Mitigating the "Element of Surprise" in Model-Based Robot Planning

Edward Schmerling

Joint work with Karen Leung, Boris Ivanovic, Prof. Mo Chen, Prof. Marco Pavone, et al.

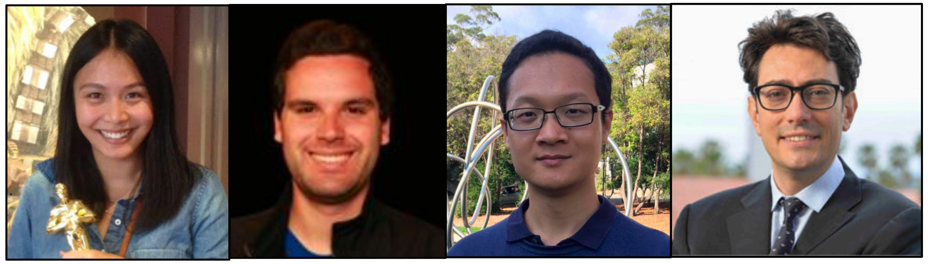

Stanford University

Autonomous Systems Lab

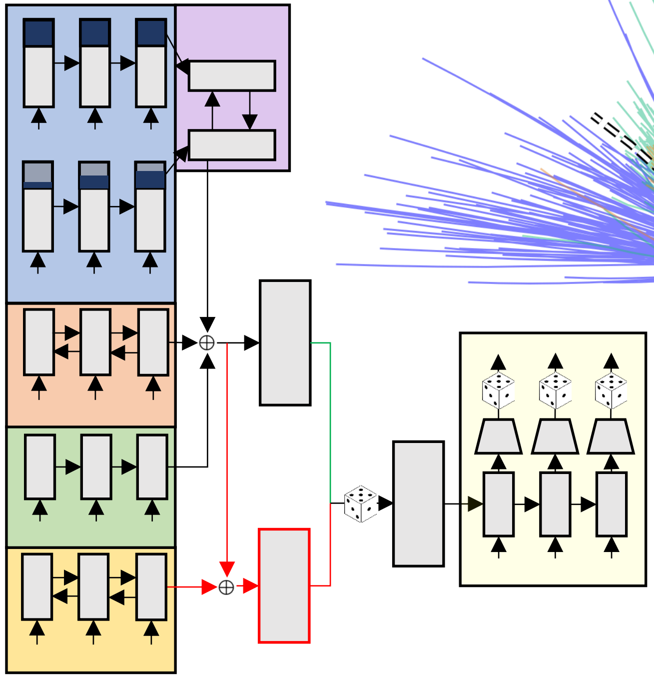

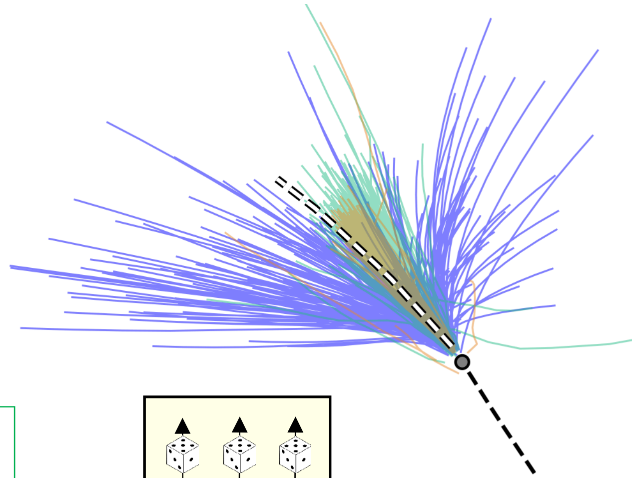

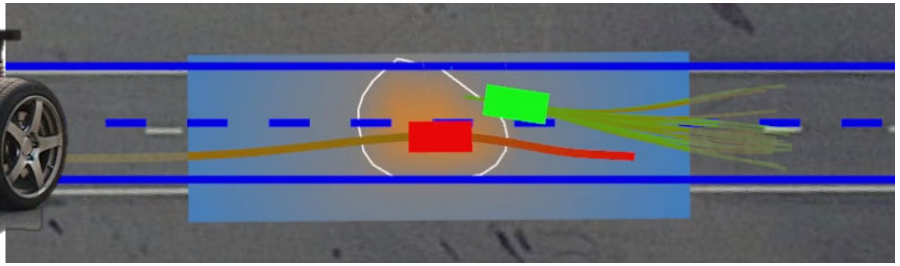

Geoff Caddick - https://www.brinknews.com/wp-content/uploads/2017/05/ron_harbour_int-631893050-1025x685.jpg

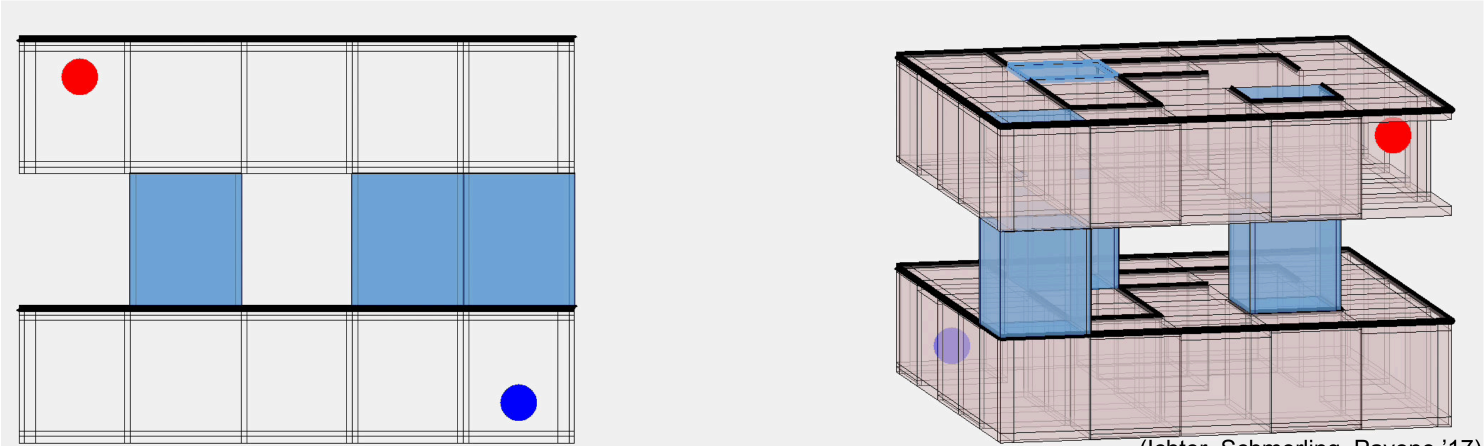

(Ichter, Schmerling, Pavone '17)

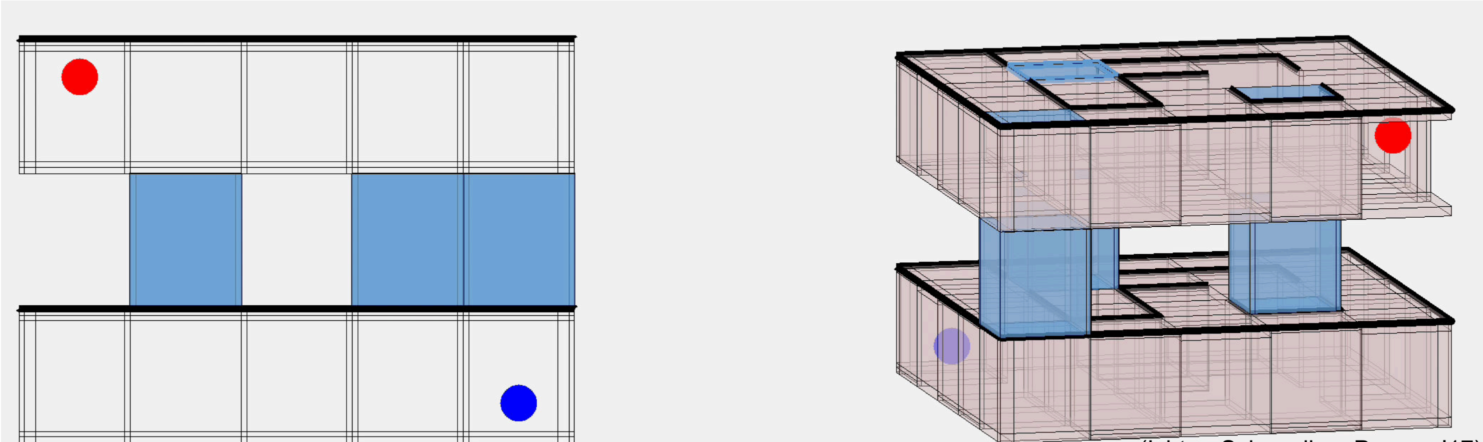

(Ichter, Schmerling, Pavone ’17)

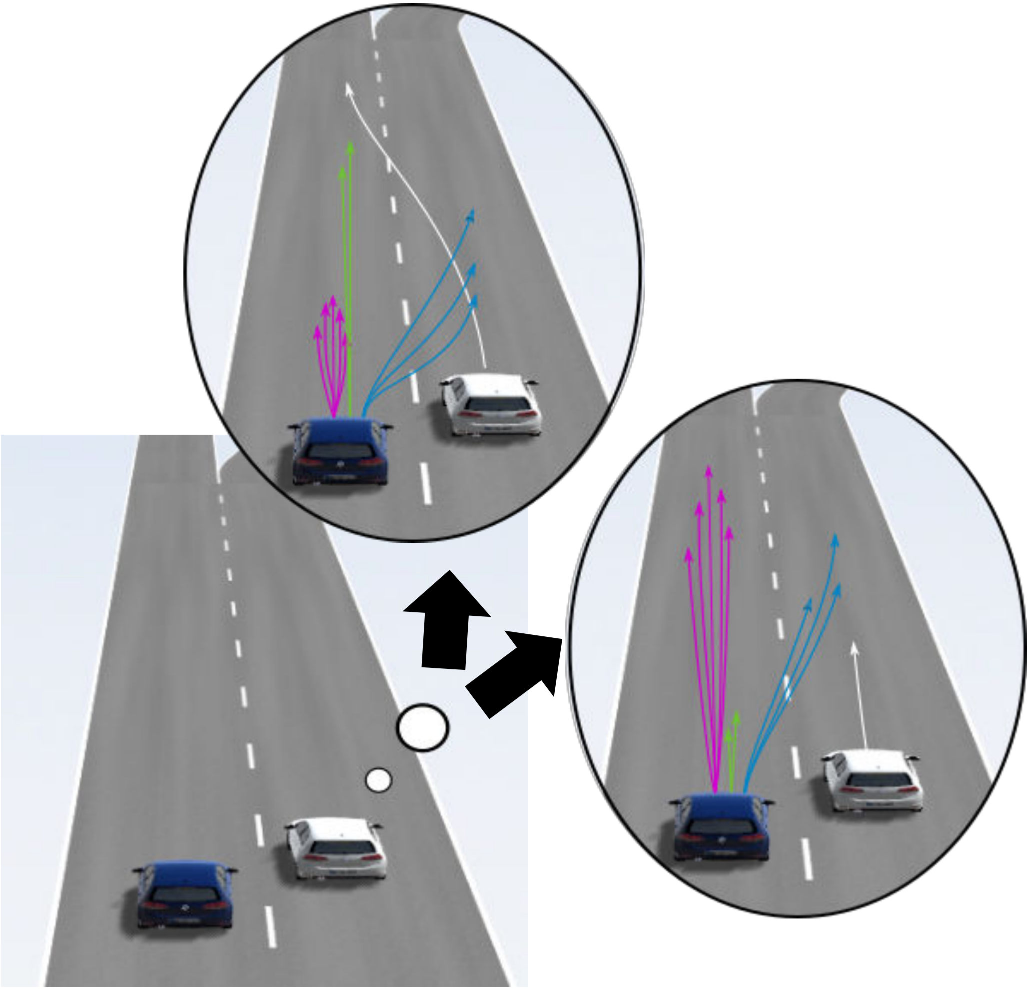

The “Element of Surprise” in HRI

- Intent-ambiguous scenarios
  Multiple highly distinct possible outcomes

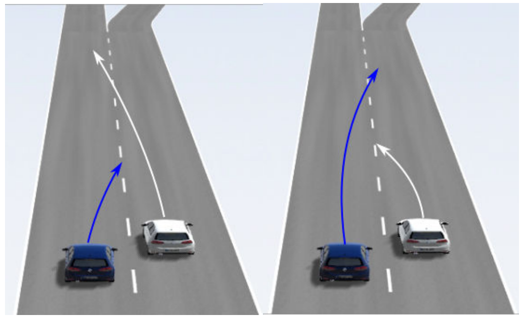

The “Element of Surprise” in HRI

- Intent-ambiguous scenarios
  Multiple highly distinct possible outcomes

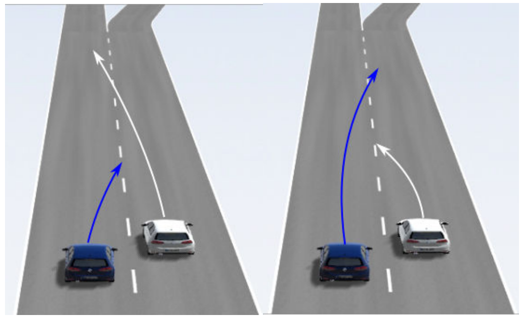

- Dynamically evolving scenarios
  Variable number of agents coming into and out of relevance

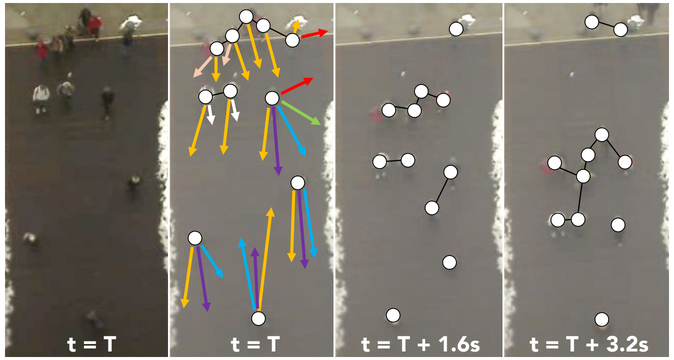

The “Element of Surprise” in HRI

- Intent-ambiguous scenarios
  Multiple highly distinct possible outcomes

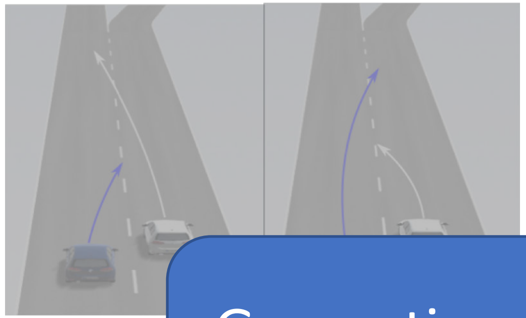

- Dynamically evolving scenarios
  Variable number of agents coming into and out of relevance

Generative modeling of interaction dynamics

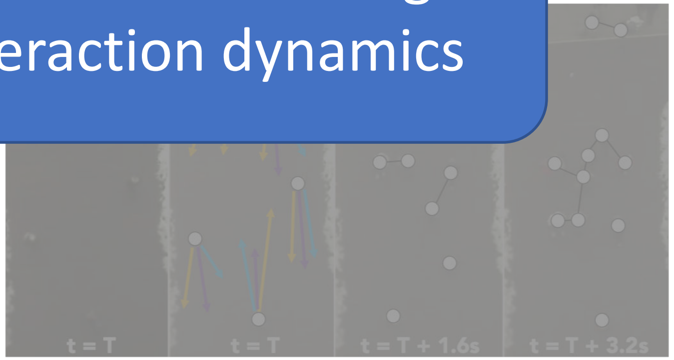

The “Element of Surprise” in HRI

- Intent-ambiguous scenarios
  Multiple highly distinct possible outcomes

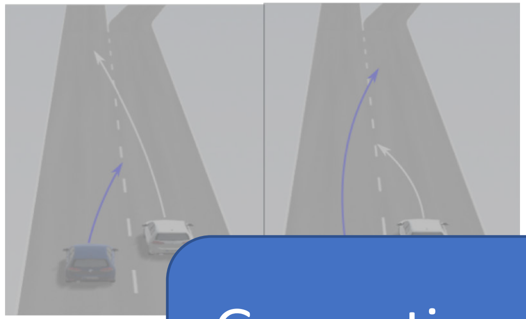

Generative modeling of interaction dynamics

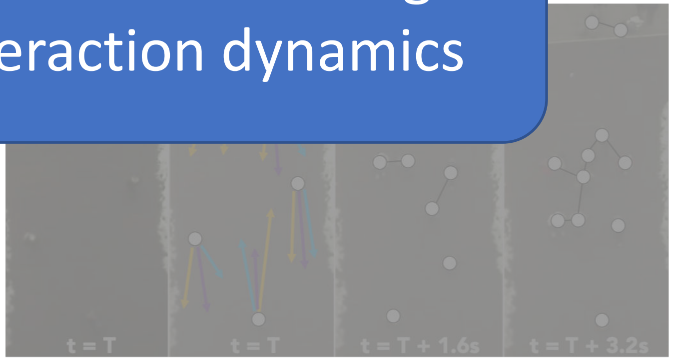

- Dynamically evolving scenarios
  Variable number of agents coming into and out of relevance

- Scenarios where these predictive models fail?
  Humans that consistently defy a robot’s expectations

The “Element of Surprise” in HRI

- Intent-ambiguous scenarios
  Multiple highly distinct possible outcomes

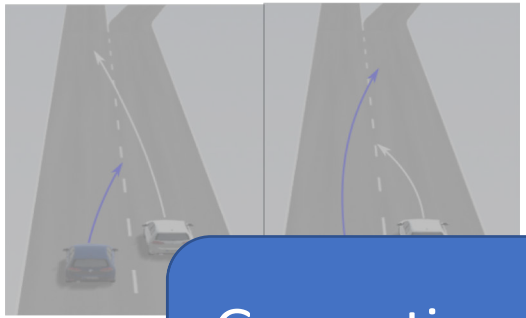

Generative modeling of interaction dynamics

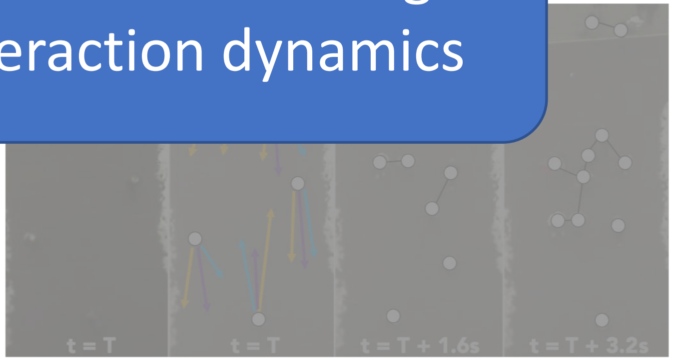

- Dynamically evolving scenarios
  Variable number of agents coming into and out of relevance

- Scenarios where these predictive models fail?
  Humans that consistently defy a robot’s expectations

Deterministic safety controller

Probabilistic Agent Modeling

We consider generative models of human action distributions conditioned on

- Joint interaction history of all agents in the scene
- Candidate robot future action sequence

$$x_H^{(t+1)} = f_H(x_H^{(t)}, u_H^{(t)}) \quad \text{(human)}$$

$$x_R^{(t+1)} = f_R(x_R^{(t)}, u_R^{(t)}) \quad \text{(robot)}$$

$$x^{(t)} = (x_H^{(t)}, x_R^{(t)}) \quad \text{(joint state)}$$

$$u^{(t)} = (u_H^{(t)}, u_R^{(t)}) \quad \text{(joint control)}$$

$$\underbrace{p(u_H^{(t+1:t+N)}}_{\mathbf{y}} \mid \underbrace{x^{(0:t)}, u^{(0:t)}, u_R^{(t+1:t+N)}}_{\mathbf{x}}) = \prod_{i=1}^{N} p(u_H^{(t+1)} \mid x^{(0:t)}, u^{(0:t)}, u_R^{(t+1:t+N)}, u_H^{(t+1:t+i-1)})$$

$\mathbf{y}$: Human future

$\mathbf{x}$: Joint history, robot future

Probabilistic Agent Modeling

We consider generative models of human action distributions conditioned on

- Joint interaction history of all agents in the scene
- Candidate robot future action sequence

$$x_H^{(t+1)} = f_H(x_H^{(t)}, u_H^{(t)}) \quad \text{(human)} \qquad x^{(t)} = (x_H^{(t)}, x_R^{(t)}) \quad \text{(joint state)}$$

$$x_R^{(t+1)} = f_R(x_R^{(t)}, u_R^{(t)}) \quad \text{(robot)} \qquad u^{(t)} = (u_H^{(t)}, u_R^{(t)}) \quad \text{(joint control)}$$

$$\underbrace{p(u_H^{(t+1:t+N)}}_{\mathbf{y}} \mid \underbrace{x^{(0:t)}, u^{(0:t)}, u_R^{(t+1:t+N)}}_{\mathbf{x}}) = \prod_{i=1}^{N} \underbrace{p(u_H^{(t+1)} \mid x^{(0:t)}, u^{(0:t)}, u_R^{(t+1:t+N)}, u_H^{(t+1:t+i-1)})}_{\text{Apply neural network function approximator}}$$

$\mathbf{y}$: Human future

$\mathbf{x}$: Joint history, robot future

Apply neural network function approximator

Probabilistic Agent Modeling

We consider generative models of human action distributions conditioned on

- Joint interaction history of all agents in the scene
- Candidate robot future action sequence

$$\underbrace{p(u_H^{(t+1:t+N)}}_{\mathbf{y}} \mid \underbrace{x^{(0:t)}, u^{(0:t)}, u_R^{(t+1:t+N)}}_{\mathbf{x}}) = \prod_{i=1}^{N} p(u_H^{(t+1)} \mid x^{(0:t)}, u^{(0:t)}, u_R^{(t+1:t+N)}, u_H^{(t+1:t+i-1)})$$

For robot planning purposes we want models that are:

- Capable of capturing the full breadth of future outcomes (i.e., not regression)
- Amenable to computationally efficient sampling
- Flexible with respect to evolving multi-agent scenarios yet scalable to tens of possibly relevant agents → we seek a modular model architecture

Multi-Agent Trajectory Prediction

Take the following highway scene as an example:

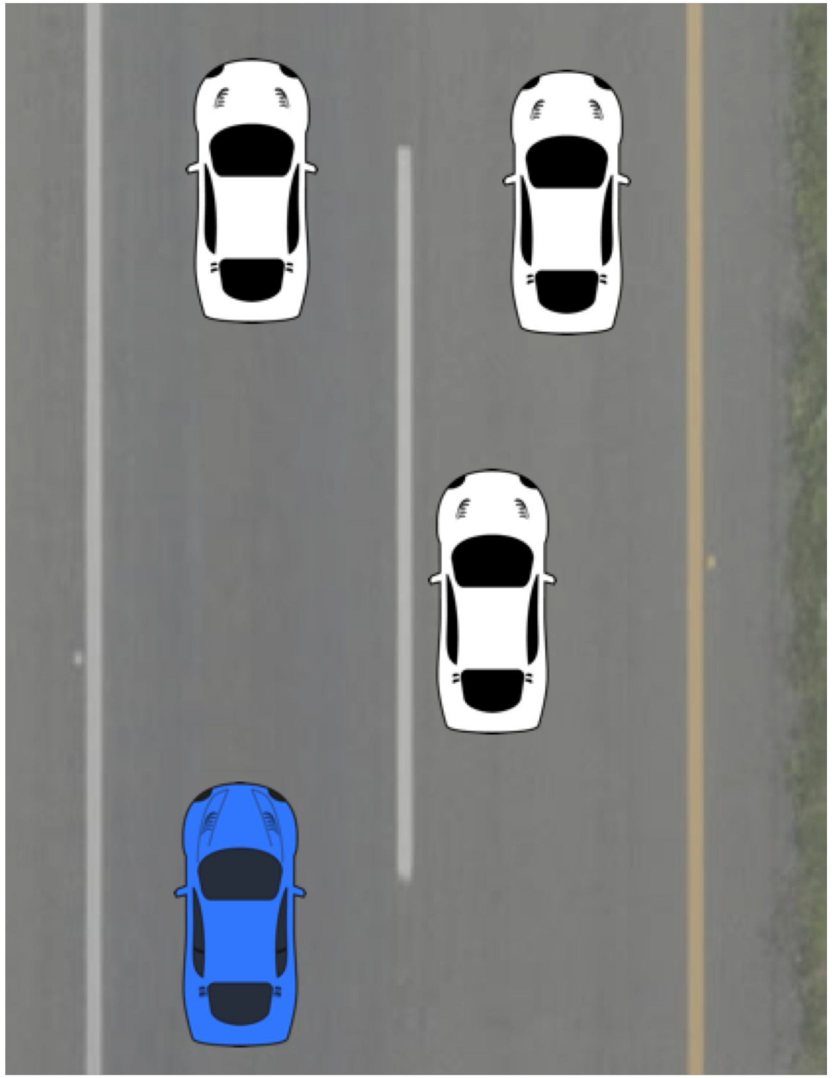

Multi-Agent Trajectory Prediction

Take the following highway scene as an example:

Entities in the scene are represented as **nodes** in a spatiotemporal graph.

Their interactions are encoded as **edges** (formed according to spatial proximity).

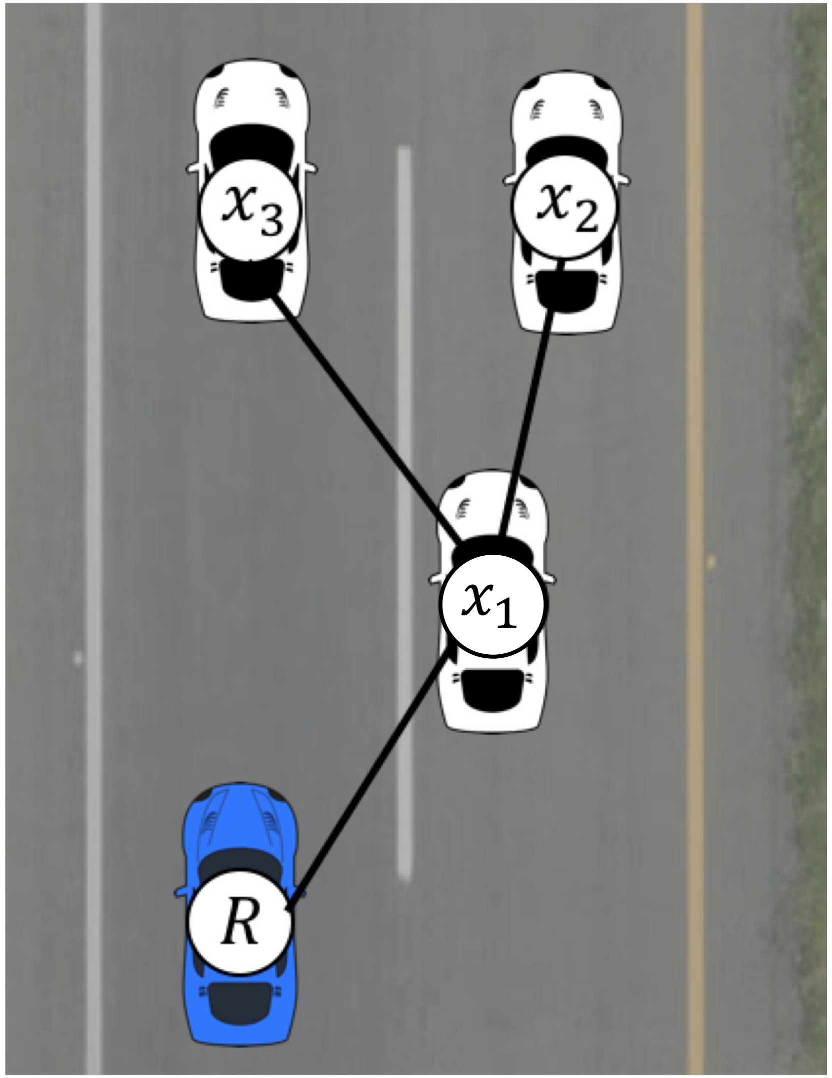

Multi-Agent Trajectory Prediction

Now we can focus on modeling the resulting spatiotemporal graph.

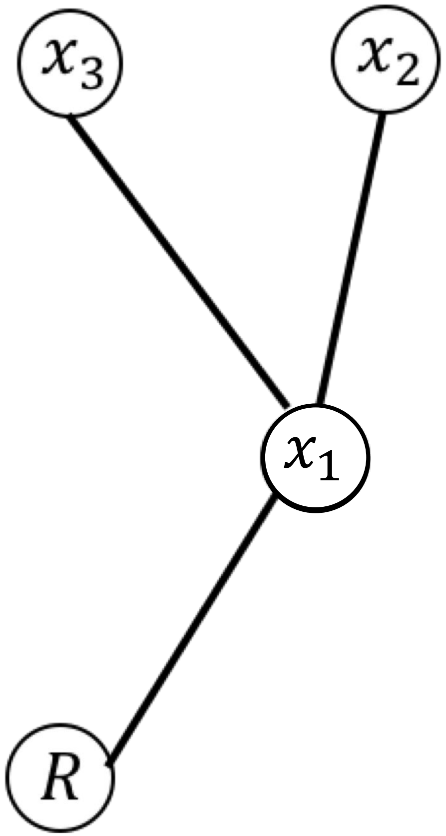

Multi-Agent Trajectory Prediction

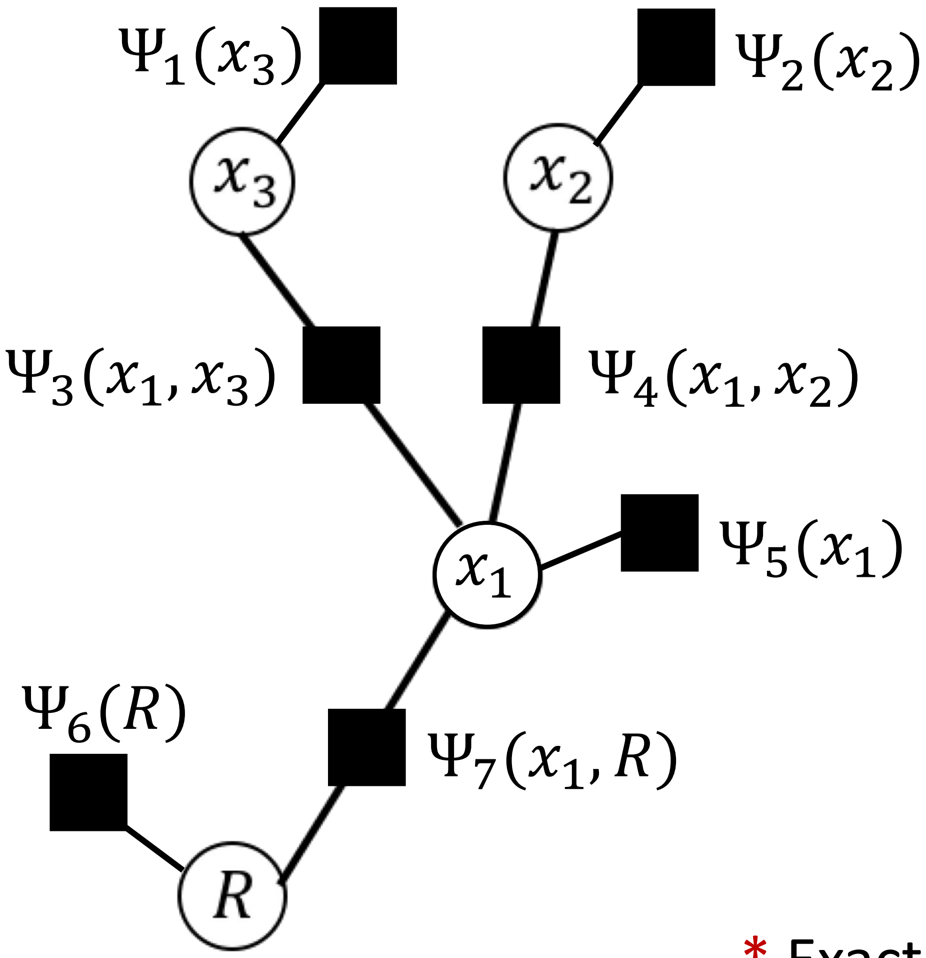

$$P(x_1 | x_3, x_2, R) \propto \prod_{i=1}^{7} \Psi_i(\dots)$$

Our approach is inspired* by Conditional Random Fields (CRFs)

* Exact inference/sampling on CRFs can be very expensive

(Ivanovic, Schmerling, et al. 2018)

Multi-Human Modeling with Spatiotemporal Graphs

Conditional Variational Autoencoder (CVAE)

- We introduce a discrete latent variable **z**

$$p(\mathbf{y}|\mathbf{x}) = \sum_{\mathbf{z}} p(\mathbf{y}|\mathbf{x}, \mathbf{z}) p(\mathbf{z}|\mathbf{x})$$

Multi-Human Modeling with Spatiotemporal Graphs

Conditional Variational Autoencoder (CVAE)

- We introduce a discrete latent variable **z**

$$p(\mathbf{y}|\mathbf{x}) = \sum_{\mathbf{z}} p(\mathbf{y}|\mathbf{x},\mathbf{z})p(\mathbf{z}|\mathbf{x})$$

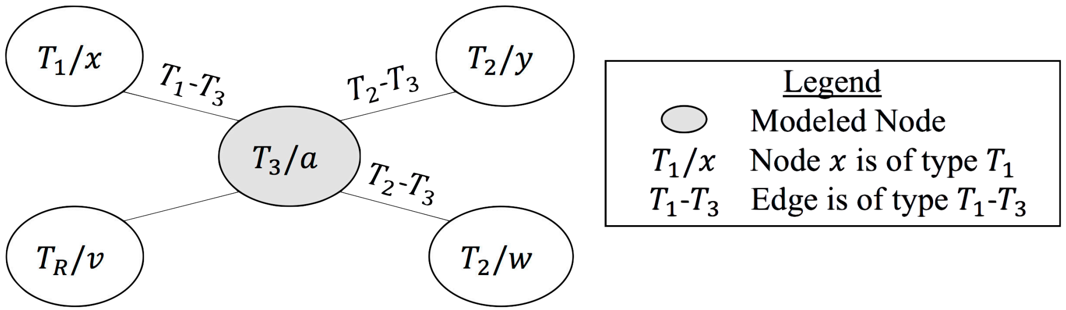

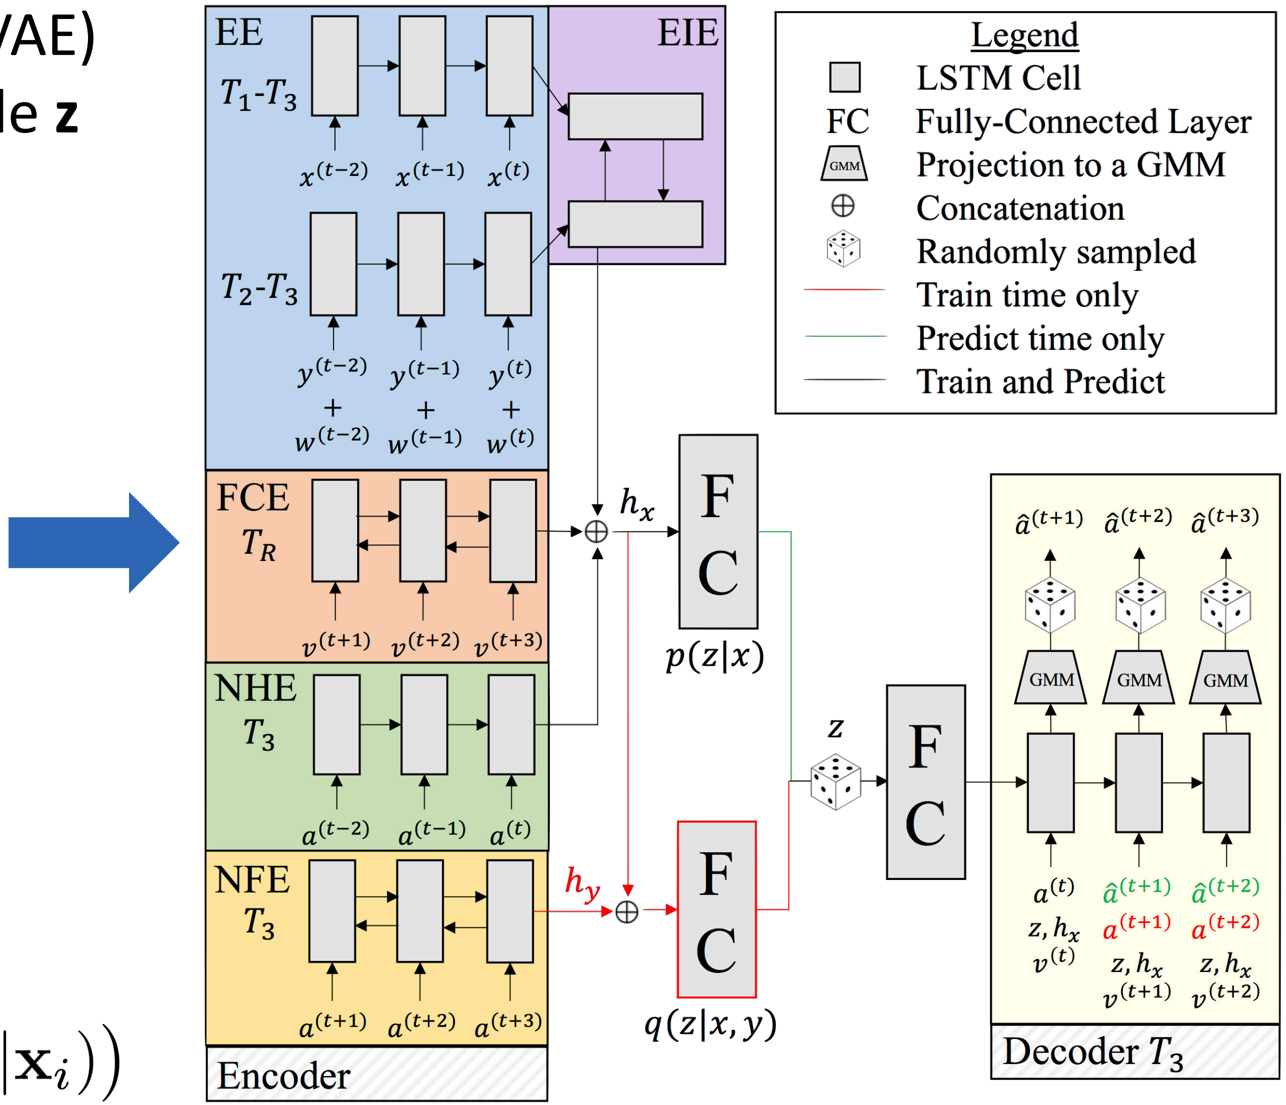

$$\max_{\phi,\theta,\psi} \sum_{i=1}^{N} \mathbb{E}_{z\sim q_\phi(z|\mathbf{x}_i,\mathbf{y}_i)}\left[p_\psi(\mathbf{y}_i|\mathbf{x}_i, z)\right] - D_{KL}\left(q_\phi(z|\mathbf{x}_i,\mathbf{y}_i)\parallel p_\theta(z|\mathbf{x}_i)\right)$$

Multi-Human Modeling with Spatiotemporal Graphs

Conditional Variational Autoencoder (CVAE)

- We introduce a discrete latent variable **z**

$$p(\mathbf{y}|\mathbf{x}) = \sum_{\mathbf{z}} p(\mathbf{y}|\mathbf{x},\mathbf{z})p(\mathbf{z}|\mathbf{x})$$

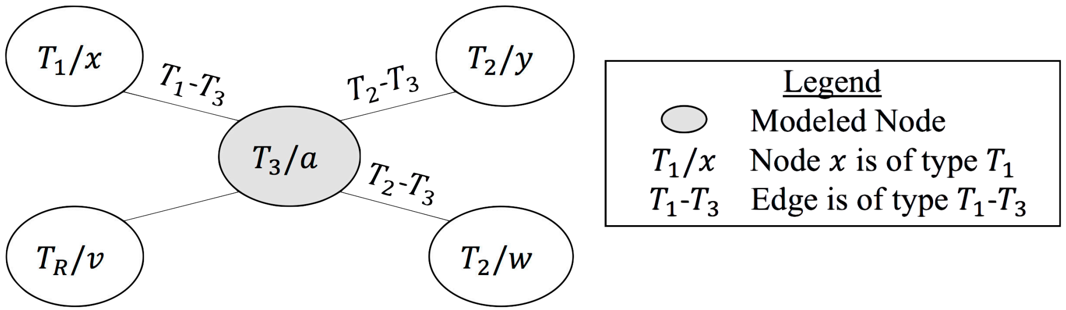

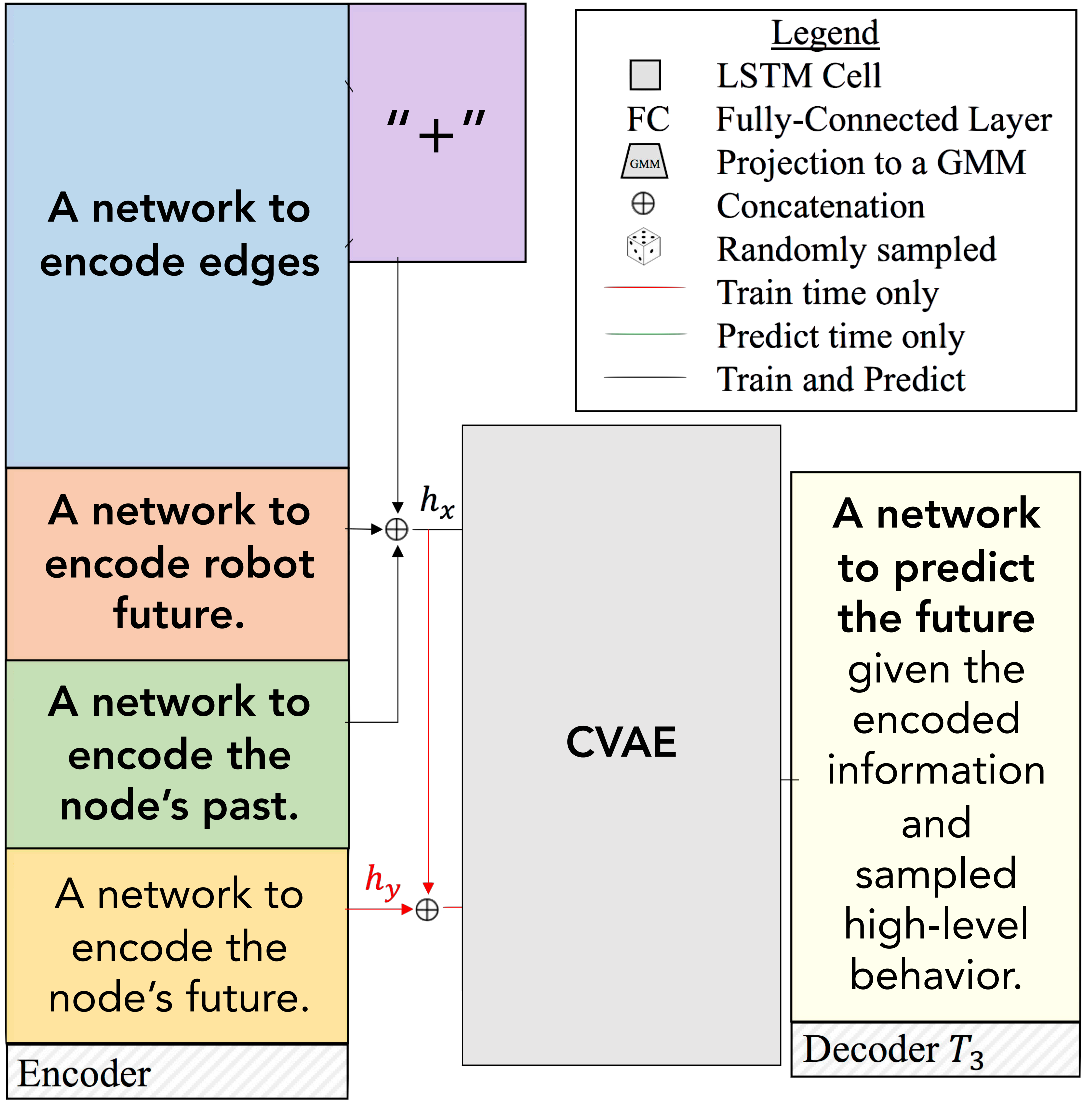

$$\max_{\phi,\theta,\psi} \sum_{i=1}^{N} \mathbb{E}_{z\sim q_\phi(z|\mathbf{x}_i,\mathbf{y}_i)}\left[p_\psi(\mathbf{y}_i|\mathbf{x}_i,z)\right] - D_{KL}\left(q_\phi(z|\mathbf{x}_i,\mathbf{y}_i)\parallel p_\theta(z|\mathbf{x}_i)\right)$$

Multi-Human Modeling with Spatiotemporal Graphs

What if the scene structure changes?

ETH Pedestrians Dataset

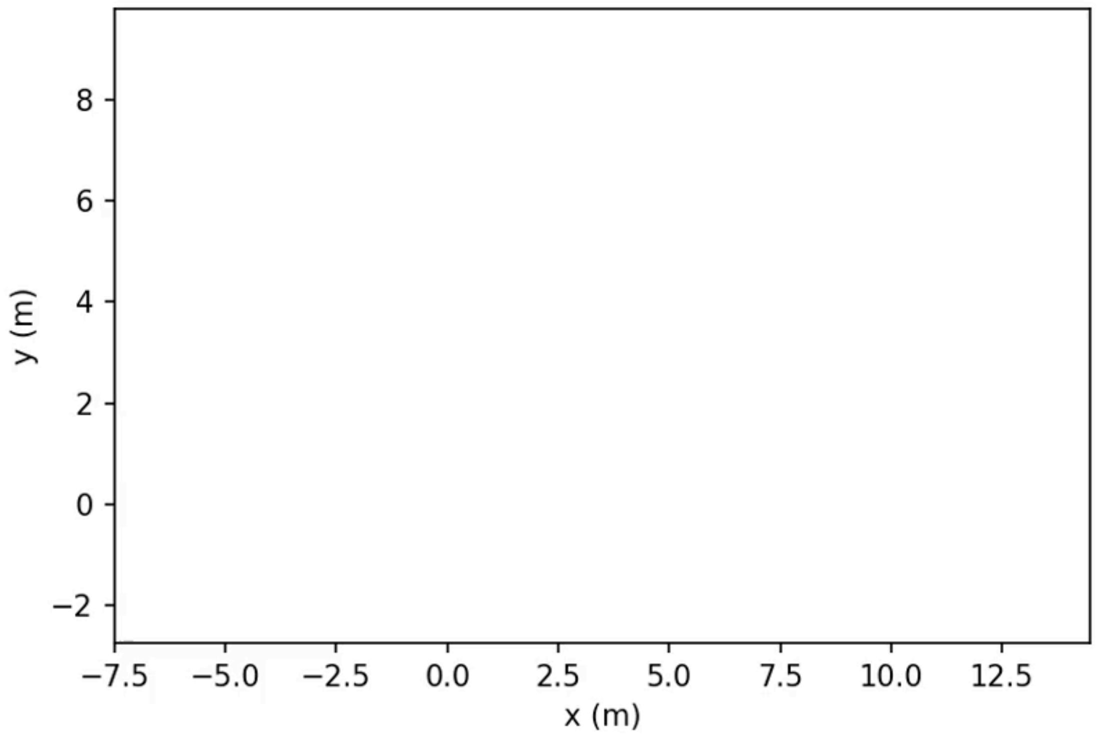

NBA Basketball Players Dataset

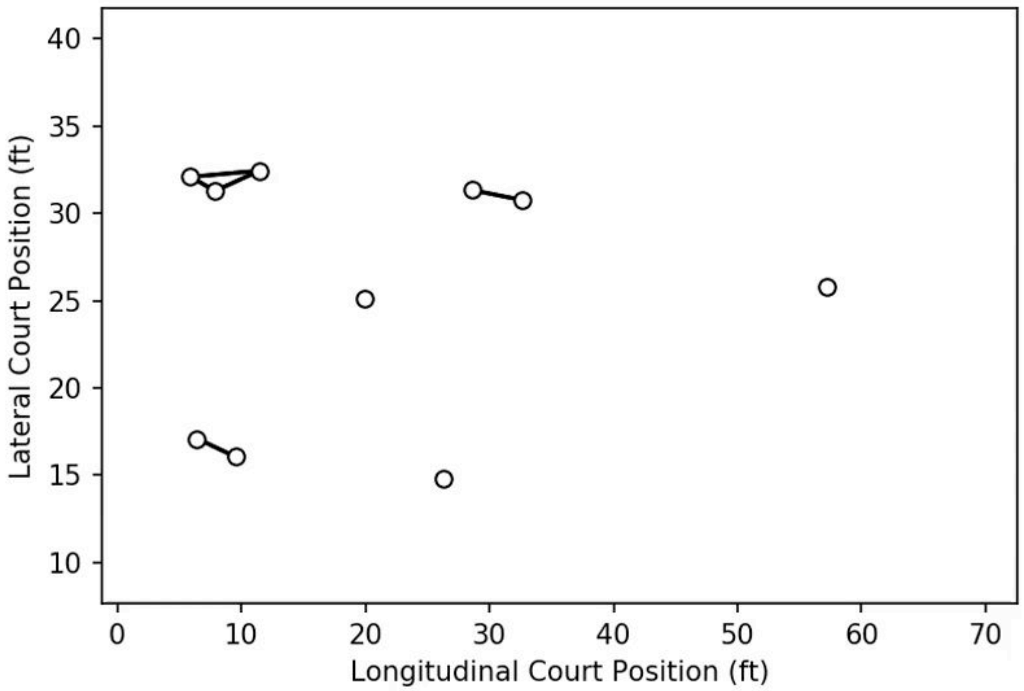

Dynamic Spatiotemporal Graphs

(Ivanovic and Pavone, under review)

How do you model edges/nodes being added/removed from your graph?

1. Adding/removing edge/node models is critical for computational efficiency
2. We therefore modulate edge/node strengths to smoothly "interpolate" between graph changes

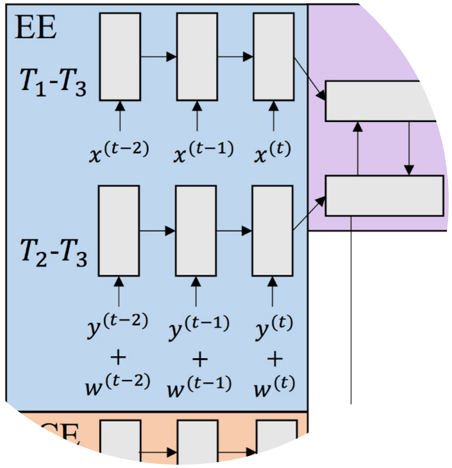

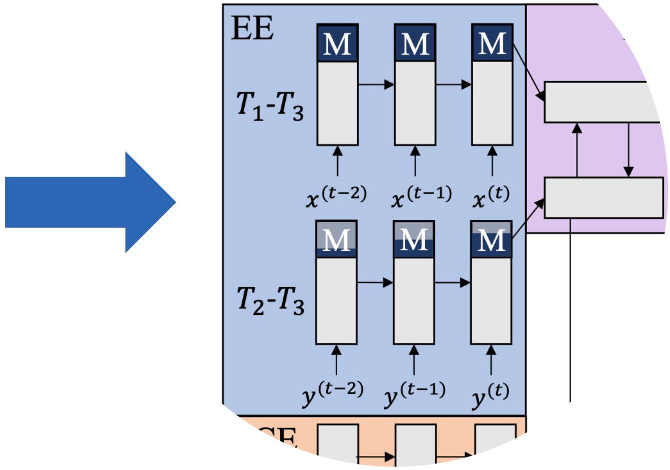

Extra scalar multiplication and learned attention layer

(Ivanovic and Pavone, under review)

The "Trajectron"

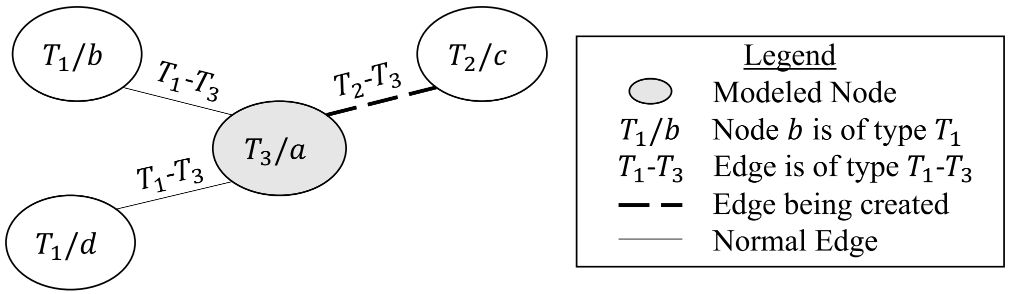

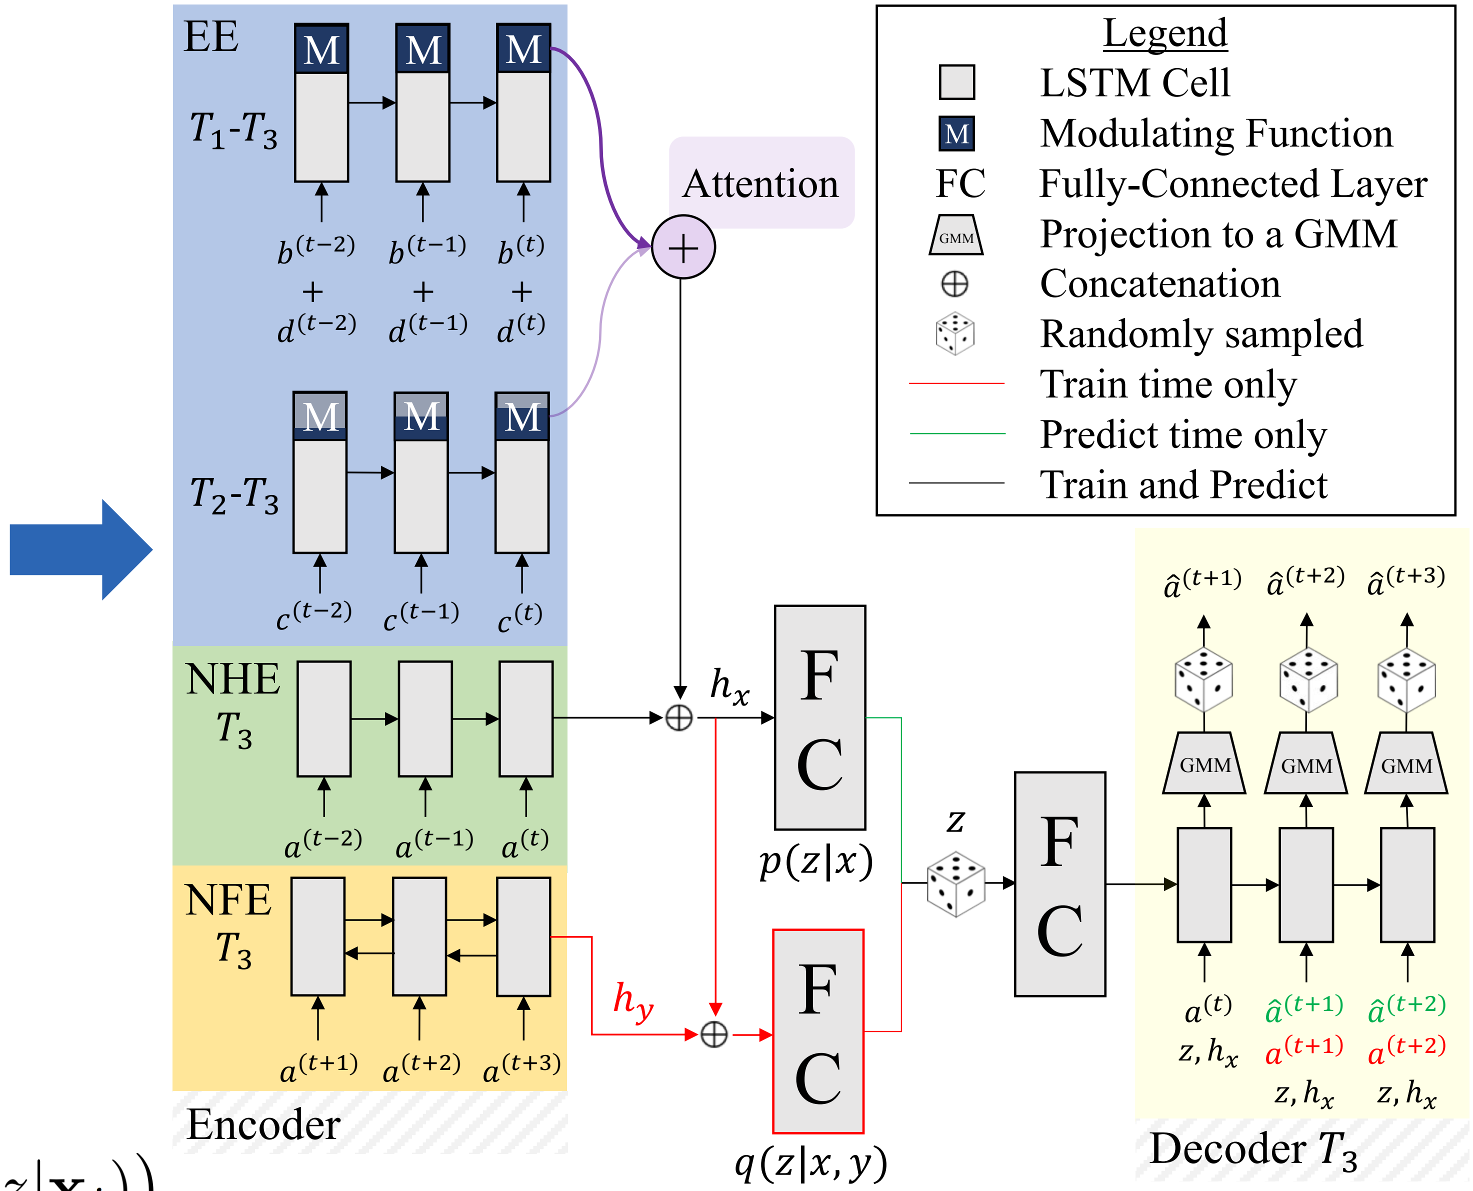

$$\max_{\phi,\theta,\psi} \sum_{i=1}^{N} \mathbb{E}_{z\sim q_\phi(z|\mathbf{x}_i,\mathbf{y}_i)}\left[p_\psi(\mathbf{y}_i|\mathbf{x}_i, z)\right] - D_{KL}\left(q_\phi(z|\mathbf{x}_i,\mathbf{y}_i)\,\|\,p_\theta(z|\mathbf{x}_i)\right)$$

(Ivanovic and Pavone, under review)

The “Trajectron”

Edges are now **fully dynamic**.

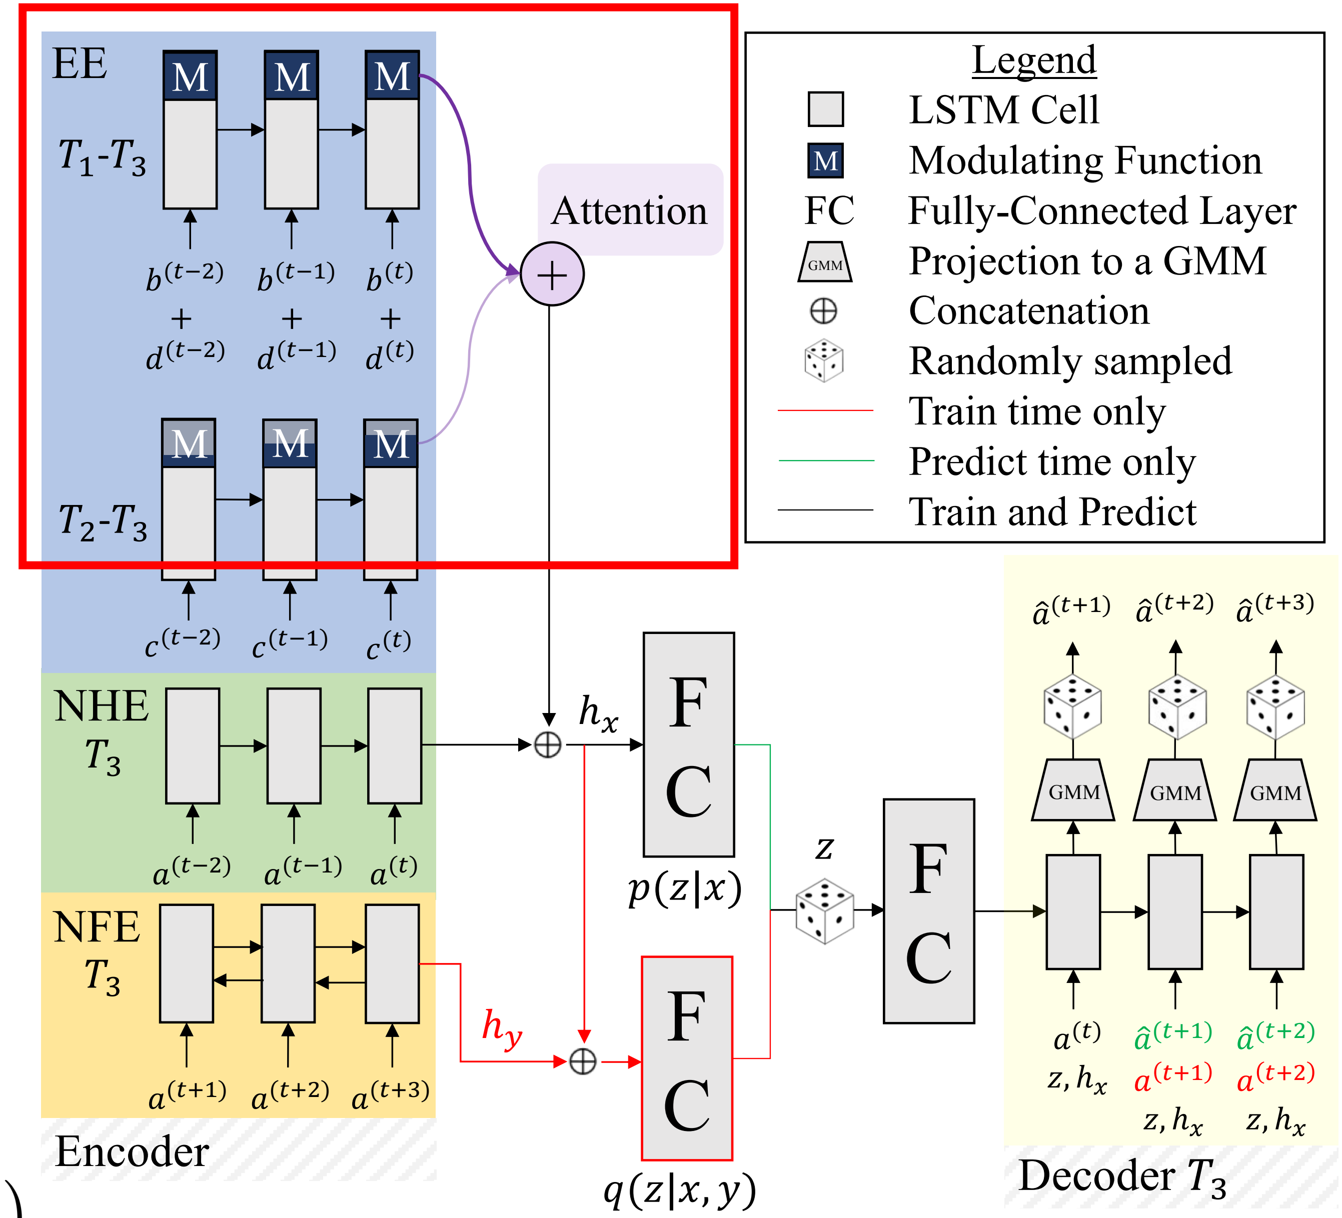

$$\max_{\phi,\theta,\psi} \sum_{i=1}^{N} \mathbb{E}_{z\sim q_\phi(z|\mathbf{x}_i,\mathbf{y}_i)}\left[p_\psi(\mathbf{y}_i|\mathbf{x}_i,z)\right] - D_{KL}\left(q_\phi(z|\mathbf{x}_i,\mathbf{y}_i)\parallel p_\theta(z|\mathbf{x}_i)\right)$$

(Ivanovic and Pavone, under review)

Results

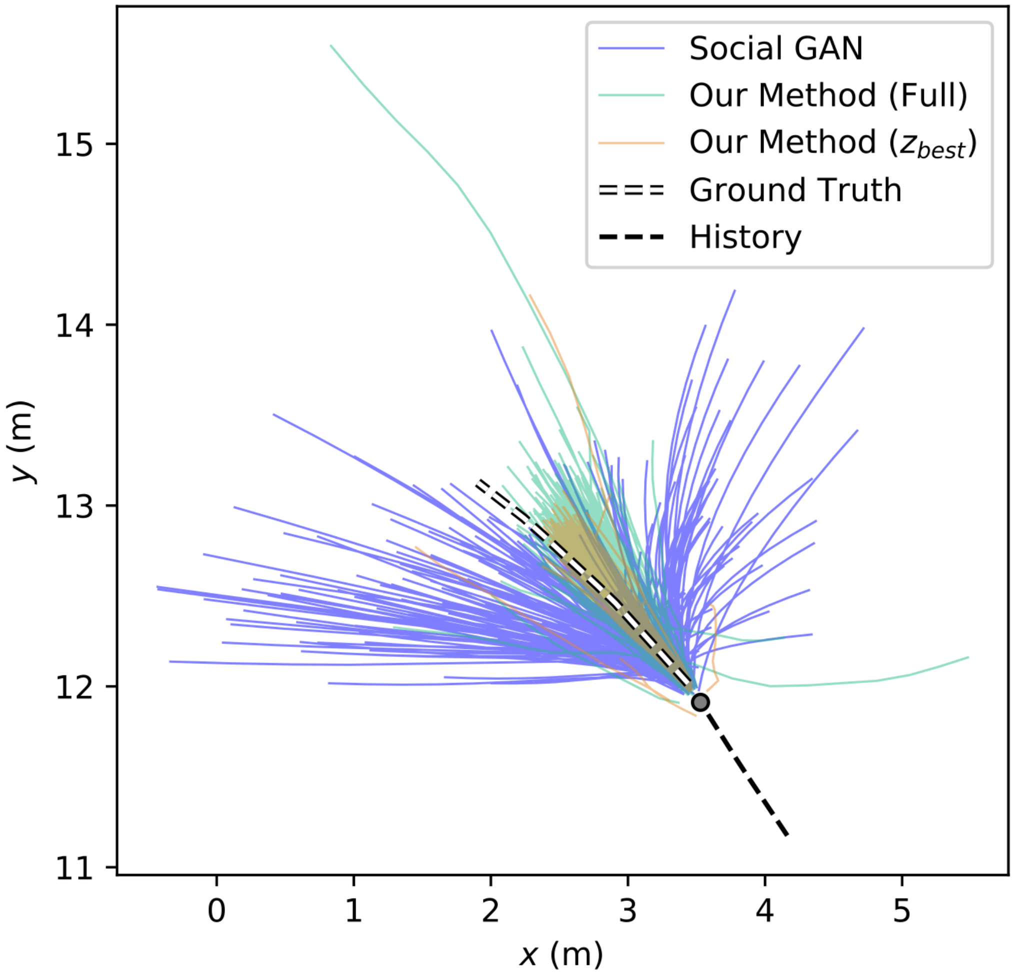

8 steps of history (3.2s), 12 steps of prediction (4.8s)

More concentrated predictions compared to the previous state of the art (Gupta et al. 2018).

Can choose between accuracy (z_{best}) and coverage of potential futures (Full).

$$z_{best} = \arg\max_z p_\theta(z \mid x), y \sim p_\psi(y \mid x, z_{best})$$

$$z \sim p_\theta(z \mid x), y \sim p_\psi(y \mid x, z)$$

(Ivanovic and Pavone, under review)

Results – Negative Log-Likelihood

For comparability we use kernel density estimation (KDE) to construct pdfs from sets of sampled trajectories

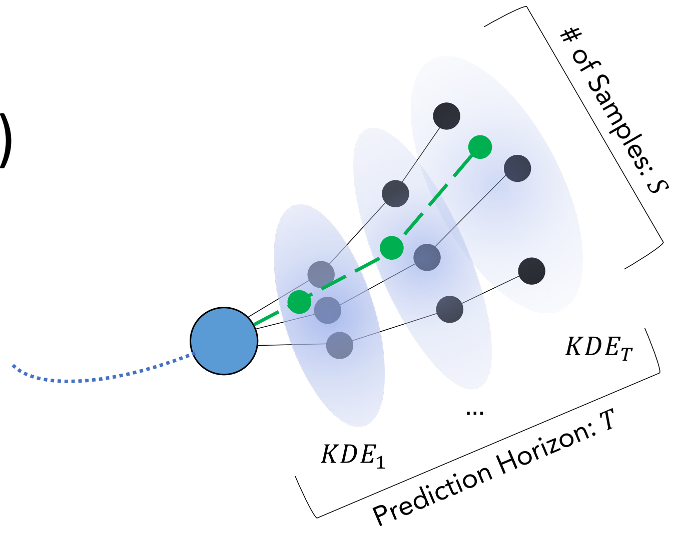

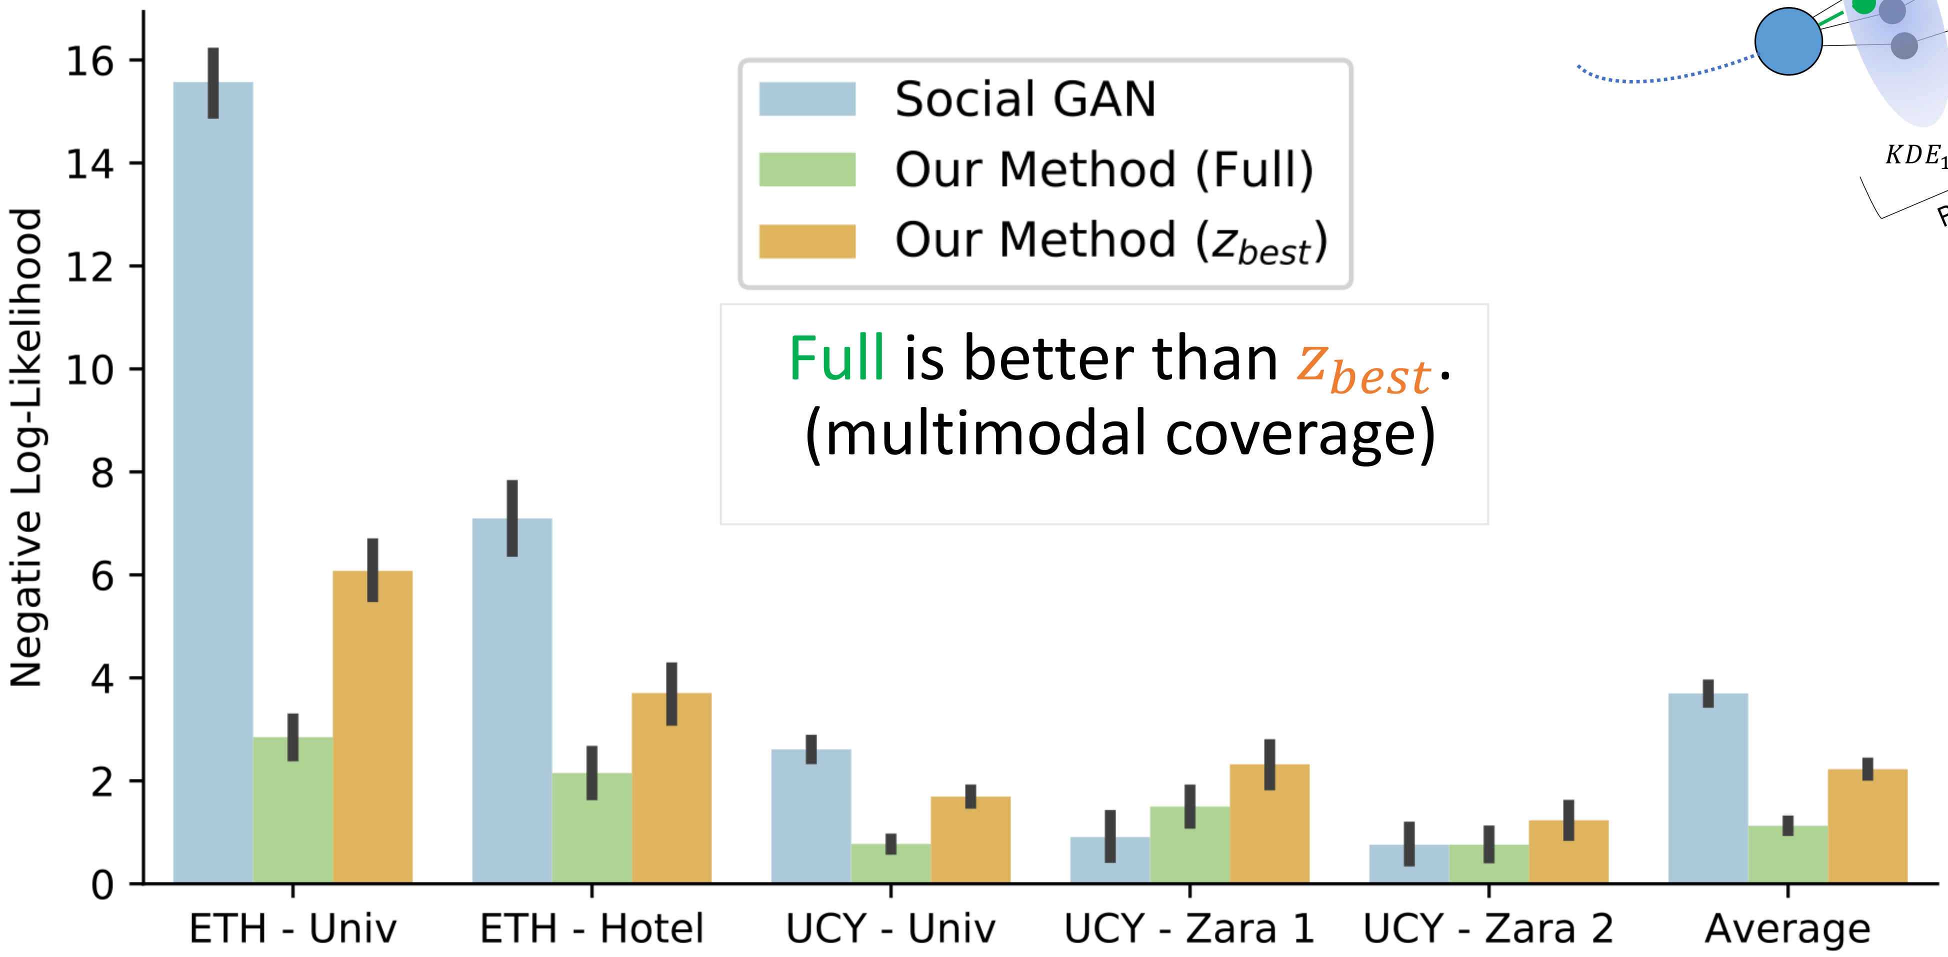

Full is better than z_{best}.
(multimodal coverage)

Results – Average Displacement Error

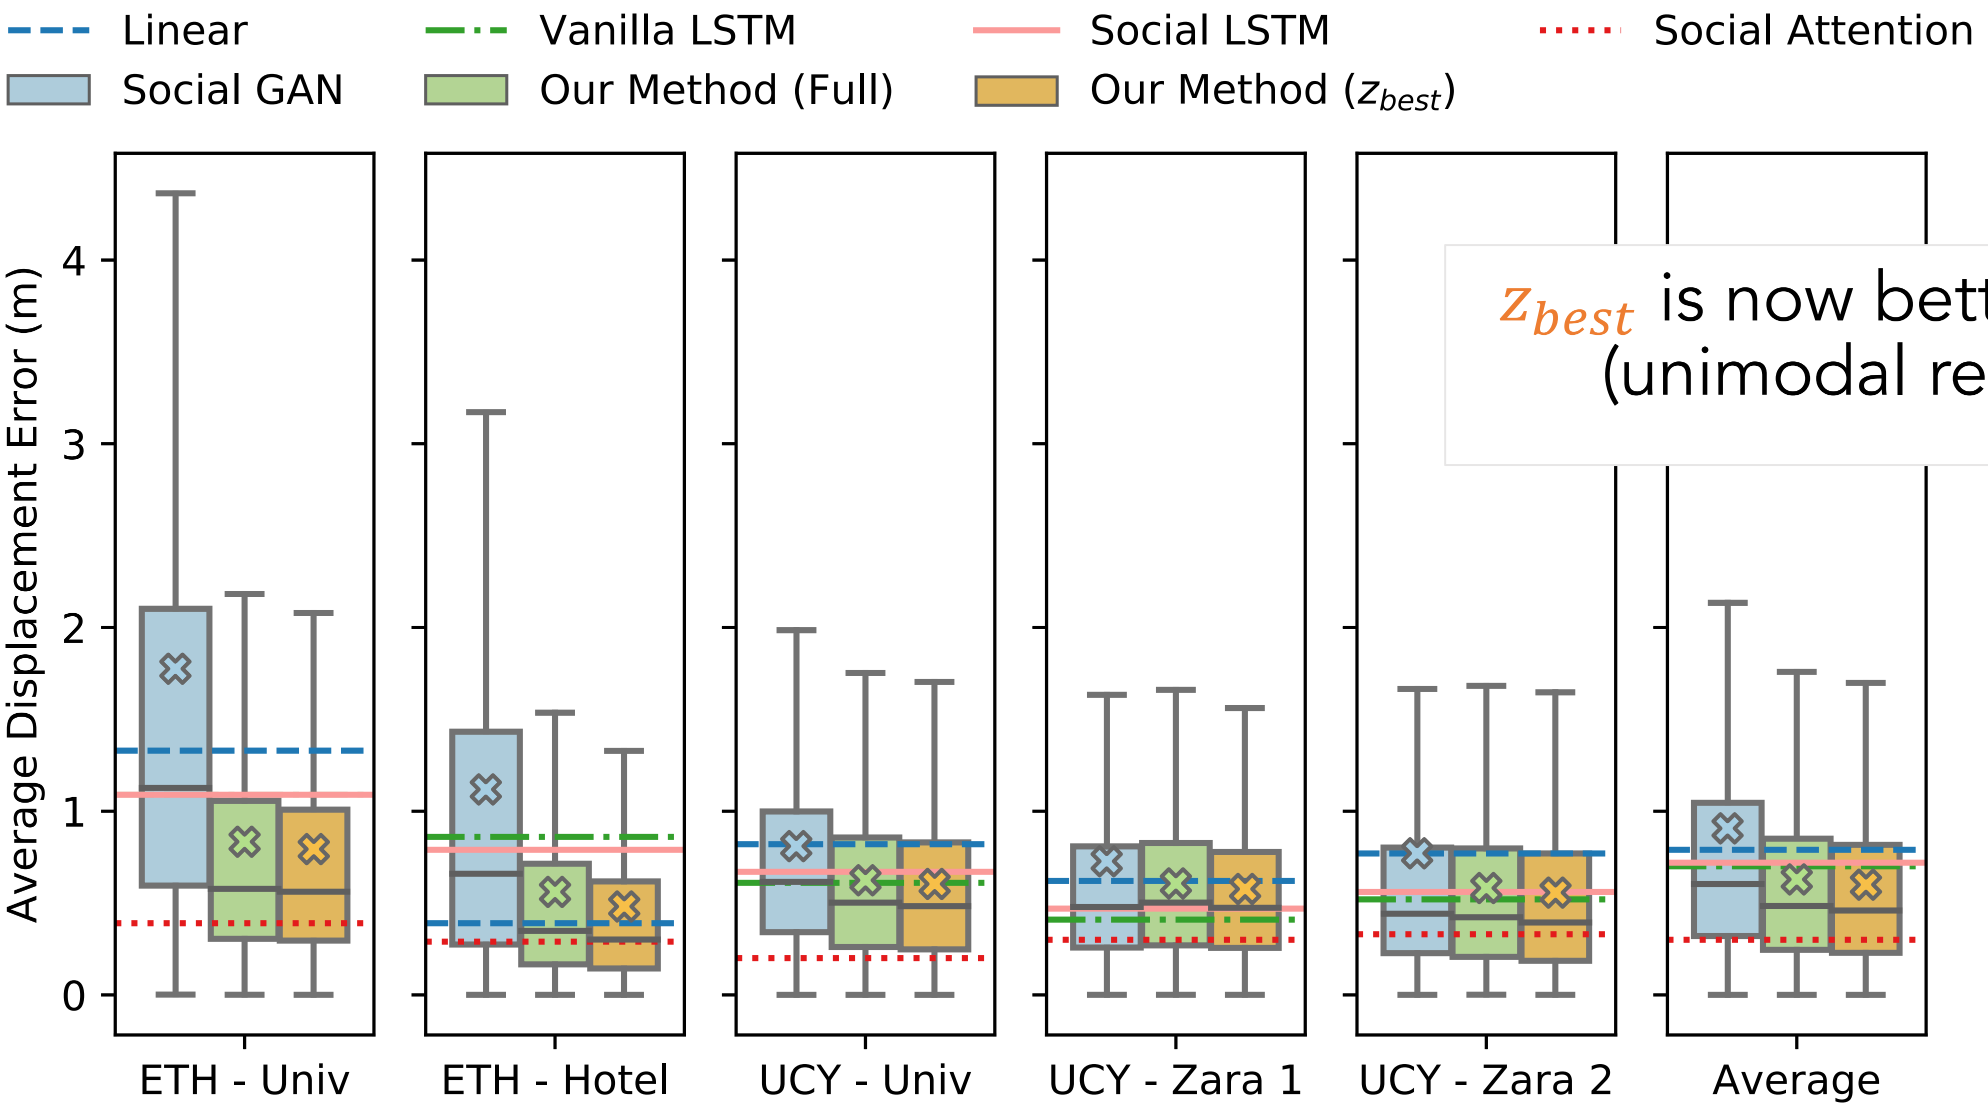

z_{best} is now better than Full.
(unimodal regression)

(Ivanovic and Pavone, under review)

The Trajectron

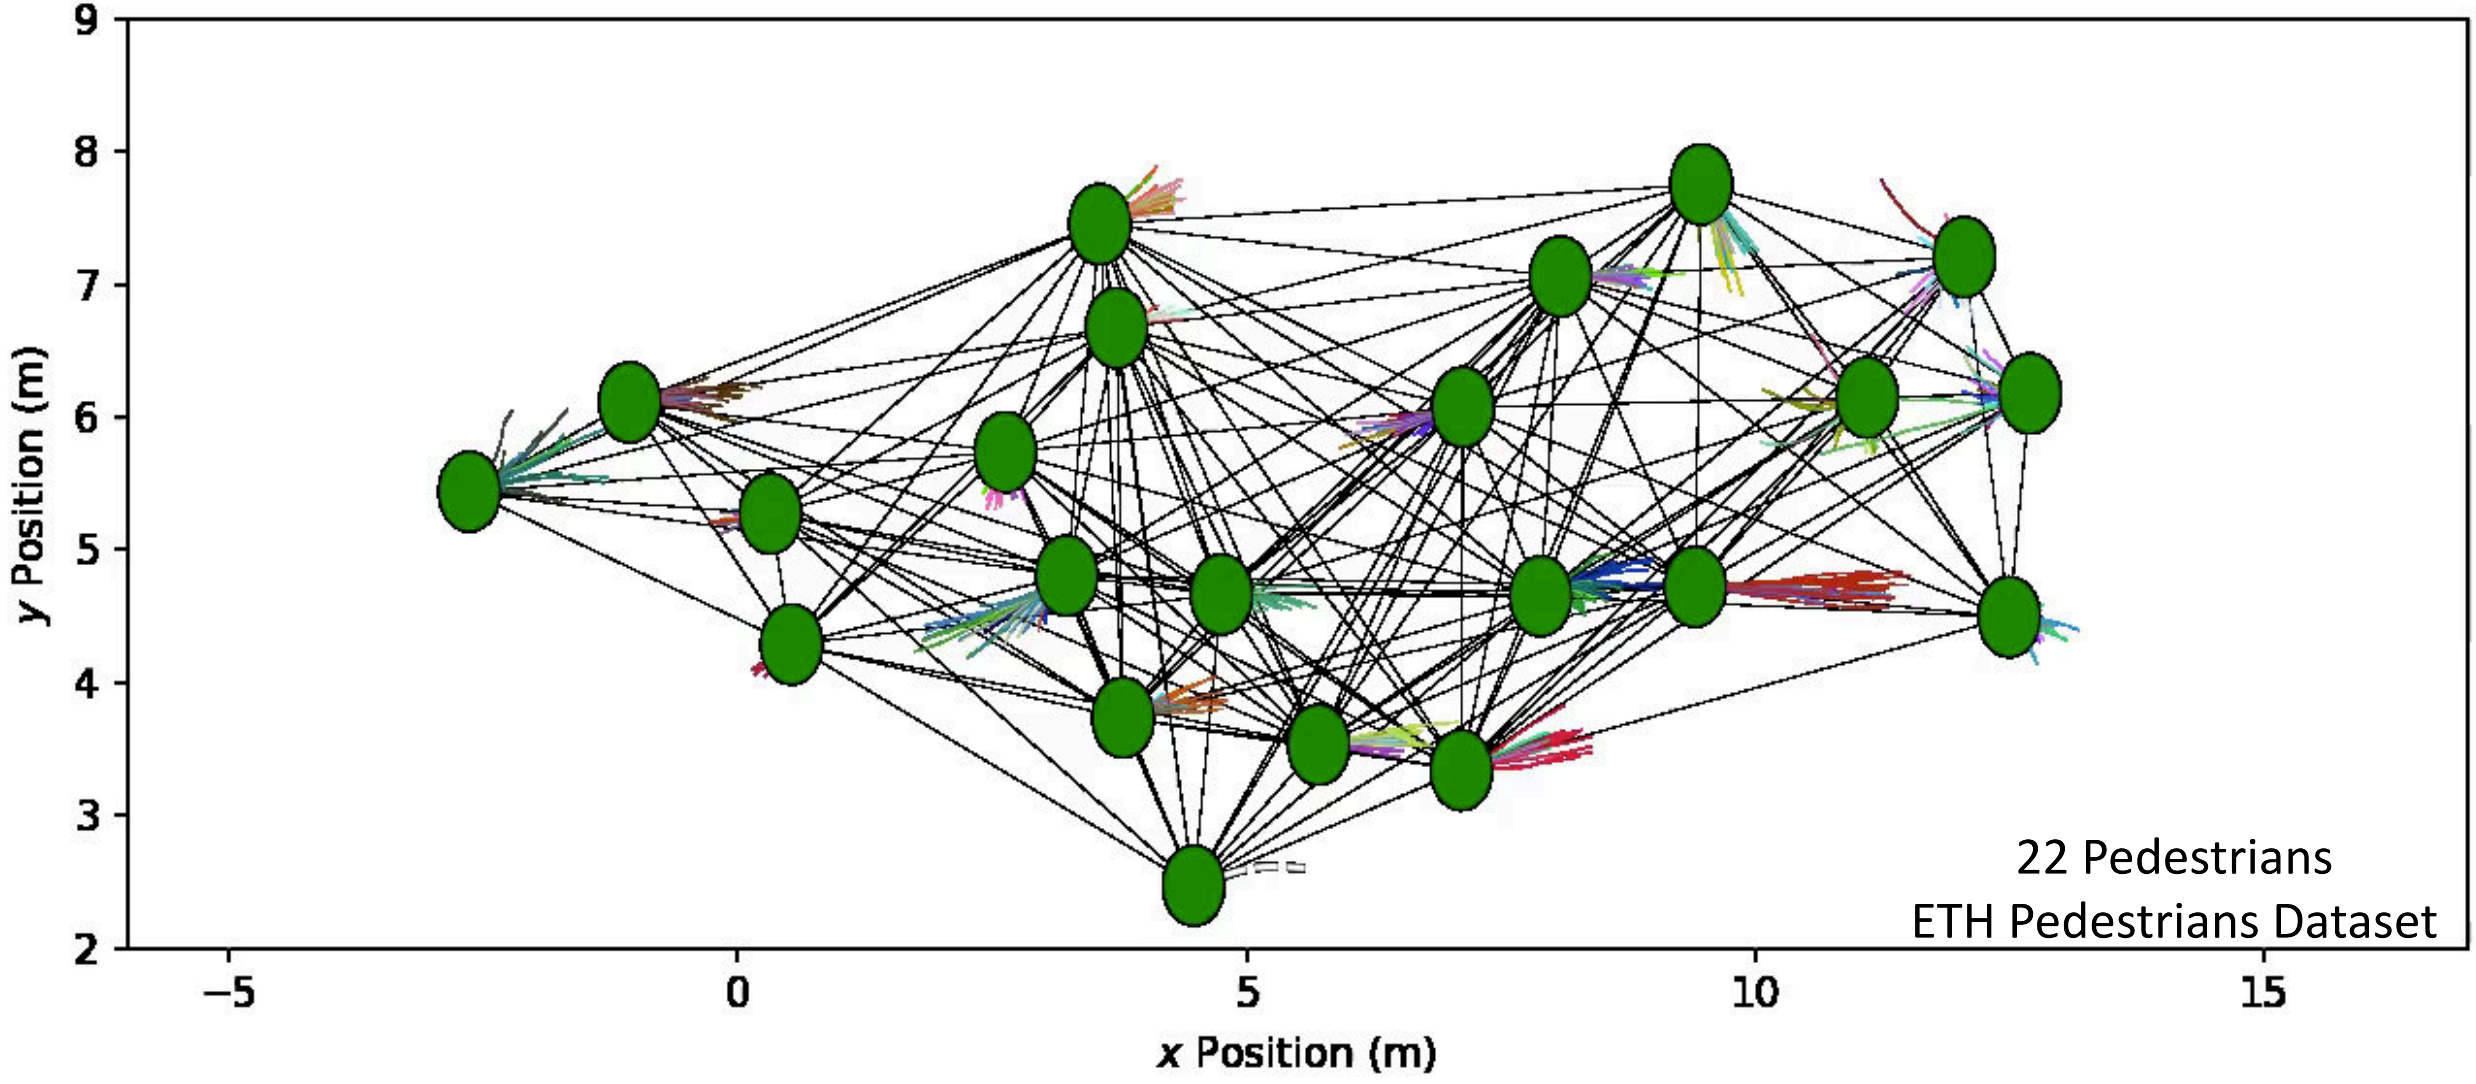

Ongoing/Future Work

- Spatiotemporal graph-based modeling
  - Here, we're only using the graph at $t < T$ to make a prediction about $t > T$

    Modeling future graph evolution → enable truly long term motion prediction
  - Moving beyond purely dynamic-state-based representations
- Model-based planning

(Schmerling, Leung et al. 2018)

Model-Based Planning

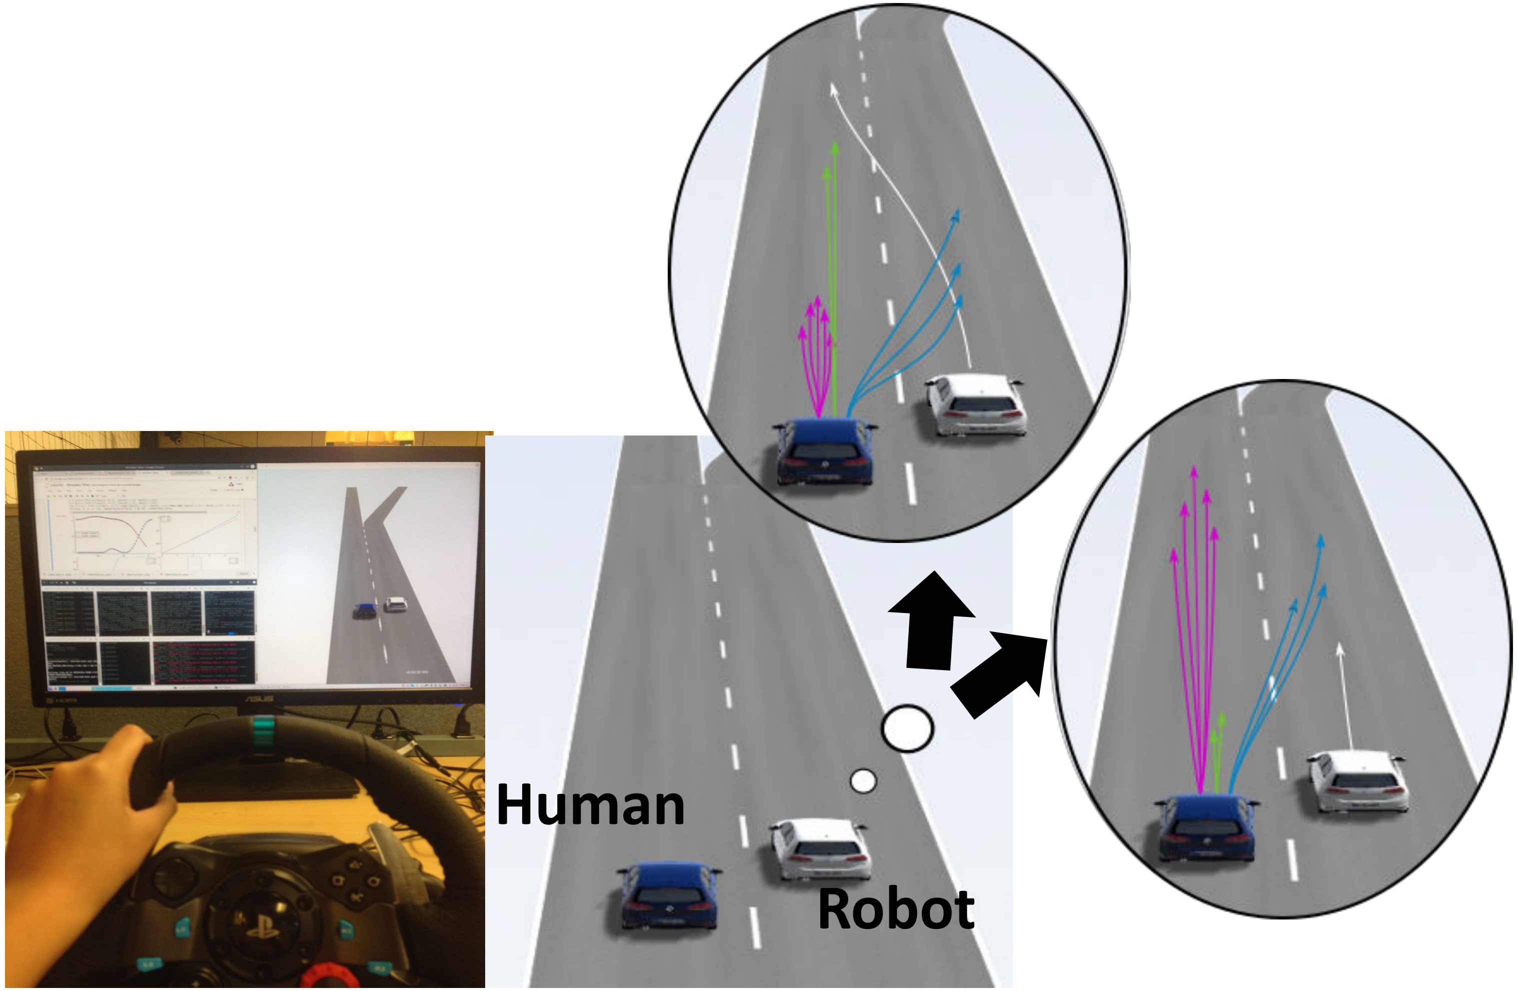

(Schmerling, Leung et al. 2018)

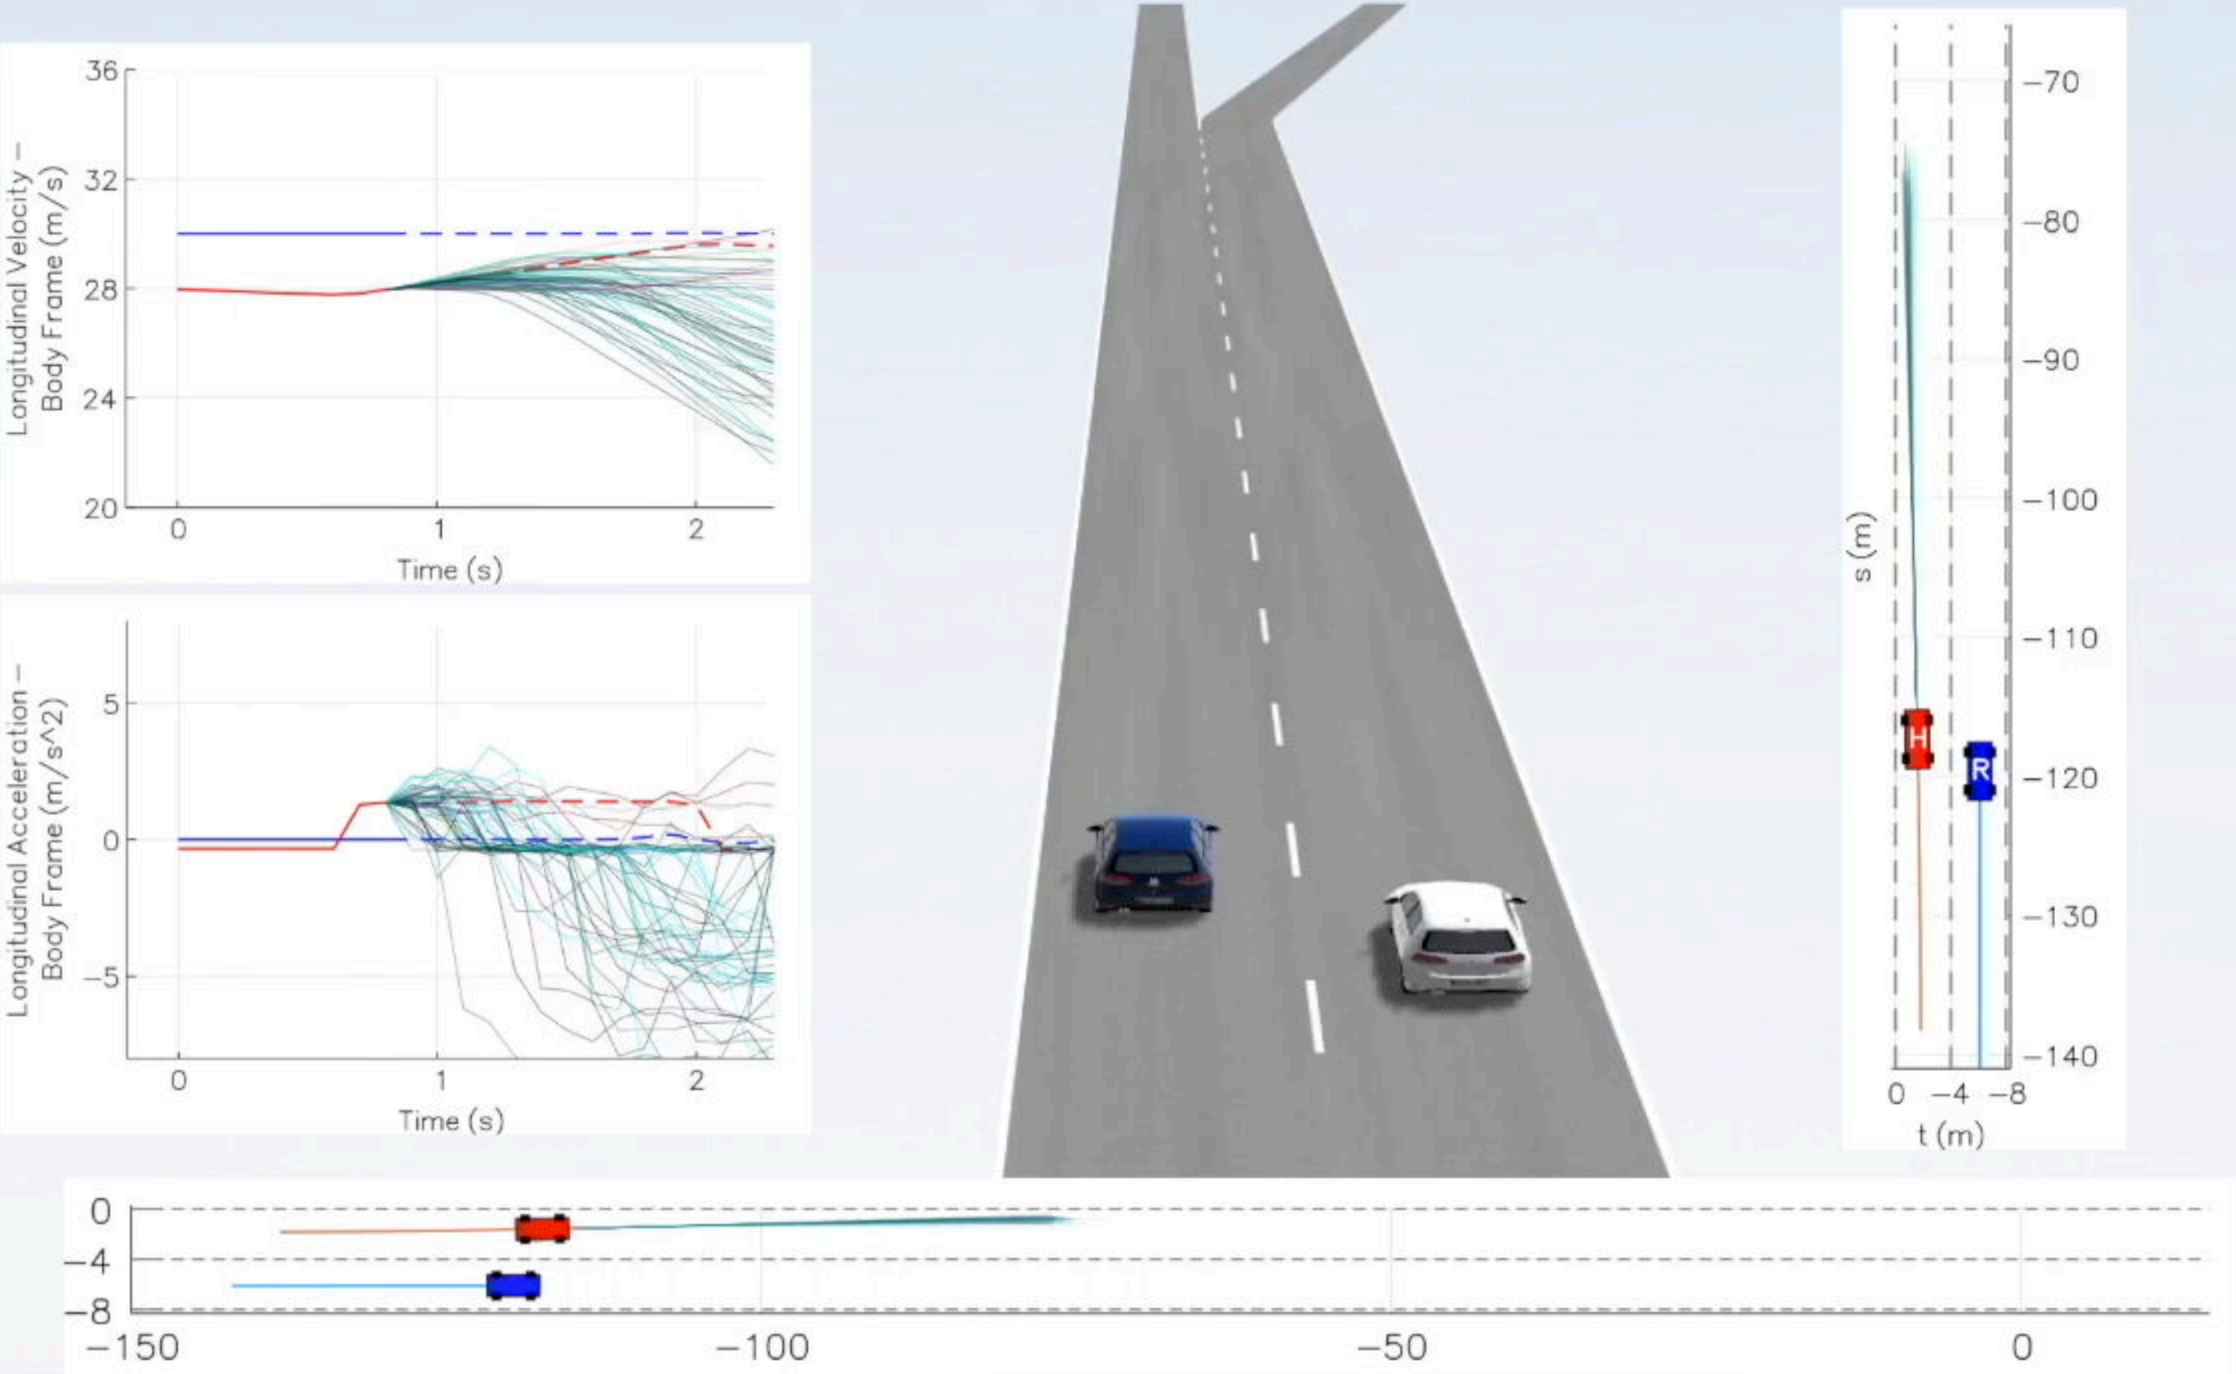


Disaster?!

- Replanning is too slow to ensure safety
- Probabilistic model may “get it wrong”, especially with low-likelihood events

Hamilton-Jacobi Reachability Analysis

(Leung*, Schmerling*, et al. 2018)

"If I want to avoid this set of states in the future, what is the set I should avoid now?"

Backward Reachable Set

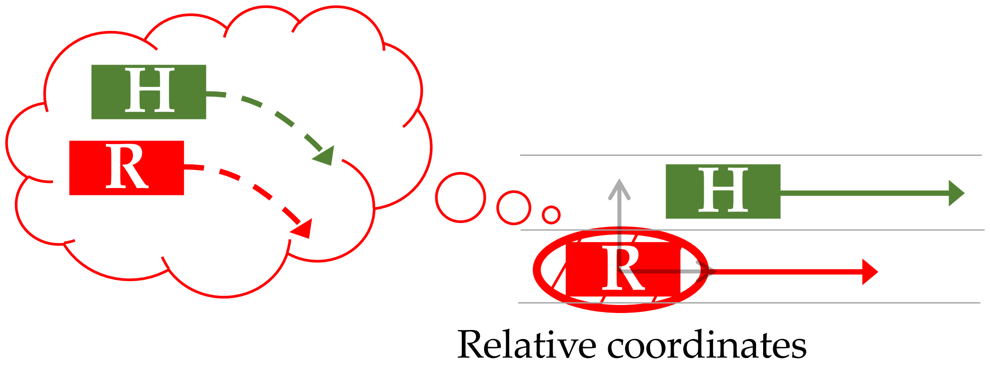

(Leung*, Schmerling*, et al. 2018)

Hamilton-Jacobi Reachability Analysis

"If I want to avoid this set of states in the future, what is the set I should avoid now?”

Backward Reachable Set

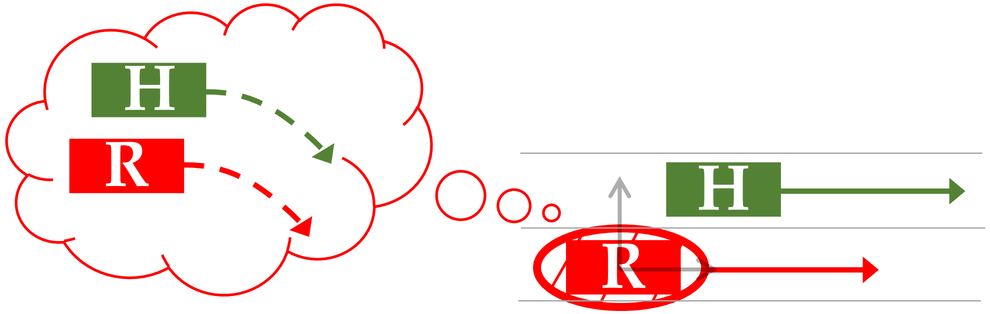

Relative coordinates

Forward Reachable Set

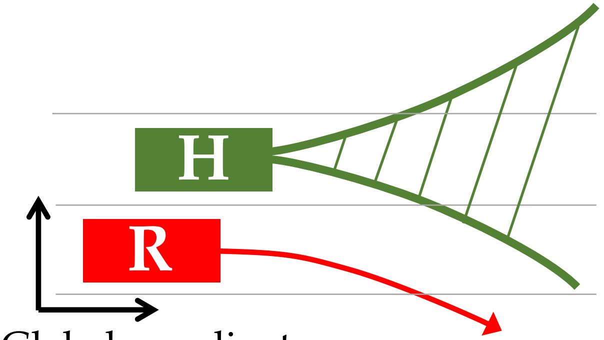

Global coordinates

(Leung*, Schmerling*, et al. 2018)

Hamilton-Jacobi Reachability Analysis

"If I want to avoid this set of states in the future, what is the set I should avoid now?"

Backward Reachable Set

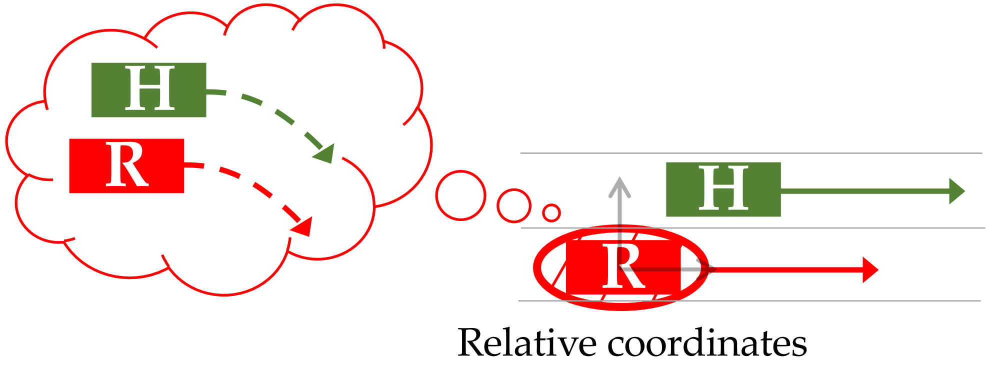

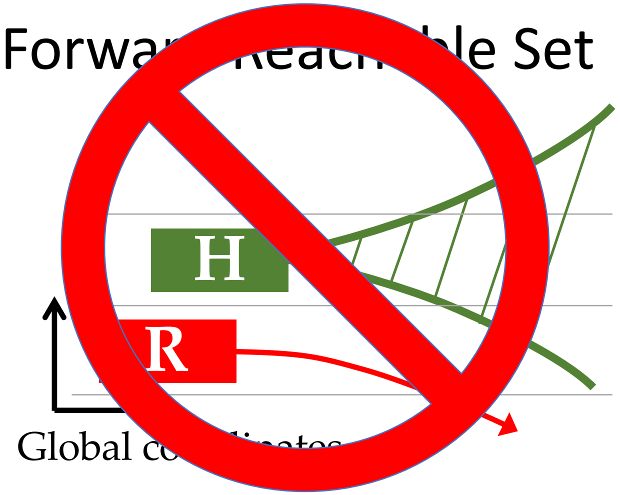

(Leung*, Schmerling*, et al. 2018)

Hamilton-Jacobi Reachability Analysis

"If I want to avoid this set of states in the future, what is the set I should avoid now?"

Backward Reachable Set

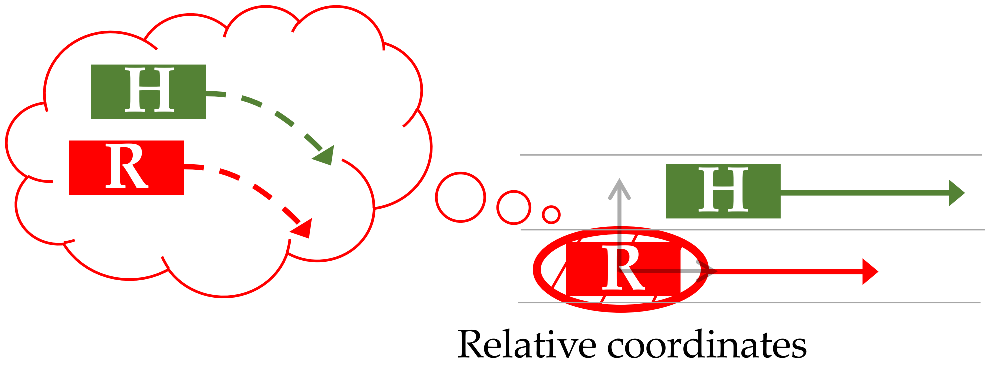

Forward Reachable Set

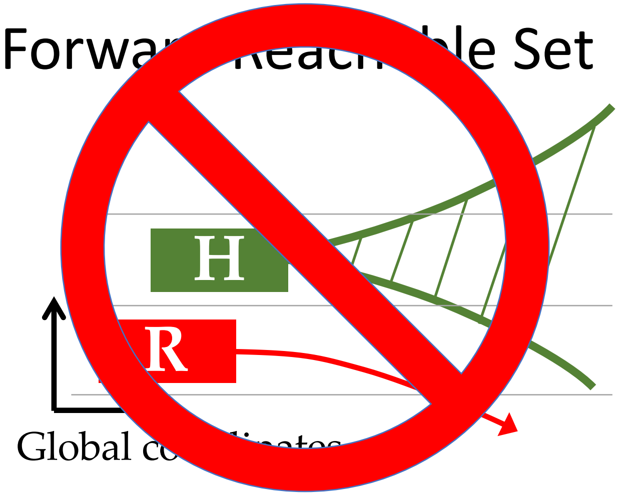

(Leung*, Schmerling*, et al. 2018)

Our Contribution

How to integrate **safety assurance** within a **performance-centric** planning framework?

(Leung*, Schmerling*, et al. 2018)

Our Contribution

How to integrate **safety assurance** within a **performance-centric** planning framework?

Incorporate HJ reachability as a constraint for a low-level MPC tracking controller

(Leung*, Schmerling*, et al. 2018)

Our Contribution

How to integrate **safety assurance** within a **performance-centric** planning framework?

Incorporate HJ reachability as a constraint for a low-level MPC tracking controller

Does this **guarantee** safety on the road?

Our Contribution

(Leung*, Schmerling*, et al. 2018)

How to integrate **safety assurance** within a **performance-centric** planning framework?

Incorporate HJ reachability as a constraint for a low-level MPC tracking controller

Does this **guarantee** safety on the road?

HJ analysis is highly dependent on model fidelity but results are still interpretable

(Leung*, Schmerling*, et al. 2018)

Relative Vehicle Dynamics

Robot vehicle dynamics

$$\mathbf{x}_R = \begin{bmatrix} x_R \\ y_R \\ \psi_R \\ U_{x_R} \\ U_{y_R} \\ r_R \end{bmatrix}, \quad \dot{\mathbf{x}}_R = \begin{bmatrix} \dot{x}_R \\ \dot{y}_R \\ \dot{\psi}_R \\ \dot{U}_{x_R} \\ \dot{U}_{y_R} \\ \dot{r}_R \end{bmatrix} = \begin{bmatrix} U_{x_R}\cos\psi_R - U_{y_R}\sin\psi_R \\ U_{x_R}\sin\psi_R + U_{y_R}\cos\psi_R \\ r_R \\ \frac{1}{m}(F_x + F_{x_{drag}}) + r_R U_{y_R} \\ \frac{1}{m}(F_{yf} + F_{yr}) - r_R U_{x_R} \\ \frac{1}{I_{zz}}(aF_{y_f} - bF_{y_r}) \end{bmatrix}, \quad \mathbf{u}_R = \begin{bmatrix} F_x \\ \delta \end{bmatrix}$$

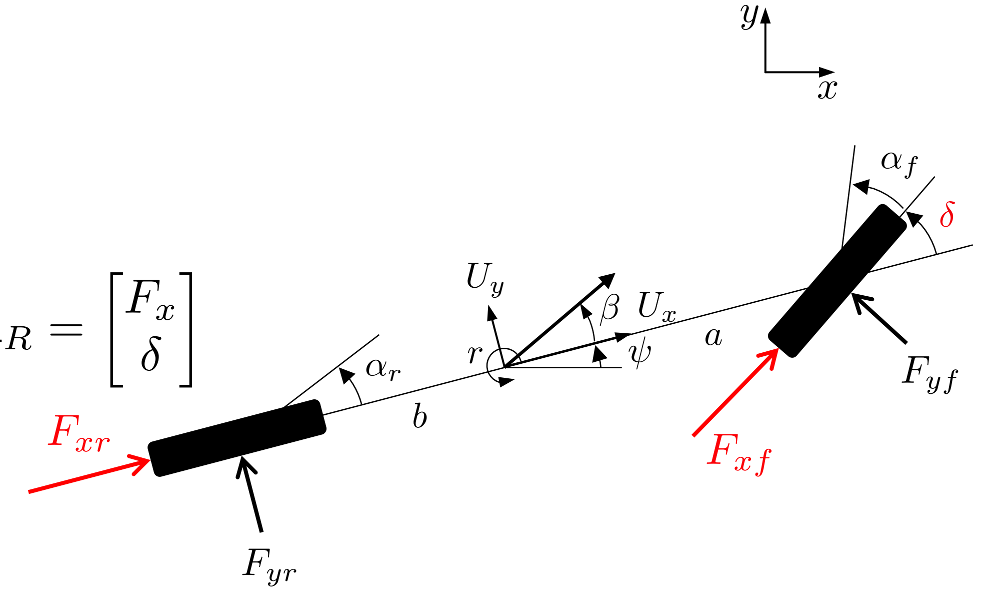

Human vehicle dynamics

$$\mathbf{x}_H = \begin{bmatrix} x_H \\ y_H \\ \psi_H \\ v_H \end{bmatrix}, \quad \dot{\mathbf{x}}_H = \begin{bmatrix} \dot{x}_H \\ \dot{y}_H \\ \dot{\psi}_H \\ \dot{v}_H \end{bmatrix} = \begin{bmatrix} v_H\cos\psi_H \\ v_H\cos\psi_H \\ \omega \\ a \end{bmatrix}, \quad \mathbf{u}_H = \begin{bmatrix} \omega \\ a \end{bmatrix}$$

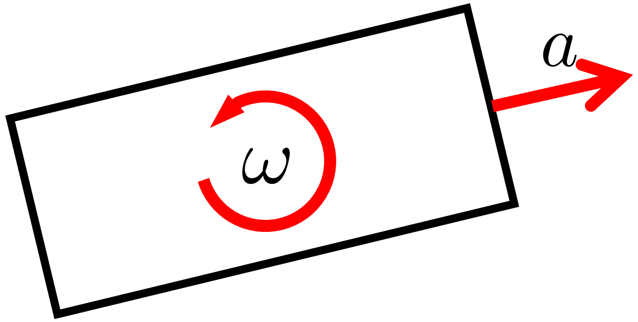

Relative Vehicle Dynamics

(Leung*, Schmerling*, et al. 2018)

$$\begin{bmatrix} x_{rel} \\ y_{rel} \end{bmatrix} = R_{-\psi_R} \begin{bmatrix} x_H - x_R \\ y_H - y_R \end{bmatrix}$$

$$\psi_{rel} = \psi_H - \psi_R$$

$$\begin{bmatrix} \dot{x}_{rel} \\ \dot{y}_{rel} \\ \dot{\psi}_{rel} \\ \dot{U}_{x_R} \\ \dot{U}_{y_R} \\ \dot{v}_H \\ \dot{r}_R \end{bmatrix} = \begin{bmatrix} v_H \cos\psi_{rel} - U_{x_R} + y_{rel} r_R \\ v_H \sin\psi_{rel} - U_{y_R} - x_{rel} r_R \\ \omega - r_R \\ \frac{1}{m}(F_x + F_{x_{drag}}) - r_R U_{x_R} \\ \frac{1}{m}(F_{y_f} + F_{y_r}) - r_R U_{x_R} \\ a \\ \frac{1}{I_{zz}}(aF_{y_f} - bF_{y_r}) \end{bmatrix}$$

Relative vehicle dynamics
(equal control authority)

(Leung*, Schmerling*, et al. 2018)

Value Function

- Terminal set is the zero sub-level set of a **value function**
- Value function varies as relative states changes

(Leung*, Schmerling*, et al. 2018)

Value Function

- Terminal set is the zero sub-level set of a **value function**
- Value function varies as relative states changes

$$\mathcal{A} = \{x \,|\, V(x) \leq 0\}$$

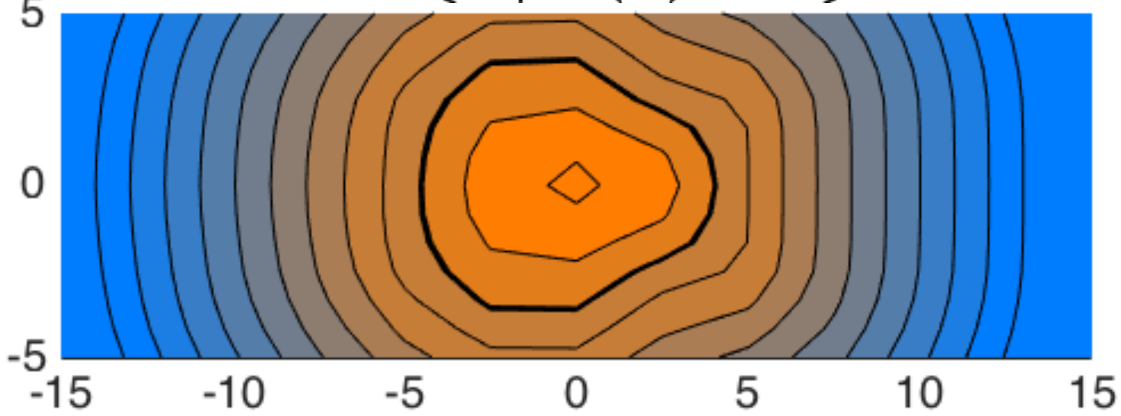

Value function contour plot

$$u_R^* = \arg\max_{u_R} \min_{u_H} \nabla V(x_{\mathscr{R}}) \cdot f(x_{\mathscr{R}}, u_R, u_H)$$

Value Function Contour Plots

(Leung*, Schmerling*, et al. 2018)

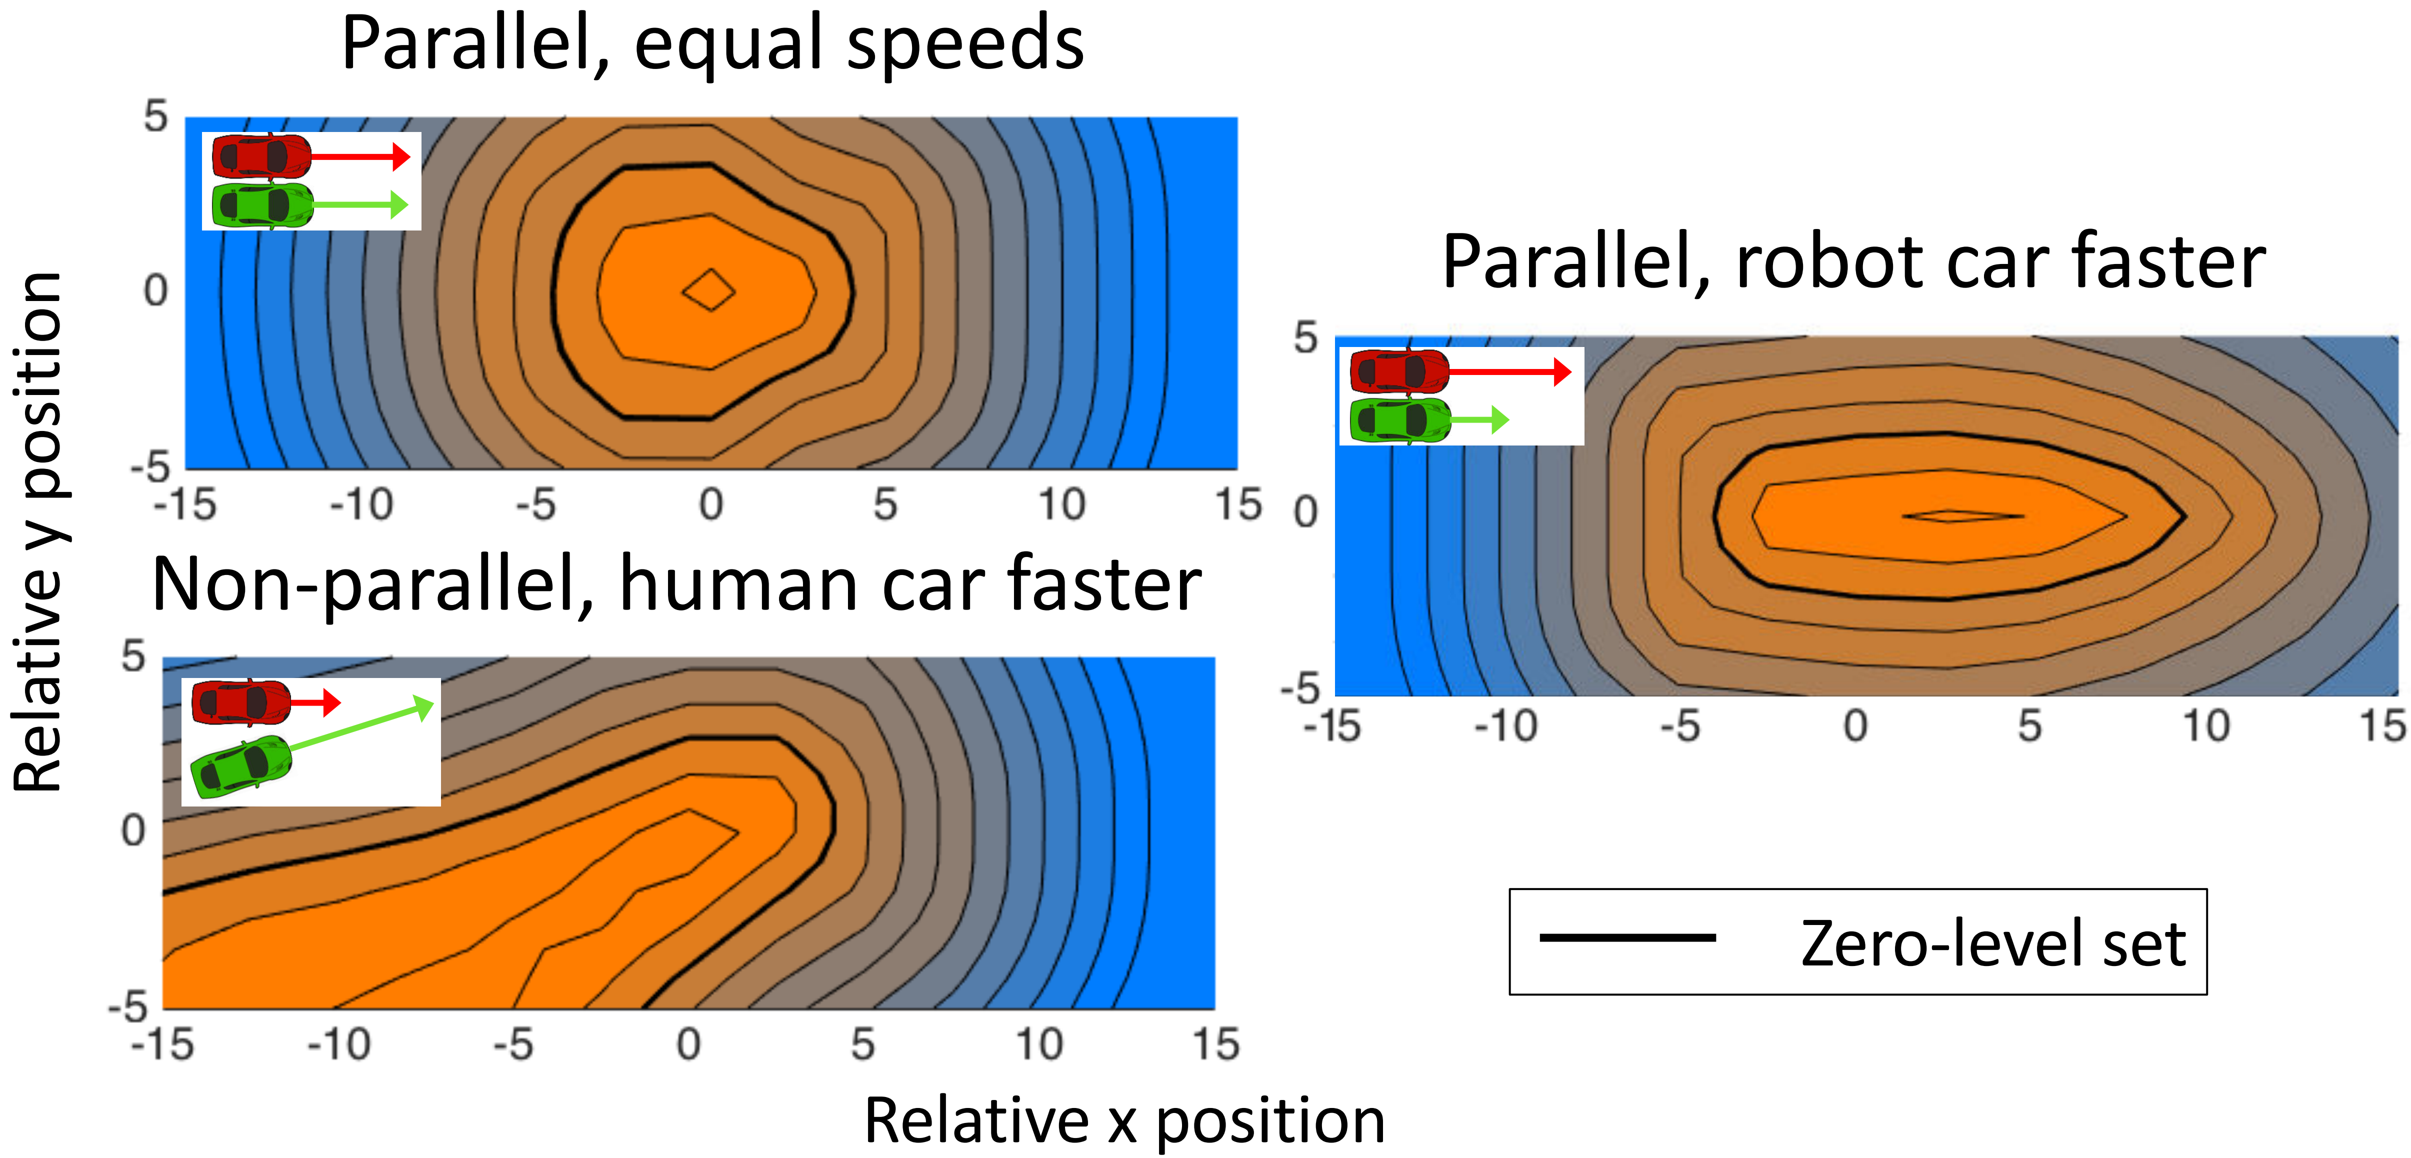

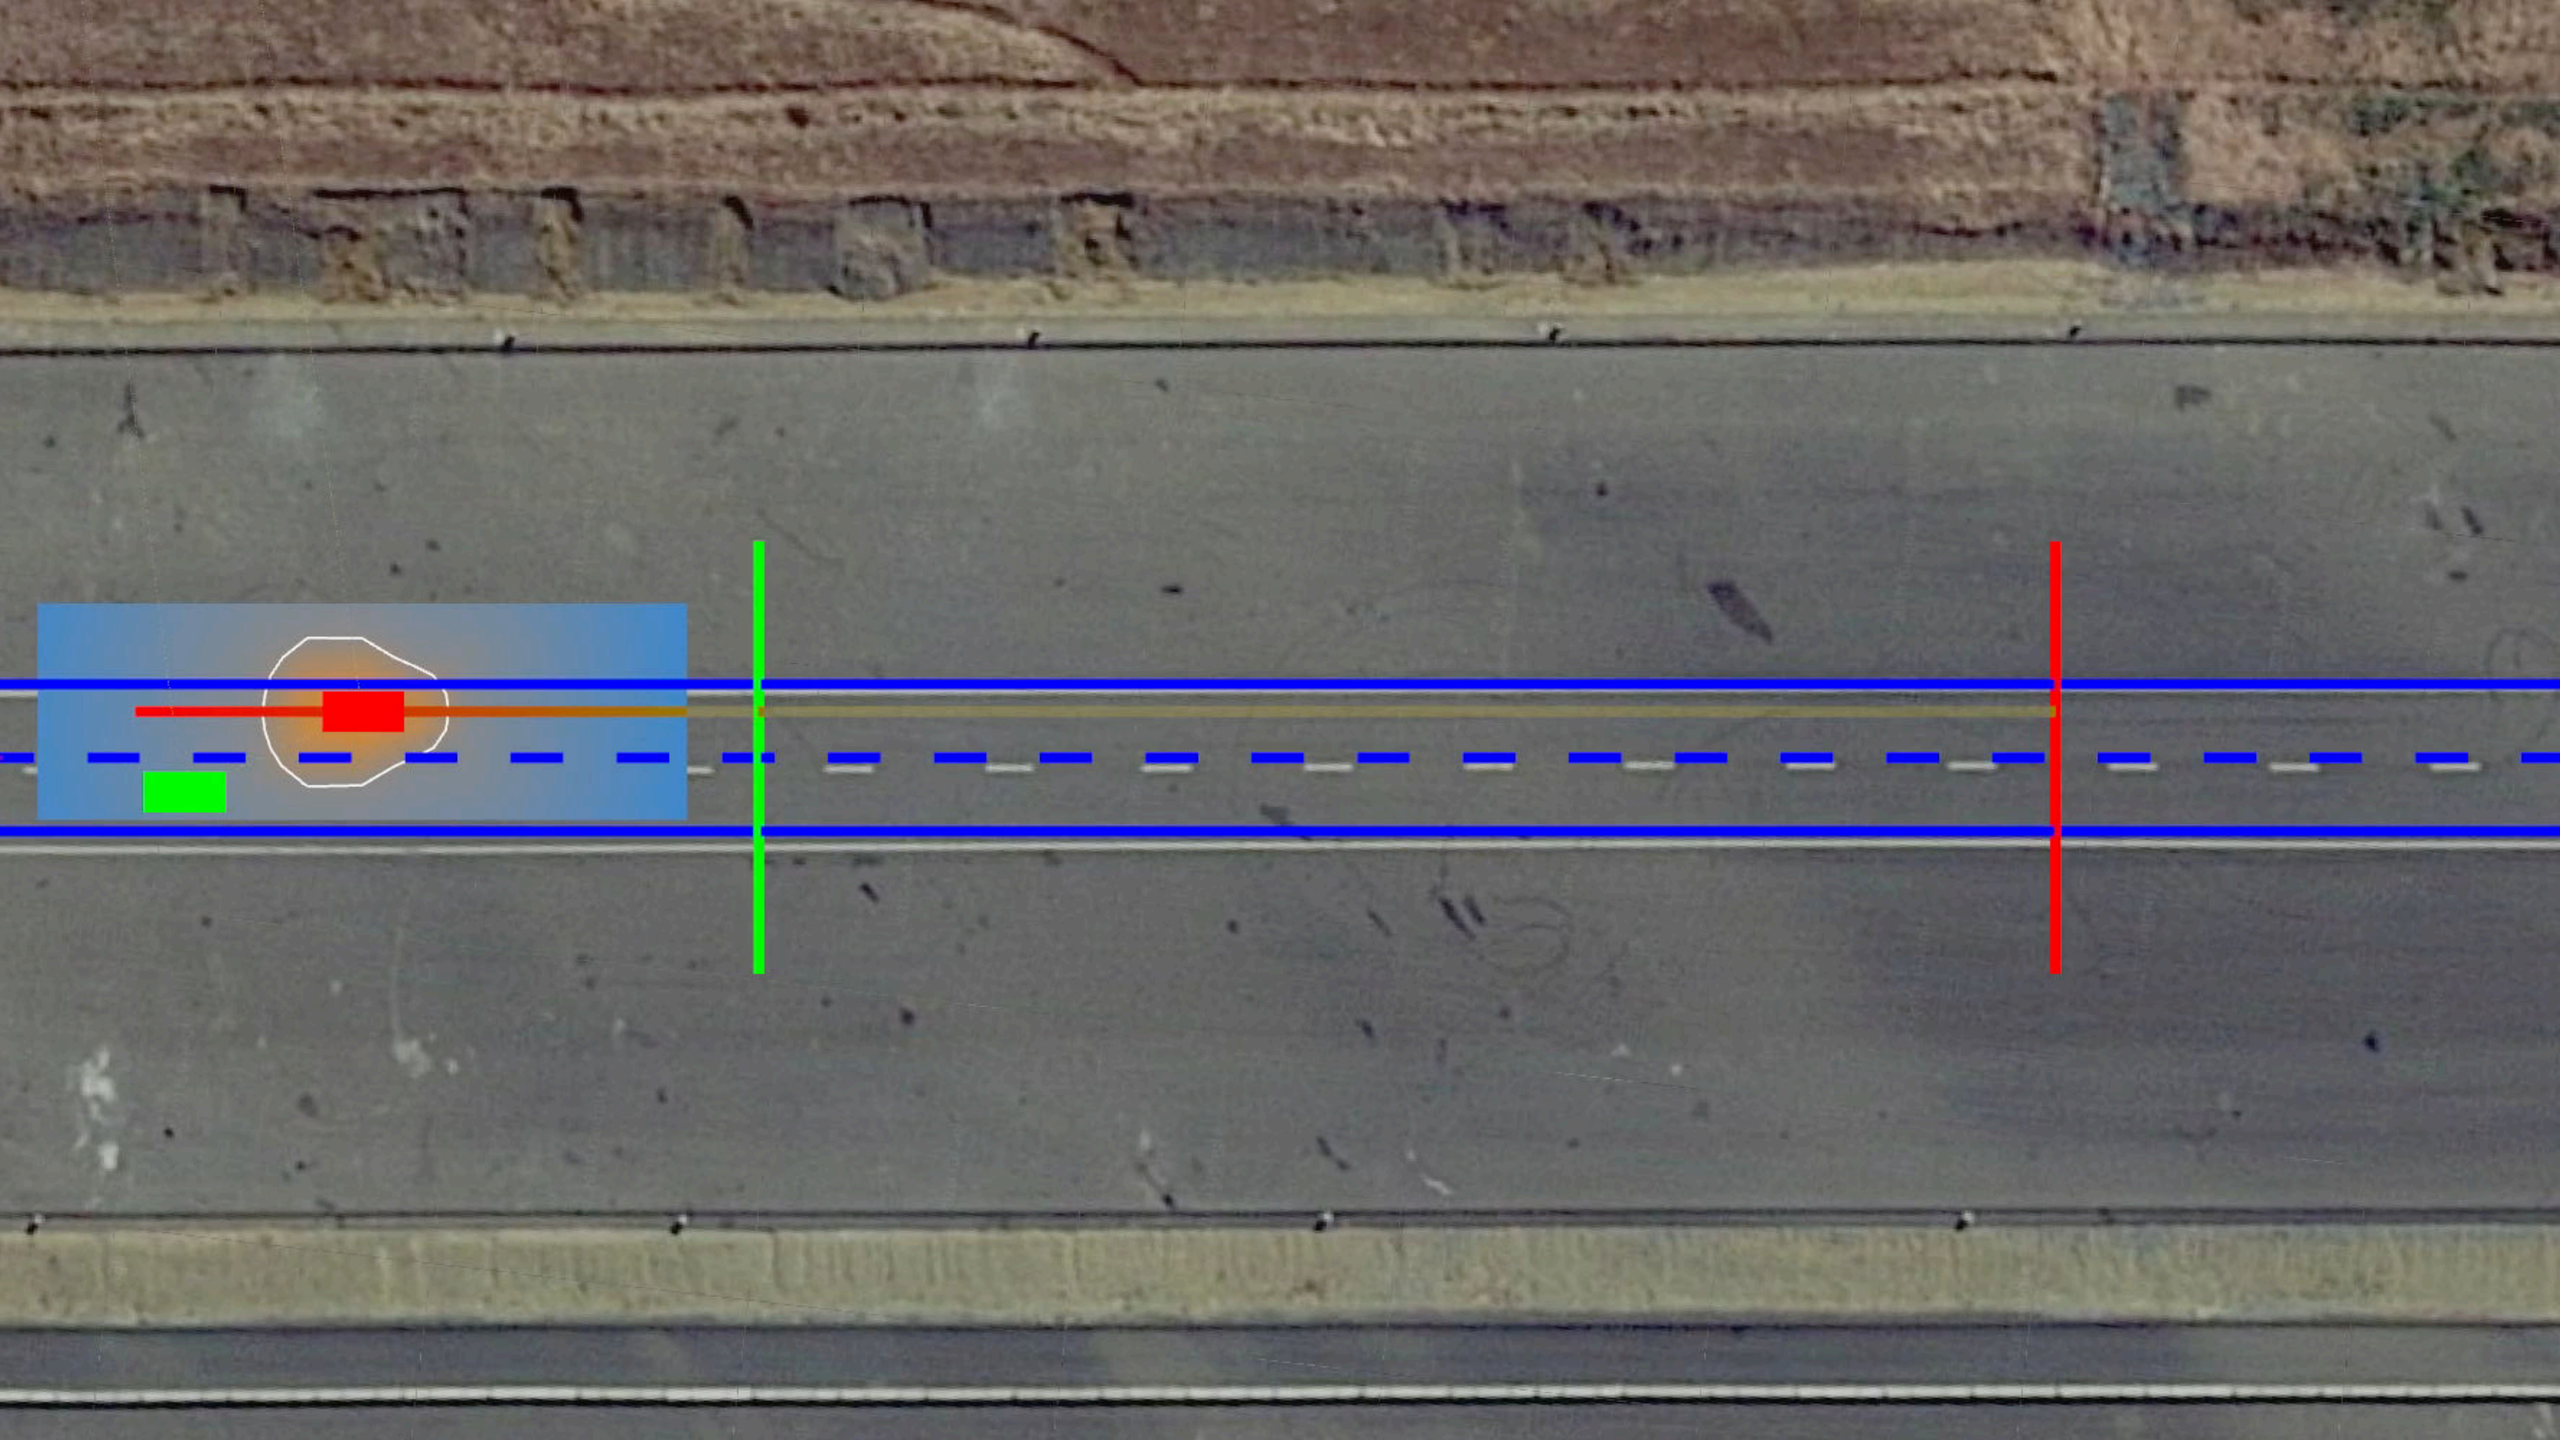

MPC Tracking Controller

(Leung*, Schmerling*, et al. 2018)

- Quadratic program
- Centimeter accuracy
- 100Hz

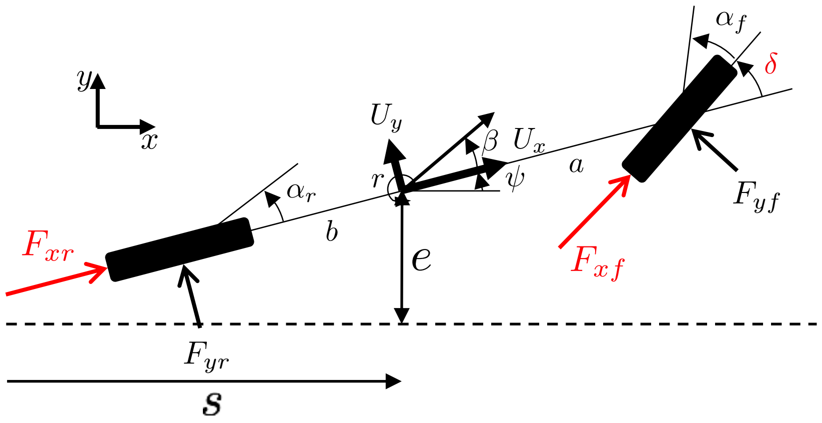

MPC Tracking Controller

(Leung*, Schmerling*, et al. 2018)

- Quadratic program
- Centimeter accuracy
- 100Hz

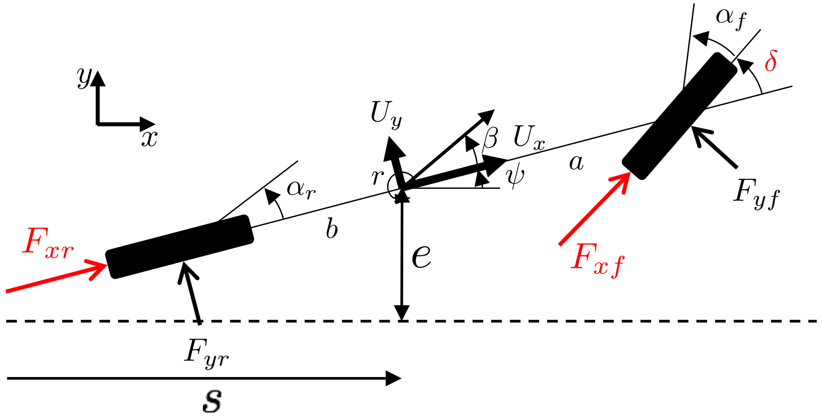

$$\underset{q,u,\sigma,\sigma_{HJI},\Delta\delta,\Delta F_x}{\text{minimize}} \sum_{k=1}^{T} \Delta s_k^T Q_{\Delta s} \Delta s_k + \Delta\psi_k^T Q_{\Delta\psi} \Delta\psi_k + e_k^T Q_e e_k + \Delta\delta_k^T R_{\Delta\delta} \Delta\delta_k +$$

$$\Delta F_{x,k}^T R_{\Delta F_x} \Delta F_{x,k} + W_\beta \sigma_{\beta,k} + W_r \sigma_{r,k} + W_{HJI} \sigma_{HJI,k}$$

$$\text{subject to} \quad \delta_{k+1} - \delta_k = \Delta\delta_k, \quad \Delta\delta_{min} \le \Delta\delta_k \le \Delta\delta_{max}, \quad \delta_{min} \le \delta_k \le \delta_{max}$$

$$F_{x,k+1} - F_{x,k} = \Delta F_{x,k}, \quad V_{min} \le U_{x,k} \le V_{min}, \quad F_{x,min} \le F_{x,k} \le F_{x,max}$$

$$\sigma_{1,k} \ge 0, \quad \sigma_{2,k} \ge 0, \quad \sigma_{HJI,j} \ge 0$$

$$H_k \begin{bmatrix} U_{y,k} \\ r_k \end{bmatrix} - G_k \le \begin{bmatrix} \sigma_{\beta,k} \\ \sigma_{r,k} \end{bmatrix}, \quad A_k q_k + B_k^- u_k + B_k^+ u_{k+1} + c_k = q_{k+1}$$

$$q_1 = q_{curr}, \quad u_1 = u_{curr}, \quad M_{HJI} u_j + b_{HJI} \ge -\sigma_{HJI}$$

$$\text{for } j = 1, \ldots, T_{HJI}, \quad k = 1, \ldots, T$$

MPC Tracking Controller

(Leung*, Schmerling*, et al. 2018)

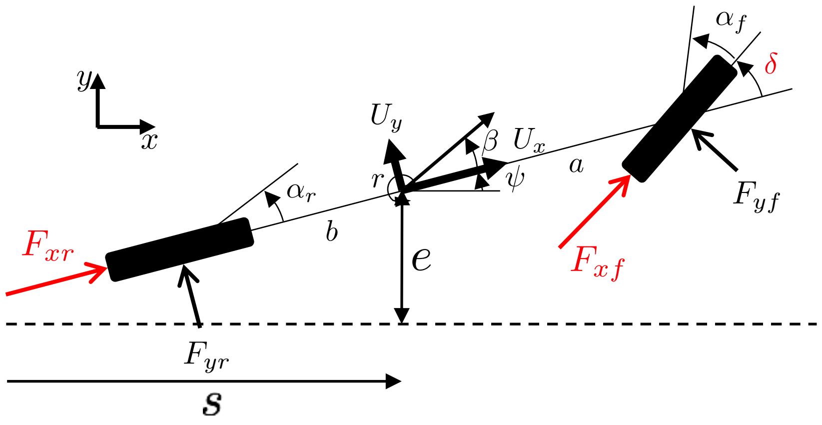

- Quadratic program
- Centimeter accuracy
- 100Hz

$$\underset{q,u,\sigma,\sigma_{HJI},\Delta\delta,\Delta F_x}{\text{minimize}} \quad \sum_{k=1}^{T}$$

Quadratic cost: position error, control, control rate, slack variables

$$\begin{aligned}
\text{subject to} \quad & \delta_{k+1} - \delta_k = \Delta\delta_k, \quad \Delta\delta_{min} \le \Delta\delta_k \le \Delta\delta_{max}, \quad \delta_{min} \le \delta_k \le \delta_{max} \\
& F_{x,k+1} - F_{x,k} = \Delta F_{x,k}, \quad V_{min} \le U_{x,k} \le V_{min}, \quad F_{x,min} \le F_{x,k} \le F_{x,max} \\
& \sigma_{1,k} \ge 0, \quad \sigma_{2,k} \ge 0, \quad \sigma_{HJI,j} \ge 0 \\
& H_k \begin{bmatrix} U_{y,k} \\ r_k \end{bmatrix} - G_k \le \begin{bmatrix} \sigma_{\beta,k} \\ \sigma_{r,k} \end{bmatrix}, \quad A_k q_k + B_k^- u_k + B_k^+ u_{k+1} + c_k = q_{k+1} \\
& q_1 = q_{curr}, \quad u_1 = u_{curr}, \quad M_{HJI} u_j + b_{HJI} \ge -\sigma_{HJI} \\
& \text{for } j = 1, \ldots, T_{HJI}, \quad k = 1, \ldots, T
\end{aligned}$$

MPC Tracking Controller

(Leung*, Schmerling*, et al. 2018)

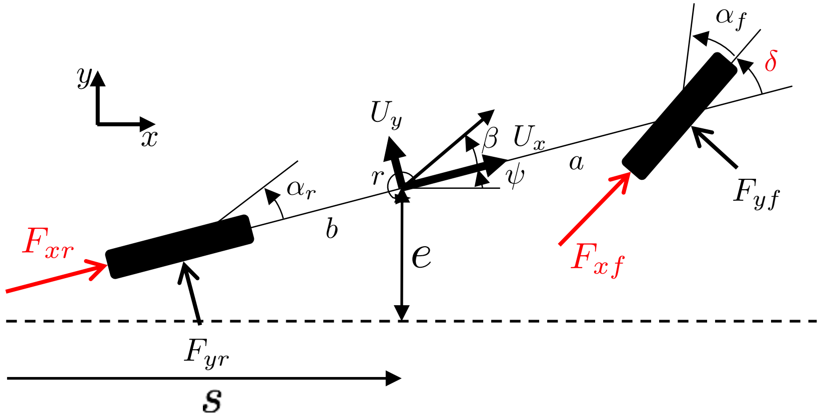

- Quadratic program
- Centimeter accuracy
- 100Hz

$$\underset{q,u,\sigma,\sigma_{HJI},\Delta\delta,\Delta F_x}{\text{minimize}} \sum_{k=1}^{T}$$

Quadratic cost: position error, control, control rate, slack variables

subject to

Constraints: Continuity, saturation, positivity in slack variables, stability and environmental envelope, linearized dynamics, initial conditions

$$\text{for } j = 1, \ldots, T_{HJI}, \quad k = 1, \ldots, T$$

(Leung*, Schmerling*, et al. 2018)

MPC Tracking Controller

- Quadratic program
- Centimeter accuracy
- 100Hz

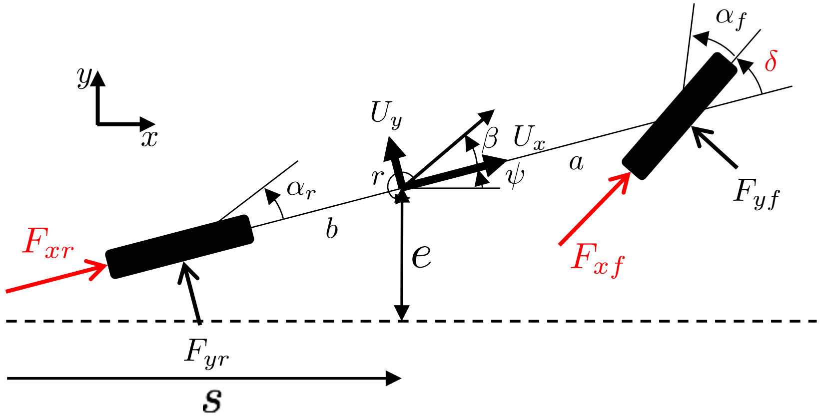

$$\underset{q,u,\sigma,\sigma_{HJI},\Delta\delta,\Delta F_x}{\text{minimize}} \sum_{k=1}^{T}$$

Quadratic cost: position error, control, control rate, slack variables

subject to

Constraints: Continuity, saturation, positivity in slack variables, stability and environmental envelope, linearized dynamics, initial conditions

Reachability **constraint**

for $j = 1, \dots, T_{HJI}, \quad k = 1, \dots, T$

Probabilistic Model-based Planning

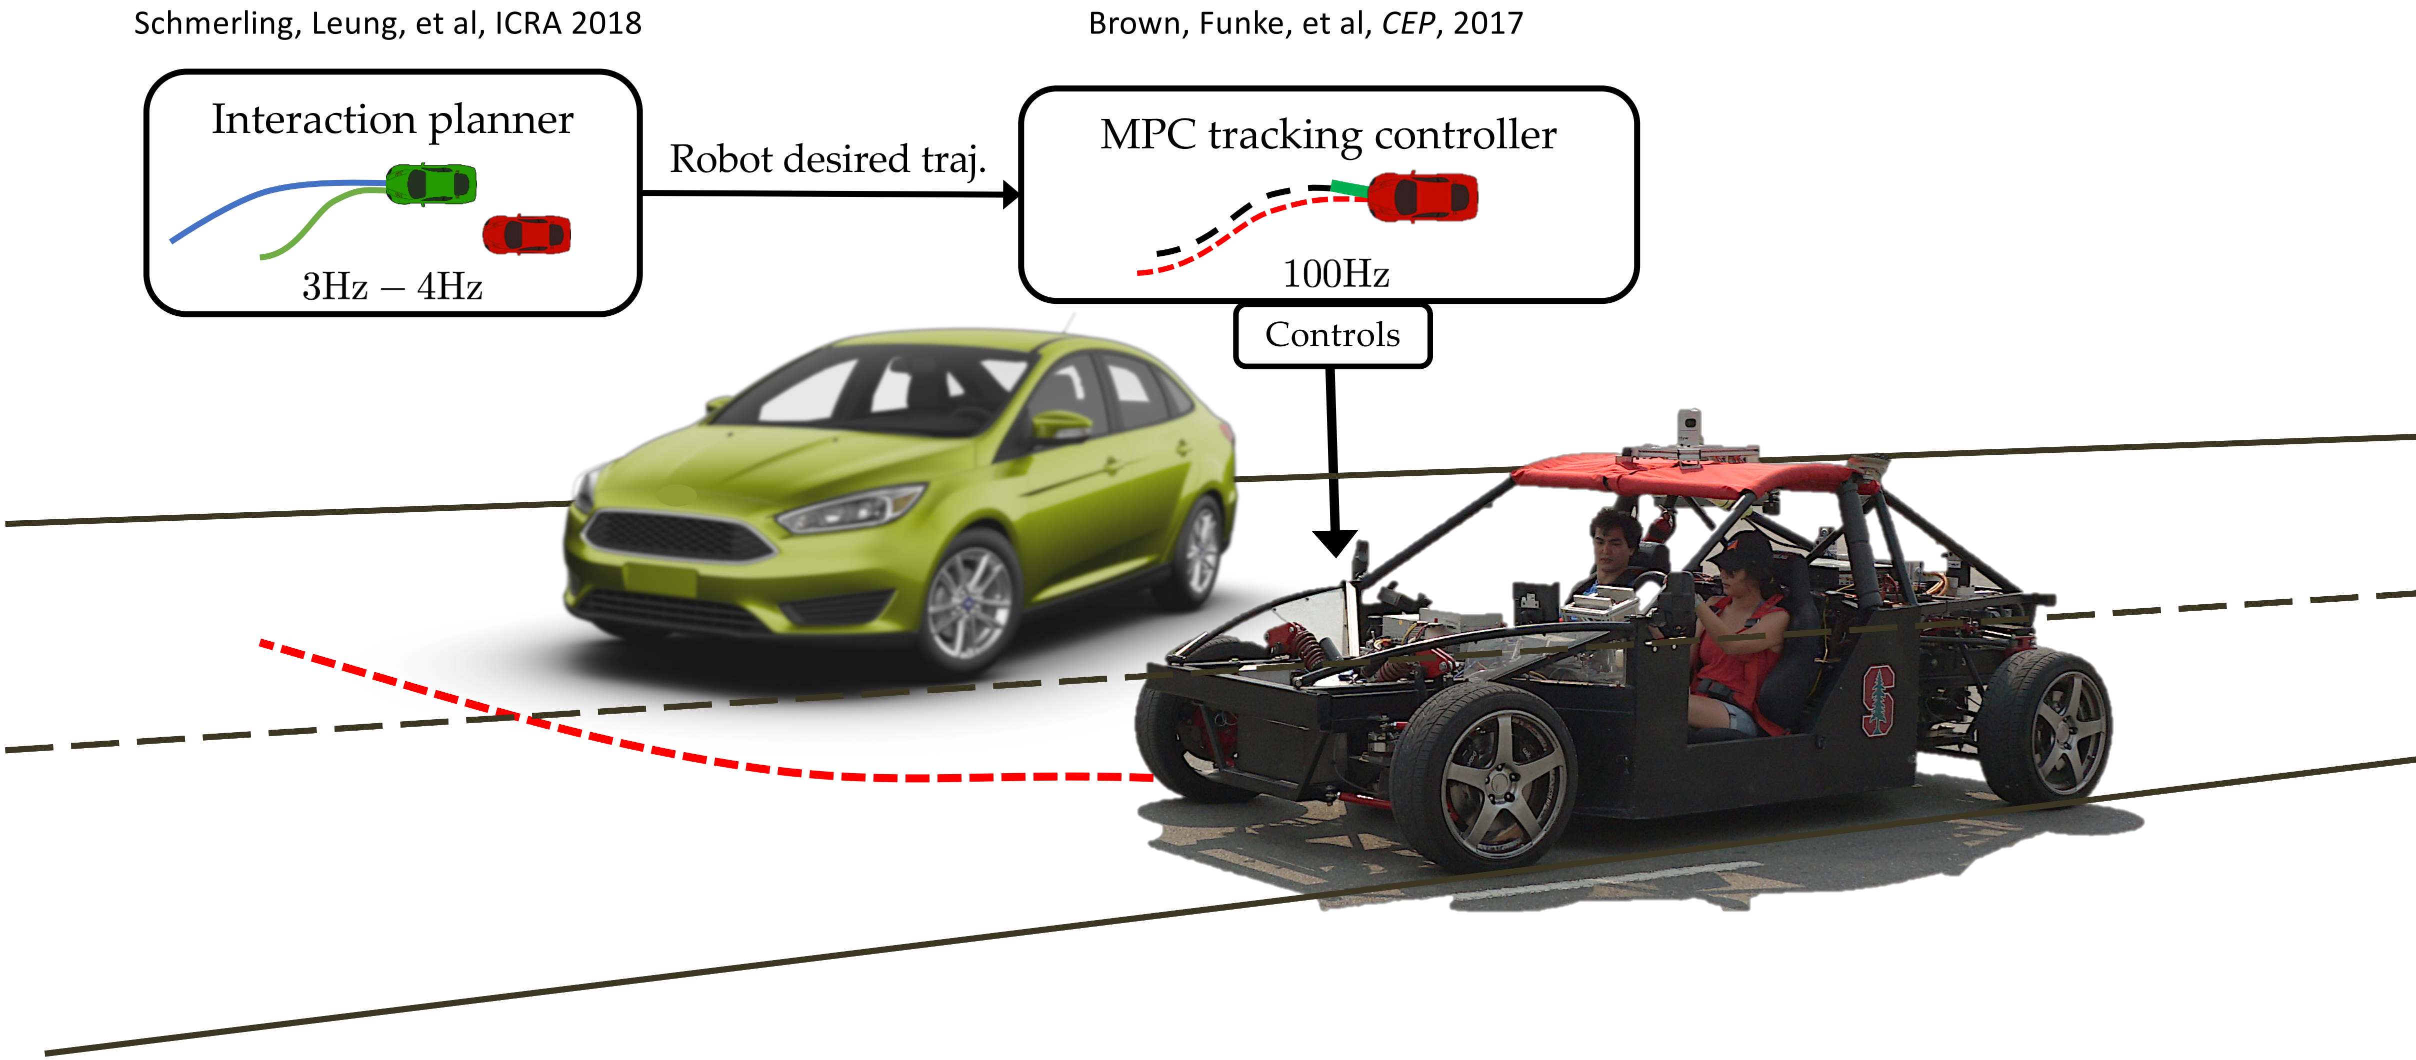

(Leung*, Schmerling*, et al. 2018)

Probabilistic Model-based Planning with **Safety Assurance**

Schmerling, Leung, et al, ICRA 2018

Brown, Funke, et al, *CEP*, 2017

Chen and Tomlin, *AR of CRAS*, 2018.

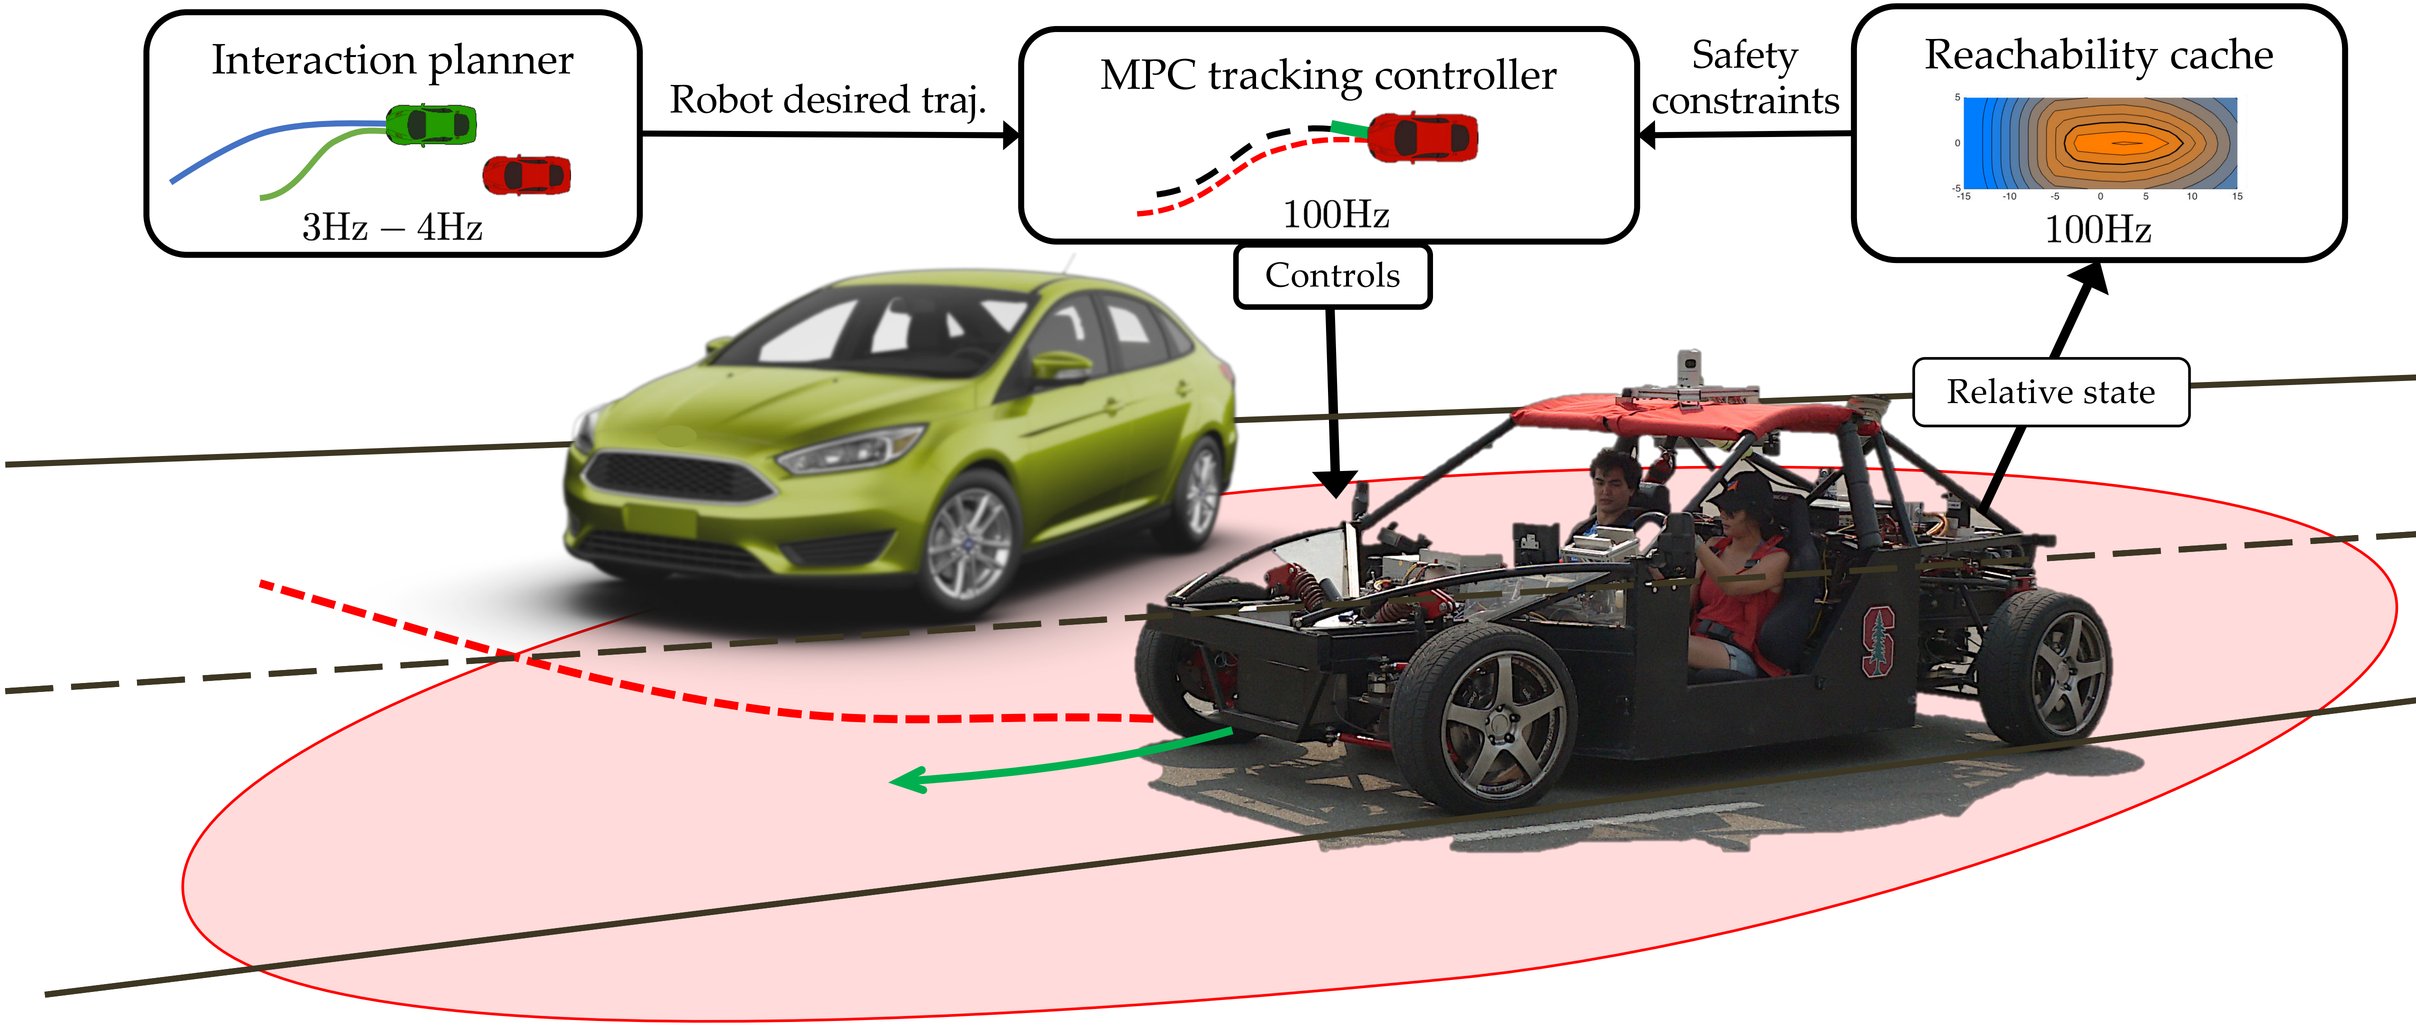

(Leung*, Schmerling*, et al. 2018)

(Leung*, Schmerling*, et al. 2018)

Experimental Platform (X1)

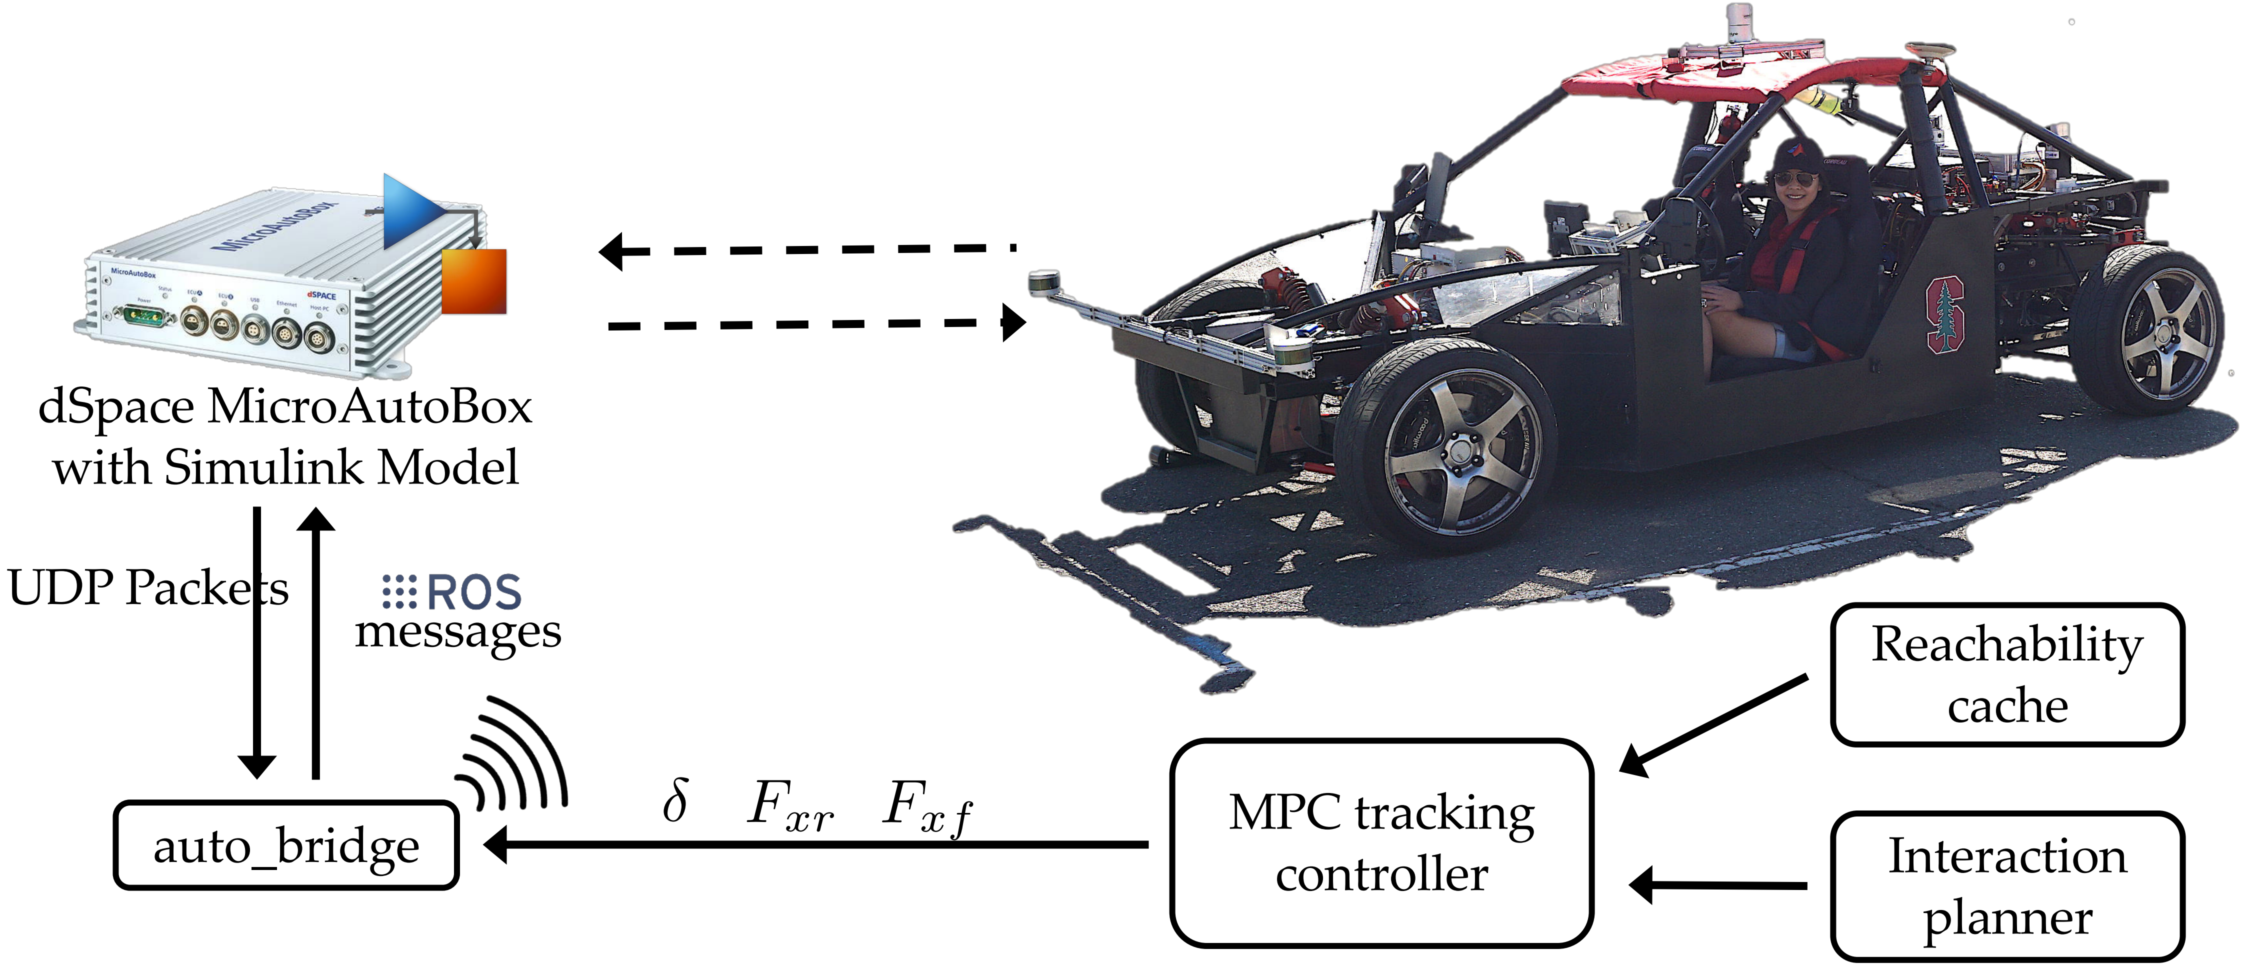

(Leung*, Schmerling*, et al. 2018)

Experimental Results

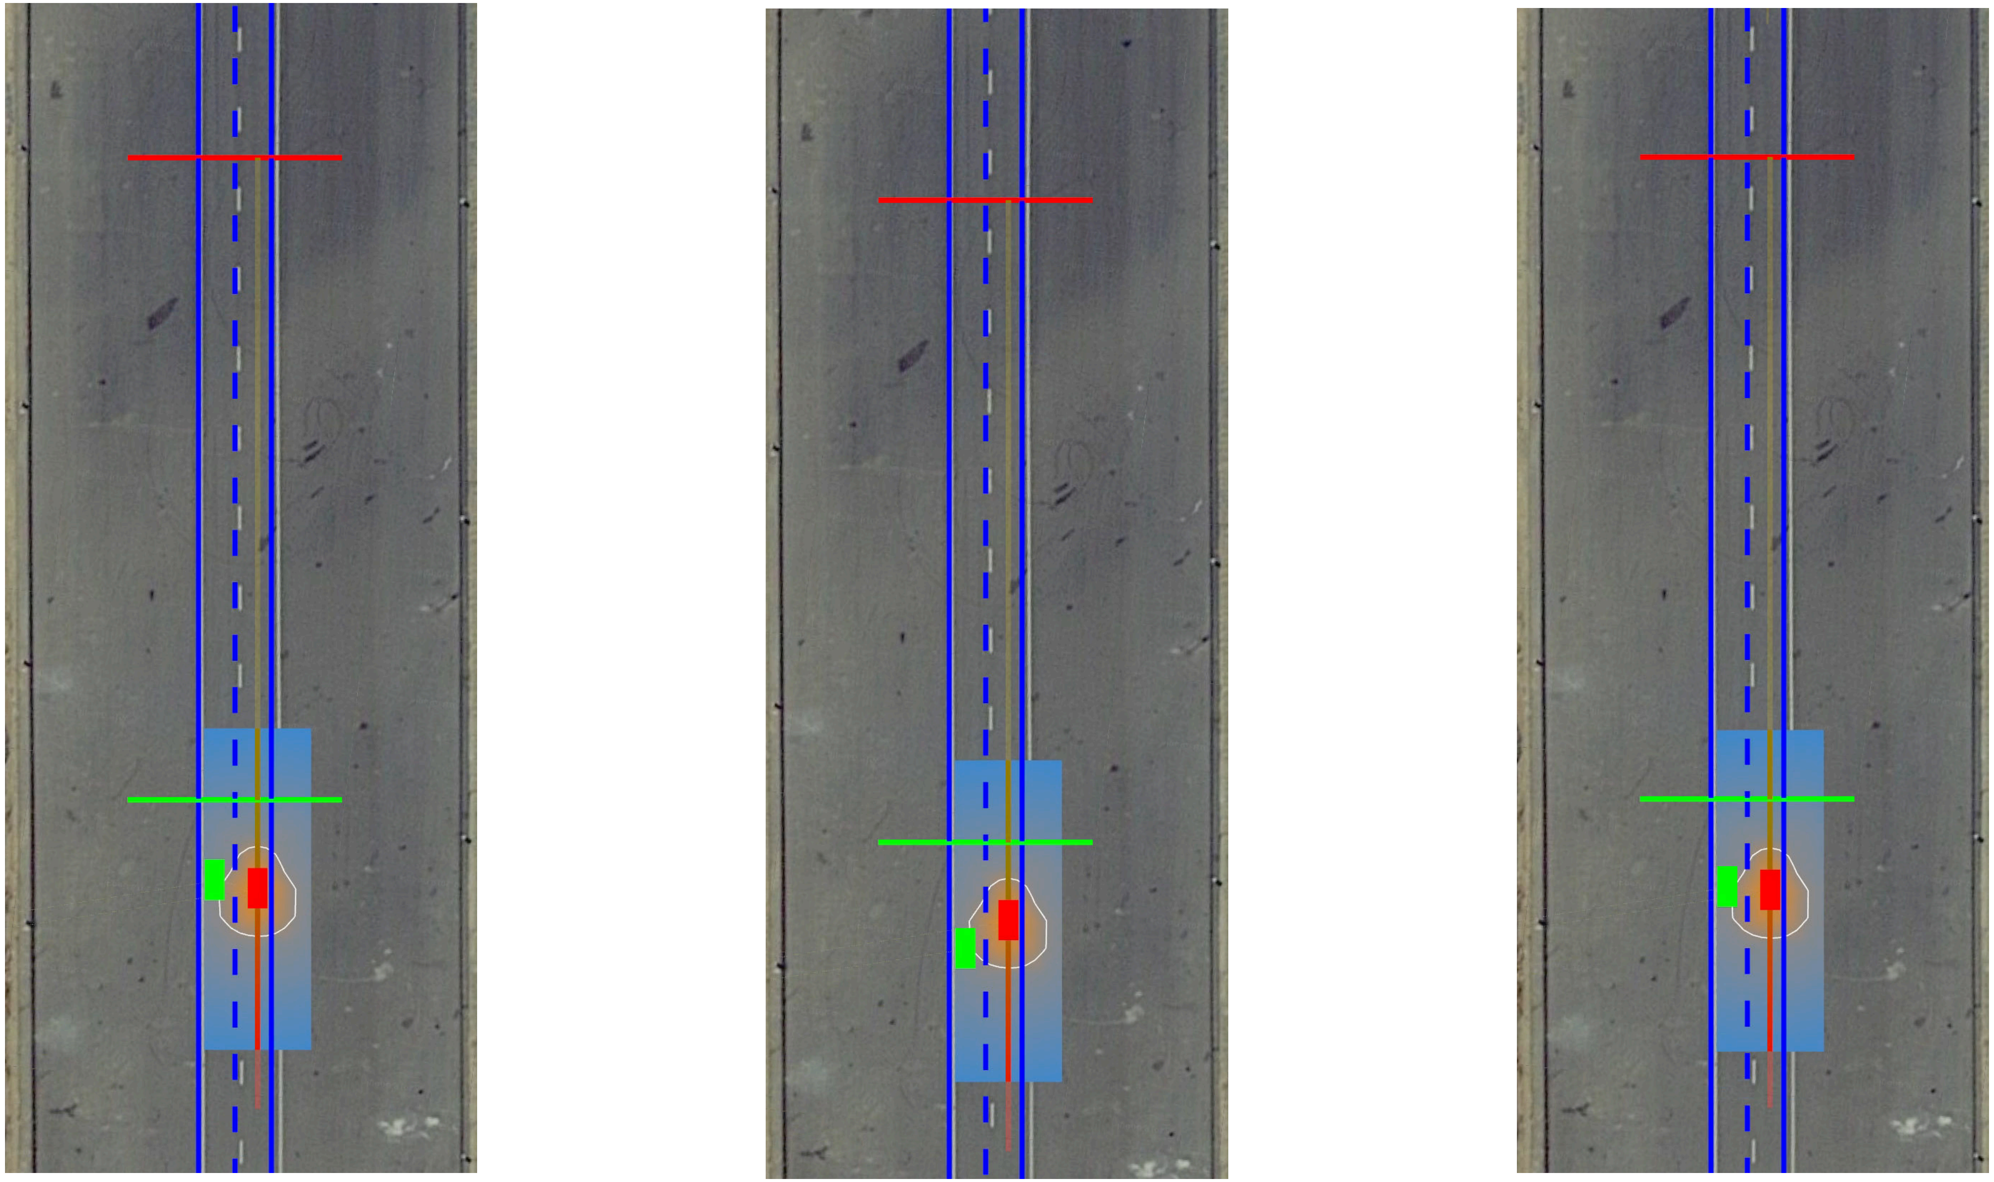

(Leung*, Schmerling*, et al. 2018)

Experimental Results

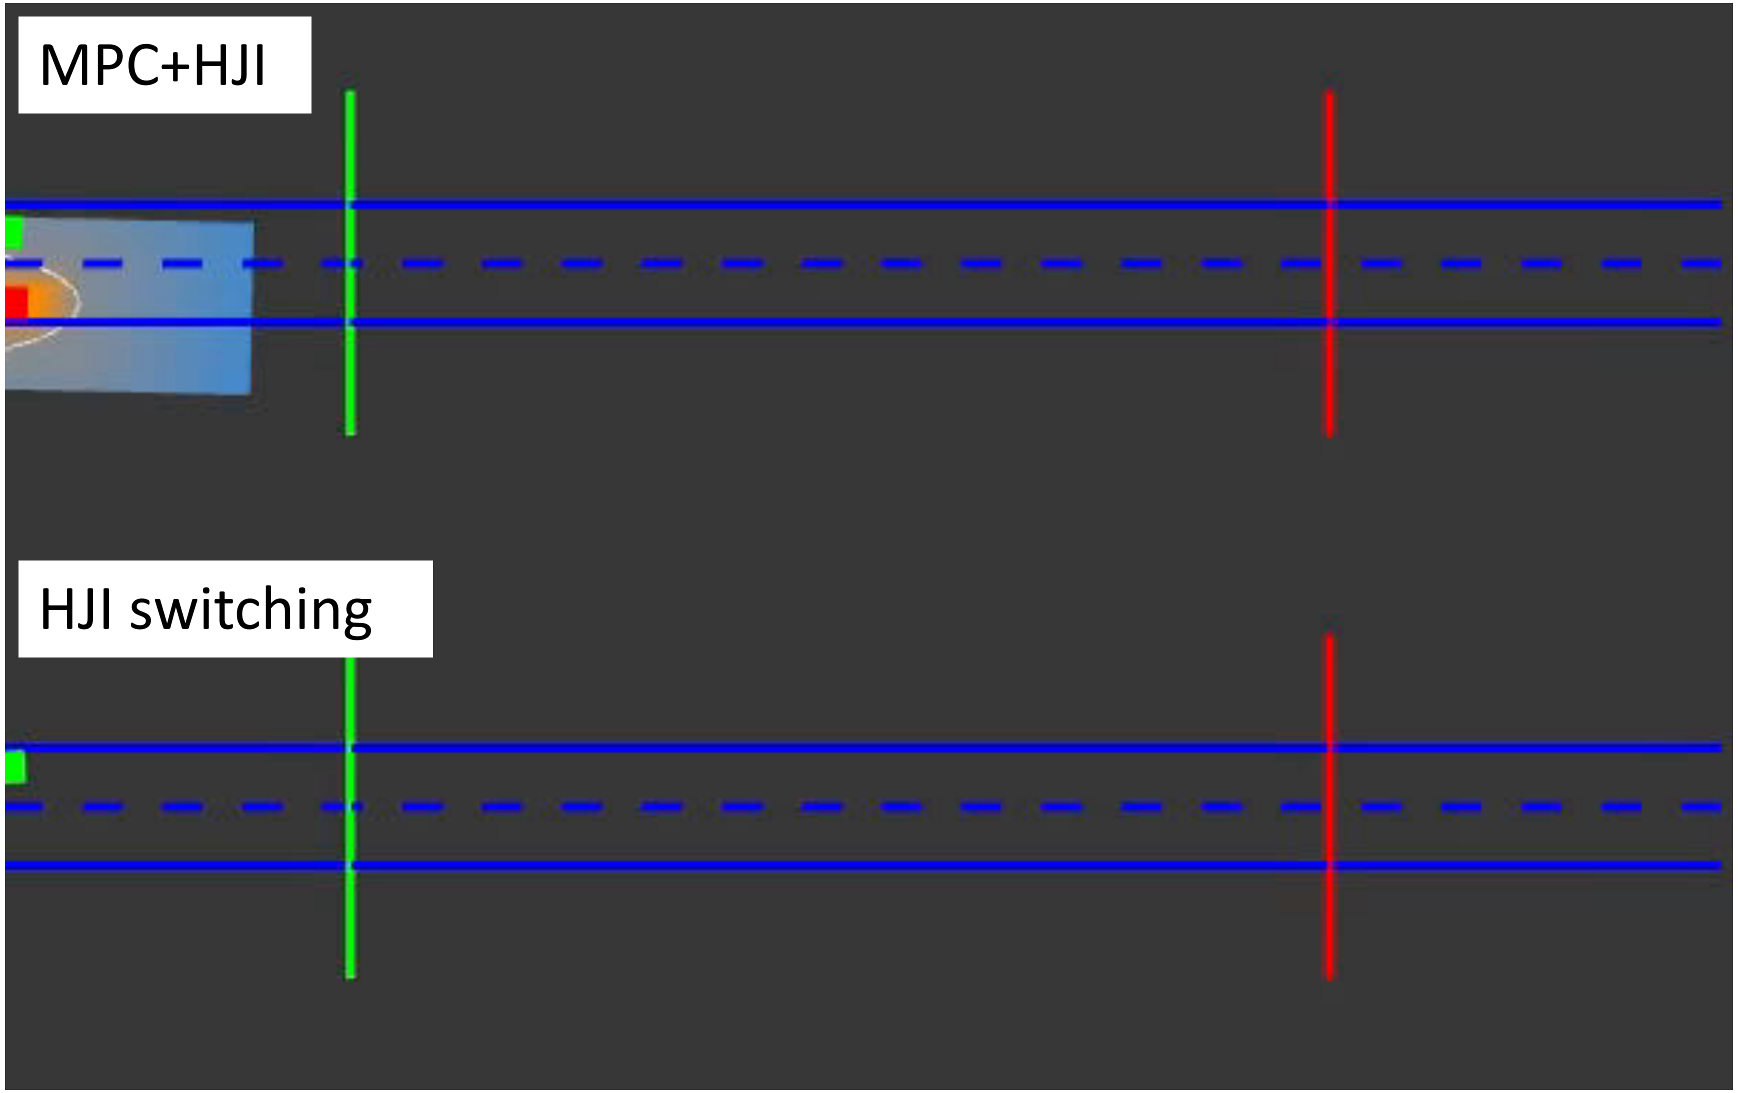

(Leung*, Schmerling*, et al. 2018)

Experimental Results

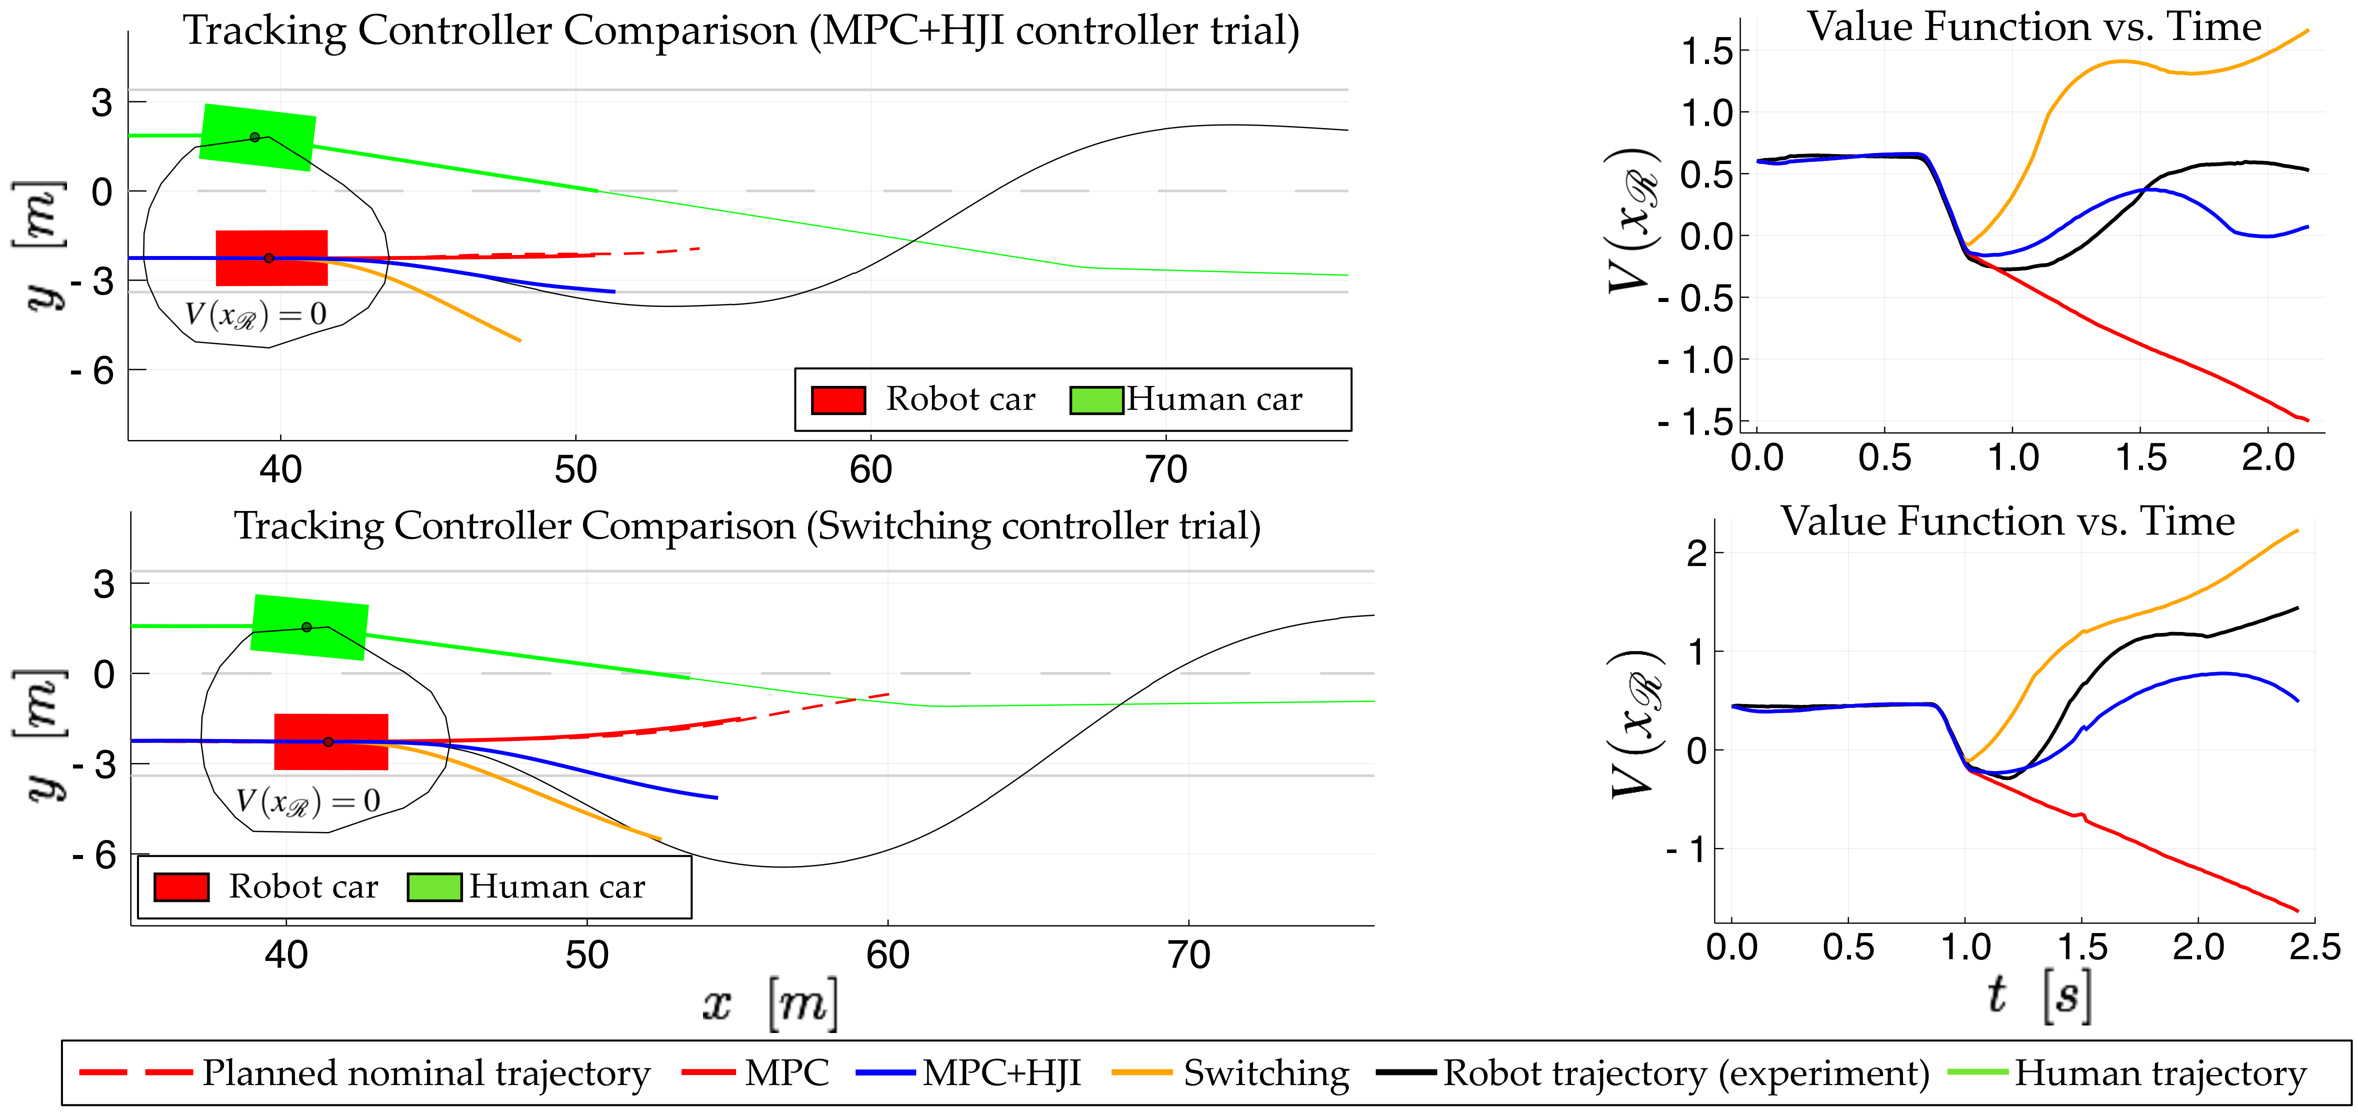

Experimental Insights

(Leung*, Schmerling*, et al. 2018)

- Unmodeled steering angle slew rate causes us to dip into danger
- Negative value function represents worst-case collision penetration
  - Alternative: collision severity
- Interpretability of value function is key for more realistic scenarios (barriers, multiple other agents)
- Ongoing experimental work: incorporating static obstacles as MPC constraints

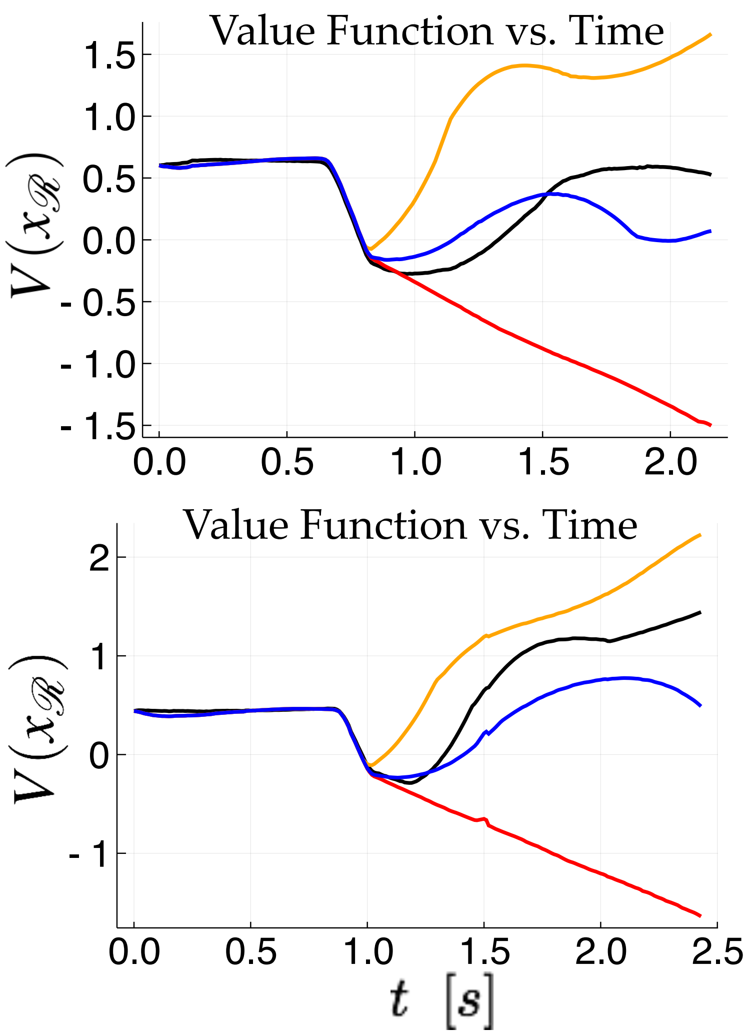

1/10 scale RC car

Referenced Papers

- B. Ivanovic, E. Schmerling, K. Leung and M. Pavone, “Generative Modeling of Multimodal Multi-Human Behavior,” in IEEE/RSJ Int. Conf. on Intelligent Robots & Systems, 2018. Available at https://arxiv.org/abs/1803.02015.

- B. Ivanovic and M. Pavone, “The Trajectron: Probabilistic Multi-Agent Trajectory Modeling with Dynamic Spatiotemporal Graphs,” Arxiv Preprint. Available at https://arxiv.org/abs/1810.05993.

- E. Schmerling, K. Leung, W. Vollprecht and M. Pavone, “Multimodal Probabilistic Model-Based Planning for Human-Robot Interaction,” in Proc. IEEE Conf. on Robotics and Automation, 2018. Available at https://arxiv.org/abs/1710.09483.

- K. Leung, E. Schmerling, M. Chen, J. Talbot, J.C. Gerdes and M. Pavone, “On Infusing Reachability-Based Safety Assurance within Probabilistic Planning Frameworks for Human-Robot Vehicle Interactions,” in Int. Symp. on Experimental Robotics, 2018. Available at https://arxiv.org/abs/1812.11315.