Scalable End-to-End Autonomous Vehicle Testing via Rare-event Simulation

Matthew O'Kelly*, Aman Sinha*, Hongseok Namkoong*, John Duchi, & Russ Tedrake University of Pennsylvania, Stanford University, MIT
mokelly@seas.upenn.edu, {amans, hnamk, jduchi}@stanford.edu, russt@mit.edu

Introduction

- **Real-world testing** of autonomous vehicles (AVs) is expensive and dangerous
- **Formal verification** of AV "correctness" is intractable and subjective
- We consider a **risk-based framework**: we evaluate the probability of an accident under a distribution of standard traffic behavior
- Estimate generative models of traffic behavior from data via **imitation learning**
- Use **rare-event simulation** techniques to efficiently find probability
- Our system achieves 10-300P$\times$ speedup over real-world testing (P = # of processors)

Rare-event simulation

- Given: continuous measure of safety $f : \mathcal{X} \to \mathbb{R}$, threshold level γ, and distribution P_0 of traffic behavior
- Goal: Evaluate probability of bad events $p_\gamma := \mathbb{P}_0(f(X) \le \gamma)$
- Monte Carlo sampling is too slow (relative error of estimate $\propto \sqrt{1/p_\gamma}$)

Cross-entropy method

- Adaptive importance sampling technique that iteratively tilts sampling distribution P_θ to estimate $\widehat{p}_{N,\gamma} := \frac{1}{N}\sum_{i=1}^{N} \frac{p_0(X_i)}{p_\theta(X_i)} \mathbf{1}\{f(X_i) \le \gamma\}$
- Intuitively makes bad events more frequent. Iteratively tries to find $\theta^\star \in \mathrm{argmin}_{\theta\in\Theta} D_{\mathrm{kl}}(P^\star \| P_\theta)$, where $p^\star \propto \mathbf{1}\{f(X_i) \le \gamma\}\, p_0$
- Iterations are convex optimization problems for exponential families $\{P_\theta\}$

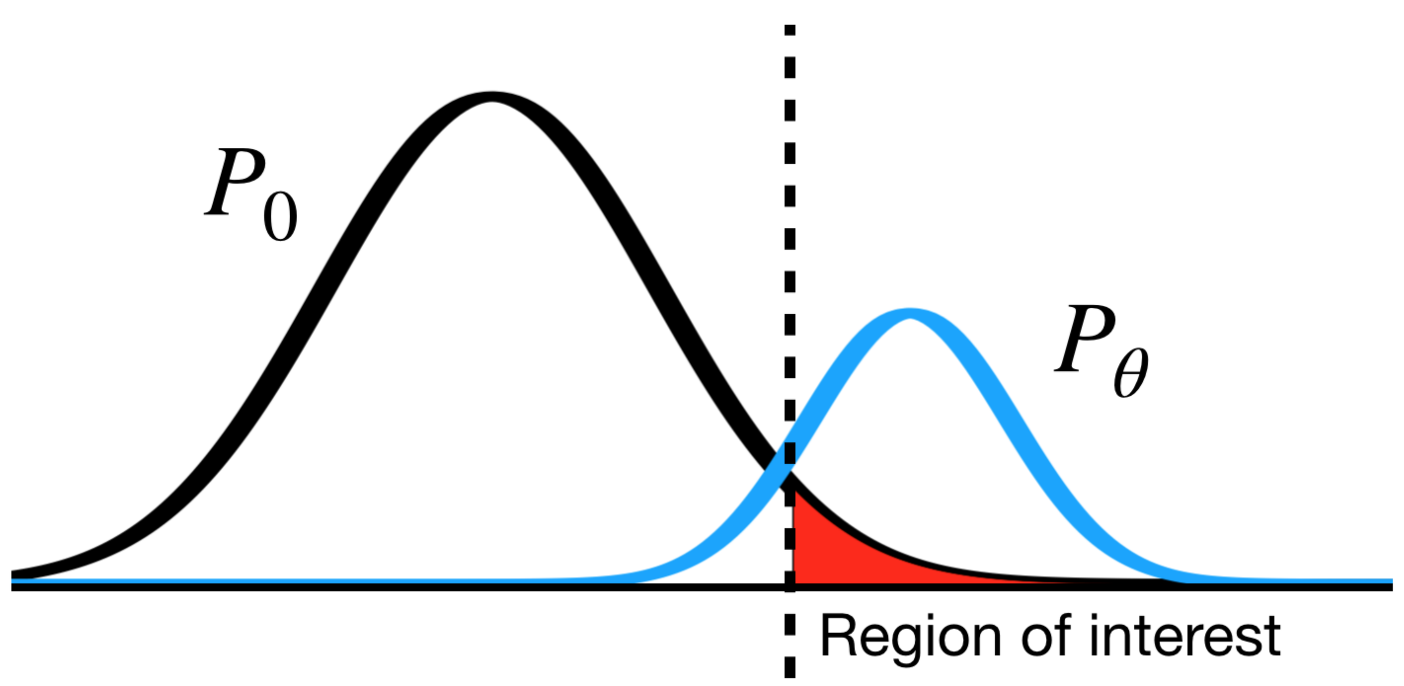

Data-driven generative models

- Use data of real highway traffic to build behavior of traffic cars P_0 [1]
- Employ generative adversarial imitation learning [number]
- Model-based variant [2] allows fully differentiable training
- Use parametric bootstrap of many GAIL agents to build P_0

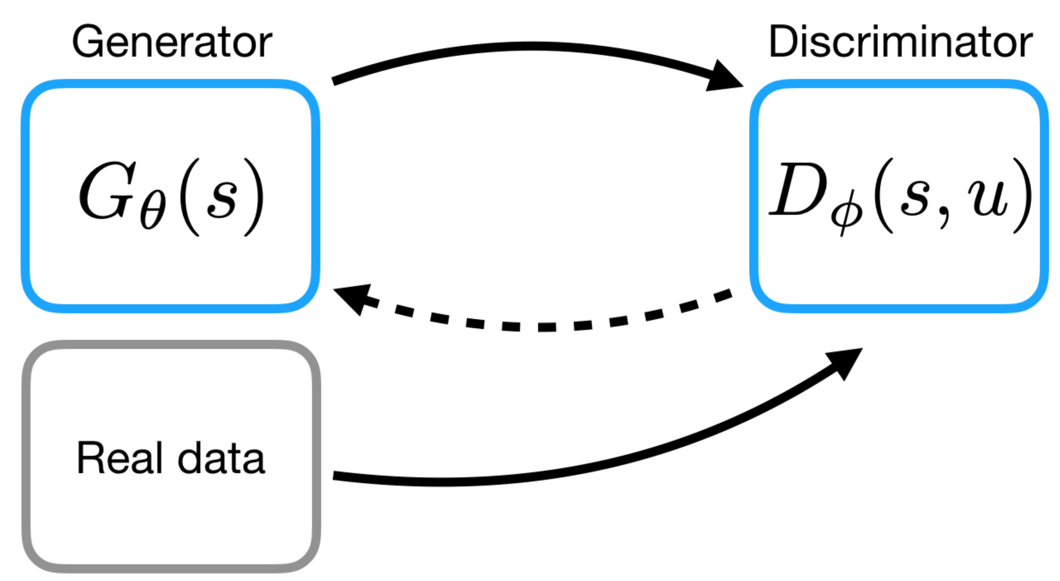

Simulation stack

- Simulate ego-agent (AV algorithm to be tested) amongst traffic cars from P_0
- Modular and fully distributed architecture: separate physics and perception (Unreal) engines with communication between processors via ZMQ

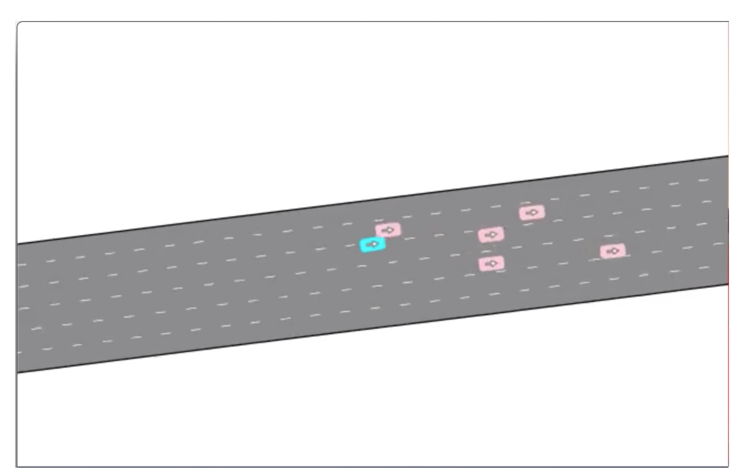
Overhead view


Dashcam RGB

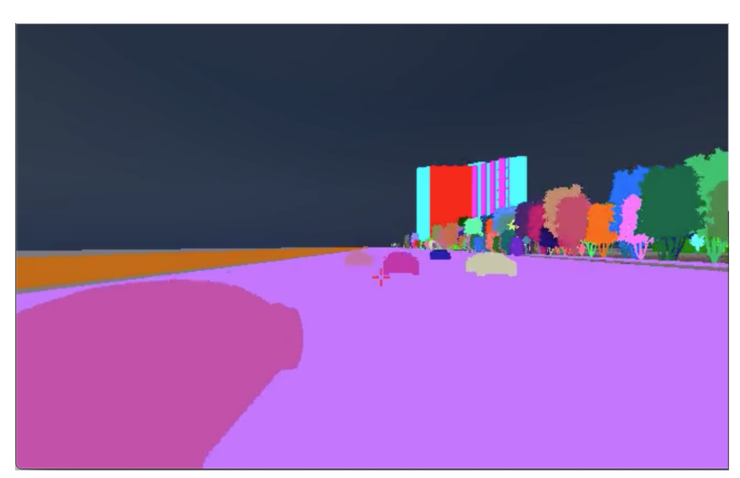
Dashcam segmentation


Dashcam LIDAR

Experiments

- Multi-agent highway scenario
- Test framework with both non-vision and vision-based ego policies
- Search over vehicle poses and behaviors of environment agents (i.e. weights of GAIL generator networks)
- Tune hyperparameter ρ (controls aggressiveness of cross-entropy method) on non-vision policy

Non-vision

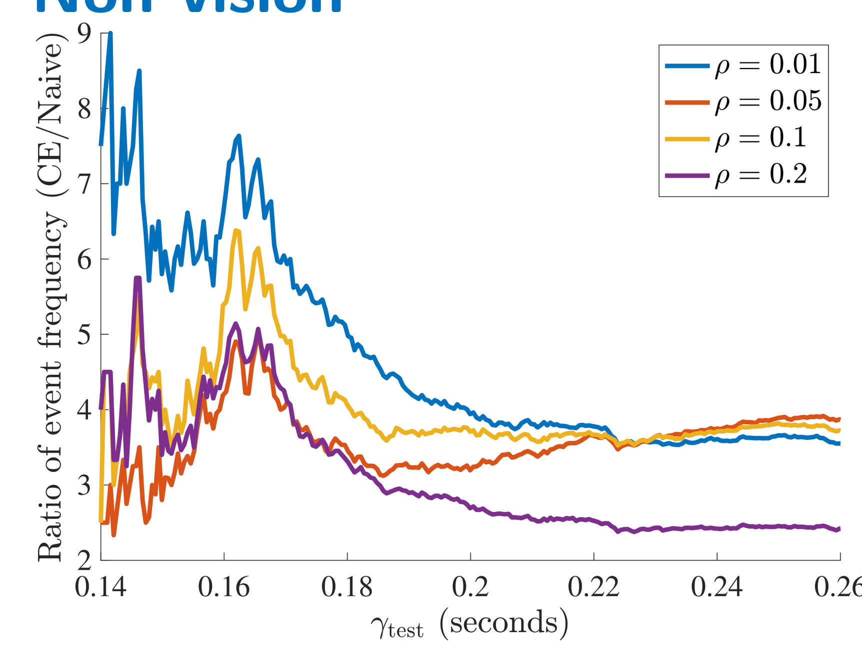

Ratio of # of rare events vs. threshold

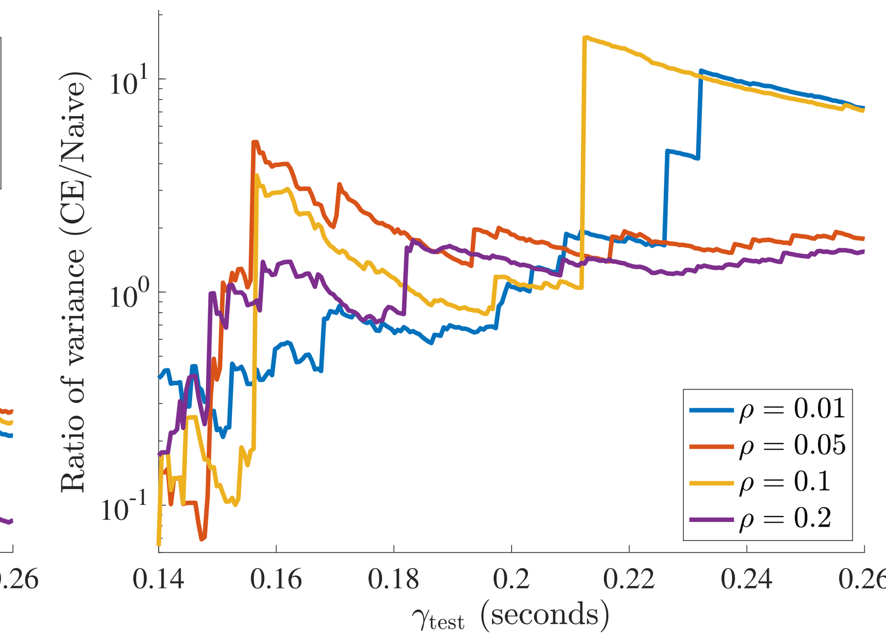

Ratio of variance vs. threshold

Vision

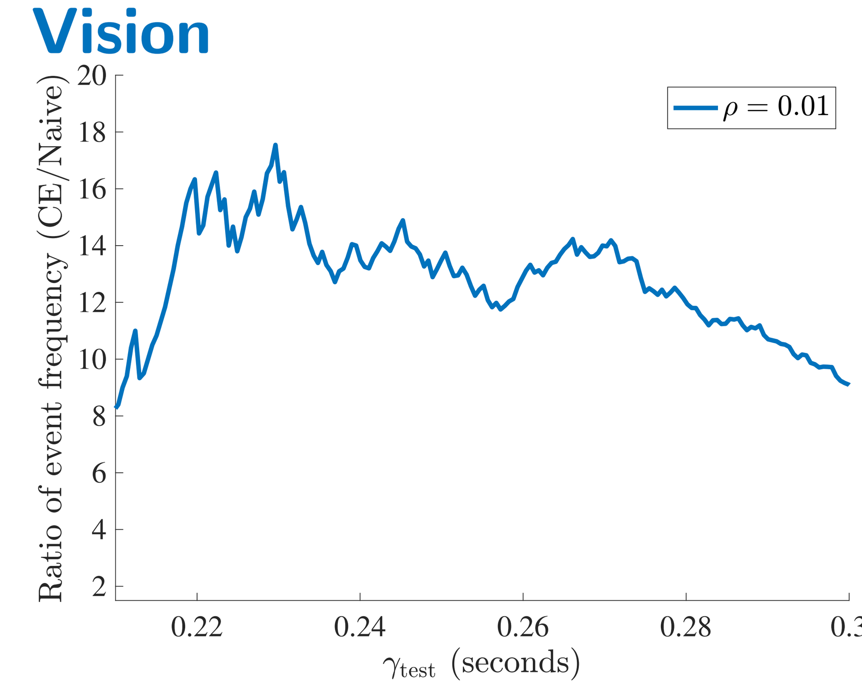

Ratio of # of rare events vs. threshold

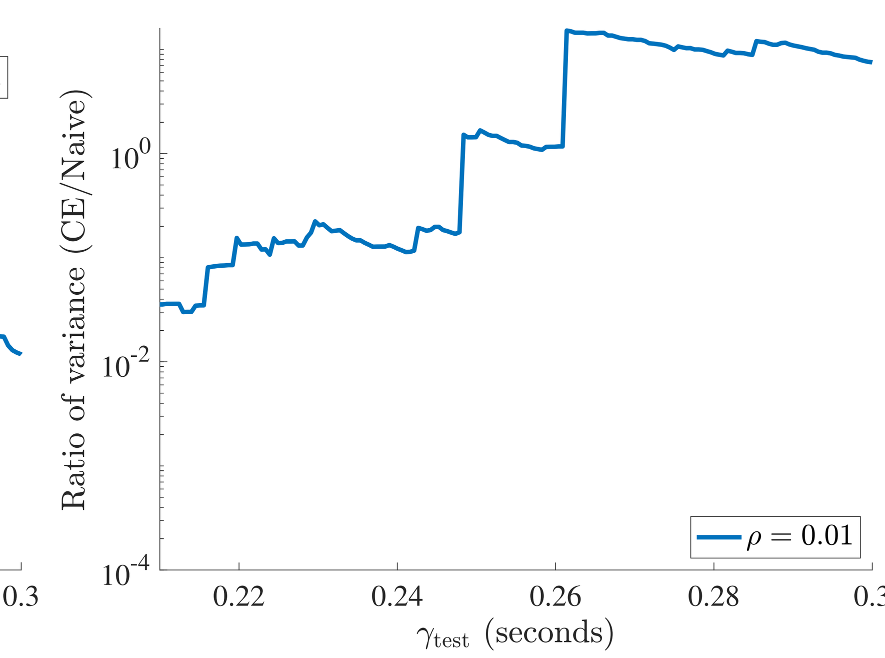

Ratio of variance vs. threshold

[1] US. Dept. of Transportation FHWA. Ngsim next generation simulation, 2008.
[2] N. Baram, O. Anschel, I. Caspi, and S. Mannor. End-to-end differentiable adversarial imitation learning. ICML, pages 390399, 2017.

Link to code: https://github.com/travelbureau/RareSim