Supporting Electric Vehicle Planning in African Cities

Tanmay Agrawal, Bob Muhwezi
tanmaygrawl@umass.edu, bmuhwezi@umass.edu
University of Massachusetts, Amherst

Introduction:

Several developing countries have started exploring the adoption of electric vehicles as an important component for meeting climate goals. However, in order to formulate the relevant policies and invest in electric vehicle related infrastructure (e.g charging stations), governments and other stakeholders need more information about **vehicle use**, electric grid capacity, and vehicle purchase patterns; Data on all of which is severely limited for many developing countries. In this project, we use a combination of publicly available satellite imagery data, and off-the-shelf image classification and object detection models to estimate the spatial distribution of vehicles in Nairobi city over several years.

Data:

To model our problem, we mainly relied on 3 datasets;

- Car overhead with context (COWC) data. This is a set of publicly available high resolution aerial imagery provided by the Lawrence Livermore National Laboratory. It consists of more than 33,000 unique car annotations from six different cities over three continents.
- High resolution satellite imagery for Kenya. This dataset consists of 2 sets of images from Digitalglobe and Google Earth Pro with a ground sample distance (GSD) of 0.3-0.5 meters. We annotated more than 6000 vehicles for fine tuning and testing the model.

Modeling Approach:

We experiment with two different approaches. A purely image classification and a standard object detection pipeline. Both the approaches require some manual data transformation and engineering tricks to be useful and give us different insights.

- Classification is a much faster and a more tractable problem.
- Localization is slower and complicated with images that don't have bounding boxes annotated in the training data, but it provides a deeper spatial insight and is more extensible to other applications.

Classification approach:

In this approach, we reduce the problem of "counting" into a multi-class classification problem.We start by slicing the labeled images into smaller ones where each image is associated with a label from 0 to 9 corresponding to the number of cars in that slice. We then train a ResNet50 model to classify these slices and stitch them together to get a full count. This approach makes a reasonable assumption that there will be no more than 9 cars per slice.

Object detection approach:

We start by transforming the centroid labels of the COWC dataset into bounding boxes by assuming a median car size of 3 meters. We downsample the hi-res images of the COWC dataset from 15 cm to 30 cm and 50 cm ground sample distance (GSD) to match the resolution of our satellite imagery dataset. We then train a YOLOv3 object detection model on cropped images from this pipeline.

Results:

1. **Classification:**
   The results from classification on the COWC test set (Utah AGRC) was ~98% but this model failed on low resolution imagery and further experiments with more training data is needed to use it on DigitalGlobe/Google Earth imagery from Africa.

2. **Detection:**
   The results from the object detection approach are illustrated by the image below where a bounding box is drawn on each instance of a car with opacity of the bounding box representing confidence level for the prediction.