Deep Prediction for Uber Self Driving Cars

Jeff Schneider
CMU RI and MLD

Summary

Uber has been very successful at developing self-driving cars. They're vehicles have traveled over 3 million miles in autonomous mode and have given over 50,000 rides to members of the public who called for an Uber through the regular app and got a self-driving Uber (with a safety driver).

An important and not yet perfect part of the autonomy stack is the prediction system. It is tasked with making predictions about the future movements of all the relevant actors in the scene. These are mainly vehicles, bicycles, and pedestrians, but also include many other objects. The system should provide a full distribution of future movements through time that may be multi-modal (e.g. at intersections an actor may turn or go straight).

Recent work has demonstrated performance improvements on the prediction problem through the use of CNN's and Uber has begun to use these methods on its cars [Djuric *et al.*, 2018, Cui *et al.*, 2018]. This project will investigate further improvements in the performance of these deep learning approaches.

Research Directions

We are interested in the following potential approaches for the deep prediction problem:

- **Recurrent architectures.** Predictions are time-indexed distributions, or distributions over trajectories. The prior work handles this by using a set of output layer nodes that index time. However, the time series nature very naturally maps onto recurrent architectures, such as social LSTMs [Alahi *et al.*, 2016].
- **Trajectory distributions.** Prior work has investigated both the use of dropout to produce (unimodal) distributions for each predicted trajectory and the use of multiple output nodes to produce multi-modal distributions. We would like to explore novel approaches to obtaining predictive distributions.
- **Connecting to raw sensors.** Prior work starts from processed sensor data. We will investigate utilizing raw sensor data directly to improve the quality of predictions.

Background and resources

Participants should be proficient with Python and either TensorFlow or PyTorch. We will have access to Uber log data for training and testing as well as benchmarks on current performance. Participants will meet weekly with Jeff Schneider and occasionally with the Uber authors cited above.

References

[Alahi *et al.*, 2016] A. Alahi, K. Goel, V. Ramanathan, A. Robicquet, F. Li, and S.Savarese. Social lstm: Human trajectory prediction in crowded spaces. In *Computer Vision and Pattern Recognition (CVPR)*, 2016.

[Cui *et al.*, 2018] H. Cui, V. Radosavljevic, F.-C. Chou, T.-H. Lin, T. Nguyen, T.-K. Huang, J. Schneider, and N. Djuric. Multimodal Trajectory Predictions for Autonomous Driving using Deep Convolutional Networks. *ArXiv e-prints*, September 2018.

[Djuric *et al.*, 2018] N. Djuric, V. Radosavljevic, H. Cui, T. Nguyen, F.-C. Chou, T.-H. Lin, and J. Schneider. Short-term Motion Prediction of Traffic Actors for Autonomous Driving using Deep Convolutional Networks. *ArXiv e-prints*, August 2018.