Automated Building Energy Consumption Estimation From Aerial Imagery

Duke UNIVERSITY | BASS CONNECTIONS

Researchers: Mitchell Kim (BSE '18), Sebastian Lin (BS '18), Sophia Park (BSE '17), Eric Peshkin (BS '18), Nikhil Vanderklaauw (BSE '18), Yue Xi (BS '19)

Project Managers: Samit Sura (MA '17), Hoël Wiesner (MEM '18)

Faculty Advisors: Kyle Bradbury (Energy Initiative), Leslie Collins (Pratt), Timothy Johnson (Nicholas)

Bass Connections in Energy

Overview

Detailed building-level energy consumption data for cities is rare due to the prohibitive collection costs, but could be used to identify and plan infrastructure and policy developments. Autonomous detection of objects such as buildings, roads, power lines, and pipelines can be useful for policy makers to map infrastructure, track development patterns over time, find indicators of economic activity, or quickly assess environmental damages. Recent advances in computation for big data and image processing now make it possible to learn about energy use in a fast and automated manner using machine learning.

This project aims to estimate building-level energy consumption from high resolution aerial imagery by: (1) identifying buildings and extracting their properties (size, perimeter, etc.) and (2) inputting these properties into an energy consumption estimation model designed using existing energy consumption data from the Department of Energy.

To find an effective building detection technique, we implement both a traditional and a state of the art deep learning classifier. We then apply our workflow to Gainesville, FL to assess its effectiveness.

Classifier Performance Comparison

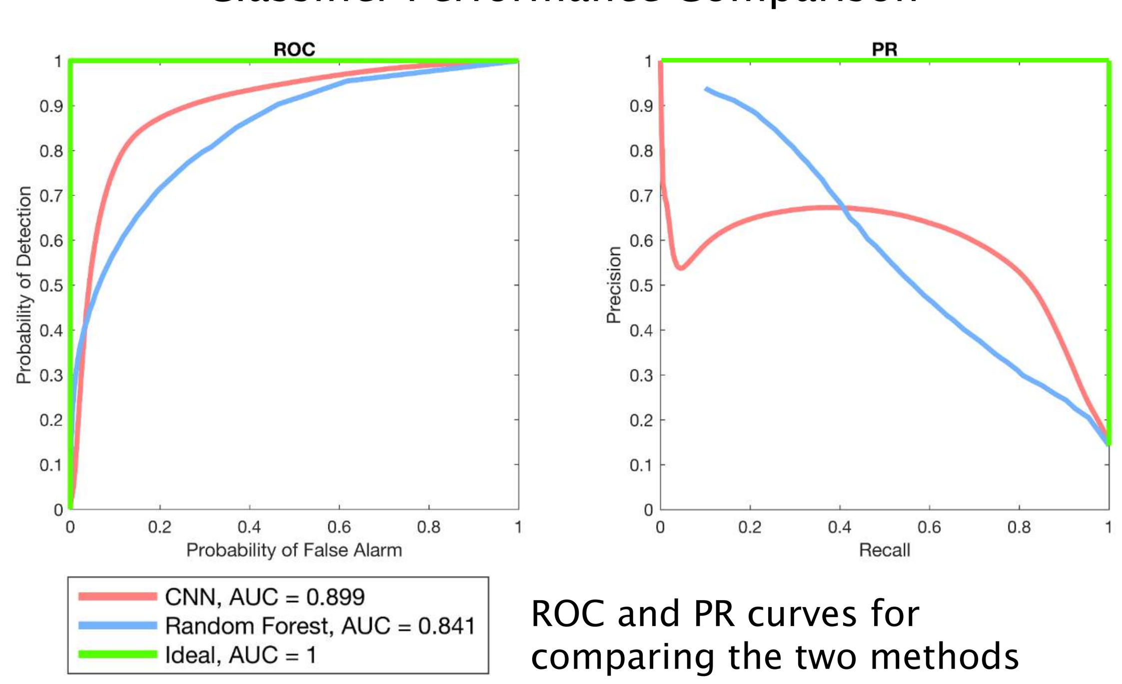

ROC and PR curves for comparing the two methods

Process Summary

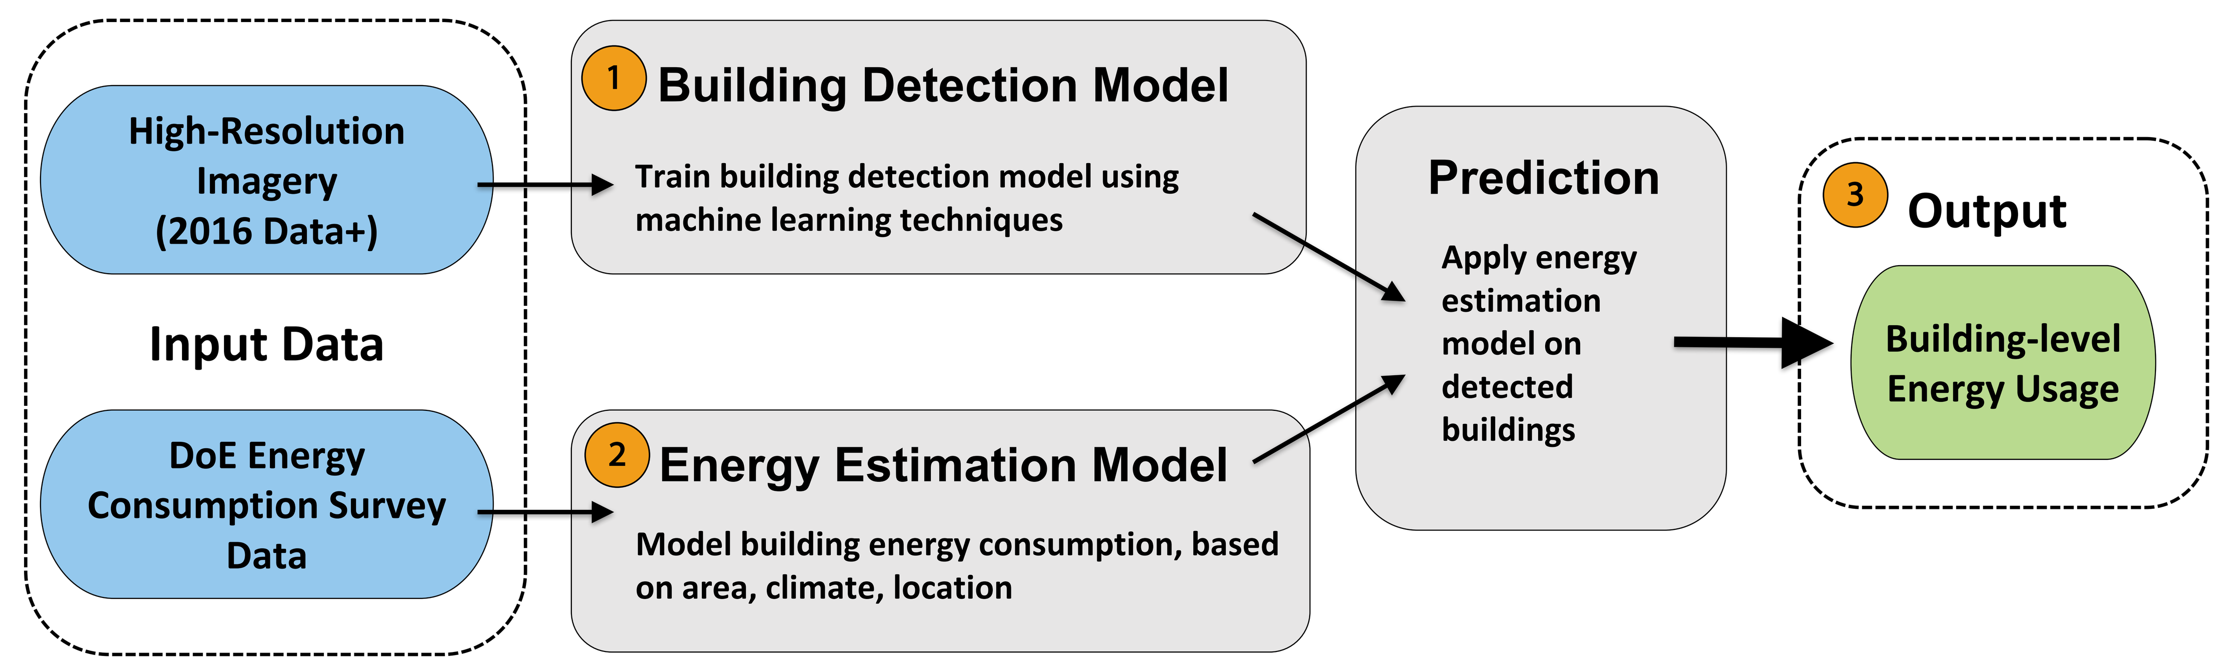

1 Building Detection

To develop our building detection algorithms, we use a dataset of 25 high-resolution (0.3m or finer) orthorectified images from 9 different US cities[1], compiled by a 2016 Data+ research team. We developed an algorithm to detect buildings by training a random forest (RF) classifier to identify building pixels in images. To compare performance between the classifiers, we then applied a convolutional neural network (CNN)[3] to generate estimates from a cutting-edge machine learning algorithm. The detected buildings from the CNN are compared to the RF and the ground truth below by testing the algorithms on an image of Norfolk, VA. The CNN performed better, evidenced by a higher AUC score, and as seen in the ROC and PR curves below. The CNN tends to classify buildings in the right location but the borders are often inaccurate, while the random forest tends to have more regions containing false alarms with truer borders on targets. Once we calculate the area of each of the detected buildings, we feed it into our energy-consumption estimation model.

Ground truth building outlines, i.e., the ideal classification output

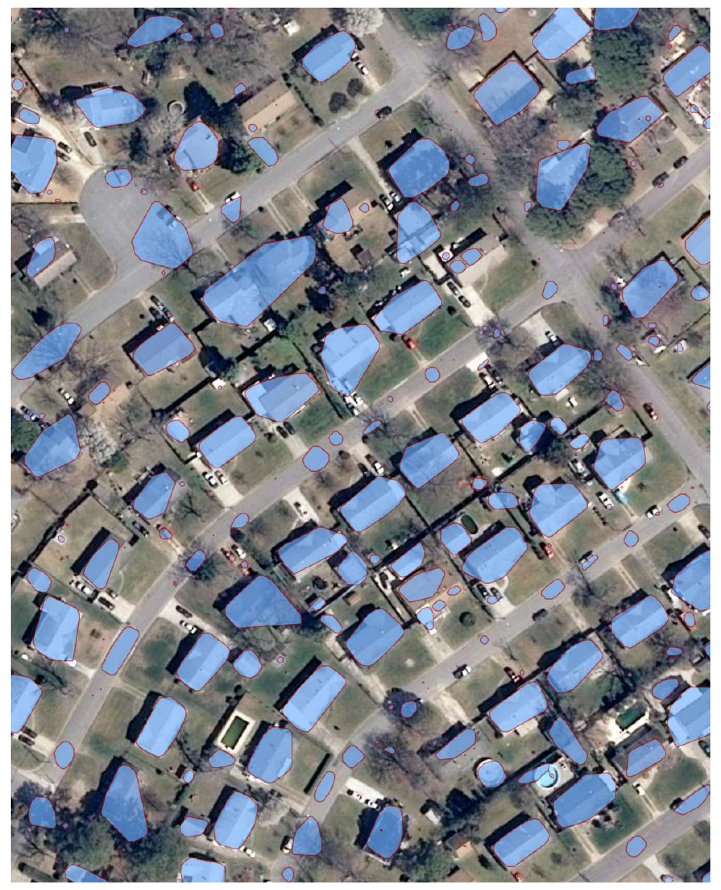

Building outlines detected by random forest classification

Building outlines detected by convolutional neural network

2 Energy Estimation

The energy consumption model is built using data from the Department of Energy's Residential and Commercial Building Energy Consumption Surveys[6] (RECS and CBECS) and from the City of Gainesville, FL. It takes as primary input the overhead building area obtained as described in the "Building Detection" section. We build a multiple linear regression model using RECS and CBECS to estimate coefficients. The model then outputs yearly energy demand in total kilowatt-hours for the household. The model can scale up from one variable (overhead area) depending on the availability of other data, such as climate, population density, building height and income.

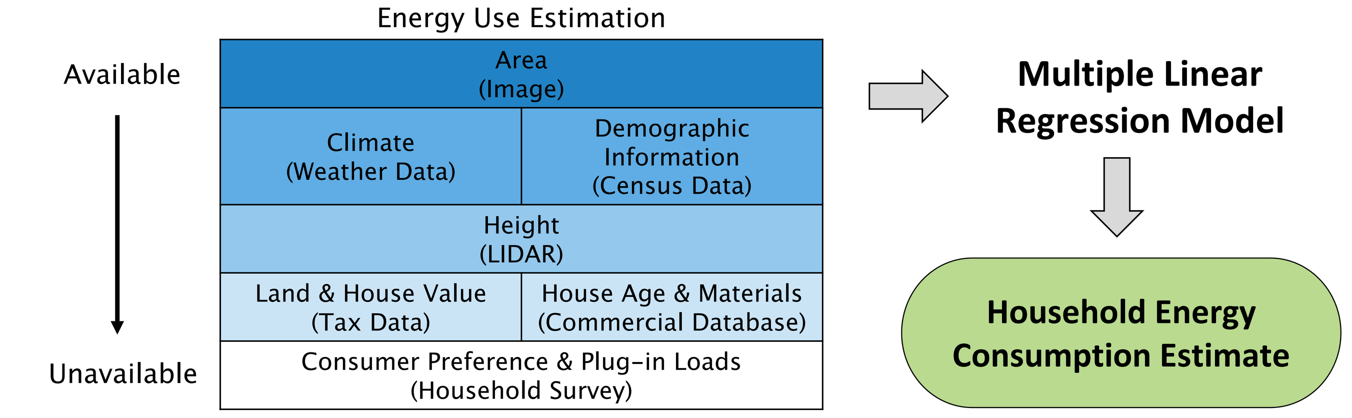

We use publicly available household energy consumption data from the city of Gainesville, FL to assess the performance of both our energy estimation model, and our energy estimation model applied to building shapes detected by our CNN. Applying our energy estimation model directly to the true building size yields a slight overestimate of building electricity consumption (B below). We then apply our CNN building detection algorithm to the Gainesville image and apply our energy estimation model to the buildings detected automatically (C below). The result is a slight underestimation of aggregate energy consumption over the whole image.

3 Gainesville, FL Test Region

	A	B	C
	Actual Energy Consumption	Actual Buildings & Estimated Consumption	Detected Buildings & Estimated Consumption
Number of Buildings	388	388	299
Average Energy Use (kWh/yr)	10,237	11,977	12,405
Total Energy Consumption (kWh/yr)	3,971,802	4,647,068	3,709,057
Total Energy Estimation Error (%)	-	17%	-7%

Actual Energy Consumption — Actual Buildings & Estimated Consumption — Detected Buildings & Estimated Consumption

Sources

1. Bradbury et. al. (2016): Aerial imagery object identification dataset for building and road detection, and building height estimation. figshare.https://dx.doi.org/10.6084/m9.figshare.c.3290519.v1
2. H. Mayer, "Object extraction in photogrammetric computer vision," ISPRS Journal of Photogrammetry and Remote Sensing, vol. 63, no. 2, pp. 213–222, Mar. 2008.
3. J. Yuan, "Automatic Building Extraction in Aerial Scenes Using Convolutional Networks," arXiv preprint arXiv:1602.06564, 2016
4. S. B. Kotsiantis, D. Kanellopoulos, and P. E. Pintelas, "Data preprocessing for supervised leaning," International Journal of Computer Science, vol. 1, no. 2, pp. 111–117, 2006.
5. T. Blaschke, "Object based image analysis for remote sensing," ISPRS Journal of Photogrammetry and Remote Sensing, vol. 65, no. 1, pp. 2–16, Jan. 2010.
6. U.S. Energy Information Administration, "Residential Energy Consumption Survey," 2009.

Conclusions and Future Work

We have demonstrated an approach to estimate building-level energy consumption, given only high-resolution aerial orthoimagery. Our overall energy estimation results for this study resulted in 7% error for a 2.25 km² region. Our building detection approach identifies over 80% of building pixels with fewer than 10% false detections. Our approach could be improved by refining the building detection method through decision fusion, merging results from the random forest and CNN classifiers. Energy estimation could be improved by incorporating information on building height and roofing material. Finally, this process could be compiled into a user-friendly tool that can be applied to any area in the world with available high resolution aerial imagery.

Acknowledgements: We would like to thank the Bass Connections program for sponsoring this project, the Energy Initiative for its mentoring and computing resources, Duke Library's The Edge for providing a meeting space, the summer 2016 Data+ team and mentors for laying the groundwork of this research, and Jordan Malof for his lecture on machine learning.