Classifying Road Quality Using Satellite Imagery in Metro Detroit

Nabil Ahmed, David Karamardian, Trent McMullen

Department of Statistics, Stanford University

Stanford
CS231N Spring 2022

Abstract

Our project utilizes transfer learning to classify the quality of roads in Metro Detroit. We cultivate a new dataset of 40,000 satellite images queried from the Google Static Maps API and combine them with publicly available PASER road evaluation data. After addressing significant levels of noise in the PASER road quality labels, we experiment with a pretrained version of ResNet-50 and different loss functions to achieve over 70% classification accuracy on our validation set. Finally, we visualize our model results and feature embeddings through the use of saliency maps and t-SNE cluster analysis.

Stanford University

Creating the Dataset

We pulled Southeast Michigan road quality evaluations (PASER condition dataset) from 2017-2019 and calculated midpoint latitude and longitude coordinates for each road segment. These coordinates were then used to query satellite images from the Google Static Maps API to provide coverage of an expansive set of roadways throughout Metro Detroit with ”poor”, “fair”, and “good” road quality labels.

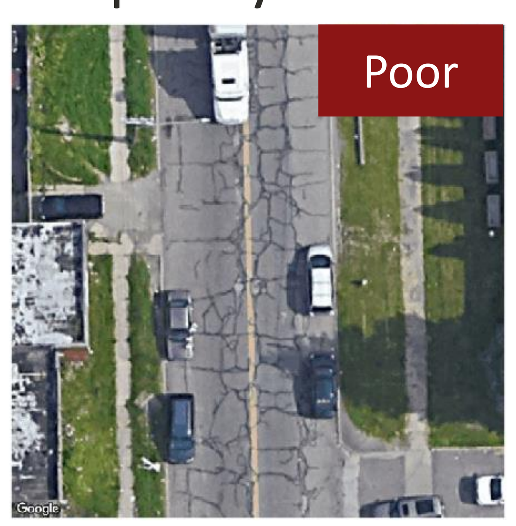

Unfortunately, gaps in timing between the road evaluations and date of photo capture led to misaligned data labels. A human expert only agreed with the PASER rating on 58.7% of images from a randomly sampled set of 150. To address this issue, the team hand-labeled 2700 images for training and validation.

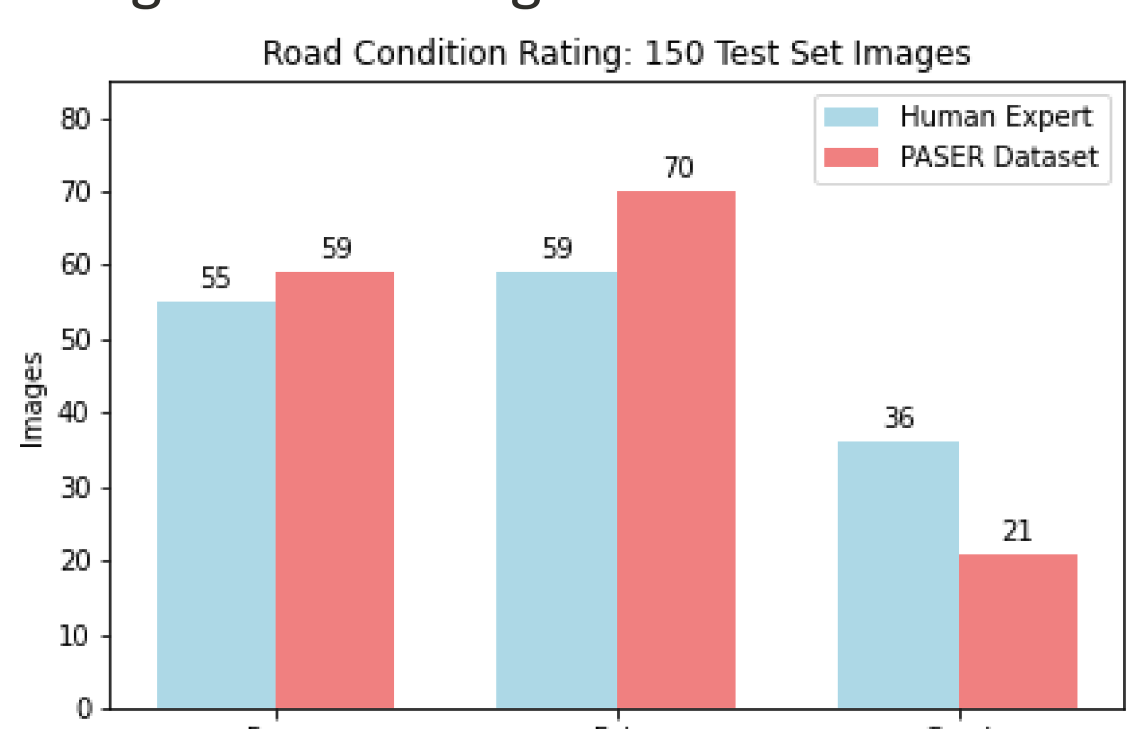

Model Development and Training

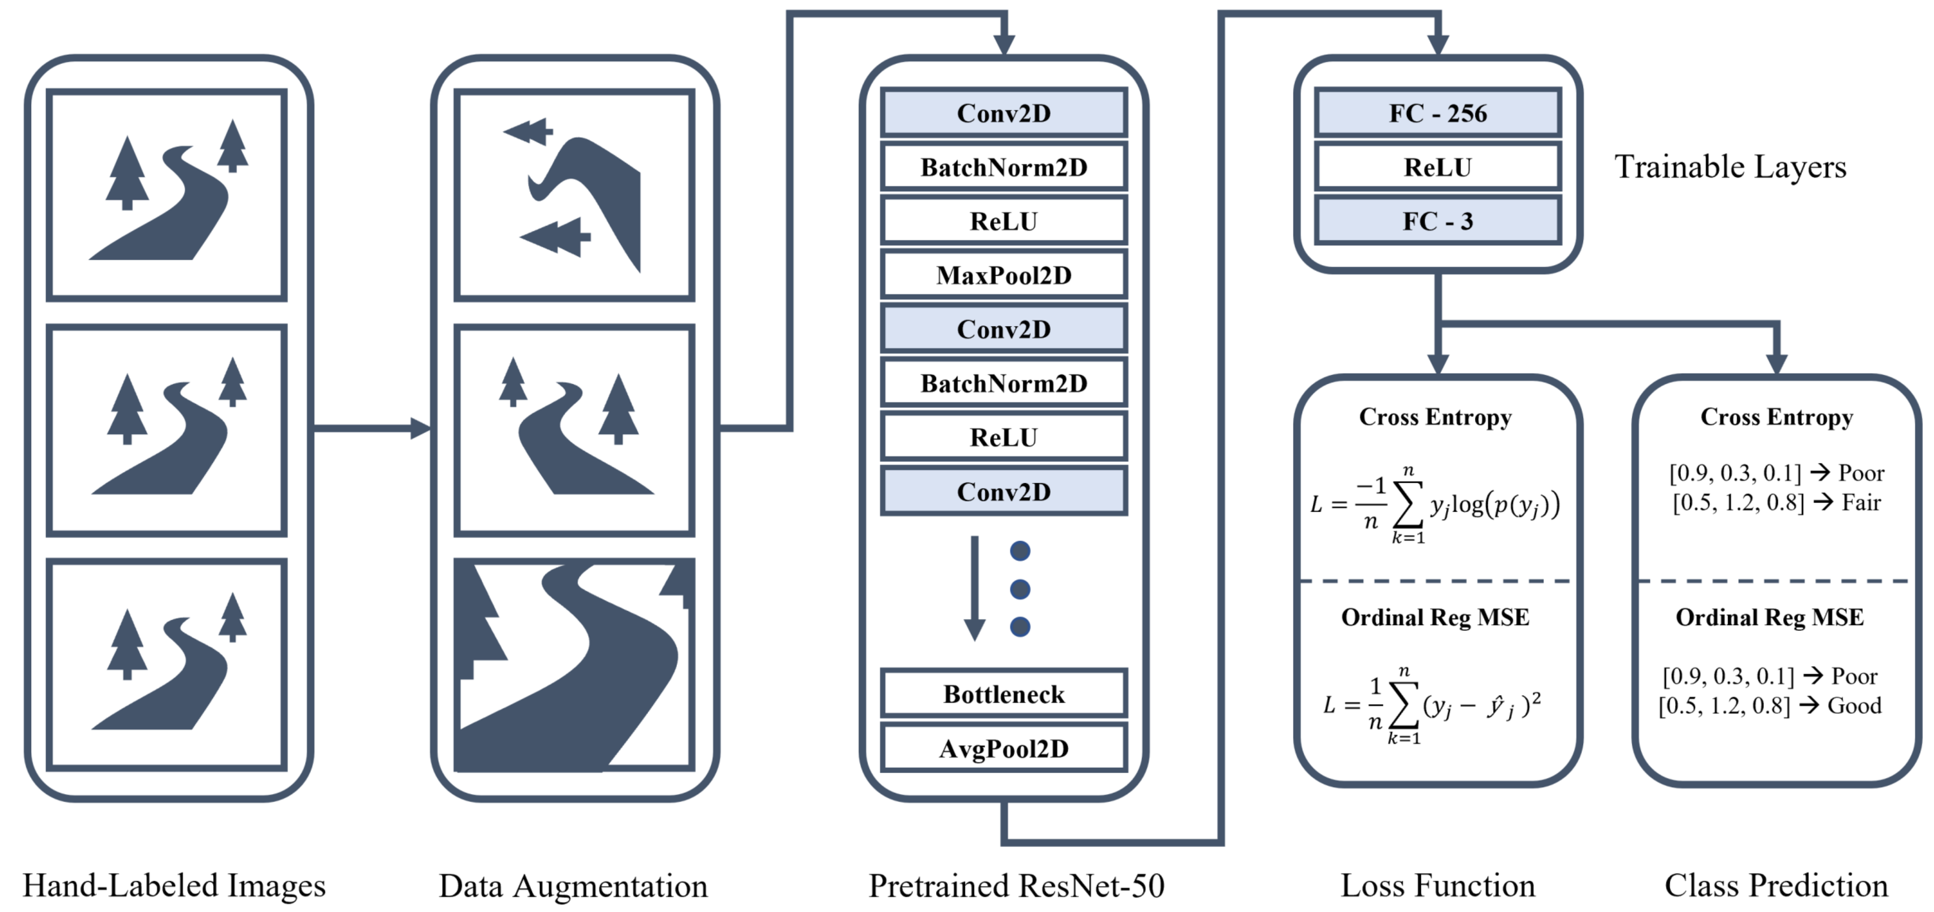

After hand-labeling, data augmentation was used to enhance the diversity of images presented to the model during training:

- Center-cropping (256x256)
- Resizing (224x224)
- Random Flips / Rotations
- Image Normalization

We leveraged transfer learning by employing a pre-trained version of ResNet-50 and modifying the final layers to produce 3 class scores.

Validation Classification Matrix

Actual (Hand Label)	ResNet-50 CE Predicted: Poor	Fair	Good	ResNet-50 ORD Predicted: Poor	Fair	Good
Poor	27.2%	5.0%	0.8%	31.4%	1.8%	0.0%
Fair	6.1%	17.5%	9.6%	10.5%	18.2%	4.6%
Good	1.3%	4.6%	27.6%	3.5%	7.7%	22.4%
	72.4% Accuracy			71.9% Accuracy		

Two loss functions were trialed, each with its own advantages.

Our model using cross-entropy loss (CE) achieved **72.4%** validation accuracy, while the ordinal regression loss model (ORD) attained **71.9**%!

One-Hot Target Labels with Cross–Entropy Loss:
Poor → [1, 0, 0] Fair → [0, 1, 0] Good → [0, 0, 1]
Ordinal Regression Target Labels with MSE Loss:
Poor → [1, 0, 0] Fair → [1, 1, 0] Good → [1, 1, 1]

Results & Visualizations

While both models had similar raw classification accuracy on the training and validation sets, the confusion matrix demonstrates that the ORD model is superior at classifying truly poor-quality roads. This would be desirable in applications that involve flagging roads in immediate need of repair.

Saliency maps aid in understanding which pixels within an image drive the model’s class prediction. We can see that the model has learned to associate unbroken black lanes with good quality roads while ignoring ancillary features like road markings.

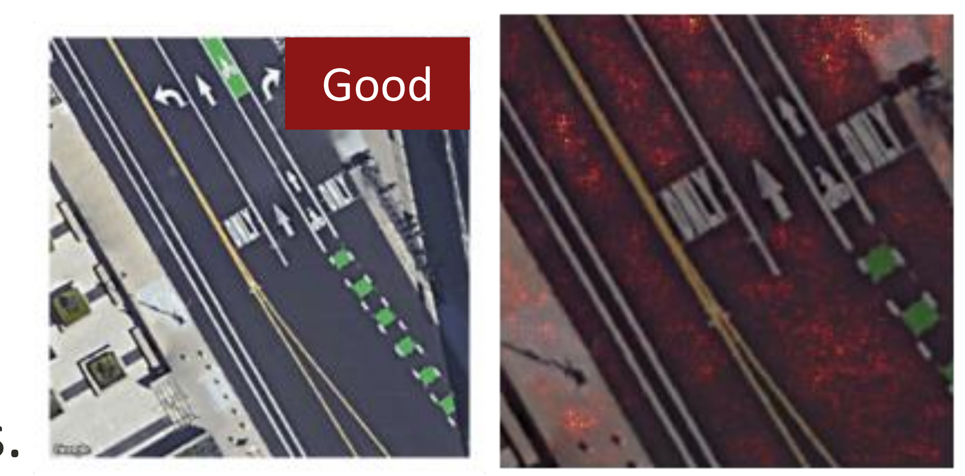

The feature embeddings stored in the penultimate layers of our models can be visualized in t-SNE plots. We can see distinct clusters, with the ORD model exhibiting a greater distance between the “poor” and “good” road quality clusters in the train and test sets.