Leveraging Human Input for Self Driving Cars

Gokul Swamy



BERKELEY ARTIFICIAL INTELLIGENCE RESEARCH



With humans in the loop, learning algorithms should be able to learn *quickly* from data that is both *easy* and *safe* for people to provide.

We can accomplish these goals by taking advantage of the *structure* present in specific problems.

Weakly-Supervised 3D Intersection Structure Estimation (slides hidden from public version)

Modeling Interaction for Self-Driving Cars

Up Till Today: Driving in Isolation



Up Till Today: Driving in Isolation



(will talk about some cool perception work if I have time)

The Good

The Bad

How Does a Waymo Merge into Traffic?





The Ugly



What was the problem here?



- Approach V0: Constant Velocity
 - Will work decently well on highways
 - Or for things like drones



- Approach V1: Isolated Prediction
 - Fit network to map from state to action of car.
 - **Q**: Why will this not work well?
 - A: Need to incorporate information about other cars.



- Approach V2: Full Prediction
 - Fit network to map from state of all cars to action of a particular car
 - **Q**: Why might this not work well?
 - A: Your actions influence the actions of other cars



- **Problem**: other drivers *react* to the positions and actions of your car
 - Predictions \rightarrow Your Actions \rightarrow Other Actions \rightarrow Pred.
 - So if you update your policy, you change input distribution to network
- **Solution**: learn a function that models the person's response to a robot's action
 - Question: what kind of function to learn?

Joint work w/



Black-Box Models





Theory of Mind



Inferring Utility Functions

- Approach: Inverse Optimal Control / Reinforcement Learning
 - RL: Given reward function, find best actions
 - IRL: Given actions, find reward function that would have produced these actions
- High level blueprint: Fix a set of features a person driving could care about. Figure out what weights on these features would produce observed behavior on average.

Theory of Mind Results

- Can learn model from recorded human-human driving data
 - Much safer than throwing faulty car on road





Scaling to the Real World

- Unfortunately, running a nested optimization for each other car on the road is computationally infeasible to do online
 - Have only accounted for subset of pairwise interactions
- **Best of both worlds**: At training time, use ToM as data source to train network. At deployment time, use network for fast inference via a single forward pass.

Predictive Network Architecture: Waymo's MultiPath



Leveraging Problem Structure

- **Given**: we have history of state / action pairs of other agents from perception + other state information (lights)
- Roughly, where can a car go in the following situation?



Predicting Offsets

- **Q**: Why is this a good idea?
 - A: Lower Variance
- **Q**: Do we just want to predict an offset?
 - A: No, also want our network to be able to express *uncertainty*



Detour: Multivariate Gaussian

- Univariate Gaussian: fully determined by mean and variance
 - Maximum entropy distribution with these constraints makes the fewest further assumptions
- Multivariate Gaussian: mean becomes vector, variance becomes matrix
 - 2D for our use case.



Model Output: Mixture of Gaussians

• Single Trajectory Anchor Distribution (predict μ & Σ):

$$\phi(s_t^k | \mathbf{a}^k, \mathbf{x}) = \mathcal{N}(s_t^k | a_t^k + \mu_t^k(\mathbf{x}), \Sigma_t^k(\mathbf{x}))$$

- Full distribution over possible future states (predict π): $p(\mathbf{s}|\mathbf{x}) = \sum_{k=1}^{K} \pi(\mathbf{a}^{k}|\mathbf{x}) \prod_{t=1}^{T} \phi(s_{t}|\mathbf{a}^{k}, \mathbf{x})$
- Loss function:

$$\ell(\theta) = -\sum_{m=1}^{M} \sum_{k=1}^{K} \mathbb{1}(k = \hat{k}^m) \Big[\log \pi(\mathbf{a}^k | \mathbf{x}^m; \theta) + \sum_{t=1}^{T} \log \mathcal{N}(s_t^k | a_t^k + \mu_t^k, \Sigma_t^k; \mathbf{x}^m; \theta) \Big]$$

Model Output Visualized



Model Output Snazzily Visualized



Predictive Network Architecture: Waymo's MultiPath



Expanded Self-Driving Stack



Scaled Autonomy

Problem Setup

- Consider a fleet of autonomous robots, each independently executing an identical suboptimal policy
- At any time, the operator can step in and teleoperate a single robot
- Which robot should the operator choose to teleoperate?

Core Problem: Teleoperator Performance Degrades as Fleet Sizes Grow



When presented with a small fleet of robots, skilled teleoperators can supervise them effectively and step in where their assistance is needed most. However, teleoperators do not have limitless attention which makes it difficult for them to supervise large fleets.

Key Idea

Learn from user interventions given a few robots *to anticipate* which of many robots need help the most.

 This is easier for people to provide good data for because they can simultaneously focus on all robots.
Imitating these switches puts the person in situations where they can best provide demonstrations to improve the robot policy.

Method



Results: User Study



	Unassisted	Assisted	F(1,11)	<i>p</i> -value
Q1: On average, it was easy to guide the robots to their goals.	2.92	4.50	17.49	< 0.01
Q2: I was successful at guiding the robots.	2.25	3.92	13.75	< 0.01
Q1, after objective measures revealed	2.67	4.17	17.47	< 0.01
Q2, after objective measures revealed	2.92	3.25	0.88	0.37



Results: Real Robots





Questions?