A Deep Learning Approach for Population Estimation from Satellite Imagery

Caleb Robinson, Fred Hohman, Bistra Dilkina

ACM SIGSPATIAL Workshop on Geospatial Humanities
Redondo Beach, California, United States

Georgia Tech

Knowing Where People Live is Important

Knowing Where People Live is Important

Urban development

Knowing Where People Live is Important

- Urban development
- Disaster response

Knowing Where People Live is Important

- Urban development
- Disaster response
- Infectious disease containment

Sustainable Development Goals

Most rely on population counts for measuring outcomes

Need to know population of a country

Need to measure population mortality rates

Need to know where people live in relation to services

Geospatial Humanities Applications

What is the influence of physical or geographical space on human behavior and cultural development? [1]

Using detailed population data and other accessory pieces of data we can learn relationships that may be useful in historic or cultural contexts.

Relationship between:

- ***land-cover*** and ***population*** with the Historic Land Dynamics Assessment
- ***population distribution in cities*** - ***cultural differences in cities***

[1] Bodenhamer, David J., John Corrigan, and Trevor M. Harris, eds. *The spatial humanities: GIS and the future of humanities scholarship*. Indiana University Press, 2010.

Censuses are Hard!

Expensive, time consuming, varying degrees of accuracy, varying degrees of spatial resolution, ...

2010 round - **5 countries** without a census

2000 round - **27 countries** without census

Lebanon hasn't taken a census since 1932

Much can Change Between Censuses!

Time-lapse of Las Vegas, Nevada

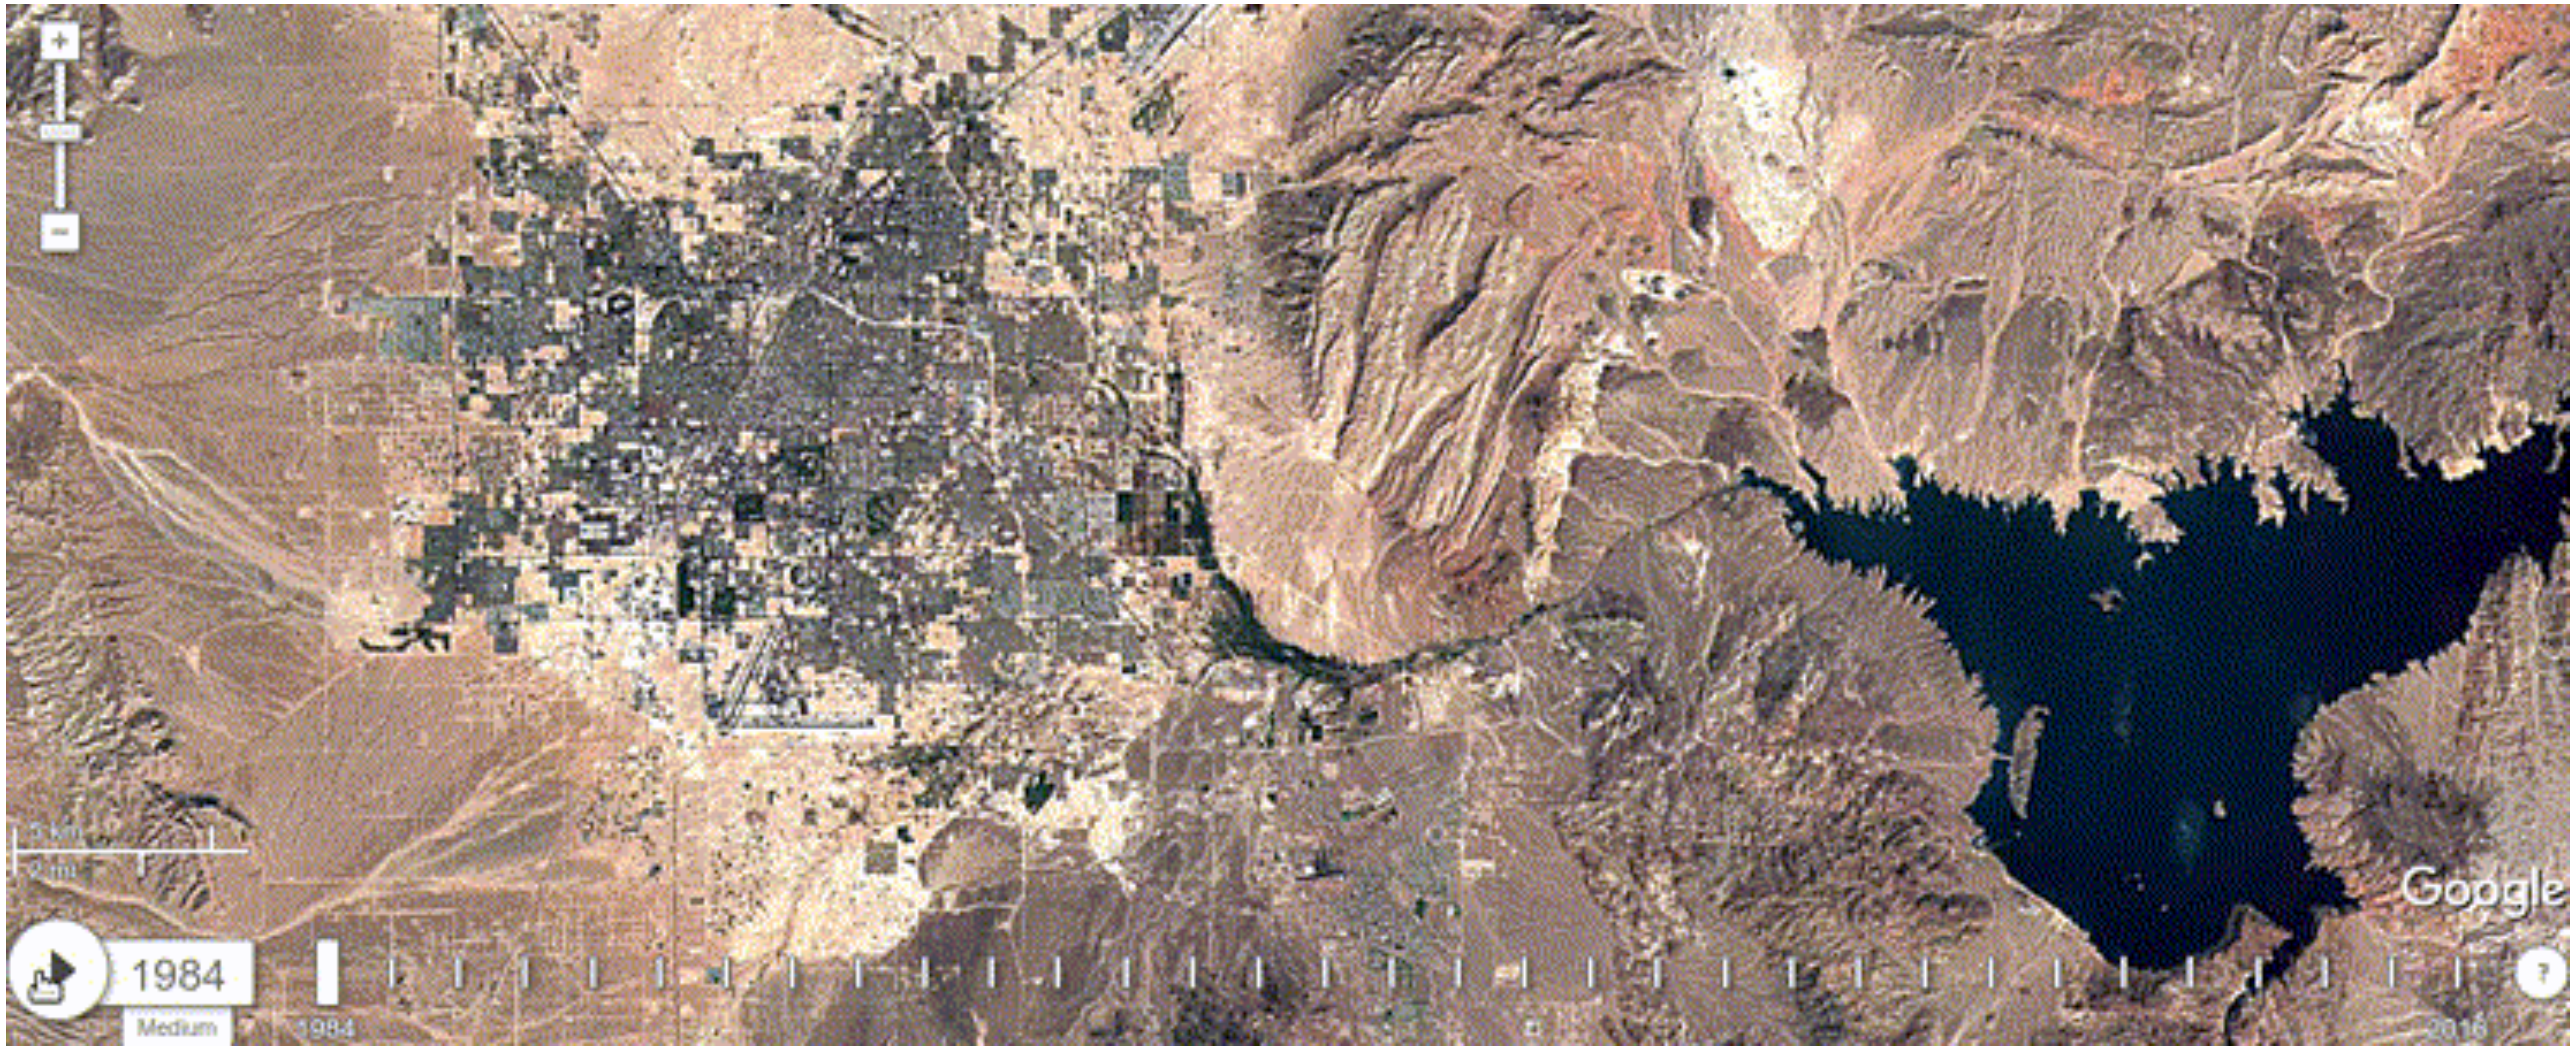

Made with: https://earthengine.google.com/timelapse/

How Satellite Data Can Help

Models based on satellite data are:

- *Much faster* than waiting on census
- Can be used in-between census years
- Can be applied to other countries without population counts
- Highly scalable

Landsat 7: https://earthobservatory.nasa.gov/Newsroom/NasaNews/2002/2002072210307.html

Landsat

The Landsat project (NASA and USGS) has been continuously collecting satellite images of the entire earth since 1972

Landsat 7 has been in operation since 1999, taking images of the entire globe every 16 days

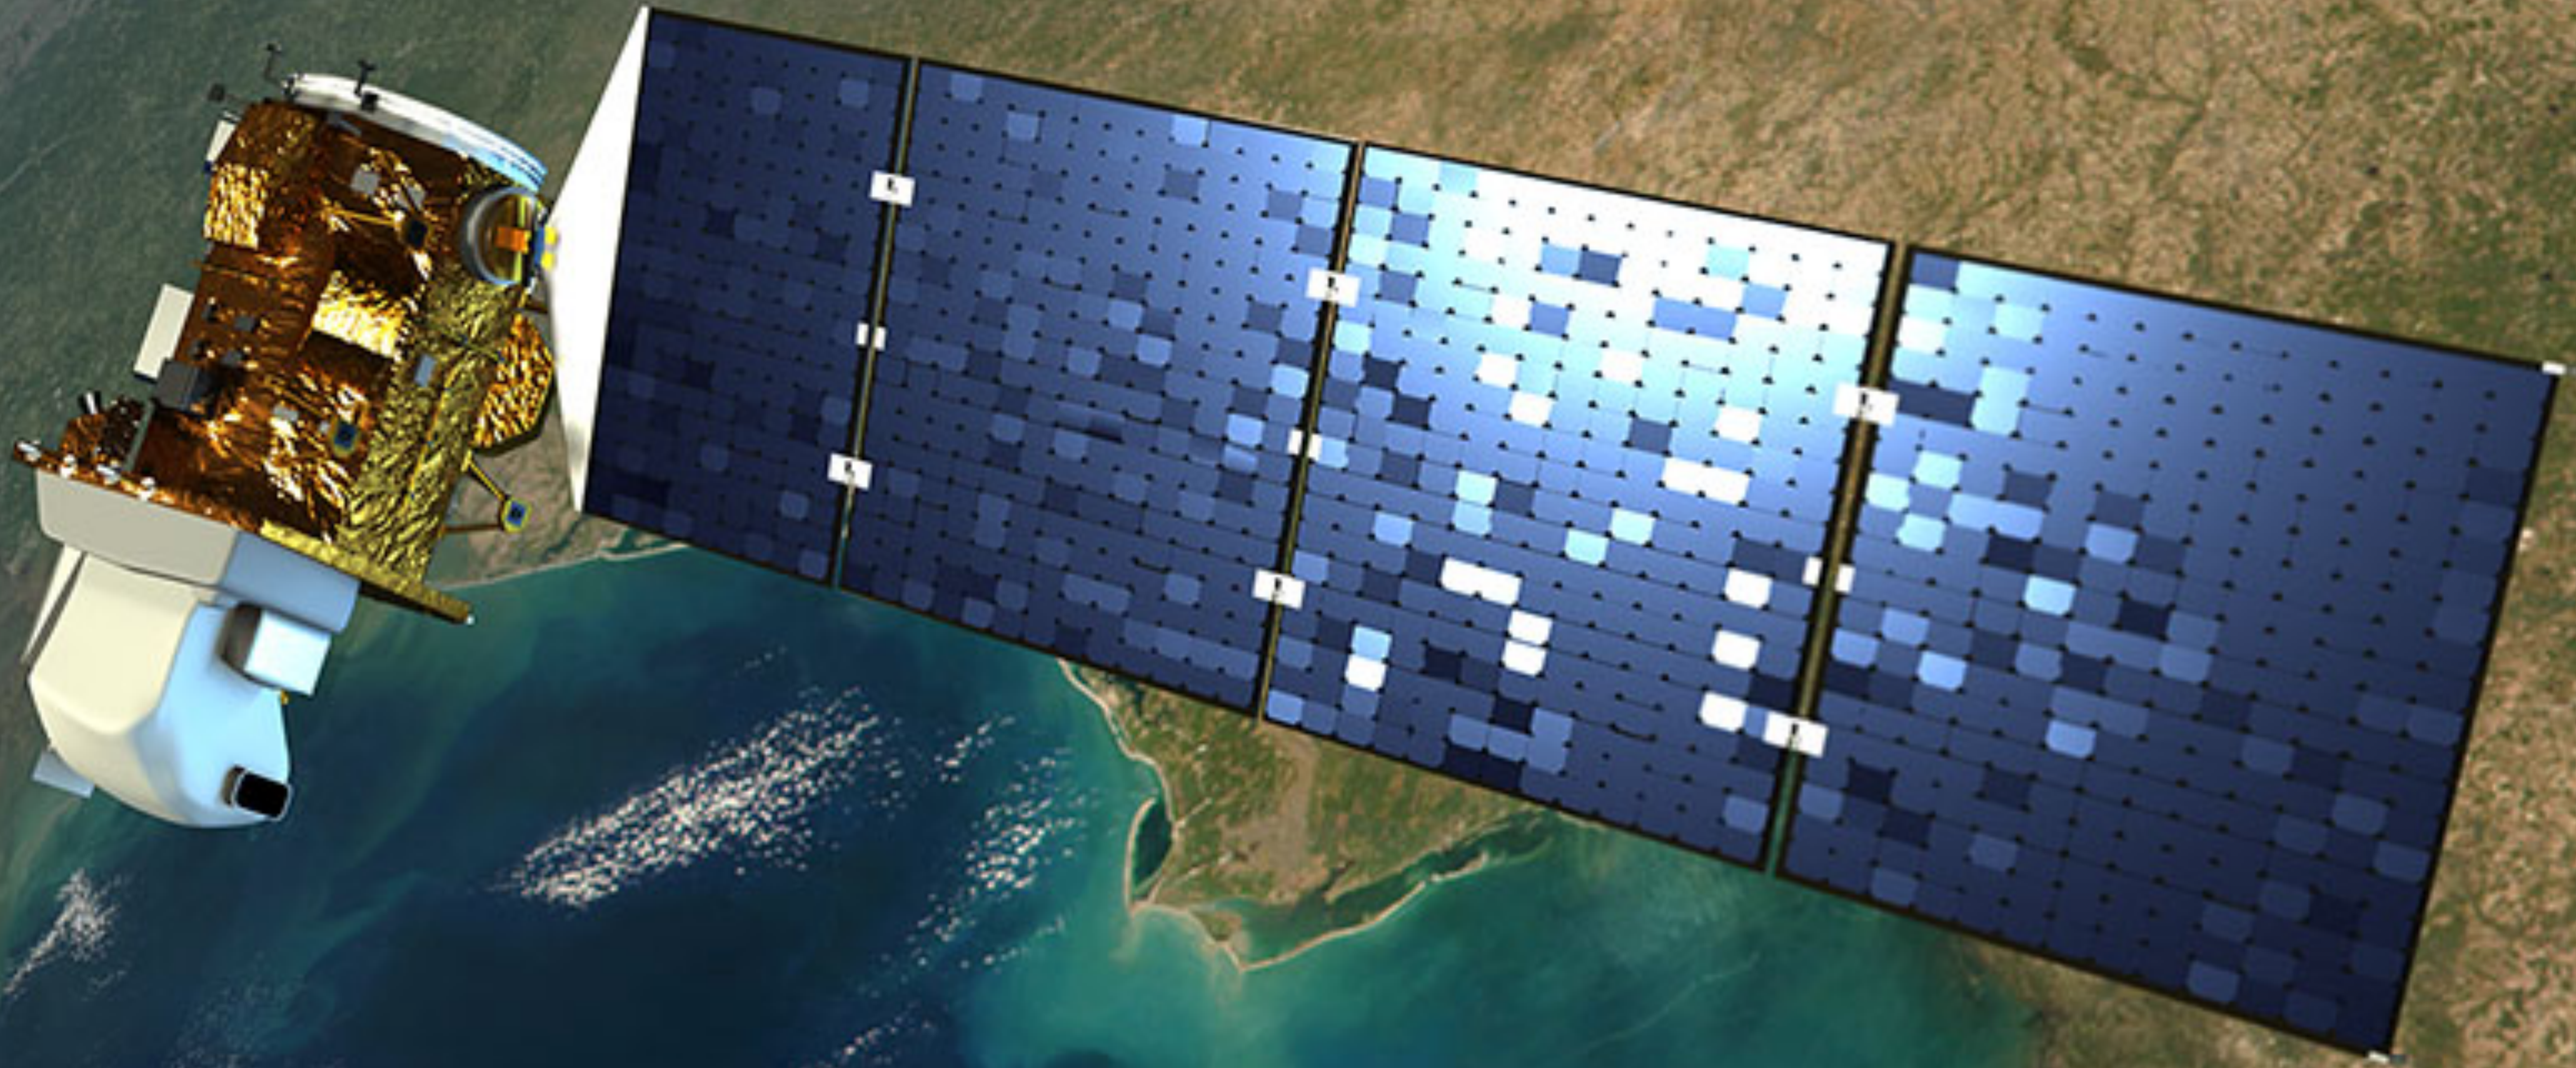

Landsat

The Landsat project (NASA and USGS) has been continuously collecting satellite images of the entire earth since 1972

Landsat 7 has been in operation since 1999, taking images of the entire globe every 16 days

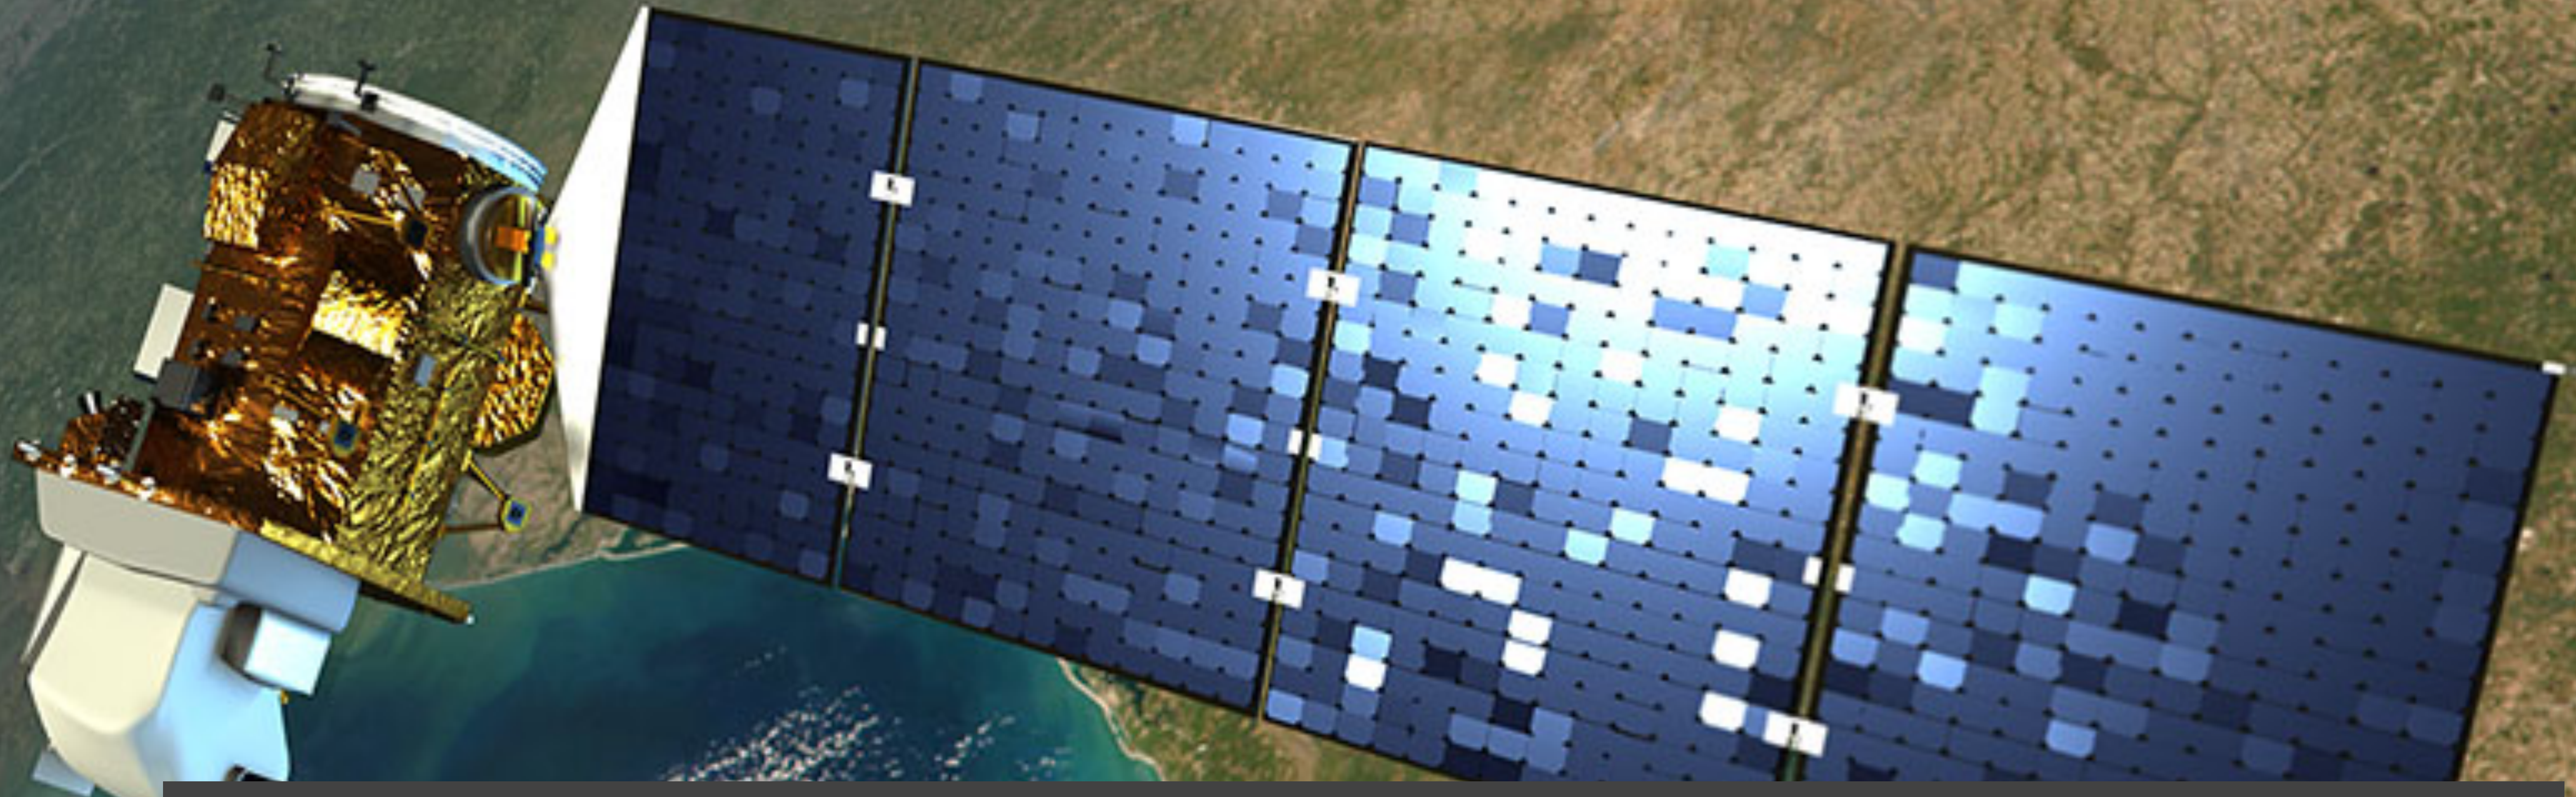

What can we learn from all of this data?

- *What is in a single satellite image?*
- *Where people live*

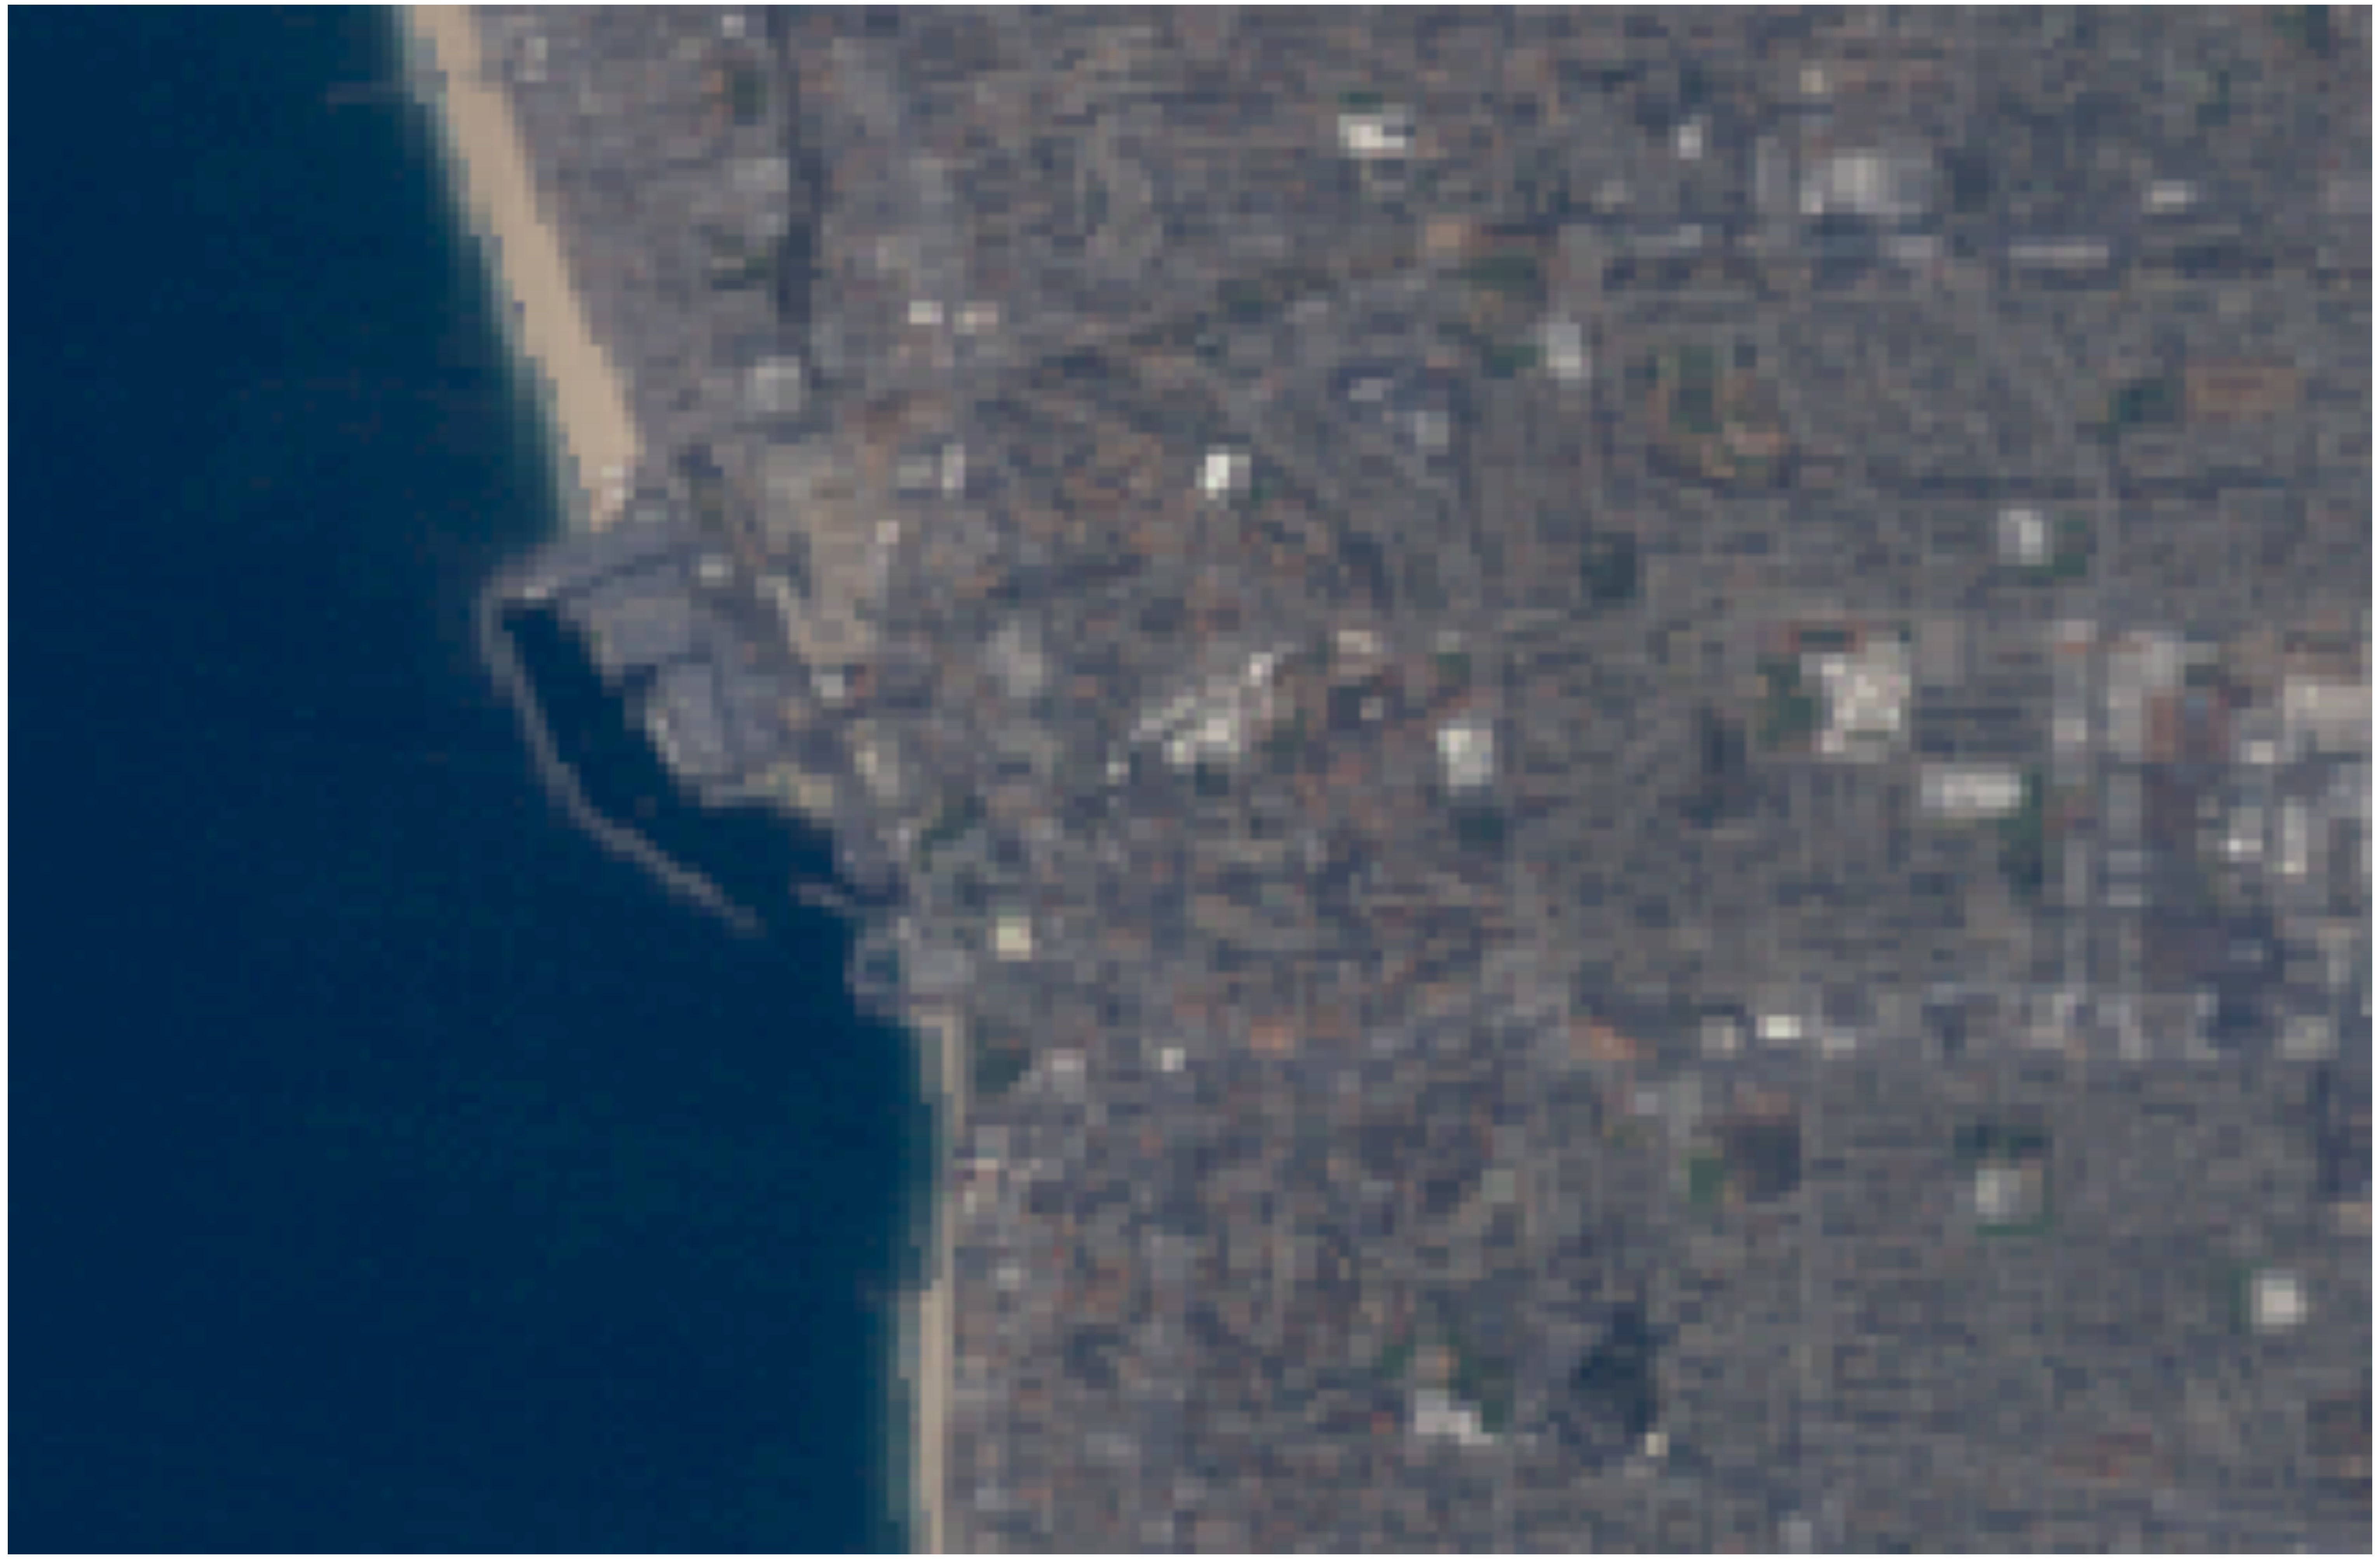

Landsat 7 imagery of Redondo Beach (2010)

We are learning to estimate population from satellite imagery

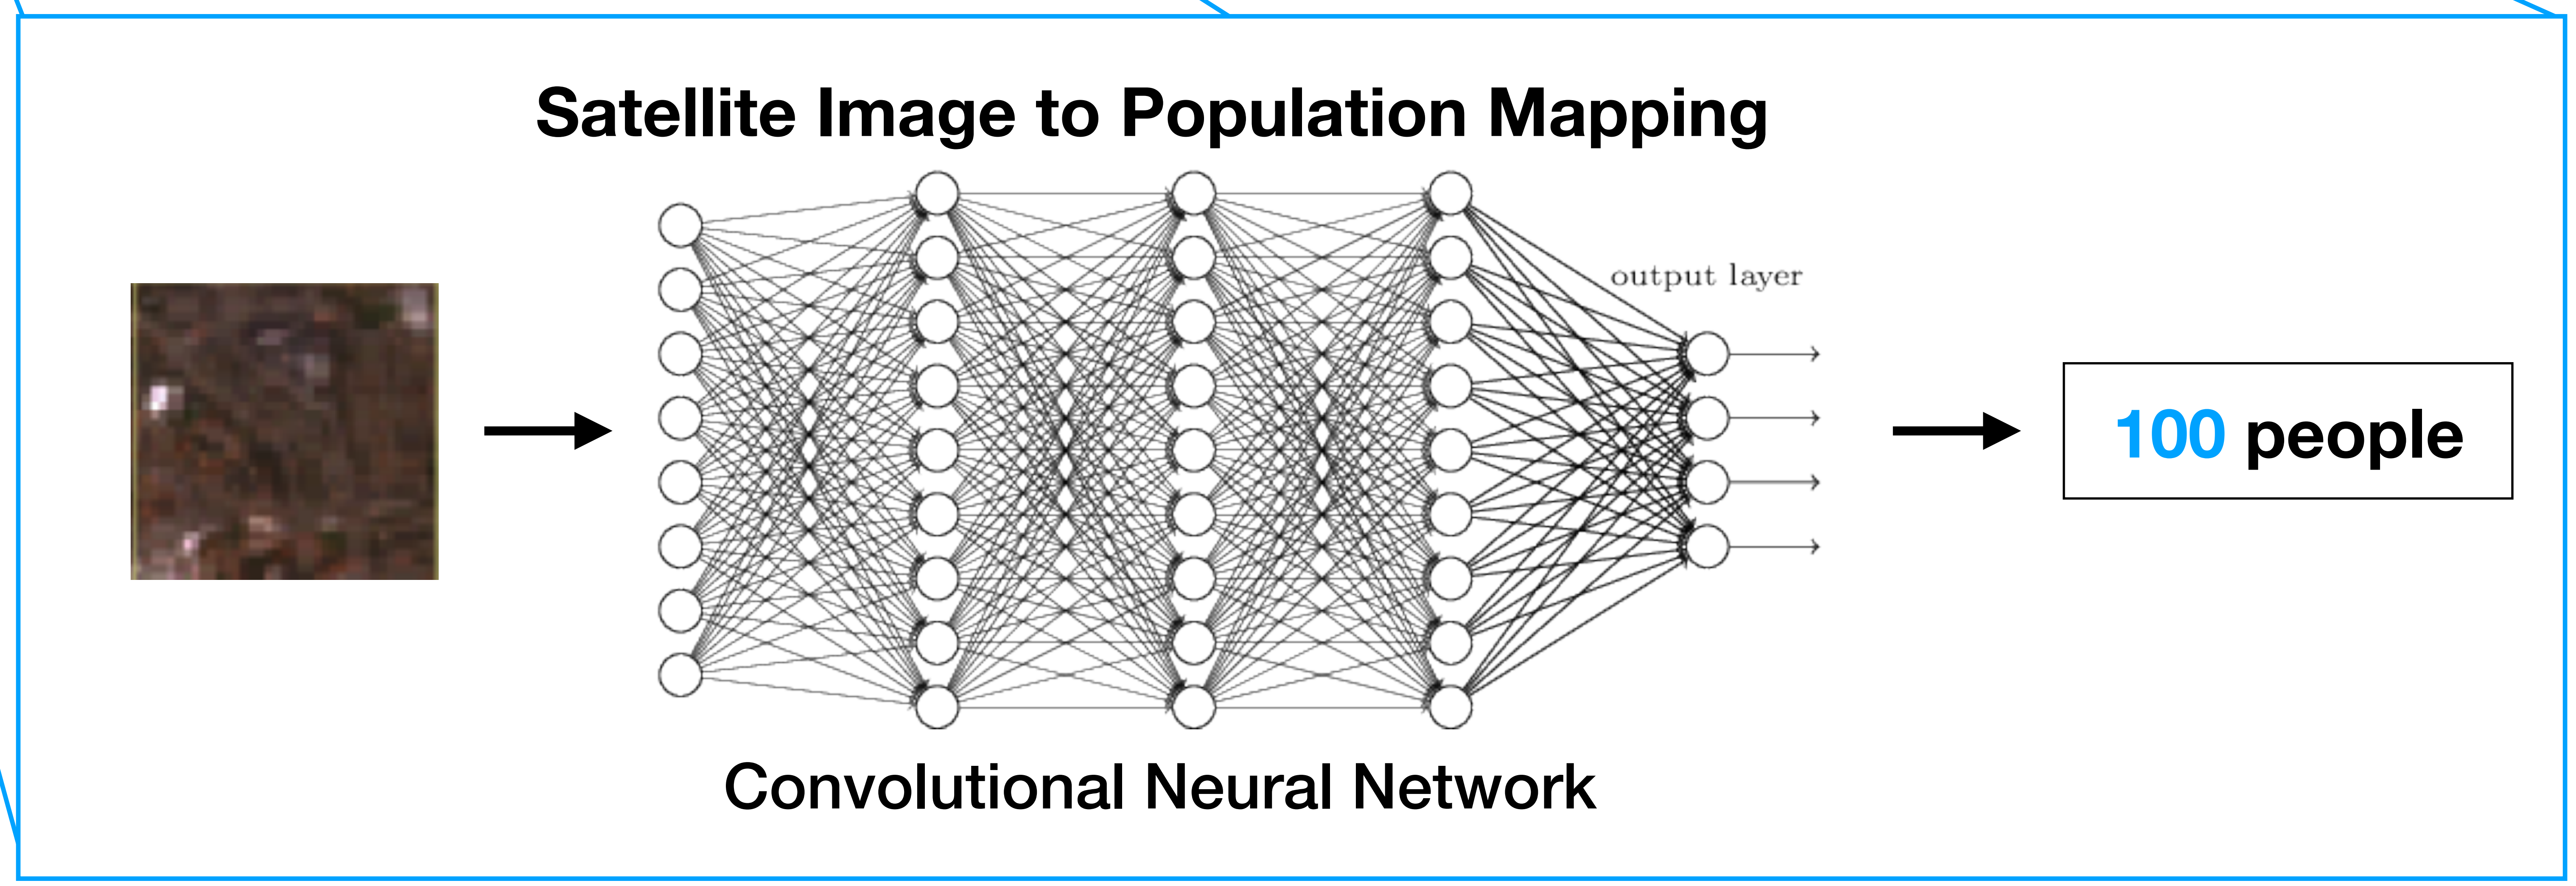

Gridded Population Data Products

Gridded Population of the World LandScan

Global Rural-Urban Mapping Project WorldPop Facebook Labs

Population Ground Truth

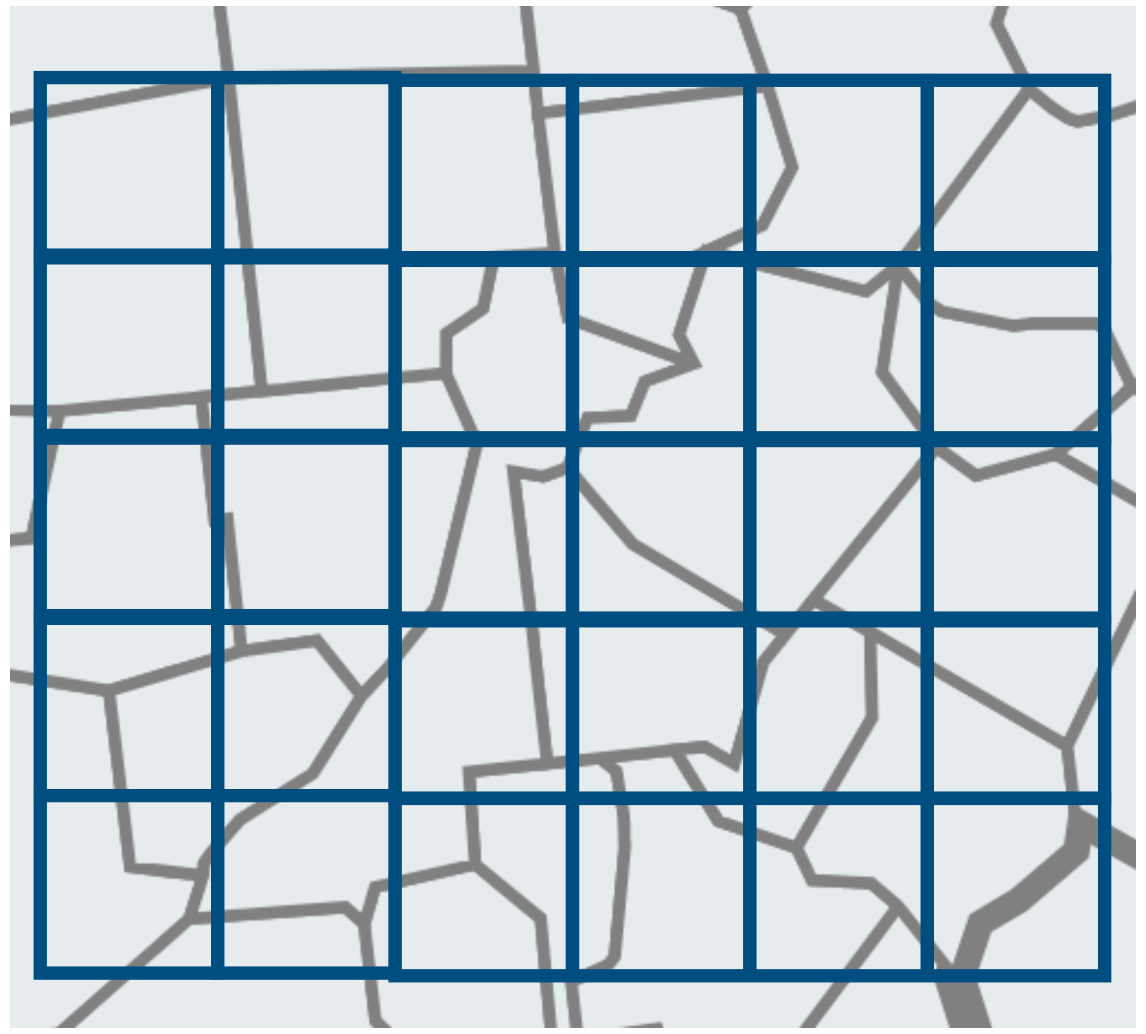

Disaggregated to grid

Gridded Population Data Products

Gridded Population of the World LandScan

Global Rural-Urban Mapping Project WorldPop Facebook Labs

Population Ground Truth

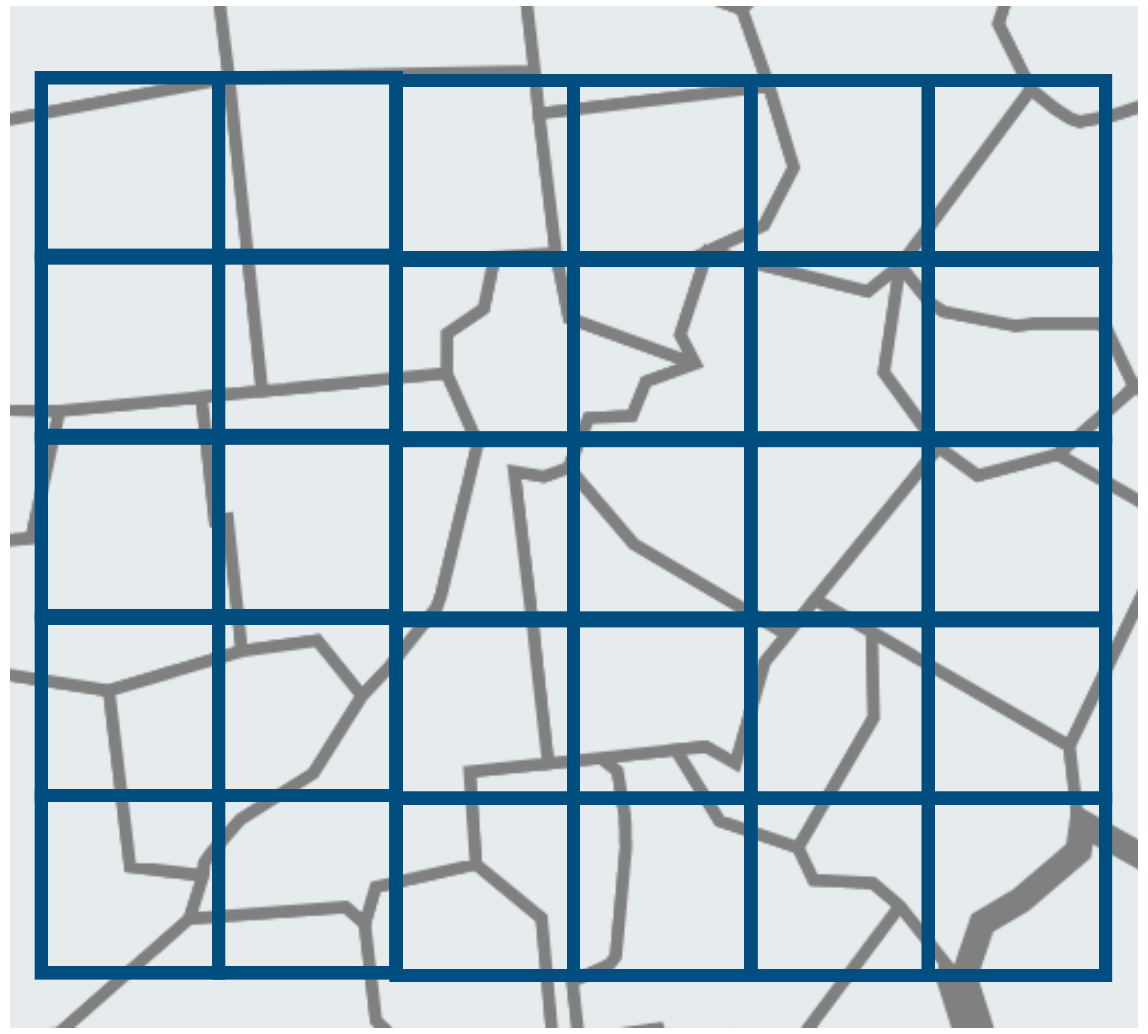

Disaggregated to grid

Need to know ground truth!

Gridded Population Estimations

Equitable development through deep learning: The case of sub-national population density estimation

Patrick Doupe
Arnhold Institute for Global Health
New York, NY, USA
patrick.doupe@mssm.edu

Emilie Bruzelius
Arnhold Institute for Global Health
New York, NY, USA
emilie.bruzelius@mssm.edu

James Faghmous
Arnhold Institute for Global Health
New York, NY, USA
james.faghmous@mssm.edu

Samuel G. Ruchman
Office of the UN Secretary-General's Special Envoy for Health in Agenda 2030 and for Malaria
New York, NY, USA
sruchman@healthenvoy.org

ABSTRACT

High-resolution population density maps are a critical component for global development efforts, including service delivery, resource allocation, and disaster response. Traditional population density efforts are predominantly survey driven, which are laborious, prohibitively expensive, infrequently updated, and inaccurate – especially in remote areas. Furthermore, these maps are developed on a regional-basis where the methods used vary region to region, hence introducing notable spatio-temporal heterogeneity and bias.

The advent of global-scale satellite imagery provides us with an unprecedented opportunity to create inexpensive, accurate, homogeneous, and rapidly updated population maps. To fulfill this vision, we must overcome both infrastructure and methodological obstacles. We propose a convolutional neural network approach that addresses some of the methodological challenges, while employing a publicly available, albeit low resolution, remote sensed product. The method converts satellite images into population density estimates. To explore both the accuracy and generalizability of our approach, we train our neural network on Tanzanian imagery and test the model on Kenyan data. We show that our method is able to generalize to unseen data and we improve upon the current state of the art by 177 percent.

1. INTRODUCTION

The 2015 adoption of the Sustainable Development Goals has created an acute need to collect accurate subnational statistics to monitor progress towards these targets [14]. Furthermore, subnational data are a critical component of equitable development since national averages routinely "smooth out" the most marginalized communities. In fact, one of the most significant limitations to achieving the Millennium Development Goals was that, while national averages may have showed overall gains towards certain goals, these gains veneered significant disparities, where the poorest people did not necessarily benefit from the overall progress [38].

One of the most basic, yet essential, of these statistics is the accurate and timely assessment of where people live. This information is "one of the primary sources of data needed for formulating, implementing and monitoring the effectiveness of policies and programmes aimed at inclusive socioeconomic development and environmental sustainability" [40, p.2]. These census data determine resource allocation, such as where to invest in hospitals, schools and infrastructure, and may be used to define legislative districts and other important functional areas of government. Basic population level data also serve as an essential benchmark for measuring progress toward the the attainment of the Sustainable Development Goals, and other national and international objectives (see Figure 1).

Traditional population density estimates are derived from census surveys. However the utility of censuses is hampered by several widely recognized quality issues. First, they are "among the most complex and massive peacetime exercises a nation undertakes" [42, p.5]. Many countries still lack the capacity, both in terms of financial and human resources, to collect data regularly [7]. As a result, population maps in many low income countries are outdated or of poor quality [27]. For instance, in several low- and middle-

[1] Doupe, Patrick, et al. "Equitable development through deep learning: The case of sub-national population density estimation." Proceedings of the 7th Annual Symposium on Computing for Development. ACM, 2016.

Gridded Population Estimations

Equitable development through deep learning: The case of sub-national population density estimation

Patrick Doupe
Arnhold Institute for Global Health
New York, NY, USA
patrick.doupe@mssm.edu

Emilie Bruzelius
Arnhold Institute for Global Health
New York, NY, USA
emilie.bruzelius@mssm.edu

James Faghmous
Arnhold Institute for Global Health
New York, NY, USA
james.faghmous@mssm.edu

Samuel G. Ruchman
Office of the UN Secretary-General's Special Envoy for Health in Agenda 2030 and for Malaria
New York, NY, USA
sruchman@healthenvoy.org

ABSTRACT

High-resolution population density maps are a critical component for global development efforts, including service delivery, resource allocation, and disaster response. Traditional population density efforts are predominantly survey driven, which are laborious, prohibitively expensive, infrequently updated, and inaccurate – especially in remote areas. Furthermore, these maps are developed on a regional-basis where the methods used vary region to region, hence introducing notable spatio-temporal heterogeneity and bias.

The advent of global-scale satellite imagery provides us with an unprecedented opportunity to create inexpensive, accurate, homogeneous, and rapidly updated population maps. To fulfill this vision, we must overcome both infrastructure and methodological obstacles. We propose a convolutional neural network approach that addresses some of the methodological challenges, while employing a publicly available, albeit low resolution, remote sensed product. The method converts satellite images into population density estimates. To explore both the accuracy and generalizability of our approach, we train our neural network on Tanzanian imagery and test the model on Kenyan data. We show that our method is able to generalize to unseen data and we improve upon the current state of the art by 177 percent.

1. INTRODUCTION

The 2015 adoption of the Sustainable Development Goals has created an acute need to collect accurate subnational statistics to monitor progress towards these targets [14].

Furthermore, subnational data are a critical component of equitable development since national averages routinely "smooth out" the most marginalized communities. In fact, one of the most significant limitations to achieving the Millennium Development Goals was that, while national averages may have showed overall gains towards certain goals, these gains veneered significant disparities, where the poorest people did not necessarily benefit from the overall progress [38].

One of the most basic, yet essential, of these statistics is the accurate and timely assessment of where people live. This information is "one of the primary sources of data needed for formulating, implementing and monitoring the effectiveness of policies and programmes aimed at inclusive socioeconomic development and environmental sustainability" [40, p.2]. These census data determine resource allocation, such as where to invest in hospitals, schools and infrastructure, and may be used to define legislative districts and other important functional areas of government. Basic population level data also serve as an essential benchmark for measuring progress toward the the attainment of the Sustainable Development Goals, and other national and international objectives (see Figure 1).

Traditional population density estimates are derived from census surveys. However the utility of censuses is hampered by several widely recognized quality issues. First, they are "among the most complex and massive peacetime exercises a nation undertakes" [42, p.5]. Many countries still lack the capacity, both in terms of financial and human resources, to collect data regularly [7]. As a result, population maps in many low income countries are outdated or of

How can we trust these estimates?

[1] Doupe, Patrick, et al. "Equitable development through deep learning: The case of sub-national population density estimation." Proceedings of the 7th Annual Symposium on Computing for Development. ACM, 2016.

Remaining Outline

- Deep learning overview
- Data
- Methods
- Results
- Future work

Deep Learning

Deep learning is a category of state-of-the-art algorithms that has recently been shown to be very effective for image-based learning

Example: convolutional neural networks (CNNs)

Goal: learn a mapping between satellite images and population counts

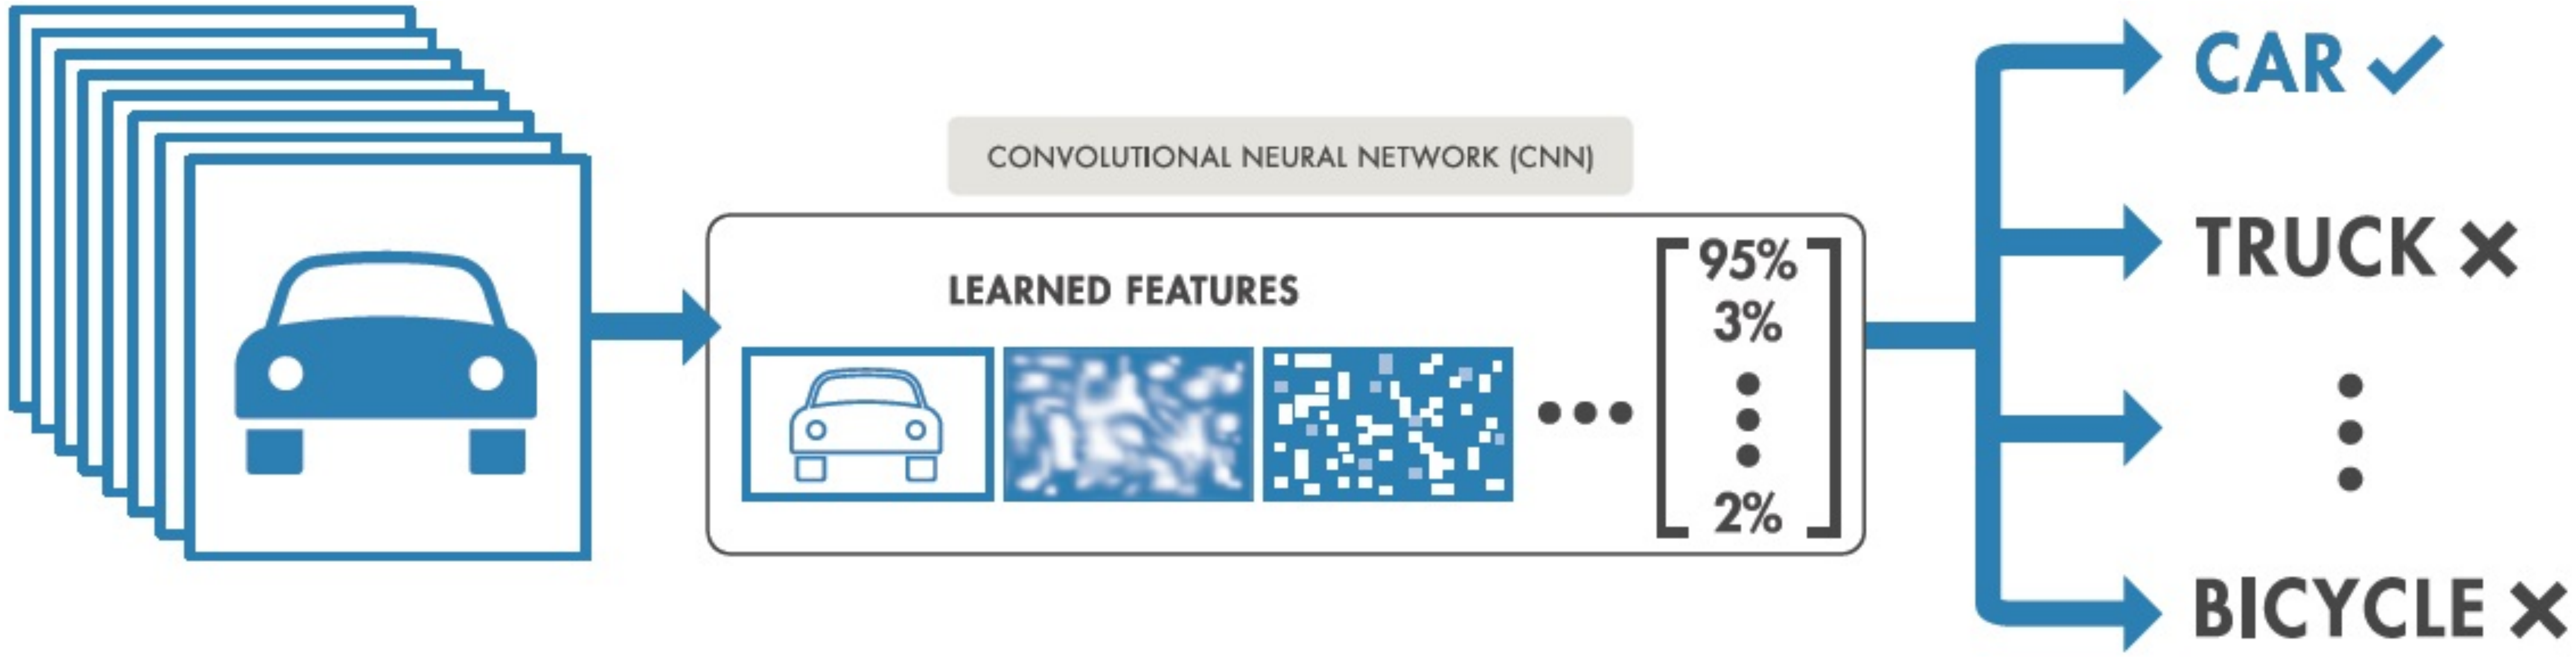

https://www.mathworks.com/discovery/convolutional-neural-network.html

Data

Landsat 7 images

- `74 x 74 x 7` "images" representing 1km^2 patches of land

Gridded US Census population count data

- Disaggregate block group data onto grid
- Raw population counts
- Population ranges

For every 1km^2 patch of land in the US we have:
a satellite image, existing population estimate for years 2000 and 2010

Methods

Estimating population ranges instead of counts (classification v. regression)

Class 0: 0 people, **Classes 1-7**: few people, **Classes 8-13**: many people

Train on data from 2000, test on data from 2010

How do we validate our model?

- Aggregate gridded estimations at county level - compare to ground truth
- Show that it makes reasonable predictions

Methods

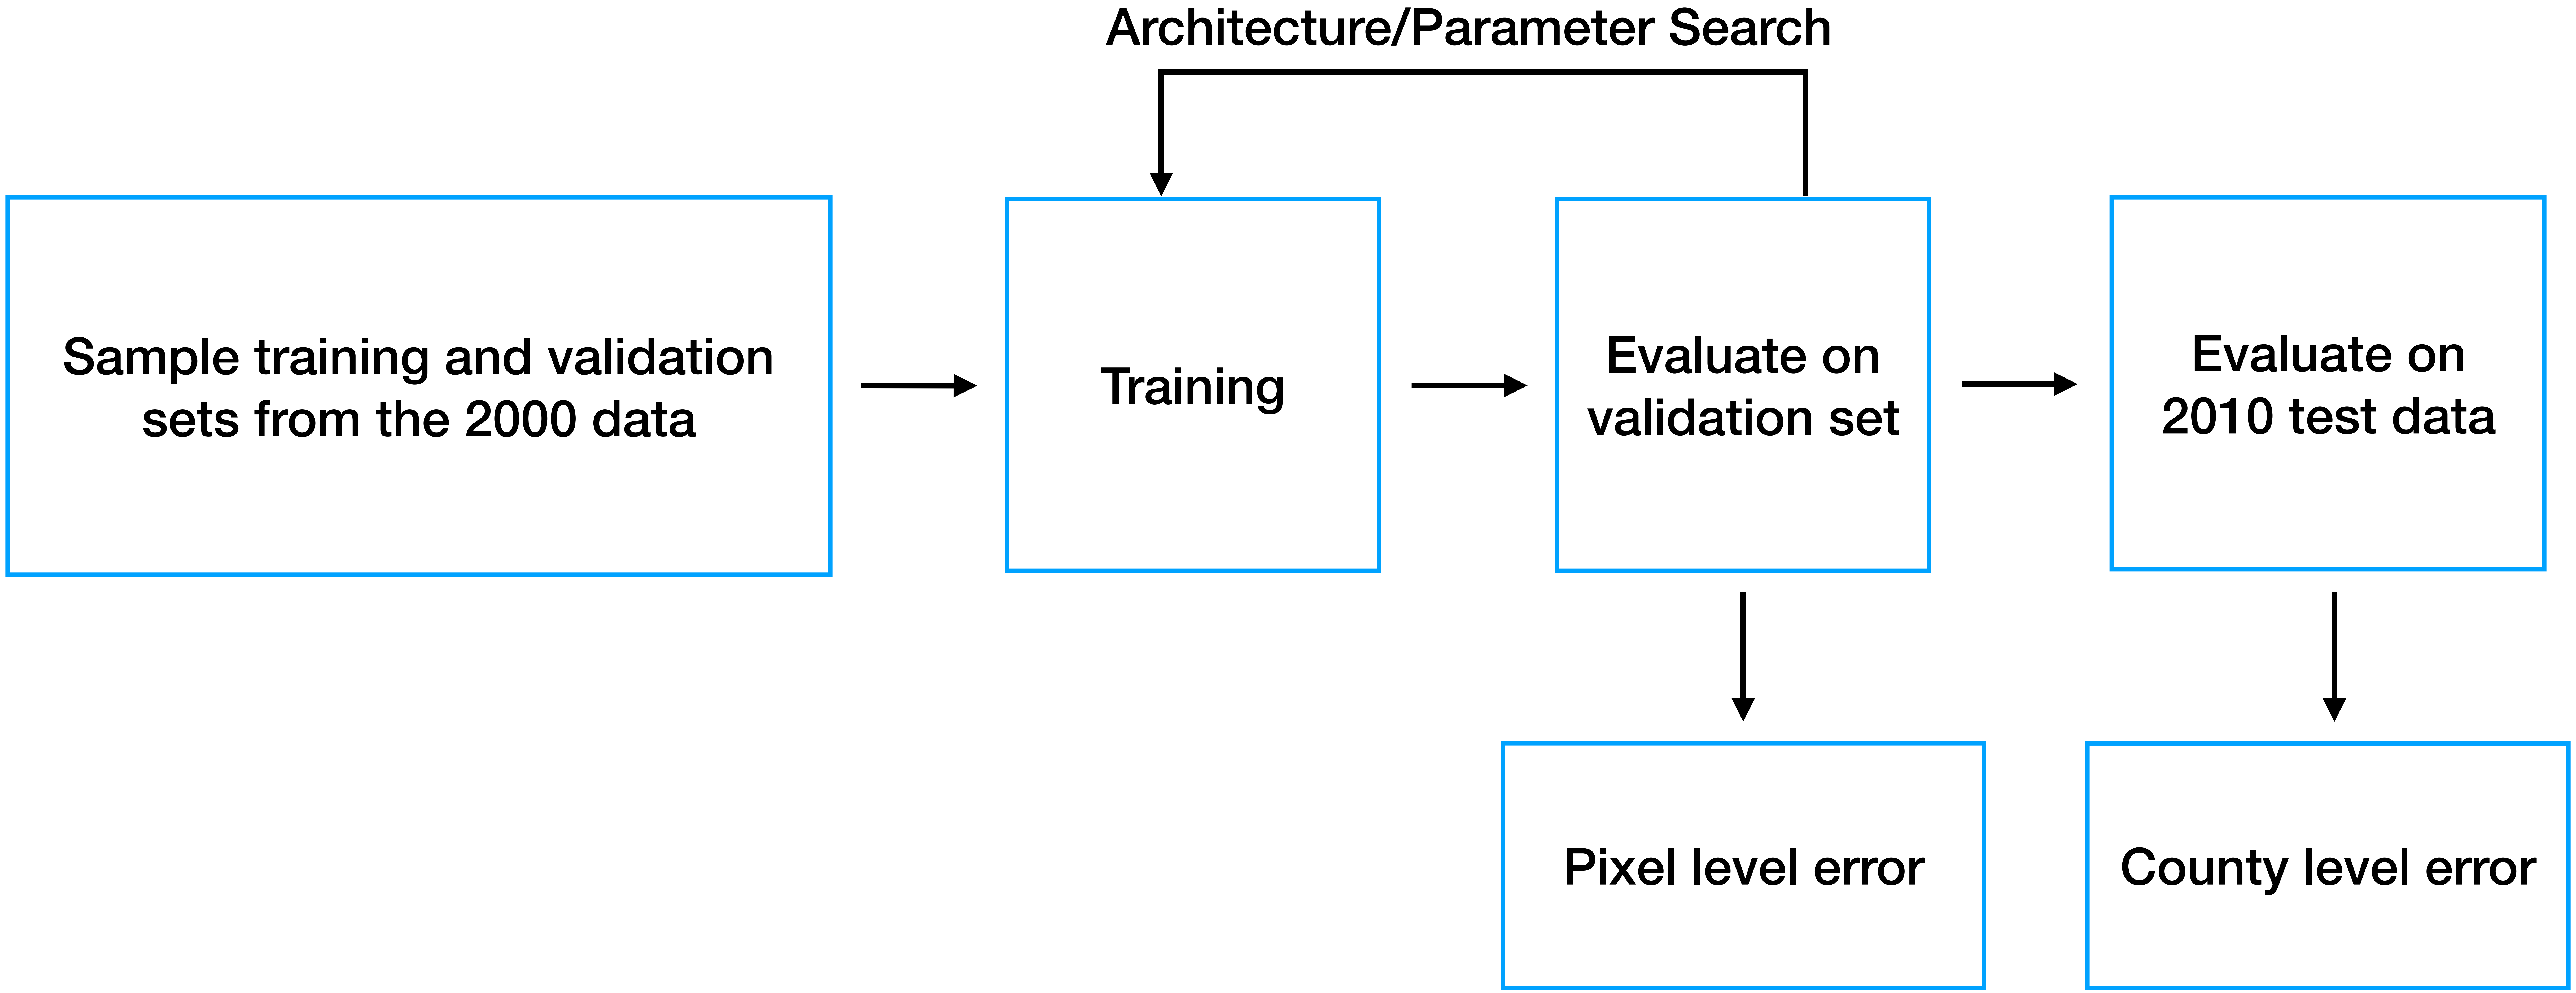
Architecture/Parameter Search
Sample training and validation sets from the 2000 data
Training
Evaluate on validation set
Evaluate on 2010 test data
Pixel level error
County level error

Neural Network Architecture

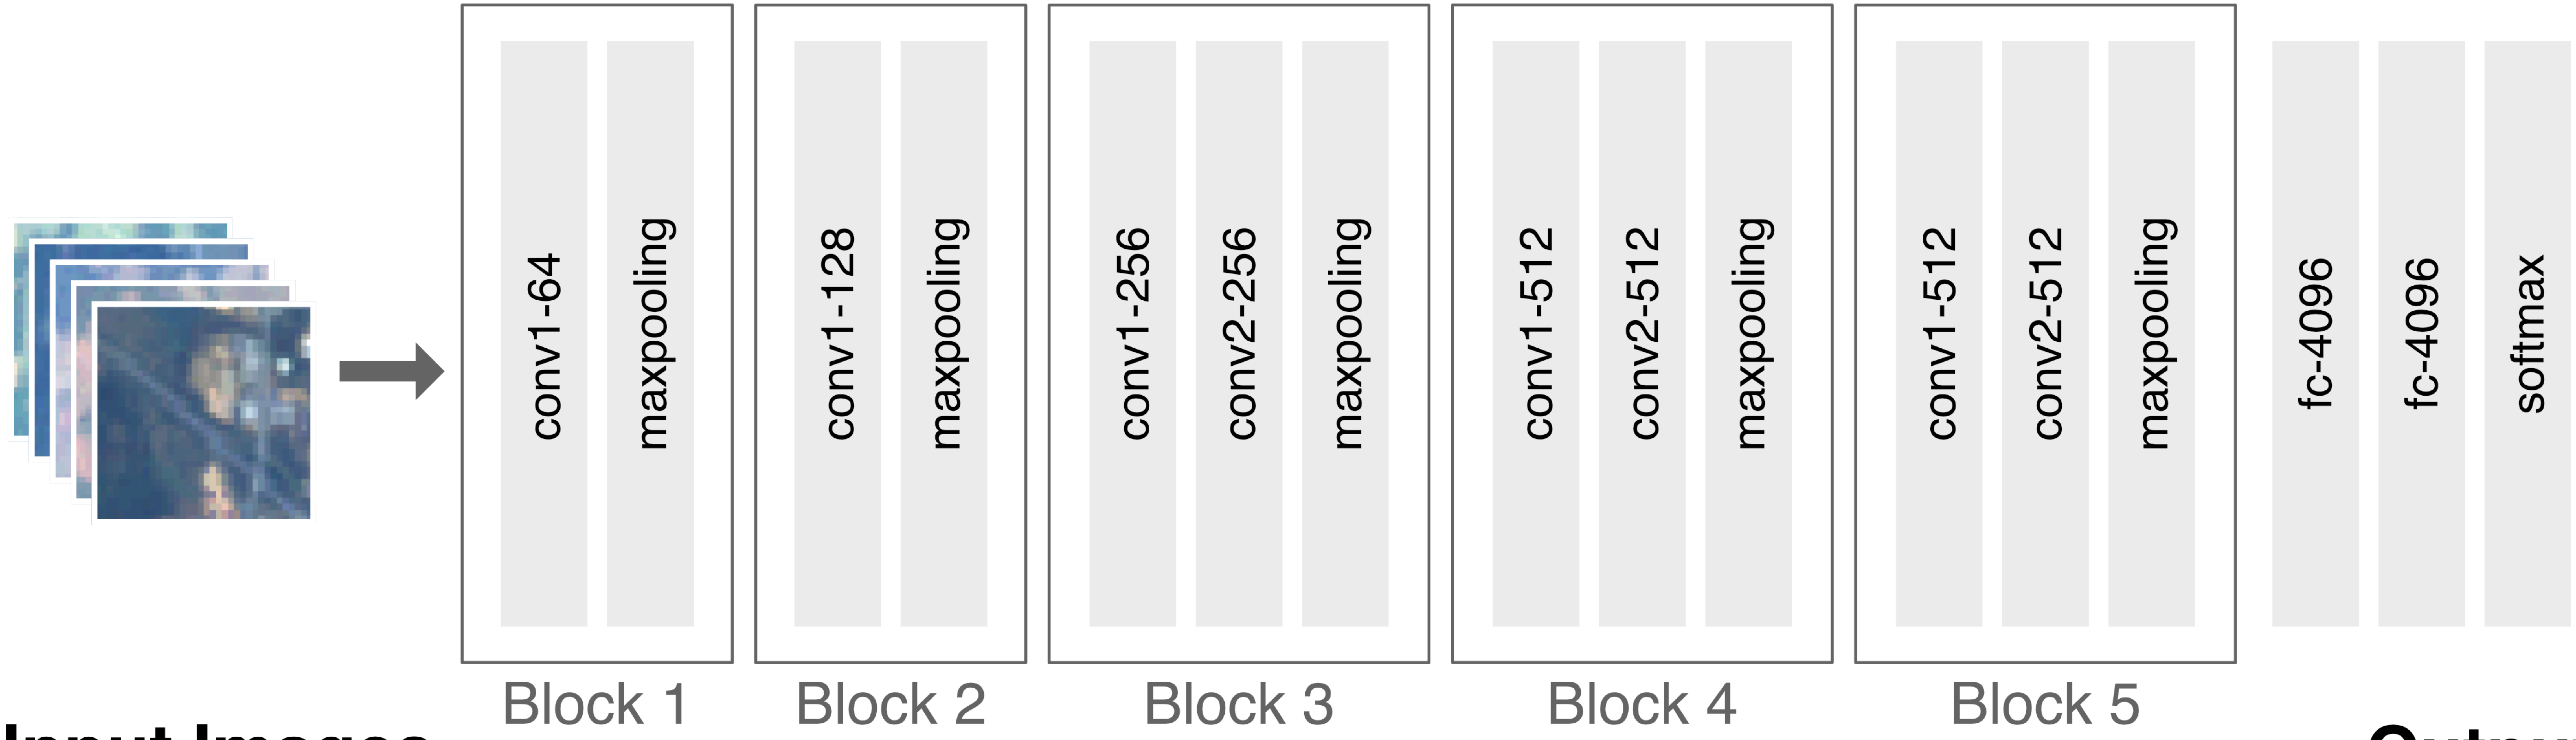

Input Images
Each image
of size (74, 74, 7)

Modified VGG-A Neural Network

Output
Classification
vector of size (17)

Quantitative Results

Results are at the county level

- Our models: **CONVRAW** and **CONVAUG**
- Census Bureau estimates: **POSTCENSAL** and **ACS5YR**

Method	Mean Absolute Error	Median Absolute Error	r^2	MAPE
CONVRAW	23,005	6,357	0.91	73.78
CONVAUG	19,484	**4,642**	0.94	49.82
POSTCENSAL	2,020	559	0.99	3.09
ACS5YR	1,704	**214**	0.99	34.44

2010 Mean Population: 97,018

2010 Median Population: 25,930

Difference between
Actual 2010 and **CONVAUG 2010**

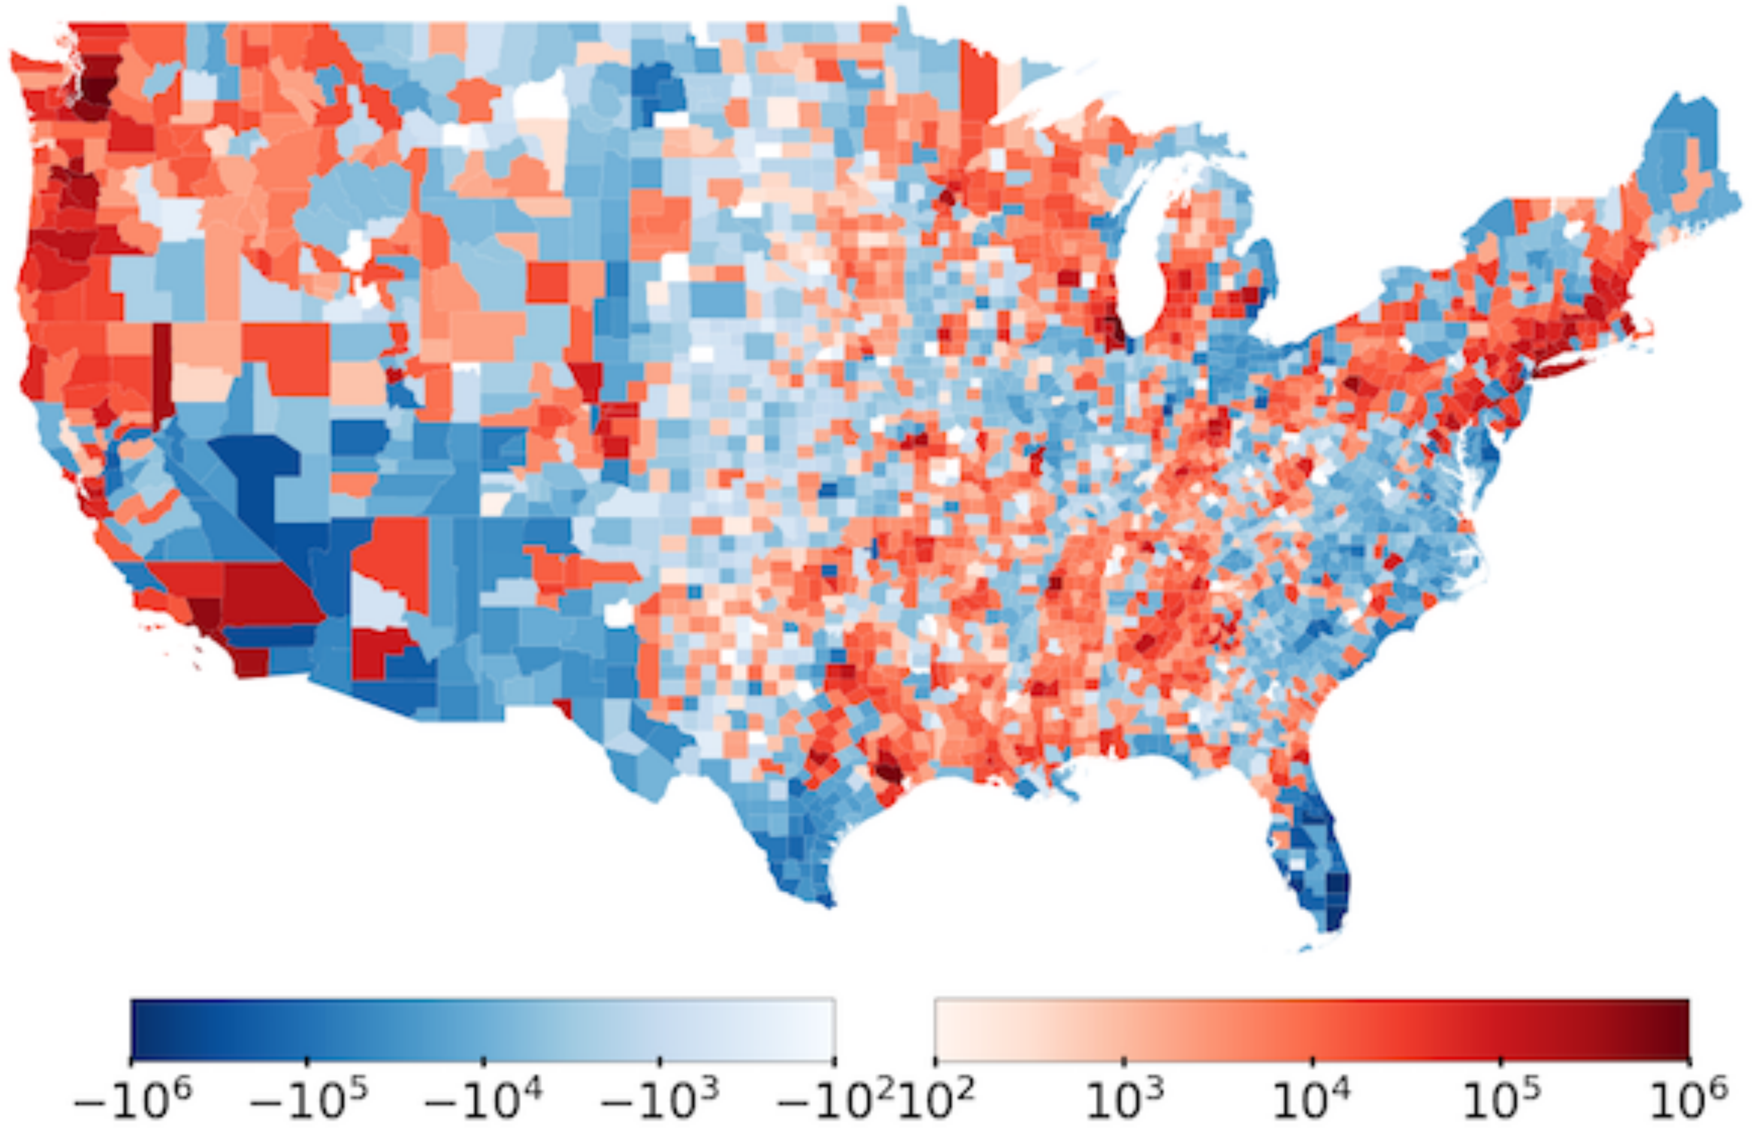

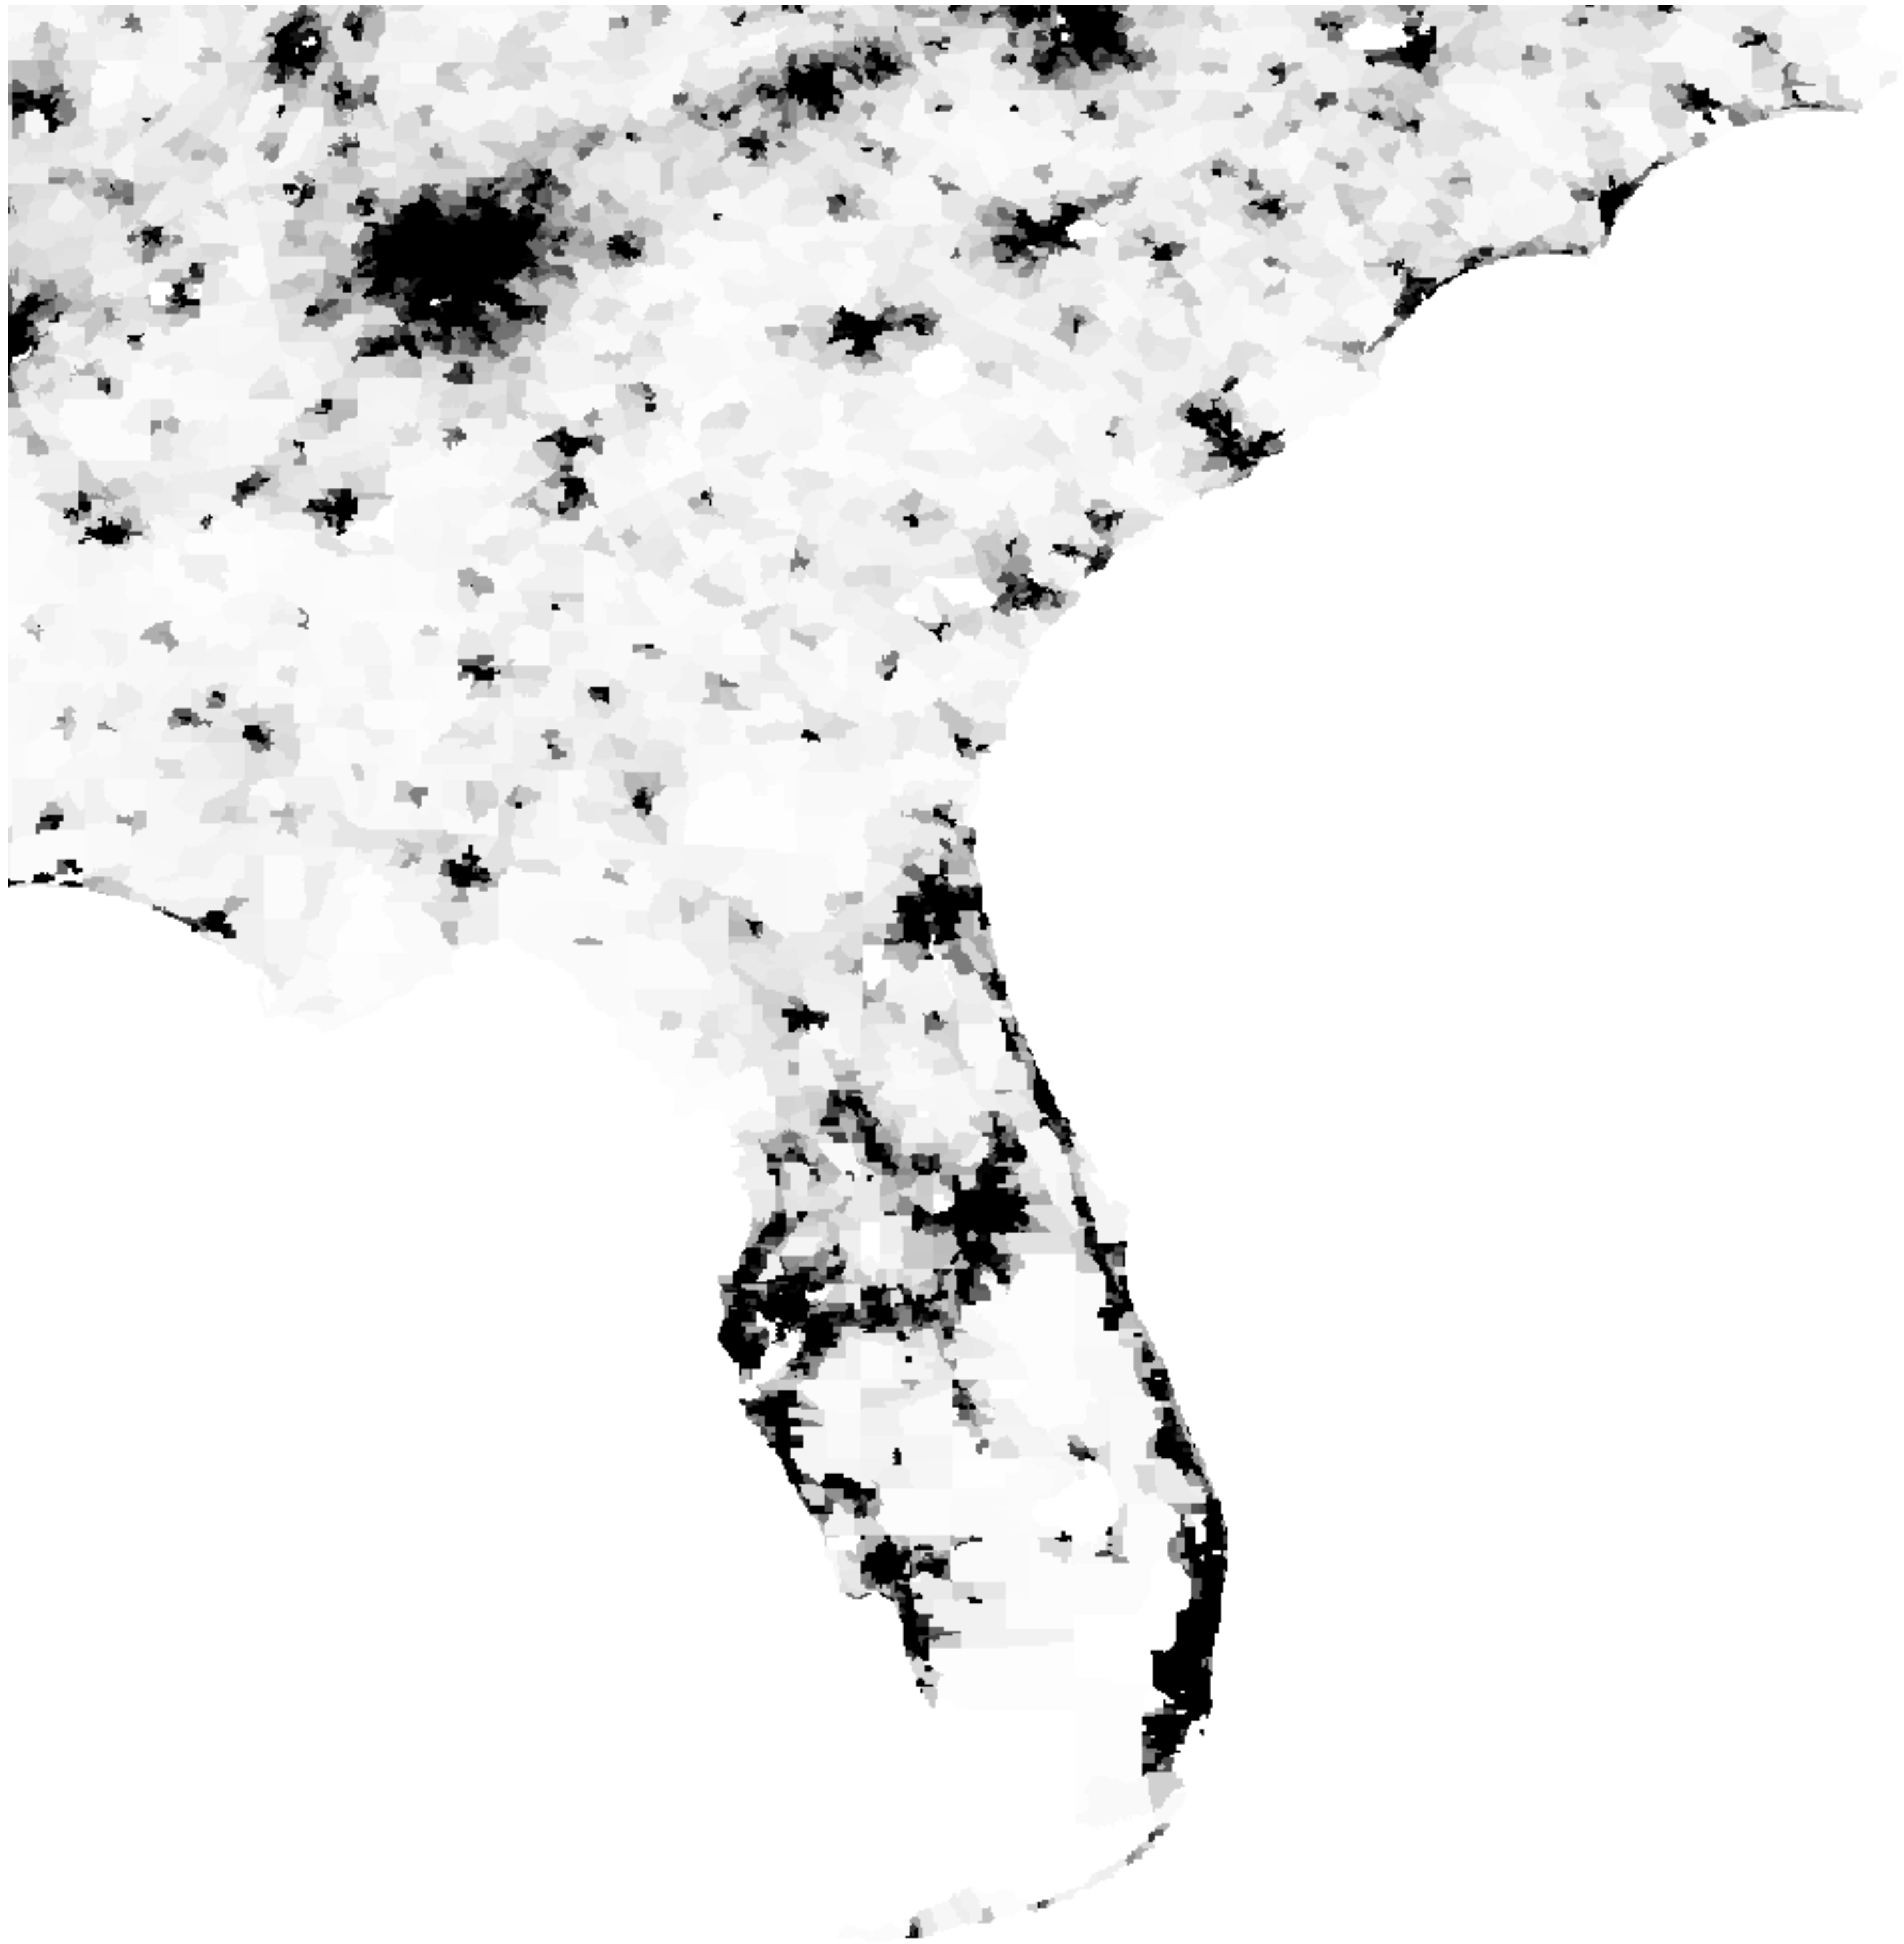

Census ground truth data, 2010

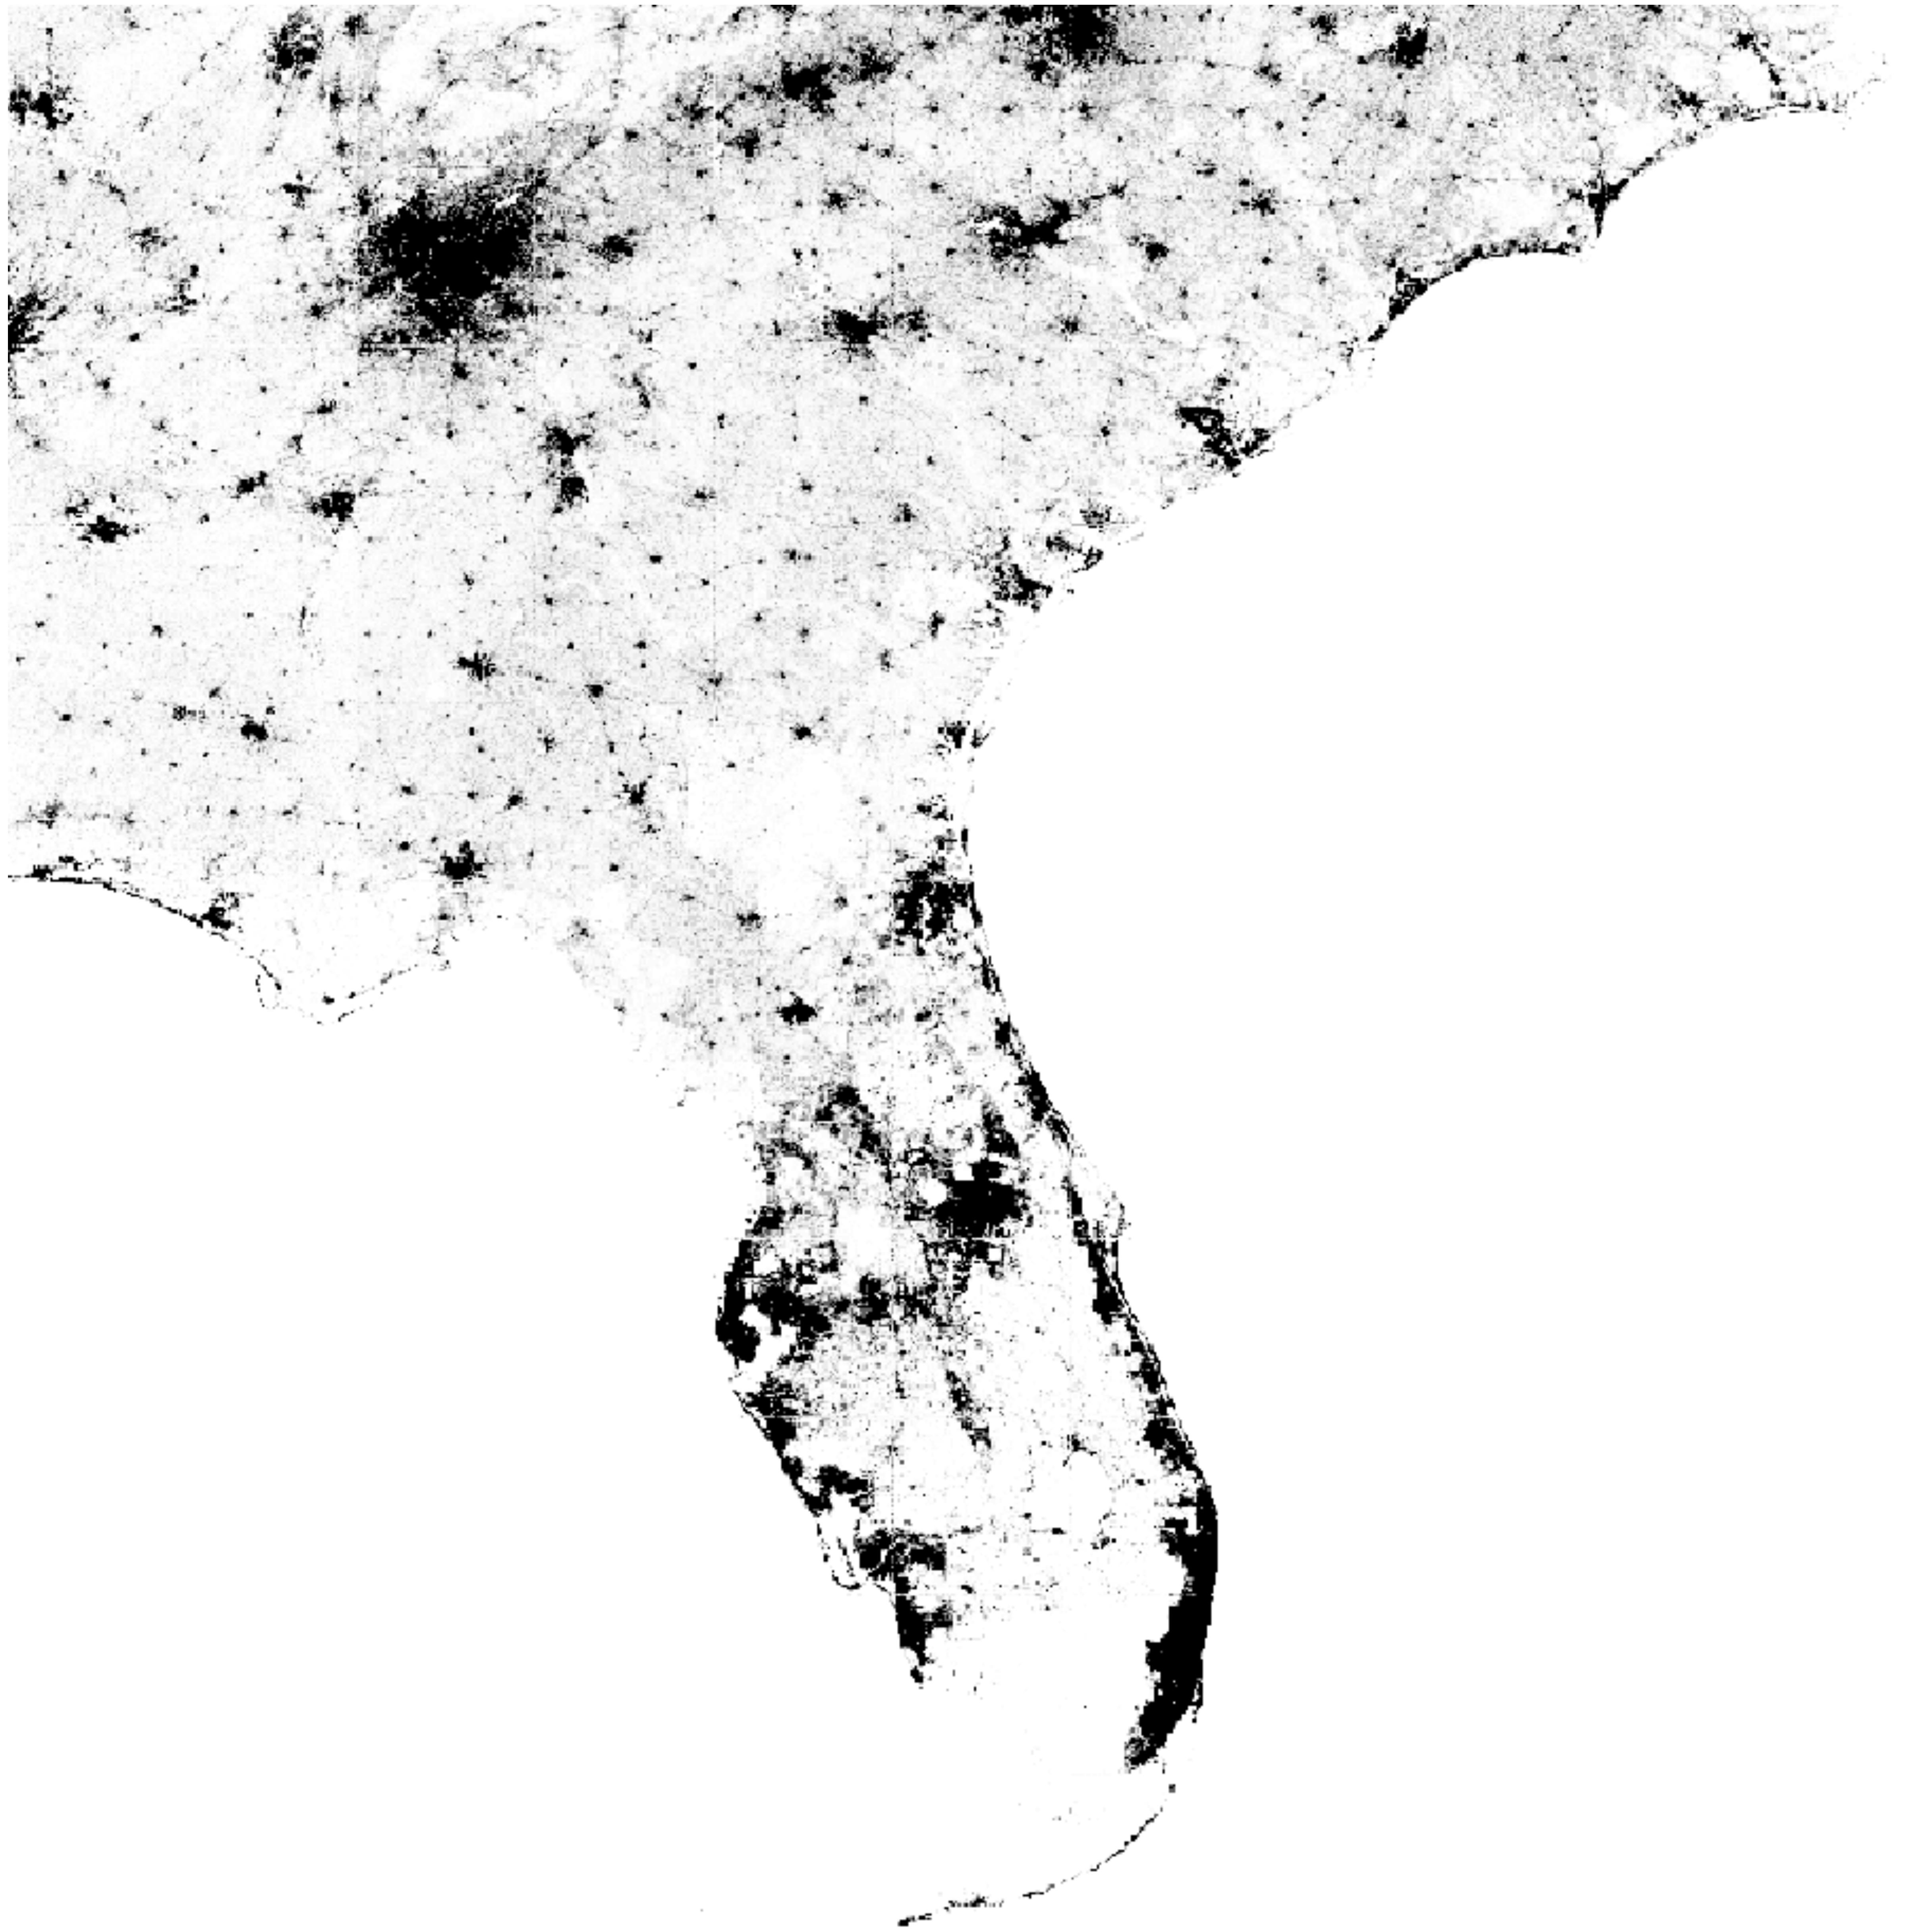

CONVRAW predictions, 2010

Qualitative Interpretation

1. *Which satellite images are the models most confident about?*

2. *Where do the models predict low —> high population areas are?*

3. *What errors are the models making?*

1. *Which satellite images are the models most confident about?*

Top 8 most confident test images for each population class.

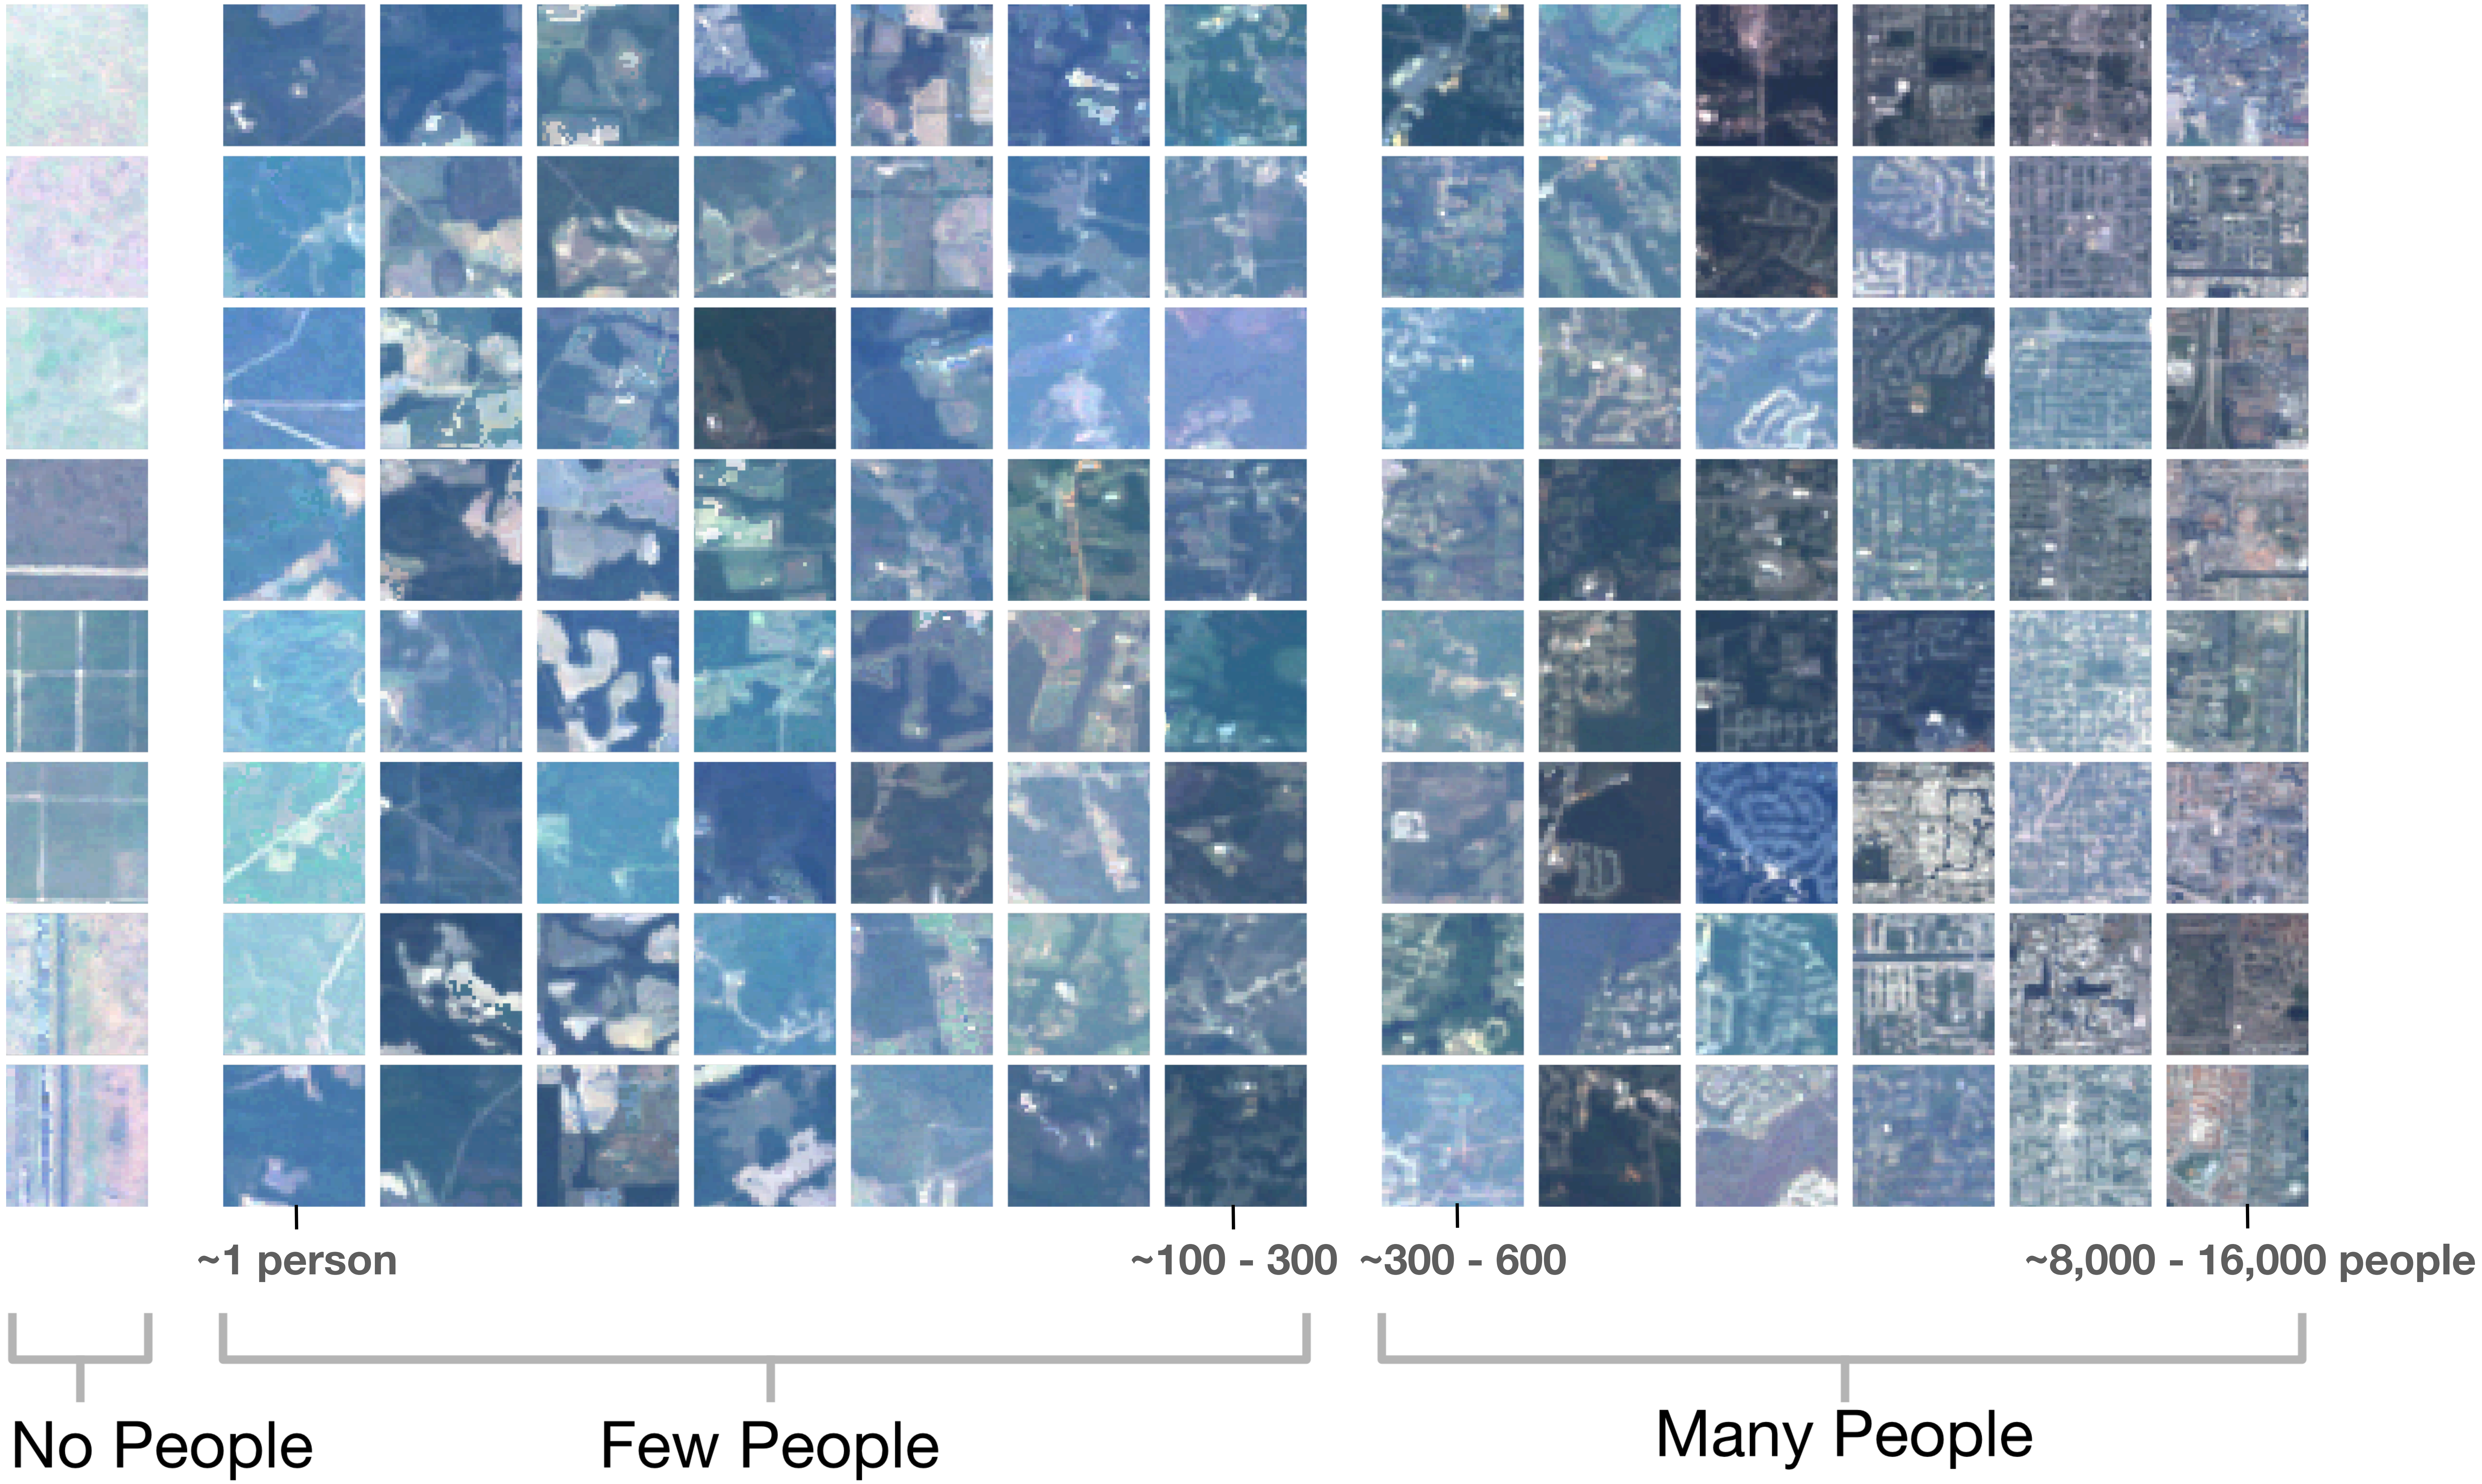

2. Where do the models predict low —> high population areas are?

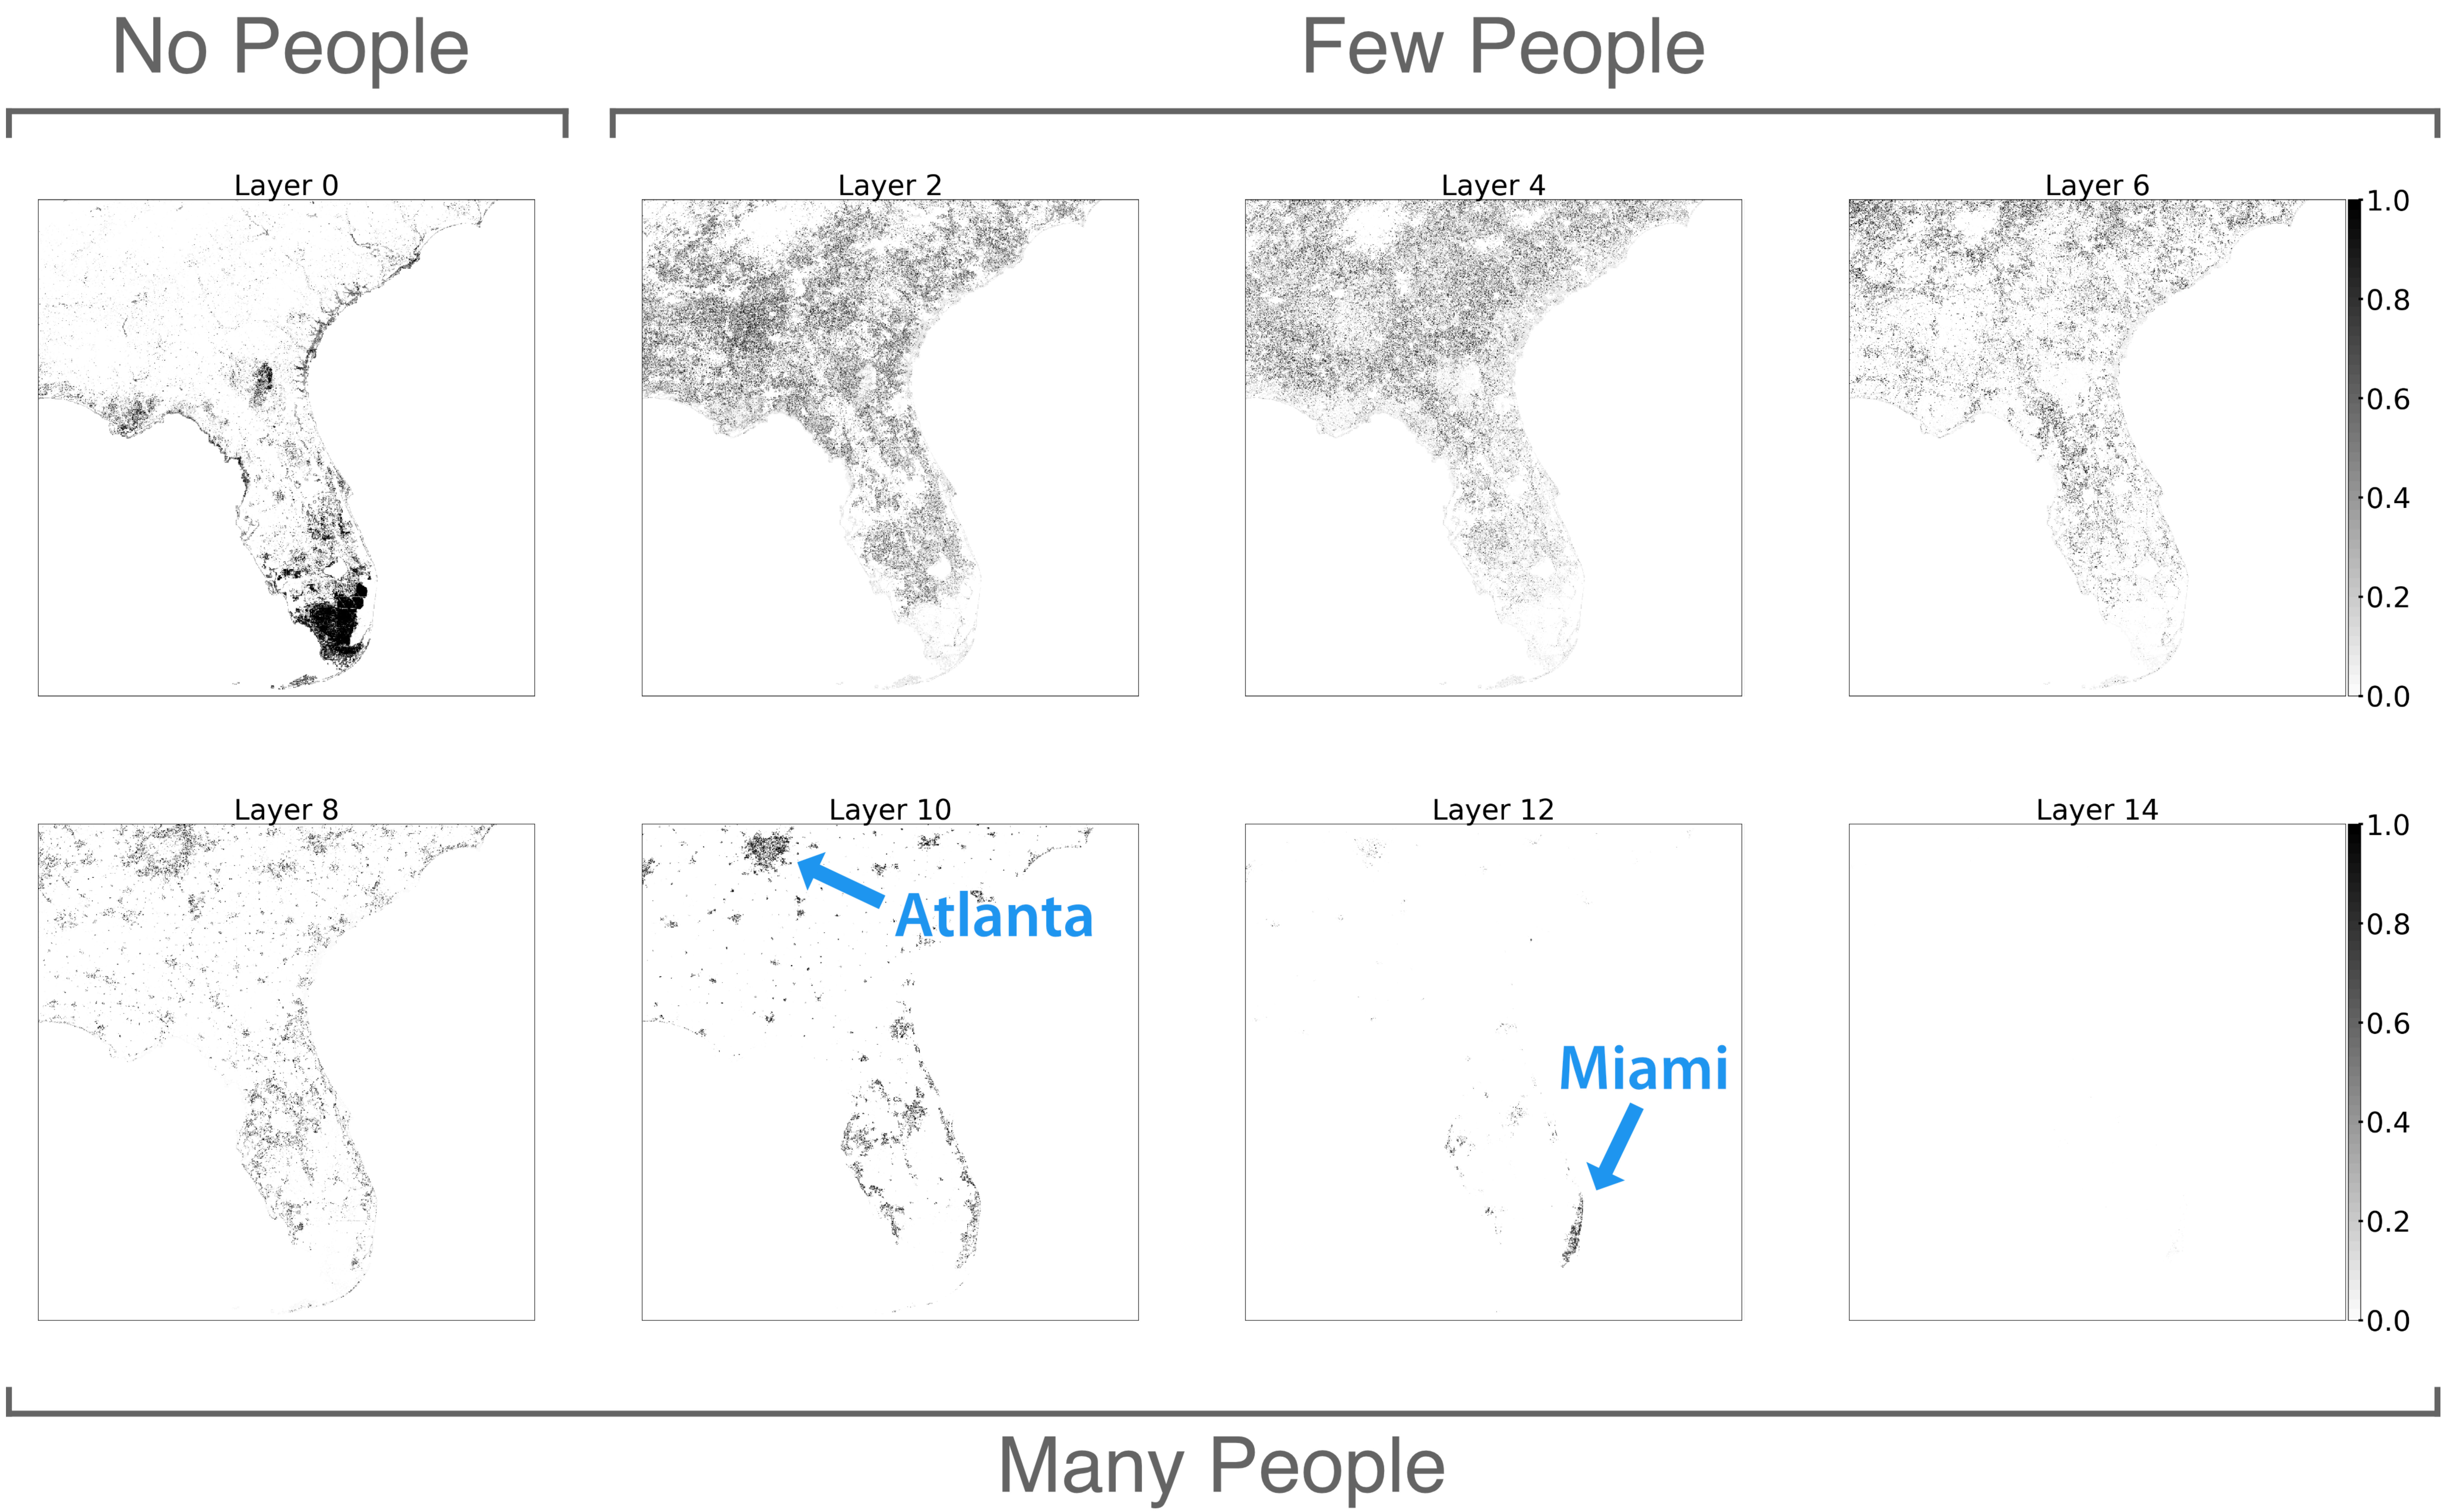

3. What errors are the models making?

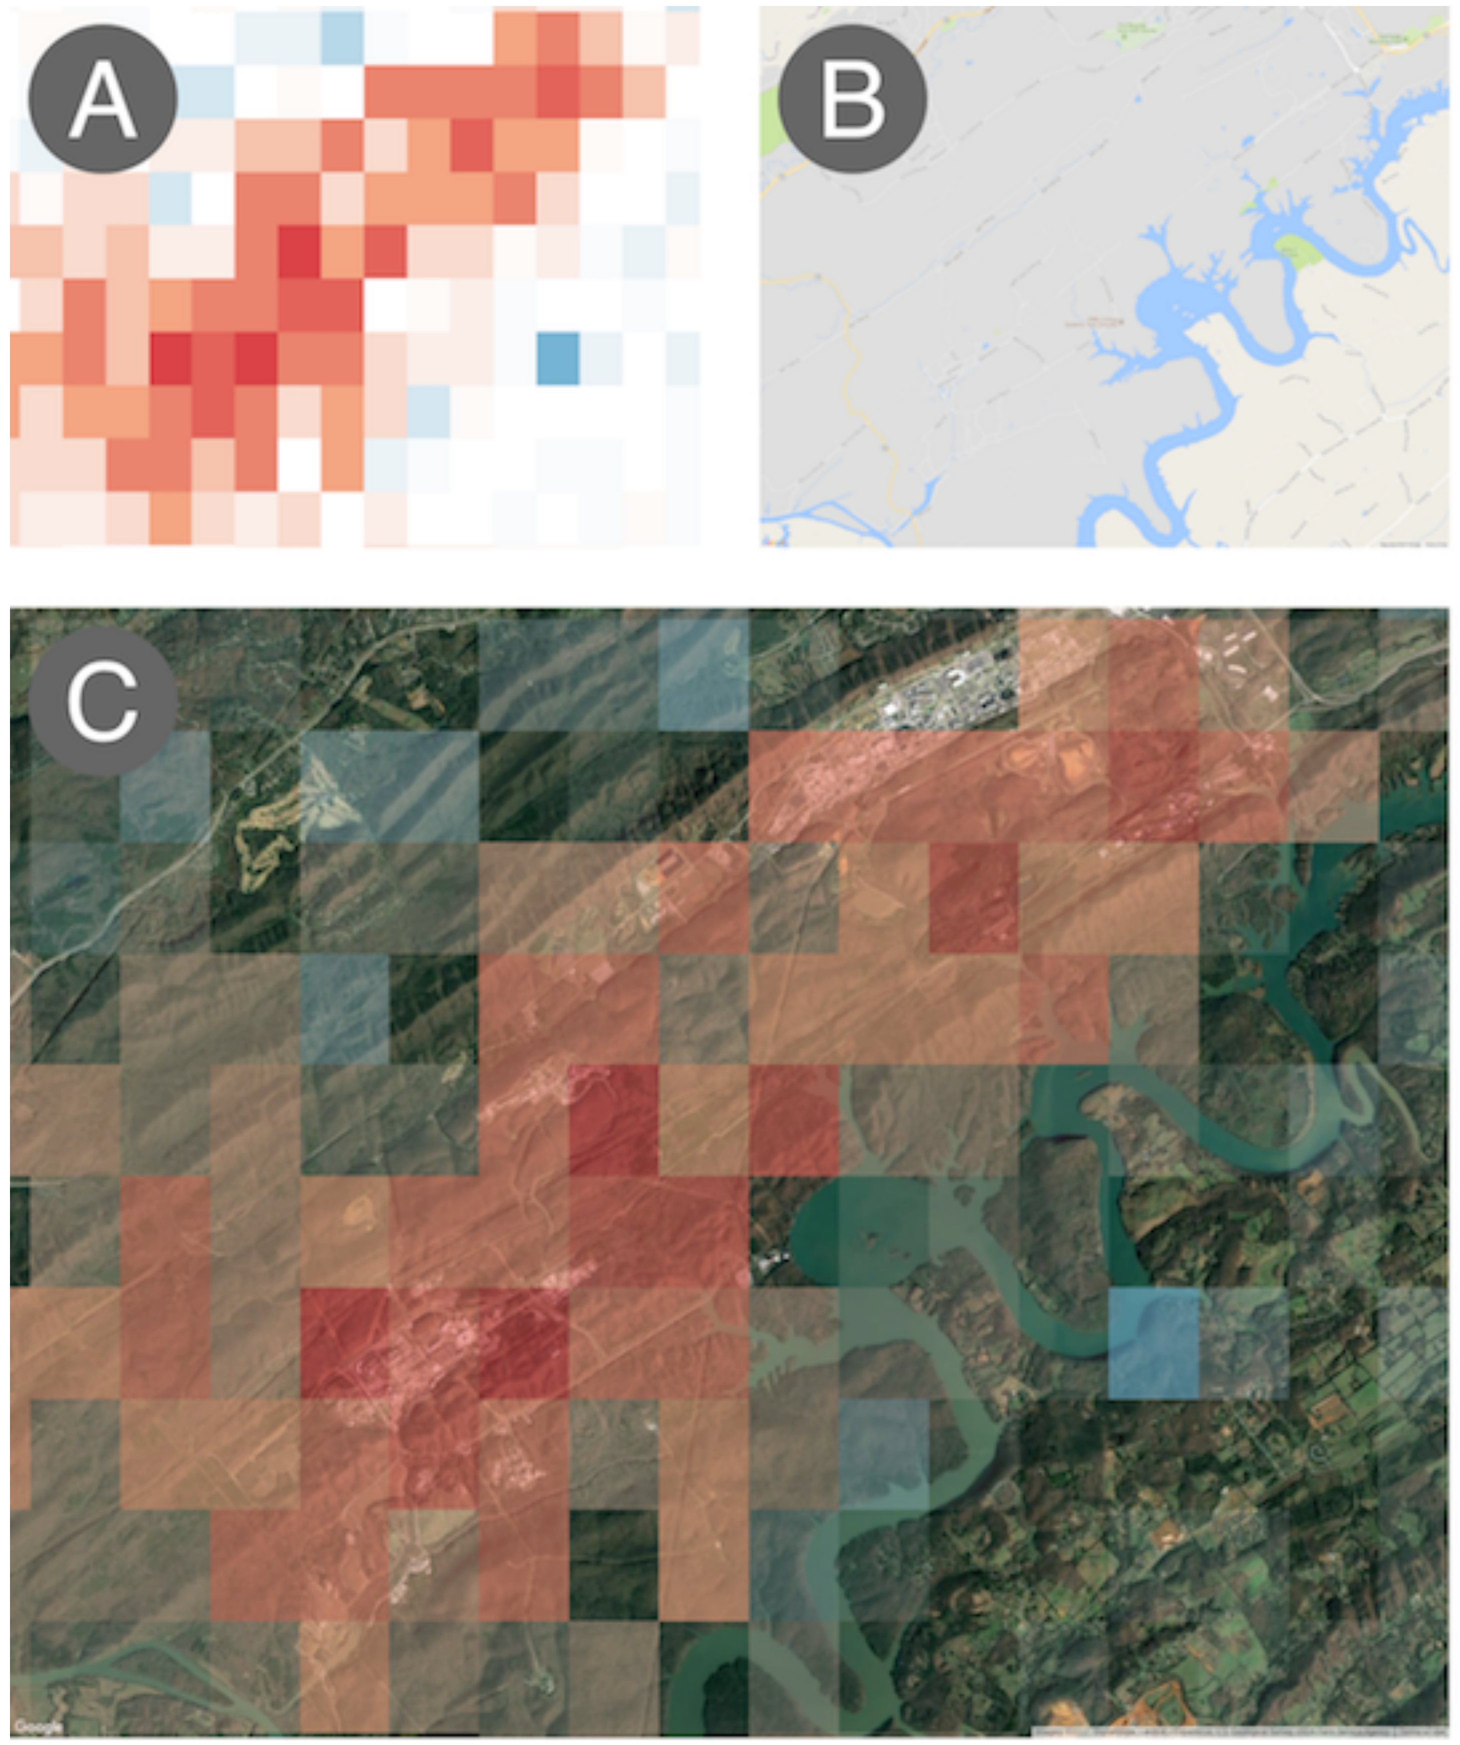

Oak Ridge National Lab
Oak Ridge, TN
Small Scale

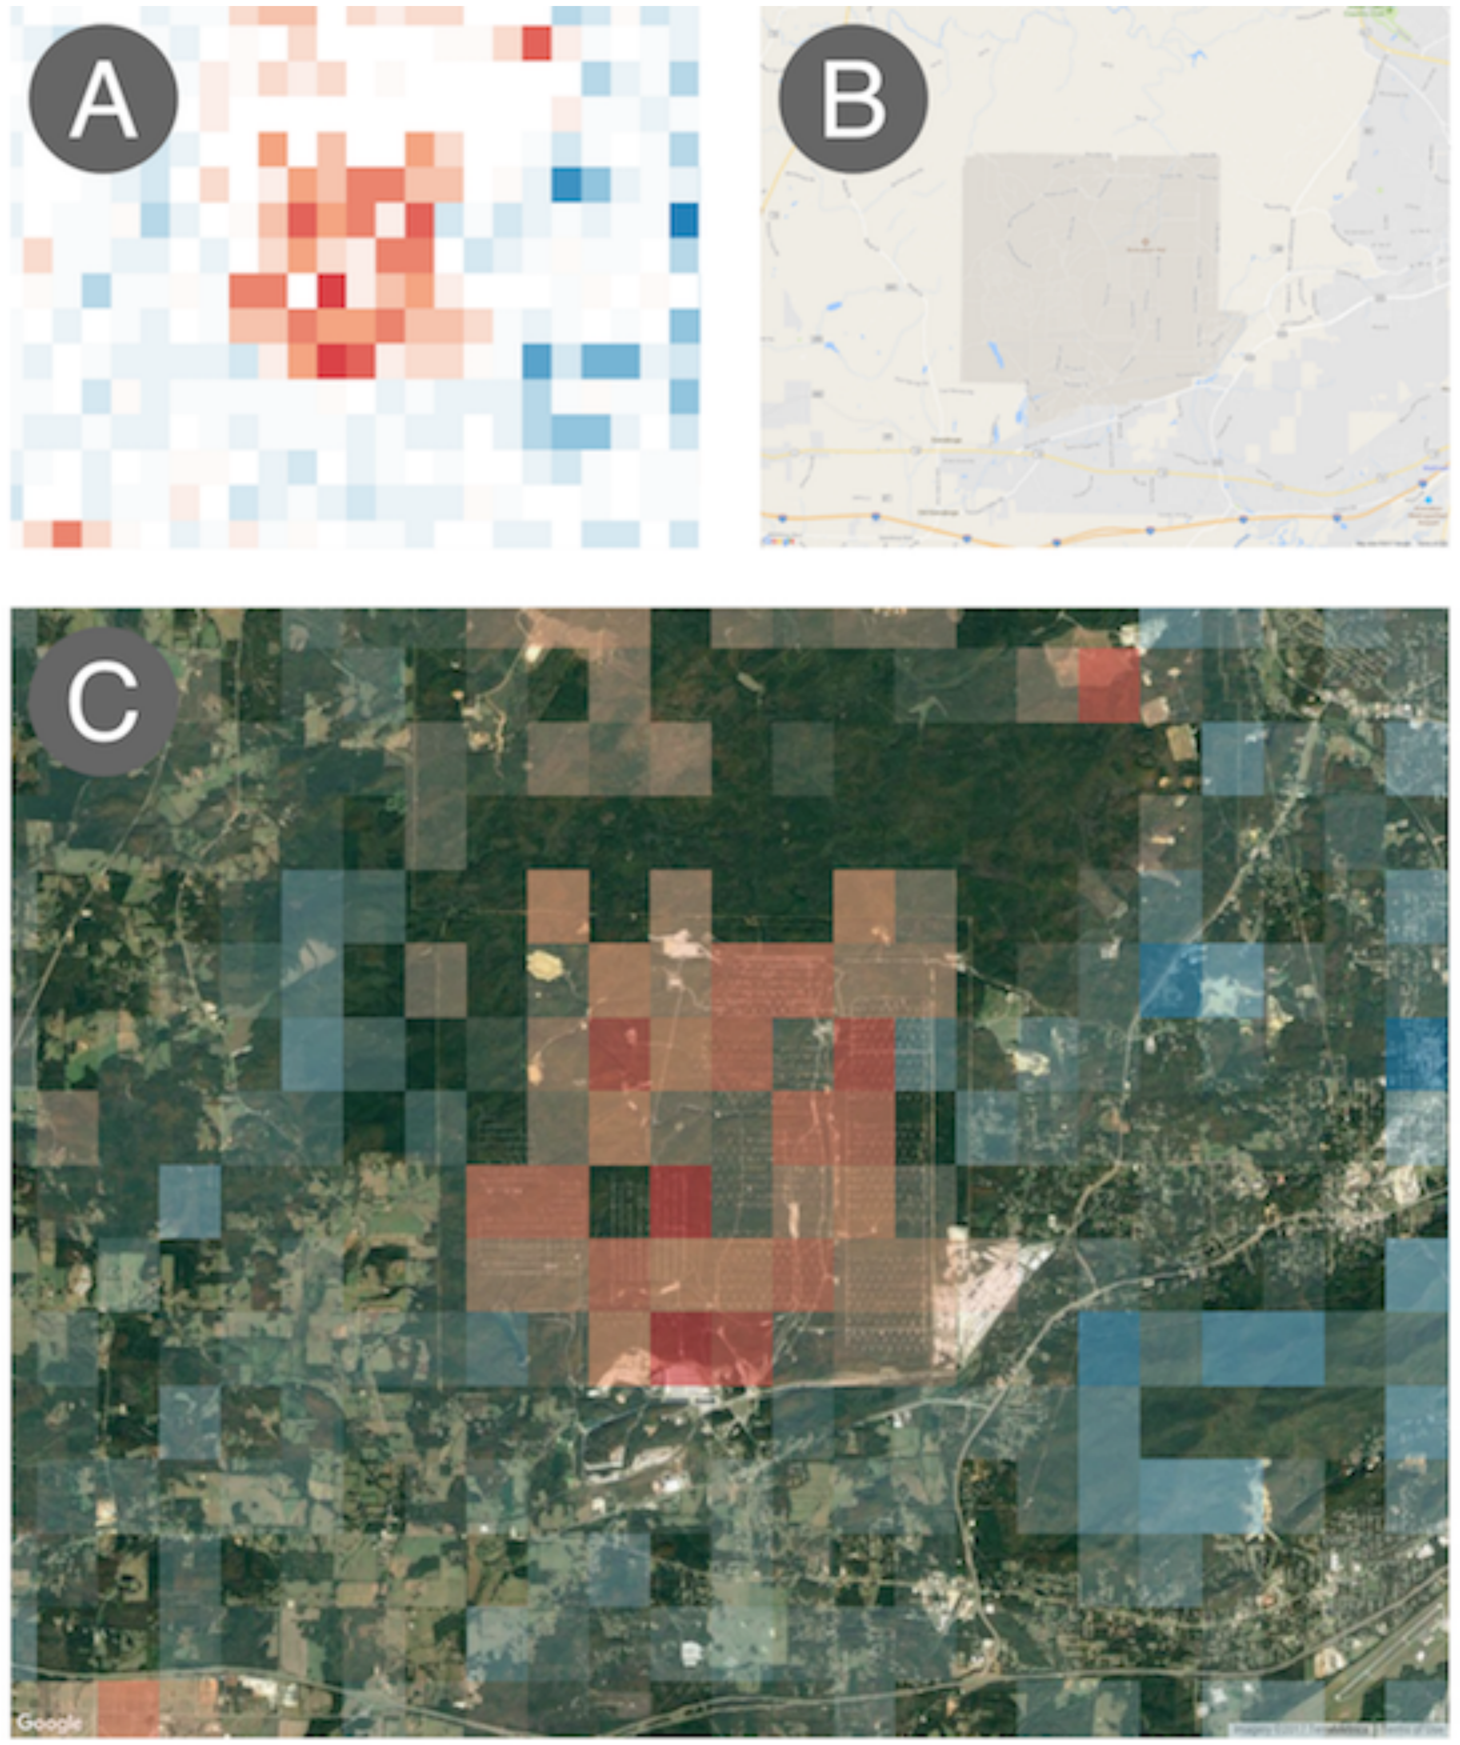

Anniston Army Depot
Anniston, AL
Medium Scale

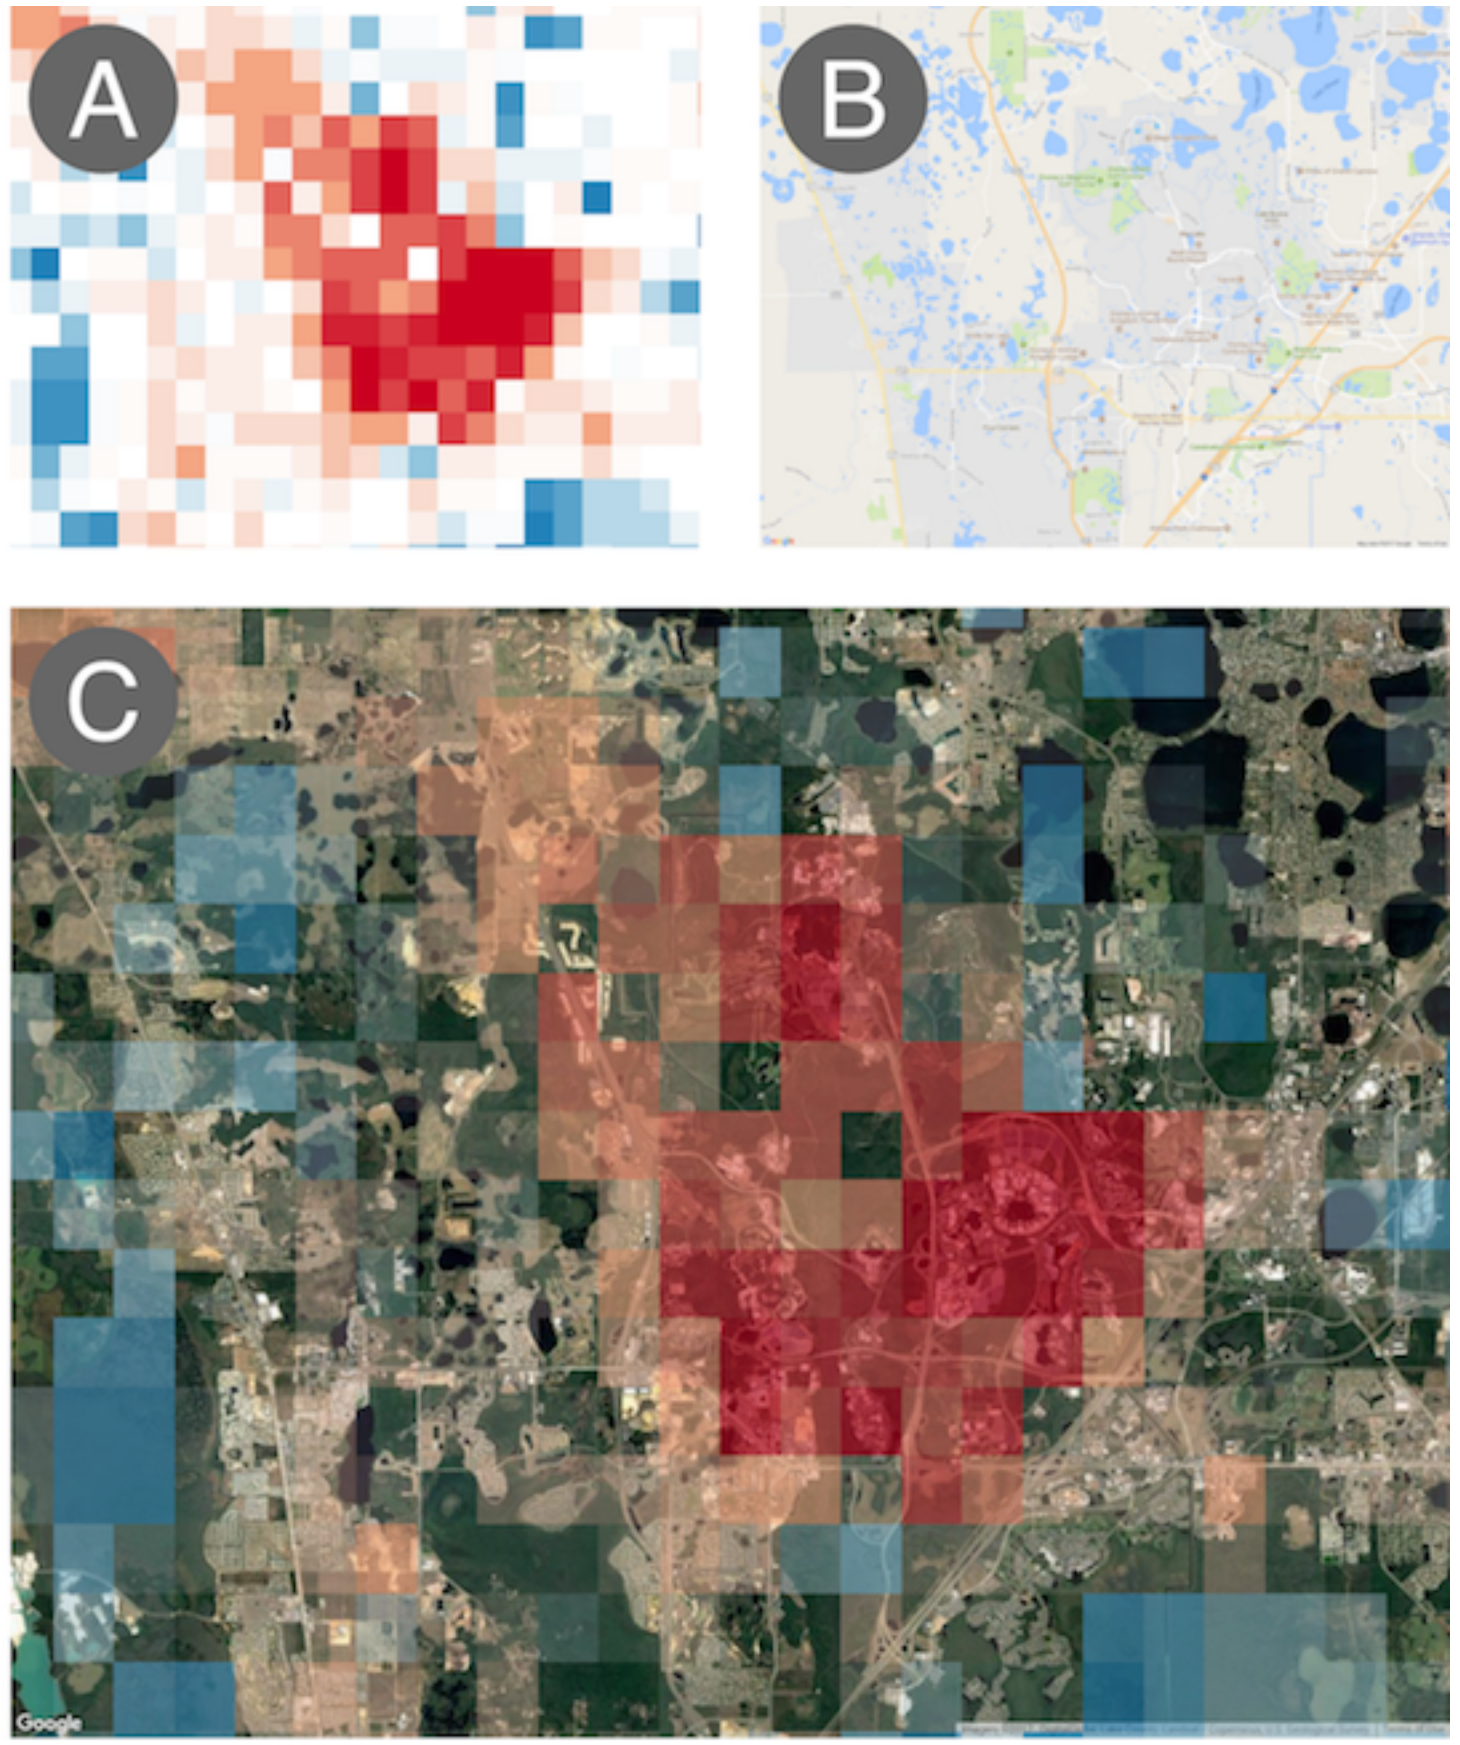

Walt Disney World
Orlando, FL
Large Scale

Future Work

Can we improve the accuracy of the models?

- More recent neural network architectures
- Effects of imperfect training data
- Loss function improvements

Do our models generalize in space?

- Test a model trained on the US in different countries?

Can we predict other socio-economic values?

A Deep Learning Approach for
Population Estimation from Satellite Imagery

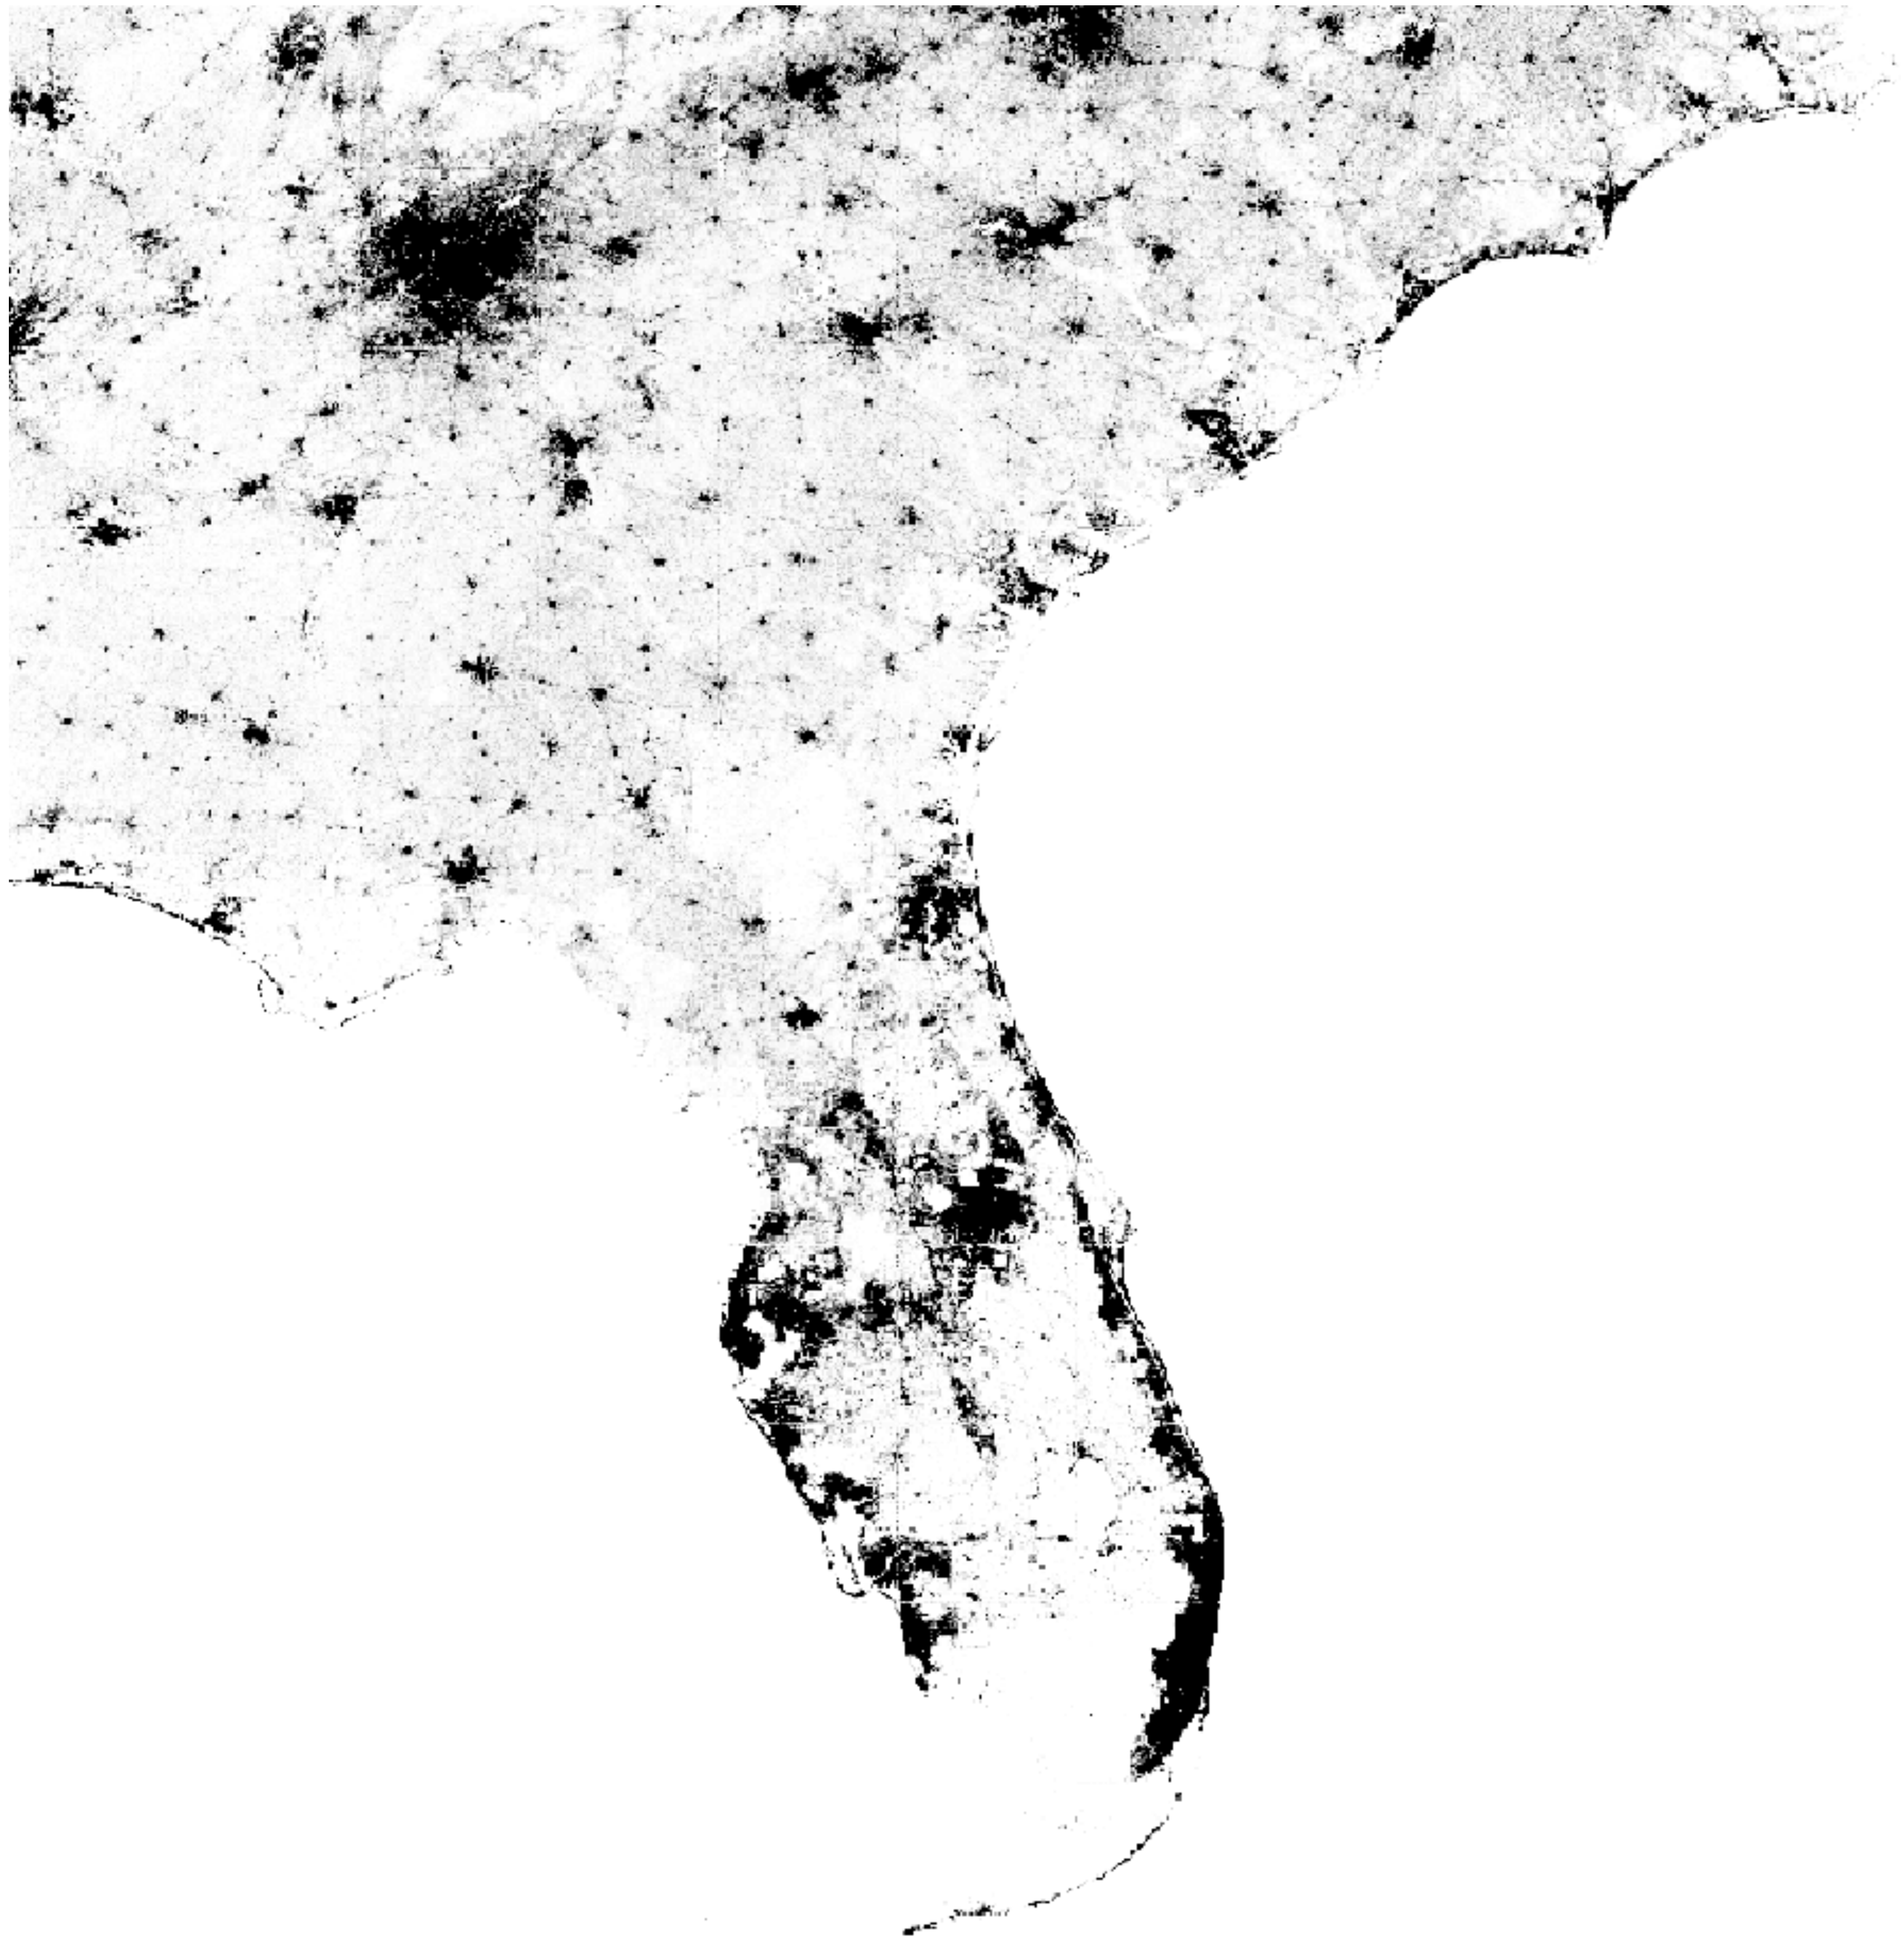

Thanks!

Caleb Robinson
calebrob.com
dcrobins@gatech.edu

Fred Hohman
fredhohman.com
fredhohman@gatech.edu

Bistra Dilkina
cc.gatech.edu/~bdilkina
bdilkina@cc.gatech.edu

We thank NSF for funding this work and the anonymous reviewers for their constructive feedback.

Bonus

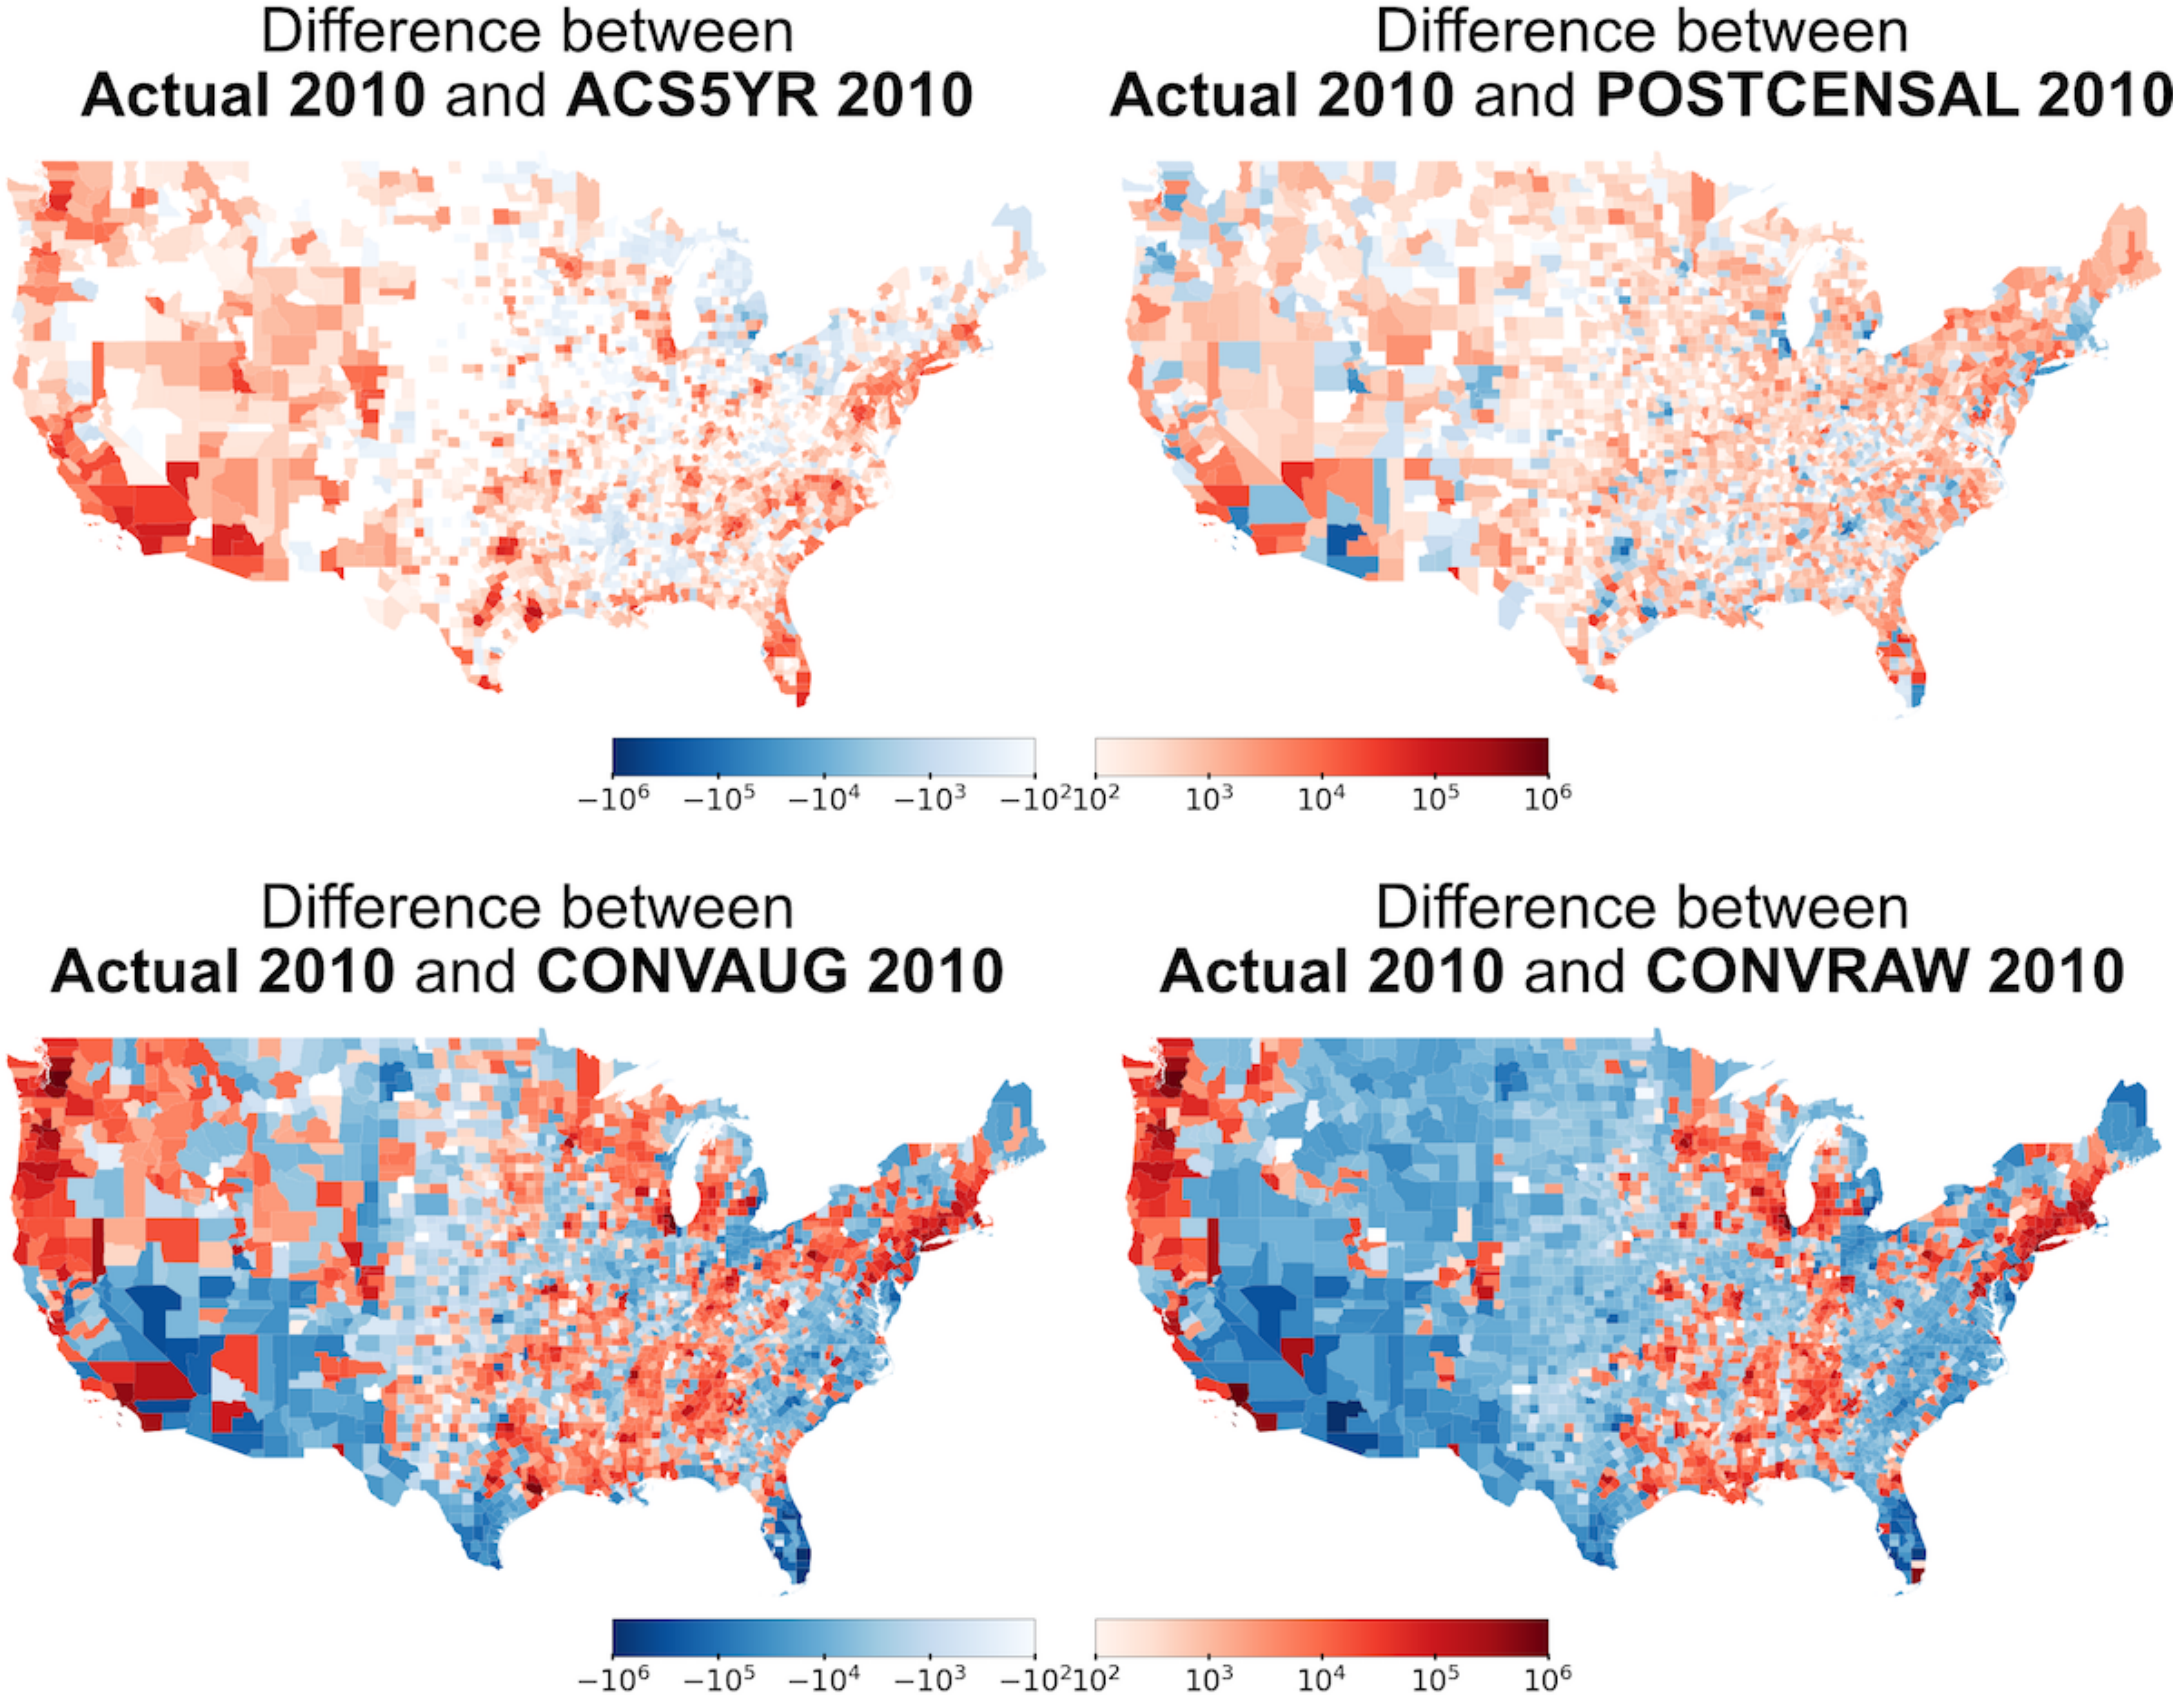
Difference between
Actual 2010 and ACS5YR 2010
Difference between
Actual 2010 and POSTCENSAL 2010
-10^6 -10^5 -10^4 -10^3 -10^2
10^2 10^3 10^4 10^5 10^6
Difference between
Actual 2010 and CONVAUG 2010
Difference between
Actual 2010 and CONVRAW 2010
-10^6 -10^5 -10^4 -10^3 -10^2
10^2 10^3 10^4 10^5 10^6