the review of income and wealth

Review of Income and Wealth
Series 68, Number S2, December 2022
DOI: 10.1111/roiw.12552

USING SATELLITE DATA TO GUIDE URBAN POVERTY REDUCTION

BY THOMAS PAVE SOHNESEN*

Independent Consultant

PETER FISKER

Development Economics Research Group, Department of Economics, University of Copenhagen

AND

DAVID MALMGREN-HANSEN

Independent Researcher

Poverty reduction in low- and middle-income countries is increasingly an urban challenge, and a challenge that continues to be constrained by lack of data, including data on the spatial distribution of poverty within cities. Utilizing existing household survey data in combination with Convolutional Neural Networks (CNN) applied to high-resolution satellite images of cities, this study shows that existing data can generate detailed neighborhood-level maps providing key targeting information for an anti-poverty program. The approach is highly automatic, applicable at scale, and cost-effective. The method also provides direct support for policy development, as illustrated by the case study, where the Government of Mozambique is implementing an urban social safety net program, targeting poor urban neighborhoods, utilizing the estimated poverty maps.

Keywords: poverty, social protection, remote sensing, convolutional neural networks, image recognition

1. INTRODUCTION

As stated in the UN's Sustainable Development Goals, poverty reduction is a key global challenge. Poverty reduction requires sustained efforts and often the combination of multiple policies. In the poorest countries, the challenge is especially difficult as policies are designed with very limited available data. The second-best solution is often to fill the information gap through household surveys. Unfortunately, such surveys are expensive and can therefore only be used to a limited extent. The World Bank estimates that monitoring and tracking poverty in the poorest countries will cost USD 945 million between 2016 and 2030 (Kilic *et al.*, 2017), while they rarely provide current and continuous data. Further, such surveys normally estimate vital statistics at urban/rural and regional levels, while they provide no guidance on planning or differences within specific urban environments.

The rapid expansion of available high-resolution satellite images provides an opportunity to close some of this information and data gap. Satellite images of cities by themselves, as seen by the human eye, are often very informative of the general living situation in any given location. However, human observations are an

*Correspondence to: Thomas Pave Sohnesen, Kløverbladsgade 24, 2500 Valby, Denmark (sohnesen@gmail.com).

insufficient tool for policy development, but satellite images combined with image recognition and machine learning methods do have potential to reduce the data gap and thereby pave the road for more efficient anti-poverty policies, as well as many other applications.

The literature on poverty estimation using satellite images is still scant, but recent progress includes the use of nighttime light data and daytime images to estimate village poverty rates in five African countries (Jean *et al*., 2016). Object identification of buildings and their quality as well as cars, combined with geospatial data on roads and farmland, are used to predict municipal poverty rates in Sri Lanka (Engstrom *et al*., 2017). Similarly, object identification of water source, roof quality and lighting source are used to predict poverty at sub-district levels in Uttar Pradesh, India (Pandey *et al*., 2018). Identification of roof quality, but not poverty itself, also provides the basis of allocation of anti-poverty transfers in some villages in Uganda and Tanzania (Abelson *et al*., 2014). Prediction of municipal poverty rates in Mexico and Philippines, on the other hand, are done by relating poverty rates directly to images without specifying any objects (Orbital Insight, 2017, Mesina *et al*., 2019), while neighborhood poverty rates are predicted by extraction of different image features in Ghana (Engstrom *et al*., 2019). Even single households' poverty status are predicted using remote sensing data (Watmough *et al*., 2019). Beyond poverty, local urban wealth, based on an asset index, has also been modeled based on satellite-derived land-use and cover (Georganos *et al*., 2019) and pockets of deprivation (Wang *et al*., 2019).

This study contributes to the nascent literature of poverty estimation based on satellite images by estimating citywide poverty scores at very small locations (smaller than 115 × 115 m). It is also the first study to compare the direct approach, relating poverty directly to images, and the indirect approach of predicting poverty based on identified objects correlated with poverty. Importantly, unlike the estimates at household level (Watmough *et al*., 2019), that rely on manually measured structure footprints as a predictor, this study employs methods that work automatically and can be applied at large scale, utilizing data available in most countries.

Continuous urbanization amplifies the need for new and detailed urban data. For instance, in Mozambique, the urban population grew by more than 50 percent during the last decade, and is expected to grow roughly another 50 percent every decade for the next three decades (United Nations 2018). This trend is similar to many other developing countries, leading to a general urbanization of poverty (Ravallion *et al*., 2007). As in many other countries, there is no or very limited detailed data on where and how this urban growth is changing the city landscapes in-between the decennial population censuses. In many cases, the extent of cities and the relative living standards across neighborhoods are not recorded in systematic ways. Such shortages impede efficient social policies targeting those most in need, as well as other public policies. Standard household surveys, that are used for collection of many other vital statistics, do not cover cities in sufficient detail to provide much guidance on planning, and for most areas there is very limited administrative data available.

To address urban poverty, the Government of Mozambique is expanding an incipient urban social safety net program. In the Productive Safety Net Program (PASP), beneficiaries receive a small subsidy for supplying their labor to the production of a public good for the benefit of the local community. Hence, the program both aims to improve the livelihood of poor households as well as improve poor neighborhoods. As in many other African countries, the program targets households and neighborhoods by first selecting poor neighborhoods and then secondly selecting households for participation. Households are eligible to join the program based on a number of observable characteristics that result in a poverty score, also known as Proxy Means Test (PMT) score. Households with a poverty score below a certain threshold are eligible to participate in the program. However, currently the program has no mechanism to efficiently target pockets of poverty within cities. In order to reduce poverty most efficiently, the government wishes to rank neighborhoods according to average poverty scores. Poverty scores' correlation to consumption expenditures underlying households' poverty status is not without critics (Brown *et al.*, 2018, Gazeaud, 2020); however, for program targeting the poverty scores based on the PMT model is the key eligibility criteria. See also recent reflections on how to design urban social protection programs (Cuesta *et al.*, 2021).

2. Data and Methods

To estimate the neighborhood- and sub-neighborhood level poverty scores, two different methods are tested:

- *Method 1* relates households' poverty scores to satellite images centered at the households' location using a Convolutional Neural Network (CNN).
- *Method 2* relates households' poverty scores to other variables using a random forest model. The data in each location include various geo-spatial data as well as the outputs from a CNN object detection, detecting density and quality of structures.

2.1. *Data*

Unit of Analysis

All predictions and most of the analysis are based on a grid of 115 × 115 m cells.[1] The size of cells was chosen to match the resolution of image sized 400 × 400 pixels. After filtering out areas with no residential buildings, like airports, military areas and industrial zones, the sample of interest has 57,540 cells covering five cities in Mozambique. The results at cell level are also aggregated to bairro (neighborhood) level. A bairro is the lowest level of government in urban areas in Mozambique and the PASP program operates through these. In the analysis there are 180 bairros in the five cities, with a combined population of more than 4.5 million people. Figures five though seven illustrate the size and distribution of bairros across the five cities.

[1]The size of cells, measured in meters at surface level, varies with a negligible amount due to the curvature of the earth between each of the cities.

Poverty Score (PMT) from Survey Data

The government uses a poverty score to rank households' eligibility for the program, based on a PMT model. To capture households' poverty scores this model is applied to the latest nationally representative household survey (Inquérito sobre Orcamento Familiar 2014/2015) as well as observations from the first round of interviews for the social protection program. In both sources, the poverty score is calculated from survey questions on ownership of durable assets and other household characteristics and both data sources include households' locations measured by latitude and longitude. The PMT model proxy households' log consumption expenditure level (model can be seen in supplemental appendix), which is also the base for monetary poverty in Mozambique. The two surveys have 15,033 urban households in total, covering 3,927 cells.

Geospatial Data

The following spatial data are utilized: (a) *Road data from Open Street Map (OSM)*. OSM is a volunteer-driven platform for creation of open access map data (OpenStreetMap contributors). This data includes the total length of—and distance to—primary, residential and other roads at the cell level. The quality of this data depends on the voluntary contributors, and can therefore vary and is hard to judge. However, the coverage and quality were found to be much higher than the alternative sources. (b) *Distance to city center*. To capture remoteness, the distance from the center point of each cell to the city center (as birds fly) was calculated using GIS software. (c) *Estimates of structure density and quality*. Based on tagged images of building structures, a Convolutional Neural Network produced estimates of structure density and quality of buildings for each cell across cities in Mozambique (Fisker *et al.*, 2020). It's a two part CNN, that first classifies if a structure is present or not and secondly classifies observed structures into colored roof (associated with higher quality), grey roofs (associated with low quality tin roofs), non-residential structure, and structure under construction.

Images

Satellite images of each cell were downloaded under Google's "Fair use" agreement, using Google Maps' API in November 2018. Each image covers one *unit of analysis* aligning estimation and predictions into to the cell grid.

2.2. *Methods*

To estimate the cell-level average poverty scores, two different methods are tested. For both methods, the full sample is 3,927 cells of which 10 percent (392 cells) of data is excluded for evaluation, leaving 3,535 cells in the training data set. 14,252 out of the 15,033 households are found within the 3,535 cells used for training and models. The excluded sample is stratified over cities, with 10 percent of cells randomly selected within each city.

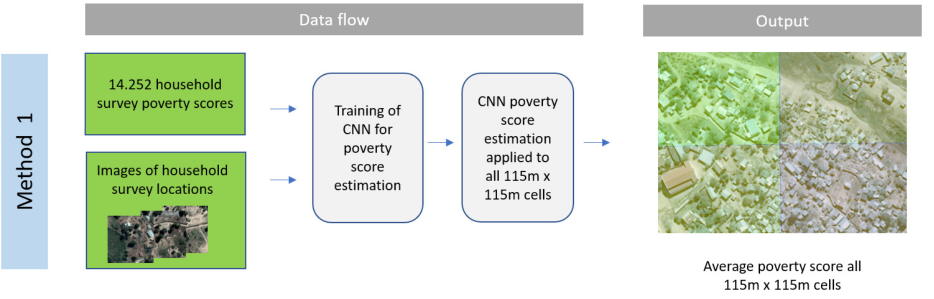

Figure 1. Figure Shows the Data Flow of Method 1
Note: The target variable for training of the CNN is the household poverty score. [Colour figure can be viewed at wileyonlinelibrary.com]

Method 1: Convolutional Neural Network

The first method directly combines households' poverty scores and images of all households' locations using a CNN regression model (Figure 1). A first attempt at training the CNN based on the available sample for households in Mozambique was unsuccessful. A second attempt was made using transfer learning. The CNN architecture Xception (Chollet, 2016), which is pre-trained to classify images in the ImageNet (Russakovsky *et al.*, 2014) dataset, was used as the basis for training. In the training, the top layer was excluded and replaced by a Global Average Pooling operation and a single output fully connected layer with no activation function. Further, in order to increase the amount of training data, five crops of each image were used, top-left, top-right, bottom-left, bottom-right and center, and each image was flipped, and a random scaling of its' brightness level of ±15 percent was applied. The CNN-regressor was trained with the ADAM weight update scheme with parameters as recommended in Kingma and Ba (2017) for 20 epochs. A learning rate decay was also applied reducing the learning rate with 2 percent for each epoch. The CNN regressor is implemented in Python using the Keras library.

Method 2: A Random Forest Model

The second method predicts poverty scores at cell level using a random forest model that has the poverty score as the target variable and the available geospatial data; structure density and quality, the distance to city center, and length of/distance to roads as predictors (Figure 2). Capturing that not only the immediate surroundings can be important for the poverty score, spatial lags of all variables are included with weight given up to 500 m from each neighborhood center point. The predictions were implemented in the Python anaconda package using the RandomForestRegressor module from sklearn.ensemble. Mean square errors was used as decision criteria, minimum leaf size was set at 10, and predictions are based on 150 trees. Application of random forest is recent and still scant for poverty predictions, though evaluation of the method compared to alternatives have been found favorable (McBride and Nichols, 2017; Sohnesen and Stender, 2017; Rusnita *et al.*, 2019).

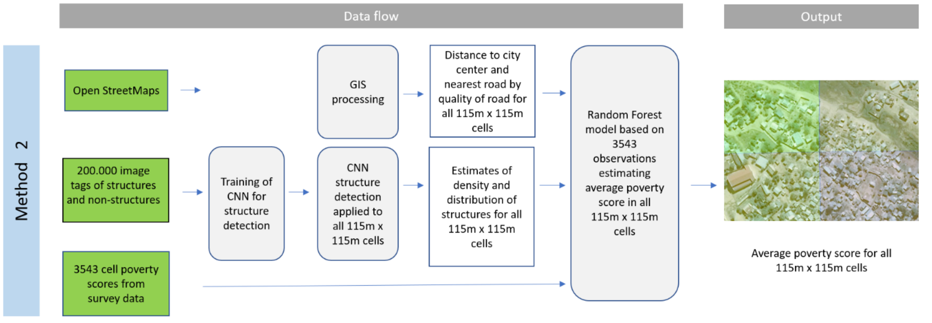

Figure 2. Figure Shows the Data Flow of Method 2
Note: The target variable for training of the CNN is the presence and quality of structures. [Colour figure can be viewed at wileyonlinelibrary.com]

The first method is the least data demanding method, as it relies on household survey data with households' poverty status and GPS location only. Such data is available in a large majority of developing countries. This method could also have difficulties as two similar in appearance locations could have different associations to poverty. Imagine two similar locations, one close to downtown, another far from the city center. Such two locations are likely to have different associations to poverty, which is unknown to the CNN. The second method utilize data from several sources, and is more demanding on required data.

3. Results

Using the direct approach in Method 1 (Figure 1) does provide insights into spatial distribution of the poverty score. The out-of-sample Spearman rank coefficient and r-square are 0.62 and 0.38, respectively (Figure 3). Though not a great prediction model, this can be seen as a breakthrough for the approach. To the authors' knowledge, such direct estimation of poverty, based on images and survey data only, has only been attempted by two previous studies (Orbital Insight, 2017, Mesina *et al.*, 2019). Mesina *et al.*'s attempt at estimating poverty rates at sub municipal level was deemed too poor for policy application and Orbital Insight (2017) concluded that *"the algorithm did not learn a meaningful representation of poverty."*

Poverty score predictions based on Method 2 (Figure 2) are notably more accurate than those based on Method 1 (out-of-sample r square of 0.58 compared to 0.38) (Figure 4). The model predictions reveal that the city variable for Maputo has the highest importance score, next is distance to city center, and the share of high-quality roofs. In general, the estimates of structure density and quality of structures (and their spatial lags capturing the wider area) are prominent among the geospatial variables most important for poverty score predictions, while the road data is less prominent (Figure 3). The geospatial model (Figure 2) highlights that information unavailable to Method 1, like city and distance to city center, are important for predictions of poverty scores. The dominance of the dummy for

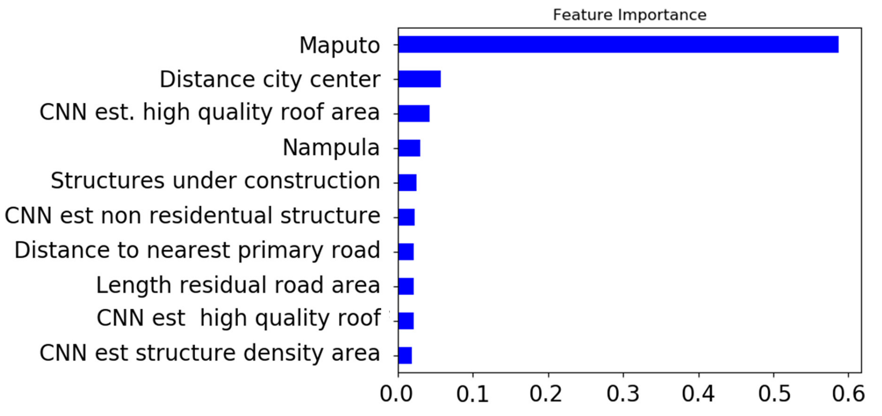

Figure 3. Importance Scores for Method 2 [Colour figure can be viewed at wileyonlinelibrary.com]

Maputo city in Importance score could indicate that the high r square is disproportionally driven by the large difference between the capital Maputo and other cities. However, replicating Method 2, but excluding the city dummies give almost identical results (out-of-sample r square of 0.574 compared to 0.576 when including the city variable) showing the approach picks up much more than just differences in averages across cities.

Operationally, the PASP program is implemented through *bairros* (the Portuguese word for neighborhood), a low-level administrative unit within cities in Mozambique. A bairro poverty score can be obtained by averaging cell estimates for within each bairro. Evaluating model accuracy at bairro level shows a remarkable out-of-sample r-square of 0.58 for Method 1 and 0.66 for Method 2, and spearmen ranking coefficients of 0.76 and 0.83, respectively (Figure 3). Figure 3 illustrates the correlation between predicted and observed poverty scores at cell and bairro level, including Method 2's better performance.

To the authors' knowledge there are currently no other studies that make estimates as such a granular level as this study, but Jean *et al.* (2016)'s predictions at cluster level are reasonably close. Their cross validated predictions explain between 37 percent and 59 percent of the variation in various countries.

Of general interest is the degree to which models are transferable, *i.e.* whether a model trained on one city can be applied in a different city. If not, the model would not be applicable to cities not included in the household surveys. Testing this, by excluding a full city from the training data and predicting the poverty score for the excluded city, utilizing Method 2, reveals that the models are not transferable. Predictions for the excluded city are poor ranging in R square from −1 to 0.04, with spearman rank correlations from 0.03 to 0.09. This is consistent with the city variables as well as the distance to city center variable being prominent among most important variables (Figure 3), and indicates that each city has its own prediction model. For utilization on the urban social program in Mozambique, this is not a problem as the household survey covers all major cities and the model accuracy is similar in all cities. However, it could be a limitation in other applications.

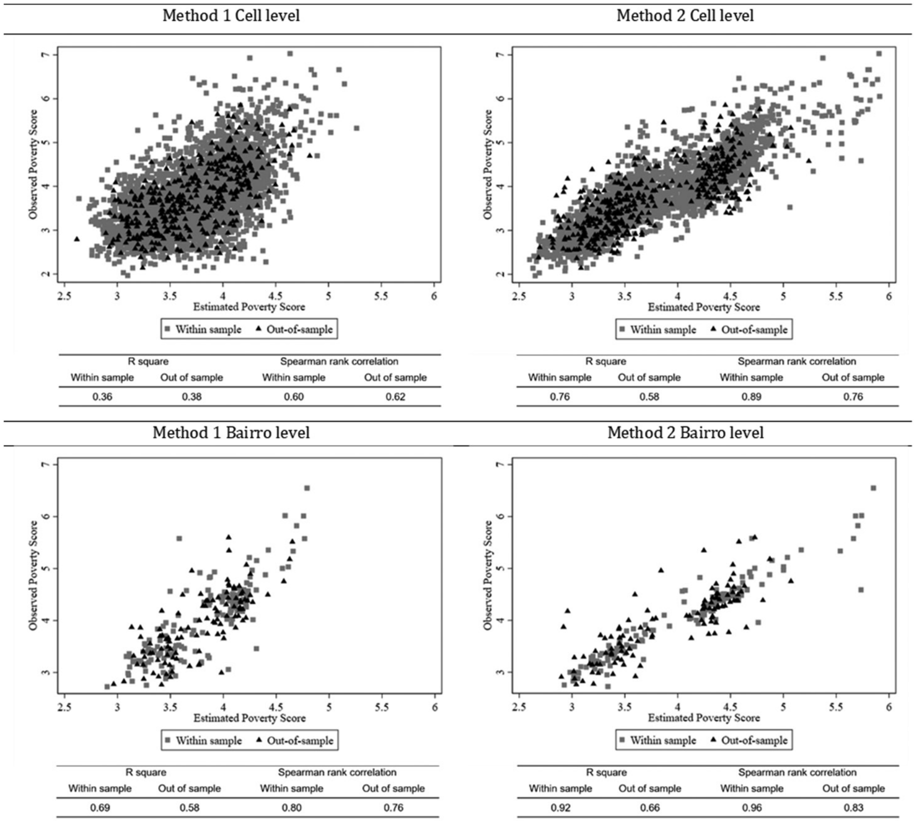

Method 1 Cell level

R square		Spearman rank correlation	
Within sample	Out of sample	Within sample	Out of sample
0.36	0.38	0.60	0.62

Method 2 Cell level

R square		Spearman rank correlation	
Within sample	Out of sample	Within sample	Out of sample
0.76	0.58	0.89	0.76

Method 1 Bairro level

R square		Spearman rank correlation	
Within sample	Out of sample	Within sample	Out of sample
0.69	0.58	0.80	0.76

Method 2 Bairro level

R square		Spearman rank correlation	
Within sample	Out of sample	Within sample	Out of sample
0.92	0.66	0.96	0.83

Figure 4. Images and Supporting Tables Show Accuracy of Predictions from Method 1 and 2 as Defined in Figures 1 and 2

Notes: R^2 is explained variance and is based on the 115 × 115 m neighborhood. Bairro estimates are averages of cell estimates. The random forest model in Method 2 is based on 150 trees and has an out-of-bag r^2 of 0.58.

Other studies have found model performing reasonably well across areas. In Jean *et al*. (2016) models trained on three countries, predicts a fourth country reasonable well. Similarly, Engstrom *et al*. (2017) find that the rank correlation for a random forest model predicting subdistrict areas, that are dropped from the model sample, results in spearman's ρ estimated at between 0.74 and 0.76. Related, both Orbital Insight (2017) and Engstrom *et al*. (2017) find limited impact from reducing the number of observations within same locations, indicating that using relative small samples from each location is sufficient for model building.

Illustrating how this approach provides support for the Government of Mozambique, Figure 5 shows predicted poverty scores at the cell level for the city of Maputo. The highlighted extract shows an overlay of predictions in Google Earth, illustrating the differences between relatively poor and affluent areas. It's notable that the method is capable of categorizing cells by average poverty score, even when in close spatial proximity. Figure 6 further illustrates the difference in

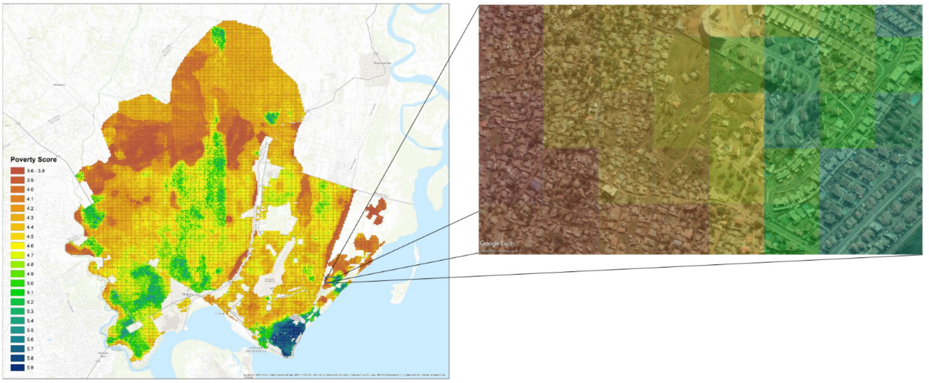

Figure 5. Left Image Shows Poverty Scores in the City Maputo Based on Predictions from Model 2 Colored from Red, Poor, to Blue, Rich

Note: Right image shows enlargement of section of city, illustrating the poverty score predictions for cells (115 × 115 m). [Colour figure can be viewed at wileyonlinelibrary.com]

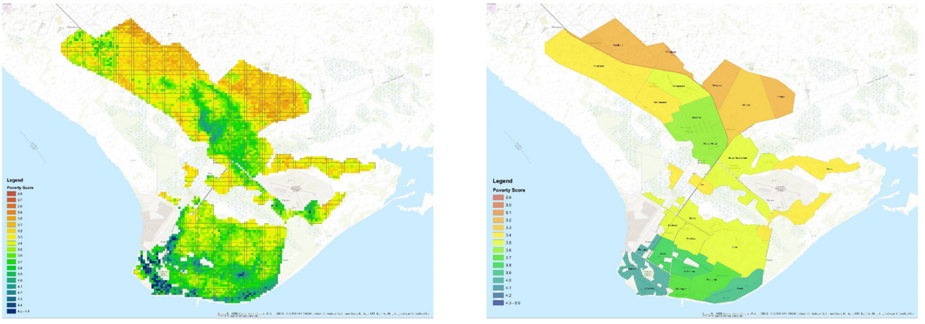

Figure 6. Estimated Poverty Scores at Cell and Bairro Level for the City of Beira [Colour figure can be viewed at wileyonlinelibrary.com]

resolution, when targeting bairros as opposed to cells, while Figure 7 illustrates these results in all five cities covered by the PASP social protection program.

As mentioned, the PASP program has already been initiated without spatial data on distribution of poverty scores neither between bairros nor within them. Bairros have been selected based on local knowledge, and households have been selected for interviews based on recommendations by community councils in those bairros. Utilizing the predicted average poverty scores, one can evaluate if selected neighborhoods as well as the cells of the selected households are indeed among the poorest. Figure 8 shows that selected cells are indeed among the poorest areas, indicating that local knowledge is useful for targeting (left image). However, Figure 8 also shows that local knowledge selects poor areas, but fails to target the very poorest cells (right image, Figure 8).

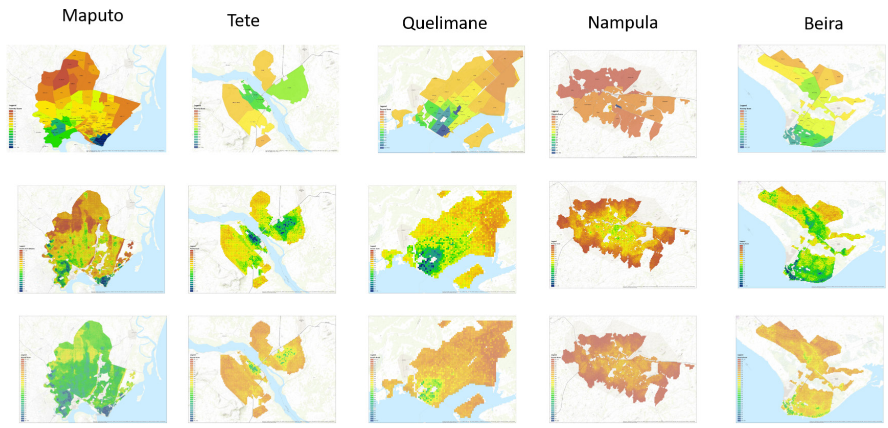

Figure 7. Poverty Scores for Bairros and Cells in the Five Cities Being a Part of PASP [Colour figure can be viewed at wileyonlinelibrary.com]

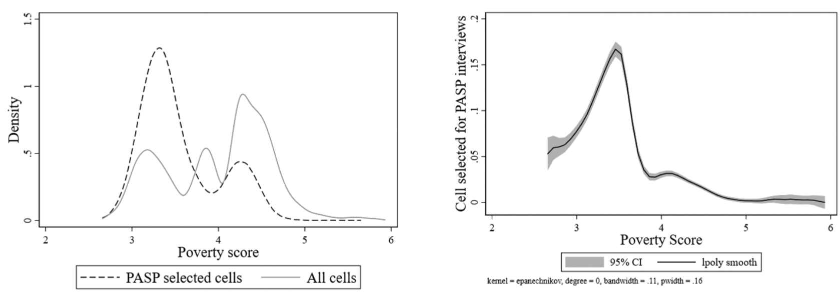

Figure 8. Left Image Shows Distribution of the Cell Average Poverty Scores for those Cells that Had Household Interviews for the PASP Program and the Distribution of All Cells

Notes: The graph is generated as a kernel density plot. Right image shows the selected cells for PASP interviews as a share of the cells across the distribution of cell poverty scores. The graph is generated using local polynomial smoothing.

4. Discussion

The core results are both a significant step forward in application of image recognition in urban settings and in use of machine learning methods, showing that they can provide key support to countries' poverty reducing operations. However, the method, application, and data are not without challenges, as highlighted in this section.

Survey data has measurement errors and can be susceptible to systematic variation in reporting over time due to variations in design and implementation (Kilic and Sohnesen, 2019) while data extracted from optical satellite images is susceptible to biases from cloud cover and seasonality, as well as changes in satellite technology and post data processing. The nighttime light data was one of the first satellite-based indicators to be used in economics and short comings in this data has been highlighted (Gibson *et al.*, 2020). Another common application of satellite data in economics is drought monitoring. Here, various challenges have also been highlighted in regards to the raw data itself, but even more so on how such data is processed into a meaningful economic indicator (Bachmair *et al.*, 2016; Sohnesen, 2020).

The data source in this application also have challenges, some of which some are surmountable should other researchers undertake similar studies. The images utilized in this study are available for free for research, but also come with a number of limitations. First, the exact source of the image is not shared by Google Maps. This means that we do not know neither the exact date of the image, nor do we know which high resolution satellite the image comes from. Second, Google updates the images, but it is unknown when this happens and neighboring locations can have images from different points in time. Third, when downloading data from the API there is no choice of data, only what is available that day. Hence, it is not possible to discern in detail the origin of the data. Google selects images to upload that have little cloud cover, but we have no knowledge of the data generating process behind this selection. The groundbreaking application by Jean *et al.* (2016) used the same data source with same shortcomings and had the additional challenge that the true location of their clusters were added up to 10 km random noise to the exact location to protect privacy of the respondents. Many of these challenges are surmountable. For instance, one can purchase the images from the satellite image vendors, which will provide control over most aspects of the data generating process. The very first steps of the data generating process, including the technology onboard the satellite, still rests with the satellite company.

The frequency and quality of satellite images are increasing rapidly, leading to a large potential for continuous updates of the poverty scores maps. Applying the same model based on new images, implicitly assumes that the same model is applicable over time. Though critical to the possibility of tracking SDGs and other data relevant for social targeting and planning this assumption has never, to the authors' knowledge, been tested empirically. There is some experience on model stability across time from prediction of poverty based on household characteristics across surveys. Such models rely on richer data including household asset ownership, location, family composition and other characteristics. The experiences are mixed with some successful applications as well as failures, with little clarity on why some succeed while others fail (Sohnesen and Stender 2017; Dang *et al.*, 2019).

The study also suffers from a time gap in the training data between the household observations and the date of the image. Households could have moved or made changes to the structure that would lead to measurement errors. Further, the length of the time gap itself is unknown as the exact date of the image from Google Maps is unknown, as is the origin of the image. Also, there are measurement errors in both the OSM data and likely in some of the GPS data. A few cases were identified, where more than one household interview was identified at the same GPS location. In this case it is assumed that several interviews took place outside at the local community gathering point, as opposed to at households' own house, which they were supposed to. The technical accuracy of GPS measurement can also be a challenge in linking poverty scores to the right household (REF), though this study did have access to the exact GPS data, and not a scrambled version often available in published data. As with most types of measurement errors, it is not possible to say much about the size or any systematic correlations of these uncertainties in the data. The measurement errors will be a part of the models and contribute to prediction errors. In light of this, the out of sample R^2 of 0.58 for cells and 0.66 for bairros is good news, as there is room for further improvement with better data.

5. Conclusion

This study contributes to the Innovations in Poverty Measurement by estimating citywide poverty scores (an estimate of log consumption) at a very detailed resolution (about 115 × 115 m), that provides direct policy support for the implementation of a social protection program aimed at poverty reduction. It is achieved through use of survey and GIS data, optical satellite images, as well as methods commonly applied in data science and in the field of image recognition. Survey data, and partly GIS data, are commonly used in various aspects of poverty analysis, while the use of optical satellite images analyzed through image recognition, and the use of prediction methods from data science are not. We find that pulling in the new data sources and applying new methods have allowed us to contribute with new and innovative findings that are useful for implementing poverty reducing policies. The partnership between authors with backgrounds in different fields of research was necessary, but also contributed to a better and broader understanding.

References

Abelson, B., K. R. Varshney, and J. Sun, "Targeting Direct Cash Transfers to the Extremely Poor," Paper presented at the Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, New York, 2014.

Bachmair, S., C. Svensson, J. Hannaford, L. J. Barker, and K. Stahl, "A Quantitative Analysis to Objectively Appraise Drought Indicators and Model Drought Impacts," *Hydrol Earth Syst Sci Copernicus Publications*, 20, 2589–609, 2016.

Brown, C., M. Ravallion, and D. van de Walle, "A Poor Means Test? Econometric Targeting in Africa," *Journal of Development Economics*, 134, 109–24, 2018.

Chollet, F., "Xception: Deep Learning with Depthwise Separable Convolutions," arXiv:1610.02357, 2016.

Cuesta, J., S. Devereux, A. G. Abdulai, J. Gupte, L. Ragno, K. Roelen, R. Sabates-Wheeler, and T. Spadafora, "Urban Social Assistance. Evidence, Challenges, and the Way Forward, with Application to Ghana," *Development Policy Review*, 39, 360–80, 2021.

Dang, H.-A., D. Jolliffe, and C. Carletto, "Data Gaps, Data Incomparability, and Data Imputation: A Review of Poverty Measurement Methods for Data Scarce Environments," *Journal of Economic Surveys*, 33, 757–97, 2019.

Engstrom, R., J. S. Hersh, and D. L. Newhouse, "Poverty from Space: Using High-Resolution Satellite Imagery for Estimating Economic Well-Being," World Bank Policy Research Working Paper 8284, 2017.

Engstrom, R., D. Pavelesku, and T. Tanaka, "A Mapping Poverty and Slums Using Multiple Methodologies in Accra, Ghana," *2019 Joint Urban Remote Sensing Event (JURSE)*, 2019.

Fisker, P., D. Malmgren-Hansen, and T. P. Sohnesen, "TASIM: Tagged Structures in Mozambique," *Mendeley Data*, 2020.

Gazeaud, J., "Proxy Means Testing Vulnerability to Measurement Errors?," *Journal of Development Studies*, 56, 2113–33, 2020.

Georganos, S., A. N. Gadiaga, C. Linard, T. Grippa, S. Vanhuysse, N. Mboga, E. Wollf, S. Dujardin, and M. Lennert, "Modelling the Wealth Index of Demographic and Health Surveys within Cities Using Very High-Resolution Remotely Sensed Information," *Remote Sensing*, 11(21), 2543, 2019.

Gibson, J., S. Olivia, and G. Boe-Gibson, "Night Lights in Economics: Sources and Uses," *Journal of Economic Surveys*, 34, 955–80, 2020.
Jean, N., M. Burke, M. Xie, W. M. Davis, D. B. Lobell, and S. Ermon, "Combining Satellite Imagery and Machine Learning to Predict Poverty," *Science*, 353, 790–94, 2016.
Kilic, T., U. Serajuddin, H. Uematsu, and N. Yoshida, "Costing Household Surveys for Monitoring Progress Toward Ending Extreme Poverty and Boosting Shared Prosperity," Policy Research Working Paper 7951, 2017.
Kilic, T. and T. P. Sohnesen, "Same Question But Different Answer: Experimental Evidence on Questionnaire Design's Impact on Poverty Measured by Proxies," *Review of Income and Wealth*, 65, 144–65, 2019.
Kingma, D. P. and J. Ba, *Adam: A Method for Stochastic Optimization*, arXiv:1412.6980, 2017.
McBride, L. and A. Nichols, "Retooling Poverty Targeting Using Out-of-Sample Validation and Machine Learning," *World Bank Economic Review*, 32, 531–50, 2017.
Mesina, J., J. L. Isanan, and C. V. Maderazo, "Poverty Incidence Identification of Cities and Municipalities Using Convolutional Neural Network as Applied to Satellite Imagery," *IOP Conference Series: Materials Science and Engineering*, 482, 012044, 2019.
OpenStreetMap contributors, *Planet Dump*, https://planet.osm.org, December 2018.
Orbital Insight, *World Bank: Poverty Mapping in Mexico*, Unpublished Manuscript, 2017.
Pandey, S. M., T. Agarwal, and N. C. Krishnan, "Multi-Task Deep Learning for Predicting Poverty from Satellite Images," Paper presented at the Thirtieth AAAI Conference on Innovative Applications of Artificial Intelligence, 2018.
Ravallion, M., S. Chen, and P. Sangraula, "New Evidence on the Urbanization of Global Poverty," *Population and Development Review*, 33, 667–701, 2007.
Rusnita, I., A. B. Azuraliza, and N. S. Sani, "Does Artificial Inteligence Prevail in Poverty Measuerement?," *Journal of Physics: Conference Series*, 1529, 042082, 2019.
Russakovsky, O., J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei, "ImageNet Large Scale Visual Recognition Challenge," *International Journal of Computer Vision*, 115(3), 211–52, 2015.
Sohnesen, T. P., "Two Sides to Same Drought: Measurement and Impact of Ethiopia's 2015 Historical Drought," *Economics of Disasters and Climate Change*, 4, 83–101, 2020.
Sohnesen, T. P. and N. Stender, "Is Random Forest a Superior Methodology for Predicting Poverty? An Empirical Assessment," *Poverty & Public Policy*, 9, 118–33, 2017.
United Nations DoEaSAPD, *World Urbanization Prospects: The 2018 Revision*. Online Edition, 2018.
Wang, J., M. Kuffer, D. Roy, and K. Pfeffer, "Deprivation Pockets Through the Lens of Convolutional Neural Networks," *Remote Sensing of Environment*, 234, 111448, 2019.
Watmough, G. R., C. L. J. Marcinko, C. Sullivan, K. Tschirhart, P. K. Mutuo, C. A. Palm, and J. C. Svenning, "Socioecologically Informed Use of Remote Sensing Data to Predict Rural Household Poverty," *Proceedings of the National Academy of Sciences of the United States of America*, 116(4), 1213–8, 2019.

Supporting Information

Additional supporting information may be found in the online version of this article at the publisher's web site:

Table S2. PMT model used for poverty scores for the urban social protection program