Transfer Learning from Deep Features for Remote Sensing and Poverty Mapping

Michael Xie, Neal Jean, Stefano Ermon

Department of Computer Science, Stanford University

Marshall Burke, David Lobell

Department of Earth Systems Science, Stanford University

AAAI-16

Why Poverty Mapping?

- #1 (of 17) UN Sustainable Development Goals
- Targeted non-profit efforts (GiveDirectly)
- Informed policy-making
- Understand poverty dynamics

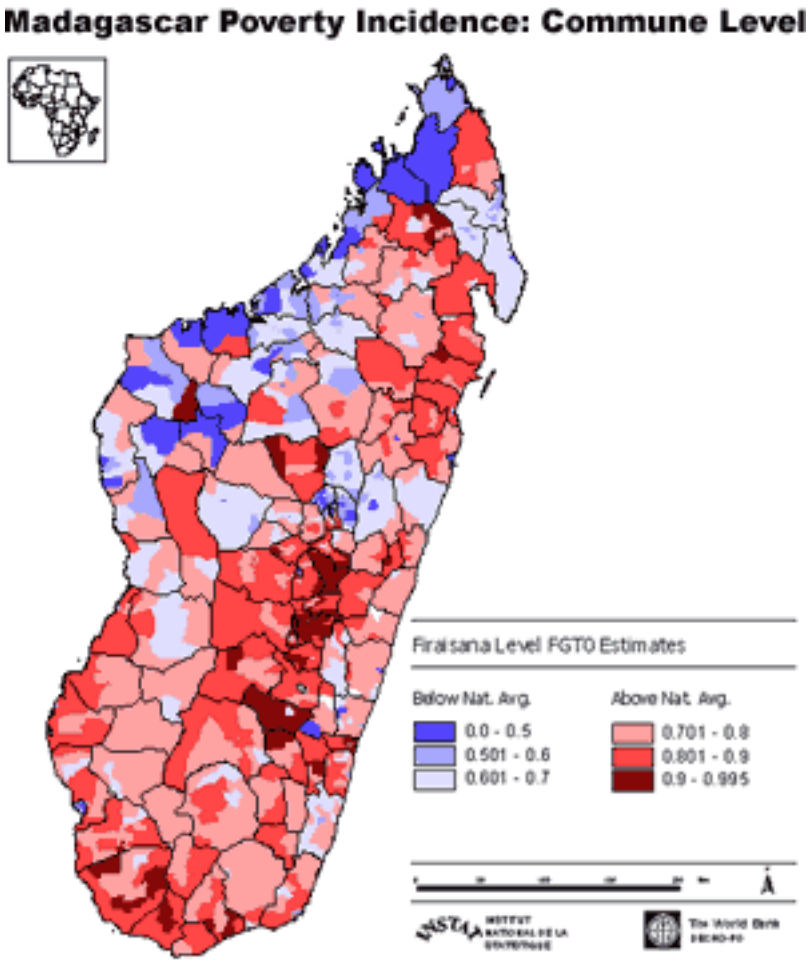

Mistiaen (2002), "Putting Welfare on the Map in Madagascar." World Bank.

Challenges with Poverty

- **Expensive** to conduct surveys ($400,000 to $1.5 million)
- **Poor** spatial and temporal **resolution** (Uganda dataset with 2,716 households)
- **Data scarcity**: lack of ground truth

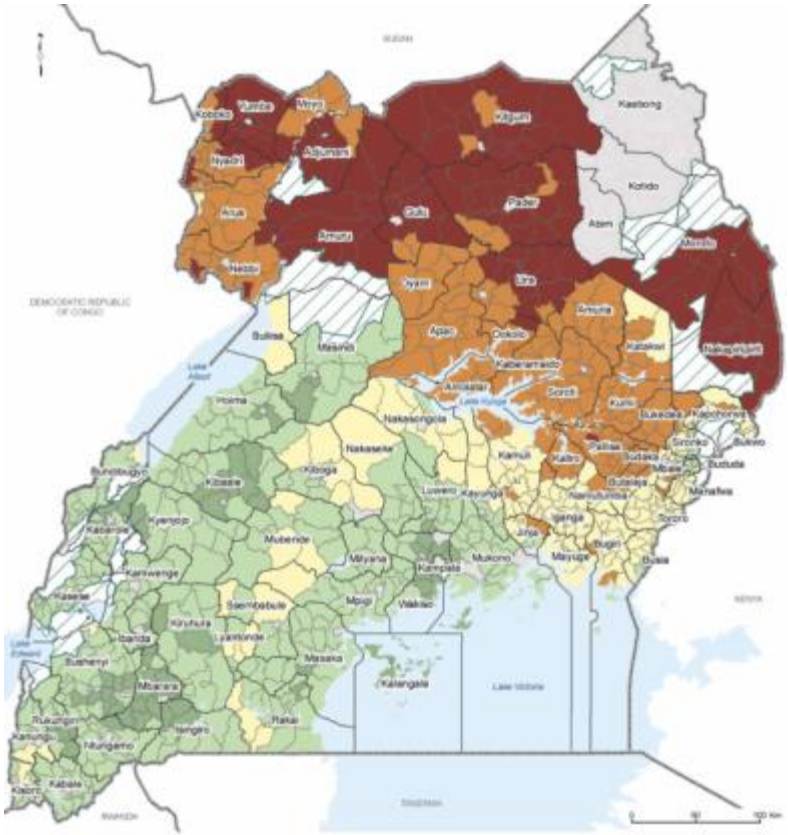

World Resources Institute (2009). “Mapping a better future: How spatial analysis can benefit wetlands and reduce poverty in Uganda.”

Remote Sensing

- Remote sensing (e.g. satellite imagery) is cost-effective and global-scale

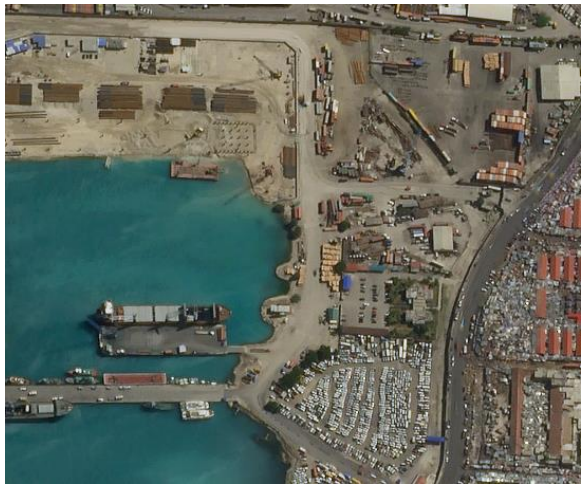

All satellite images from Google Static Maps.

- **Increasingly accurate** and **cheap** (DigitalGlobe, PlanetLabs, SkyBox).

Can we infer socioeconomic indicators (poverty, child mortality, etc.) from large-scale remotely sensed data?

Low-Cost, High-Resolution Remote Sensing

Infer poverty measures from satellite imagery

Example:  vs. 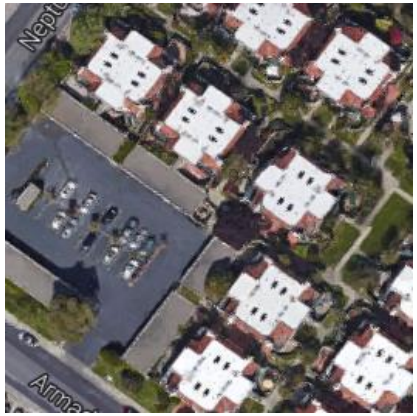

Do this **at scale, accurately, cost-effectively** and with **unprecedented spatial resolution**:

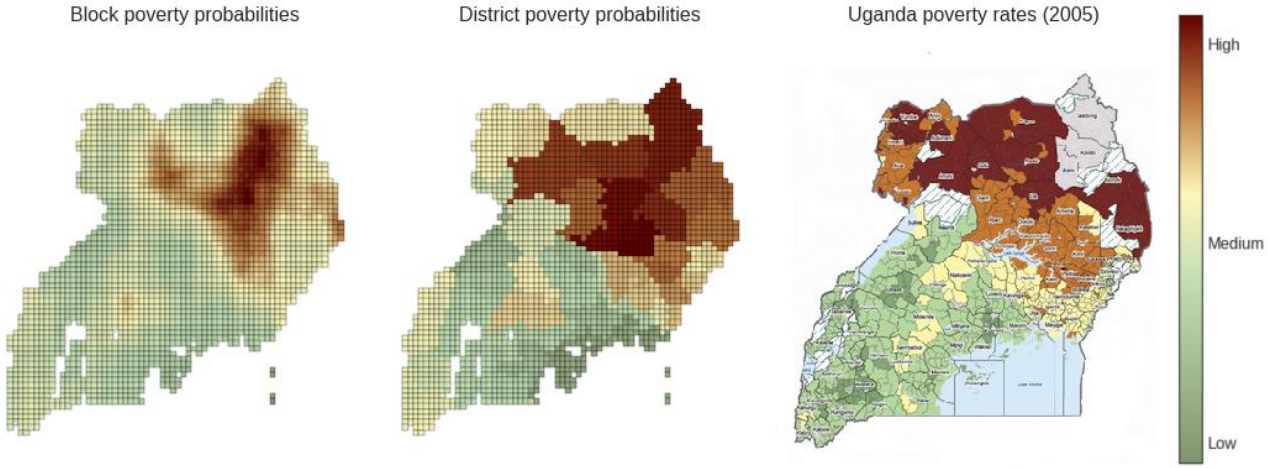

Addressing Data Scarcity

→ Socio-economic outcomes

- Standard supervised approach cannot be directly applied:
  - Very little training data (few thousand data points)
  - Nontrivial for humans (hard to crowdsource labels)

- **Transfer learning:** store knowledge gained while solving one learning problem and use it to solve a different (but related) learning problem

Night-time Lights: a Proxy for Economic Activities

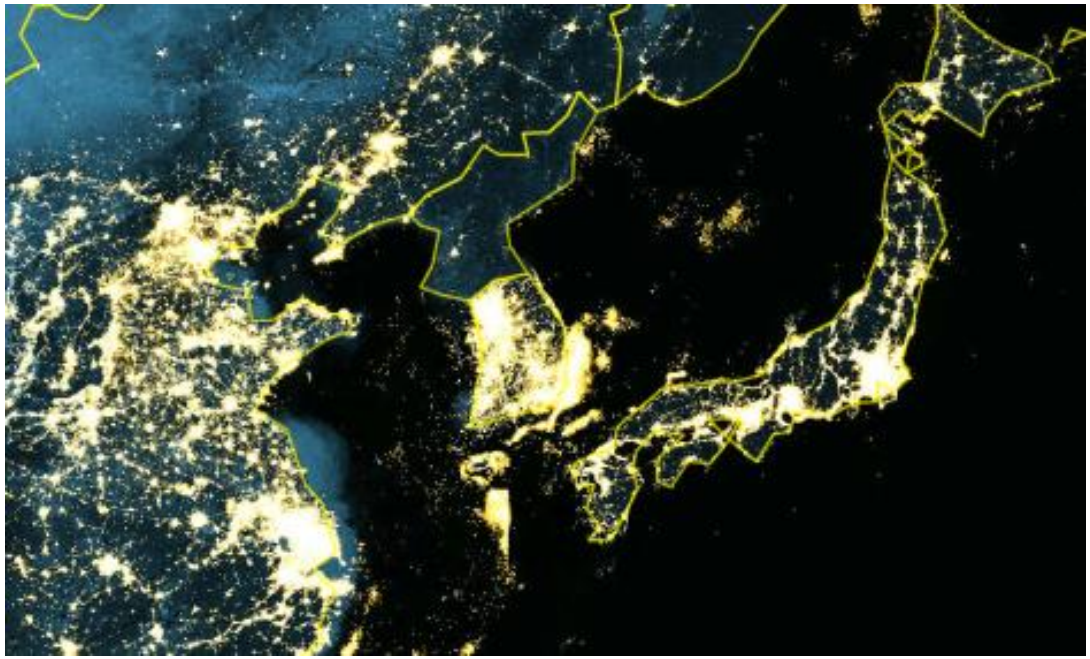

- **Idea: use nighttime light intensity as a proxy for economic development**
- Typically the high-resolution remote sensing data of choice for economists

Transfer Learning Approach

Our approach: predict **nighttime light intensity** from **daytime images**

( , Low nighttime light intensity)

…

( , High nighttime light intensity)

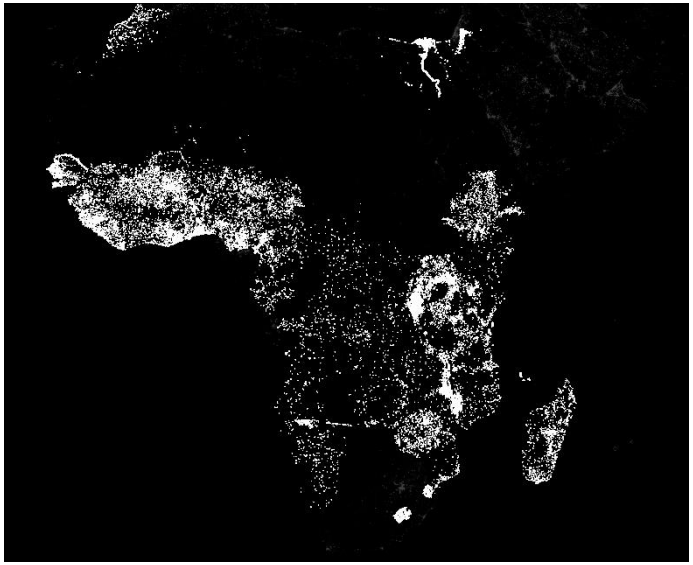

>300,000 training images

Data-rich proxy: global-scale **labeled** training data available

Goal: Automatically learn features useful for poverty task by first training a Deep Learning model for predicting nighttime light intensity

Our Model

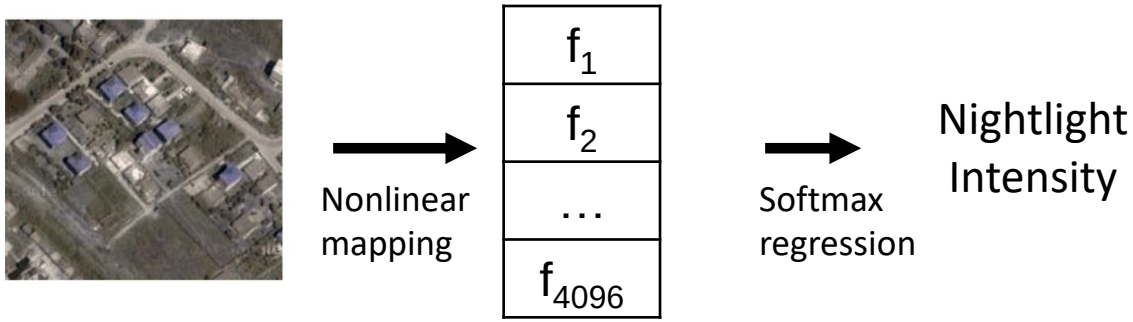

- Key aspects:
  - Spatial context is important
  - Interested in an aggregate measure over the entire image
- Developed a **fully convolutional CNN** architecture:
  - Can handle arbitrary sized images (no cropping needed)
  - Average pooling layer in the last layer
  - 55 million parameters vs. 621 million parameters for a 400x400 image
- (Final accuracy: 73%)

Have we learned to identify useful features?

Learned Features: Roads

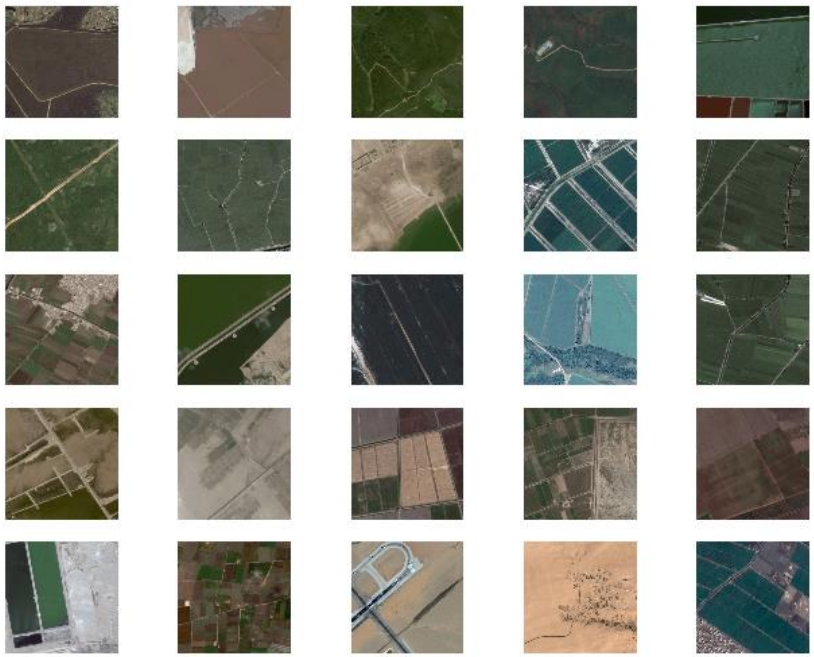

25 Maximally activating images

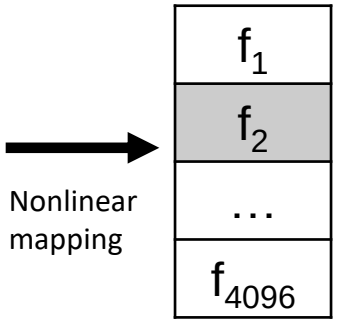

No supervision beyond nighttime lights - no labeled example of what a road looks like was provided!

Learned Features

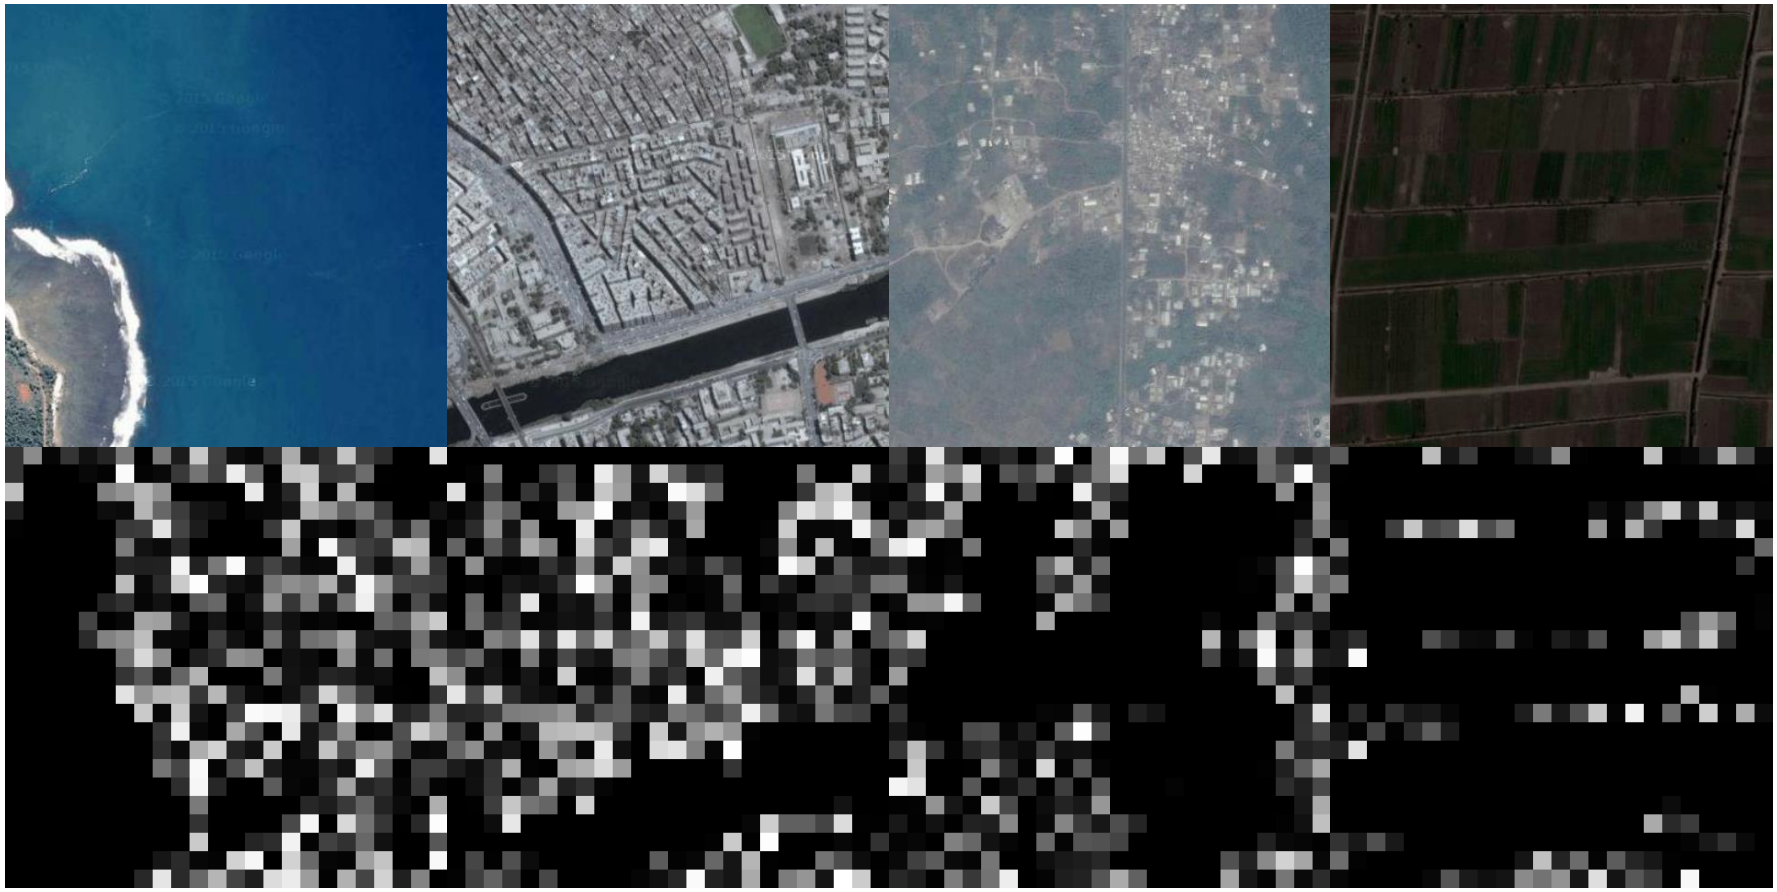

Nightlights to Poverty Estimation

- Living Standards Measurement Survey (LSMS) data in Uganda (World Bank):
  - ~700 household clusters
  - Data on expenditures, above/below poverty line, noisy locations
- **Task**: predict if the majority of households in a cluster are above or below the poverty line from corresponding images

…

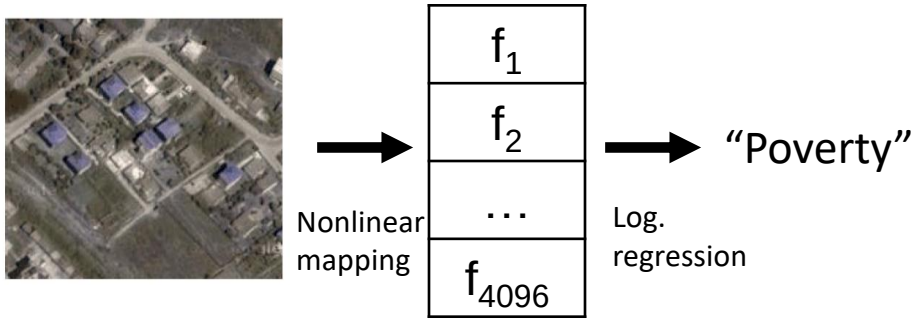

…

Features transferred from night-light task

	Transf
Accuracy	**0.71**
Precision	**0.39**
AUC	**0.76**

Nightlights to Poverty Estimation

- Compare against survey features which can feasibly be detected by satellite images (roof type, # of rooms, distance to population center)

…

Nonlinear mapping

f_1
f_2
…
f_{4096}

Log. regression

"Poverty"

…

Features transferred from night-light task

	Transf	Survey
Accuracy	**0.71**	0.75
Precision	**0.39**	0.45
AUC	**0.76**	0.78

Stanford

Nightlights to Poverty Estimation

- Compare against survey features which can feasibly be detected by satellite images (roof type, # of rooms, distance to population center)
- Dramatic improvement over the Night-Lights features currently used by economists

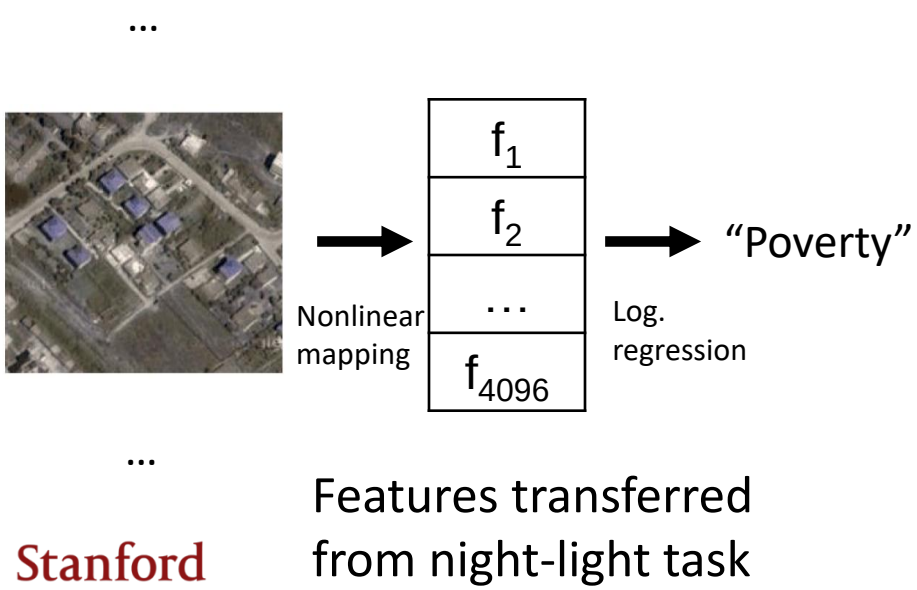

	Transf	Survey	Lights
Accuracy	**0.71**	0.75	0.53
Precision	**0.39**	0.45	0.30
AUC	**0.76**	0.78	0.72

High Resolution Poverty Maps

Running model on about 400,000 images from Uganda:

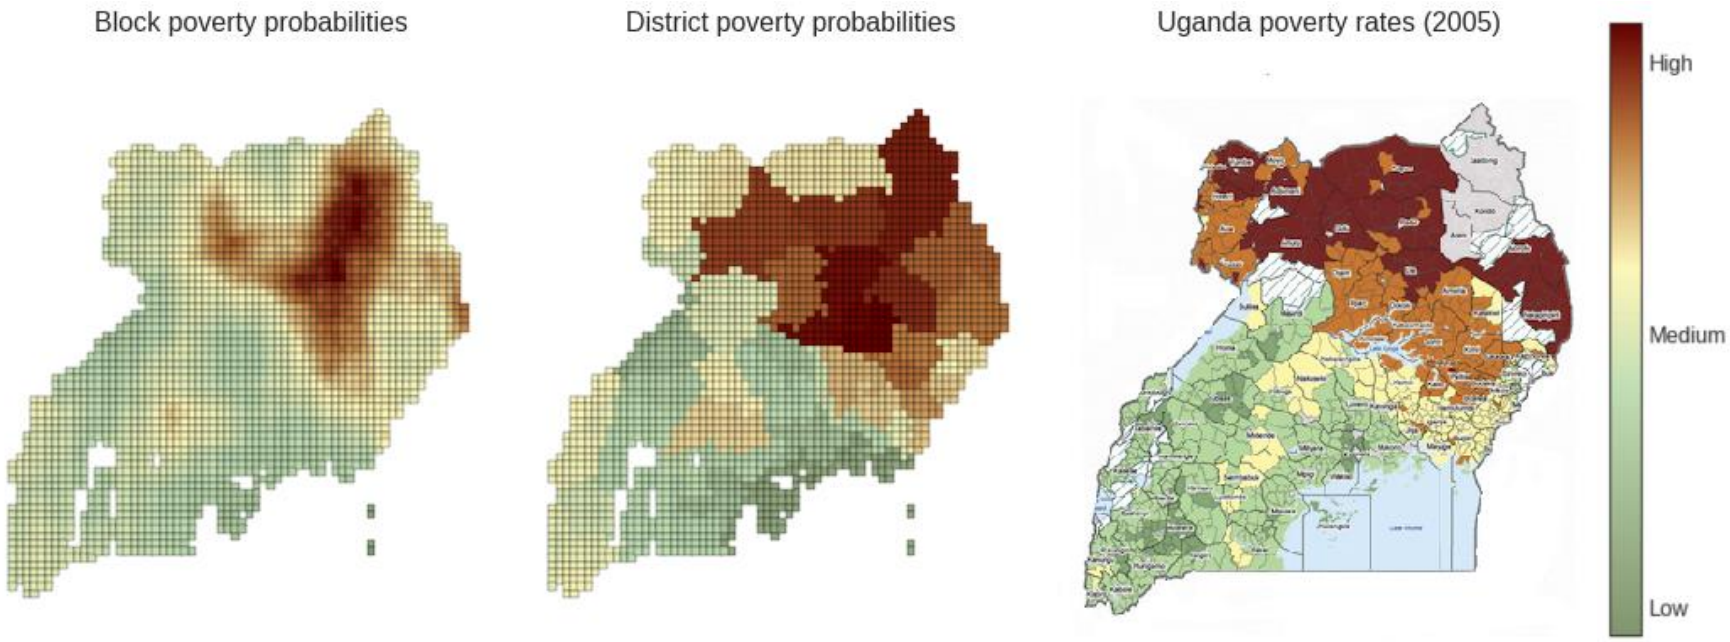

Most up-to-date map

Conclusion and Future Work

- **Scalable and inexpensive approach to generate high resolution maps.**
- Predicting real-valued consumption expenditure and asset index for more countries
- Spatial Graphical Models to encode spatial correlations

Thanks!