

# On Safe and Efficient Human-robot Interactions via Multimodal Intent Modeling and Reachability-based Safety Assurance


Marco Pavone

Autonomous Systems Laboratory  
Department of Aeronautics and Astronautics  
Stanford University

IPAM Workshop: Safe Operation of Connected and  
Autonomous Vehicle Fleets


October 29th, 2020

[pavone@stanford.edu](mailto:pavone@stanford.edu)






Stanford  
University





Courtesy of NASA JPL



# From reactive to proactive decision making



“Merging into traffic during rush hour is an exercise in negotiation.”

— Google Report, 2016

## Proactive decision making:


1. proactively interact with other agents to infer their intents, while **concurrently**
2. exploiting this information to take actions that account for agent responses

# High-level considerations

**Approach:** *model-based* decision making for pairwise interaction


## Key considerations:

1. Conditional prediction of human intent
2. Time scales on the order of  $\sim 1s$
3. Uncertainty is generally multimodal
4. History-dependent predictions
5. Interpretability




# Data-driven probabilistic modeling

[Schmerling, Leung, Vollprecht, Pavone, ICRA '18 & Ivanovic, Schmerling, Leung, Pavone, IROS '18]



# Data-driven probabilistic modeling

[Schmerling, Leung, Vollprecht, Pavone, ICRA '18 & Ivanovic, Schmerling, Leung, Pavone, IROS '18]



Phenomenological approach: can we learn action distributions directly from experience, **without reasoning about motivations?**

# Generative model of human action distributions

[Schmerling, Leung, Vollprecht, Pavone, ICRA '18 & Ivanovic, Schmerling, Leung, Pavone, IROS '18]

Generative model of human action **distributions** conditioned on:

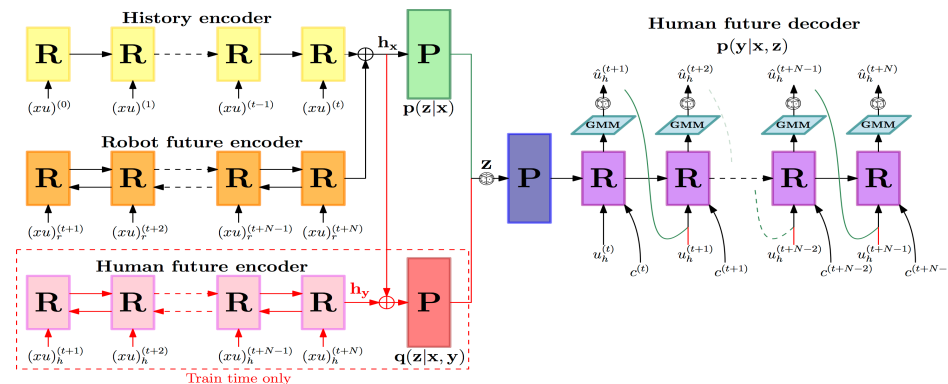
- Joint interaction history
- Candidate robot future action sequence

$$U_H^{(t+1)} \sim p(u_H^{(t+1)} \mid x^{(0:t)}, u^{(0:t)}, u_R^{(t+1)})$$

Interaction History

Robot's Next Action

Learnt via a CVAE-based model

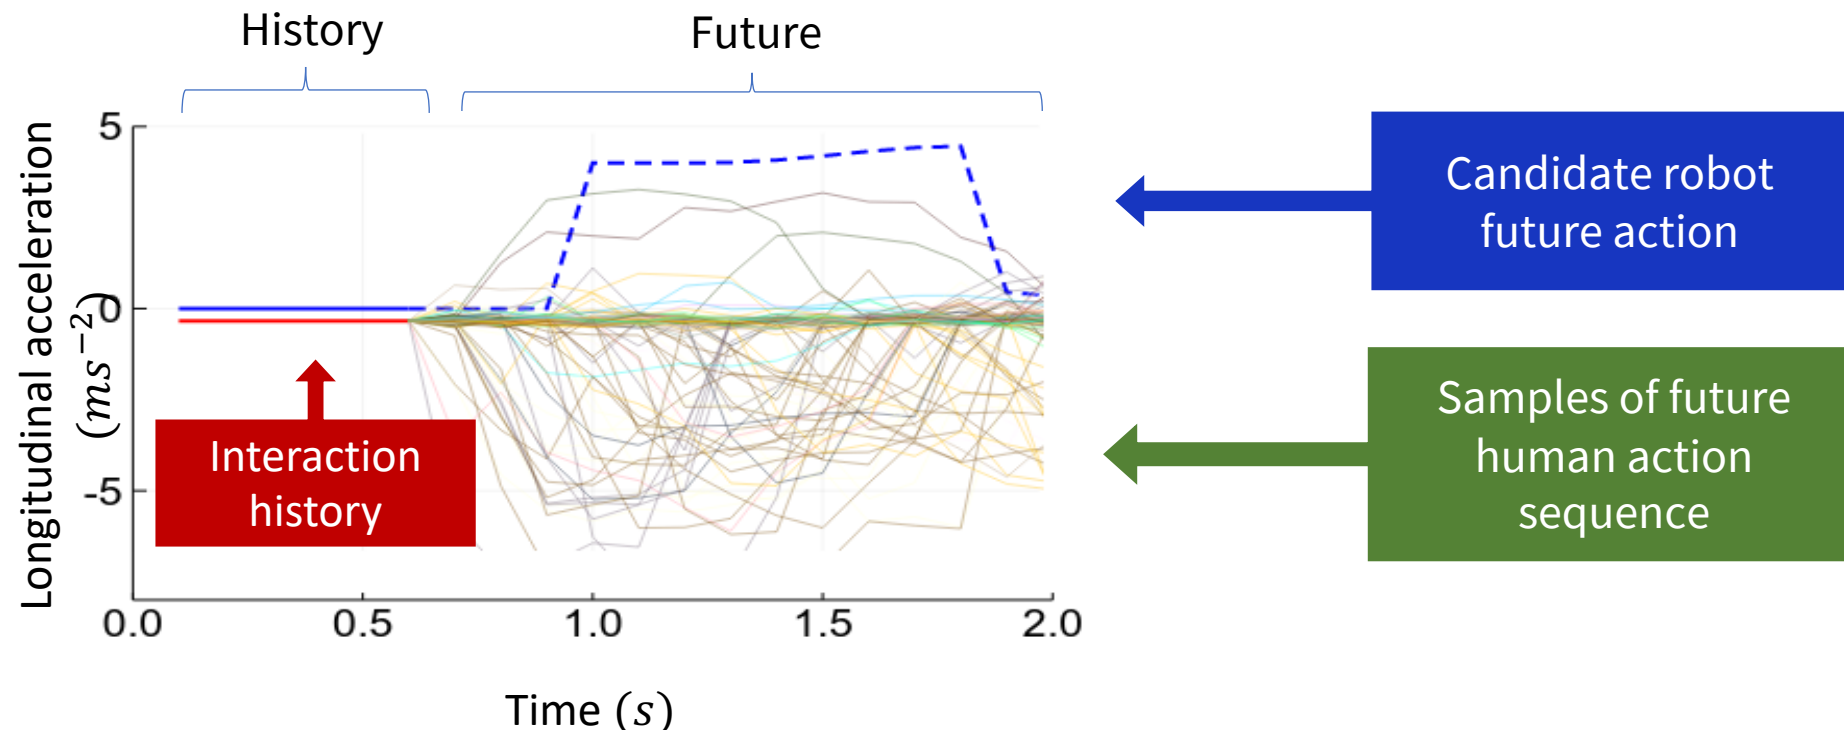


# Generative model of human action distributions

[Schmerling, Leung, Vollprecht, Pavone, ICRA '18 & Ivanovic, Schmerling, Leung, Pavone, IROS '18]

Generative model of human action **distributions** conditioned on:

- Joint interaction history
- Candidate robot future action sequence

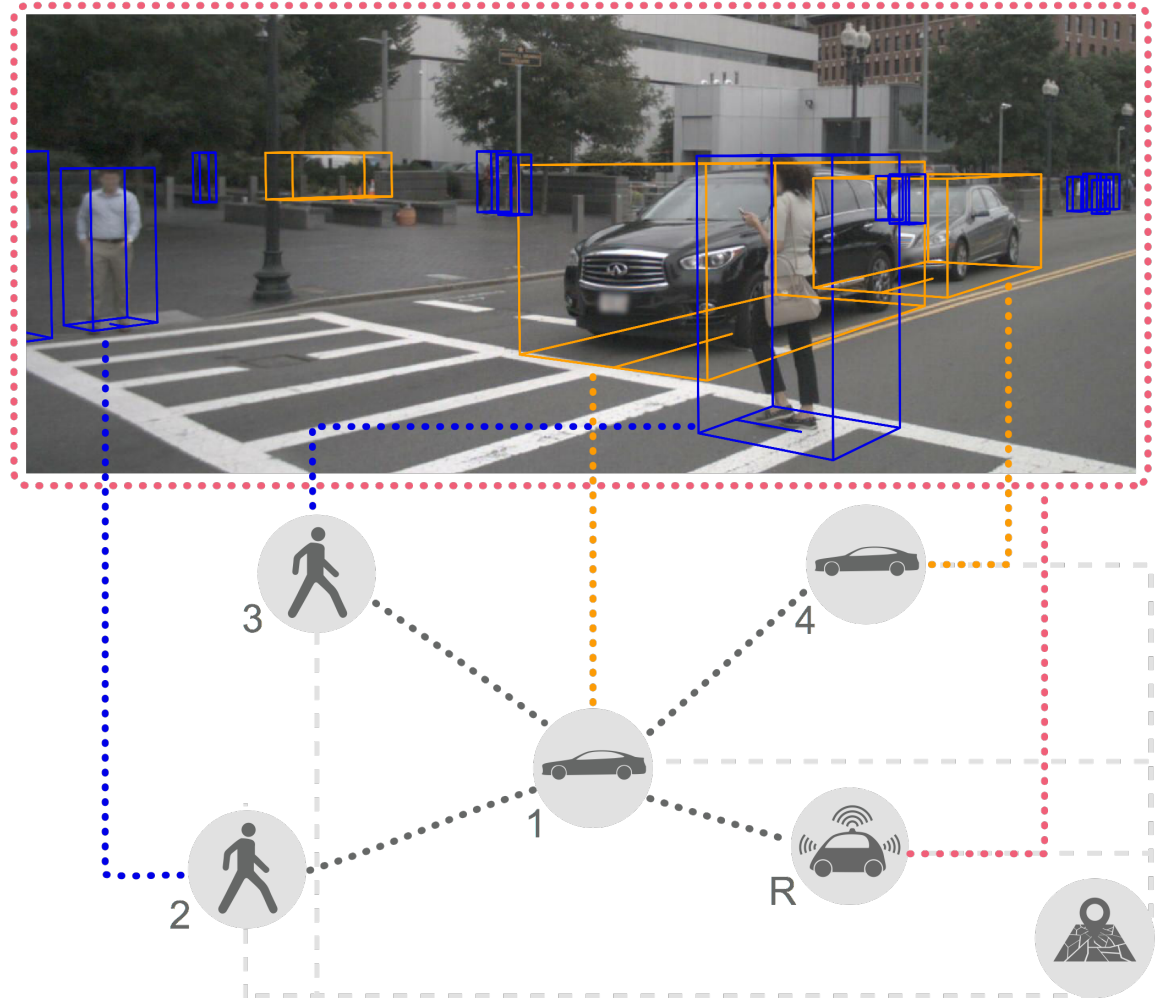


# General multi-agent trajectory modeling

[Ivanovic, Pavone, ICCV '19 & Salzmann, Ivanovic, Chakravarty, and Pavone, ECCV '20]

Towards a general multi-agent model that can ingest “everything”

1. Predictions for **any kind of agent**
2. Accounting of **dynamics constraints**
3. Conditioning on **heterogenous data**

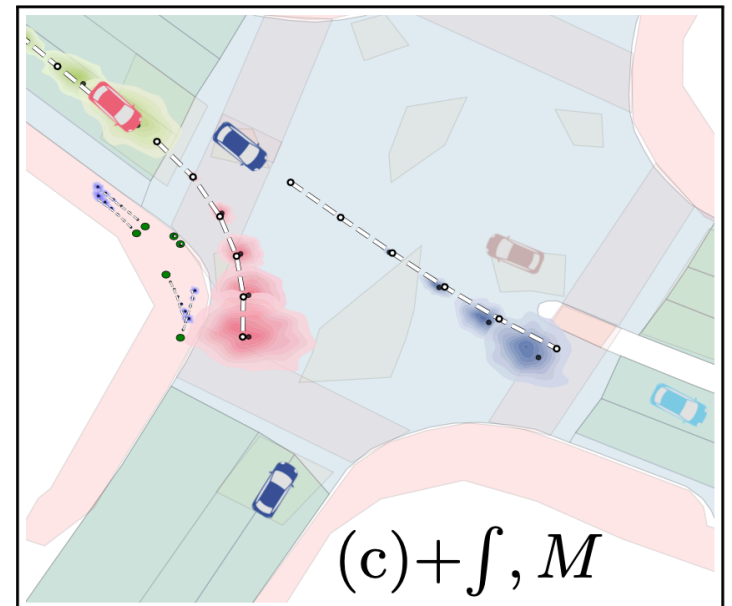
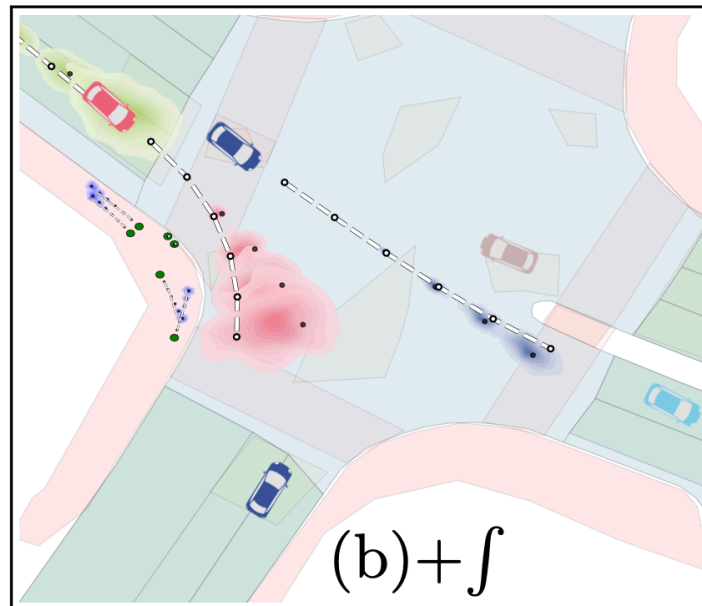
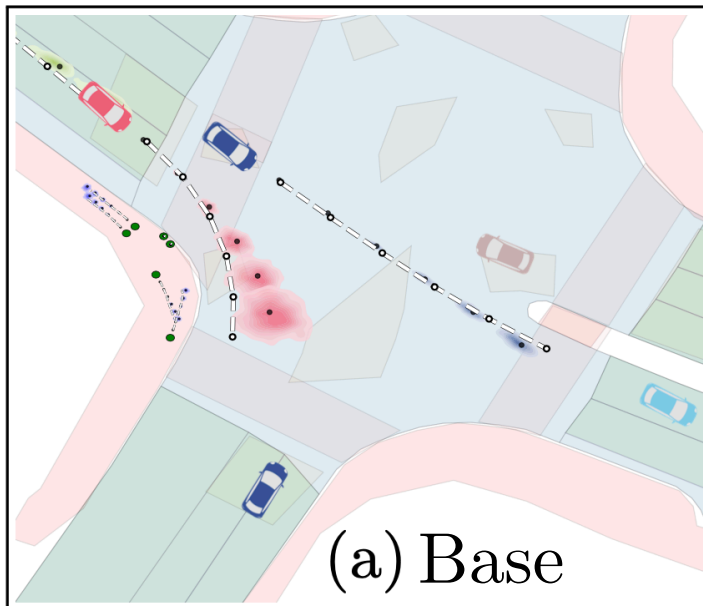


# Trajectron++

[Salzmann, Ivanovic, Chakravarty, and Pavone, ECCV '20]



State-of-the-art generative model that explicitly incorporates agent dynamics and heterogeneous data




All code, models, and data available at <https://github.com/StanfordASL/Trajectron-plus-plus>

Takeaway message: **phenomenological models** (and, in particular, deep generative models) achieve state-of-the-art performance


# Robot policy construction


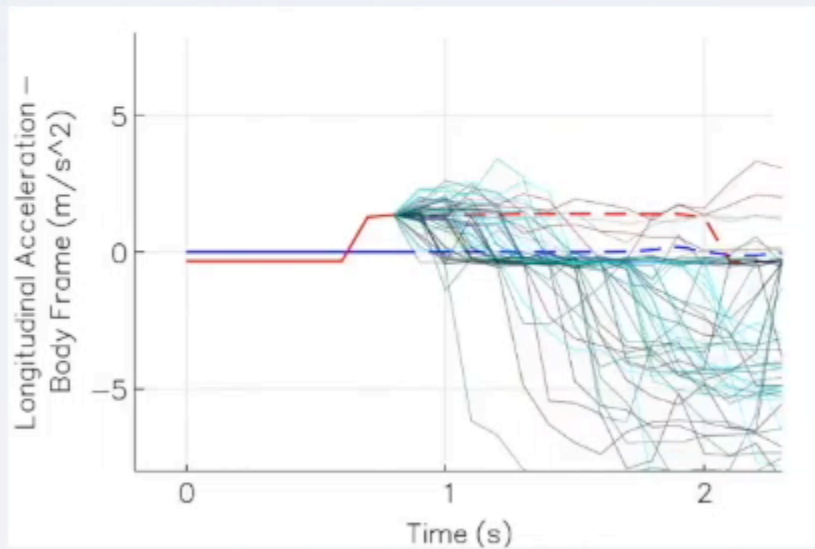

- Robot builds an action tree (~4,000 action sequences)
- Model used to sample human responses (>100K!)
- Action sequences scored via an *aggregate* cost function



$$\mathbb{E} \left( \sum_i J_i \right) \xrightarrow{?} \rho \left( \sum_i J_i \right)$$

# Human-in-the-loop testing

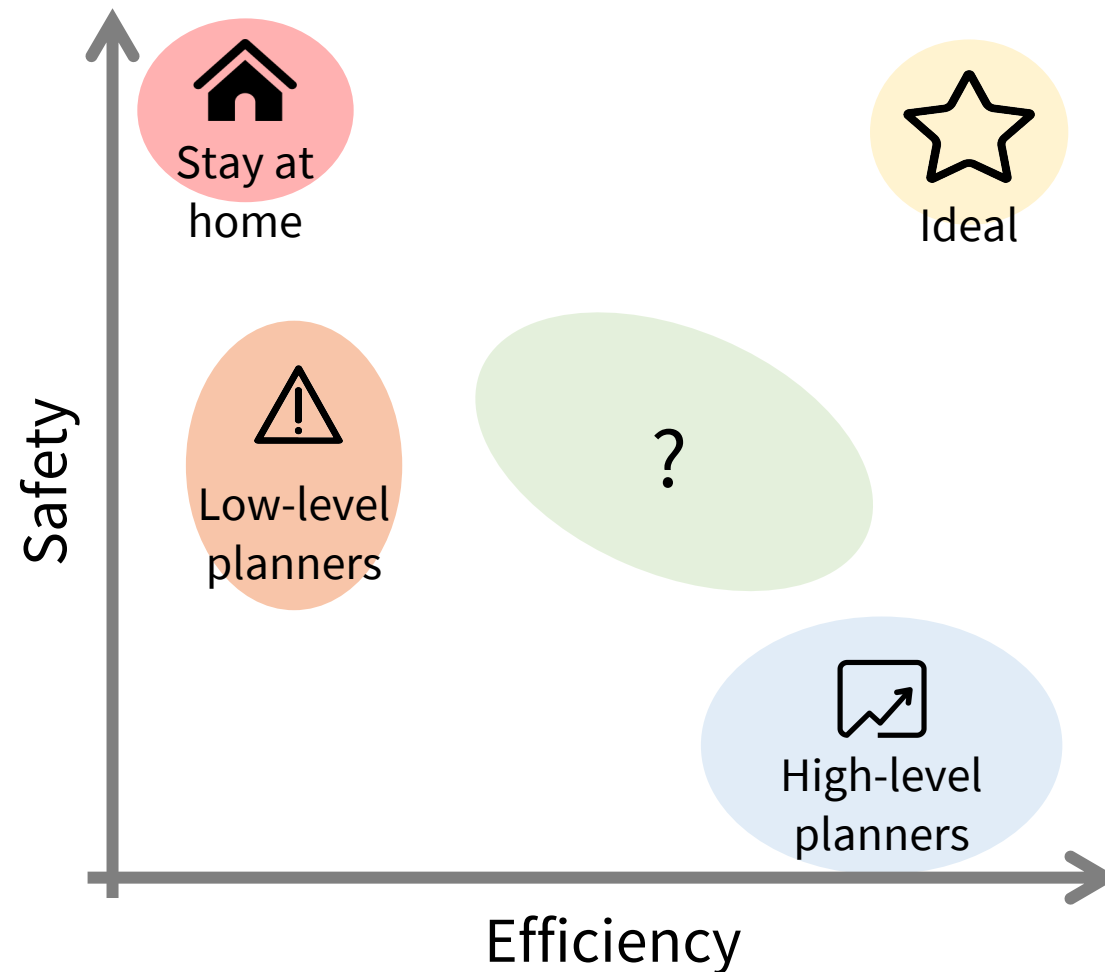






- Probabilistic model may “get it wrong”
- Incorporating collision avoidance as a penalty can cause conflicting objectives
- Replanning at 3Hz is ultimately too slow to ensure safety


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



# On Infusing reachability-based safety assurance

[Leung, Schmerling, Chen, Talbot, Gerdes, Pavone, IJRR '20; Wang, Leung, Pavone IROS '20]

- Several approaches available to promote safety
  - Hand-designed artificial potential fields (e.g., [Wolf et al., 2008])
  - Constrained optimization (e.g., [Funke et al., 2017])
  - Precomputed emergency maneuver libraries (e.g., [Arora et al., 2015])
  - Interaction with a sentient agent is generally not accounted for
- HJI reachability for safety assurance (e.g., [Mitchell et al., 2005])



# On Infusing reachability-based safety assurance

[Leung, Schmerling, Chen, Talbot, Gerdes, Pavone, IJRR '20; Wang, Leung, Pavone IROS '20]


- BRS  $\mathcal{A}$ : set of *avoid states* for which there exists strategies for the human that for all inputs of the robot will lead to collisions

$$\mathcal{A} = \{x_{\text{rel}} : V(x_{\text{rel}}) \leq 0\}$$

↑ computed (offline) by solving HJI

- Optimal avoidance control ([Mitchell et al., 2005]):

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



# Backward reachability as a constraint

[Leung, Schmerling, Chen, Talbot, Gerdes, Pavone, IJRR '20; Wang, Leung, Pavone IROS '20]

- For less extreme behavior, reason in terms of **safety-preserving controls**

$$\mathcal{U}_R(x_{\text{rel}}) = \left\{ u_R : \min_{u_H} \nabla V(x_{\text{rel}}) \cdot f(x_{\text{rel}}, u_R, u_H) \geq 0 \right\}$$

- To solve at 100Hz, linearize:  $\mathcal{U}_R(x_{\text{rel}}) \approx \{u_R : M_{HJI} \cdot u_R + b_{HJI} \geq 0\}$

tracking error (vs. planner trajectory)

control effort

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

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

slack variable penalties

subject to

$$\delta_{k+1} - \delta_k = \Delta \delta_k, \quad \Delta \delta_{\min} \leq \Delta \delta_k \leq \Delta \delta_{\max}, \quad \delta_{\min} \leq \delta_k \leq \delta_{\max}$$

control constraints

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

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

stable handling envelope

$$H_k \begin{bmatrix} U_{y,k} \\ r_k \end{bmatrix} - G_k \leq \begin{bmatrix} \sigma_{\beta,k} \\ \sigma_{r,k} \end{bmatrix}$$

$$A_k q_k + B_k^- u_k + B_k^+ u_{k+1} + c_k = q_{k+1}$$

linearized vehicle dynamics

$$q_1 = q_{\text{curr}}, \quad u_1 = u_{\text{curr}},$$

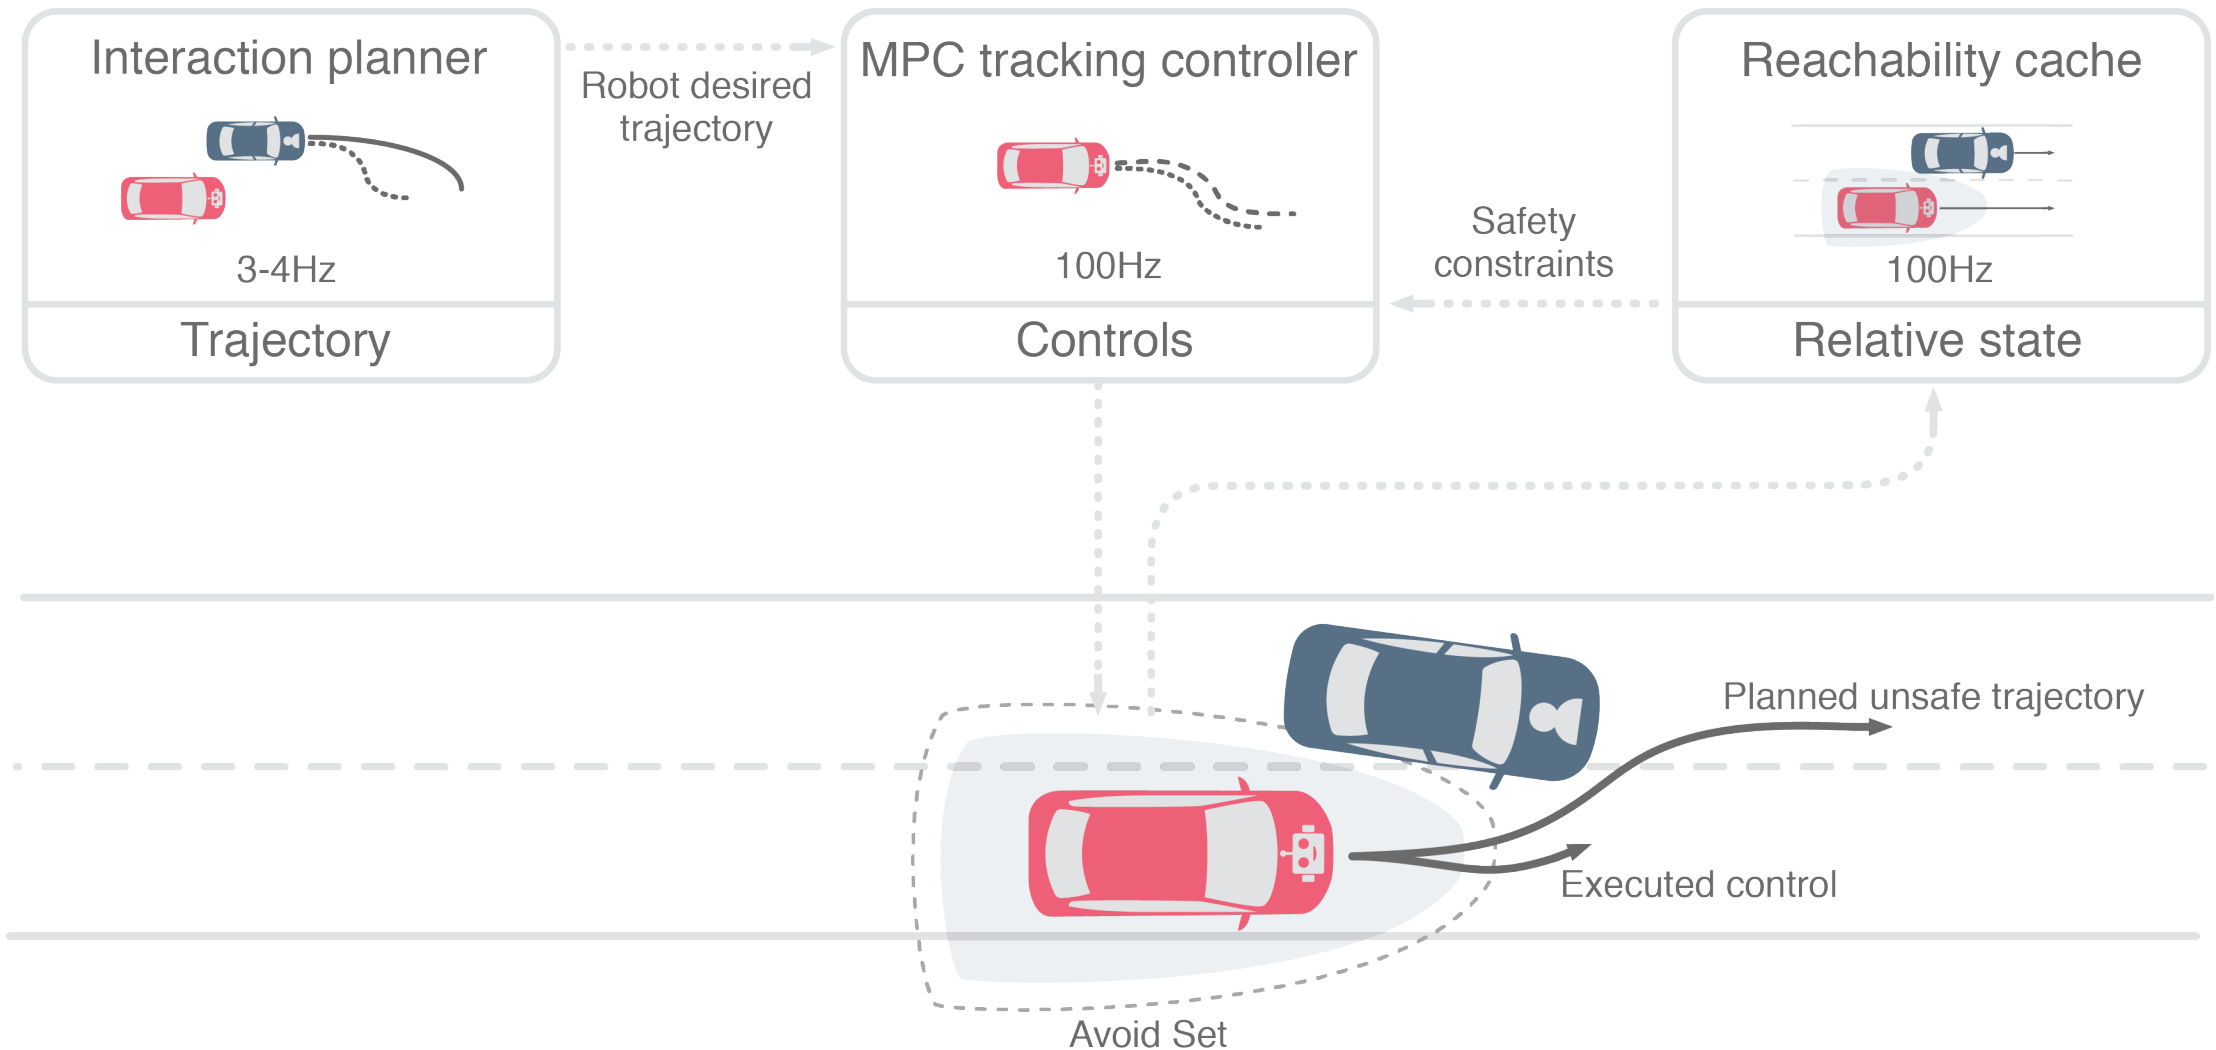
$$M_{HJI} u_j + b_{HJI} \geq -\sigma_{HJI}$$


instantaneous interaction safety constraint (new!)

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

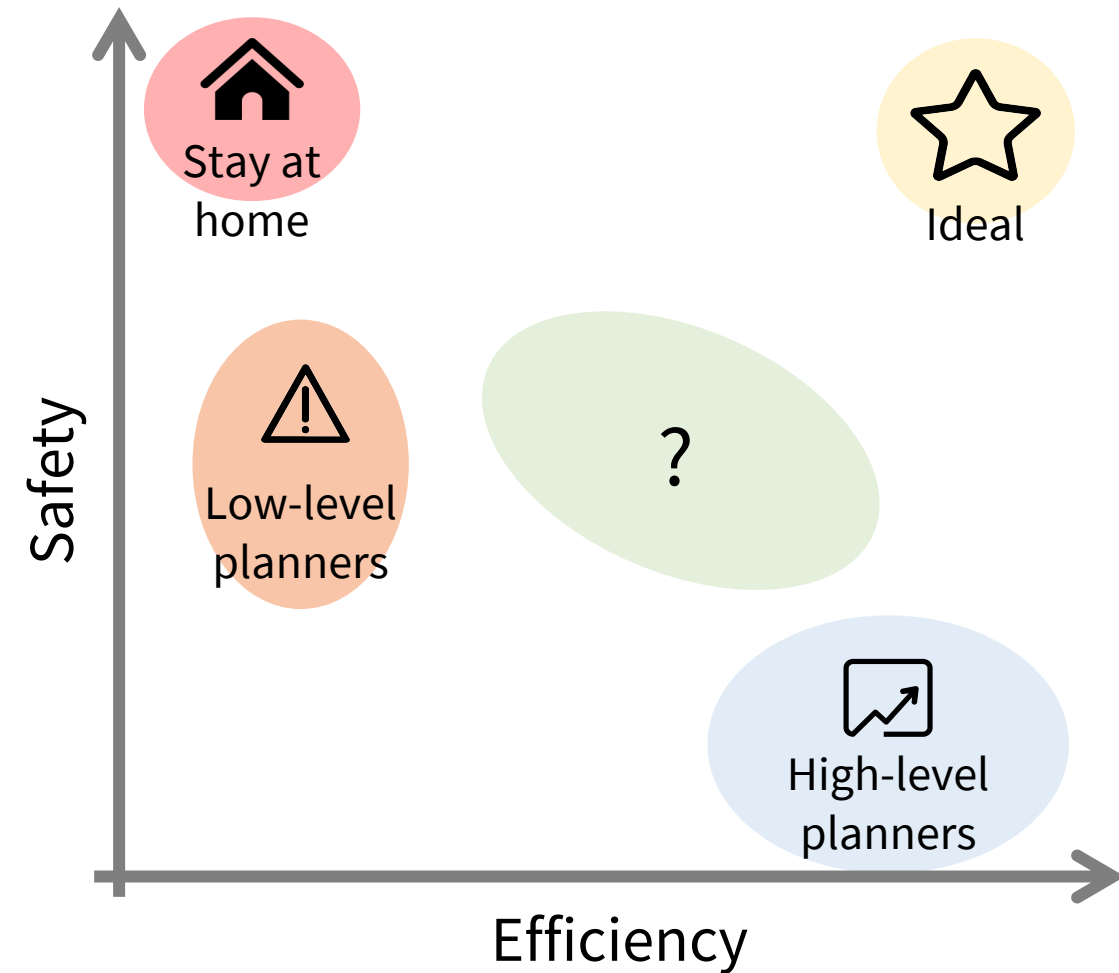
# Full decision-making and control stack

[Leung, Schmerling, Chen, Talbot, Gerdes, Pavone, IJRR '20]





# On the safety vs. efficiency trade-off



Takeaway message: safe and efficient  
interaction-aware decision making requires  
principled blending of **probabilistic and adversarial planning**

# Efficient representations for decision making

[Ivanovic, Elhafi, Rosman, Gaidon, Pavone, CoRL '20]

- Prediction of individual trajectories good for evaluation...
- ...but difficult to use in downstream decision making
- Idea: reason about **prediction representations cognizant of downstream control applications**


$$x^{(t+1)} = A^{(t)} x^{(t)} + B^{(t)} u_R^{(t)} + C^{(t)} w \quad (w \sim \mathcal{N}(0, 1))$$



# Integrating perception & trajectory forecasting


## Uncertainty Propagation

Uncertainty in perception should lead to uncertainty in prediction



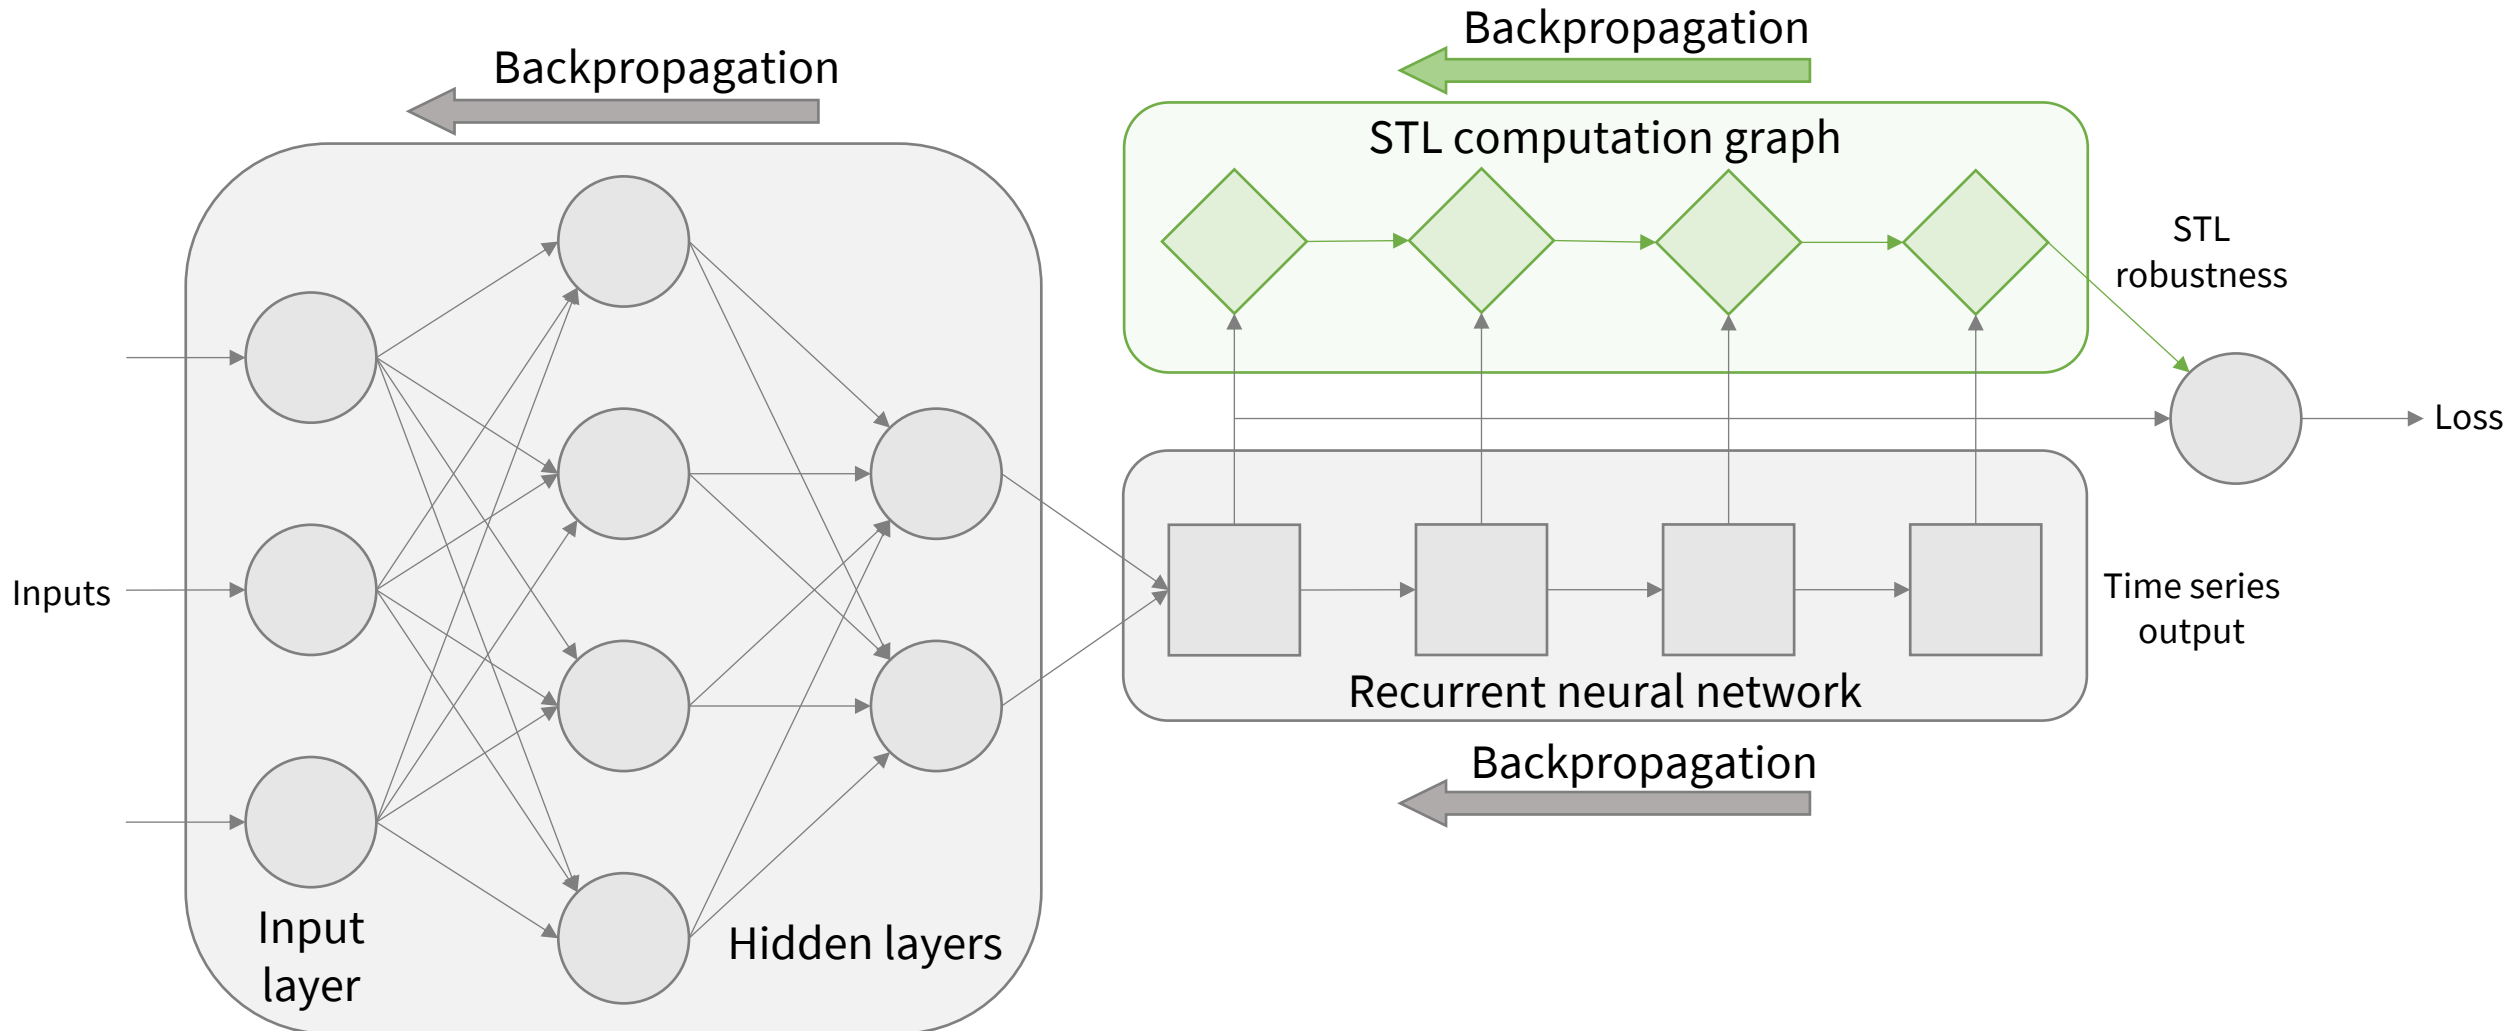
## Robustness to Perception

Object classification is noisy, prediction models should be robust to this



# Adding structure into predictions

[Leung, Arechiga, Pavone, IVS '19 & WAFR '20, patent pending]



Takeaway message: next-generation trajectory forecasting methods should account for **downstream control applications and structure**

# Conclusions

- Generative models becoming SoA tool for trajectory prediction

**Multimodal Deep Generative Models for Trajectory Prediction:  
A Conditional Variational Autoencoder Approach**



Boris Ivanovic<sup>1\*</sup>, Karen Leung<sup>1\*</sup>, Edward Schmerling<sup>2</sup>, Marco Pavone<sup>1</sup>

Also: <http://ai.stanford.edu/blog/trajectory-forecasting/>

- ...but quite difficult to integrate within autonomy stack
  1. reason about prediction representations **cognizant of downstream control applications**
  2. tighter **integration** between perception and trajectory prediction
  3. Add **logical structure** to models

All code available at: <https://github.com/StanfordASL>

# Acknowledgements



[pavone@stanford.edu](mailto:pavone@stanford.edu)

<http://asl.stanford.edu/>

<https://twitter.com/StanfordASL>



*NORTHROP GRUMMAN*

