



# Predicting poverty from satellite imagery

---

Neal Jean, Michael Xie, Stefano Ermon  
Department of Computer Science  
Stanford University

Matt Davis, Marshall Burke, David Lobell  
Department of Earth Systems Science  
Stanford University

# Why poverty?

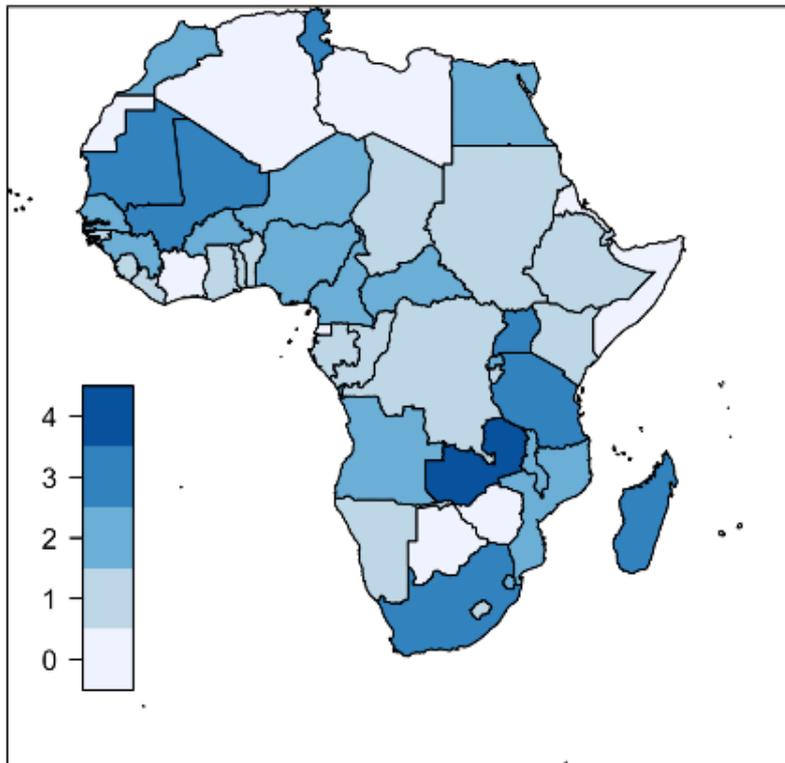


- **#1** UN Sustainable Development Goal
  - Global poverty line: **\$1.90/person/day**
- Understanding poverty can lead to:
  - Informed policy-making
  - Targeted NGO and aid efforts

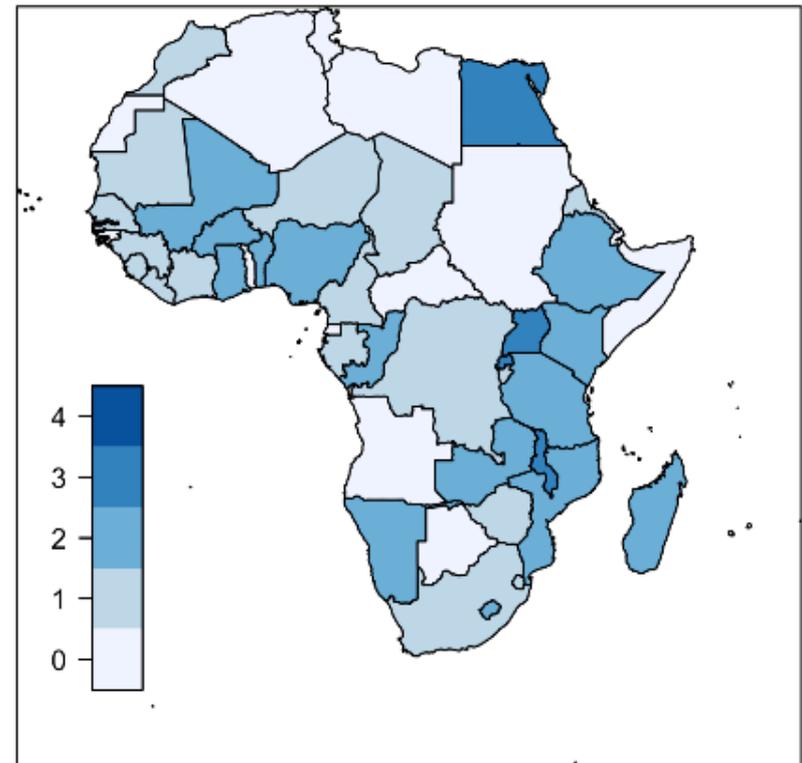
# Data scarcity



Consumption/Income Survey Availability, 2000-2010



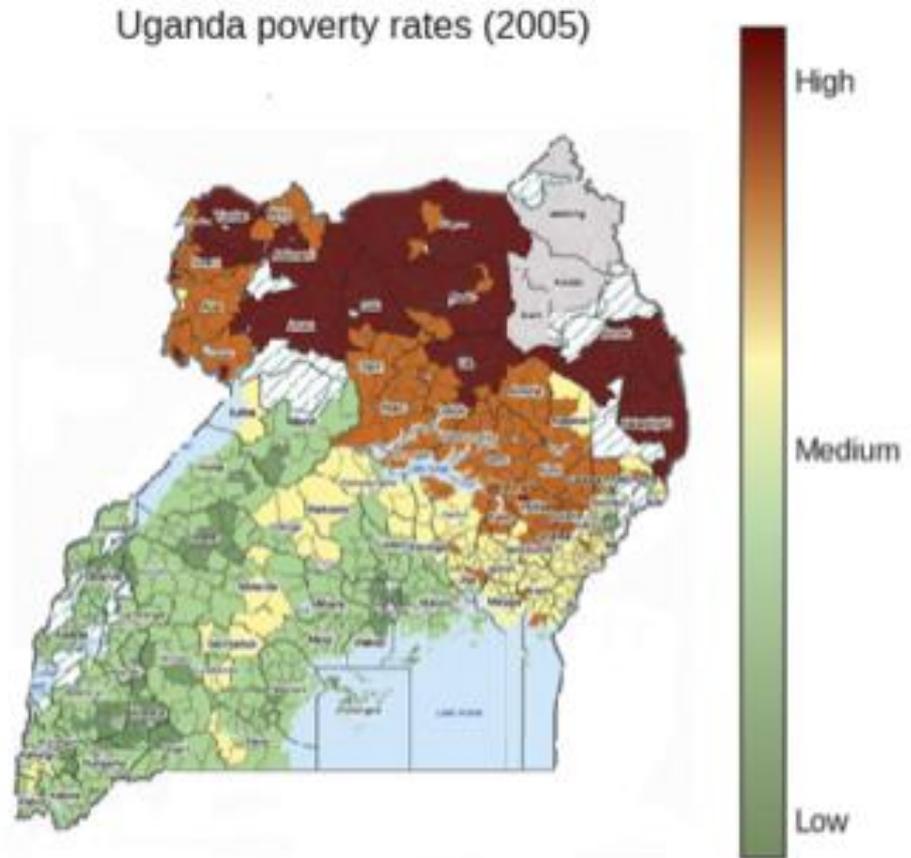
Wealth Survey Availability, 2000-2010



# Lack of quality data is a huge challenge



- **Expensive** to conduct surveys:
  - \$400,000 to \$1.5 million
- **Data scarcity:**
  - <0.01% of total households covered by surveys
- **Poor** spatial and temporal **resolution**



# Satellite imagery is low-cost and globally available

---



Simultaneously becoming **cheaper** and **higher resolution**  
(DigitalGlobe, Planet Labs, Skybox, etc.)

# What if...

---

we could **infer** socioeconomic indicators from large-scale, remotely-sensed data?

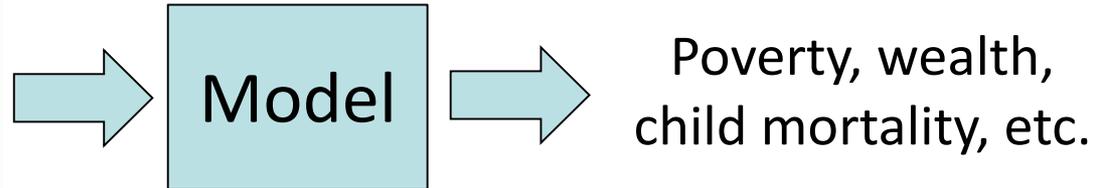
# Standard supervised learning won't work

---

Input



Output



- **Very little training data** (few thousand data points)
- **Nontrivial for humans** (hard to crowdsource labels)

# Transfer learning overcomes data scarcity

**Transfer learning:** Use knowledge gained from one task to solve a different (but related) task

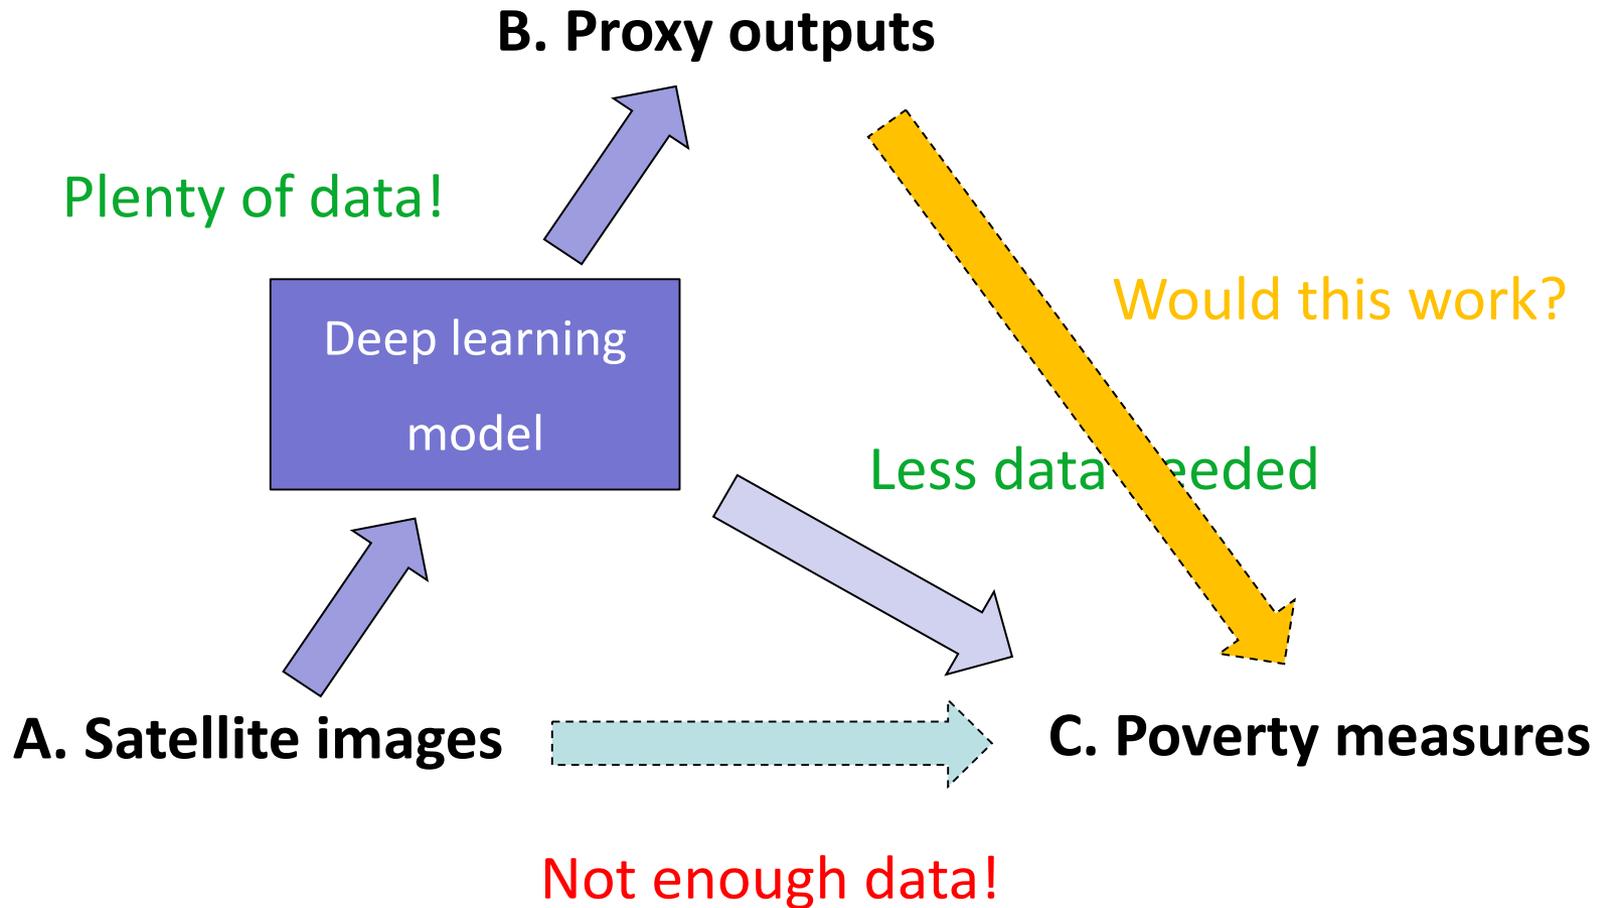
Train here



Perform here



# Transfer learning bridges the data gap



# Nighttime lights as proxy for economic development

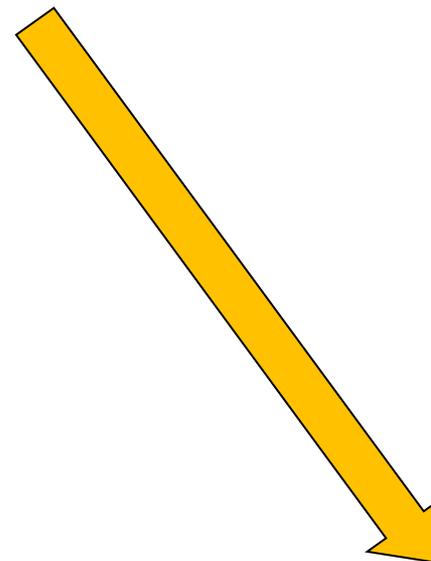


# Why not use nightlights directly?

---



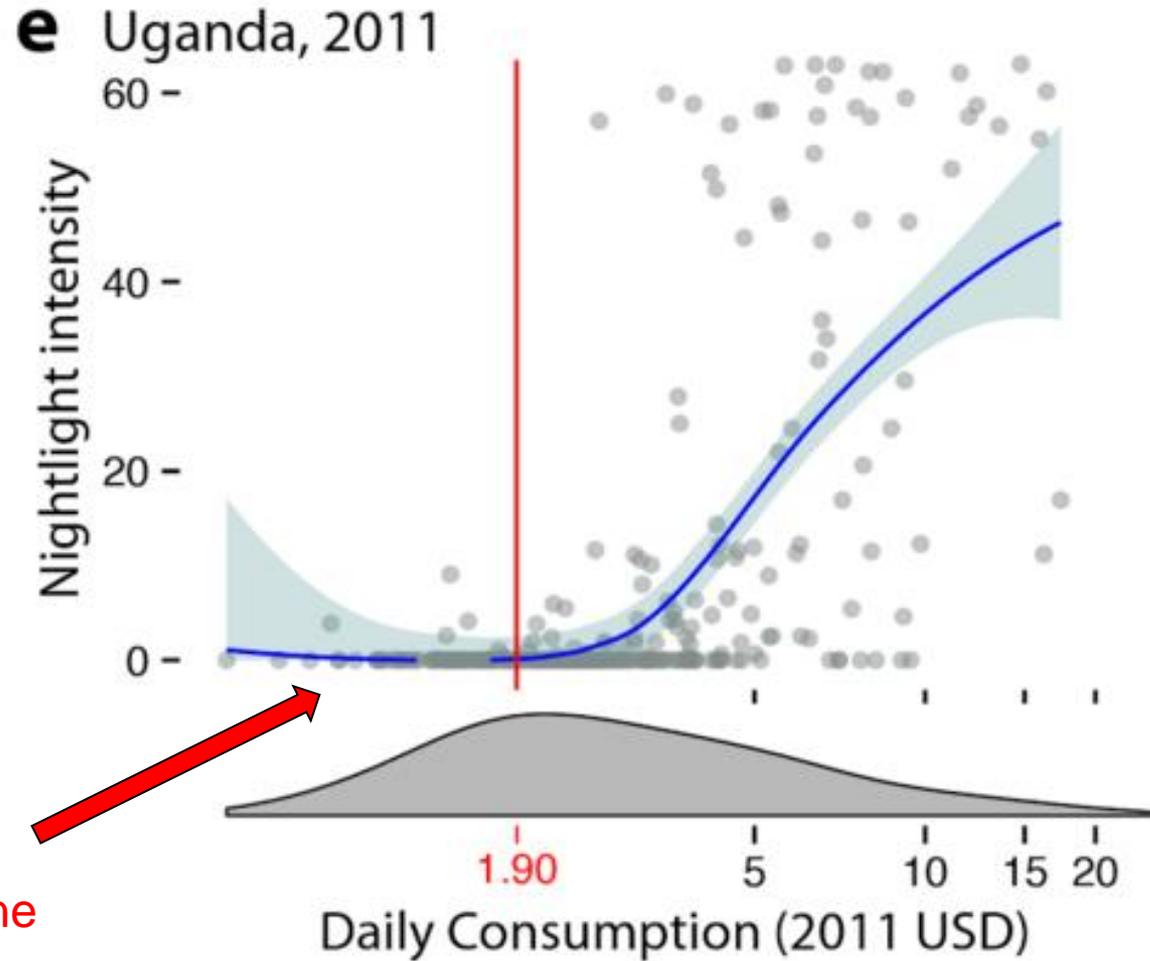
**B. Nighttime light intensities**



A. Satellite images

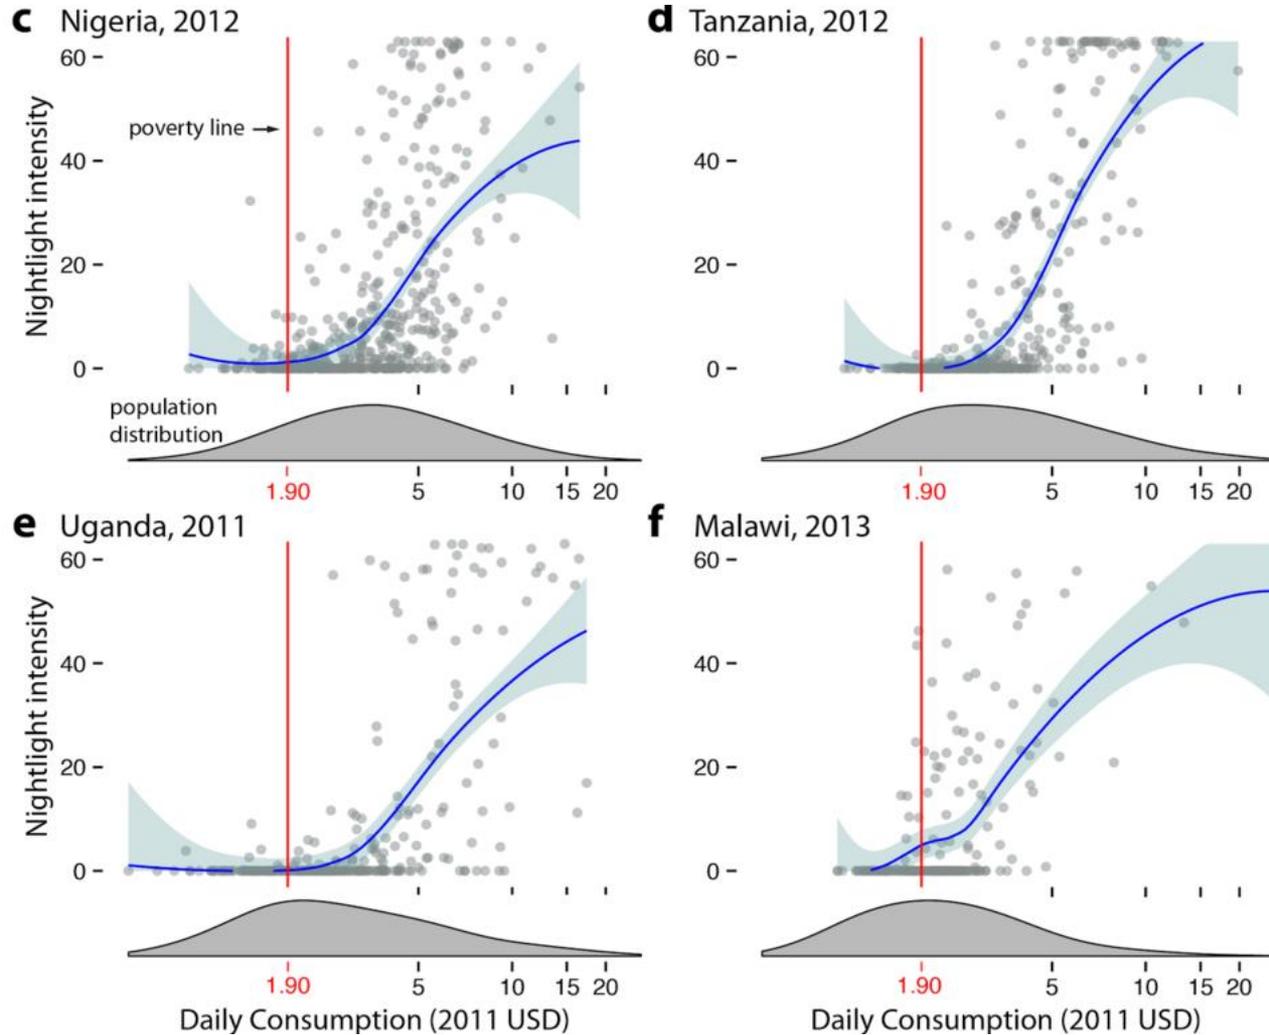
**C. Poverty measures**

# Not so fast...



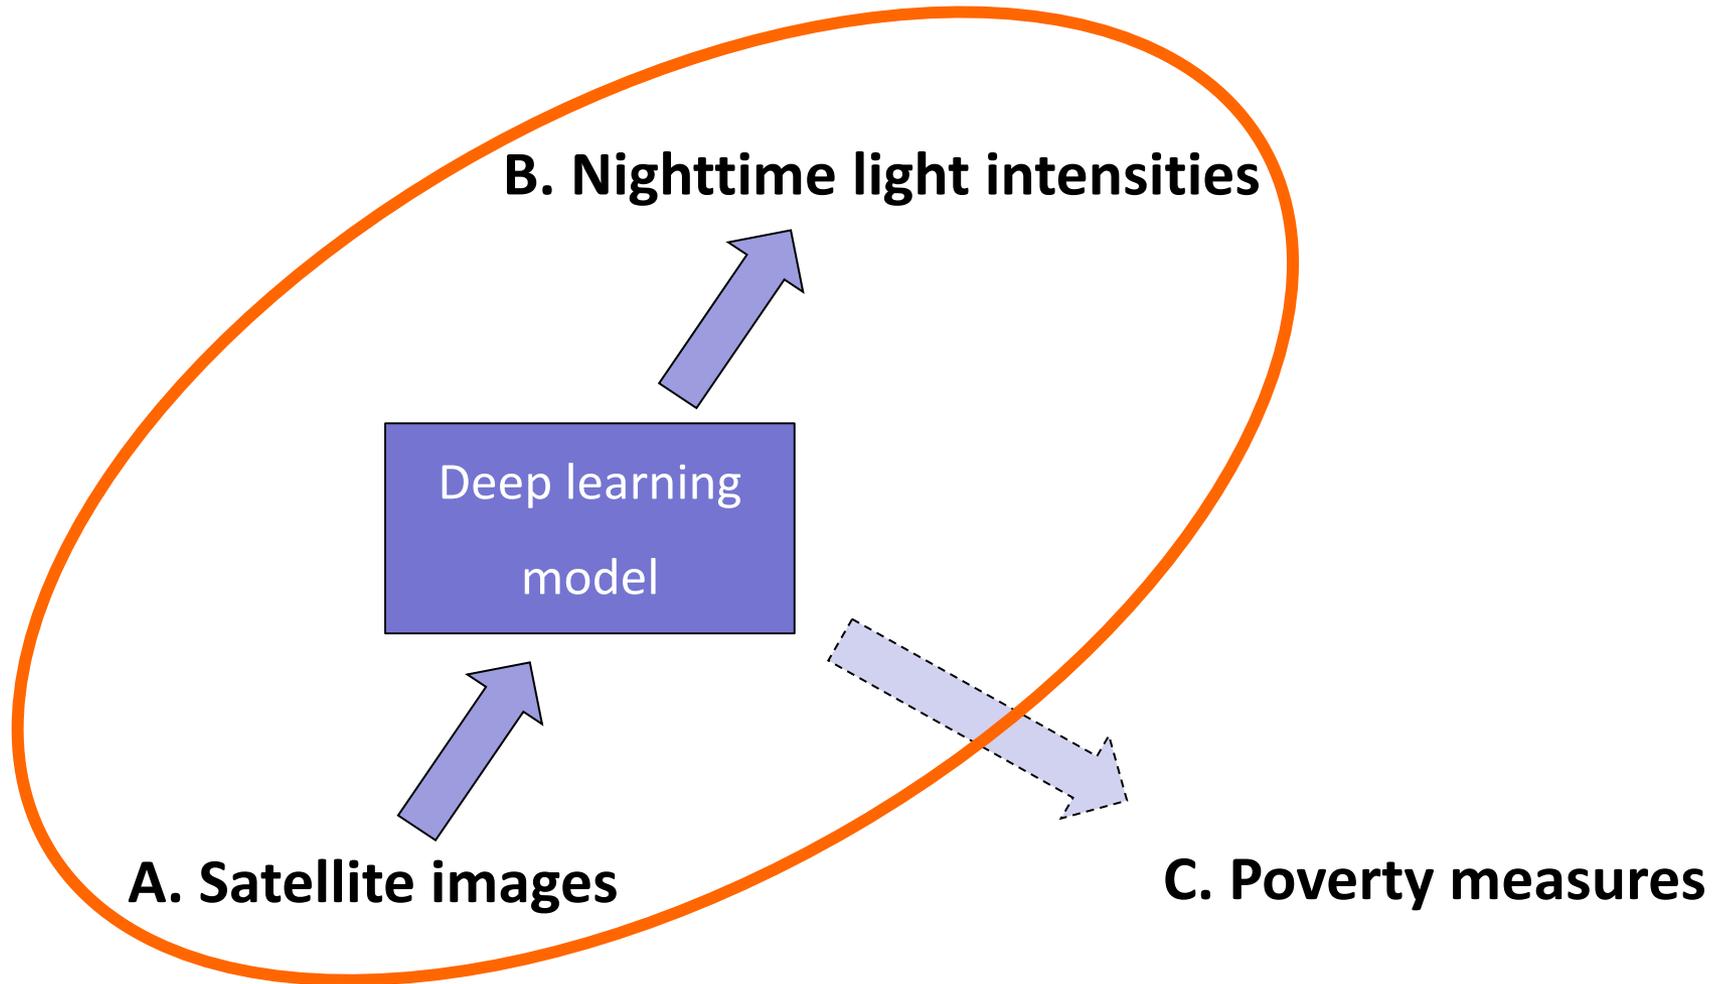
Almost no variation below the poverty line

# Lights aren't useful for helping the poorest



# Step 1: Predict nighttime light intensities

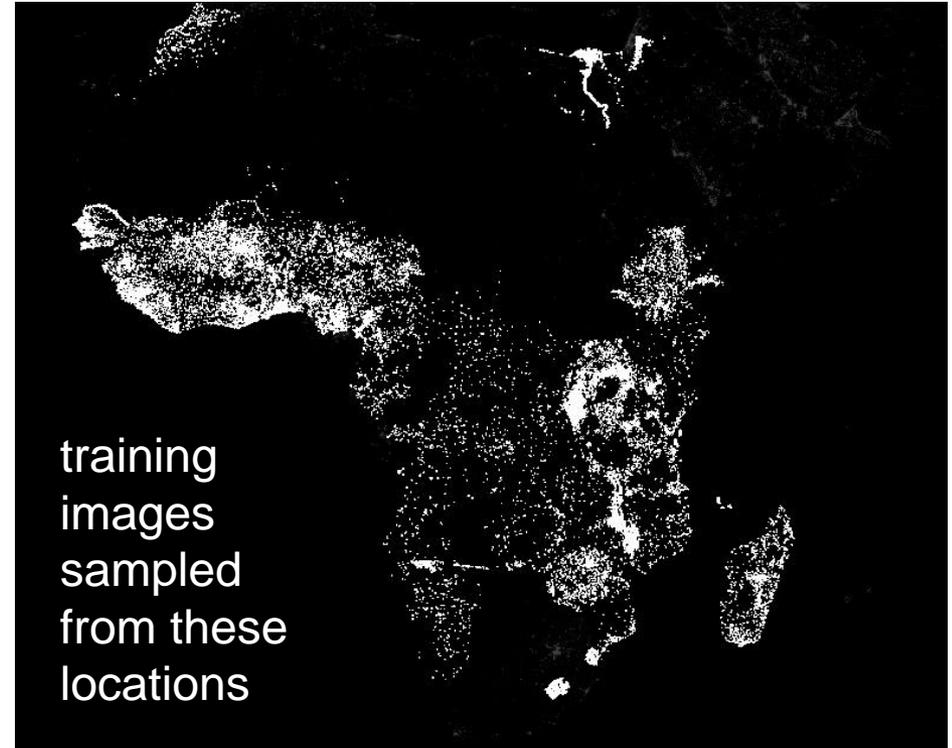
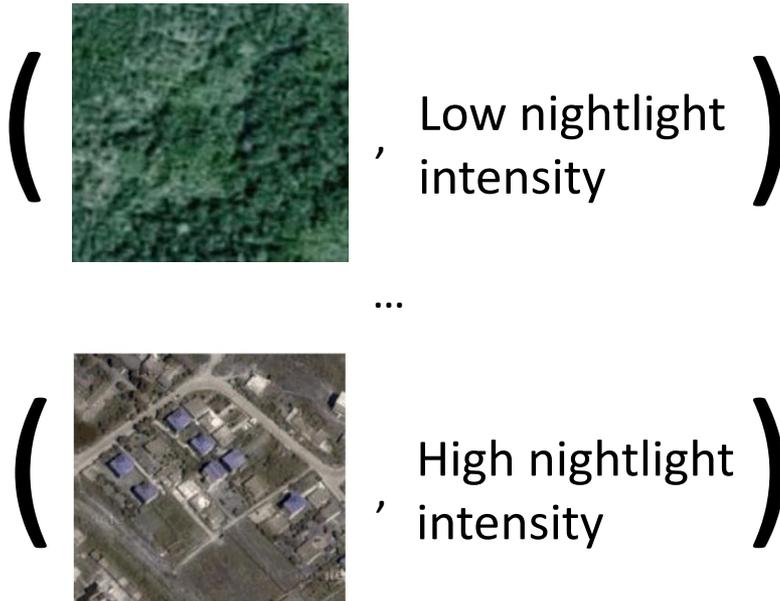
---



# Training data on the proxy task is plentiful

---

## Labeled input/output training pairs



Millions of training images

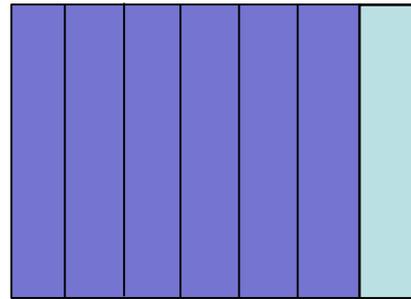
# Images summarized as low-dimensional feature vectors



**Inputs:** daytime satellite images

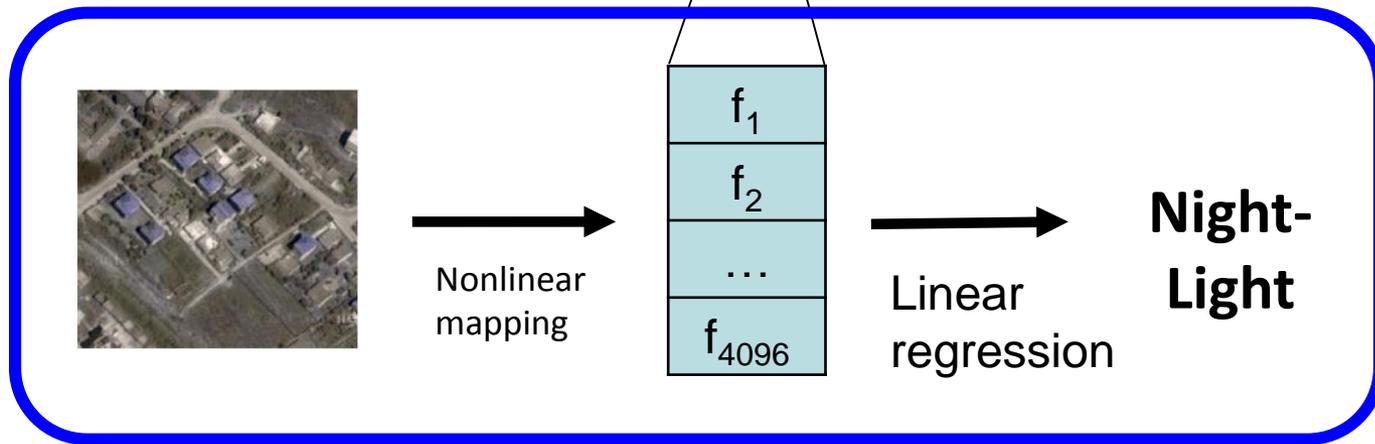


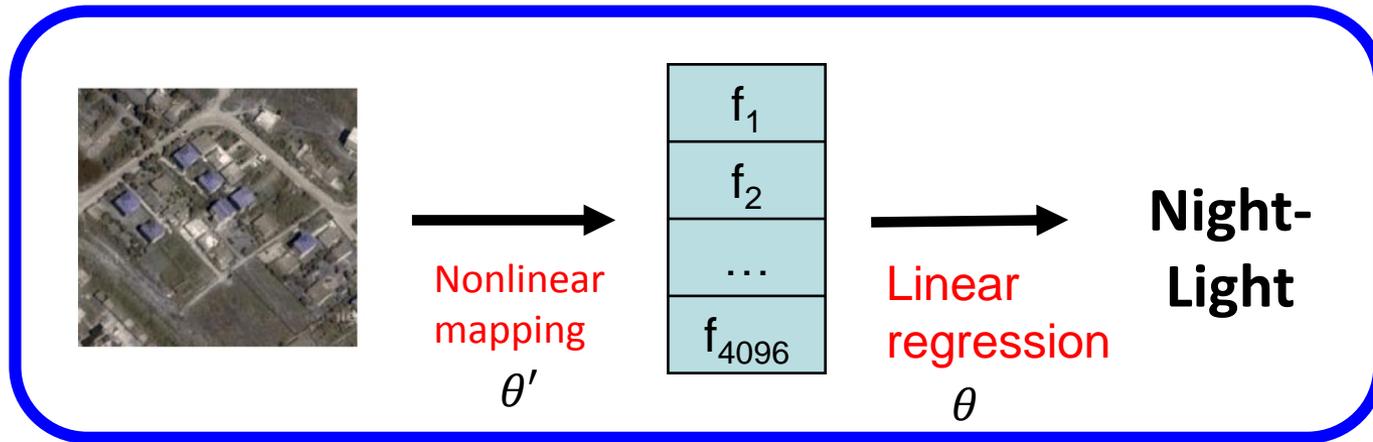
**Convolutional Neural Network (CNN)**



**Outputs:** Nighttime light intensities

{Low, Medium, High}





$$\min_{\theta, \theta'} \sum_{i=1}^m l(y_i, \hat{y}_i) = \min_{\theta, \theta'} \sum_{i=1}^m l(y_i, \theta^T f(x_i, \theta'))$$

Over 50 million parameters to fit

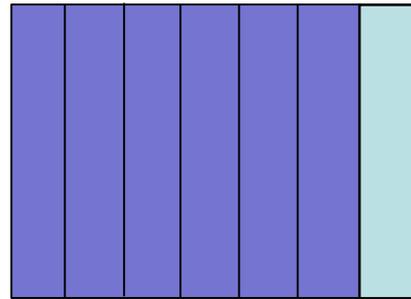
Run gradient descent for a few days

# Transfer Learning

**Inputs:** daytime satellite images

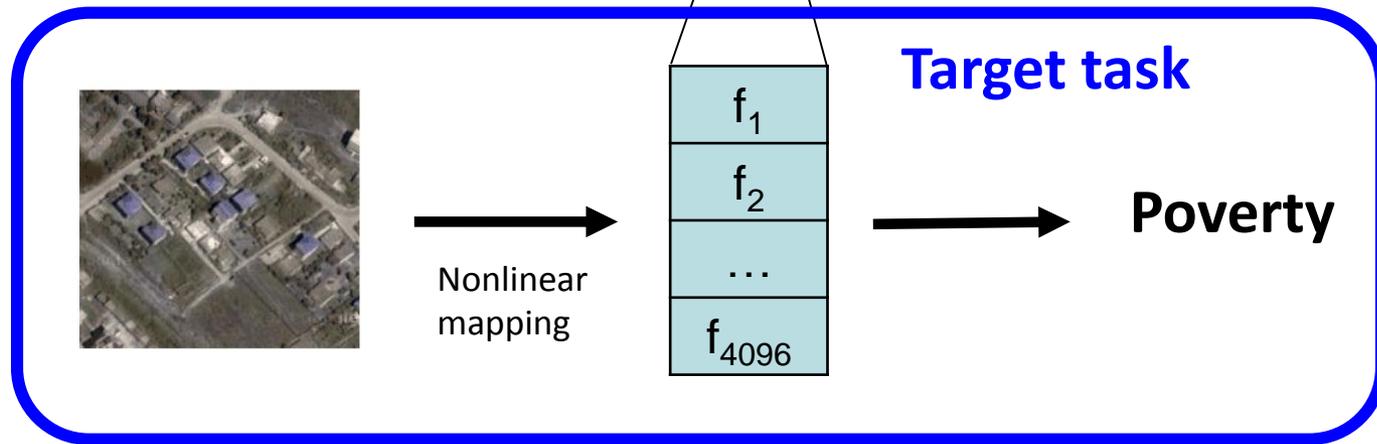


**Feature Learning**



**Outputs:** Nighttime light intensities

{Low, Medium, High}



**Have we learned to identify useful features?**

# Model learns relevant features automatically



**Satellite image**



**Filter activation map**

**Overlaid image**

# Target task: Binary poverty classification

---

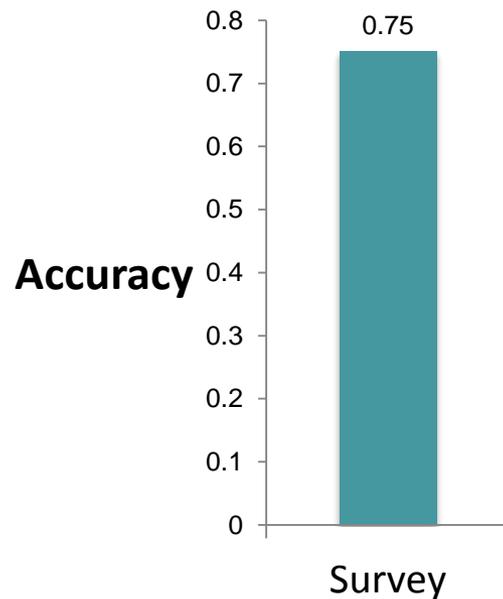
- Living Standards Measurement Study (LSMS) data in Uganda (World Bank)
  - Collected data on household features
    - Roof type, number of rooms, distance to major road, etc.
  - Report household consumption expenditures
- **Task**: Predict if the majority of households in a cluster are above or below the poverty line

# How does our model compare?



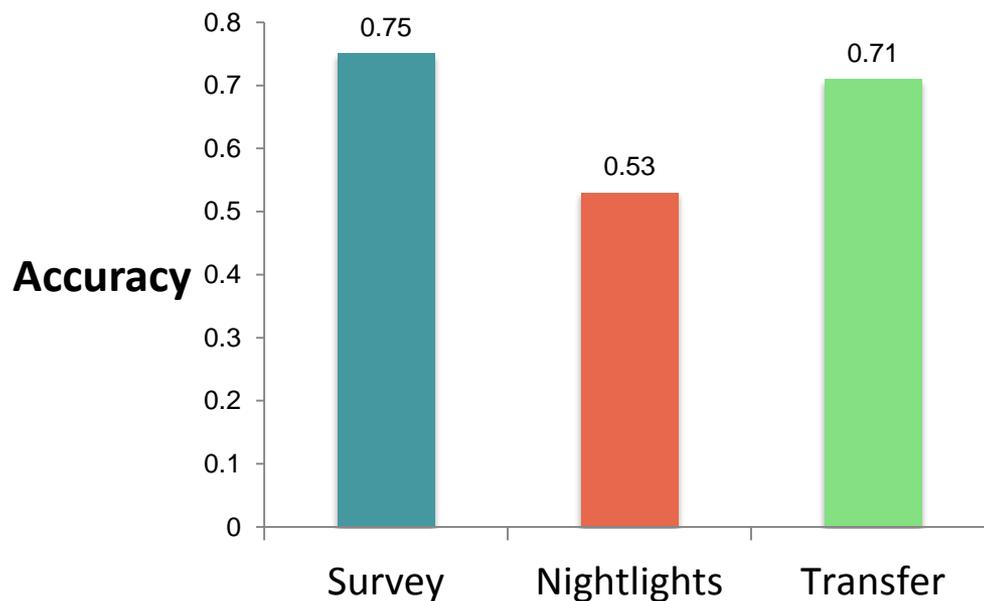
Survey-based model is the gold standard for accuracy but...

- Relies on expensively collected data
- Is difficult to scale, not comprehensive in coverage



## Advantages of transfer learning approach:

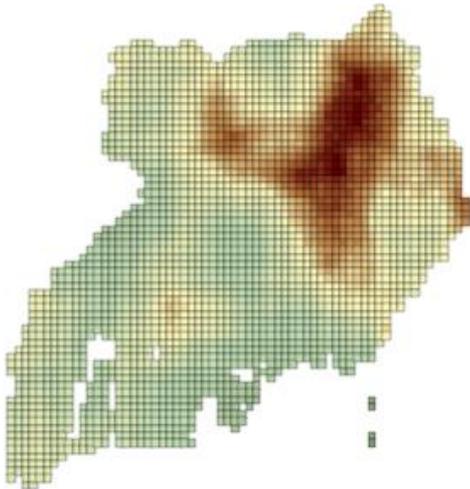
- Relies on inexpensive, publicly available data
- Globally scalable, doesn't require unifying disparate datasets



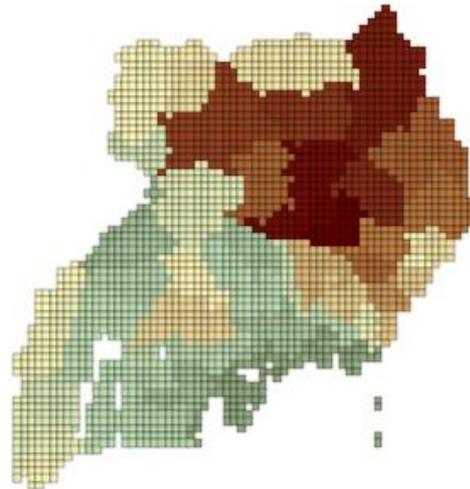
# Our model maps poverty at high resolution



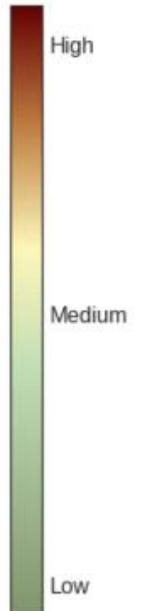
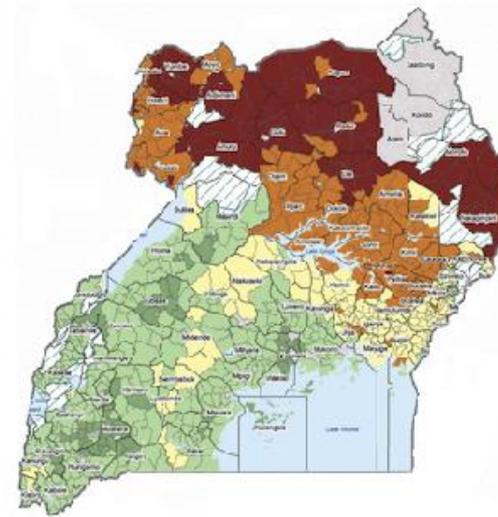
Smoothed predictions



District aggregated



Official map (2005)



## Case study: Uganda

- Most recent poverty map over a decade old
- Lack of ground truth highlights need for more data

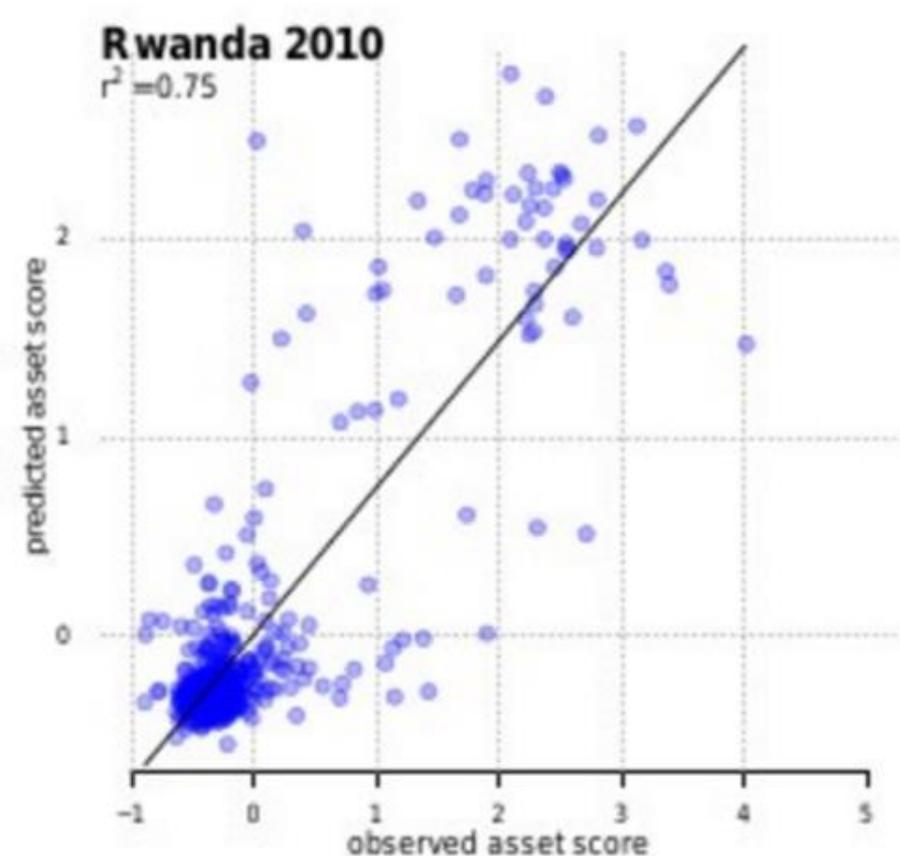
# We can differentiate different levels of poverty

## 2 continuous measures of wealth:

- Consumption expenditures



- Household assets



**We outperform recent methods  
based on mobile call record data**

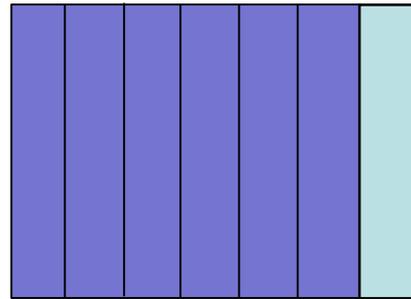
Blumenstock et al. (2015) Predicting Poverty and Wealth from Mobile Phone Metadata, *Science*

# Transfer Learning

**Inputs:** daytime satellite images

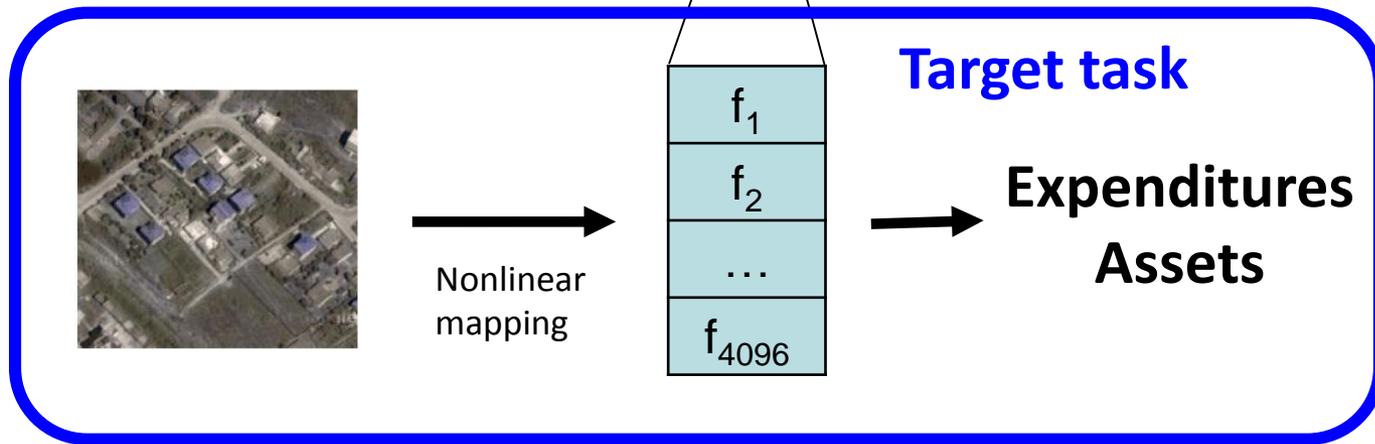


**Feature Learning**

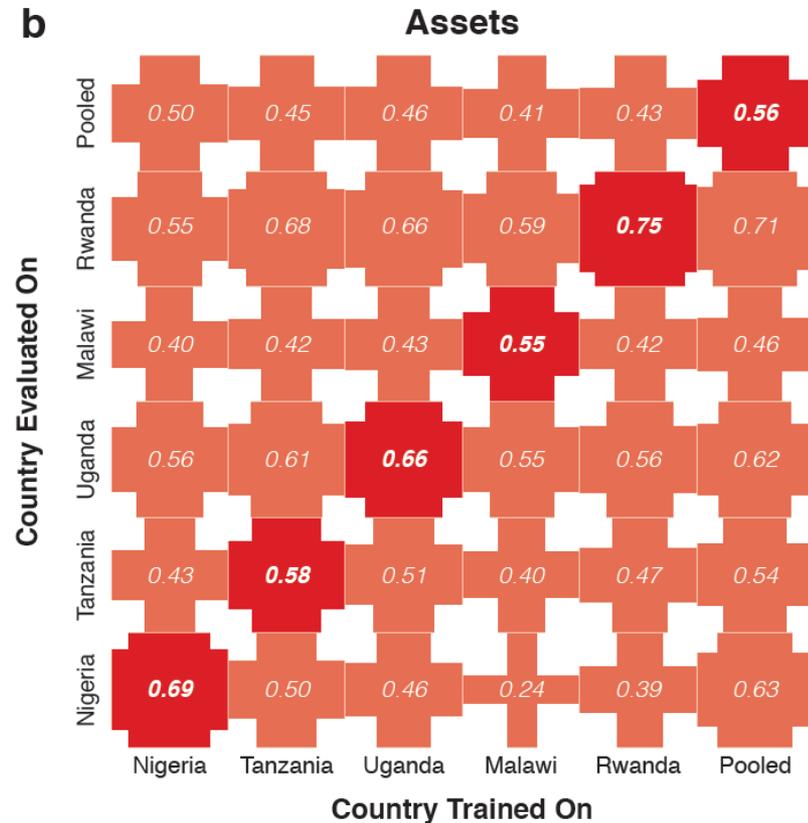
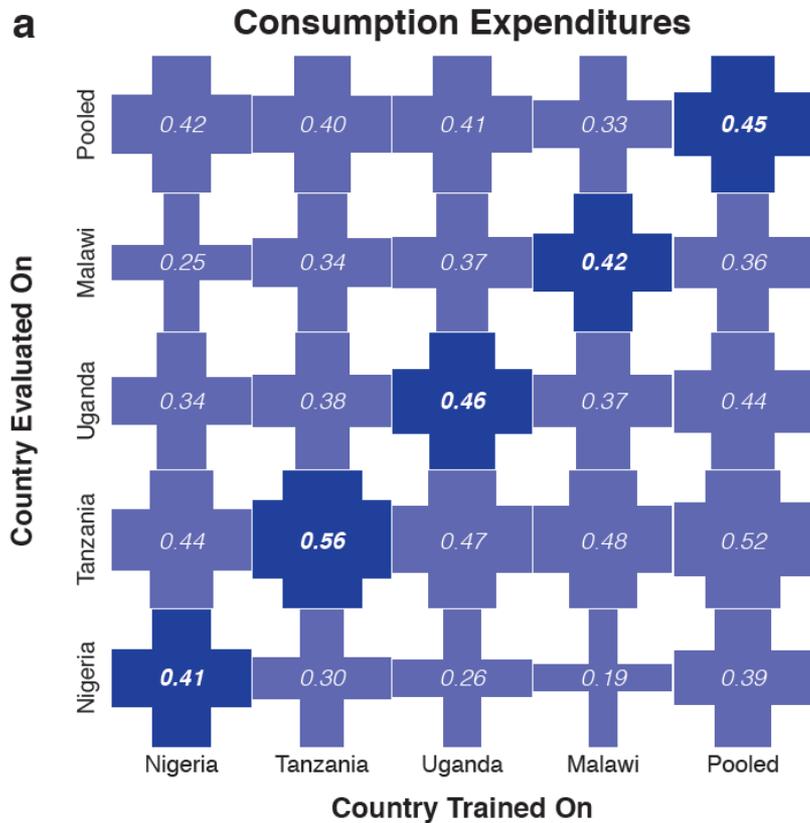


**Outputs:** Nighttime light intensities

{Low, Medium, High}



# Models travels well across borders



Models trained in one country perform well in other countries



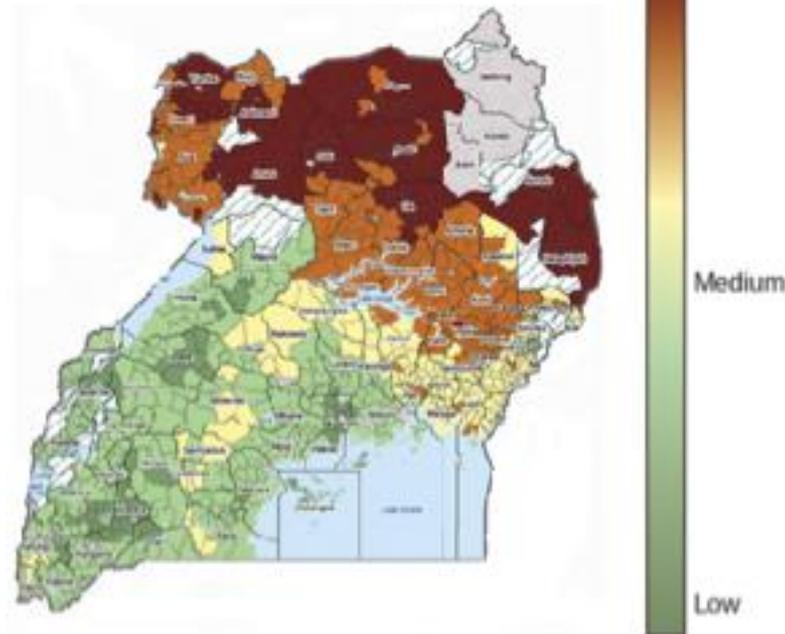
Can make predictions in countries where no data exists at all

# What do we still need?



- Develop models that account for spatial and temporal dependencies of poverty and health

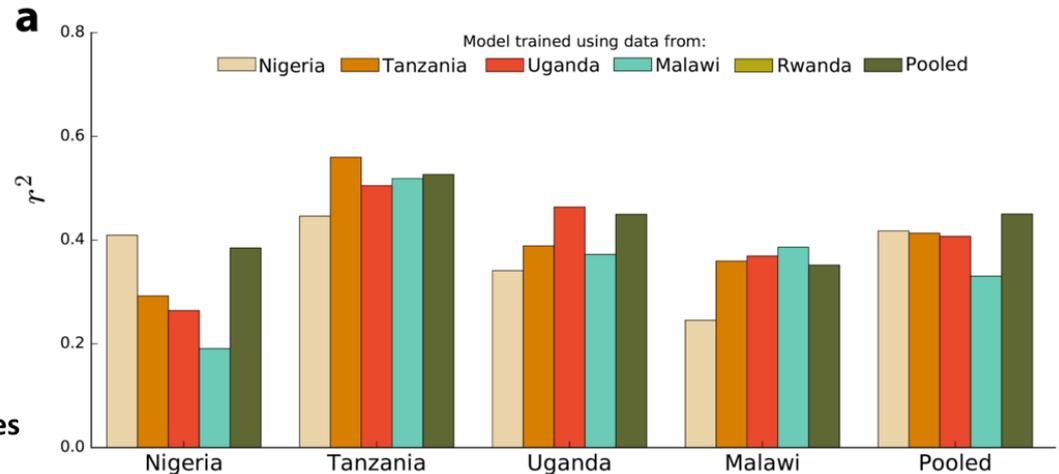
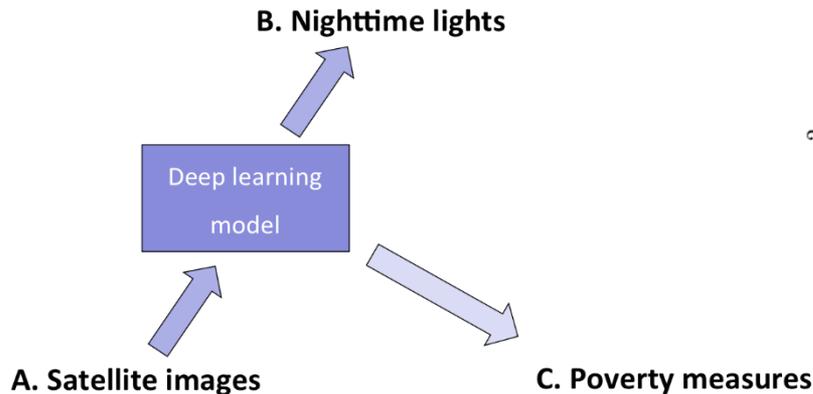
Official map (2005)



# Take full advantage of incredible richness of images



We have introduced an **accurate**, **inexpensive**, and **scalable** approach to predicting poverty and wealth



## Satellite Images Can Pinpoint Poverty Where Surveys Can't

Economic View

By SENDHIL MULLAINATHAN APRIL 1, 2016

