Diverse Challenging Driving Tests

Farid Soroush
Mechanical Engineering Department
fsoroush@stanford.edu

Predicting

Autonomous vehicles industry has been growing so fast in the past few years. Despite this progress, there still exists a lack of robust tests for driving platforms. On the other hand, real-world testing not only is expensive and dangerous for public, but also due to the rare nature of dangerous scenarios, will require billions of miles in order to statistically validate performance claims [1]. First, we implement DPPs to create the most diverse and challenging driving scenarios to minimize required testing time of autonomous cars. Then, we train a non-linear and high dimensional stochastic human driving model using GAIL algorithm to validate performance of our sampling results using experiments. Using the trained model, we show that DPPs sampling generates much more diverse and challenging scenarios for driving test than naive Monte Carlo.

Data

The data is from the Next-Generation Simulation (NGSIM) data set. NGSIM contains highway driving trajectories for US Highway 101 and Interstate 80. The raw data set consists of 11,850,526 rows and 25 columns. Data was collected through a network of synchronized digital video cameras.

Features

For each vehicle, 66 features are extracted. The policy outputs longitudinal acceleration and turn rate values as the vehicle action. Features include 20 lidar distances, 20 lidar distance time derivatives, and several ego and leading vehicle parameters (e.g. relative heading, velocity, vehicle length, vehicle width, lane curvature, longitudinal and lateral acceleration, jerk, local and global turn rate, etc).

Models

The goal is to imitate human behaviour to learn his policy in driving. We consider the input data as a sequence of state-action pairs that result from a policy. Although the true reward is unknown, a surrogate reward may be learned directly from data. Generative Adversarial Imitation Learning (GAIL) algorithm [2] tries to match a generator (G) with state occupancy distribution of the human policy. In GAIL, we have a discriminator D, which tries to distinguish the learned policy from the human policy. So the objective function is

$$\min_{\theta} \max_{\psi} E_{\pi_E} log(D_{\psi}(s,a)) + E_{\pi_{\theta}} log(1 - D_{\theta}(s,a))$$

After sufficient training, the generator (G) learns to imitate the human policy.

Results

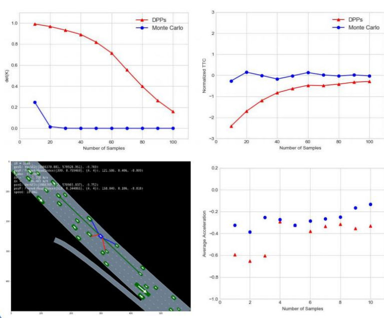

Discussion

For sampling and generating scenarios from the data set, DPPs and Monte Carlo were used. TTC [3] and diversity are calculated and for both methods. DPPs perform better than naive Monte Carlo, regardless of the sample size. A generator is trained and the learned weights for policy has been used to drive a car. The generator is used to compare results of two sampling methods in the simulated environment. Samples from each method are fed to the environment and the learned weights were used to drive a car in each case. Average acceleration of the car in the first 0.3 second of driving is calculated for both methods. Hard-brake cases (high negative acceleration) occurred more in DPPs samples. Therefore, higher performance of DPPs in generating more challenging scenarios is validated using the trained generator (G) that imitates human behaviour.

Future

An ensemble of generators can be trained and a distribution over the weights can be modeled using a (multivariate normal) parametric bootstrap. The models are high dimensional and graphical lasso can be used to fit the inverse covariance matrix for the ensemble. In this way, a new stochastic model for generator can be obtained and the covariance matrix can be sampled using DPPs in order to get the most diverse driving behaviours.

References

[1] N Kalra and SM Paddock. How many miles of driving would it take to demonstrate autonomous vehicle reliability. Driving to Safety, 2016.

[2] Jonathan Ho and Stefano Ermon. Generative adversarial imitation learning. In Advances in Neural Information Processing Systems, pages 4565–4573, 2016.

[3] Katja Vogel. A comparison of headway and time to collision as safety indicators. Accidentanalysis & prevention, 35(3):427–433, 2003