Measuring Changes in Poverty with Deep Learning and Satellite Images

Lukas Kondmann and Xiao Xiang Zhu

Technical University of Munich (TUM) & German Aerospace Center (DLR)

Contact: xiaoxiang.zhu@dlr.de

Twitter: @kondmann, @xiaoxiang_zhu

DLR


1
NO
POVERTY

Daytime Images and Deep Learning can improve on mapping development with Nighttime Lights (NTLs) especially in poor regions (Jean et al. 2016)

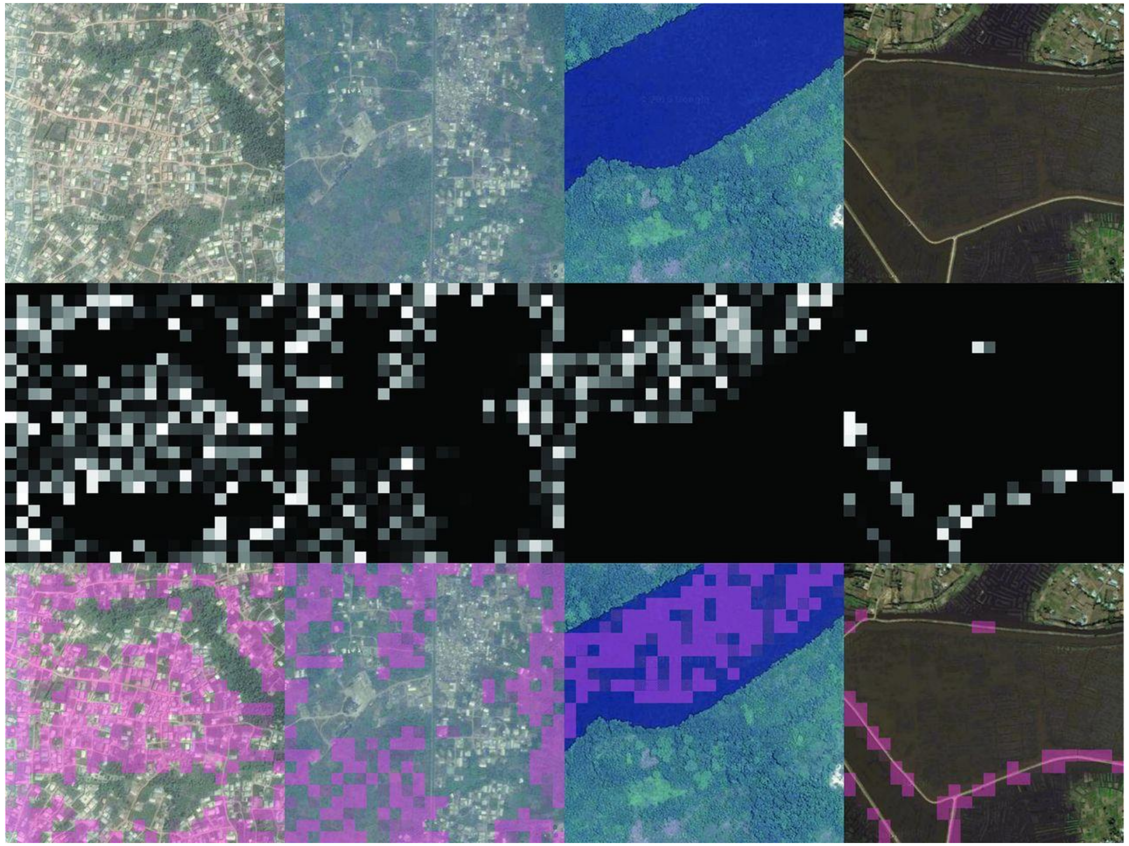

The Issue: Change Detection Abilities are unclear

‘A lot of this work is using satellite data […] to try to predict poverty at a granular level for entire countries or continents. At the moment such data **seems useful for descriptive work**, but it is **unclear whether accuracy is enough to measure changes well over time**. […] There may not be enough signal to be able to detect the impact of interventions.’

David McKenzie,
Lead Economist World Bank Development (Link)

So can this approach also detect changes in poverty over time accurately?

Research Question: Can this approach also detect changes in poverty over time?

– In this work, we test this property in Rwanda with poverty data and satellite images from 2005 to 2015.

– Data:
- Poverty Data: Household asset wealth scores in the Demographic and Health Surveys (DHS) in Rwanda from 2005, 2010 and 2015.
- Nighttime Lights: National Oceanic and Atmospheric Administration (NOAA) annual time series on worldwide nighttime luminosity (DMSP-OLS) from 1992 to 2013
- Daytime Images: Yearly composite images of Landsat 7 as input images for poverty mapping

– Approach:
1. Train on 2005 images and poverty data and obtain predictions for 2010 and 2015
2. Take temporal differences of predicted values ➔ Predicted change in poverty
3. Compare predicted changes to on the ground data and predicted change with nighttime lights (NTLs)

Poverty Ground Truth: DHS Household Wealth Data

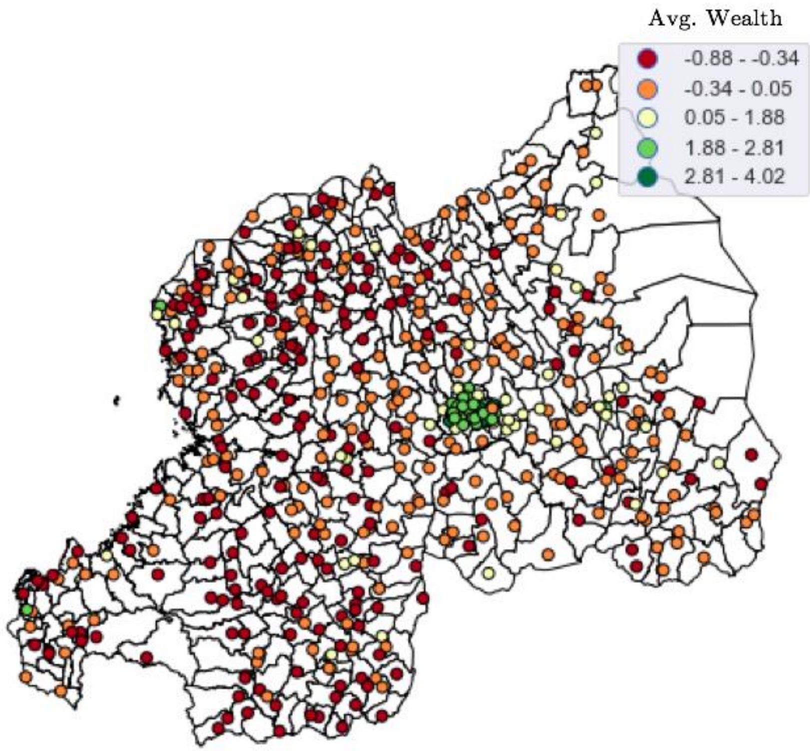

- Poverty baselines are taken from household asset wealth scores in the Demographic and Health Surveys (DHS) in Rwanda from 2005, 2010 and 2015.
- Households are surveyed in geographic clusters such that all households within a cluster share the same coordinates.
- The wealth index is the result of taking the first principal component of survey responses about asset ownership and normalizing the component within each country and year
- These scores are hence merely a relative ranking of households within a survey.

NTLs data: DMSP-OLS annual time series

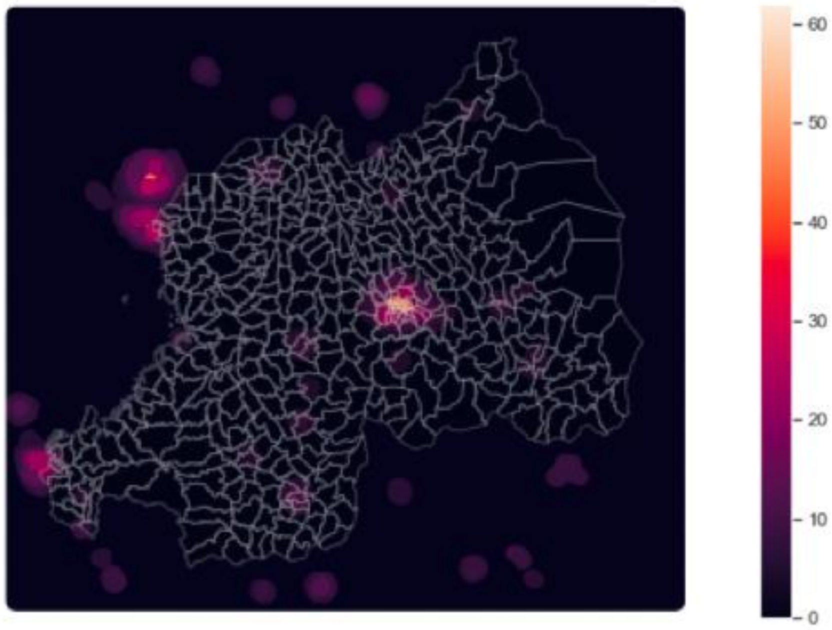

- We use the v4 DMSP-OLS stable annual time series of nightlight activity from 1992 to 2013 on 30 arc second grids which is equivalent to roughly 1km×1km

- Every cell has an average nightlight intensity value on a scale from 0 to a maximum of 63.

- Nightlights activity in Rwanda is mostly 0 and the maximum is the capital Kigali in the center.

- We use nightlight activity from 2013 as a proxy for 2015 which is the time of data collection of the survey data and the satellite imagery.

Daytime Images: Landsat 7

Google Static Maps

2.4m resolution per pixel

Public Landsat 7

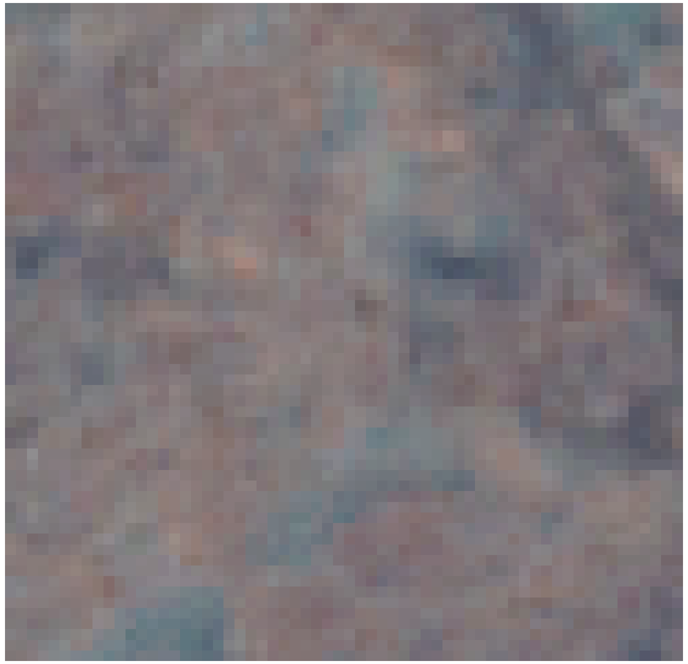

15m resolution per pixel
(pan-sharpened)

Public Landsat 7 raw

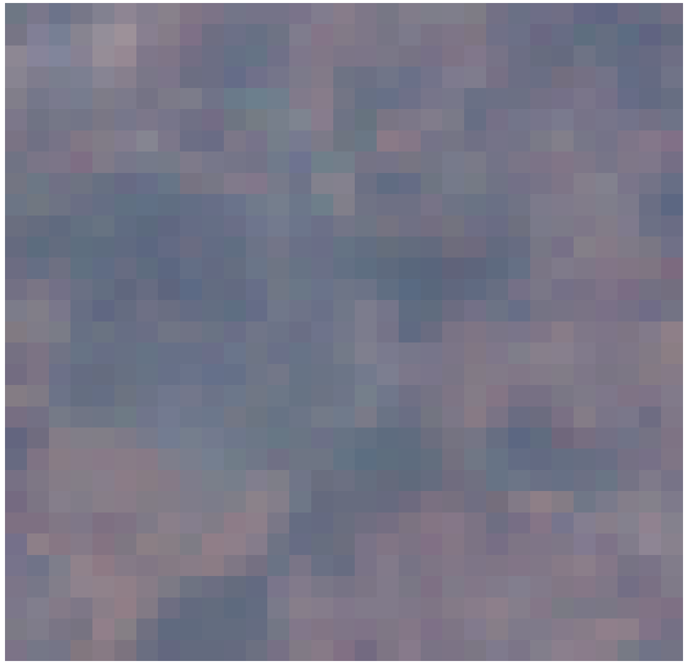

30m resolution per pixel

Daytime Images: Landsat 7

- Jean et al. (2016) use Google Static Maps Imagery for their transfer learning strategy but **this is unsuitable for time-series analysis.**
- Hence, we follow Perez et al (2017) and use yearly composite images of Landsat 7 as multi-temporal input for poverty mapping.
- These 1km x 1km composites are the result of taking the median of all cloud-free pixels of the images within the year in question.
- Training locations are sampled from all over West Africa.
- The feature extractor is fine-tuned with 65.000 images in total of which 40.000 are randomly sampled from Rwanda.
- We train with Adam set to defaults for 25 epochs and choose the model with the best validation accuracy.
- In this work, we only use the RGB channels of the images.

Method: The Transfer Learning Strategy

	Input		Output
Transfer Learning	Daytime Images	Pre-trained ResNet 50 →	Classify NTL Activity
Feature Extraction	Cluster Images	Fine-tuned ResNet 50 →	Feature Vector
Poverty Mapping	Feature Vector	Ridge Regression →	Predicted Local Wealth

Results (1): Predicted county poverty aligns well with cluster poverty levels

A Nightlight Activity

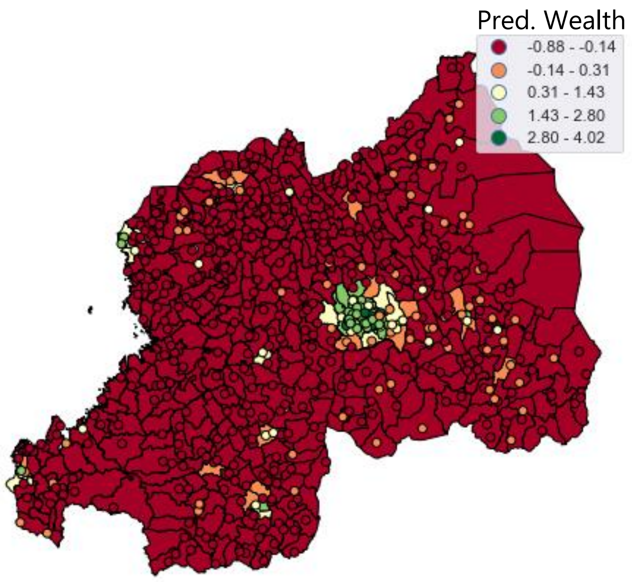

B Transfer Learning

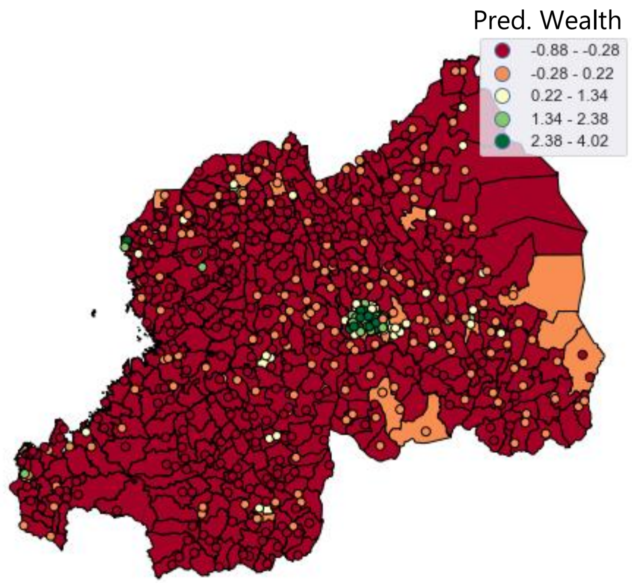

Results (2): Predicted change in county poverty creates an inconsistent picture

A Nightlight Activity

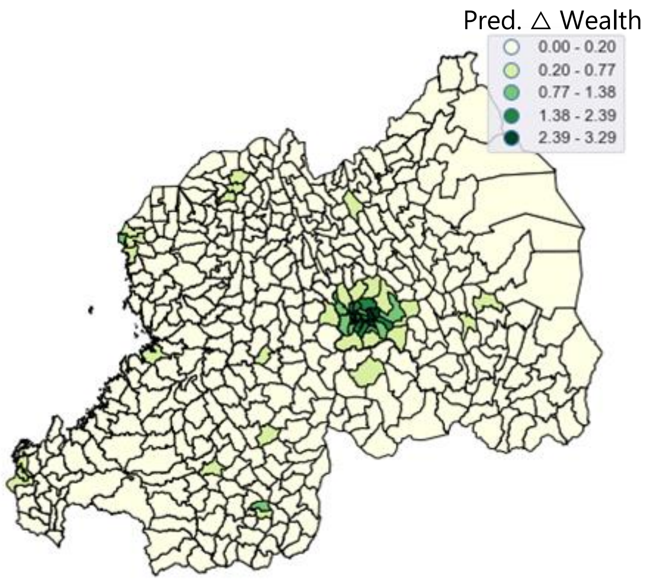

B Transfer Learning

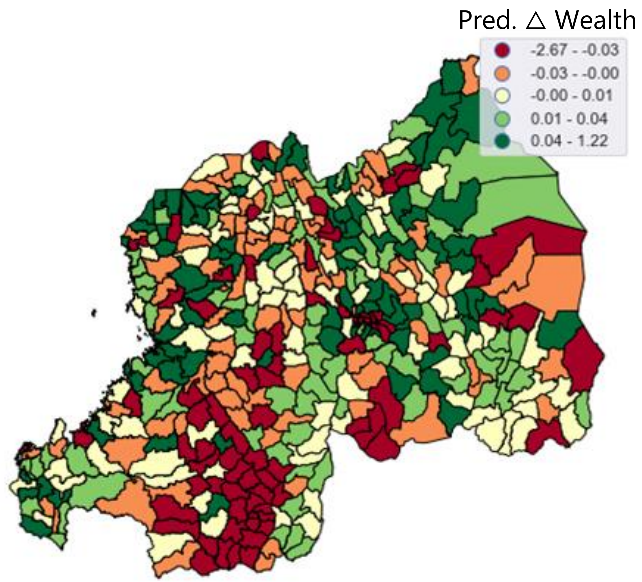

Results (3): TL seems to be do especially bad in urban areas for poverty change detection

Table 3: Predicted Average Absolute Change 2005-15 on Province Level

Time	Change TL	Change NTL	Poverty Reduction
Kigali City	-0.04	0.83***	0.11
Eastern Province	0.01***	0.01***	0.16
Northern Province	0.02***	0.01***	0.09
Southern Province	-0.04***	0.04***	0.21
Western Province	0.02***	0.02***	0.12

Significance level: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$
Poverty reduction rates for 2005-2010 taken from the World Bank (2015) and for 2010-15 from the National Institute of Statistics of Rwanda (NISR)

– The true change in Kigali is potentially overestimated by NTL because they tend to overglow in urban regions. The direction is, however, consistently estimated.
– Hence,if anything, the transfer learning strategy would predict an even lower country level growth rate if adjusting for population in Table 2.
– Results for the Eastern, Northern and Western Province seem to align but also results for the Southern Province are discouraging

Discussion: What could be reasons for this pattern?

1. Change signal is more subtle and harder to extract compared to simply mapping levels of poverty
   - Aggregates of satellite images are relatively stable over time by design
   - If changes are small compared to the levels of local wealth, they are hard to recognize because the signal to noise ratio is a lot worse

2. Low resolution of images prevents fine-grained change analysis
   - Might recognize urban vs. rural region but can not distinguish *better* roofs or roads
   - This is especially relevant for the 15m/pixel resolution of our images

3. Changes in wealth manifest not directly in structures recognizable from space
   - The TL strategy might not detect change instantly since the changes do not lead to improvements straight away
   - A longer timeframe than 10 years might show the ability of this strategy to detect changes

Conclusion

– Our results indicate that pioneering methods that map poverty from satellite images with deep learning may struggle to capture trends in economic development over time in their current form.

– Results are limited to the setting of Rwanda and might fail to generalize. But changes should be most visible in Rwanda, because
  - There is a clear trend in economic development on the ground
  - The ability of the TL strategy to map levels of Poverty in Rwanda is well established

– Careful validation of these outcomes in other countries and with other imagery is necessary to speak to the robustness of this finding.

– Further refinements of these methods to map changes could involve:
  - Training on recognizing changes in NTLs directly
  - Data fusion with other sources that have been beneficial for poverty mapping such as Open Street Map or Land Cover data

References

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

2. Anthony Perez, Christopher Yeh, George Azzari, Marshall Burke, David Lobell, and Stefano Ermon.Poverty prediction with public landsat 7 satellite imagery and machine learning.arXiv preprintarXiv:1711.03654, 2017

3. The World Bank. Rwanda poverty assessment, 2015

4. National Institute of Statistics of Rwanda (NISR). Poverty trend analysis report Rw