sekara², I. Weber³, M. Fatehkia³, M. García-Herranz², D. Kim²

¹ Thinking Machines Data Science, Twenty-Four Seven McKinley, Taguig, Philippines – (issa, ardie, kevin, stef)@thinkingmachin.es

² UNICEF Office of Innovation, 3 United Nations Plaza, New York, USA – (vsekara, mgarciaherranz, dokim)@unicef.org

³ Qatar Computing Research Institute, Education City, Doha, Qatar – (iweber, mfatehkia)@hbku.edu.qa

Commission IV

KEY WORDS: Poverty estimation, Machine learning, Deep learning, Transfer learning, Remote sensing, GIS, Volunteered geographic information, OpenStreetMap

ABSTRACT:

Mapping the distribution of poverty in developing countries is essential for humanitarian organizations and policymakers to formulate targeted programs and aid. However, traditional methods for obtaining socioeconomic data can be time-consuming, expensive, and labor-intensive. Recent studies have demonstrated the effectiveness of combining machine learning and satellite images to estimate wealth in sub-Saharan African countries (Xie et al., 2016, Jean et al., 2016). In this study, we investigate the extent to which this method can be applied in the context of the Philippine archipelago to predict four different socioeconomic indicators: wealth level, years of education, access to electricity, and access to water. We also propose an alternative, cost-effective approach that leverages a combination of volunteered geographic information from OpenStreetMap and nighttime lights satellite imagery for estimating socioeconomic indicators. The best models, which incorporate regional indicators as predictors, explain approximately 63% of the variation in asset-based wealth. Our findings also indicate that models trained on publicly available, volunteer-curated geographic data achieve the same predictive performance as that of models trained using proprietary satellite images.

1. INTRODUCTION

Despite best efforts in implementing poverty alleviation programs, the Philippines still lags behind its Southeast Asian neighbors in terms of poverty eradication, with approximately 22 million Filipinos living below the national poverty line (Philippine Statistics Authority, 2018). A major challenge in fighting poverty today is the lack of reliable socioeconomic data, which is often expensive, time-consuming, and labor-intensive to collect. Conducting on-the-ground household surveys in the Philippines can cost up to 1.5M USD. Such surveys are done only every 3 to 5 years and are often aggregated to the regional or provincial level when reported to the public (Juan-Albacea, 2009). Without more granular and up-to-date data to guide their policies and programs, development organizations and government agencies risk allocating their limited resources in the wrong areas.

In recent years, major advancements in computer vision research and an increasing availability of geospatial resources have enabled novel methods for estimating socioeconomic indicators (Jean et al., 2016, Babenko et al., 2017, Engstrom et al., 2017). To tackle the problem of poverty eradication, we look towards combining machine learning with geospatial information as a fast, low-cost, and scalable means of providing granular poverty estimates. In this study, we examine the extent to which geospatial data including nighttime lights, daytime satellite imagery, human settlement data, and crowd-sourced information can be used to estimate socioeconomic well-being in the Philippines.

To summarize, our work primarily seeks to answer the following questions: (1) Are remote sensing-based methods developed for poverty prediction in other countries applicable within the Philippine context? and (2) How well do predictive models trained on publicly available crowd-sourced geospatial information compare against state-of-the-art remote sensing-based methods for poverty estimation in the Philippines?

2. RELATED WORK

In recent years, more research works have started focusing on the development of novel techniques using a variety of geospatial data sources for estimating socioeconomic development.

Study	Approach	Countries	Out-of-sample R^2
Engstrom et al., 2017	Derived features	Sri Lanka	0.60
Jean et al., 2016	Transfer learning	Malawi	0.55
		Tanzania	0.57
		Nigeria	0.68
		Uganda	0.69
		Rwanda	0.75
Head et al., 2017	Transfer learning	Haiti	0.51
		Nepal	0.64
Babenko et al., 2017	End-to-end learning	Mexico	0.57

Table 1. A summary of the out-of-sample R^2 results of past studies on poverty prediction using remote sensing techniques.

* Corresponding author

In 2016, Jean et al. proposed a method that combines deep learning and satellite imagery for poverty prediction. The study uses convolutional neural networks to map daytime satellite images to their corresponding nighttime light intensity levels in order to implicitly identify important patterns in the satellite imagery (e.g. roads, buildings, land cover) that are indicative of wealth. The study obtained R^2 results of 0.55 - 0.75 for asset-based wealth estimation in five different sub-Saharan African countries.

In 2017, the study was followed by that of Head et al. which demonstrated the replicability of the results and extended the study to estimate a range of socioeconomic indicators in countries outside of Africa, including Haiti and Nepal. The study investigated the model's efficacy in predicting other human development indicators, including education, health, access to electricity, access to water, and anthropomorphic measures. Ultimately, the study showed that the method cannot be trivially applied to estimate other socioeconomic indicators outside of wealth, especially if these indicators do not exhibit any clear relationships with nighttime lights.

More recently, Engstrom et al. implemented an approach that involves explicitly deriving features from daytime satellite images using CNNs, semi-automated classification models, and texture-based approaches. The extracted features were then used to estimate poverty in 1,291 administrative units in Sri Lanka. Their study shows that wealth in Sri Lanka can be reconstructed with an R^2 of 0.60 using object and texture features derived from satellite images. Features extracted include density and count of buildings, the presence of building shadows (a proxy for building height), the number of cars, road length and density, agricultural land cover, and roof type or material.

Babenko et al. proposed an end-to-end deep learning approach that uses CNNs to directly map medium and high resolution satellite images to poverty rates in Mexico. Using this method, they were able to achieve R^2 results of up to 0.57 when incorporating satellite-estimated land use into the model and 0.61 when modeling poverty only in urban settings. We summarize the results of these remote sensing-based studies in Table 2.

Other non-remote sensing-based methods for predicting have also been proposed. Previous researchers (Blumenstock et al., 2015, Fernando et al., 2018, Hernandez et al., 2017) have conducted studies using anonymized call detail records (CDRs) from mobile phone networks to predict the poverty measures derived from census data in Rwanda, Sri Lanka, and Guatemala respectively. The studies were able to demonstrate the efficacy of reconstructing high-resolution maps of wealth distribution using only CDR.

3. DATA AND PRE-PROCESSING

In this section, we describe the five different data sources and preprocessing steps done in this study. We summarize the datasets in Table 2.

Demographic and Health Survey (DHS)	
Year collected	2017
No. of clusters	1,213
No. of households	27,496
Nighttime Lights Data (VIIRS DNB)	
Year collected	2016
Number of ~0.25 km ² pixels	134,540

Low intensity pixels	94,680
Moderately low intensity pixels	27,415
Medium intensity pixels	9,389
Moderately high intensity pixels	2,161
High intensity pixels	895
Daytime Satellite Imagery (Google Static Maps)	
Year collected	2018
No. of images in training set	150,000
No. of images in validation set	13,454
Human Settlement Data (HRSL)	
Year collected	2017
No. of road features	16
No. of building features	28
No. of POI features	101

Table 2. Primary datasets used in this study.

3.1 Demographic and Health Survey

We used the 2017 Philippine Demographic and Health Survey (DHS) as a measure of ground truth for the different socioeconomic indicators. Conducted every 3 to 5 years, the Philippine Statistical Authority (PSA) collects nationally representative information on social, economic, and health-related outcomes across hundreds of households, which are grouped into clusters of 2 to 44 households (Philippine Statistics Authority, 2018, Burgert et al., 2013). In line with the United Nations Sustainable Development Goals, we focused our analysis on a subset of survey questions and derived the following socioeconomic indicators from the DHS dataset:

Wealth Index. Our primary measure of socioeconomic well-being is the "wealth index", which is computed as the first principal component of attributes related to common asset ownership (e.g., roof material, television, housing material) on a per-household level. We get the mean wealth index per cluster as it is reported in the DHS dataset and do no further transformations.

Education completed. The DHS captures information on the number of years of education completed by household members over 6 years old. We aggregated this by computing the mean years of education completed across all households per cluster.

Access to Electricity. The DHS dataset contains information on the number of affirmative responses to the survey question related to access to electricity. We aggregated this information by getting the proportion of households with electricity access per cluster.

Access to Water. The DHS dataset contains information on the total travel time in minutes to get to a water source. If water source is on-site, time is set to zero. We get the mean time to access a water source across all households per cluster.

3.2 Nighttime Luminosity Data

The nighttime lights (NTL) data is taken from the Visible Infrared Imaging Radiometer Suite Day/Night Band (VIIRS DNB) for the year 2016, produced in 15 arc-second geographic grids (NOAA National Centers for Environmental Information, 2016). The VIIRS DNB NTL data includes a continuous luminosity level from 0 to 122 for the Philippines, with 0 being the darkest pixel. By observing the histogram of the nighttime light intensities, we assigned the nighttime light intensity values into the following five distinct classes: low intensity (zero pixel values), moderately low intensity (0.05-2), medium intensity (2-

15), moderately high intensity (15-30), and high intensity (30-122).

Figure 1 illustrates the relationship between the average night time lights per cluster and the four (4) different socioeconomic indicators: wealth index, education, access to electricity, and access to water.

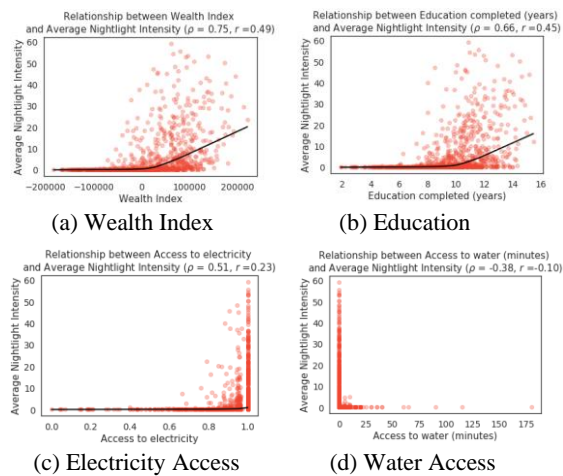


Figure 1. Relationship between average nighttime light luminosity and the Philippine socioeconomic indicators (wealth index, education completed, access to electricity, access to water) at the cluster level. Black line indicates a LOESS curve fitted to the data.

3.3 Daytime Satellite Imagery

We retrieved a number of satellite images per cluster based on the cluster centroids reported in the DHS dataset, where each cluster location is defined by the mean latitude and longitude of the households, with added noise to preserve the privacy of the households (Philippine Statistics Authority, 2018). Each cluster location is also labeled with a tag that indicates whether or not it is within a rural or urban area. We obtained up to 400 square tiles of satellite images within a 5 km radius for rural areas and up to 60 square tiles within a 2 km radius for urban areas. These tiles surround each cluster centroid and each tile corresponds to a pixel in the VIIRS DNB NTL dataset. Using Google Static Maps API, we downloaded a total of 134,540 images with a zoom level of 17, scale of 1, and pixel resolution of approximately 1.25 m. The size of each image is 400×400 pixels and matches the 0.25 km² land area covered by a single pixel of night time lights data.

3.4 High Resolution Settlement Data

The High Resolution Settlement Layer dataset by Tiecke et al. was used to identify satellite images that contained no human settlements. The dataset provides estimates of human population distribution at a resolution of 1 arc-second (30 m). The population estimates were based on recent census data and high resolution satellite imagery (0.5 m) from DigitalGlobe. Using building footprints as a proxy for human settlements, Tiecke et al. developed a CNN model to detect buildings from satellite images; population estimates were then assigned using proportional allocation over a wide coverage area, encompassing both urban and rural areas (Tiecke et al., 2017).

3.5 OpenStreetMap Data

More and more researchers are turning to volunteer-curated geographic information and open geospatial datasets to study socioeconomic development, social inequalities, and territorial conflicts (Gervasoni et al., 2018, Mahabir et al., 2018, Grippa et al., 2018). One of the more popular geospatial data crowdsourcing platforms is OpenStreetMap (OSM), a global geospatial database containing billions of entries of volunteered geographic information, maintained by a massive community of mappers from around the world all working towards the goal of curating accurate and complete geospatial data.

The community of OSM contributors typically consists of individual mappers, university researchers, volunteer communities, and non-profit organizations. Such organizations regularly organize field mapping activities, workshops, and events that promote geospatial data mapping and contribution. For local governments units and volunteer organizations such as the Humanitarian OpenStreetMap Team (HOT), OSM has been used largely for humanitarian work, disaster risk management, and emergency relief operations.

In recent years, large corporations such as Facebook, Amazon, and Microsoft as well transportation network companies like Grab and Uber have also started making major contributions to OSM (Anderson et al., 2019). Given massive computational resources, private corporations are able to use computer vision with satellite imagery to identify and map large-scale road networks. In fact, a recent study has found that user-generated road maps in OSM are approximately 83% complete as of 2016, with over 40% of countries having a fully mapped street network (Barrington-Leigh and Millard-Ball, 2017). Meanwhile, non-corporate mappers are more active contributors of building and point-of-interest data, which are typically more difficult to distinguish using satellite imagery and are thus more accurately mapped using local domain knowledge and on-the-ground field surveys.

In this study, we obtained OpenStreetMap (OSM) data for the Philippines from Geofabrik, an online repository for OSM data (Geofabrik, 2018). From this, we were able to extract information related to the roads, buildings, and points of interests present within specified areas.

4. METHODS

In this section, we describe the different methods used in predicting socioeconomic well-being. All models were evaluated using a five-fold nested cross validation scheme.

4.1 Remote Sensing-based Transfer Learning Model

We implemented the satellite-based deep learning approach proposed by Xie et al. in 2015 and later improved upon by Jean et al. in 2016, with the assumption that nighttime lights act as a good proxy for economic activity (Mellander et al., 2015). As in Head et al., we began by fine-tuning a convolutional neural network (CNN) with VGG16 architecture that has been pre-trained on the ImageNet dataset to recognize 1000 different class labels (Krizhevsky et al., 2012), with the goal of learning features that are useful for poverty prediction. We treat the problem as a classification task with five night time intensity classes: low, moderately low, medium, moderately high, and high. We set aside 90% of the images for training and used the remaining 10% for the validation set. We dealt with the class

imbalance by upsampling the minority classes (high, moderately high, and medium nighttime light intensities) and downsampling the low and moderately low light intensity classes to 30,000 images per class in order for all five classes to have the same number of training examples.

Like most models pretrained on ImageNet, the VGG16 model accepts 224×224 pixel images; meanwhile, our input images are 400×400 pixels. We proceeded to implement a fully convolutional architecture as described by Xie et al., which involves replacing the fully-connected top layers from the VGG16 model with randomly initialized fully convolutional top layers. This allows the model to accept input images of arbitrary sizes without losing information. We augmented the training set using random horizontal mirroring and began fine-tuning the full network using an Adam optimizer with an initial learning rate of 10^{-6} and a batch size of 32. We set the maximum number of epochs to 50, decreasing the learning rate by a factor of 10 whenever the validation loss began to plateau. We froze most of the layers, tuning only the last block of fully convolutional layers. At the end of training, we were able to achieve a 72% validation accuracy and 60% validation F1 score for the classification task.

For each image, we extract a 4,096-dimensional vector of activations in the top layer of the CNN, which are optimized to distinguish between different levels of night light luminosity. Each cluster has up to 400 images that we convert to feature vectors of learned representations; these feature vectors are then averaged into a single vector. Finally, we used these cluster-level feature vectors as input to a secondary regression model to predict the socioeconomic indicators. As in previous studies (Jean et al., 2016, Head et al., 2017), we used a ridge regression model to learn the mapping from cluster-level feature vectors to socioeconomic indicators.

4.2 OpenStreetMap Model

For each cluster, we engineered three types of OSM features, namely roads, buildings, and points of interest (POIs). These OSM features were extracted within a 5 km radius for rural areas and 2 km radius for urban areas, with each area centered on the cluster locations. We identified five types of roads in the dataset: primary, trunk, paved, unpaved, and intersection. For road feature engineering, we followed the pre-processing technique described by Zhao and Kusumaputri, i.e., for each type of road, we calculated the distance to the closest road, total number of roads, and total road length per cluster (Zhao and Kusumaputri, 2016).

We also identified six different types of buildings: residential, damaged, commercial, industrial, education, health. For each type, we calculated the total number of buildings, the total area of buildings, the mean area of buildings, and the proportion of the cluster area occupied by the buildings. Finally, we identified over 100 different points of interests; for each cluster, we obtained the total number of each POI within a proximity of the area, e.g., number of banks, bars, cinemas, colleges, hotels, parks, etc.

We compared the performances of random forest regression models trained on the different types of OSM features, both separately and combined, for predicting socioeconomic well-being. Furthermore, we also conducted experiments to determine the predictive performance of models trained using multiple data sources, with the hypothesis that using features from mixed data sources will bolster model performance.

Specifically, we trained random forest regression models using a combination of OSM data and nighttime lights-derived features as input. Nighttime light features consist of summary statistics and histogram-based features, including the mean, median, maximum, minimum, covariance, skewness, and kurtosis, of the nightlight luminosity pixels within each cluster.

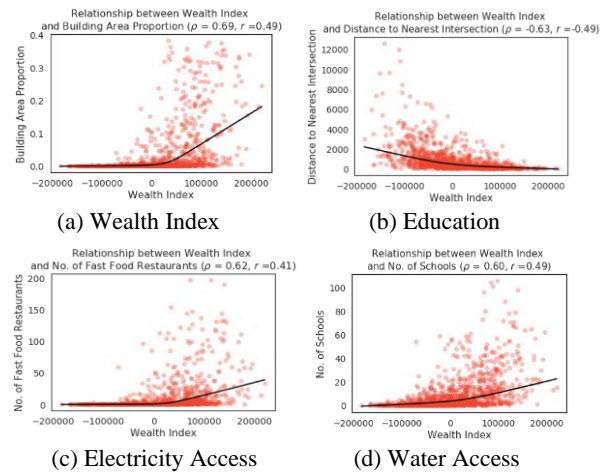


Figure 2. Relationship between cluster-level asset-based wealth and different covariates extracted from OSM. Black line indicates a LOESS curve fitted to the data.


To our knowledge, this is the first paper to study multi-dimensional poverty prediction using a combination of crowd-sourced geospatial data and satellite data in the unique context of a developing nation in Southeast Asia.

5. RESULTS AND DISCUSSION

5.1 Poverty Prediction using Satellite Imagery and Transfer Learning

Past studies have published results on using deep learning methods for predicting wealth in sub-Saharan African countries (Jean et al., 2016) as well as non-African countries (Head et al., 2017). Predictive models achieved r-squared results ranging from 0.51 to 0.75 (Haiti: 0.51; Malawi: 0.55; Tanzania: 0.57; Nepal: 0.64; Nigeria: 0.68; Uganda: 0.69; Rwanda: 0.75). In this study, we tested how well the satellite based deep learning approach performs in the Philippine setting.

Note that the Philippines, being an archipelago that consists of over 7,000 islands, required additional pre-processing steps in order to reduce noise in the dataset. Specifically, we removed satellite images composed mainly of bodies of water as well as images containing no human settlements using the High Resolution Settlement Layer (HRSL) developed by Tiecke et al.; by doing so we were able to see a notable rise in the r squared score from 0.56 to 0.59 for wealth prediction. By increasing the number of nighttime light bins from the 3 to 5 and incorporating binary regional indicators as predictors, we were able to further improve the wealth index r-squared score to 0.63. As proof of concept, we show in Figure 3 a reconstruction of provincial-level poverty maps by aggregating cluster-level wealth estimates.



(a) Ground Truth Wealth Index (b) Predicted Wealth Index

Figure 3. Ground-truth wealth indices and cross-validated Philippine poverty predictions using the satellite-based transfer learning model aggregated to the provincial level.

Our findings also indicate that the method does not generalize for other socioeconomic indicators with the same r-squared as wealth (education: 0.47, access to electricity: 0.27, access to water: 0.10). We note that these results are consistent with the conclusions reached in (Head et al., 2017), which states that high model performance cannot be expected when there is no clear relationship between the development indicator and nighttime lights.

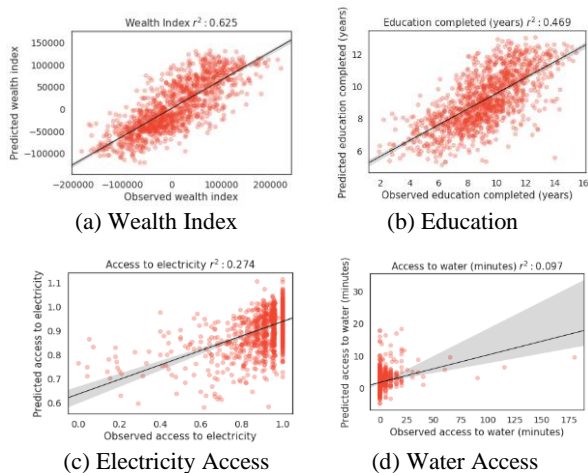



Figure 4. Ground-truth wealth index and cross-validated Philippine poverty predictions using the satellite-based deep learning model.

5.1.1 Visualizing Nighttime Light Classification Model To visualize the nighttime light classification model, we generate class saliency maps based on a given image and class (Simonyan et al., 2013). The image-specific saliency is computed by getting the magnitude of the class score derivative, which indicates the weight of each pixel in affecting the class score. We see in Figure 4 that the model identifies pixels related to roads and buildings as important for classifying medium to high nighttime light intensity; whereas, pixels related to trees and crops are given more weight for low nighttime light intensity classes.



(a) Low Nighttime Light Intensity Class



(b) Medium Nighttime Light Intensity Class



(c) Moderately High Nighttime Light Intensity Class

Figure 5. Class saliency maps for test satellite images in the nighttime light classification task. Maps were generated using a single back-propagation pass through the trained CNN.

5.2 Poverty Prediction using Crowd-sourced Geospatial Information

We trained separate random forest regression models for each type of OSM feature (road, building, or POI). We found that using roads, buildings, or points of interests alone already explain 49-55% of the variance, with roads being the best predictor (R^2 : 0.55). Training a model on all three types of OSM features results in a higher r-squared (0.59). Furthermore, by combining OSM features with nighttime lights data and binary regional indicators, we were able to obtain an r-squared of 0.63 for wealth prediction. The r-squared results for education, electricity access, water access are 0.49, 0.36, and 0.09, respectively. Since our poverty prediction approach was optimized for predicting asset-based wealth, a more indicator-specific feature engineering and feature selection process may likely bolster performance.

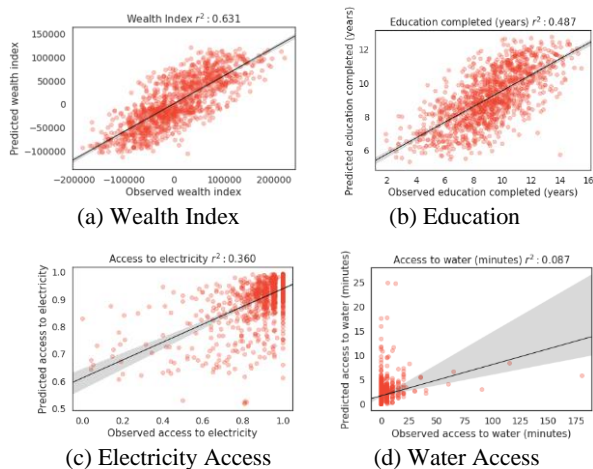


Figure 6. Ground-truth wealth index and cross-validated Philippine poverty predictions using the OpenStreetMap model.

We find that the performance of the the OSM-nightlights hybrid model (R^2 : 0.63) achieves similar results as that of state-of-the-art satellite-based transfer learning model (R^2 : 0.63). However, unlike satellite images from Google Static Maps which are proprietary and limited by licensing terms, both OSM and NTL data are publicly available and freely redistributable, making them an inexpensive alternative to daytime satellite images, which cost approximately 3,000 USD to acquire in order to generate granular poverty maps for the entire Philippines.

We also draw attention to some of the drawbacks of using OSM for poverty prediction. One major disadvantage of using OSM is that it is difficult to assess the accuracy and completeness of the data. The method is also highly dependent on human input, i.e. it requires on-the-ground expert surveyors with domain knowledge to map the area. For hard-to-access regions in far-flung or warstricken areas, obtaining reliable information can be a challenge.

While our initial work does show that despite not being a perfect data source, the OSM method does achieve good results. However, more research needs to be done to confirm the reliability of the method, and we urge researchers, practitioners, and stakeholders to consider these drawbacks, when adapting the method.

6. CONCLUSIONS

In this study, we implemented the satellite-based deep learning approach described by Xie et al. and Jean et al. in the Philippine setting. Our results confirm the applicability of the methodology, with the best model achieving an r-squared of 0.63 for estimating asset-based wealth. We also proposed an alternative cost-effective approach to poverty prediction that uses free and publicly available crowd-sourced geospatial information. Our findings indicate that a model trained on a combination of OSM and NTL-derived features also achieves an R^2 of 0.63. We conclude that both satellite images and volunteered geographic information are valuable tools for high resolution, real-time poverty mapping. Efforts in poverty mapping have great potential to help governments and humanitarian organizations better understand the spatial

distribution of poverty and implement more evidence-based targeted interventions in developing countries.

ACKNOWLEDGEMENTS

We would like to thank Neal Jean, Pia Faustino, Ram De Guzman, Lester James Miranda, and Priscilla Moraes for the insightful discussions and valuable guidance. This work was supported by the UNICEF Innovation Fund.

REFERENCES

- Anderson, J., Sarkar, D. and Palen, L., 2019. Corporate Editors in the Evolving Landscape of OpenStreetMap. *ISPRS International Journal of Geo-Information* 8(5), pp. 232.
- Babenko, B., Hersh, J., Newhouse, D., Ramakrishnan, A. and Swartz, T., 2017. Poverty Mapping Using Convolutional Neural Networks Trained on High and Medium Resolution Satellite Images, With an Application in Mexico. In: 31st Conference on Neural Information Processing Systems (NIPS) Workshop on Machine Learning for the Developing World, NeurIPS Foundation.
- Barrington-Leigh, C. and Millard-Ball, A., 2017. The world's user-generated road map is more than 80% complete. *PLOS One* 12(8), pp. 1–20.
- Blumenstock, J., Cadamuro, G. and On, R., 2015. Predicting poverty and wealth from mobile phone metadata. *Science* 350(6264), pp. 1073–1076.
- Burgert, C. R., Colston, J., Roy, T. and Zachary, B., 2013. DHS Spatial Analysis Reports 7. Technical report, ICF International, Maryland, USA.
- Engstrom, R., Hersh, J. and Newhouse, D., 2017. Poverty from Space: Using High-Resolution Satellite Imagery for Estimating Economic Well-Being. Technical report, World Bank.
- Fernando, L., Surendra, A., Lokanathan, S. and Gomez, T., 2018. Predicting population-level socio-economic characteristics using call detail records (cdrs) in sri lanka. In: Proceedings of the Fourth International Workshop on Data Science for MacroModeling with Financial and Economic Datasets, DSMM'18, ACM, New York, NY, USA, pp. 1:1–1:12.
- Geofabrik, 2018. Philippines - Geofabrik Download Server. <https://download.geofabrik.de/asia/philippines.html>. Online; accessed 18 December 2018.
- Gervasoni, L., Fenet, S., Perrier, R. and Sturm, P., 2018. Convolutional neural networks for disaggregated population mapping using open data. In: 5th IEEE International Conference on Data Science and Advanced Analytics (DSAA), IEEE, pp. 594–603.
- Grippa, T., Georganos, S., Zarougui, S., Bognounou, P., Diboulo, E., Forget, Y., Lennert, M., Vanhuyse, S., Mboga, N. and Wolff, E., 2018. Mapping Urban Land Use at Street Block Level Using OpenStreetMap, Remote Sensing Data, and Spatial Metrics. *ISPRS International Journal of Geo-Information* 7(7), pp. 246.

Head, A., Manguin, M., Tran, N. and Blumenstock, J. E., 2017. Can Human Development be Measured with Satellite Imagery? In: Proceedings of the Ninth International Conference on Information and Communication Technologies and Development, ACM, pp. 8:1–8–11.

Hernandez, M., Hong, L., Frias-Martinez, V. and Frias-Martinez, E., 2017. Estimating poverty using cell phone data: evidence from Guatemala. Technical report, World Bank.

Jean, N., Burke, M., Xie, M., Davis, W. M., Lobell, D. B. and Ermon, S., 2016. Combining satellite imagery and machine learning to predict poverty. *Science* 353(6301), pp. 790–794.

Juan-Albacea, Z. V., 2009. Small Area Estimation of Poverty Statistics. Technical report, Philippine Institute of Development Studies.

Krizhevsky, A., Sutskever, I. and Hinton, G. E., 2012. ImageNet Classification with Deep Convolutional Neural Networks. *Advances in Neural Information Processing Systems* 25 pp. 1097–1105.

Mahabir, R., Croitoru, A., Crooks, A., Agouris, P. and Stefanidis, A., 2018. News coverage, digital activism, and geographical saliency: A case study of refugee camps and volunteered geographical information. *PLOS One* 13(11), pp. 1–28.

Mellander, C., Lobo, J., Stolarick, K. and Matheson, Z., 2015. Night-Time Light Data: A Good Proxy Measure for Economic Activity? *PLoS one* 10(10), pp. 1–18.

NOAA National Centers for Environmental Information, 2016. VIIRS DNB Dataset. https://ngdc.noaa.gov/eog/viirs/download/dnb_composites.html. Online; accessed 18 December 2018.

Philippine Statistics Authority, 2018. Philippines National Demographic and Health Survey 2017. Technical report, PSA, Quezon City, Philippines.

Simonyan, K., Vedaldi, A. and Zisserman, A., 2013. Deep Inside Convolutional Networks: Visualising Image Classification Models and Saliency Maps. *CoRR*.

Tiecke, T. G., Liu, X., Zhang, A., Gros, A., Li, N., Yetman, G., Kilic, T., Murray, S., Blankespoor, B., Prydz, E. B. and Dang, H. H., 2017. Mapping the world population one building at a time. *CoRR*.

Xie, M., Jean, N., Burke, M., Lobell, D. and Ermon, S., 2016. Transfer Learning from Deep Features for Remote Sensing and Poverty Mapping. In: Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, AAAI, pp. 3929–3935.

Zhao, L. and Kusumaputri, P., 2016. OpenStreetMap Road Network Analysis for Poverty Mapping. http://www.ludazhao.com/files/poverty_mapping.pdf.