# Supervised Learning for Autonomous Driving

Greg Katz, Abhishek Roushan and Abhijeet Shenoi {gkatz,aroushan,ashenoi}@stanford.edu
Stanford University

## Abstract

In this work, we use supervised learning to predict steering and throttle inputs from raw images generated in a real-time simulation environment. We demonstrate the capability of the model by driving autonomously and comparing the behavior to the training data.

## Data Collection and Processing

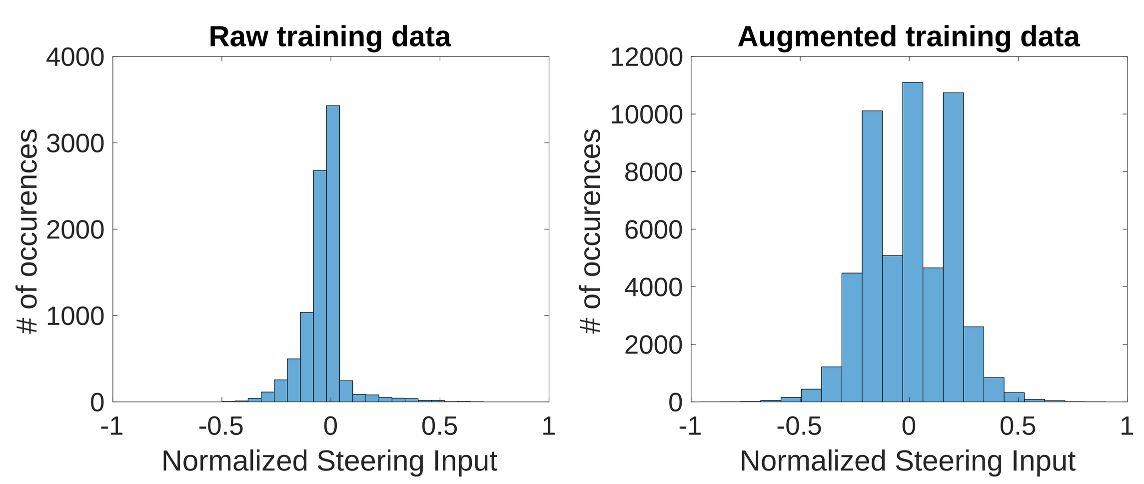


Figure: Comparison of Raw and Augmented Steering Angles

Figure: Typical simulator images and inputs to the network

The images from the simulator were cropped, and converted to grayscale before being used for training. The dataset was augmented by flipping the images horizontally and reversing the steering angles. Moreover, the simulator provides additional left and right camera views, which we used with a tunable steering correction factor. Finally, we experimented with and without an exponential smoothing of the steering joystick input.

## Model Selection

We selected from five different model types by training each on a single data set and judging the following:

- Error vs epoch (to check optimization convergence)
- Learning curve (bias vs variance indicator)
- Training and Validation error
- Cross track error (Figure of merit: model driving well on track)

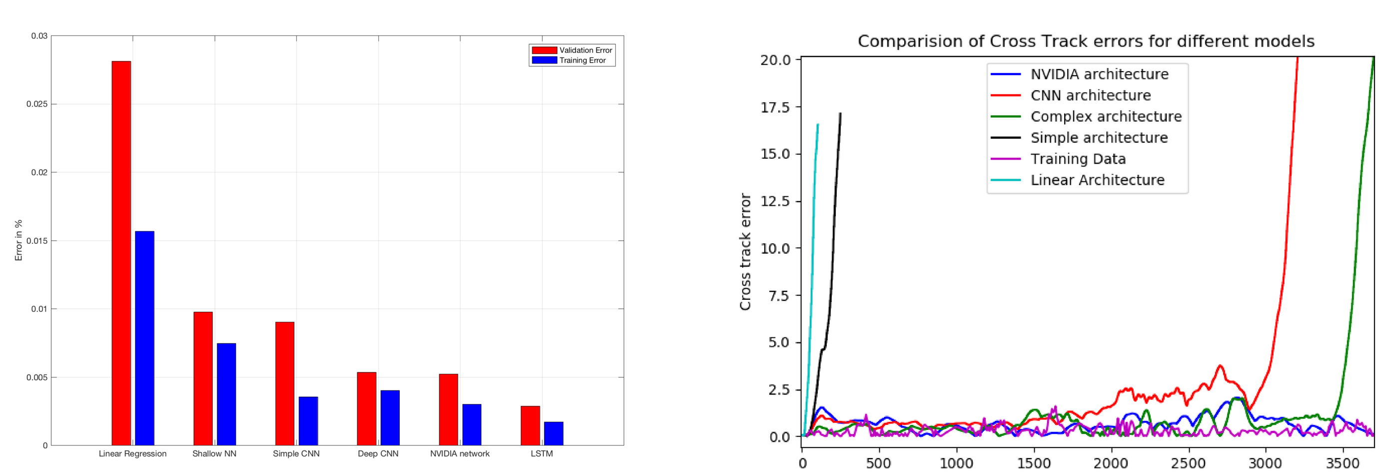


(a) Training and Validation Error (b) Cross track errors

Figure: Training and Validation Error & Cross track errors for all models

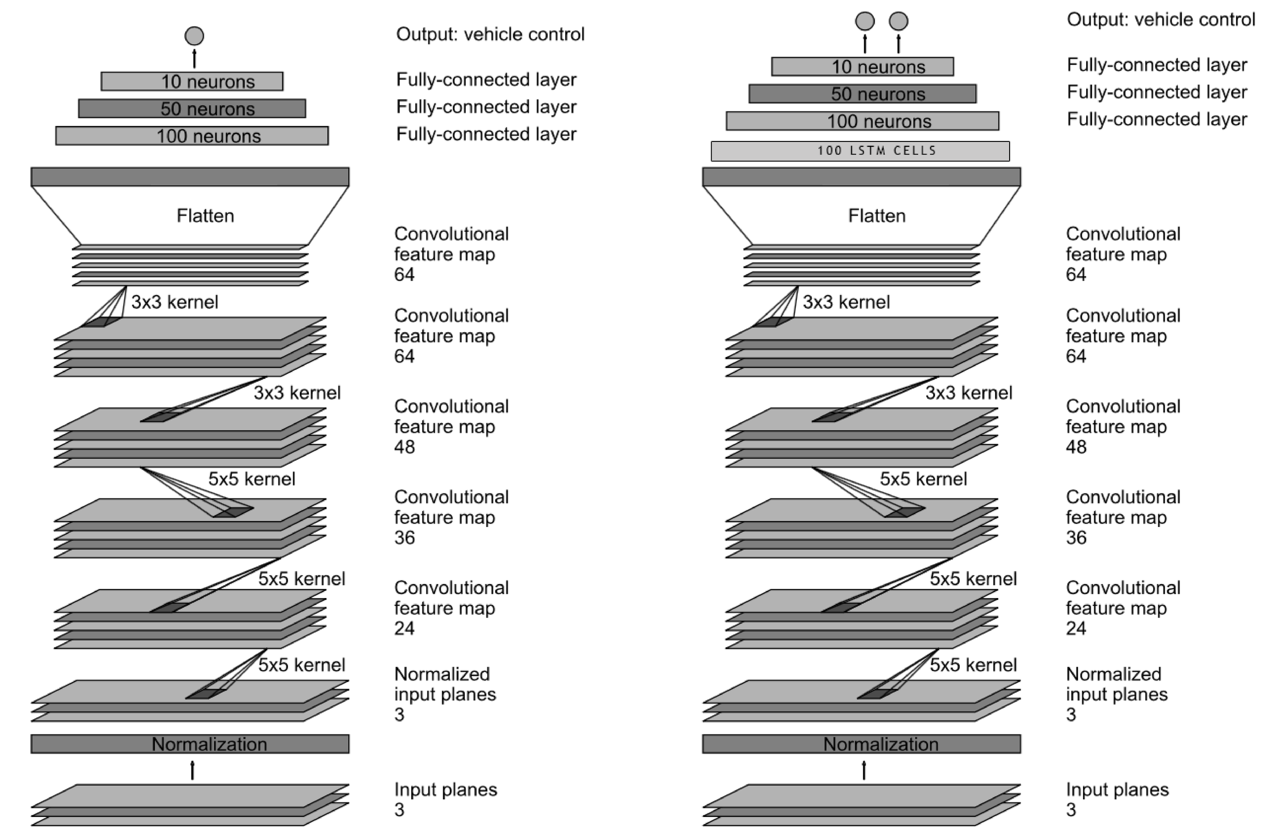


(a) NVIDIA (b) LSTM

Figure: Network Structures

## Model Training

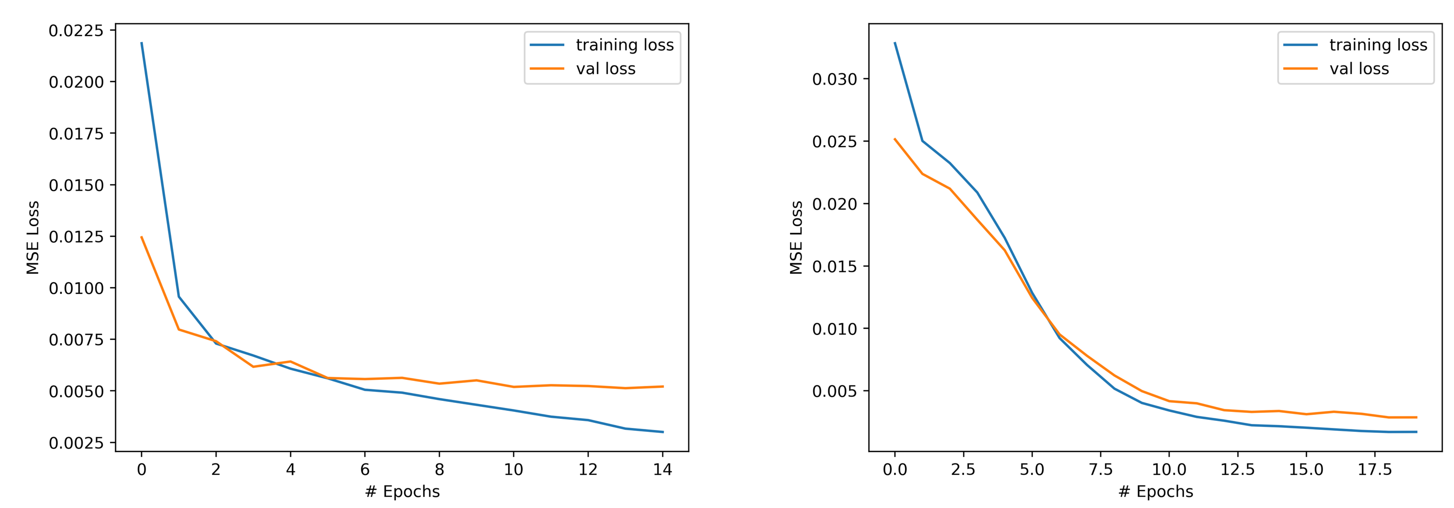


(a) NVIDIA model (b) LSTM model

Figure: Training and Validation Errors

Once we selected NVIDIA and LSTM we tuned the training to minimize validation error. Over-fitting was combated by using dropout with a probability of 0.2 or by early stopping. Training was done with mini-batch stochastic gradient descent on a mean squared error loss, using the Adam optimization method.

## Results

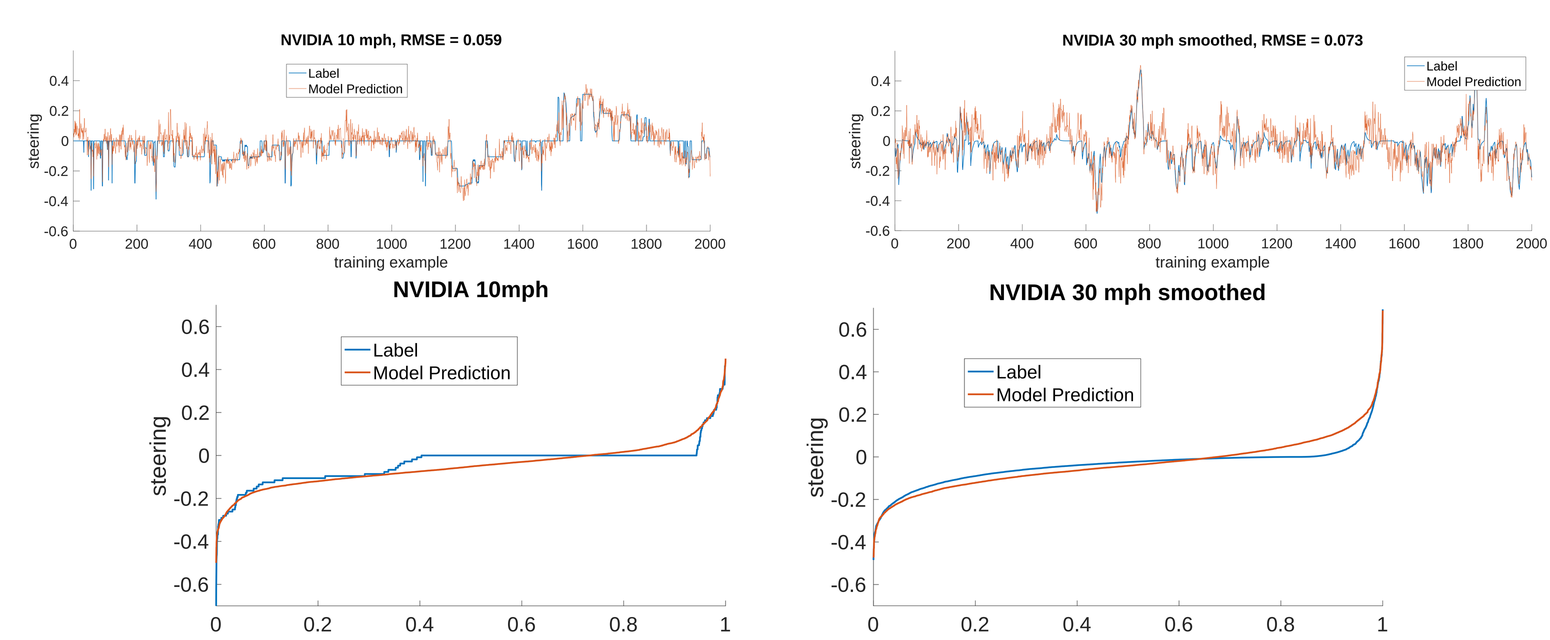


Figure: Steering angle Predictions on center camera test data

| Model | Forward Track | Reverse Track |
|---|---|---|
| NVIDIA | 20 | 15 |
| LSTM | 22 | 25 |

Table: Max stable speed (mph)

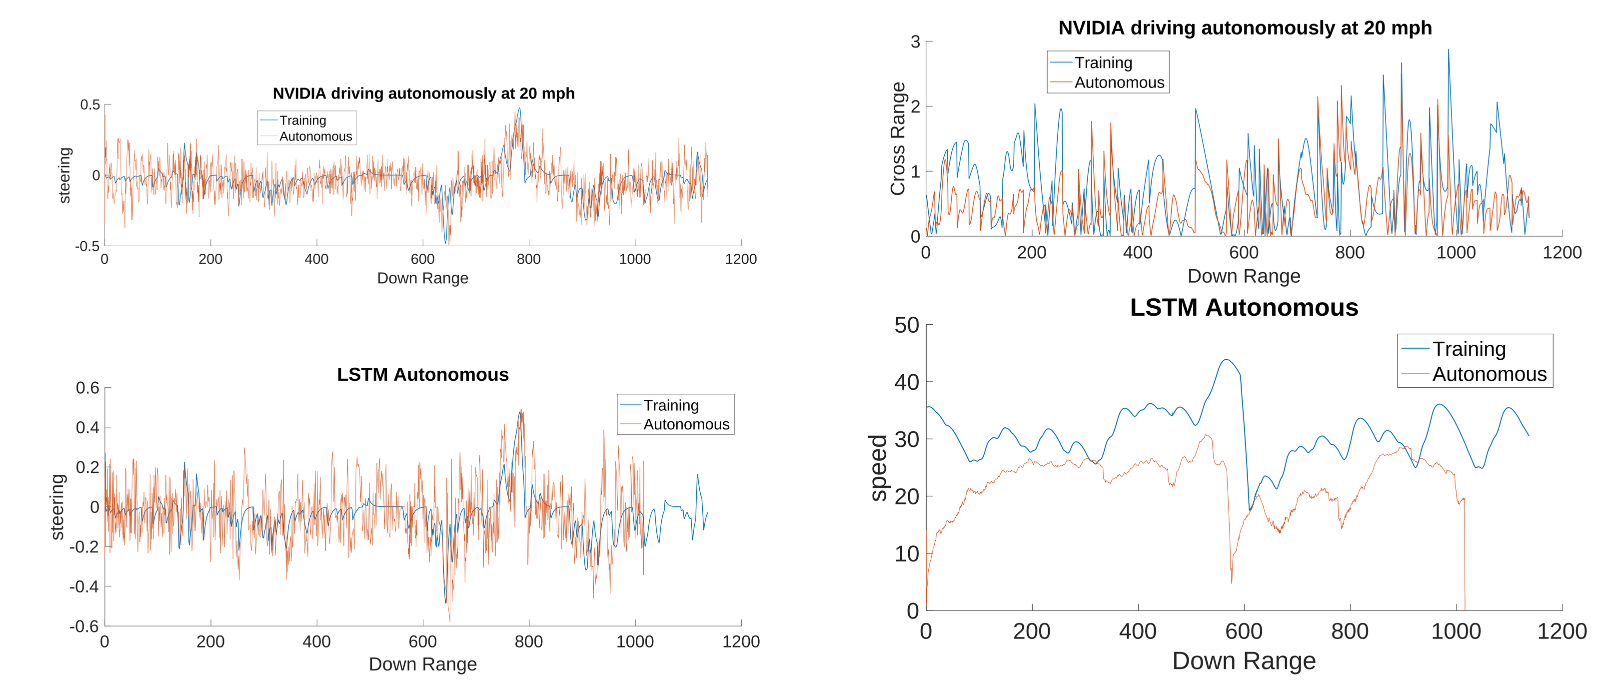


Figure: Autonomous Driving

## Observations and Inferences

- Both models were able to complete the training and test (reverse) tracks
- Although autonomous driving was a success, the steering behavior was quite noisy perhaps because of the use of left and right camera images
- Difficult to determine exact relationship between validation error and on-track performance

## Future Work

- Creating custom tracks/scenarios which thoroughly test driving capability and analyze failures effectively
- Preprocessing images to remove unnecessary background textures should allow us to generalize better
- Learning to drive is inherently suited to an RL framework, which gives us a new approach to this problem

## References

[1] Bojarski et al.
End to end learning for self-driving cars.
*arXiv preprint arXiv:1604.07316*, 2016.

[2] Udacity-an open source self-driving car.
`https://github.com/udacity/self-driving-car-sim`, 2017.

[3] François Chollet et al.
Keras.
`https://github.com/fchollet/keras`, 2015.