Convolutional Architectures for Self-Driving Cars

Dhruv Choudhary and Gaurav Bansal

Introduction

- Self-driving vehicle control system need to determine steering wheel angle, brakes, and acceleration in any driving environment.
- We focus on use of deep learning for predicting steering angle using forward-facing cameras.
- NVIDIA paper showed CNN can obtain close to human driving performance.

Datasets

Udacity Data sets from 5 drives with a drive time of 28.23 minutes. Test vehicle records data from 3 mounted cameras and steering wheel angle as driven by a human driver.

- Normalized and zero-centered Image data
- Cropped top 1/3rd of the image as it does not have information on road
- Removed images with steering angle values below a certain threshold (.03 radians)
- Image flipping augmentation

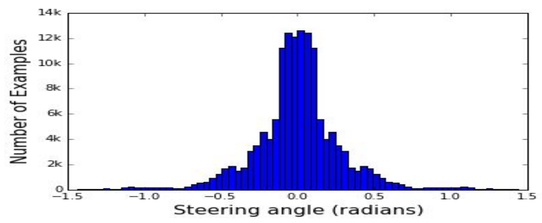

Models

- 10 layer architecture with 5 Conv layers and 5 FC layers
- Used Batch Normalization after every Conv layer and Dropout (0.2) after every FC layer.
- Adam Optimizer with LR=1e-3
- Hyper-parameter tuning using cross validation.

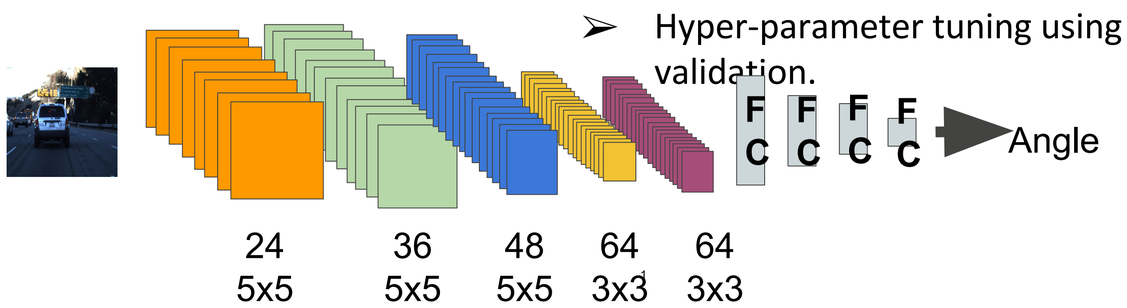

Results

- The validation MSE loss after 20 epochs = .011.
- Corresponds to 5.9 degree error in steering wheel angle prediction
- Best results reported in Udacity Challenges is 2.89 degree error
- Batch Normalization after the ReLU gives best performance.
- Dropout after FC layers helps overfitting.

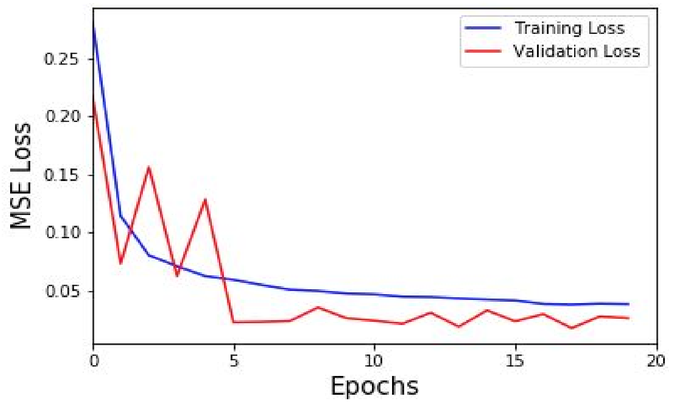

Visualization

VIsualization shows the actual and predicted steering wheel angle for three scenarios on the road.

Conclusions and Next Steps

- Deep Learning has a great potential to be used for end-to-end self driving applications.
- CNN+LSTM architectures to improve the accuracy.
- Evaluate effect of sensor data sharing between neighboring vehicles.