Using Computer Vision to further econometric analysis: A case study with predicting Urban Heat Island. By Jake Silberg

Introduction: Urban Heat Island

- Urban Heat Island (UHI), increased heat experienced by urban areas vs. nearby rural areas, is believed to have contributed to many of 10,000 heat-related deaths in the US between 2004 and 2018
- The effects are unequally distributed – People of Color live in census tracts with more UHI effects than non-Hispanic whites
- Known predictors of UHI include pavement (positively correlated), vegetation (negatively correlated)
- Few have tried to directly predict UHI directly from visible satellite bands
- Identifying how the model predicts UHI could help us understand factors that drive UHI in real-world

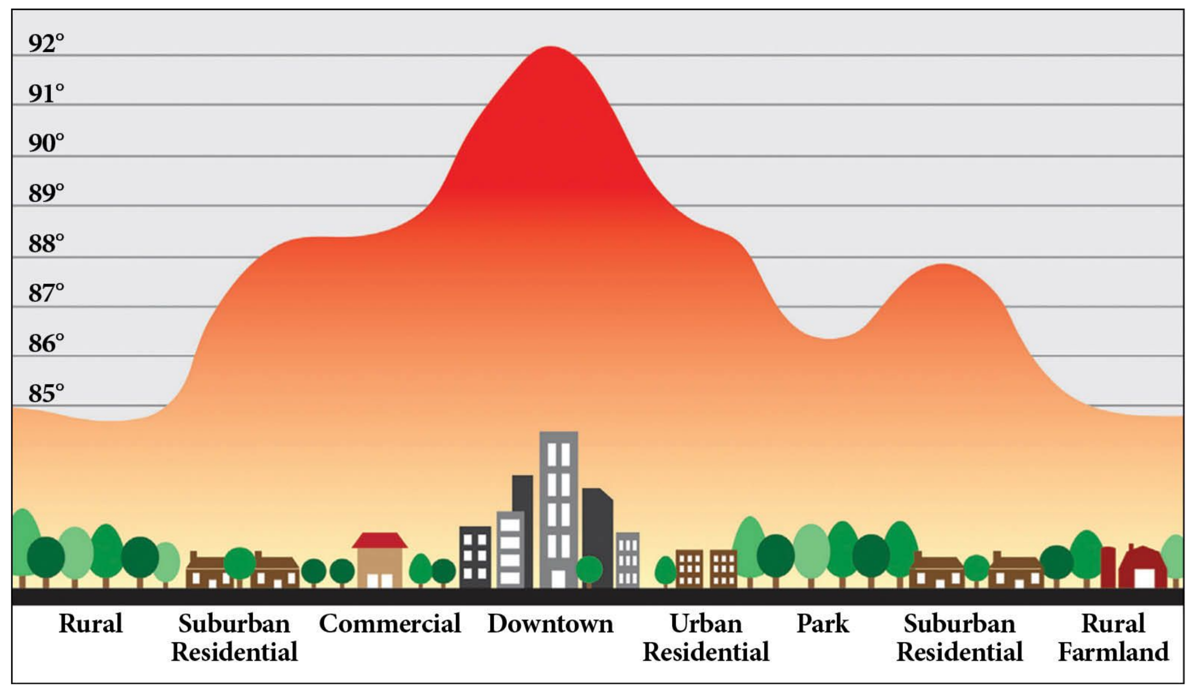

Dataset

- Yale researchers have prepared 55,000 UHI observations across the country
- I downloaded images of the 5,000 most populated areas, as these are the smallest Census Tracts so easiest to find local effects

Optimizing the model

- Regression, rather than classification problem
- Normalized UHI scores and images (based on ImageNet)
- Tested several architectures:
  - VGG19
  - ResNet50
  - Modified ResNet with additional FC layer
  - Modified ResNet with two additional FC layers
- Hypothesized that a traditional ResNet pooling layer would underperform for regression

Traditional ResNet can be used for regression:

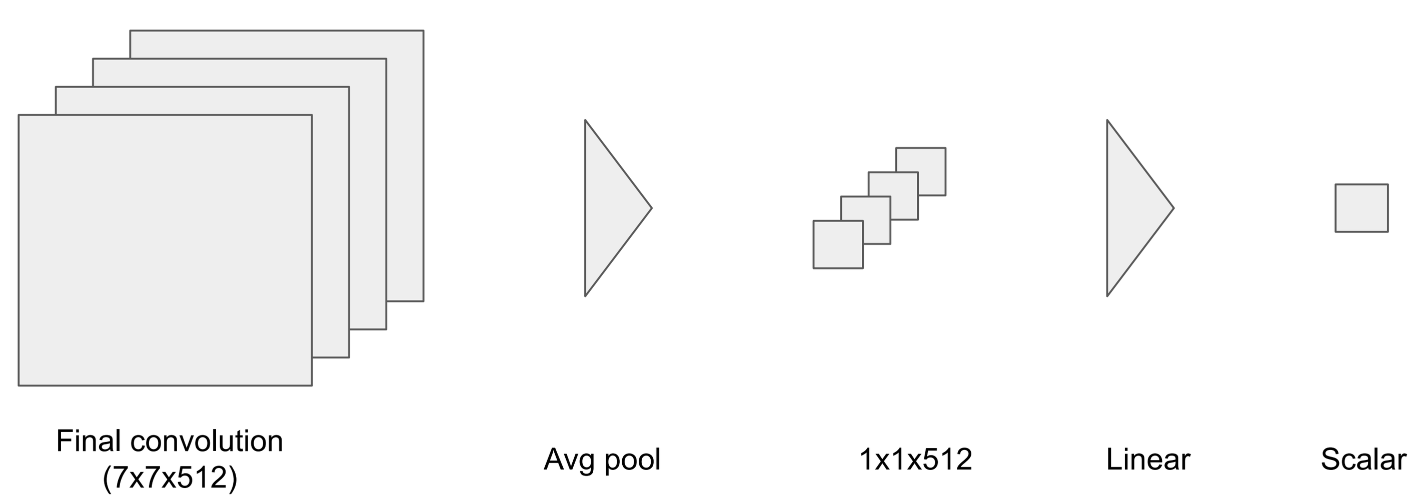

Modified ResNet for regression:

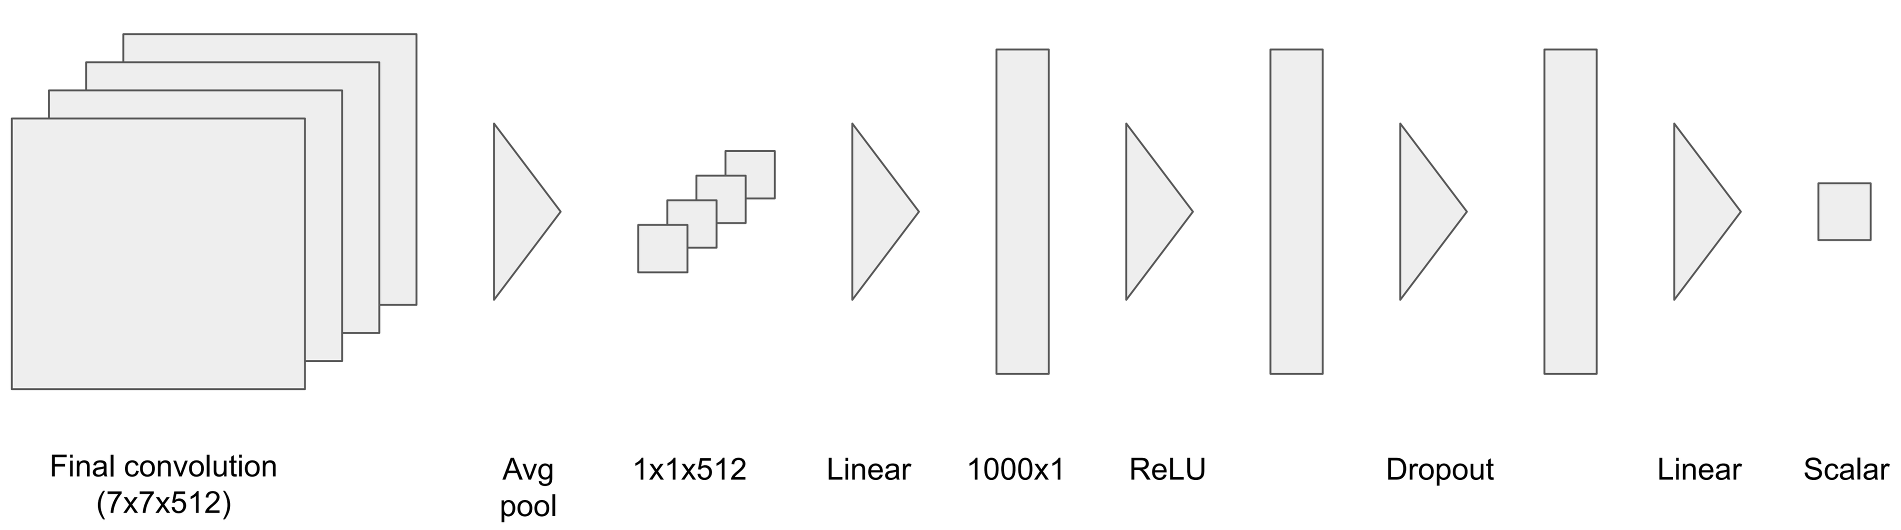

Model performance

One additional FC layer helps performance, but two layers (bottom row) hurts performance

Model	**R^2**
Color histograms	0.44059
VGGNet	0.46703
ResNet (baseline)	0.45422
Modified ResNet (512x1000 linear, ReLU, dropout, 1000x1 linear) with MSE loss	0.47013
Modified ResNet (512x1000 linear, ReLU, dropout, 1000x1 linear) with L1 loss	**0.48472**
Modified ResNet (512x1000 linear, ReLU, dropout, 1000x256, ReLU, dropout, 256x1linear) with MSE loss	0.45967

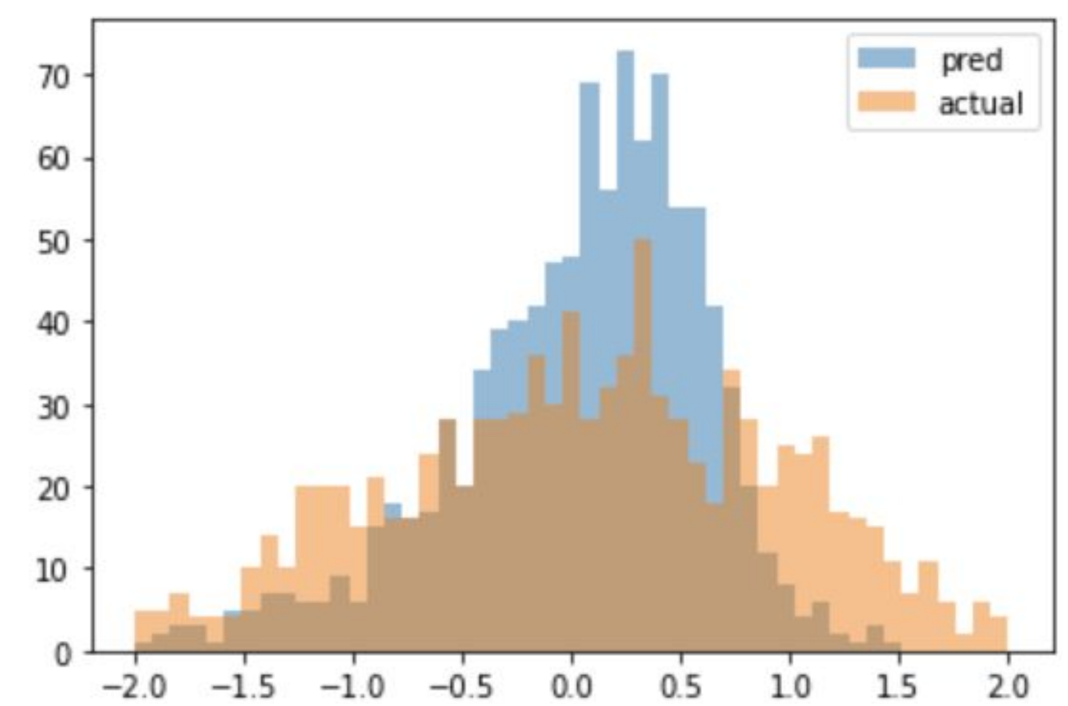

Traditional ResNet shrinks validation set predictions to 0 (above), while training distributions match (below)

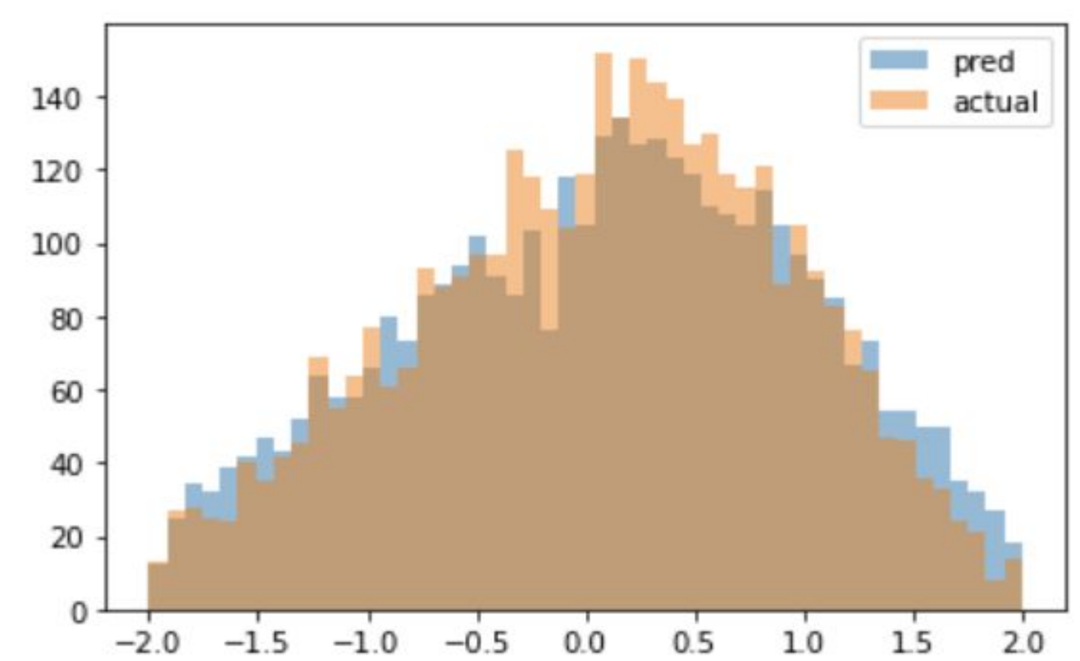

Modified ResNet with L1 loss shrinks slightly less on validation set (below)

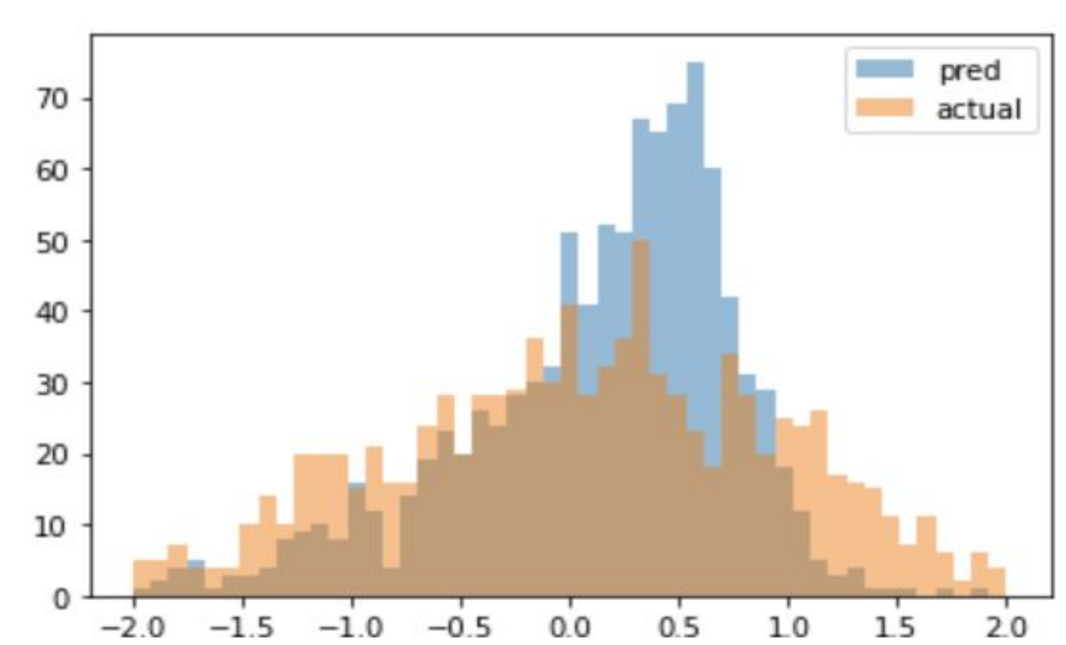

Understanding drivers of model scores

Lowest (left) and highest (right) ground-truth UHI scores in validation set

- Qualitative visualization techniques can be hard to draw conclusions from: Saliency map seems to ignore road, but LIME (perturbation-based method for understanding model decisions) shows road is key to pushing up UHI score

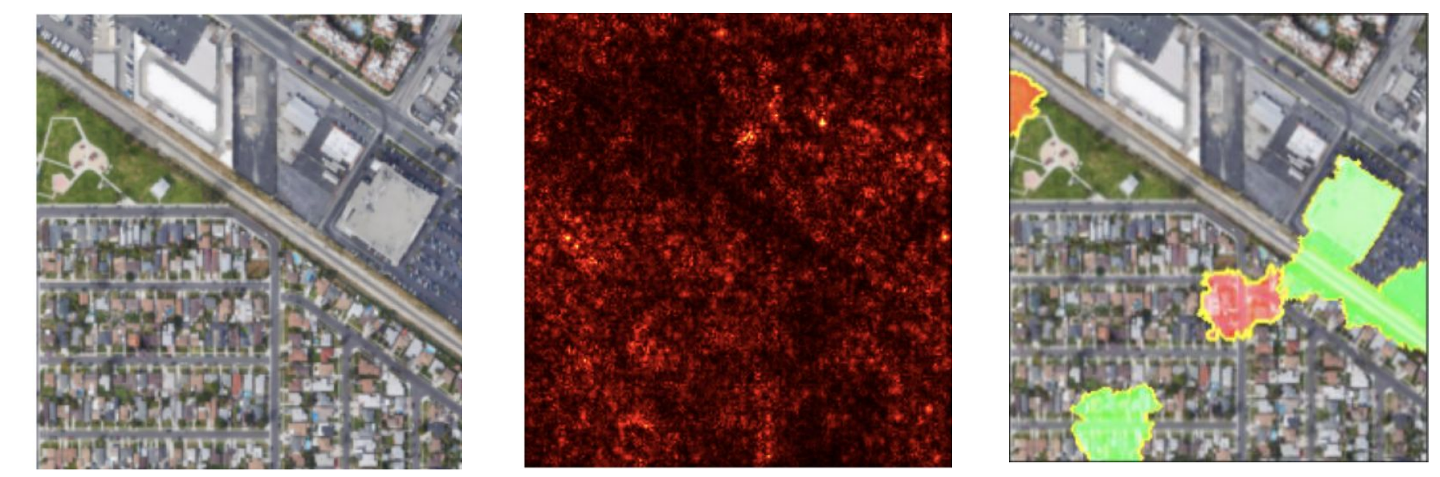

- I compared nearest neighbor images (in embedded space) with largest UHI discrepancies. Middle image has higher predicted UHI than left image, and, LIME shows the highway (pavement) is pushing up UHI score, as hoped

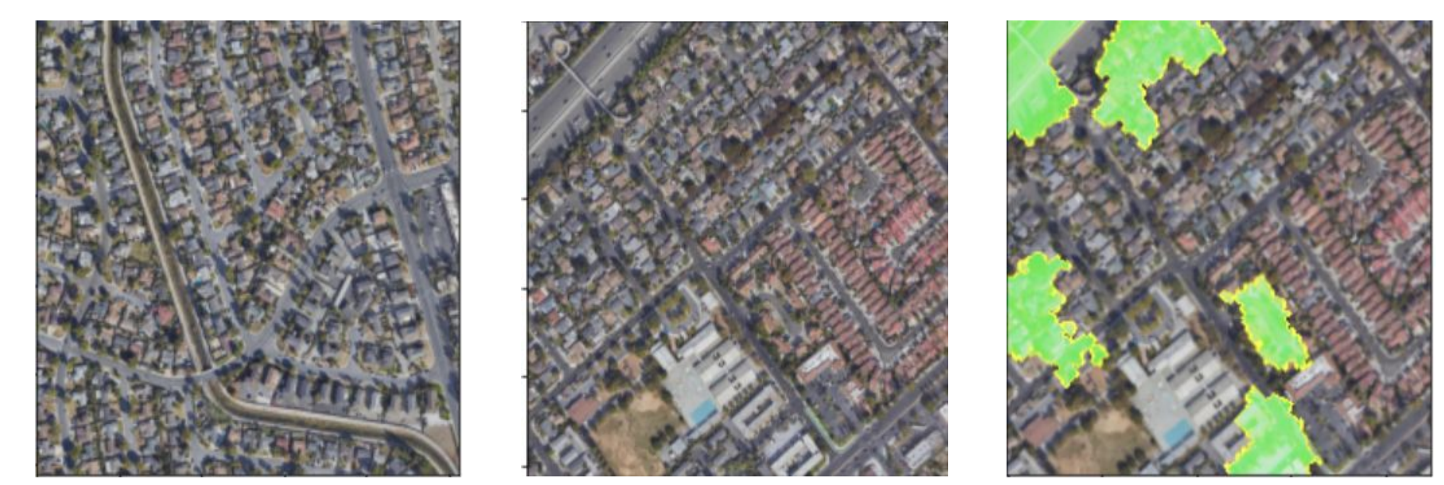

- Quantitative: amount of highway pavement in an image was NOT really correlated with either predicted or ground-truth UHI

Variable	**Coefficient**	**p-value**
Ground-truth UHI	0.0005	0.986
Predicted UHI	0.0210	0.304

Learnings

- Architectures that perform best for classification may benefit from modifications for regression
- Qualitative feature understanding techniques may produce inconsistent or even contradictory results
- More consistent quantitative measures are needed (preservation of known predictors is one approach even though inconclusive in this case)