Certificate I: Understanding AI and Machine Learning in Africa

Course AIML02: AI and Machine Learning in Africa

Module 02: Application Case Studies

Lecture 05: Socioeconomics

Carnegie Mellon University
Africa

Learning Objectives

1. Understand one role of AI and machine learning in socioeconomics in Africa

2. Identify, explain, and summarize the use of deep learning with satellite images to infer the level of regional and local poverty, particularly in rural areas

3. Understand how to evaluate AI and machine learning strategies, technologies and solutions

Lecture Contents

1. Introduction to the case study target article
2. Geographical information and human well-being
3. Options for acquiring data to generate economic indicators
4. Deep learning with satellite images
    - Implementation details
    - Results
    - Strengths, weaknesses, & possible improvements
5. Lecture summary
6. Recommended reading & references

A Case Study on AI & Machine Learning in Socioeconomics

Socioeconomics is a social science that studies how social and economic activities interact and affect one another

This socioeconomic case study is based on an article by Yeh et al. (2020): Using publicly available satellite imagery and deep learning to understand economic well-being in Africa

The article demonstrates how machine learning, specifically deep learning, can be used to predict asset wealth across approximately 20,000 African villages when trained with satellite images

Geographical Information and Human Well-being

For effective allocation of resources to different locations in a country based on human well-being, we need

- Accurate and up-to-date information
- … on key economic indicators
- … in various locations across the country

Geographical Information and Human Well-being

Asset wealth surveys provide one way to obtain the required information

However, they are:

- Carried out infrequently

Every four or more years in most African countries

- An expensive investment especially for lower-income countries

An estimated annual cost of $1 billion USD

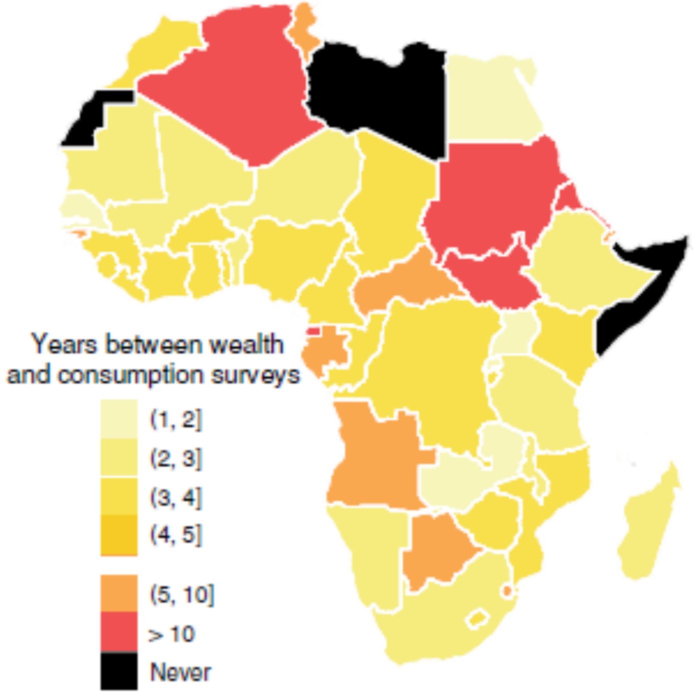

Frequency of household consumption expenditure or asset wealth surveys in Africa, 2000-2016

Geographical Information and Human Well-being

The challenge is to obtain geographical information in a timely, cost-effective manner

The target article sets out to demonstrate that data from satellite imagery can be used to provide the required information

- Local indicators well-being across geographical regions
- Measurements of change in well-being over time

Spatial variations

Temporal variations

Options for Acquiring Data to Generate Economic Indicators

Two approaches have previously been used for obtaining data for economic indicators

1. Traditional data collection methods

2. Deep learning using satellite imagery

Options for Acquiring Data to Generate Economic Indicators

Traditional data collection methods involve the use of surveys and questionnaires such as the Demographic and Health Surveys (DHS)

116	Do you own a mobile telephone?	YES .... 1 NO 2	 → 118
117	Do you use your mobile phone for any financial transactions?	YES .... 1 NO 2	
118	Do you have an account in a bank or other financial institution that you yourself use?	YES .... 1 NO 2	
119	Have you ever used the internet?	YES .... 1 NO 2	 → 122
120	In the last 12 months, have you used the internet? IF NECESSARY, PROBE FOR USE FROM ANY LOCATION, WITH ANY DEVICE.	YES .... 1 NO 2	 → 122
121	During the last one month, how often did you use the internet: almost every day, at least once a week, less than once a week, or not at all?	ALMOST EVERY DAY .... 1 AT LEAST ONCE A WEEK .... 2 LESS THAN ONCE A WEEK .... 3 NOT AT ALL 4	

Excerpt of questions from a DHS survey
https://blog.dhsprogram.com/dhs7-part-3/

Options for Acquiring Data to Generate Economic Indicators

Problems with the traditional data collection methods

- Inconsistency of continuous data collection in African countries
  - A given household may feature in one survey but not in previous or subsequent surveys
  - Difficult to track economic changes and assess the impact of intervention programs over a period of several years
- Cost of conducting the surveys

Options for Acquiring Data to Generate Economic Indicators

Alternative approach: use satellite imagery to determine the economic indicators

Coarse 1 km/pixel nighttime light imagery is used to train deep learning models to predict the asset wealth index

(Henderson, Storeygard, & Weil, 2012)

Options for Acquiring Data to Generate Economic Indicators

Alternative approach: use satellite imagery to determine the economic indicators

(Babenko, Hersh, Newhouse, Ramakrishnan, & Swartz, 2017)

Models can also be trained using high-resolution daytime images (< 1 m / pixel)

Options for Acquiring Data to Generate Economic Indicators

Alternative approach: use satellite imagery to determine the economic indicators

From left to right, four different convolutional filters are used to identify features corresponding to urban areas, nonurban areas, water, and roads.

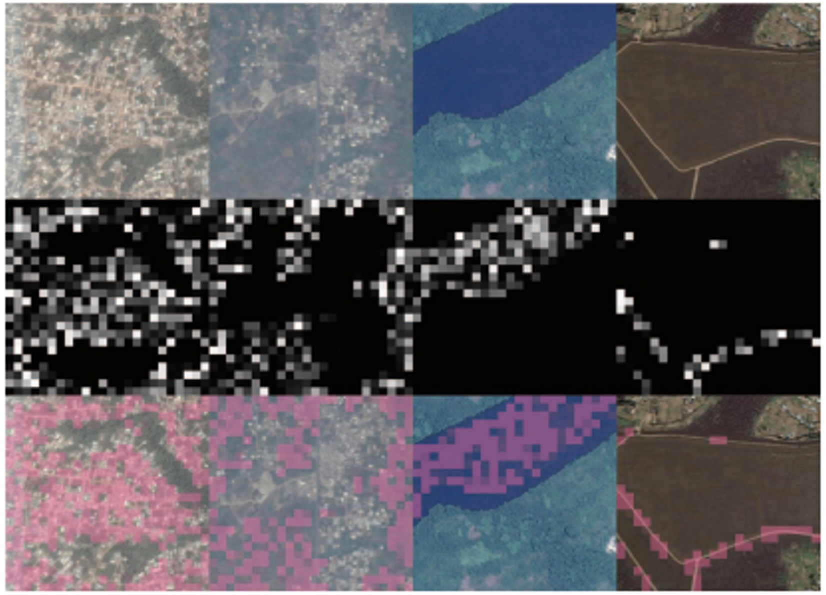

(Jean et al., 2016)

Nighttime lights are used in transfer learning, providing intermediate labels for training a convolutional neural network (CNN) feature extractor to find features relevant to predicting wealth in multi-spectral Landsat imagery

Deep Learning with Satellite Images

The case study uses cloud-free nighttime and daytime imagery to model both spatial and temporal differences in economic well-being at a local level across sub-Saharan Africa

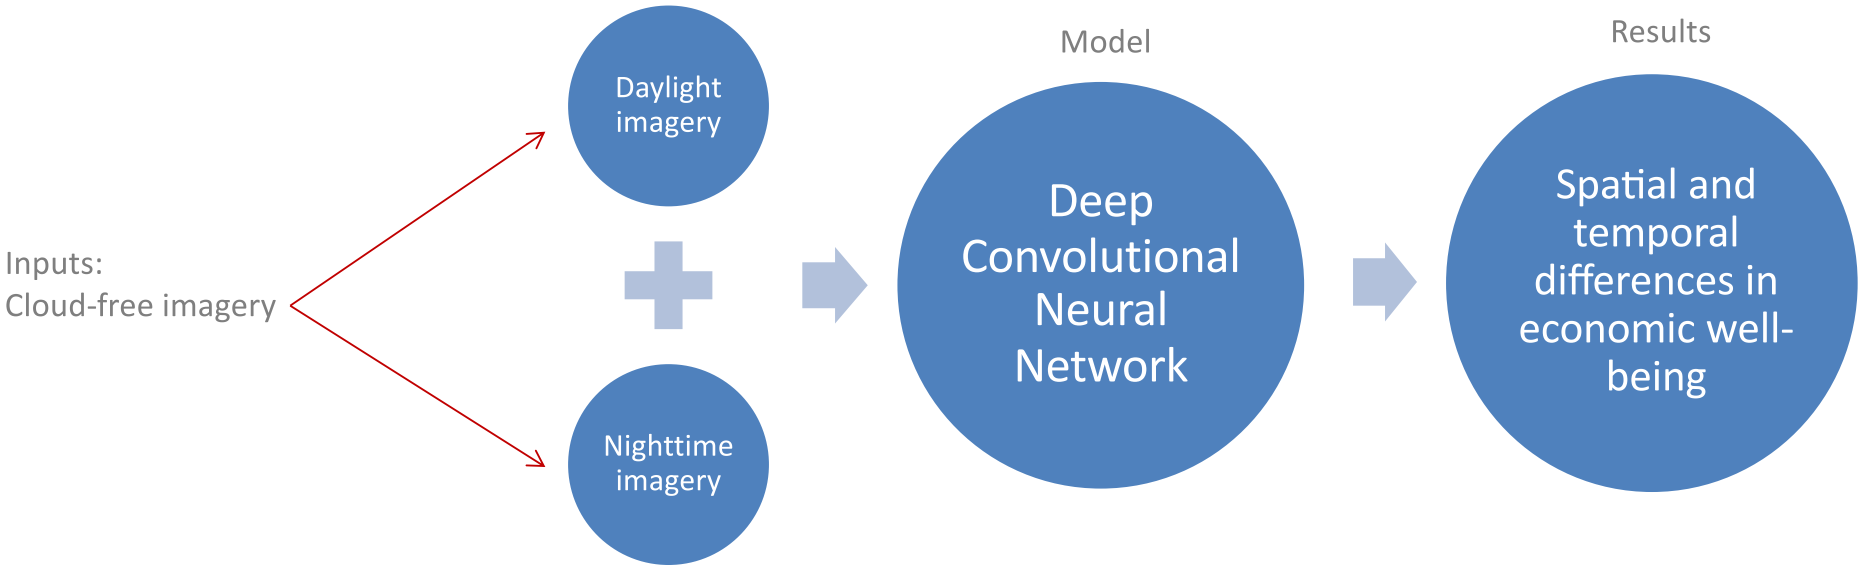

Deep Learning with Satellite Images

In this approach, the asset-wealth index is used as the welfare measurement index

- It is a less noisy measure of a longer-term economic well-being of a household
- It is widely found in publicly available African survey data
- It is also a common component of the multi-dimensional poverty index

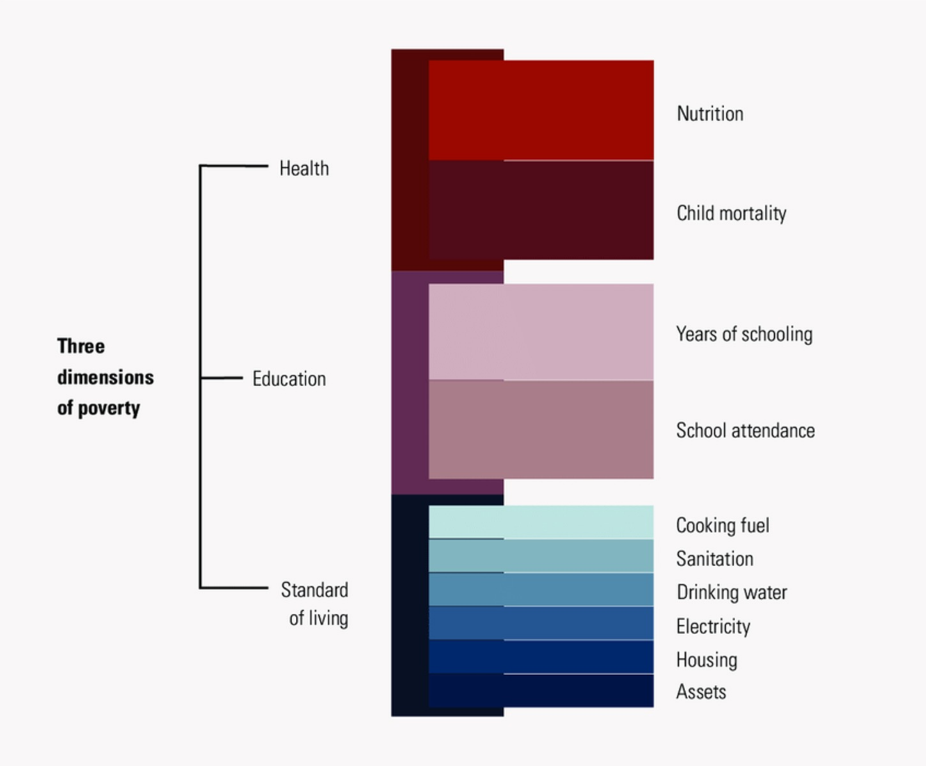

https://www.researchgate.net/figure/Structure-of-the-global-Multidimensional-Poverty-Index_fig1_342989035

Deep Learning with Satellite Images

An out-of-sample model evaluation technique was used to

- To avoid the problem of overfitting

Overfitting occurs when a model produces good results on the training data but poor results on new, unseen data

- To make the model robust when it is applied in situations where there is no ground-truth survey data

Implementation Details

Data Collection

01

43 DHS surveys carried out across 23 African countries between 2009 to 2016

Implementation Details

Data Collection

01

43 DHS surveys carried out across 23 African countries between 2009 to 2016

Data Cleaning

02

Removal of data with invalid coordinates which could not be used for the study

Implementation Details

Data Collection

01

43 DHS surveys carried out across 23 African countries between 2009 to 2016

Data Cleaning

02

Removal of data with invalid coordinates which could not be used for the study

Data Processing

03

Convert responses to asset quality scores between a range 1 to 5

Implementation Details

Data Collection	Data Cleaning	Data Processing	Data Standardization
01	02	03	04
43 DHS surveys carried out across 23 African countries between 2009 to 2016	Removal of data with invalid coordinates which could not be used for the study	Convert responses to asset quality scores between a range 1 to 5	Adjusted asset-wealth index to follow a standard normal distribution

Implementation Details

Data Collection

01

43 DHS surveys carried out across 23 African countries between 2009 to 2016

Data Cleaning

02

Removal of data with invalid coordinates which could not be used for the study

Data Processing

03

Convert responses to asset quality scores between a range 1 to 5

Data Standardization

04

Adjusted asset-wealth index to follow a standard normal distribution

Data Validation

05

PCA index pooled across countries to avoid bias from subjective and country-specific questions

Implementation Details

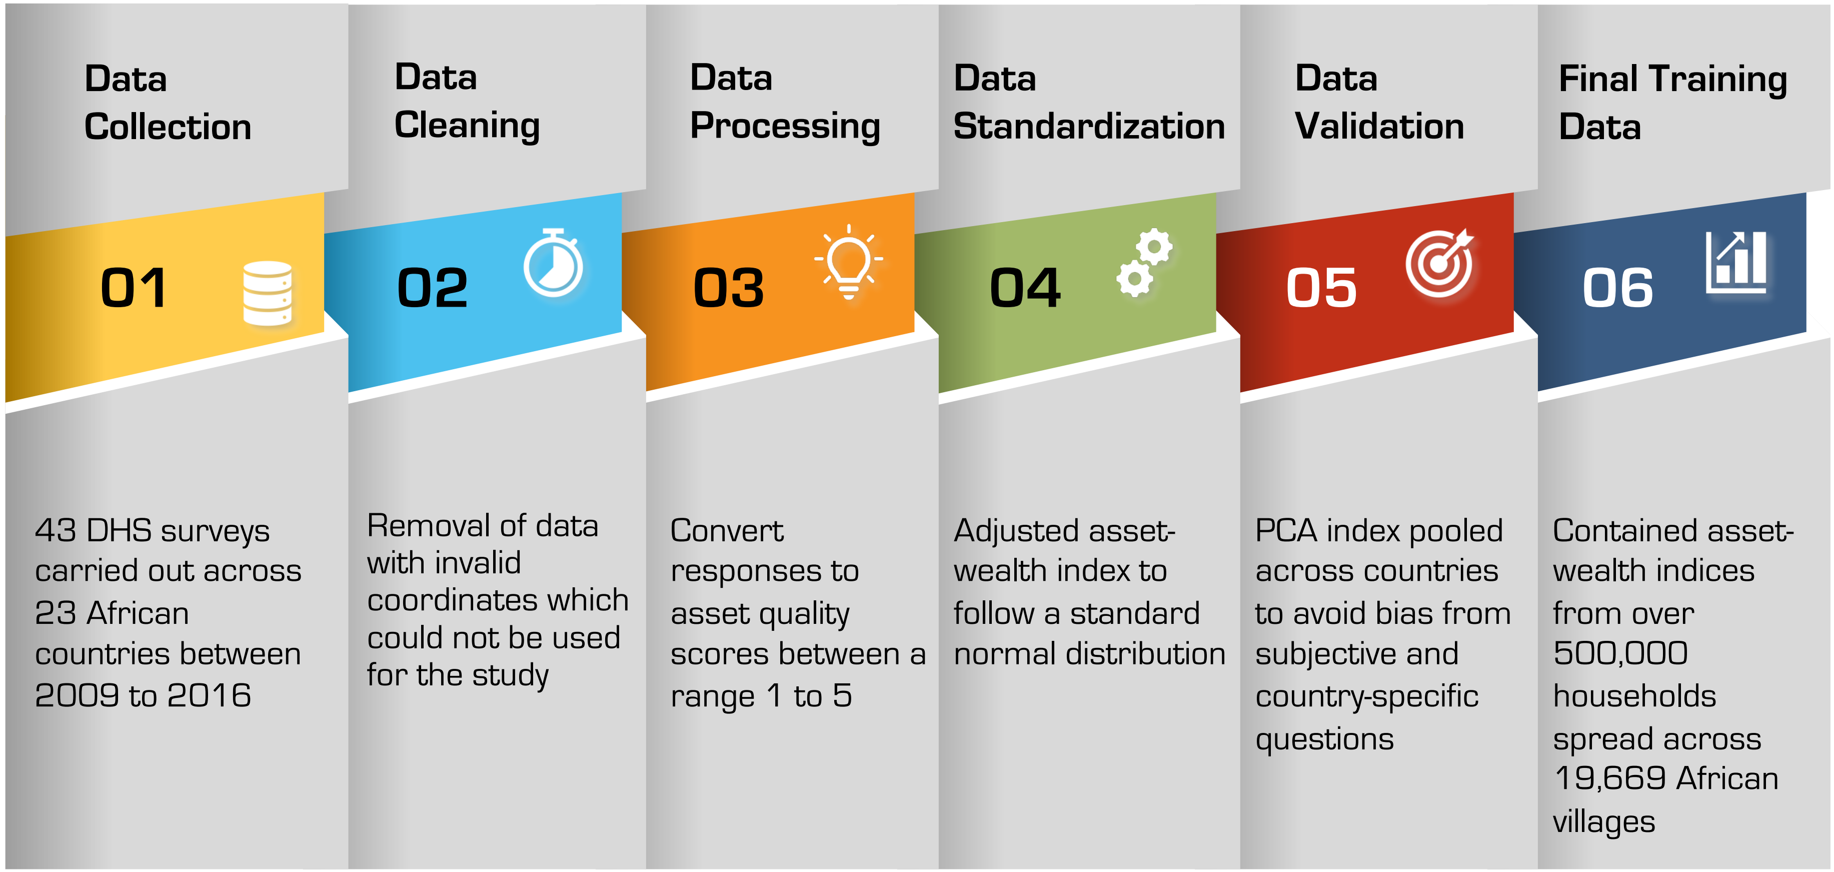

Implementation Details

For additional validation with data geolocated to the second administrative level (district or county), the asset-wealth indices obtained from the DHS surveys were compared with data from two other contexts:

- Country census data
- Living Standards Measurement Study (LSMS)

Implementation Details

Two baseline models were established

- A simple k-nearest neighbor model (KNN scalar NL) on nightlights that predicts wealth in a particular location as the average wealth over the *k* locations with nightlights values closest to the target location

- A regularized linear regression on scalar nightlights (Linear scalar NL)

Linear regression is a procedure by which we model a process that relates an independent variable, e.g., time or location, and a dependent variable, e.g., wealth, by fitting a line to the data that characterize the process. This model then allows one to predict values of the dependent variable based on the value of the independent variable, e.g., the wealth value at some required time or place.

Implementation Details

The dual-input model learns predictive features from both inputs

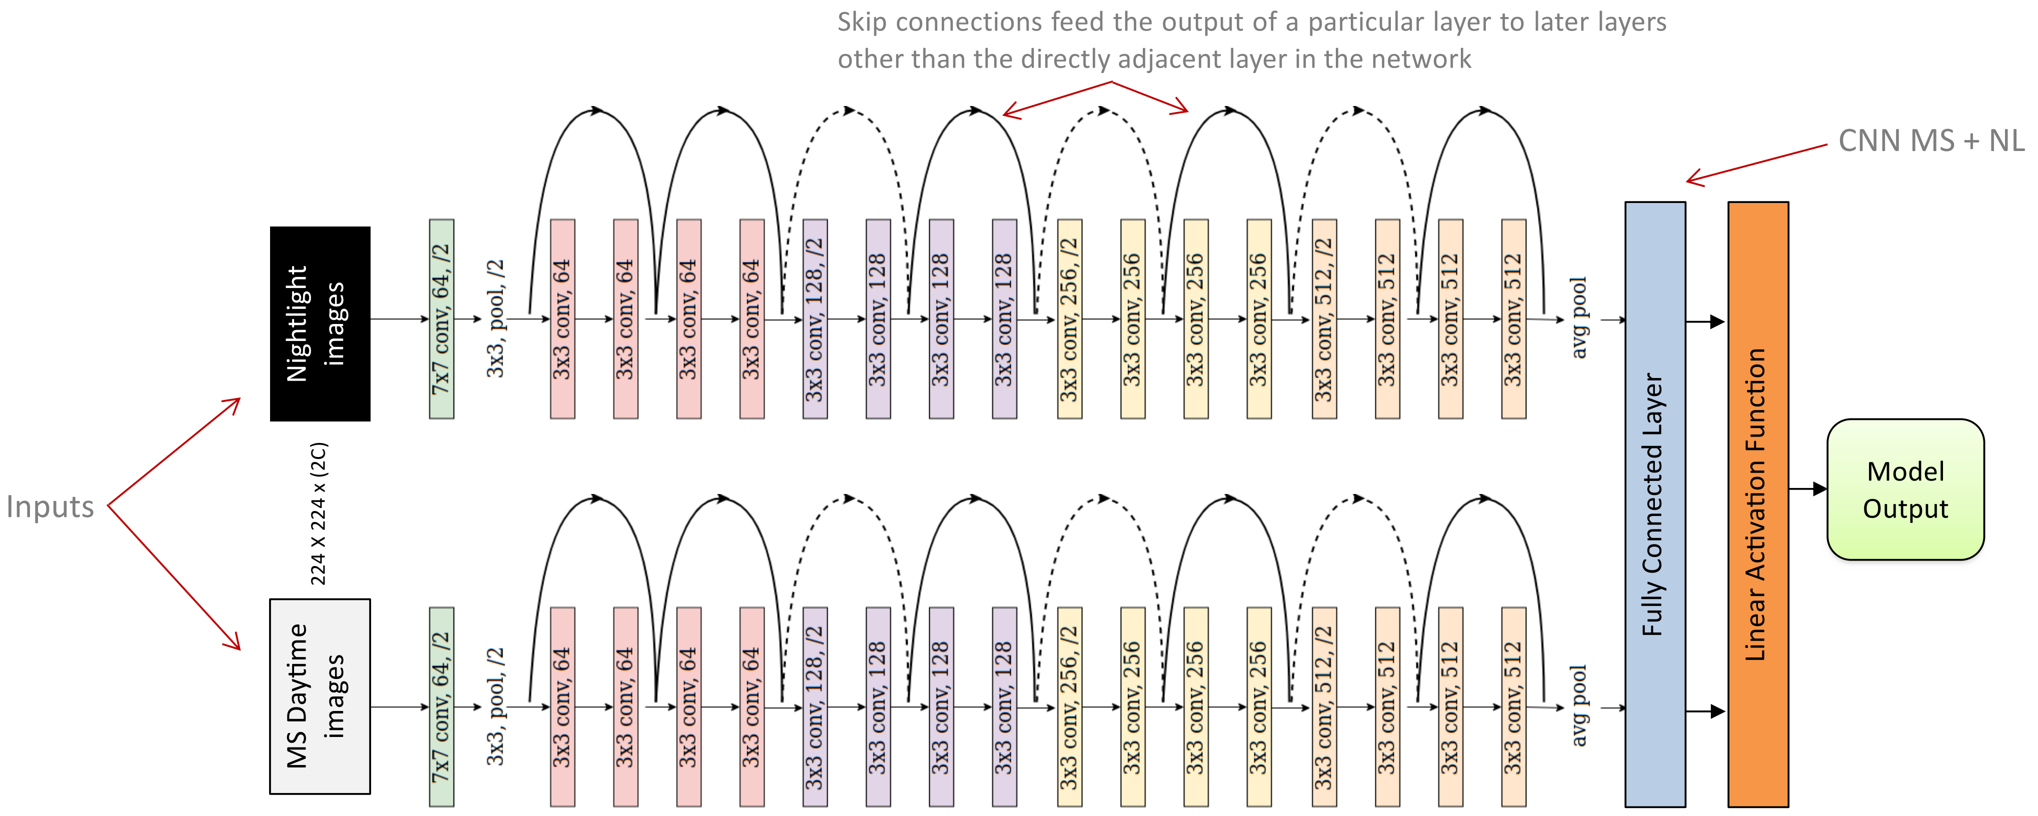

Implementation Details

Other training details:

- Adam optimizer

  Optimizers are functions or algorithms that modifies the attributes of a neural network such as weights and learning rates. Adam is an alternative to Stochastic Gradient Descent algorithm
- Mean squared-error loss function

  Loss function is used to determine how well the deep learning neural network model predicts the correct value
- Batch size of 64

  Number of samples processed before the model is updated
- Trained for 150 epochs (200 epochs for DHS out-of-country data)

  One complete processing of all the training images
- Learning rate decay factor of 0.96 after each epoch

  Reducing the learning rate to improve the model's ability to learn complex patterns
- Grid search

  To determine optimal learning rate and weight regularization
- Geometric augmentation

  Random horizontal and vertical flipping; colour based augmentation like contrast and brightness adjustment

Implementation Details

To compare performance of the dual-input model (CNN MS + NL) with current approaches, the authors also had to train models on:

- Only the daylight multispectral model (CNN MS)
- Only the nightlight model (CNN NL)
- Transfer learning approach on nightlights with Landsat imagery (CNN transfer)

Results

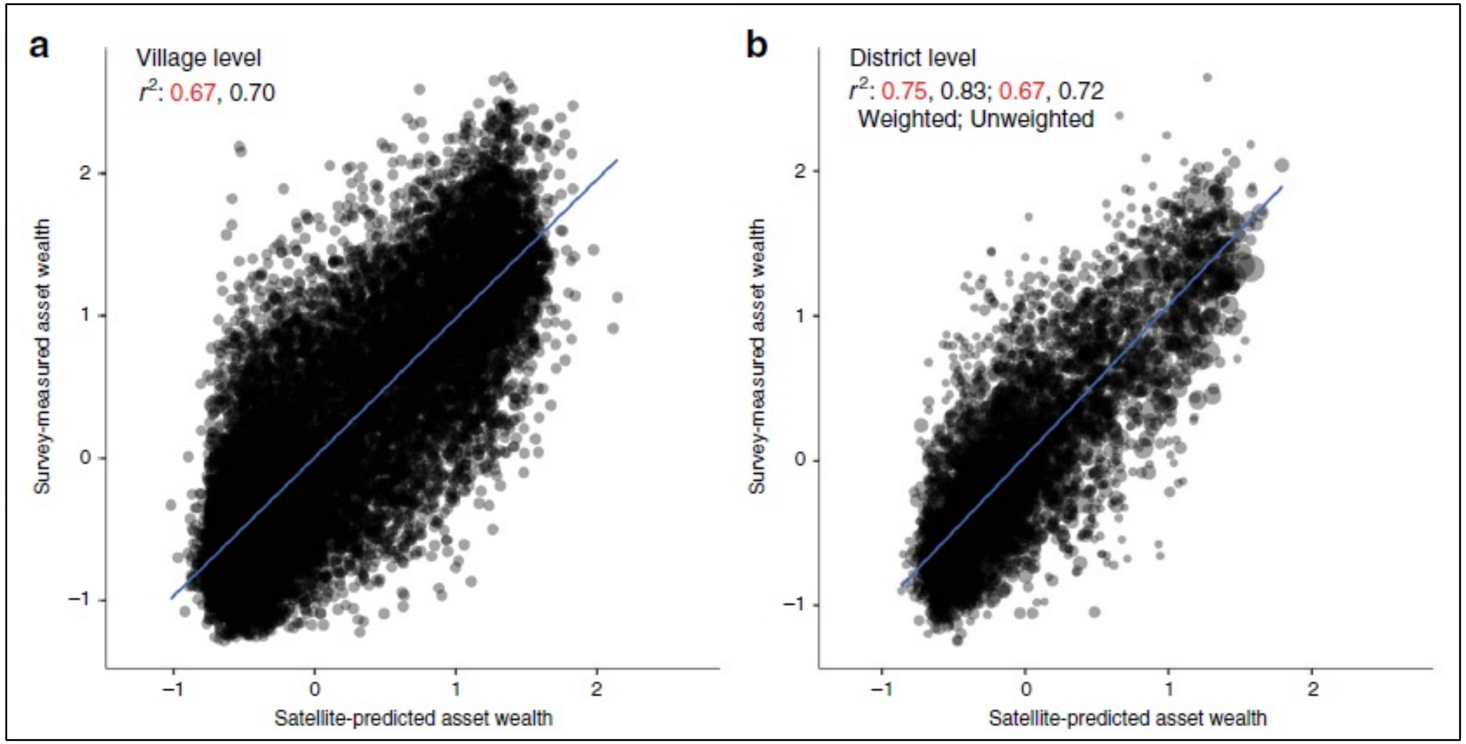

Satellite-based predictions explain the majority of variation in survey-based wealth estimates in all countries

a. Comparison of satellite-predicted wealth index to DHS survey measured wealth index across all locations for all survey years. Each point represents a village in a given survey-year, with the predictions from the dual-input model for each country from a model trained outside that country using a 5-fold cross validation. The r^2 values in red show goodness-of-fit on pooled observations while values in black show the average calculated within country-years. **b.** Shows the graph aggregated to the district level.

Results

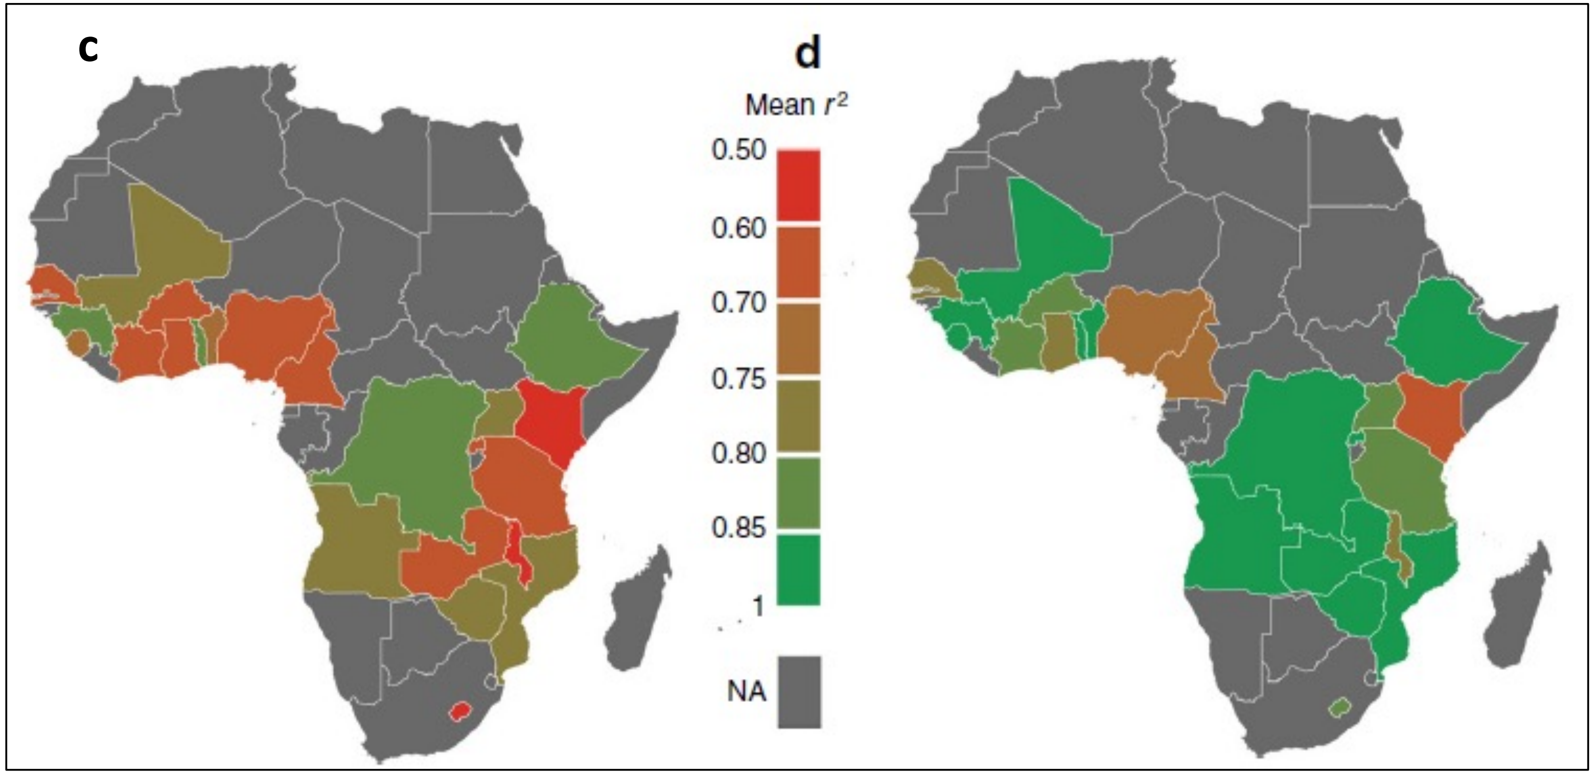

Satellite-based predictions explain the majority of variation in survey-based wealth estimates in all countries
c. and **d.** Map showing average r^2 values over survey years by country at village and district level respectively.

Results

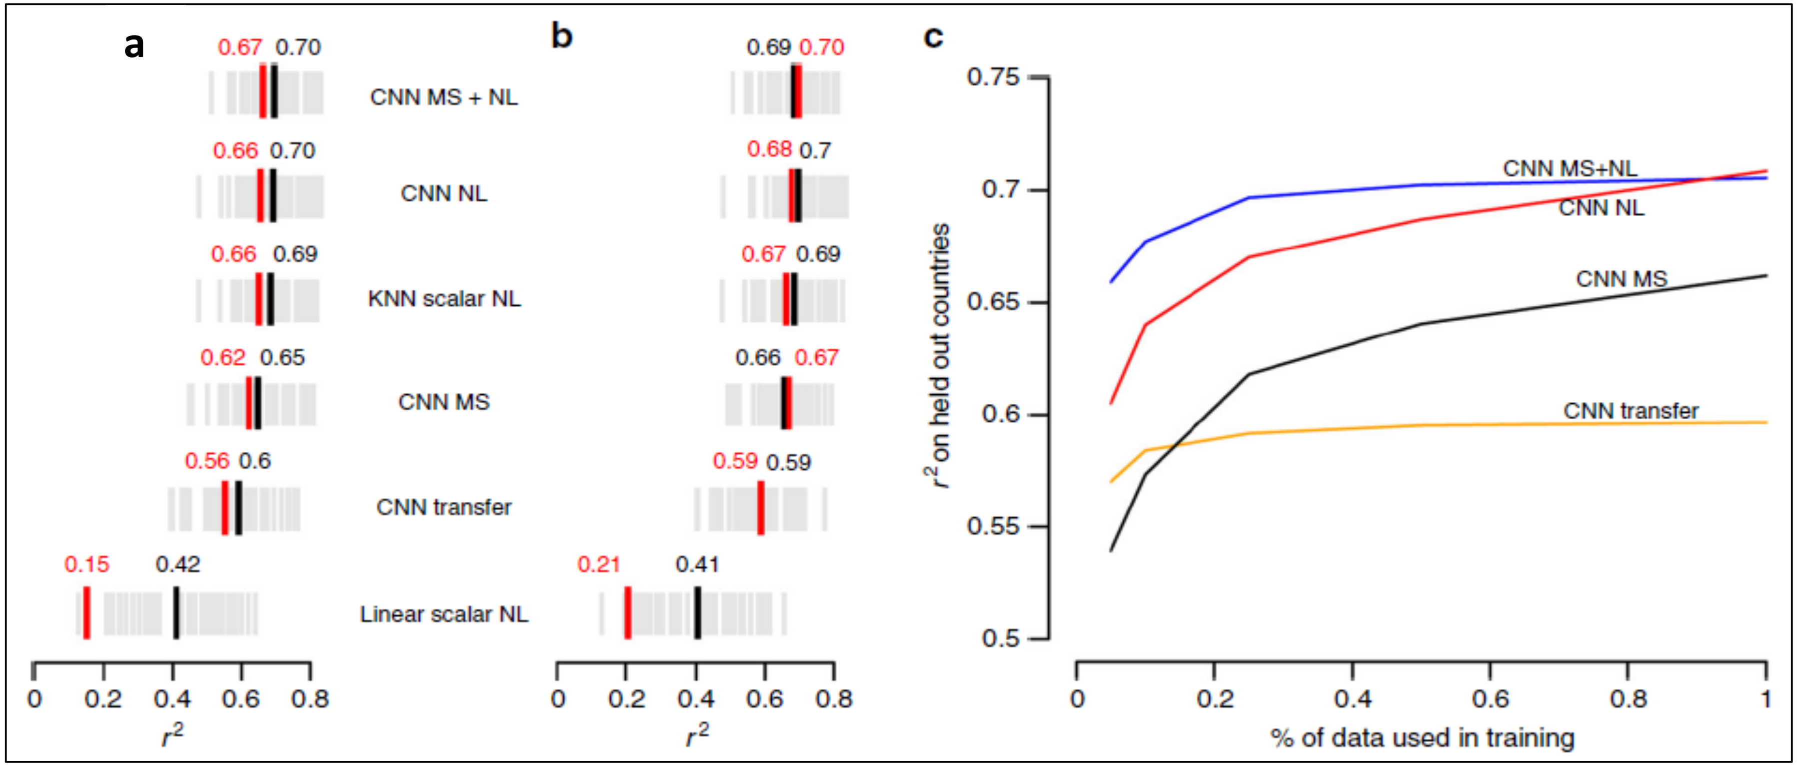

The dual-input model consistently outperforms all other approaches in predicting asset-wealth index.

a. Model performance on satellite predictions; each grey line represents the performance r2 on a held-out country year, black lines and text show average across country-years and red lines show the r2 on the pooled sample. **b.** similar a but for evaluation on held-out villages within the same country. **c.** shows model performance as a function of amount of training data used.

Strengths of the CNN ML+NL Model

1. Efficient way of obtaining or complementing the collection of data on wealth distribution in Africa

2. High accuracy means it can be applied to downstream research or policy tasks

3. Scalable in predicting wealth asset for locations where ground-truth information is not available

4. Robust in ignoring spurious data that are not indicative of wealth

Weaknesses of the CNN ML+NL Model

1. Poor performance in settings where within-village wealth variation is high

2. Weak explainability of deep neural networks may impede adoption by policy makers

Possible Options for Improvement

1. Combined use of higher-resolution images and other wealth indicators from passive sensors, e.g., mobile phones and social media

2. The issue with the lack of explainability of deep learning neural network architectures remains an open research question

Lecture Summary

1. Reliable and precise measures of economic well-being are essential inputs for effectively allocating resources based on needs, but this data is often unavailable at the local level in most developing countries

2. The approach described in this case study makes use of a dual-input model with satellite imagery to predict survey-based estimates of asset-wealth across approximately 20,000 African villages

3. The model outperforms previous benchmarks from high-resolution imagery and transfer learning, and performs comparably with independent wealth measurements from censuses by explaining 70% of the variation in ground-measured village wealth

Lecture Summary

4. For temporal changes in wealth aggregated across districts, the satellite-based estimates from the model explains 50% of the variation, with daytime imagery being especially helpful in this task

5. This solution shows that AI and machine learning can be used to address the issue of lack of reliable and up-to-date data when making economic decisions in African countries

Recommended Reading

Yeh, C., Perez, A., Driscoll, A., Azzari, G., Tang, Z., Lobell, D., . . . Burke, M. (2020). Using publicly available satellite imagery and deep learning to understand economic well-being in Africa. Nature Communications, 11 (1), 2583.

https://www.nature.com/articles/s41467-020-16185-w

References

Babenko, B., Hersh, J., Newhouse, D., Ramakrishnan, A., & Swartz, T. (2017). Poverty mapping using convolutional neural networks trained on high and medium resolution satellite images, with an application in Mexico. In Proc. NIPS 2017 Workshop on Machine Learning for the Developing World (pp. 994–1028-794).

https://arxiv.org/abs/1711.06323

Henderson, J. V., Storeygard, A., and Weil, D. N. (2012). Measuring economic growth from outer space. The American Economic Review, 102 (2), 994–1028-794.

https://doi.org/10.1257/aer.102.2.994

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), 790-794.

https://www.science.org/doi/abs/10.1126/science.aaf7894