



SoDa LABS

Social [science] insights from
alternative data

Est. 2018



Object Recognition for Economic Development from Daytime Satellite Imagery

Klaus Ackermann, Alexey Chernikov, Nandini Anantharama, Miethy Zaman,
Paul A Raschky

SoDa Laboratories Working Paper Series

No. 2020-02

REF

Klaus Ackermann, Alexey Chernikov, Nandini Anantharama, Miethy Zaman, Paul A Raschky (2020), SoDa Laboratories Working Paper Series No. 2020-02, Monash Business School, available at <http://soda-wps.s3-website-ap-southeast-2.amazonaws.com/RePEc/ajr/sodwps/2020-02.pdf>

PUBLISHED ONLINE

11 September 2020

© The authors listed. All rights reserved. No part of this paper may be reproduced in any form, or stored in a retrieval system, without the prior written permission of the author.

SoDa Laboratories

<https://www.monash.edu/business/soda-labs/>



MONASH
BUSINESS
SCHOOL



Object Recognition for Economic Development from Daytime Satellite Imagery

Klaus Ackermann,^{1,2*} Alexey Chernikov,^{1,3*} Nandini Anantharama,^{1,3}
Miethy Zaman,¹ Paul A Raschky^{1,4}

¹ SoDa Labs, Monash Business School, Monash University

² Department of Econometrics and Business Statistics, Monash University

³ Faculty of Information Technology, Monash University

⁴ Department of Economics, Monash University

{klaus.ackermann, alexey.chernikov, nandini.anantharama, miethy.zaman, paul.raschky}@monash.edu

Abstract

Reliable data about the stock of physical capital and infrastructure in developing countries is typically very scarce. This is particular a problem for data at the subnational level where existing data is often outdated, not consistently measured or coverage is incomplete. Traditional data collection methods are time and labor-intensive costly which often prohibits developing countries from collecting this type of data. This paper proposes a novel method to extract infrastructure features from high-resolution satellite images. We collected high-resolution satellite images for 5 million 1km × 1km grid cells covering 21 African countries. We contribute to the growing body of literature in this area by training our machine learning algorithm on ground-truth data. We show that our approach strongly improves the predictive accuracy. Our methodology can build the foundation to then predict subnational indicators of economic development for areas where this data is either missing or unreliable.

Introduction

The efficient allocation of limited governmental funds from local governments as well as international aid organizations crucially depends on reliable information about the level of socioeconomic indicators. These indicators (e.g. income, education, physical infrastructures, social class etc.) are critical inputs for addressing the socioeconomic issues for researchers and policy-makers alike. Although data availability and quality for the developing countries has been improving in recent years, consistently measured and reliable data is still relatively scarce. Numerous studies have documented specifically the problems of aggregate economic accounts, in particular to Africa, where the data suffers from various conceptual problems, measurement biases, and other errors (e.g. Chen and Nordhaus 2011; Johnson et al. 2013; Jerven and Johnston 2015).

Researchers have probed into alternative options in the absence of reliable official statistics. Among this newer generation of alternative economic data research, a burgeoning literature has emerged that uses satellite imagery of nighttime luminosity as a proxy for economic activity. Work by Sutton and Costanza (2002), Elvidge et al. (2009), Chen and Nordhaus (2011), Henderson, Storeygard, and Weil (2012), Sutton, Elvidge, and Tilottama (2007) and

Hodler and Raschky (2014) documents a strong relationship between nighttime luminosity and gross domestic product (GDP) at the national and subnational levels. This allows researchers to generate information for any levels of regional analysis and also the likelihood of strategic, human manipulation is limited with satellite generated data. However, luminosity data as a proxy for economic activity is not free from concerns. Satellite sensors have a lower detection bound and nighttime light emissions below this bound are not captured by the satellites' readings. This leads to bottom-coding problem and this is particularly an issue in low-output and low-density regions (Chen and Nordhaus 2011), which are very often regions and countries (e.g. Africa) where official macroeconomic data is missing or unreliable as well.

Over the past few decades, some parts of the African continent have witnessed large increases in economic development. Nevertheless, the majority of regions within African nations still lacks behind. The continent faces further challenges due to localized conflicts (Berman et al. 2017), rapid urbanization (Moyo et al. 2020) as well as the impacts of the COVID-19 pandemic (Treger et al. 2020), among others. A key pre-requisite in formulating adequate strategies to address these challenges, is reliable socioeconomic data at a spatially, granular level. As of now, even data about basic infrastructure such as roads and buildings is not consistently collected across the African continent.

The purpose of this project is to overcome this data problem, by applying machine learning and artificial intelligence tool to a vast amount of unstructured data from daytime satellite imagery. Ultimately, this project aims to go beyond the use of nighttime luminosity as a proxy for economic development data and use high resolution, daytime satellite imagery to predict key infrastructure variables at national and subnational levels for less developed countries like in Africa. Daytime images contain more information about the landscape that is correlated with economic activity, but the images are highly complex and unstructured, making the extraction of meaningful information from them rather difficult. Our approach builds upon and further expands the work of (Jean et al. 2016a). The standard approach in the literature is to learn a representation out of satellite images, that allow an interpretation of pixel activation that are important for predicting night time light or other target. This represen-

*Contributed equally. Work in progress.

tation is then used to predict an aggregated wealth index. Instead, we directly predict infrastructure measures on the ground, albeit knowing that there is a wide spread scarcity of ground truth data.

Motivation and Related Works

Existing solutions for policy makers in developing countries often rely from traditional data gathering processes (i.e. surveys), which are costly and infrequent. Given the high costs, this data does not cover an entire country but only a sub-sample of geographic units. Our solution provides a low-cost method to collect valuable insights about economic development for every location in a country. Our methodology provides relevant decision makers in developing countries as well as NGOs and international organizations with very accurate counts of buildings and the length of roads for an entire country and continent. For example, accurate building counts and density can be used in natural hazard preparedness tools as an indicator for an area’s vulnerability against natural disasters. Information about roads and settlement helps infrastructure agencies to quickly identify areas that lack market access, a key determinant for economic growth in developing countries.

Although relatively new, recent studies have begun to use different daytime satellite images to conduct novel economic research (Donaldson and Storeygard 2016). Daytime images contain more information than night-time images and are thus a good alternative data source for empirical economics. Marx, Stoker, and Suri (2013) used daytime images to analyse the effects of investment on housing quality in the slums of Kibera, Kenya. Investment was calculated based on the age of a household’s roof. The results showed that ethnicity plays an important role in determining investment in housing and belonging to the same tribe as that of the local chief has a positive effect on household investment. (Engstrom et al. 2017) used daytime satellite imagery and survey data to estimate the poverty rates of 3,500 km² subnational areas in Sri Lanka. Using a convolutional neural networks algorithm, they identify object features from raw images that were predictive of poverty estimates. The features examined by the study included built-up areas (buildings), cars, roof types, roads, railroads and different types of agriculture. The results showed that built-up areas, roads and roofing materials had strong effects on poverty rates. A suite of related work has used satellite images to predict population density (e.g. Simonyan and Zisserman 2015; Doupe et al. 2016), urban sprawl (Burchfield et al. 2006), urban markets (Baragwanath et al. 2019) electricity usage (Robinson, Hohmans, and Dilkina 2017), as well as income levels (Pandey, Agarwal, and Krishna 2018). More broadly, we also relate growing body of literature that uses other passively collected data to measure local economic activity (e.g. Abelson, Varshney, and Sun 2014; Blumenstock, Cadamuro, and On 2015; Chen and Nordhaus 2011; Henderson, Storeygard, and Weil 2012; Hodler and Raschky 2014). Methodologically, our paper contributes to the large remote-sensing literature that applies high-dimensional techniques to extract features from satellite imagery (e.g. Jean et al. 2016b, 2019; Yeh et al. 2020; Ronneberger, Fischer, and Brox 2015).



Figure 1: Quality: Freely available Africa imagery extracted via the Google MAPS api

Data

In general, reliable data at a more granular spatial level is very scarce for the African continent. This poses a particular challenge if the researcher wants to apply machine learning tools that require some form of ground truth data.

To overcome this problem, we accessed data from two open-data sources. The first one is Open Street Map, a collaborative project allowing volunteers around the world to contribute georeferenced information in an open-source GIS. We utilized <http://download.geofabrik.de/> to retrieve a complete snapshot of all geo-located objects Africa in 2018. In general, OSM coverage for Africa is very sparse and often non-existent outside urban areas. Our strategy to mitigate this issue, was to build an iterative procedure that would help us select areas (1 × km) with good OSM coverage. We were then able to convert the geometric OSM data into an image mask.

Our image data was collected in 2018 via the google maps api following the exact procedure as in Jean et al. (2016b). This data set has been used in various studies (e.g. Jean et al. 2019; Sheehan et al. 2019; Uzgent et al. 2019; Oshri et al. 2018). Again the same pattern as with OSM data emerges, the image quality of these freely available African images is not as good as in other places around the world, see figure 1.

In the absence of reliable ground truth data, we selected the architecture based on data that we could make look like as if it would be from our target domain. For buildings, we employ imagery collected by drones in Africa from the “Open Cities AI Challenge: Segmenting Buildings for Disaster Resilience”¹, with the corresponding ground truth data provided and re-scaled and blurred the drone imagery. For roads, we build a model to select images with almost complete masks, albeit having missing roads and errors.

We benchmark our proposed methodology against the latest publication of poverty predictions in Africa using their provided wealth index based on DHS cluster data (Yeh et al. 2020). As it is common in this literature an index is created with a principal component analysis (PCA) out of survey respondents. Again, due to data limitations, research in this area always only performed a in-sample validation. A true out of sample comparison would require a strict separation

¹<https://www.drivendata.org/competitions/60/building-segmentation-disaster-resilience/>

between the train and test set, something that is not possible if the PCA is calculated over all data points across all countries and therefore inflating the prediction results. As such, (Yeh et al. 2020) also provided an index that is based on within country survey respondents. This enables us to benchmark against both indices.

Model Architecture

In principle, we follow the outline of the well known U-Net architecture for medical images (Ronneberger, Fischer, and Brox 2015) and modify it for satellite images creating a Satellite-U-Net (Sat-Unet). figure 2 provides a general overview of our approach. The network contains 61 layers in total, with 11 major blocks of 3 types: convolution / down-sampling block, intermediate convolution block and the de-convolution/up-sampling block. The convolution block, shown in figure 3, consists of a batch normalization layer, two convolution layers with the kernel of (3,3) and a dropout layer. The dropout layer is not used in the first down-sampling block. The number of filters in down-sampling blocks (encoder part) starts from 32 and doubles every time in the following block reaching 1024 in the intermediate convolution block, and then decreases in the up-sampling blocks (decoder) with the coefficient 0.5. The core difference from (Ronneberger, Fischer, and Brox 2015) is that instead of up-sampling layers, we are using transposed convolution layers, which performs the reverse convolution operation (Dumoulin and Visin 2016). In addition, we added drop-out layers after each convolution and de-convolution block.

In a 400x400 image the number of pixel belong to a house or a road, class one, vs the number of pixels of zero-class (non classified space) is up to 10^4 times higher, creating a severe class imbalance. We address this issue with a hybrid loss function. First, we use the loss of the sum of binary cross entropy

$$H_p(q) = -\frac{1}{N} \sum_{i=1}^N y_i \cdot \log(p(y_i)) + (1 - y_i) \cdot \log(1 - p(y_i)) \quad (1)$$

and the Sorensen-Dice coefficient:

$$S_V = \frac{2|a \cdot b|}{|a|^2 + |b|^2} \quad (2)$$

combined. As metric we used the intersection over a union (the Jaccard index)

$$J(A, B) = \frac{|A \cap B|}{|A \cup B|} = \frac{|A \cap B|}{|A| + |B| - |A \cap B|} \quad (3)$$

Data pre-processing

To make image input size is 400x400, compatible with a factor of 32 to conform the shape reduction coefficients of the network, we added a padding of 8. On average across our images, the RGB colours maximum was around 180-190 out of the maximum of 255. Color channels re-scaling has been implemented to intensify colors before feeding the image into the network. For augmentation we used rotation by 90, 180 and 270 degrees.

Model evaluation under uncertainty

The main difficulty in choosing the exact architecture for the road network was the lack of a sufficiently large amount of, error-free, ground truth data. Therefore, we used the following iterative strategy:

1. Create an initial mask with OSM data and train on them.
2. Filter out masks, where the model predicts significantly more objects than the OSM mask has.
3. Retrain the model on the filtered data-set.

Due to the large possible set of images to train from, around 22 million, we first selected a subset based on OSM data. As our main focus is to get an indicator of the economic development, the best case would be to find areas of economic activity. OSM has a classification for commercial buildings, which is rarely used (2143/22 mil). We selected areas in the same ADM2 regions of those images based on descending order of square meters occupied by buildings on a uniform grid, until we had selected a base set of 10000 masks. Next, we trained our Sat-Unet model for roads on all masks as labels we had created from OSM.

The Judge: For filtering purposes the Sat-U-Net based model Judge has been created with an additional input for the OSM mask. Using transfer learning the weights of the pretrained Sat-Unet model have been transferred to the bottom layers of the Judge for the mask creation from the original image, and top layers perform the calculations of the Index of validity, using the calculated mask and the OSM mask as inputs. Combining everything in a single GPU model allows to achieve more than 50x increase in performance comparing to CPU-based technique. 4. The index of validity is

$$\alpha = \frac{\sum_{i=1}^K i}{\sum_{j=1}^K j} \quad (4)$$

where i, j - pixel values of predicted mask and the OSM respectively. The resulting filtering model decreased the data-set approximately by 40%, filtering out instances like those presenting in figure 5. Furthermore, as our network predictions had very low values caused by model uncertainty, we re-trained several Sat-Unet model on the selected images with different random seeds. This ensemble learning significantly increased our predictive performance, as shown in figure 7. The top three are predictions on the test set, while the last image is the combination of all three with the pixels reduced to the skeleton for counting.

Building Recognition

For building recognition, we used Open Cities AI Challenge data set as ground truth data set. This data set contains imagery of several African cities in a ultra high resolution of up to 5cm per pixel. Each city is split by square areas and for each image there is a GeoJSON file with vector data describing contours of the buildings. Out of these geometric data, we created a contour layer and a centroid layer, which represents the center of every building structure. We down scaled the images to the scale of 1.53 pixel/meter to match Google

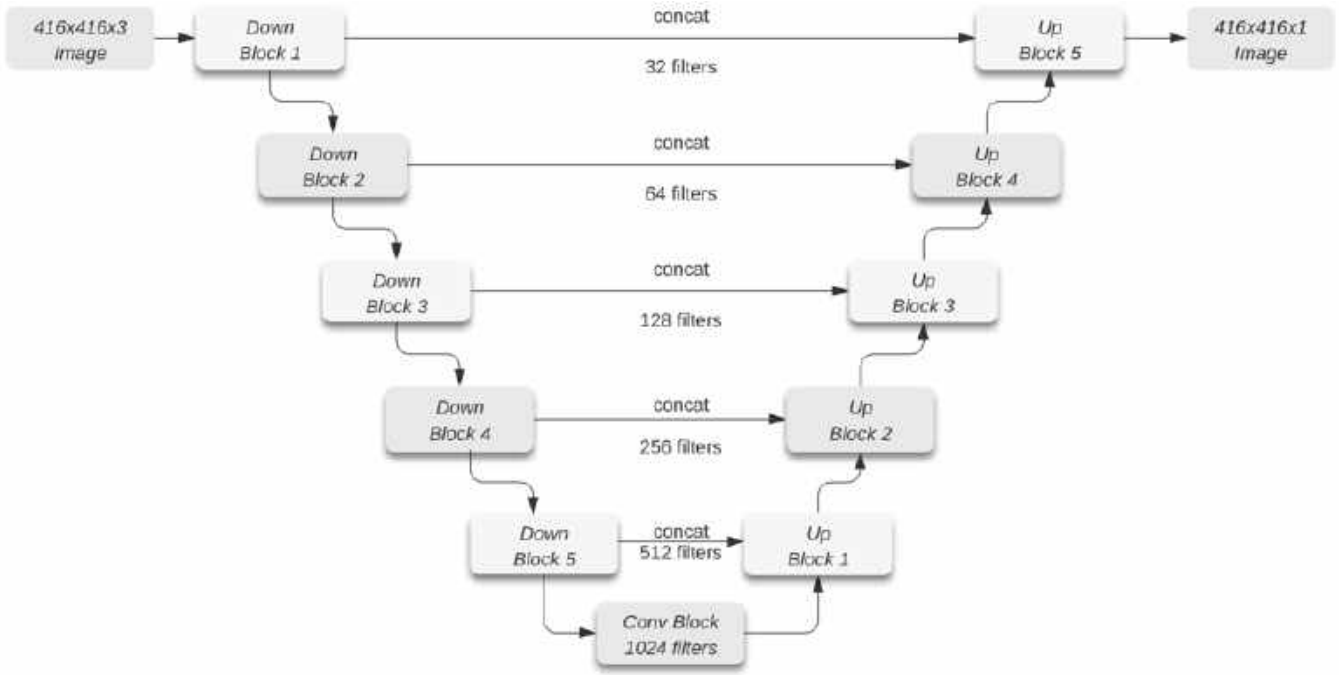


Figure 2: Satellite segmentation network (Sat-Unet) general overview.

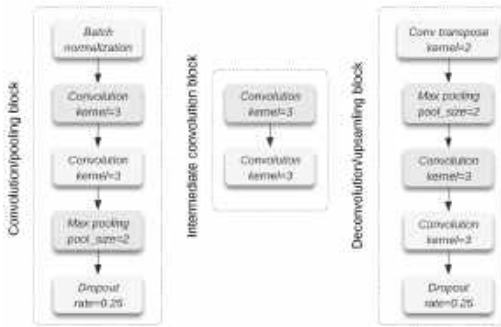


Figure 3: Blocks of the segmentation network

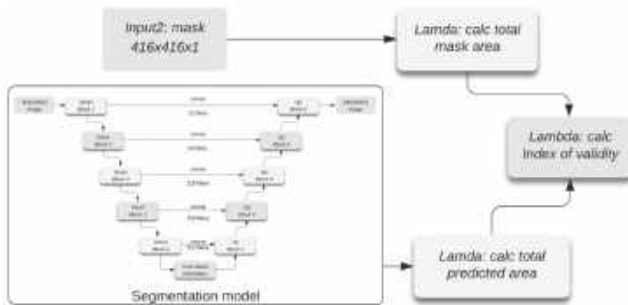


Figure 4: The model judge

Maps images and split the images into 400×400 patches. As target mask for training we used the centroid layer as we were interested in the number of houses, not necessary the full area extent.

Type	Loss	IOU	Dice
Sat-Unet	1.299	0.010	0.019
Incep3-Unet	0.895	0.282	0.329
Resnet50-Unet	0.870	0.326	0.380

Table 1: Model pre-selection results, an encoder of ResNet50 has been chosen.

This high-quality ground-truth data further allowed us to experiment with different architectures. We replaced the encoder part in the Sat-Unet model 2 with the inception v3(Xia, Xu, and Nan 2017) and the resnet50(He et al. 2016). In every instance we reset all the weights to random before training but we did not make use of any transfer learning. Table 1 presents the evaluation results on our test set. We also compared how well the network performs in counting the correct house based on the jaccard index, visually shown in figure 6.

The performance of the Incep3-Unet and Resnet50-Unet was quite similar. For the final model selection we trained both architectures on the previous selected masks out of OSM, and compared their performance in terms of their predictability on a test set. Table 2 demonstrates the effect of different thresholds on the performance. The threshold is in color intensity units (range 0-255), TP is true positive, when the predicted centroid is located inside the building contour of the mask, Pred-To-Mask coefficient is the ratio between predicted number of houses and the ground truth number, and False Positives (FP):

$$FP \leftarrow 100 \cdot \frac{TotalPred - TP}{TotalPred} \quad (5)$$

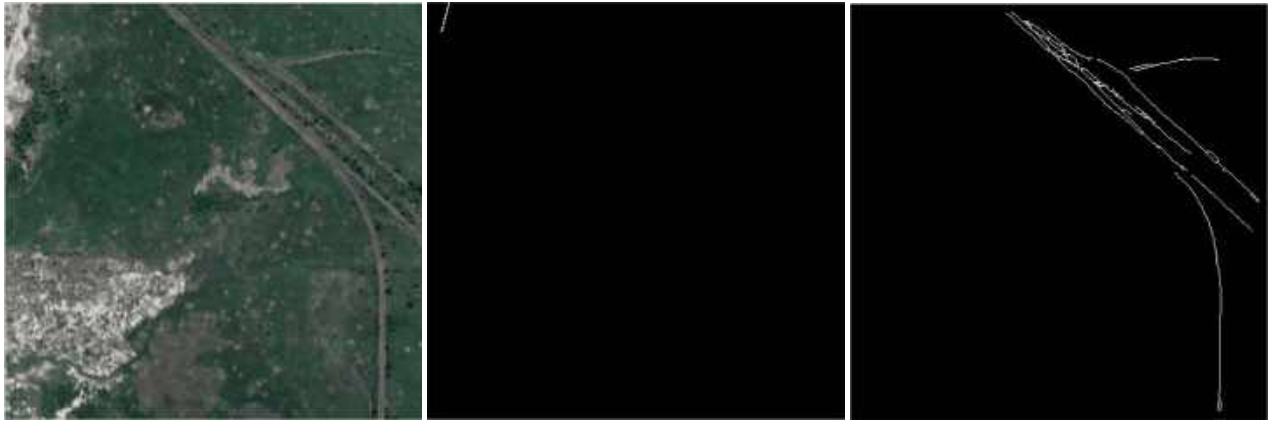


Figure 5: The filter removes incomplete OSM masks. From left to right: original image, OSM mask, prediction of the trained model on the stage 1.

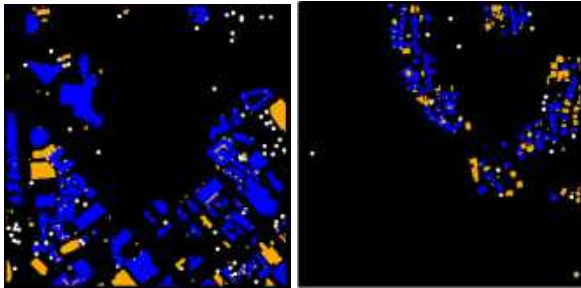


Figure 6: Africa, building model stage, contour-in-contour evaluation. Buildings that were predicted correctly are in blue, not predicted in orange. Orange areas of random shapes inside building blocks are usually courtyard areas and not considered as a wrong prediction.

A threshold of 15 has the closest Prediction-to-mask score, acceptable TP and FP rates. Therefore, we picked this threshold for further modeling. The final model Resnet50-Unet has 268 layers. To avoid the vanishing gradient problem with the depth of hundreds of layers ResNet uses skip connections, it adds input information of the convolution block to its output. In addition skip connections give the model the ability to learn the identity function which guarantee the similar performance of the lower and higher layers (He et al. 2016).

Results

Benchmark Results

We compare our prediction results of buildings and roads to the latest benchmark study in the field of poverty prediction based on DHS data (Yeh et al. 2020). The DHS data is collected in various waves across countries and years. For the comparison use of the most recent wave available for a country and use the aggregated wealth index data. Two indices are provided, the first *wealthpooled* is the PCA calculation across all years, while the second index *wealth*, is calculated

Type	Thresh	TP	Pred-To-Mask	FP
Incep3-Unet	05	47.876	126.226	61.143
Incep3-Unet	10	45.639	112.446	60.903
Incep3-Unet	15	44.243	102.651	61.529
Incep3-Unet	25	42.162	85.075	61.360
Resnet50-Unet	05	50.640	129.173	63.654
Resnet50-Unet	10	48.496	113.195	61.774
Resnet50-Unet	15	47.1438	105.009	61.644
Resnet50-Unet	25	45.099	93.966	61.750

Table 2: Threshold effect

out of the survey respondents only within a country and year, allowing out of sample comparison.

(Yeh et al. 2020) did not provide any out of sample estimates, therefore we fully replicated their method by getting all their images they used for their location and trained their combined CNN model of multi-spectrum and night lights images to determine economic well-being in Africa, with *wealth* as label for the last of their training folds (D). The performance results are almost identical in terms of R-squared. In their study they also found a high correlation of the measures to other type's of aggregation, such the sum total of all assets. Using their *wealthpool* predictions as predictor for *wealth*, the r-squared is around 0.61 vs 0.67 for *wealthpool*.

Prediction Results

The DHS location data from (Yeh et al. 2020) has 7,315 unique cluster location in the last wave of each respective country. We use a 5km radius corresponding to the possible displacement of survey measurements, as selection criteria to select our images. For every square km we predict the number of buildings, the number of roads as well as calculate the night time light (Elvidge et al. 2017) by grid cell. We then aggregate the roughly 1.5 million images into features, by building the sum, average and quantiles by cluster across all 3 input variables. In total, this leaves us with 6,112 locations. We perform LOOCV cross validation by country, as

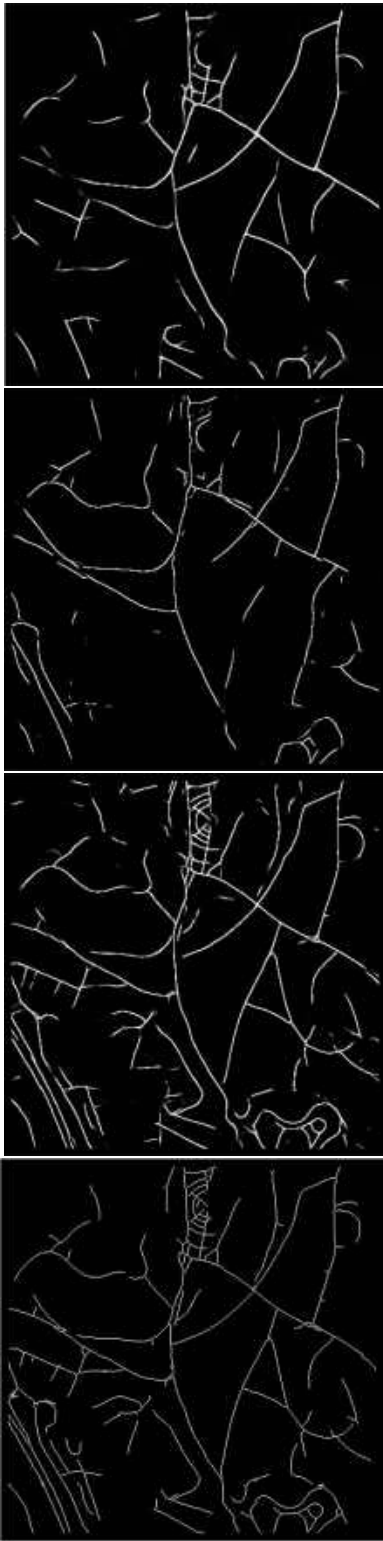


Figure 7: Ensembling: The first three are predicted masks based on different random seeds, the last one is the resulting mask.

we are interested in predicting the marginal unit if we would use the model to get data for one extra country. We also iterate over standard machine learning algorithms without any hyper parameter tuning, by only using the default settings. Table 3 and table 4 present the results for out of sample and out of country predictions, respectively. As expected, using a normalized outcome measures across all samples, inflates the performance. In comparison to the previous literature, our predictions show an increased predictive performance, both in and out of sample.

Discussion and Future Work

This paper introduces a novel and scalable method to predict road and housing infrastructure from daytime satellite imagery. Compared to existing approaches, we achieve higher predictive performance by training a U-net style architecture using ground-truth data from a subset of images. Using satellite images from 21 African countries we show how our method can be used to generate very granular information about the stock of housing and road infrastructure for regions in the world, where reliable information about the local level of economic development is hardly available. Consistently measured and comparable indicators about local economic development are crucial inputs for governments in developing countries as well as international organizations in their decision where to allocate scarce public funds and development aid.

The predictions generated by our method can be directly included in existing decision support systems. For example, international organization such as the Red Cross are using similar data at the local level to evaluate an area's vulnerability against natural hazards. Our data can be considered as more granular complements to existing measures of the local stock of physical infrastructure. Numerous charitable organizations already rely on satellite imagery to identify districts of African countries that are among the least developed (e.g. Abelson, Varshney, and Sun 2014). Our approach provides a low-cost and scalable alternative to identify areas that are in need. In addition, the Open Street Map mapping community would benefit from our findings as well. The road prediction model could be used worldwide to help completing the road network or help narrowing down possible errors in the data.

Finally, our approach is an important methodological contribution to the large group of scholars from varying disciplines working in the area of poverty measurement. The majority of the existing research focuses on predicting poverty based on aggregate household wealth. This paper shows that predicting poverty measures can also be viewed as a simple high dimensional feature representation problem. Our study is a proof-of-concept exercise to show that combining daytime satellite imagery, open source ground truth data and machine learning tools can translate unstructured image data into valuable insights about local economic development at an unprecedented scale.

Model	Buildings	Roads	Buildings & Roads	Nighttime light	Buildings, Road & Nighttime light
ridge	0.322	0.154	0.155	0.353	0.155
r-tree	0.336	0.459	0.468	0.540	0.567
r-tree (bs)	0.437	0.552	0.576	0.637	0.652
r-tree (bg)	0.495	0.609	0.639	0.678	0.723

Table 3: Predictive performance of satellite predictions, r-squared based on LOOCV on out of country and out of sample predictions by country for *wealth*

Model	Buildings	Roads	Buildings & Roads	Nighttime light	Buildings, Road & Nighttime light
ridge	0.110	0.146	0.146	0.311	0.147
r-tree	0.239	0.401	0.402	0.589	0.594
r-tree (bs)	0.359	0.494	0.525	0.678	0.692
r-tree (bg)	0.419	0.563	0.595	0.719	0.738

Table 4: Predictive performance of satellite predictions, r-squared based on LOOCV on out of country predictions by country using *wealthpooled*

Appendix

List of countries used: Angola, Benin, Burkina Faso, Cameroon, Central African Republic, Côte d’Ivoire, Democratic Republic of the Congo, Ethiopia, Ghana, Guinea, Kenya, Lesotho, Malawi, Mali, Nigeria, Rwanda, Senegal, Sierra Leone, Tanzania, Togo and Uganda

References

Abelson, B.; Varshney, K. R.; and Sun, J. 2014. Targeting Direct Cash Transfers to the Extremely Poor. *KDD '14: Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining* 1563–1572.

Baragwanath, K.; Goldblatt, R.; Hanson, G.; and Khandelwal, A. K. 2019. Detecting urban markets with satellite imagery: An application to India. *Journal of Urban Economics* 103–173.

Berman, N.; Couttenier, M.; Rohner, D.; and Thoenig, M. 2017. This Mine Is Mine! How Minerals Fuel Conflicts in Africa. *American Economic Review* 107(6): 1564–1610. doi:10.1257/aer.20150774. URL <https://www.aeaweb.org/articles?id=10.1257/aer.20150774>.

Blumenstock, J.; Cadamuro, G.; and On, R. 2015. Predicting poverty and wealth from mobile phone metadata. *Science* 350(6264): 1073–1076.

Burchfield, M.; Overman, H. G.; Puga, D.; and Turner, M. A. 2006. Causes of Sprawl: A Portrait from Space. *The Quarterly Journal of Economics* 121(2): 587–633.

Chen, X.; and Nordhaus, W. D. 2011. Using luminosity data as a proxy for economic statistics. *Proceedings of the National Academy of Sciences* 108(21): 8589–8594.

Donaldson, D.; and Storeygard, A. 2016. The View from Above: Applications of Satellite Data in Economics. *Journal of Economic Perspectives* 30(4): 171–98.

Doupe, P.; E.Bruzelius; Faghmous, J.; and S.G.Ruchman. 2016. Equitable development through deep learning: The case of subnational population density estimation. In *proc. of the ACM DEV*.

Dumoulin, V.; and Visin, F. 2016. A guide to convolution arithmetic for deep learning. *arXiv preprint arXiv:1603.07285*.

Elvidge, C. D.; Baugh, K.; Zhizhin, M.; Hsu, F. C.; and Ghosh, T. 2017. VIIRS night-time lights. *International Journal of Remote Sensing* 38(21): 5860–5879.

Elvidge, C. D.; Sutton, P. C.; Ghosh, T.; Tuttle, B. T.; Baugh, K. E.; Bhaduri, B.; and Bright, E. 2009. A global poverty map derived from satellite data. *Computers Geosciences* 35(8): 1652–1660.

Engstrom, R.; Newhouse, D.; Haldavanekar, V.; Copenhagen, A.; and Hersh, J. 2017. Evaluating the relationship between spatial and spectral features derived from high spatial resolution satellite data and urban poverty in Colombo, Sri Lanka. In *2017 Joint Urban Remote Sensing Event (JURSE)*, 1–4.

He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.

Henderson, J. V.; Storeygard, A.; and Weil, D. N. 2012. Measuring Economic Growth from Outer Space. *American Economic Review* 102(2): 994–1028. doi:10.1257/aer.102.2.994. URL <http://www.aeaweb.org/articles?id=10.1257/aer.102.2.994>.

Hodler, R.; and Raschky, P. A. 2014. Regional favoritism. *The Quarterly Journal of Economics* 129(2): 995–1033.

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

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

Jean, N.; Wang, S.; Samar, A.; Azzari, G.; Lobell, D. B.; and

- Ermon, S. 2019. Tile2Vec - Unsupervised Representation Learning for Spatially Distributed Data. *AAAI* .
- Jerven, M.; and Johnston, D. 2015. Statistical tragedy in Africa? Evaluating the database for African economic development. *Journal of Development Studies* 51(2): 111–115.
- Johnson, S.; Larson, W.; Papageorgiou, C.; and Subramanian, A. 2013. Is newer better? Penn World Table Revisions and their impact on growth estimates. *Journal of Monetary Economics* 60(2): 255 – 274.
- Marx, B.; Stoker, T.; and Suri, T. 2013. The economics of slums in the developing world. *Journal of Economic Perspectives* 27(4): 187–210.
- Moyo, S.; Endert, A. S. V.; Rudzvizo, S.; Beunderman, J.; and Treger, C. 2020. African Cities — Disrupting the Urban Future. URL <https://medium.com/@undp.innovation/african-cities-disrupting-the-urban-future-5a774b101e58>.
- Oshri, B.; Hu, A.; Adelson, P.; Chen, X.; Dupas, P.; Weinstein, J.; Burke, M.; Lobell, D.; and Ermon, S. 2018. Infrastructure Quality Assessment in Africa using Satellite Imagery and Deep Learning. In *the 24th ACM SIGKDD International Conference*, 616–625. New York, NY, USA: ACM.
- Pandey, S. M.; Agarwal, T.; and Krishna, N. C. 2018. Multi-task deep learning for predicting poverty from satellite images. In *proc.of the AAAI*.
- Robinson, C.; Hohmans, F.; and Dilkina, B. 2017. A deep learning approach for population estimation from satellite imagery. In *proc. of the ACM SIGSPATIAL Workshop on Geospatial Humanities*.
- Ronneberger, O.; Fischer, P.; and Brox, T. 2015. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, 234–241. Springer.
- Sheehan, E.; Meng, C.; Tan, M.; Uz Kent, B.; Jean, N.; Burke, M.; Lobell, D.; and Ermon, S. 2019. Predicting Economic Development using Geolocated Wikipedia Articles. In *the 25th ACM SIGKDD International Conference*, 2698–2706. New York, New York, USA: ACM Press.
- Simonyan, K.; and Zisserman, A. 2015. Very deep convolutional networks for large-scale image recognition. In *Proceedings of the ICLR*.
- Sutton, P.; Elvidge, C.; and Tilottama, G. 2007. Estimation of Gross Domestic Product at Sub-National Scales Using Nighttime Satellite Imagery. *International Journal of Ecological Economics Statistics* 8.
- Sutton, P. C.; and Costanza, R. 2002. Global estimates of market and non-market values derived from nighttime satellite imagery, land cover, and ecosystem service valuation. *Ecological Economics* 41(3): 509–527.
- Treger, C.; Johar, I.; Koulouri, K.; Beunderman, J.; Pieterse, E.; Cirolia, L.; van Endert, A. S.; and Begovic, M. 2020. The NextGenCities Africa Programme. URL <https://medium.com/@undp.innovation/the-nextgencities-africa-programme-a01014171913>.
- Uz Kent, B.; Sheehan, E.; Meng, C.; Tang, Z.; Burke, M.; Lobell, D. B.; and Ermon, S. 2019. Learning to Interpret Satellite Images using Wikipedia. *IJCAI* .
- Xia, X.; Xu, C.; and Nan, B. 2017. Inception-v3 for flower classification. In *2017 2nd International Conference on Image, Vision and Computing (ICIVC)*, 783–787. IEEE.
- Yeh, C.; Perez, A.; Driscoll, A.; Azzari, G.; Tang, Z.; Lobell, D.; Ermon, S.; and Burke, M. 2020. Using publicly available satellite imagery and deep learning to understand economic well-being in Africa. *Nature communications* 11(1): 1–11.