Risk and Social Behavior for Decision Making for Autonomous Vehicles

Daniela Rus

(with Alexander Amini, Igor Gilitschenski, Teddy Ort, Alyssa Pierson, Wilko Schwarting, Sertac Karaman)

CSAIL, MIT

SMART, MIT

Autonomous Vehicles: Where are they?

Autonomous Vehicles: Where are they?

Autonomous Vehicles: Where are they?

Autonomous Vehicles: Where are they?

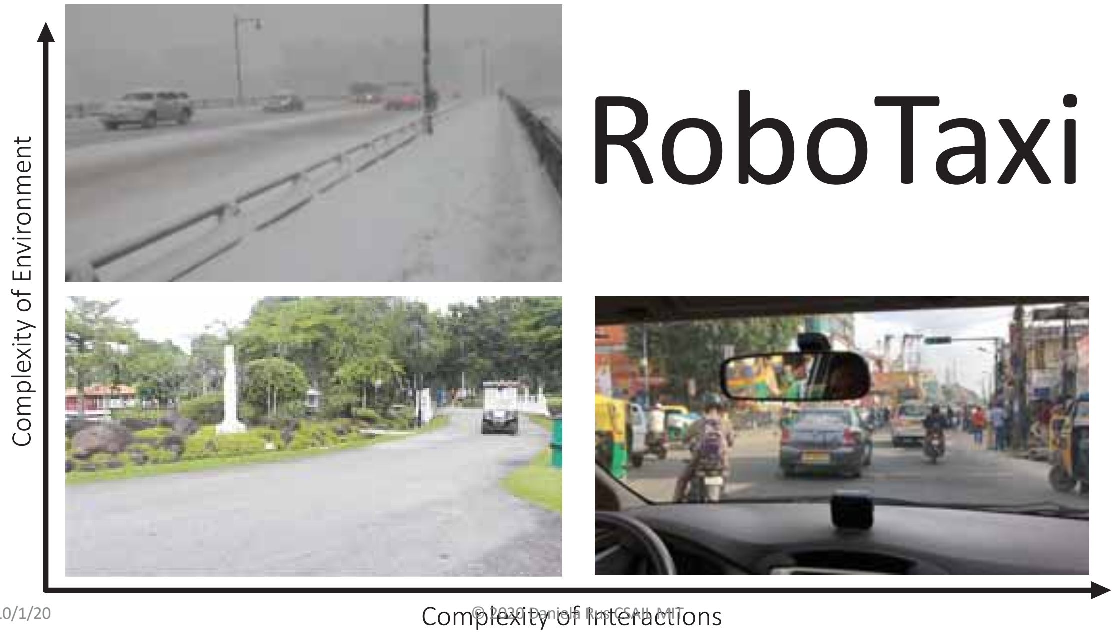

Outline

How can we enable self-driving vehicles to operate in more complex environments?

How can we incorporate risk in the control loop?

How can we handle congestion and interactions with human-driven cars?

Increased Capabilities: Learning to Drive from Humans

Classical autonomous driving pipeline

Separate problem into smaller sub-modules, tackle each independently

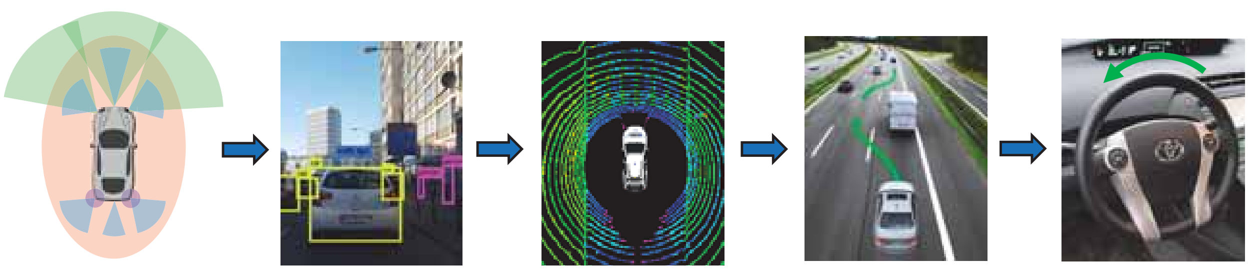

Sensor Fusion
What's happening around me?

Detection
Where are obstacles?

Localization
Where am I relative to the obstacles?

Planning
Where do I go?

Actuation
What control signals to take?

Learning our autonomous controller

Autonomous systems need the ability to handle a wide range of scenarios using raw and complex perception sensors

Night-time Driving

No Lane Markings

Rainy Weather

Leveraging large datasets, we **learn** an underlying representation of driving based on how **humans drive in similar situations**

End-to-End Learning

Learn the control directly from raw sensor data

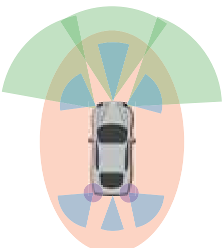

Learning Algorithm

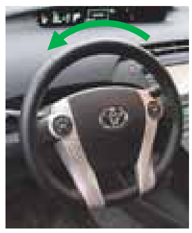

Sensor Fusion
What's happening around me?

Learned Model
Underlying representation of how humans drive

Actuation
What control signals to take?

[11, 12]

Learning to navigate

Raw Perception
I
(ex. camera)

Coarse Maps
M
(ex. GPS)

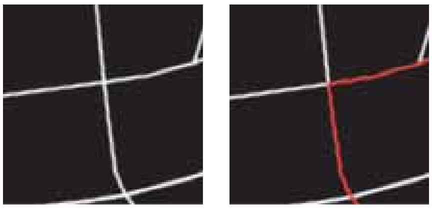

Probabilistic Control
$P(\theta|I, M)$

$$P(\theta|I, M) = \sum_{i=1}^{K} \pi_i \, \mathcal{N}(\mu_i, \sigma_i^2)$$

End-to-end optimization formulation

Learn a continuous probability distribution over the space of all control $P(\theta|I, M)$

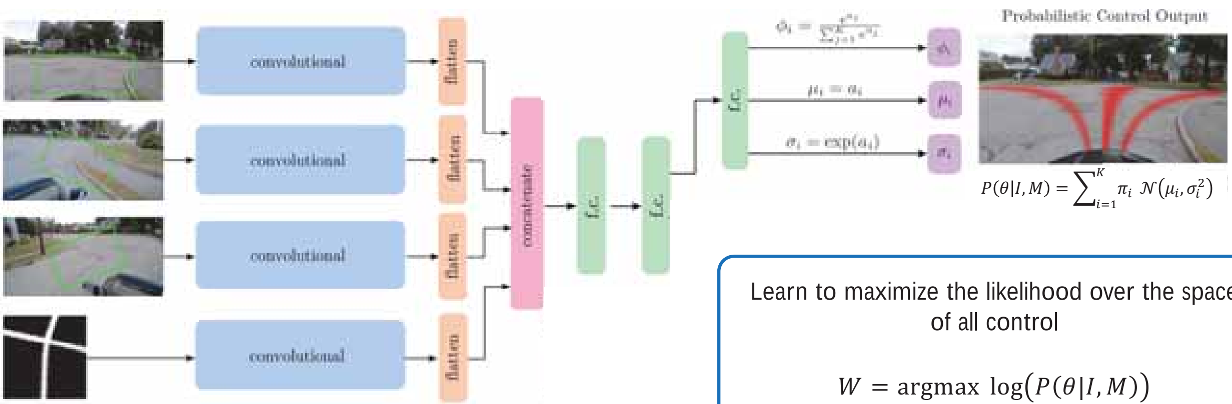

Learn to maximize the likelihood over the space of all control

$$W = \underset{w}{\operatorname{argmax}} \log\big(P(\theta|I, M)\big)$$

End-to-end optimization formulation

Input a route to compute a **deterministic control command** for navigation

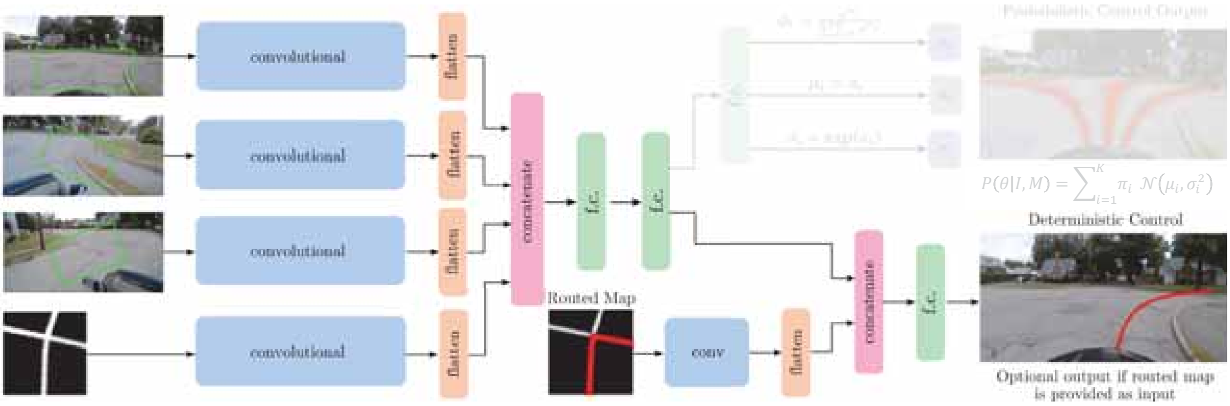

End-to-end optimization formulation

Entire model is trained end-to-end **without any human labelling or annotations**

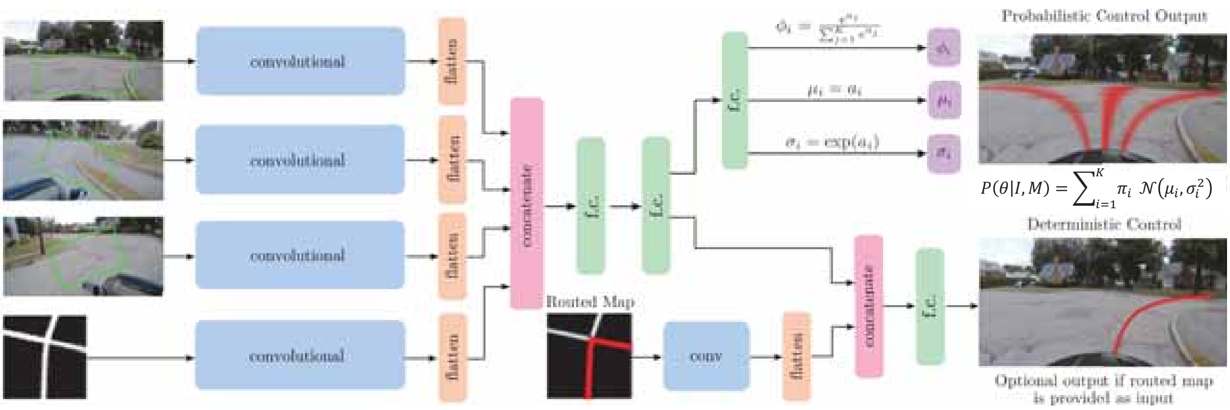

Generalization to new roads

Perception Input **Coarse Streetmap** **Control Output**

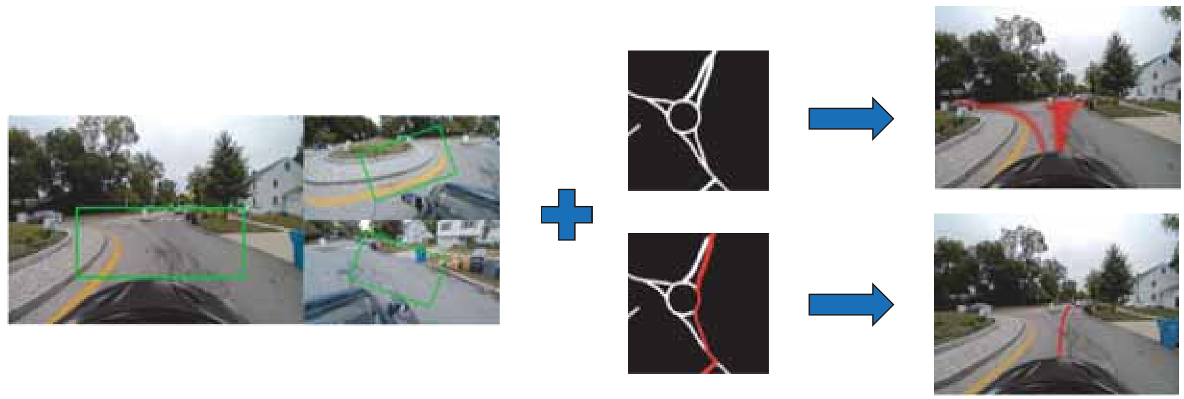

Model learned to generalize to **new roads** and even **new types of intersections** (ex. roundabouts never included in training)

Correcting pose based on visual perception

What do we do when our GPS is not accurate or even not available? $\Rightarrow \; Var[P(M)] \gg 0$

$$P(M|I) = \mathbb{E}_{\theta}\left[\frac{P(\theta|I,M)}{\mathbb{E}_{M'}P(\theta|I,M')}P(M)\right]$$

Given this image of your surroundings… … which map are you most likely in?

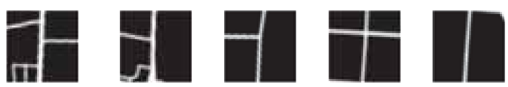

Correcting pose based on visual perception

What do we do when our GPS is not accurate or even not available? $\Rightarrow \ Var[P(M)] \gg 0$

$$P(M|I) = \mathbb{E}_\theta\left[\frac{P(\theta|I,M)}{\mathbb{E}_{M'}P(\theta|I,M')}P(M)\right]$$

Given this image of your surroundings
which map are you in?

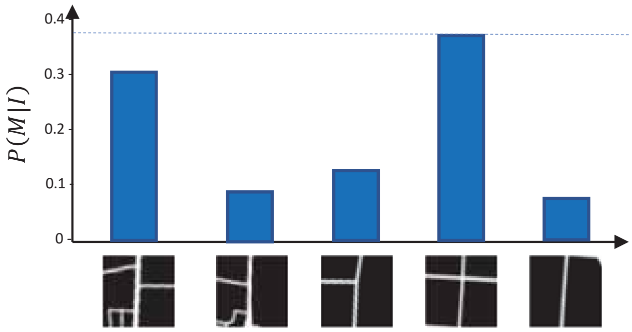

Increasing Scope of Training with Sim-to-Real

- End-to-end perception-to-control learning
- Imitate human driving through supervised learning
  - Dangerous to collect data from situations vehicles must be able to handle
  - Requires large amounts of “gold-standard” human driving
  - Difficulty in transferring to new domains, edge cases
- Allow agents to autonomously navigate and learn how to drive **without human supervision**
- Real world edge-cases and safety-critical scenarios

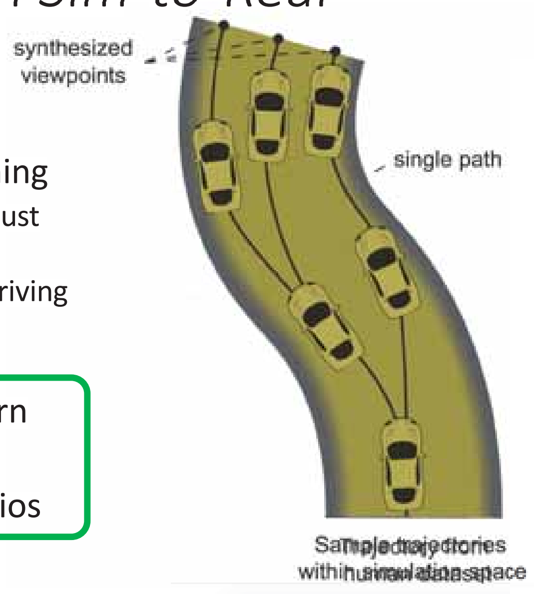

Related Works

Model-based Simulation

- Lacks photorealism
- Does not capture semantic complexity
- Does not transfer to real world (current state of art)
- [CARLA] Dosovitskiy et al (2017)
- [Torcs] Wymann et al (2000)

Domain Transfer

- Possible to transfer to real world
- Transfer limited to textures
- Lacks photorealism and semantic complexity
- [Wayve] Bewley et al (2018)
- Pan et al (2017)

Data-Driven Simulation

- Photorealistic + transferable
- Not scalable to large scale driving environments
  - [Gibson] Xia et al (2018)
- Deformities from non-realistic assumptions
  - [NVIDIA] Bojarski et al (2016)

Approach

1. **Photorealistic data-driven simulation** engine for synthesizing new control trajectories.

2. **Real-world transferable reinforcement learning**. End-to-end without human imitation.

End-to-end without human supervision

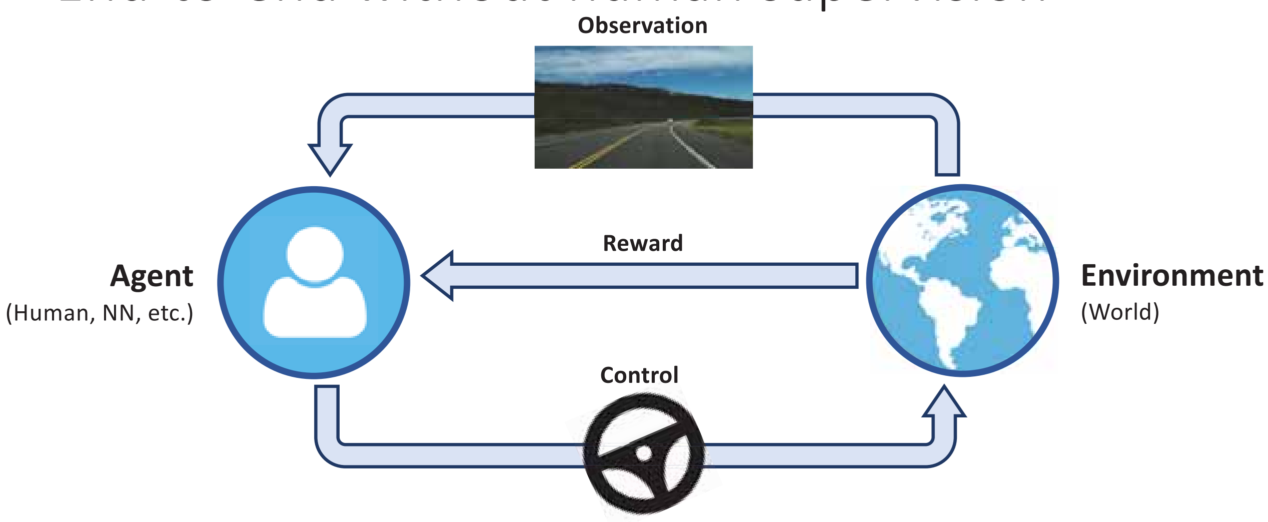

Optimizing high level reward functions

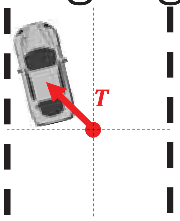

$$r_t = \begin{cases} 1 & \text{if } \|T\| < \varepsilon \\ 0 & \text{otherwise} \implies \text{"crash"} \end{cases}$$

Instead of imitating a human driver, directly **optimize the agent to maximize its own rewards**

$$\max_{\pi_\theta} \mathbb{E}_{\tau \sim \pi_\theta} \left[\sum_t \gamma^t r_t \right]$$

$$\tau = \{(a_1, s_1, r_1), (a_2, s_2, r_2), \dots\} \sim \pi_\theta(a_t | s_t)$$

End-to-end without human supervision

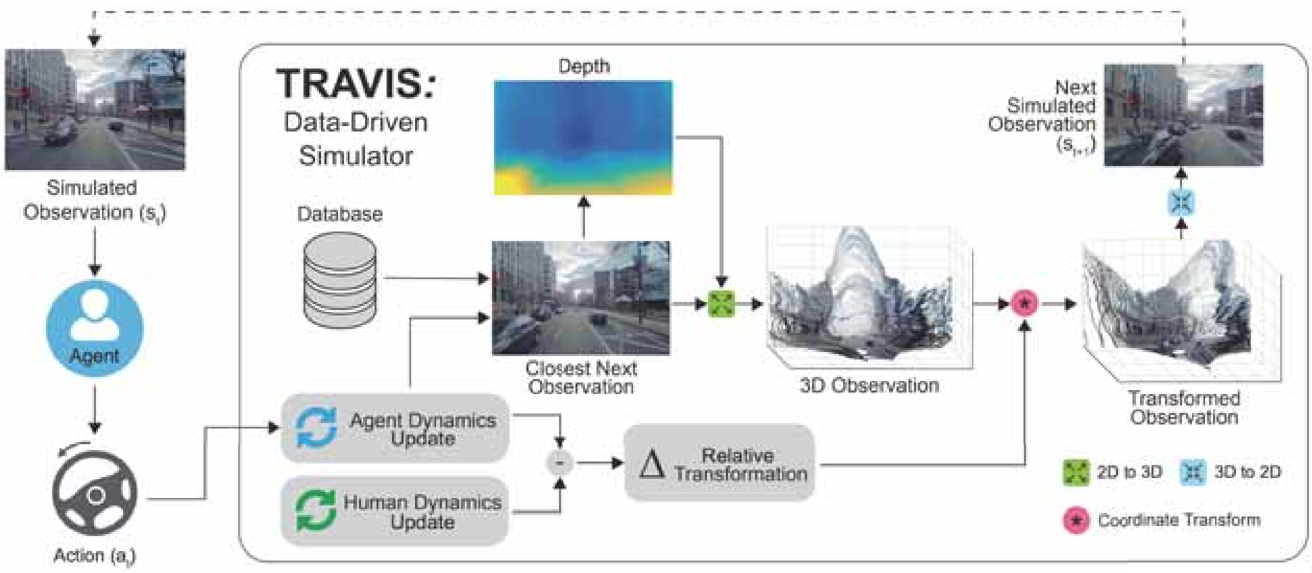

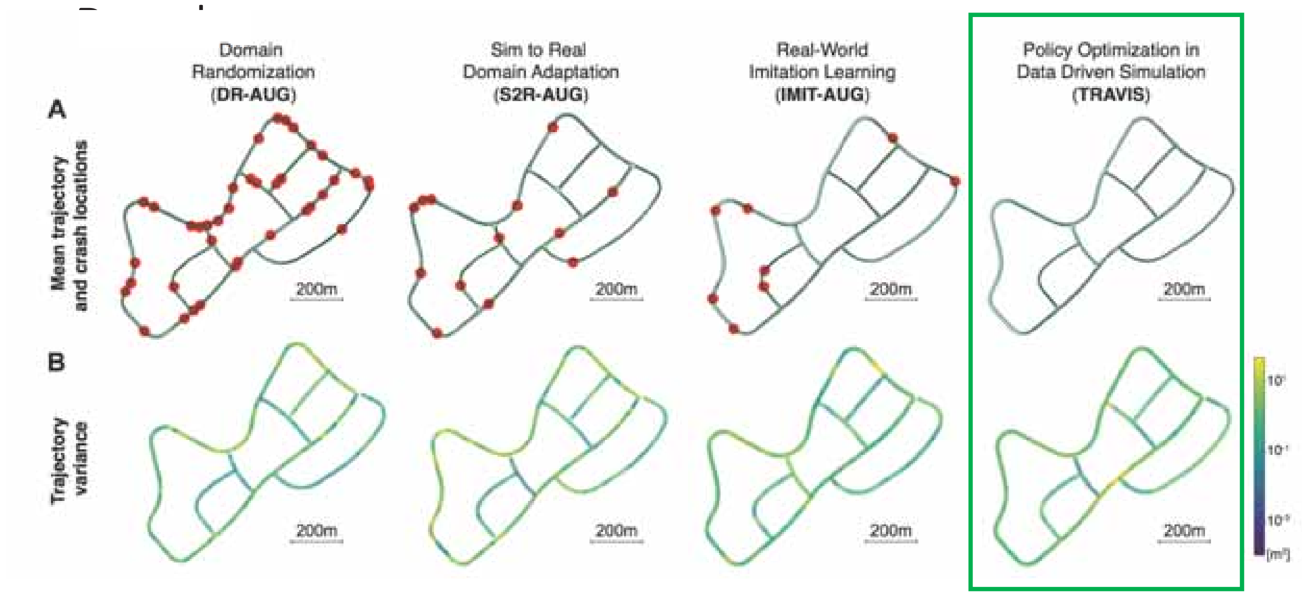

Results

Direct deployment to real-world without any adjustments

Superior robustness to recover from challenging off-orientations

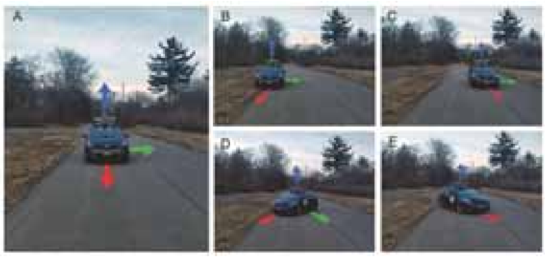

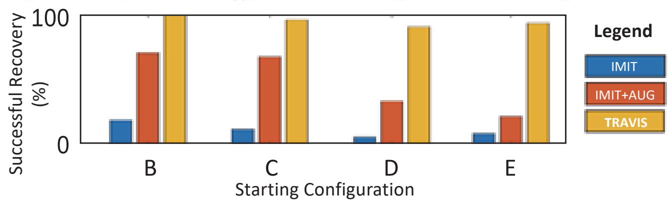

Challenging Environments: No Maps, Weather

Standard Method: High Resolution Maps

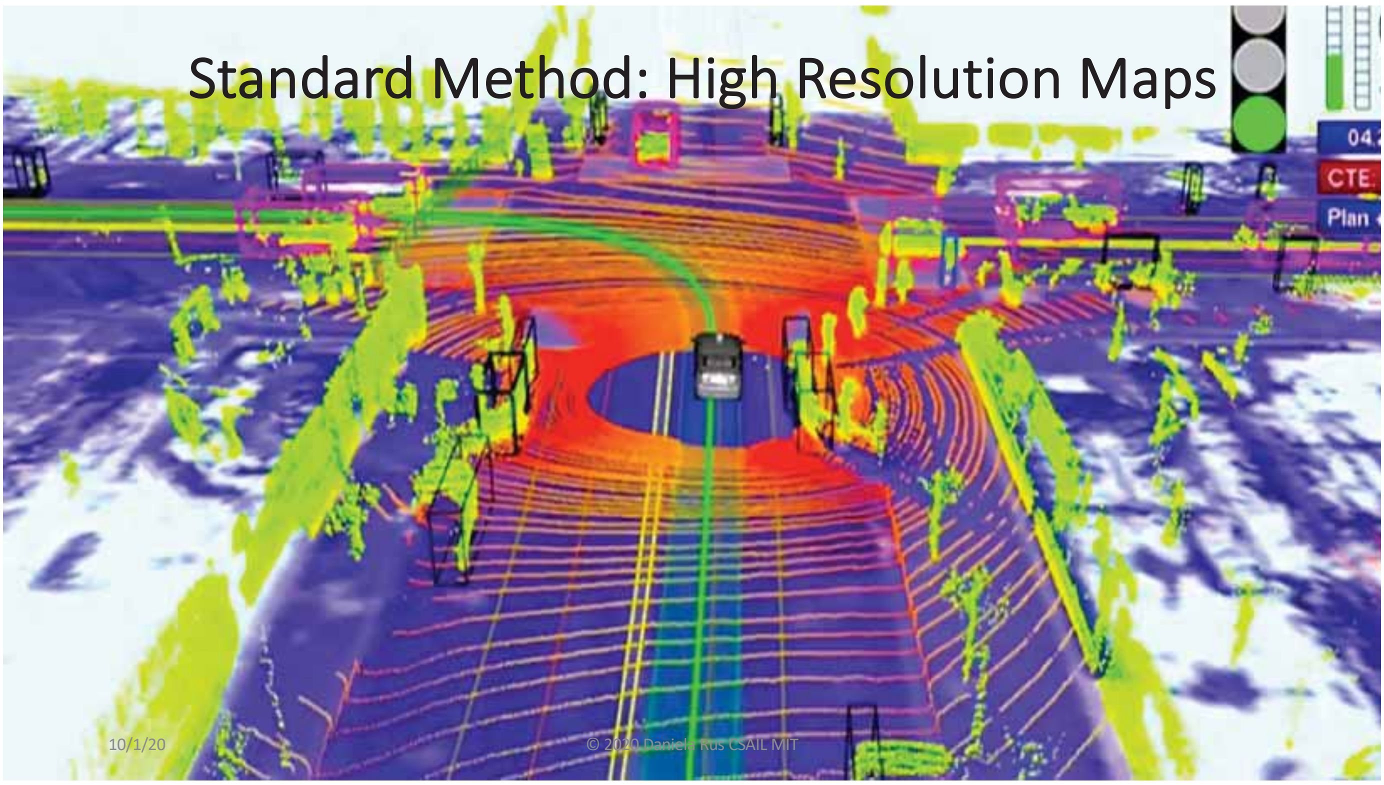

Challenges with Maps:

- Scalability
  - Maps can grow large; hard to store or transmit regions larger than small cities
- Maintenance
  - Maps must be maintained, small changes in the world can cause localization failure
- Coverage
  - Rural areas not densely populated and the landscape can change rapidly
- Features

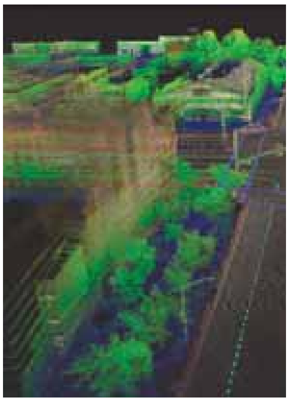

San Francisco, 4 TB
[Puttagunta, Civil Maps]

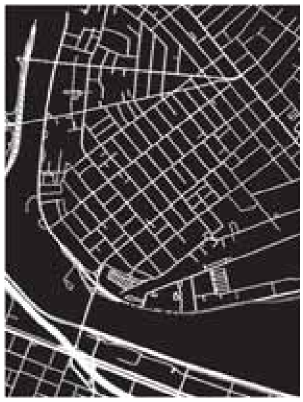

Earth, 40 GB
[Planet.osm]

Vision and LiDAR

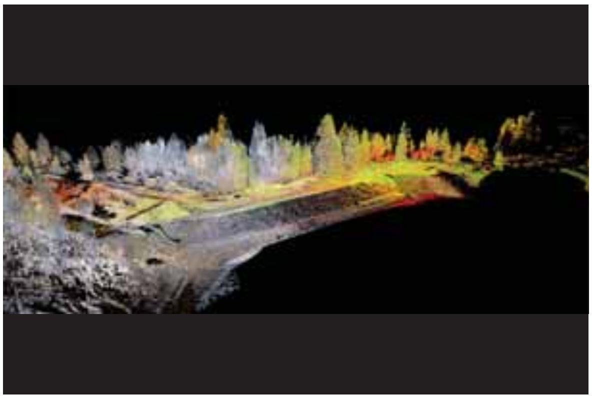

Changing Environments are Challenging

Localizing Ground Penetrating Radar

- Use Ground Penetrating Radar to build a map of underground features
- Soil content, type, layers can be reliably detected down to 2-3m
- Radar is unaffected by surface parameters like light and lidar

Related Work

Radar-based Perception

- With LiDAR
  [Rasshofer, ARS 2005]
- With Cameras
  [Mori, IV 2007]
- For SLAM
  [Schuster, Ward, IV 2016]

Appearance Modeling

- Dynamic Object Removal
  [Bescos, RA-L 2018]
- Stable Features
  [Dymczyk, 3DV 2007]
- Landmark Selection
  [Burki, JFR 2019]

Ground Penetrating Radar

- Soil Analysis
  [Rea, Water Resources Research 1998]
- Autonomous Analysis
  [Williams, IGARSS 2012]
- Localization
  [Cornick, JFR 2016]

LGPR System

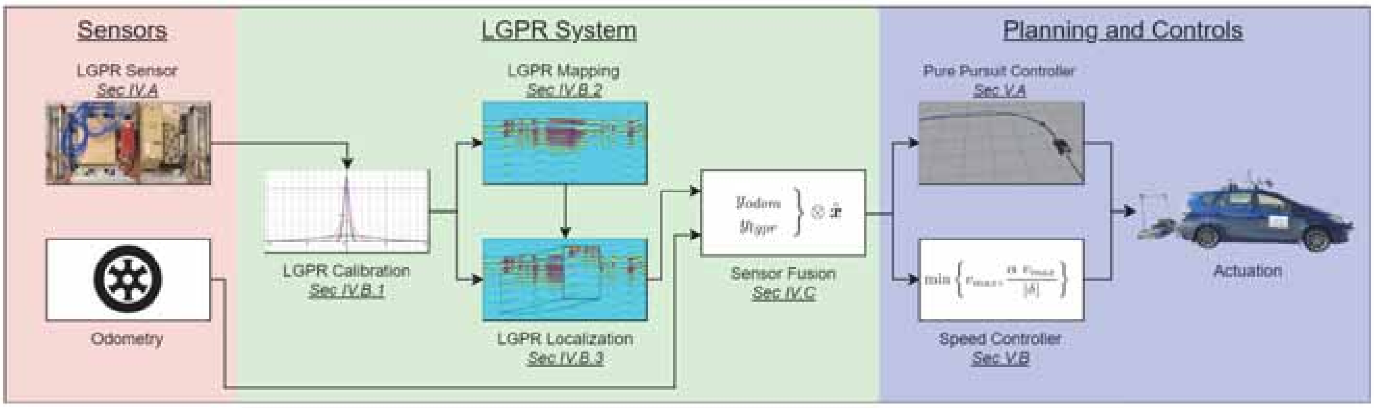

LGPR Sensor

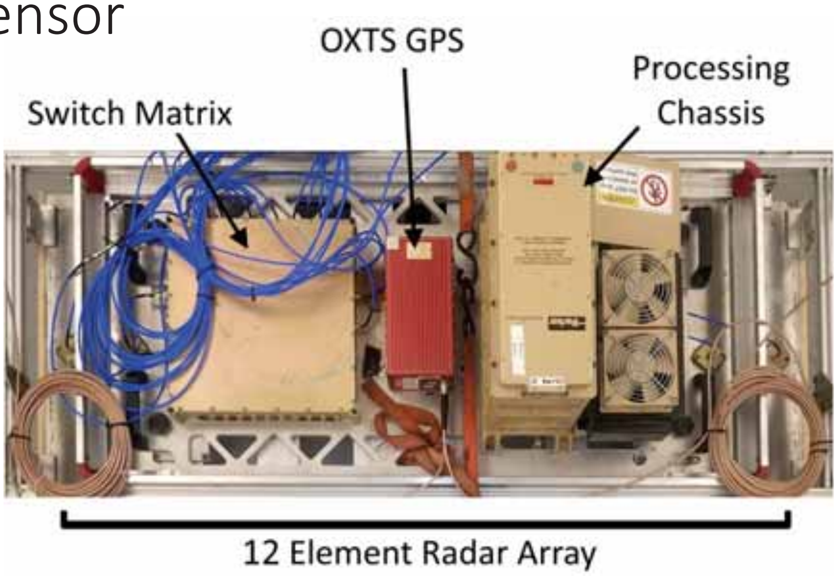

Mapping

- Sensor records 2D scans beneath the vehicle [11x369]
- Data rate is up to 126Hz
- Each scan is stored with a GPS location for localization

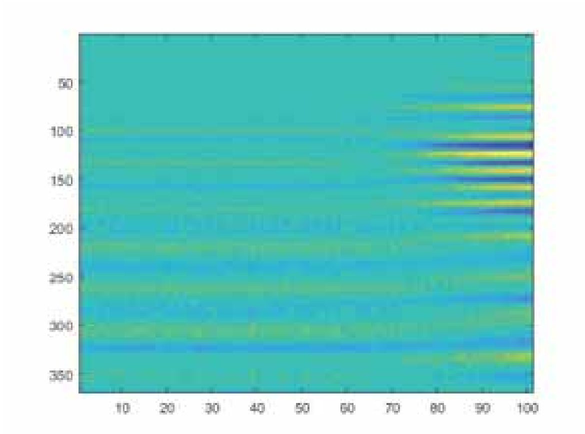

Localization

- A single scan is located in the map (5 DOF)
- GPS rough position limits search space
- Interpolation is used between map scans

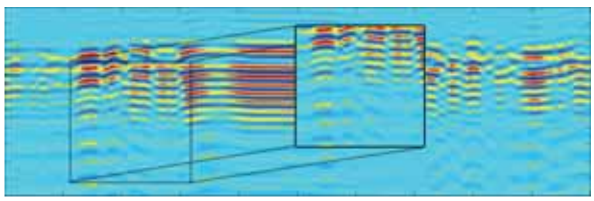

Correlation:

$$r_{A,B} = \frac{\sum_{i,d} A_{i,d} B_{i,d}}{\sqrt{\sum_{i,d} A_{i,d}^2 B_{i,d}^2}}$$

System Evaluation

- Full system implemented on autonomous Prius
- Real-Time-Kinematic GPS Inertial Navigation System for ground-truth

Relative mean error:

$$\frac{1}{n}\sum_{n}\left\|\left(T^{test}_{GNSS,i}-T^{map}_{GNSS,i}\right)-\left(T^{test}_{LGPR,i}-T^{map}_{LGPR,i}\right)\right\|$$

Driving in Weather Results

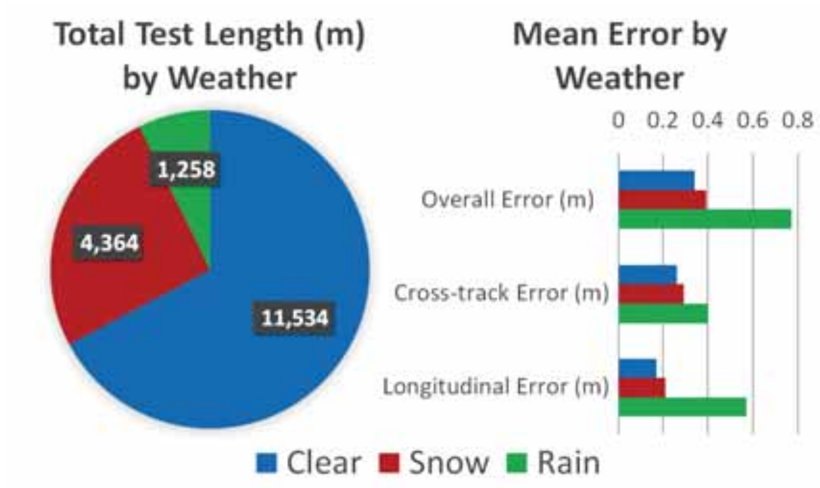

Challenging Interactions: Clutter, Human-Robot Systems

Navigating in Clutter

[ocregister.com]

[nacto.org]

- Agnostic to static and dynamic obstacles in environment
- Use density and velocity field to compute dynamic risk density

Defining a Safety Net

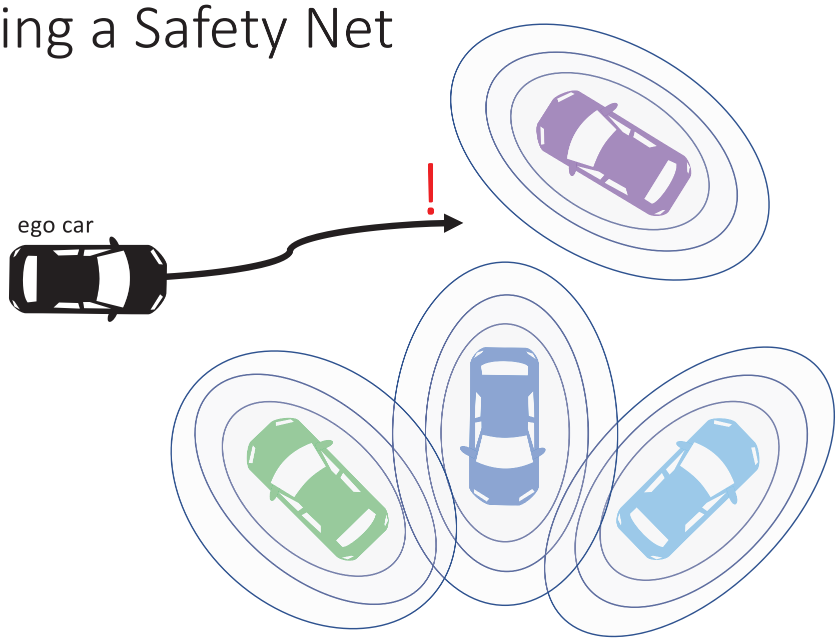

Defining a Safety Net: Risk Level Sets

- **Input:** position estimates and velocity
- **Assumptions:**
  - Other agents are self-preserving
  - Continue moving in current direction
- **Output Cost:**

$$H(q,x,t) = \sum_{i=1}^{n} \frac{\exp\left(-\left((q-x_i)^T \Omega_i (q-x_i)\right)^{\beta}\right)}{1+\exp\left(\alpha \dot{x}_i^T (q-x_i)\right)}$$

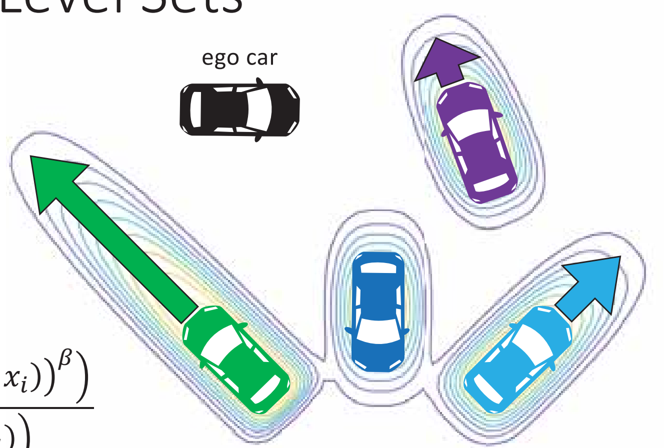

- **Risk Level Set:** $L_{\bar{p}} = \{q \mid H(q,x,t) < H_P\}$

1. **A. Pierson,** W. Schwarting, S. Karaman, and D. Rus, Navigating Congested Environments with Risk Level Sets, ICRA 2018
2. **A. Pierson,** W. Schwarting, S. Karaman, and D. Rus, Learning Risk Level Set Parameters from Data Sets for Safer Driving, IV 2019

Planning in Congestion with Risk Level Sets

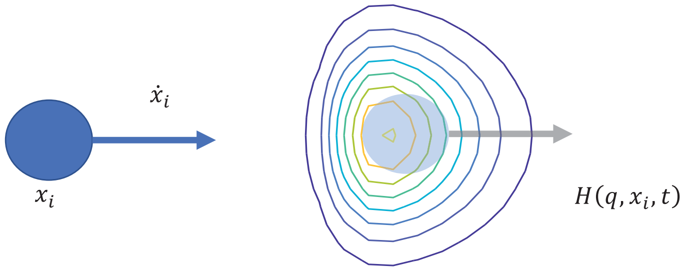

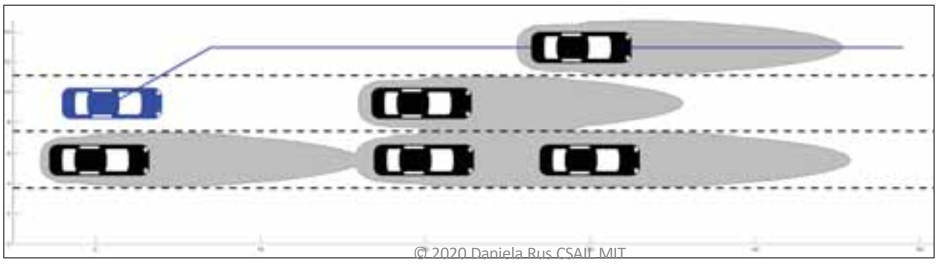

Density **Velocity Field** **Dynamic Risk Density**

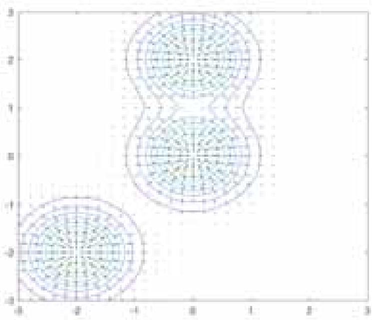

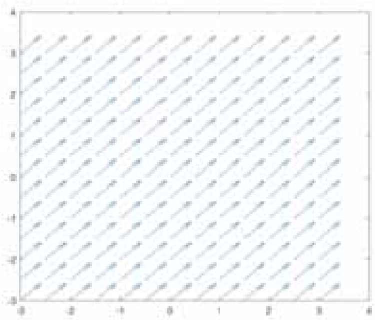

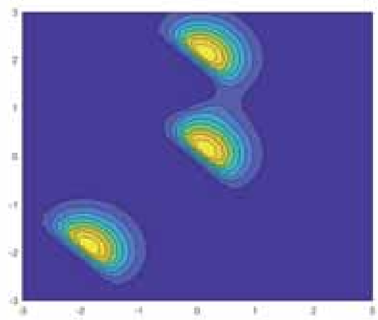

Risk Level Sets

- Congestion Cost

$$H(q,x,t) = \sum_{i=1}^{n} \frac{\exp\left(-\left((q-x_i)^T \Omega_i (q-x_i)\right)^{\beta}\right)}{1+\exp\left(\alpha \dot{x}_i^T (q-x_i)\right)}$$

- Create level set from cost

$$L_{\bar{p}} = \{q \mid H(q,x,t) < H_P\}$$

- Plan actions within $L_{\bar{p}}$
- Higher value of H_P → higher risk (**ICRA 2018**)

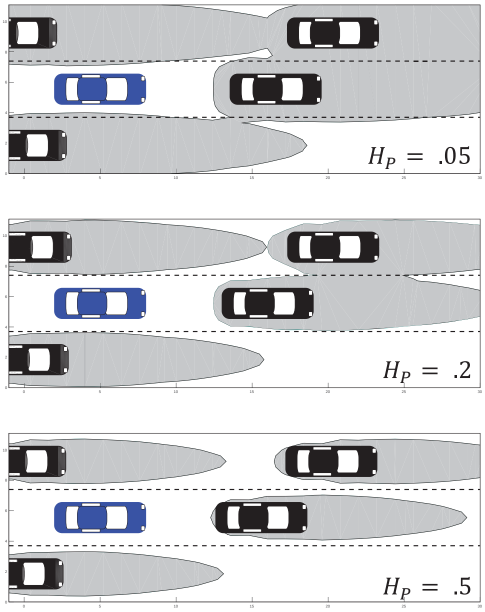

Simulation: Conservative vs Aggressive Driver

- White area: $L_{\bar{p}} = \{q \mid H(q,x,t) < H_P\}$
- Low H_P → lower risk, more conservative

Simulation: Conservative vs Aggressive Driver

- White area: $L_{\bar{p}} = \{q \mid H(q,x,t) < H_P\}$
- Low H_P → lower risk, more conservative

Simulation: Conservative vs Aggressive Driver

- Higher H_P → larger planning space $L_{\bar{p}}$
- More lane changes

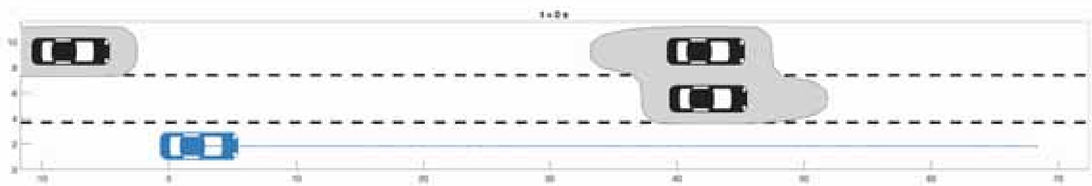

Simulation: Multiple Drivers

- Each car views other cars as obstacles
- Route planner updates to other cars changing lanes

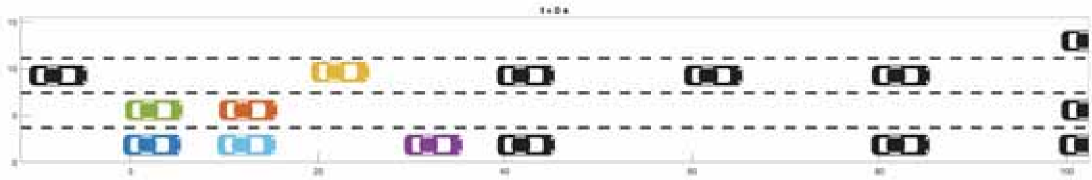

CARLA Validation: Risk Level Sets

- **Blue cars:** ego agents running risk level sets algorithm

Key Features:

- Integration into our codebase across other platforms
- Multi-ego-vehicle scenarios

Learning Risk Level Set Parameters from Data

- NGSIM and HighD data set validation (**IV 2019**)
- Quickly identify distributions of environment and driver features

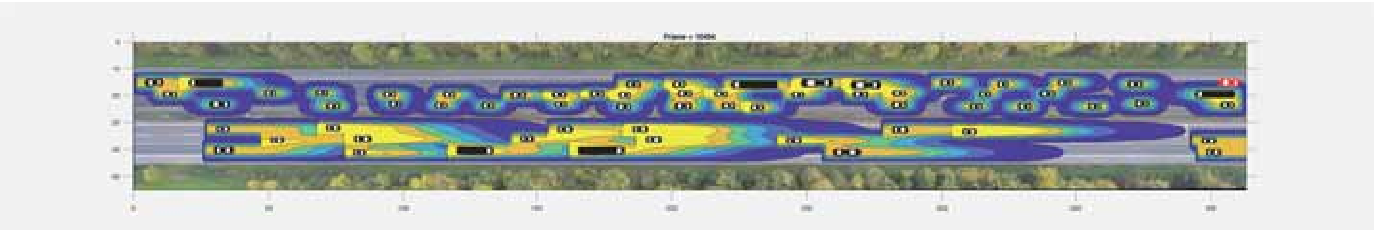

Risk Level Sets without Object Detection

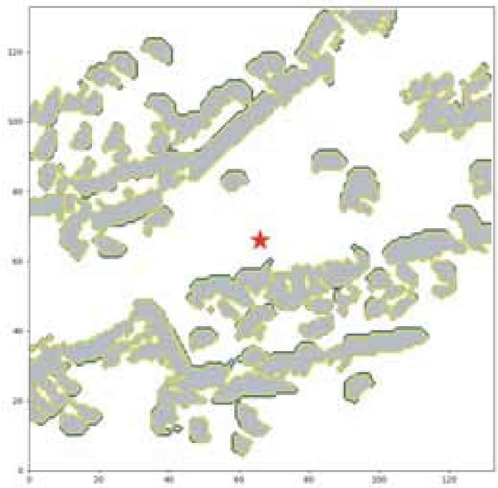

[1] Navigating Congested Environments with Risk Level Sets, ICRA 2018, patent pending
[2] Dynamic Risk Density for Autonomous Navigation in Cluttered Environments without Object Detection, ICRA 2019, submitted

Mixed Human Driven-Robot Car Systems

The Information

Waymo's Big Ambitions Slowed by Tech Trouble

CHANDLER, Ariz. – Alphabet's Waymo unit is a worldwide leader in autonomous vehicle development for suburban environments. It has said it would launch a driverless robo-taxi service to suburban Phoenix residents this year. Yet its self-driving minivan prototypes have trouble crossing the T-intersection closest to the company's Phoenix-area headquarters here.

Two weeks ago, Lisa Hargis, an administrative assistant who works at an office a stone's throw from Waymo's vehicle depot, said she nearly hit a Waymo Chrysler Pacifica minivan because it stopped abruptly while making a right turn at the intersection.

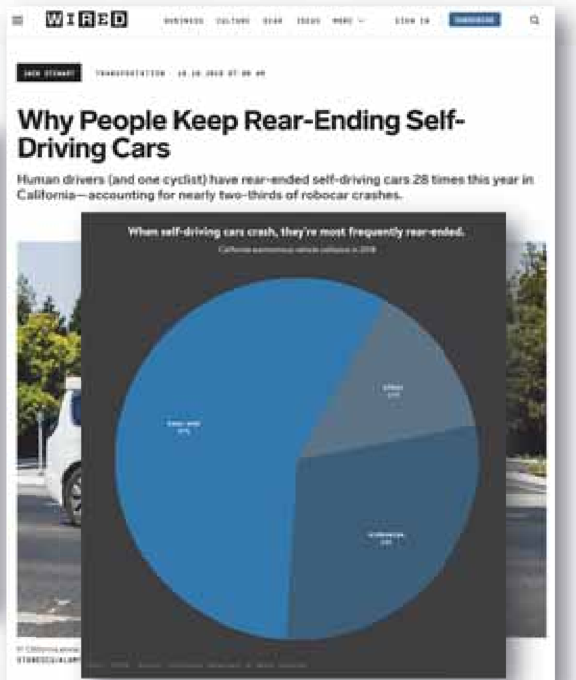

Why People Keep Rear-Ending Self-Driving Cars

Human drivers (and one cyclist) have rear-ended self-driving cars 28 times this year in California—accounting for nearly two-thirds of robocar crashes.

1. https://www.theinformation.com/articles/waymos-big-ambitions-slowed-by-tech-trouble
2. https://www.wired.com/story/self-driving-car-crashes-rear-endings-why-charts-statistics/

Autonomous Driving: Social Dilemma

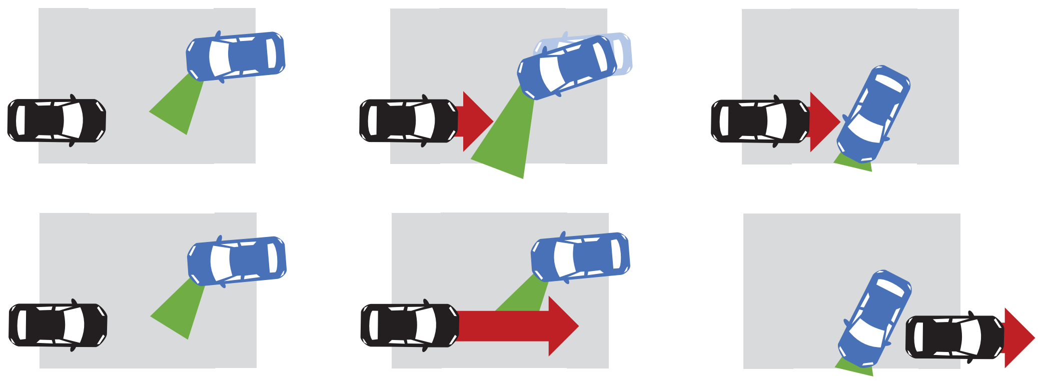

Social dilemmas: Situations in which collective interests are at odds with private interests

Social Value Orientation (SVO) Ring

Capturing Human Preferences in Social Dilemmas

Altruistic: Maximize other party's utility, without consideration of own outcome.

Prosocial: Benefiting a group as a whole.

Individualistic: Maximize their own outcome, without concern of the utility of other agents.

Competitive: Improve relative gain over others.

Cooperative: All agents are better off.

[1] W. B. G. Liebrand and C. G. McClintock, "The ring measure of social values: A computerized procedure for assessing individual differences in information processing and social value orientation," European Journal of Personality, vol. 2, no. 3, pp. 217–230, 1988.

Social Value Orientation (SVO) Ring

Studies of Human Preferences

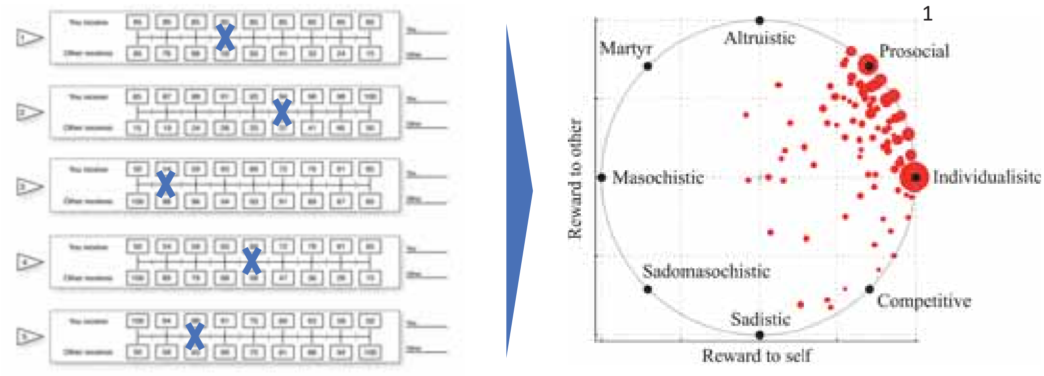

~ 90% of individuals are either prosocial (~ 50%) or individualistic (~ 40%) [2]

[1] A. Garapin, L. Muller, and B. Rahali, "Does trust mean giving and not risking? experimental evidence from the trust game," Revue d''economie politique, vol. 125, no. 5, pp. 701–716, 2015.
[2] R. O. Murphy, K. A. Ackermann, and M. Handgraaf, "Measuring social value orientation," Judgment and Decision Making, vol. 6, no. 8, pp. 771–781, 2011.

Split $100 with a stranger…

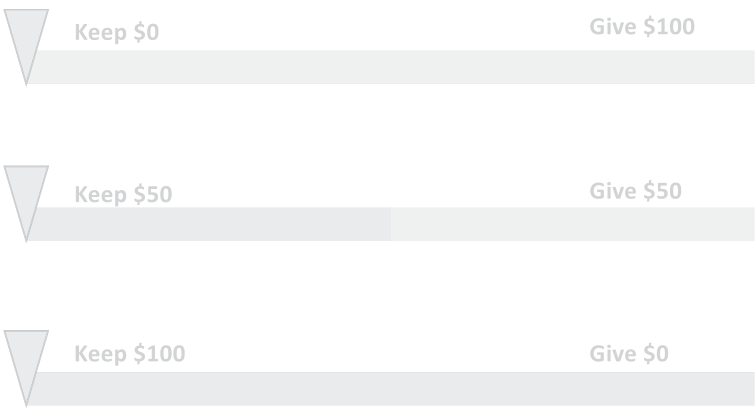

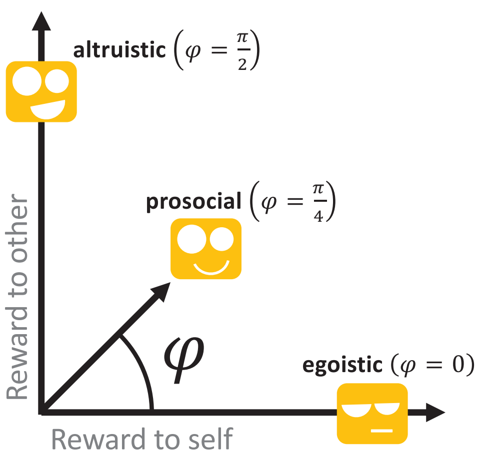

- A. Garapin, L. Muller, and B. Rahali, Does trust mean giving and not risking? experimental evidence from the trust game, *Revue d''economie politique*, 2015
- R. O. Murphy, K. A. Ackermann, and M. Handgraaf, Measuring social value orientation, *Judgment and Decision Making*, 2011

Social Value Orientation

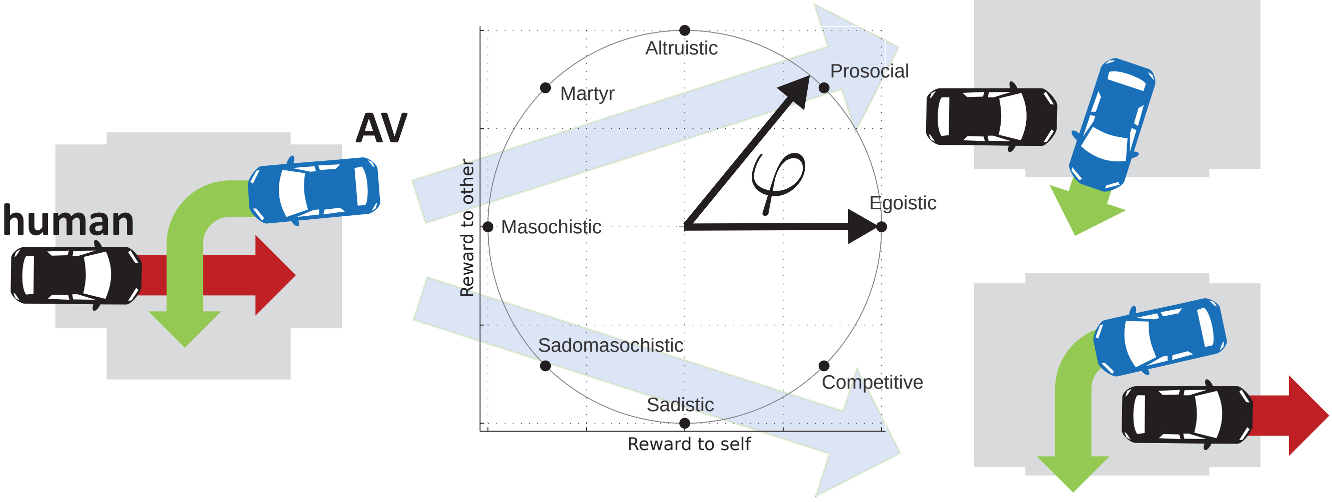

Social Value Orientation

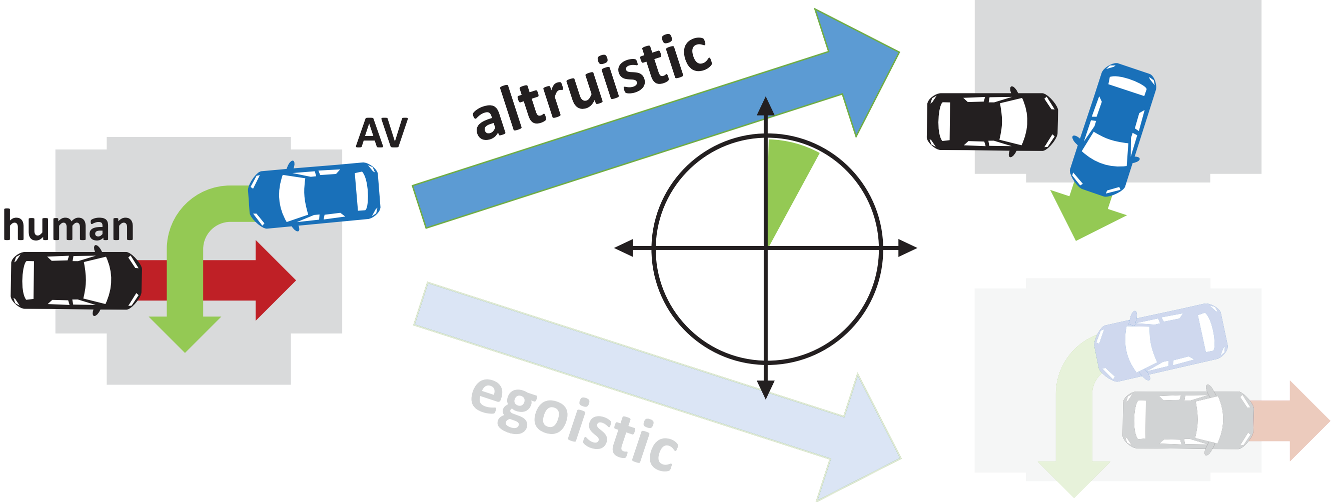

Social Value Orientation

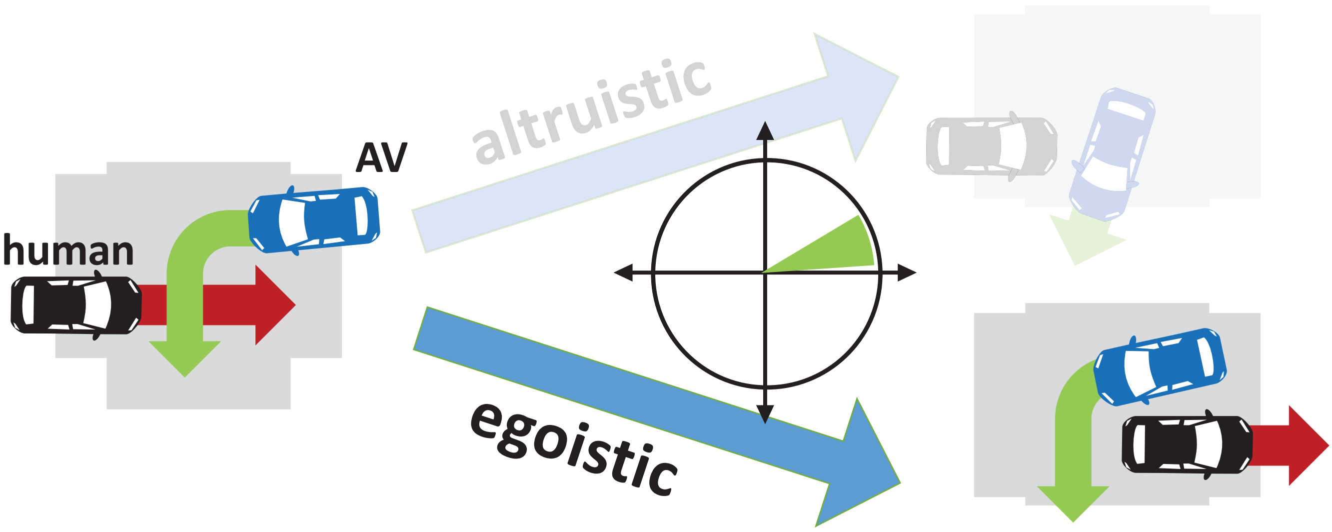

Our Approach

Social Value Orientation

- Behavior model from social psychology

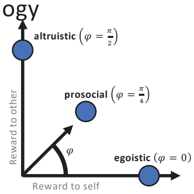

$$g_i(\cdot) = \cos\varphi_i\, r_i + \sin\varphi_i\, r_j$$

- Weight reward to self vs other

Best Response Game

- Each agent maximizes its individual utility

$$G_i(\boldsymbol{x}^0, \boldsymbol{u}, \varphi_i) = \sum_{k=0}^{N-1} g_i(\boldsymbol{x}^k, \boldsymbol{u}^k, \varphi_i) + g_i^N(\boldsymbol{x}^N, \varphi_i)$$

$$\boldsymbol{u}_i^* = \operatorname*{argmax}_{\boldsymbol{u}_i} G_i(\boldsymbol{x}^0, \boldsymbol{u}_i, \boldsymbol{u}_{\neg i}, \varphi_i)$$

- Solve for Nash Equilibrium

Learned Rewards

- Inverse Reinforcement Learning
- Calibrate rewards on NGSIM data set

1. W. Schwarting, **A. Pierson**, J. Alonso-Mora, S. Karaman, and D. Rus, Social Behavior for Autonomous Vehicles. *Proceedings of the National Academy of Sciences (PNAS), 2019*

Utility-Maximizing Policy with SVO

- Joint reward weighted by SVO

$$g_i(\cdot) = \underline{\cos\varphi_i\, r_i} + \underline{\sin\varphi_i\, r_j}$$

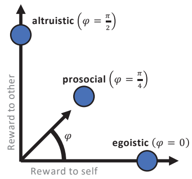

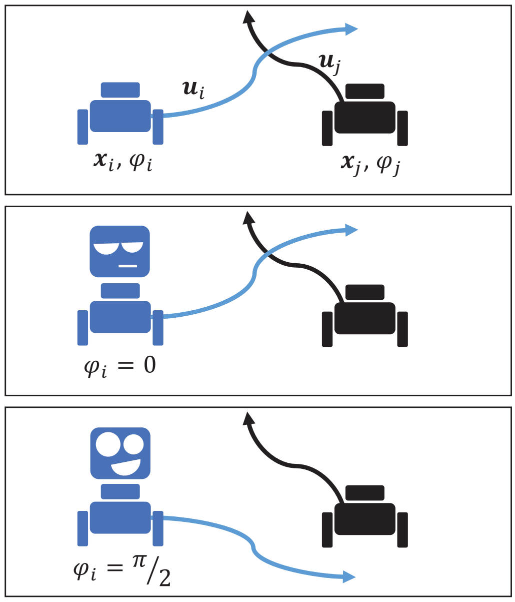

- Utility over time horizon

$$G_i(\boldsymbol{x}^0, \boldsymbol{u}, \varphi_i) = \sum_{k=0}^{N-1} g_i\left(\boldsymbol{x}^k, \boldsymbol{u}^k, \varphi_i\right) + g_i^N(\boldsymbol{x}^N, \varphi_i)$$

- Find control $\boldsymbol{u}_i^*$ that maximizes utility

$$\boldsymbol{u}_i^* = \underset{\boldsymbol{u}_i}{\operatorname{argmax}}\, G_i(\boldsymbol{x}^0, \boldsymbol{u}_i, \boldsymbol{u}_{\neg i}, \varphi_i)$$

Unprotected Left Turns

The AV must wait for an altruistic driver to yield

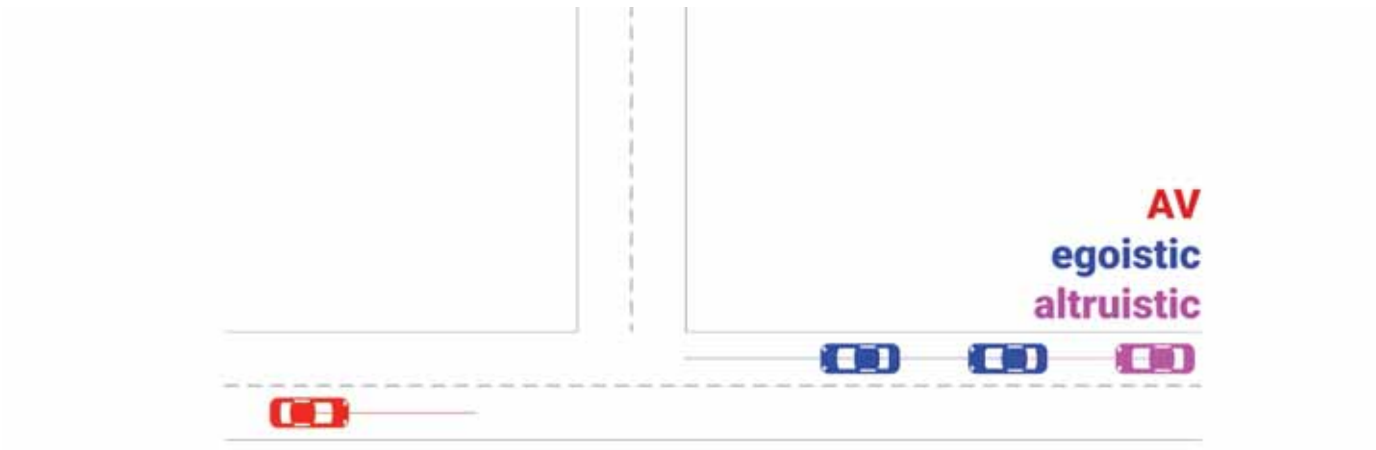

1. W. Schwarting, **A. Pierson**, J. Alonso-Mora, S. Karaman, and D. Rus, Social Behavior for Autonomous Vehicles. *Proceedings of the National Academy of Sciences (PNAS), 2019*

Egoistic Merge

Among egoistic drivers, the AV must wait to merge

1. W. Schwarting, **A. Pierson**, J. Alonso-Mora, S. Karaman, and D. Rus, Social Behavior for Autonomous Vehicles. *Proceedings of the National Academy of Sciences (PNAS), 2019*

Prosocial Merge

Prosocial drivers create a gap for the AV to merge

1. W. Schwarting, **A. Pierson**, J. Alonso-Mora, S. Karaman, and D. Rus, Social Behavior for Autonomous Vehicles. *Proceedings of the National Academy of Sciences (PNAS), 2019*

Estimate SVO online

Estimate SVO of other drivers online

Integrate into motion planner to improve decision-making and predictions

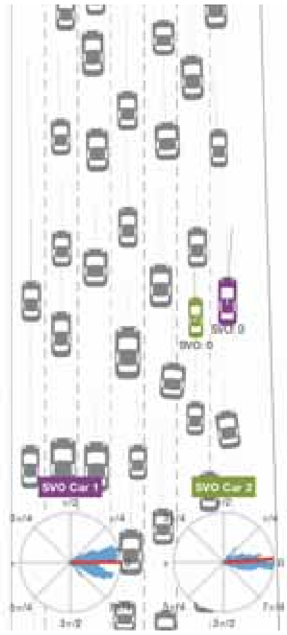

SVO Predictions on NGSIM

Estimating driver SVO improves trajectory predictions by 25%

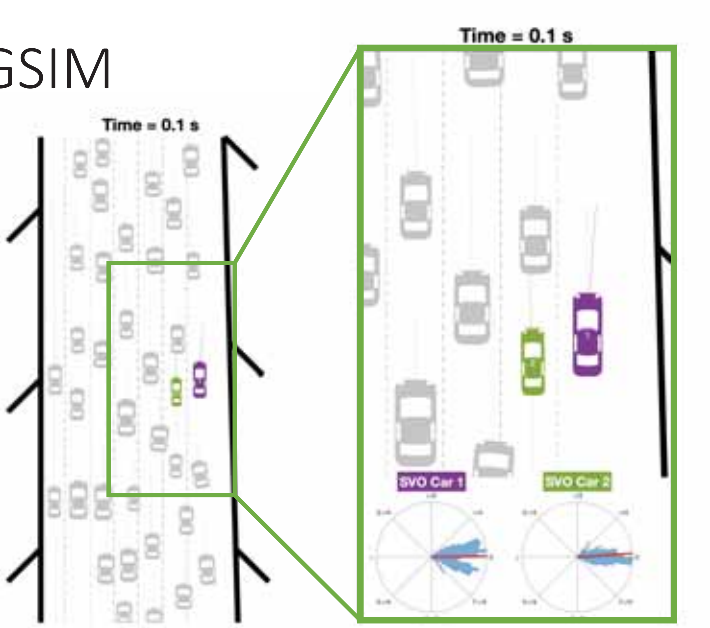

1. W. Schwarting, **A. Pierson**, J. Alonso-Mora, S. Karaman, and D. Rus, Social Behavior for Autonomous Vehicles. *Proceedings of the National Academy of Sciences (PNAS), 2019*

SVO Trends in NGSIM dataset

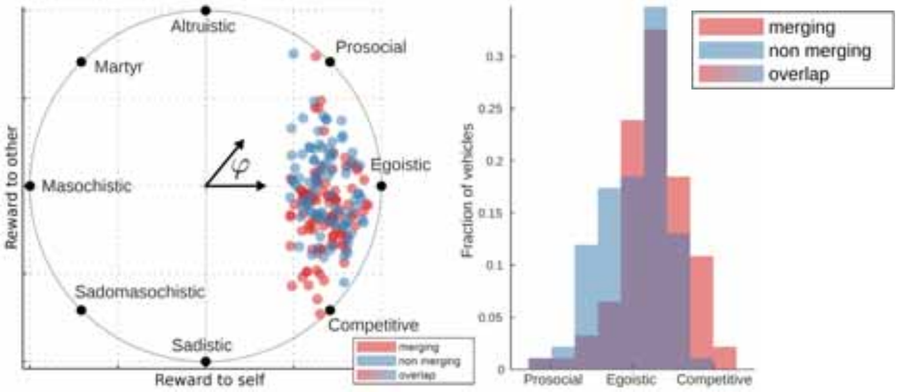

Merging vehicles are more competitive than non merging vehicles ($p < 0.002$)

1. W. Schwarting, **A. Pierson**, J. Alonso-Mora, S. Karaman, and D. Rus, Social Behavior for Autonomous Vehicles. *Proceedings of the National Academy of Sciences (PNAS), 2019*

Evaluation of SVO on NGSIM dataset

Improved prediction with dynamically estimated SVO during merges

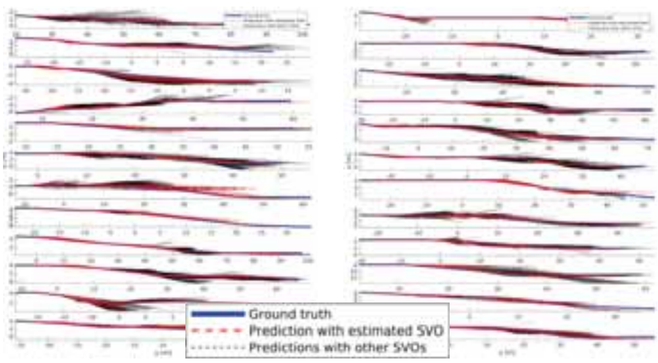

Prediction	baseline	multi-agent game theoretic		
SVO	-	egoistic	static best	estimated
MSE position	1.0	0.947	0.821	**0.753**

Table 1. Relative mean square position error (MSE) between predicted and actual trajectories, as compared to a single-agent planning baseline.

25% reduced prediction error

10/1/20
© 2020 Daniela Rus CSAIL MIT

What will come first?

Level 5 Autonomy, or

The Flying car?

Conclusions

- Today: self-driving cars at low speed in low complexity environments
- Tomorrow: increased speed and complexity, mobility as a service
- The Future: Pervasive self-driving (flying) cars, pervasive robotics