

Inference and Prediction Insights for Safe Vehicle-Pedestrian Interaction




HUMAN-CENTERED
AUTONOMY LAB

Katie Driggs-Campbell

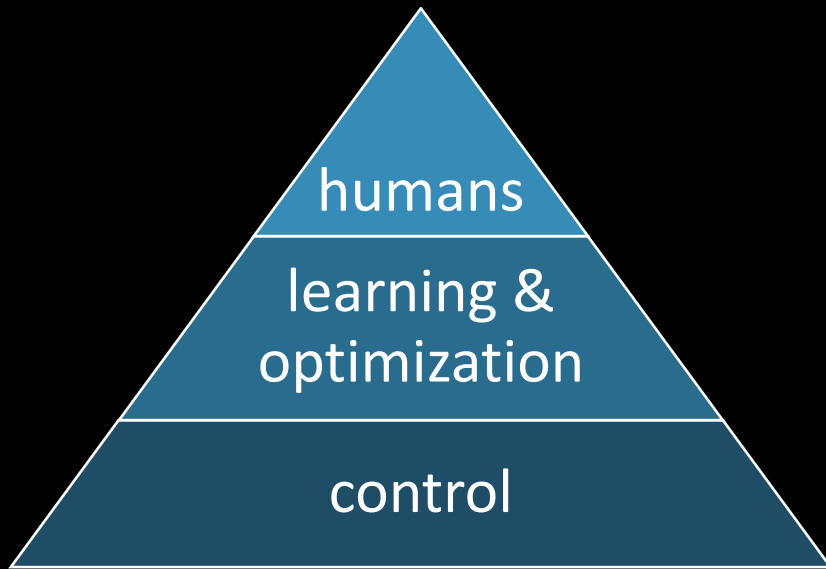
Assistant Professor

Dept. of Electrical and Computer Engineering

University of Illinois at Urbana-Champaign



How can we ensure safety in autonomous systems that operate with people in the real-world?



1.5x speed



Robust Decision-Making and Control

We wish to safely and efficiently control a vehicle,
despite uncertainty and disturbances:



$$x[k + 1] = f(x[k], u[k]) + v(x[k], u[k], d[k])$$

Robust Decision-Making and Control


We wish to safely and efficiently control a vehicle,
despite uncertainty and disturbances:

$$x[k + 1] = f(x[k], u[k]) + v(x[k], u[k], d[k])$$

Sensor Noise & Environment Uncertainty



Model Mismatch


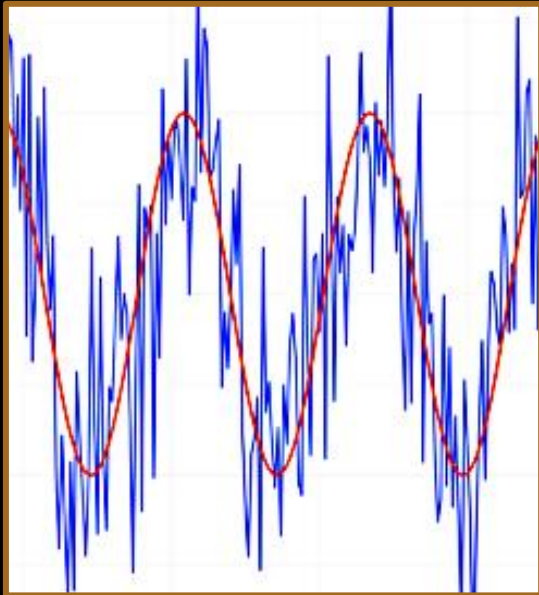


Robust Decision-Making and Control


We wish to safely and efficiently control a vehicle,
despite uncertainty and disturbances:

$$x[k + 1] = f(x[k], u[k]) + v(x[k], u[k], d[k])$$

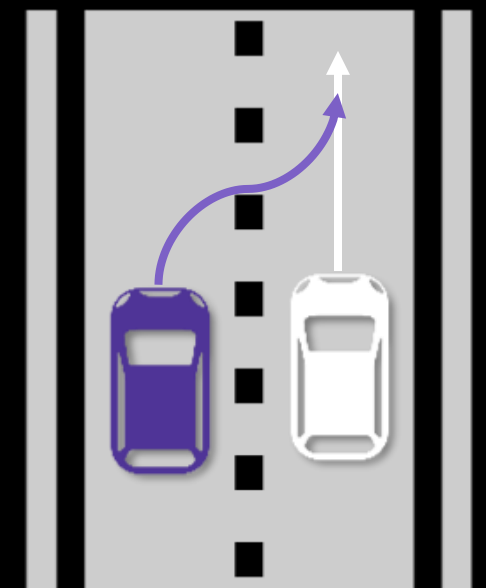
Sensor Noise & Environment Uncertainty



Model Mismatch

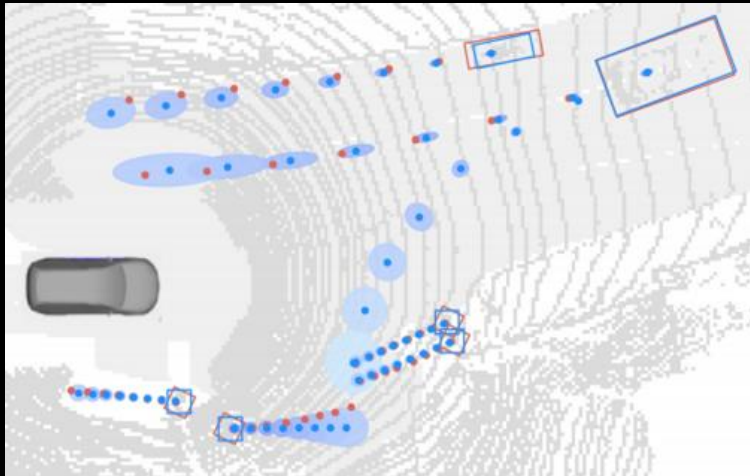



Long-Term Interactions




Are pedestrian-vehicle interactions any different?

Simulations and datasets are not readily available.

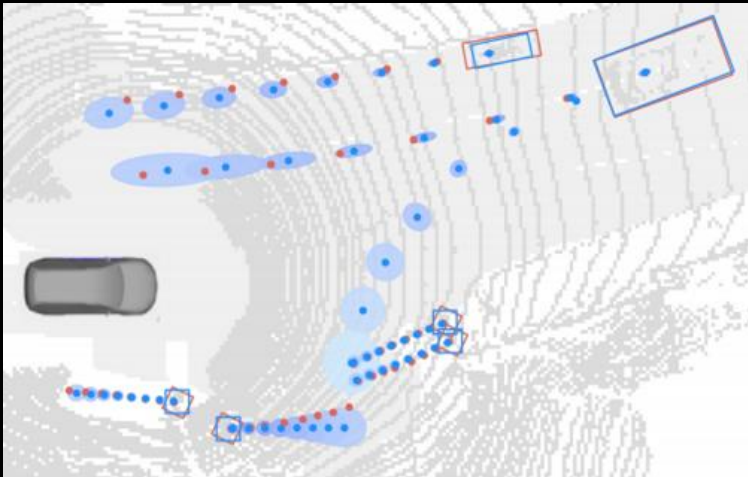


Are pedestrian-vehicle interactions any different?

Simulations and datasets are not readily available.





Experiments are a bit of a production.

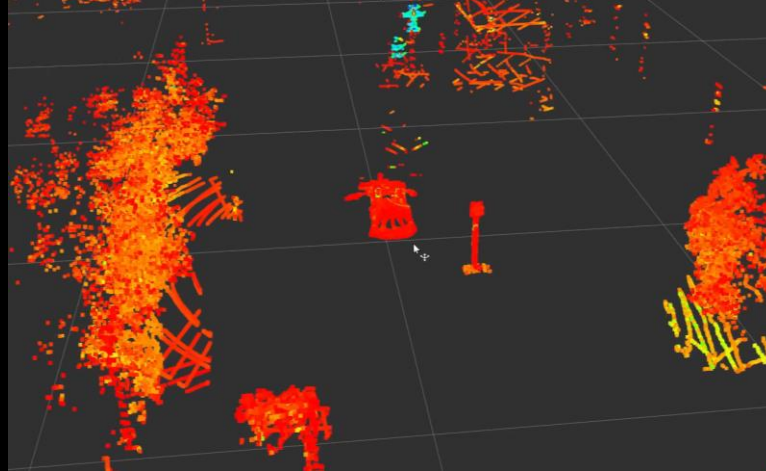


Are pedestrian-vehicle interactions any different?

Simulations and datasets are not readily available.



Experiments are a bit of a production.



In validation, pedestrians are the most challenging.

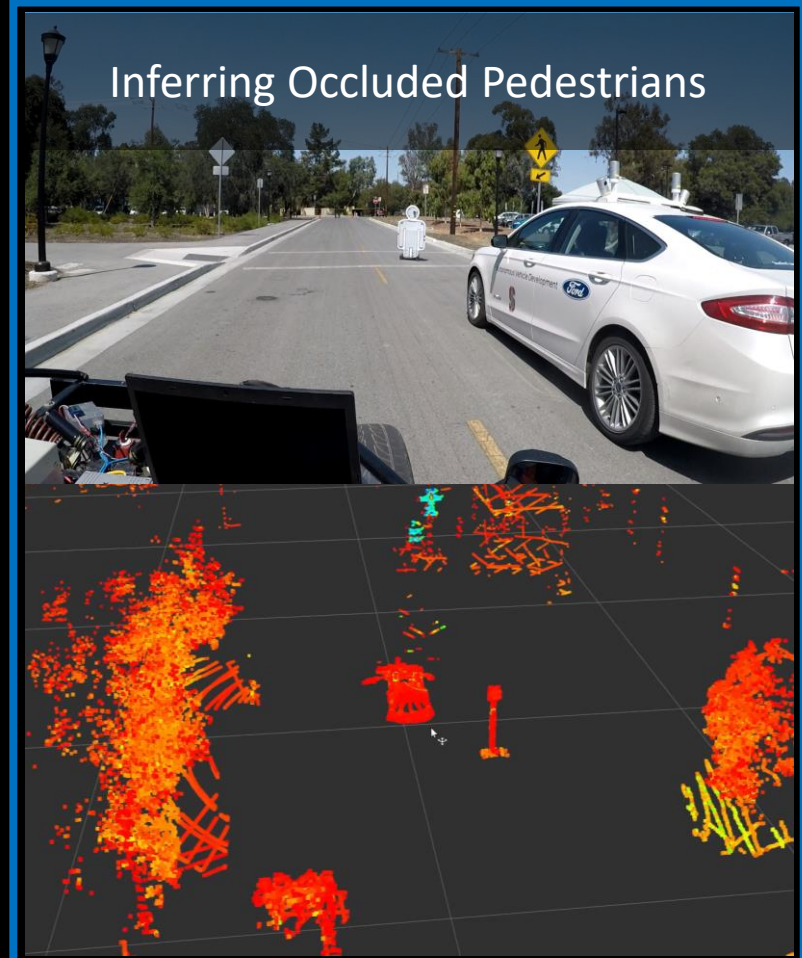


Today's Destinations

Effective Pedestrian Prediction Methods for Safe Vehicle-Pedestrian Interactions



Inferring Occluded Pedestrians



Pedestrian Prediction Methods

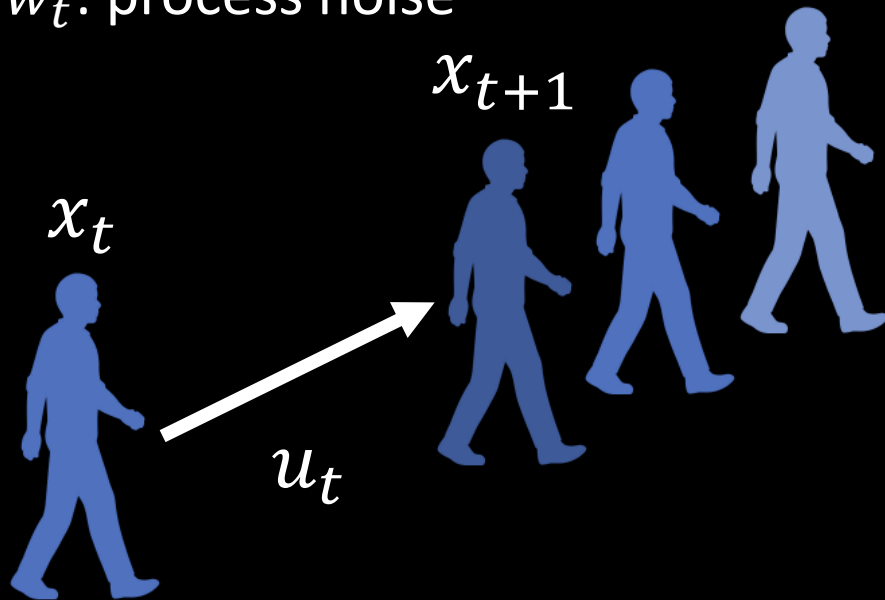
Physics-based Methods

$$\dot{x}_t = f(x_t, u_t, t) + w_t$$

x_t : human state

u_t : control input

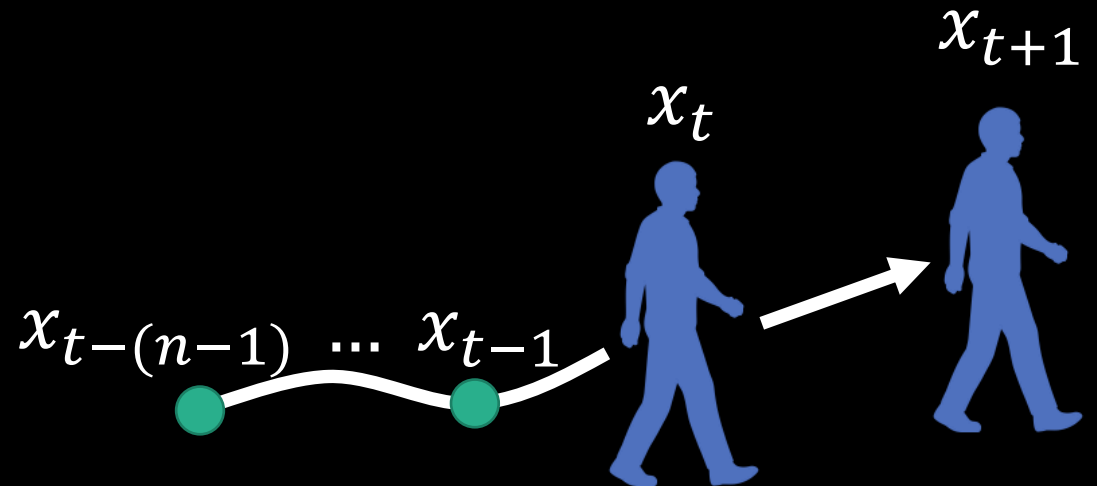
w_t : process noise




Pattern-based Methods

$$P[x_{t+1}|x_{1:t}] = P[x_{t+1}|x_{t-(n-1):t}]$$

$$x_{t+1} = f(x_{t-(n-1):t})$$



Pedestrian Prediction Methods

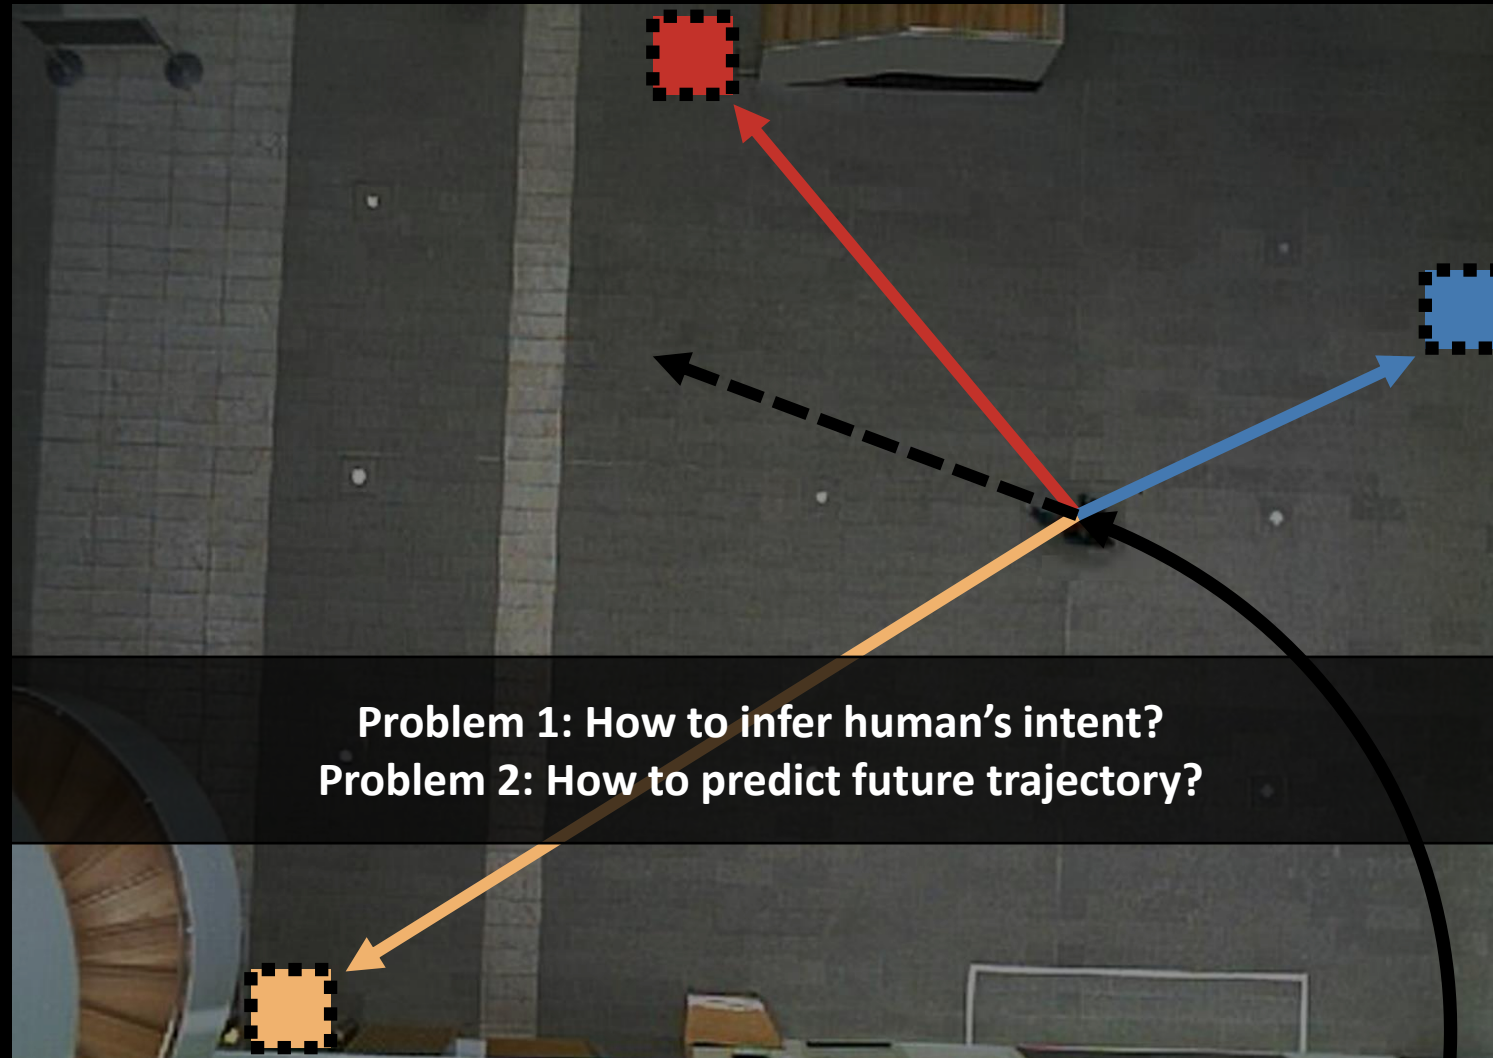


Trajectory Forecasting Benchmark

Method (rank)	Final Disp.	Mean Disp.
Vanilla LSTM (1)	1.191	0.355
Physical Comp (6)	1.229	0.366
Social Forces (10)	1.266	0.371

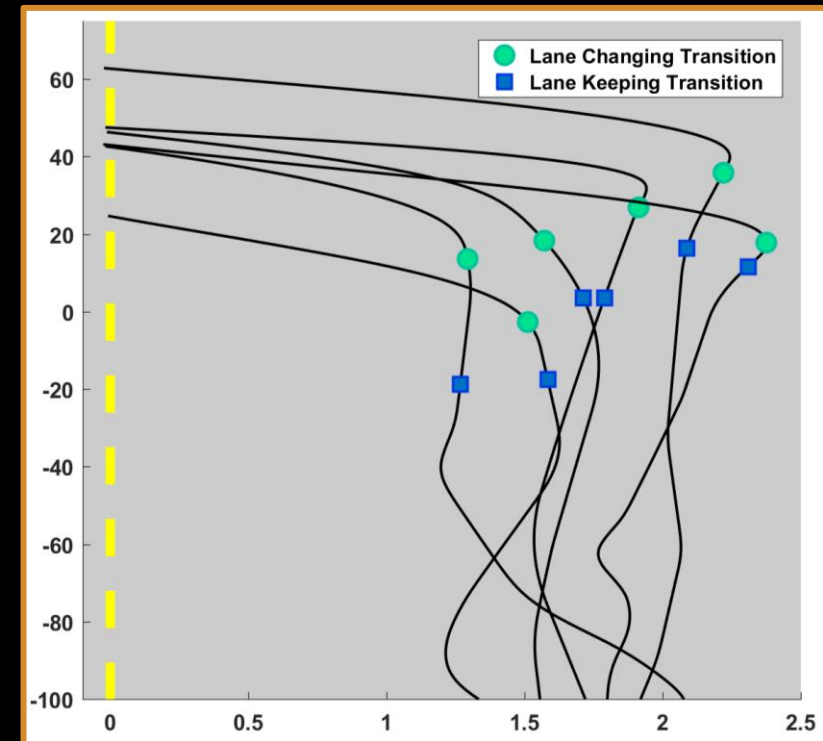


Intention-Aware Predictions

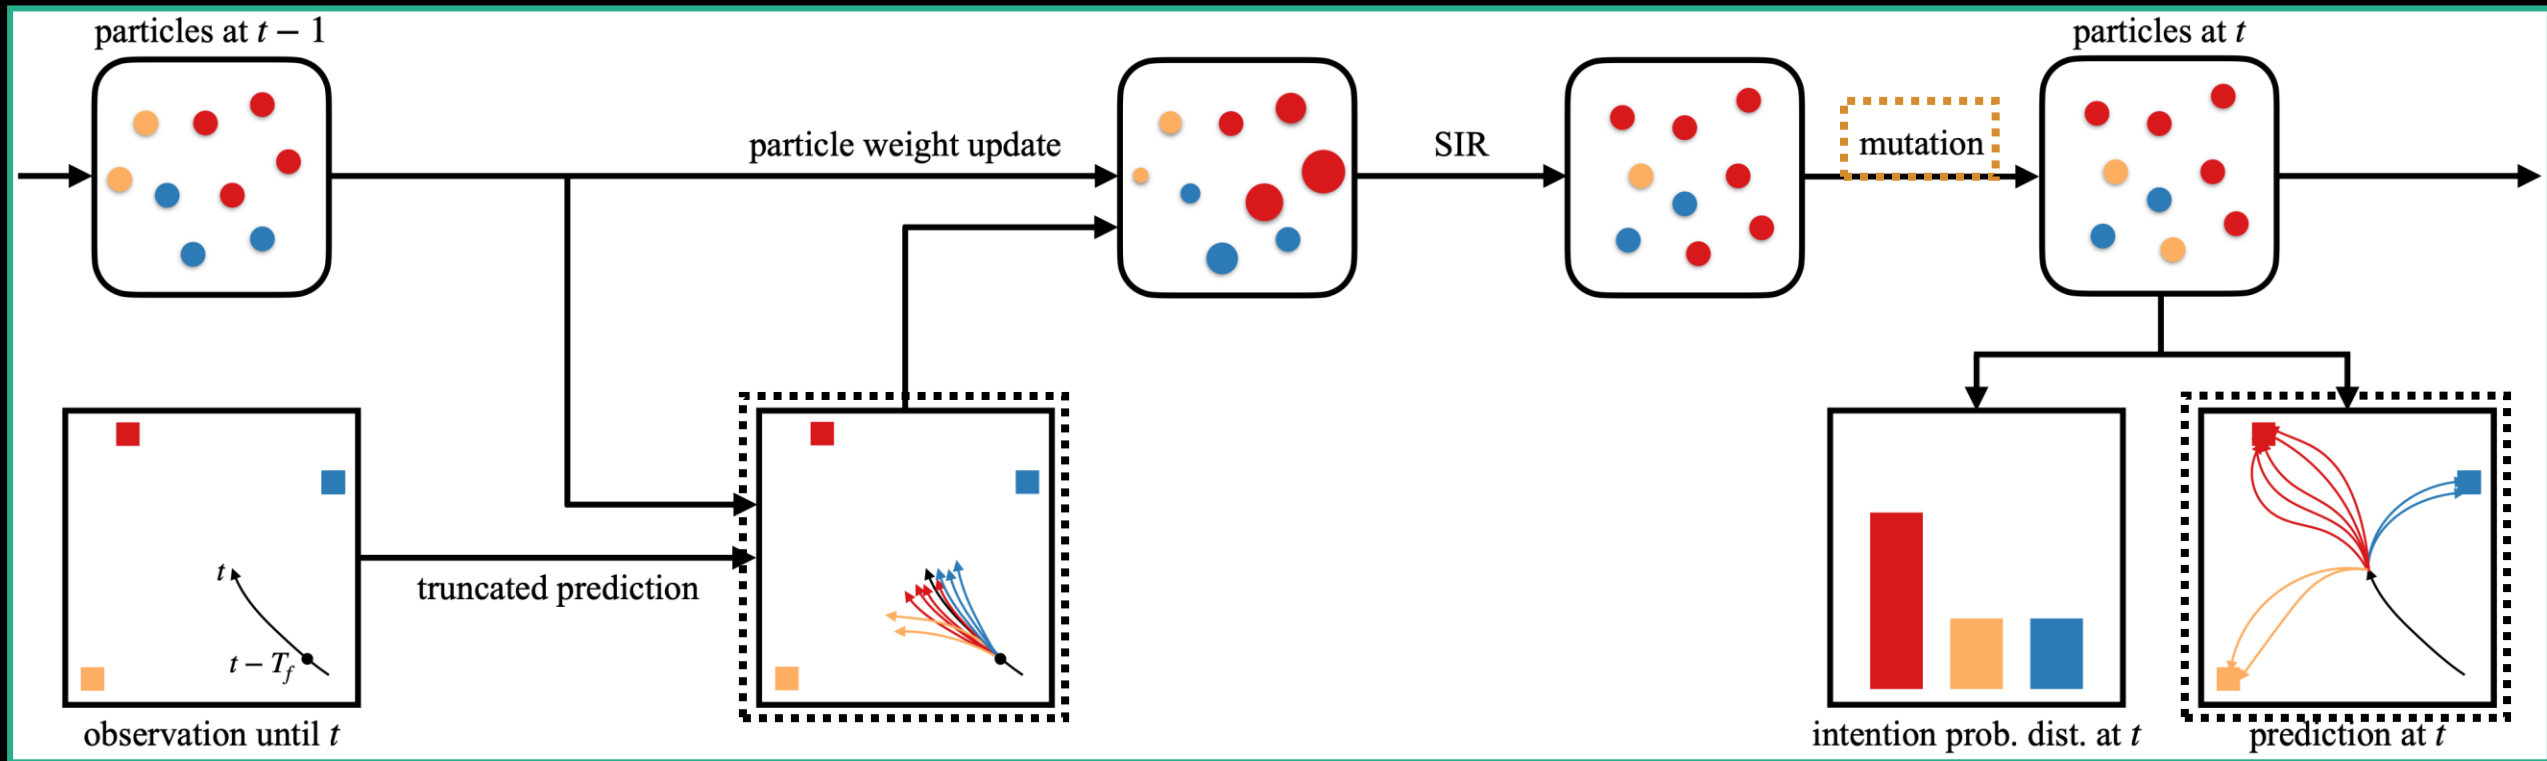


Intent-Guided Prediction

- When we observe motion, we usually care very little about the surface behaviors
 - *Intentions* determine how we understand, recall, react, and predict
- When observing continuous motion, humans often agree where **boundaries separating distinct actions** lie, corresponding to intent
- In human motion prediction, incorporating intent (or goals) significantly improves predictions


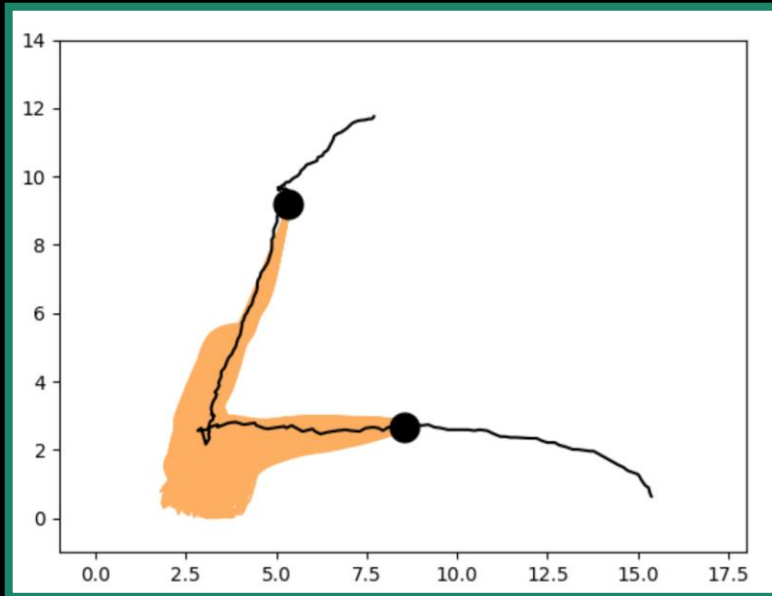


Intention-Aware Predictions


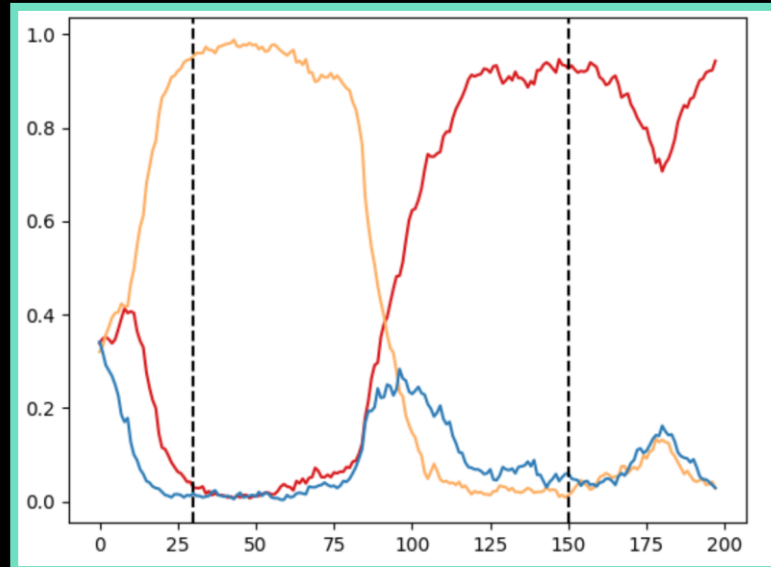



Filter Comparisons

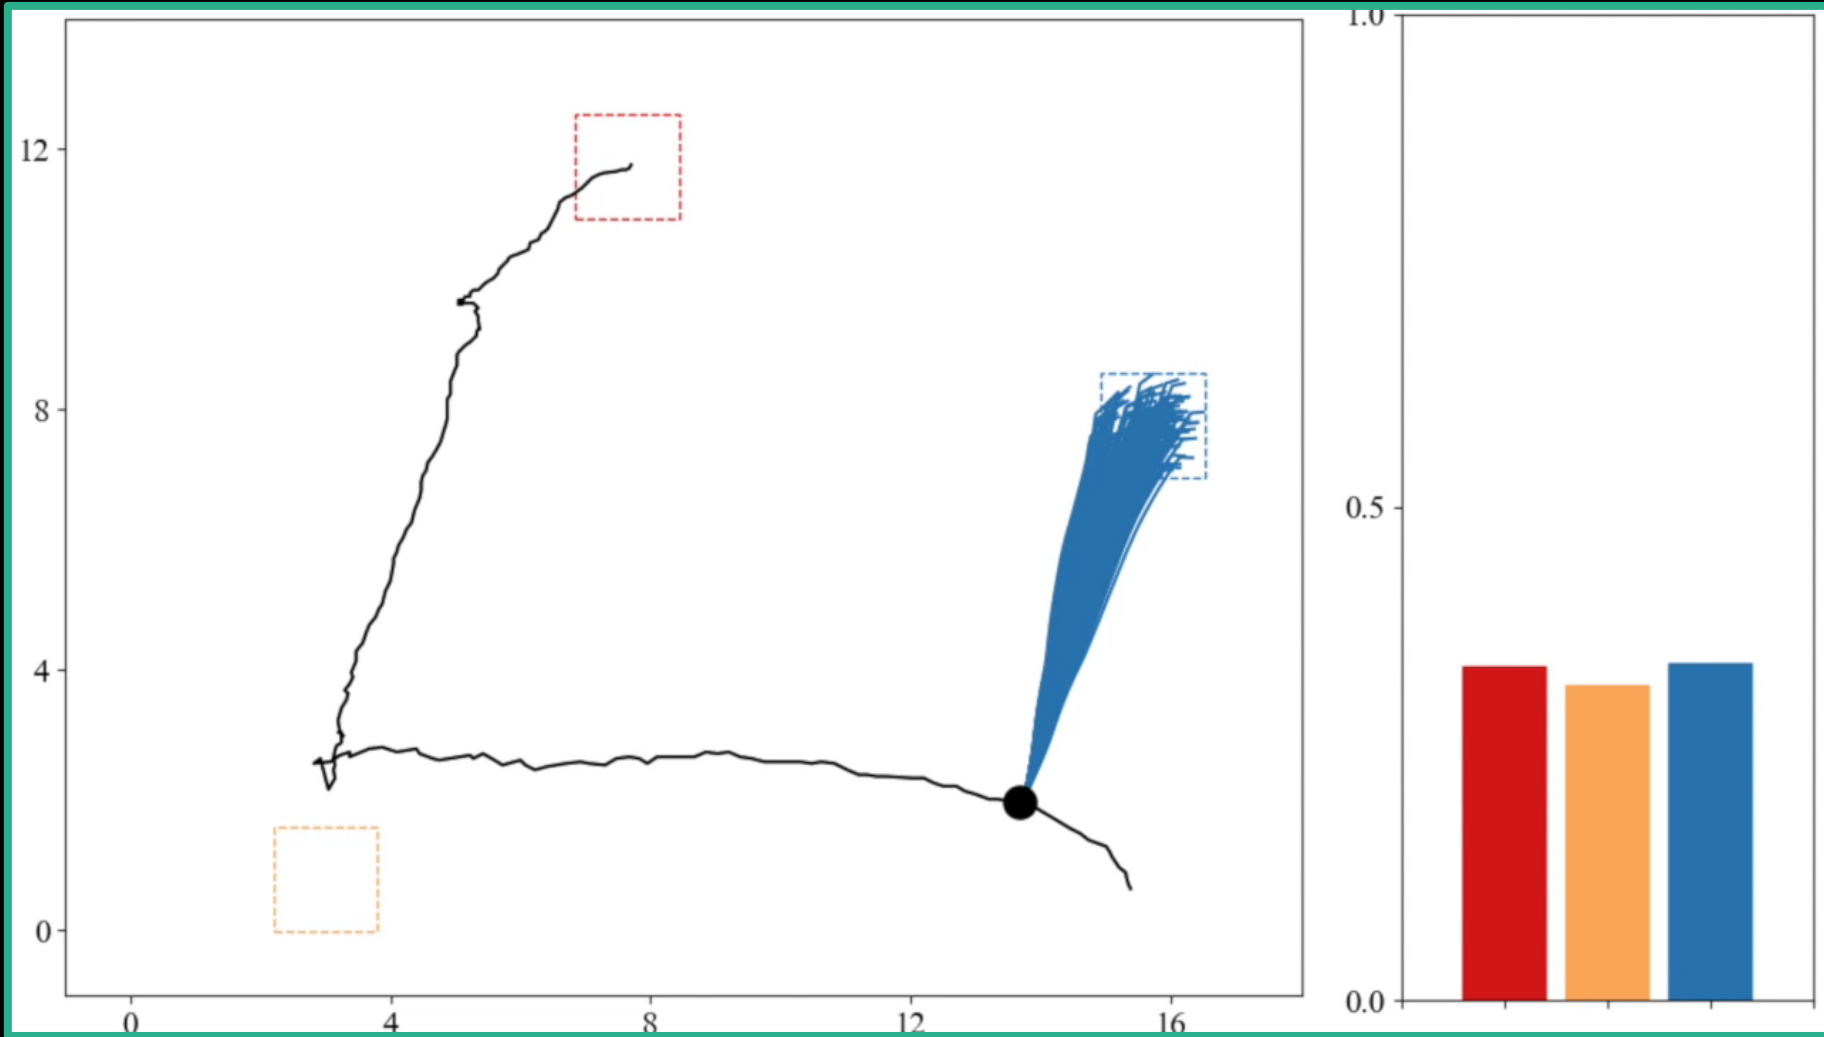
Prediction Samples




Intent Prob. Distribution



Mutable Intention Filter



Intention-Aware Predictions



Nominal pedestrian predictions are modeled using the **Generalized Potential Field Approach**:


- A set of **sources represent obstacles** in the map and generate a repulsive force
- A set of **sinks represent goal locations** and pose an attractive force on the pedestrian

The sum of the forces act as a control input to a **linear state space model** of the pedestrian:

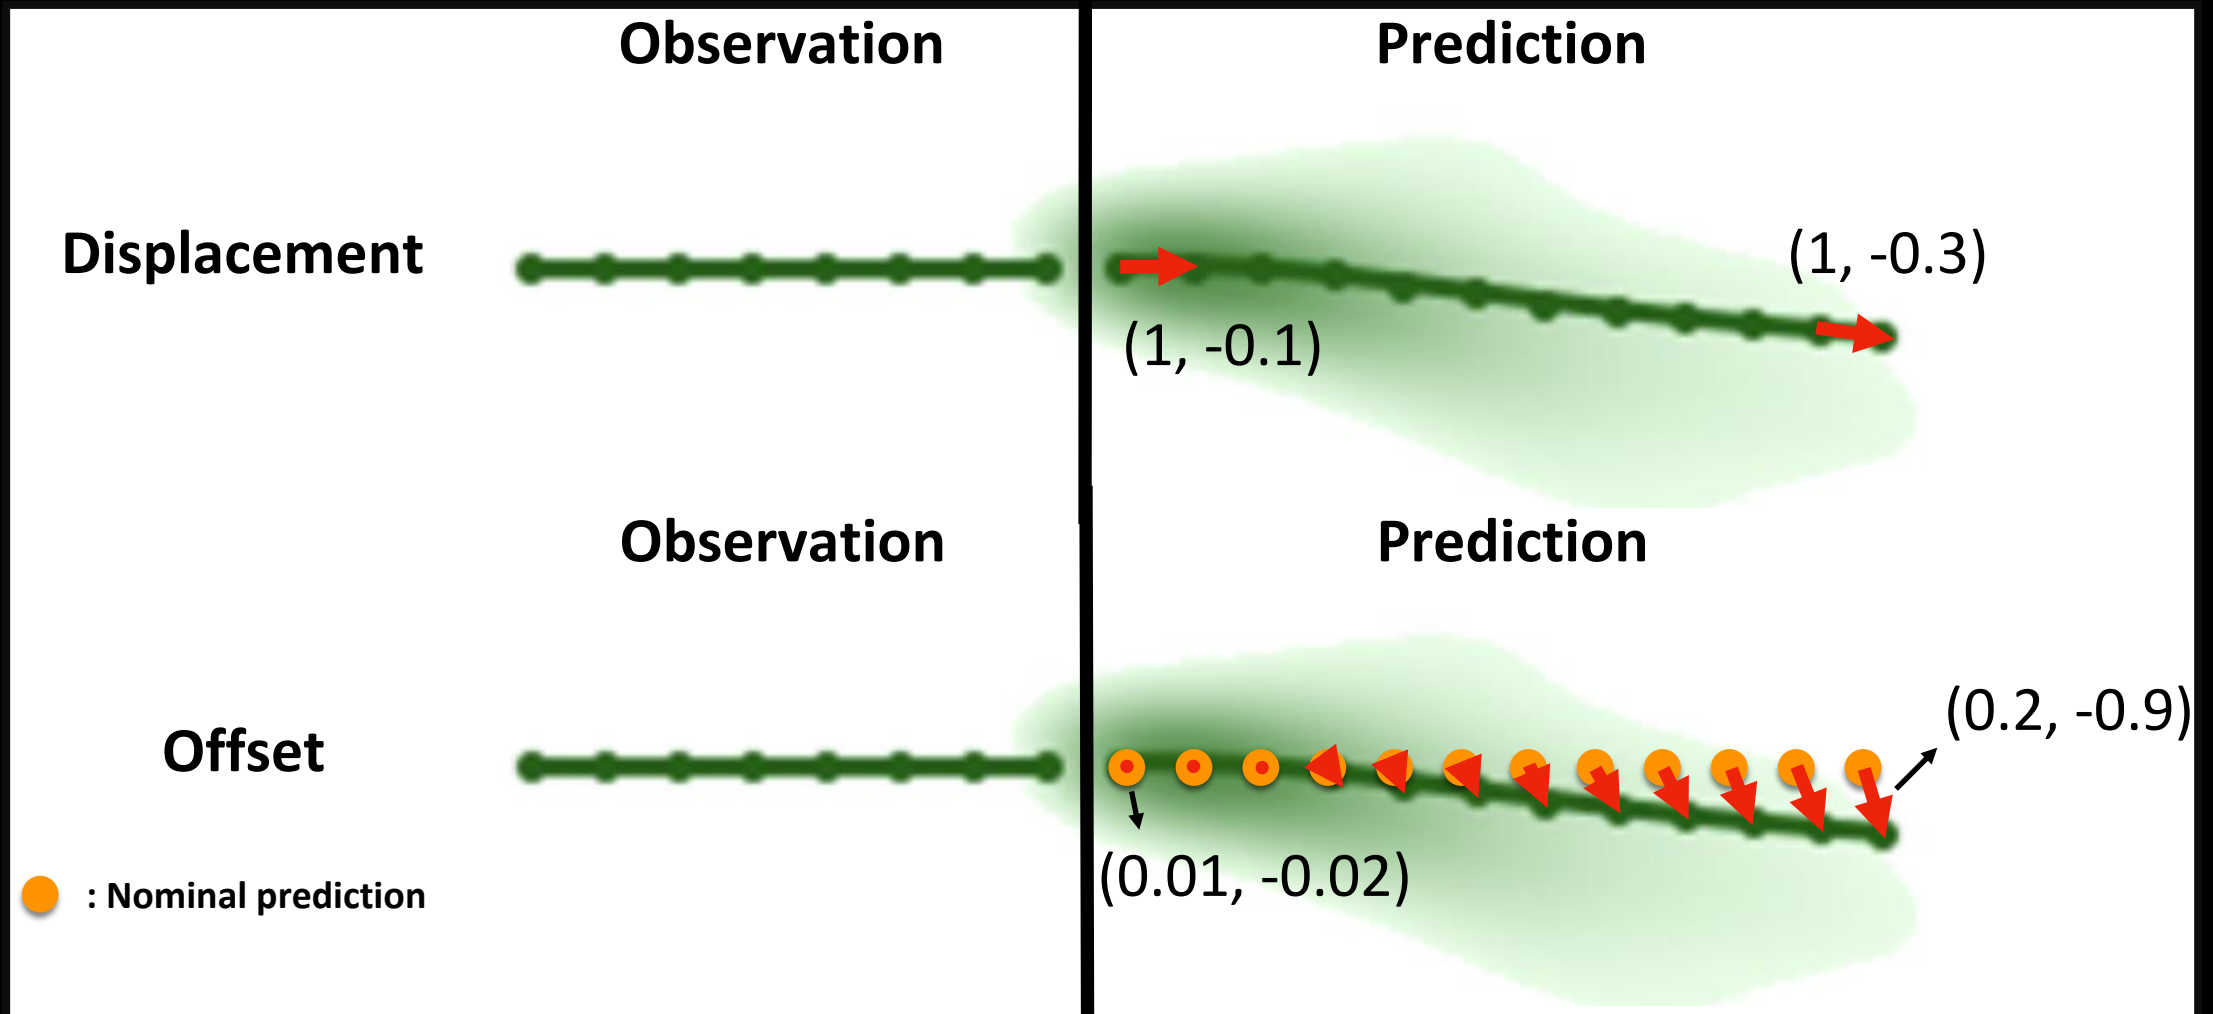
$$\begin{aligned}x_t &= Fx_{t-1} + Gu_{t-1} + w_t \\y_t &= Hx_{t-1} + v_t\end{aligned}$$

where x_t is the pedestrian state, y_t is the measurement, w_t and v_t are Gaussian white noise tuned to match real-world data


Intention-Aware Predictions



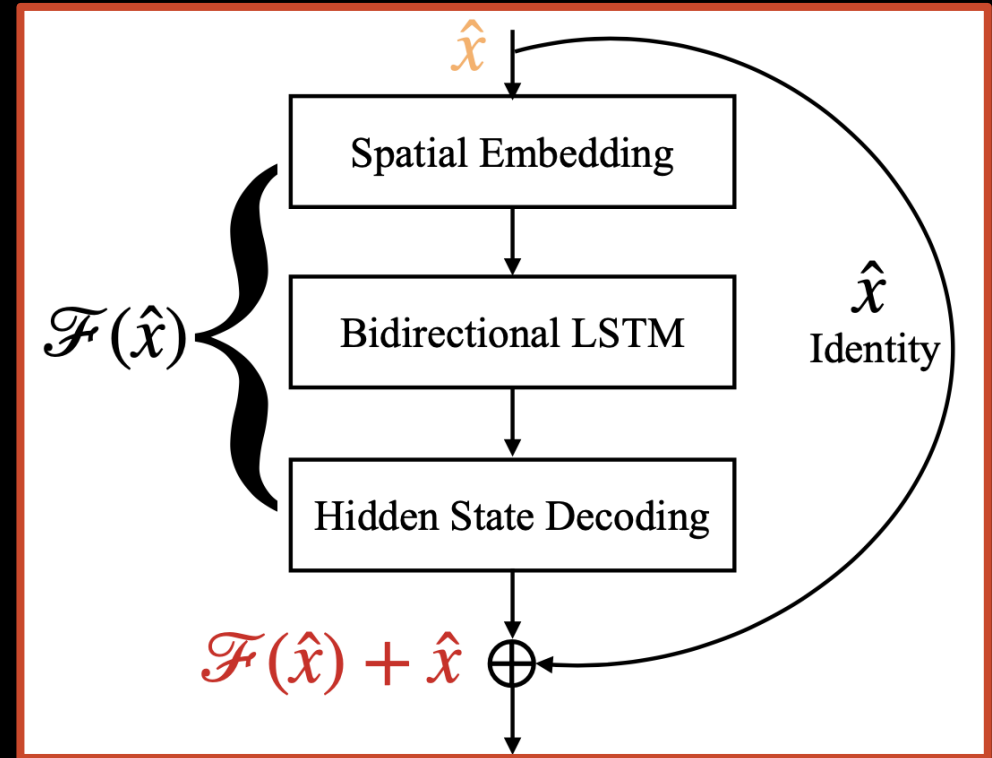
Trajectory Warping with Residuals




Residual Method: Warping LSTM



$$\hat{x}_{1:t+T_g} = [x_{1:t} \tilde{x}_{t+1:t+T_g}]$$



Intention-Aware Predictions



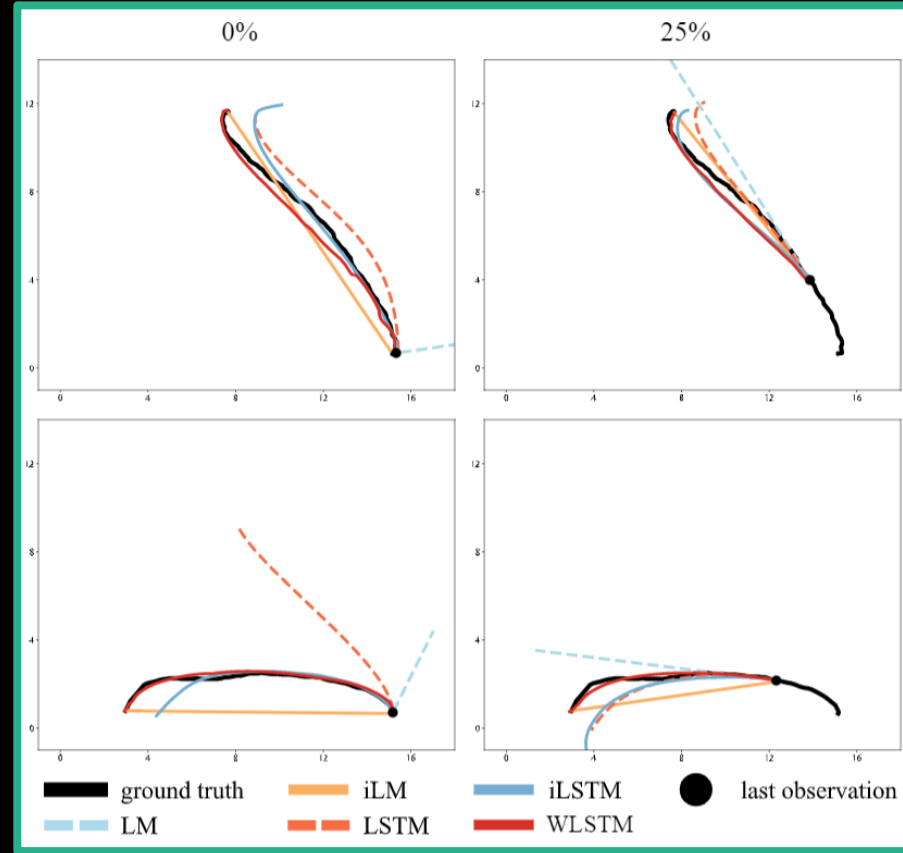


Prediction Results

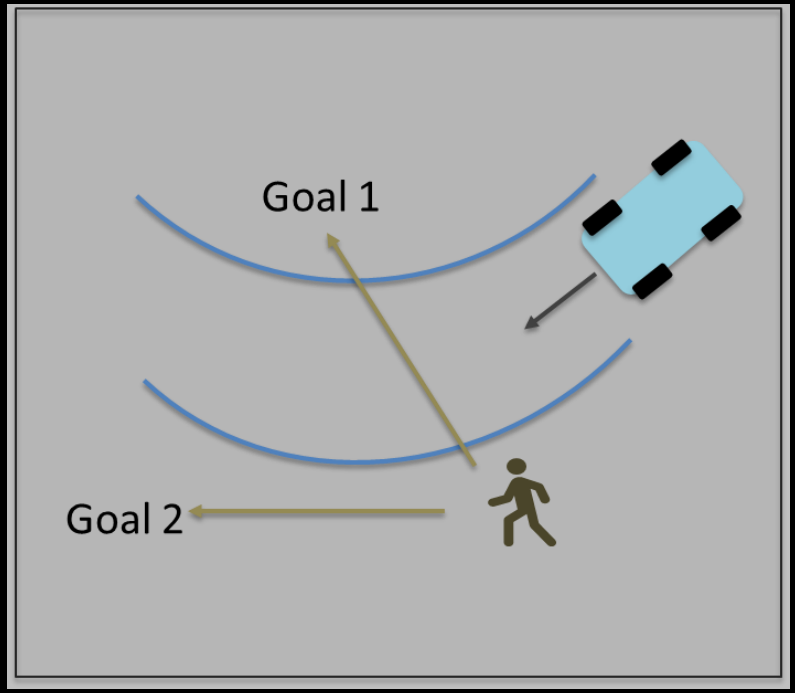
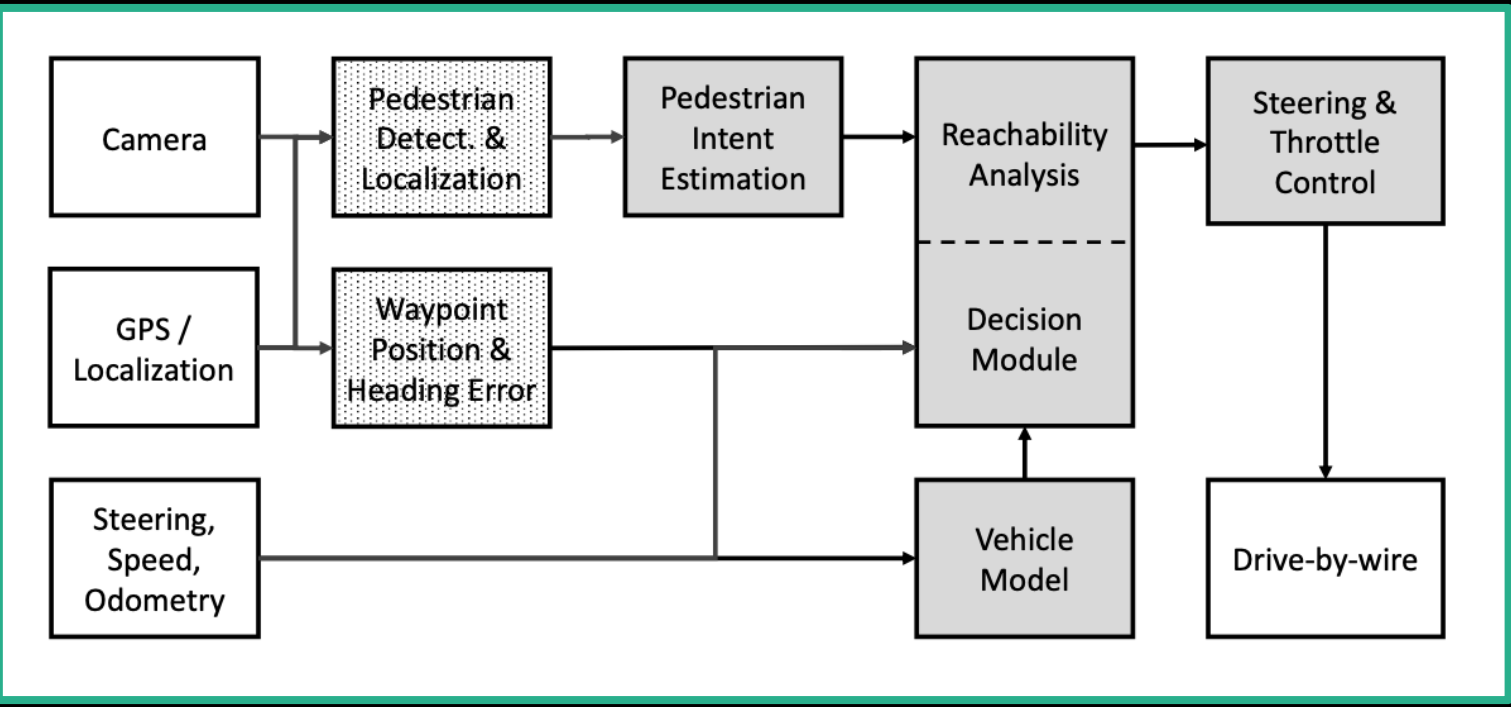
Effectively warps the nominal prediction to capture human tendencies

Simultaneously estimating intent with prediction provides improved performance and bonus utility

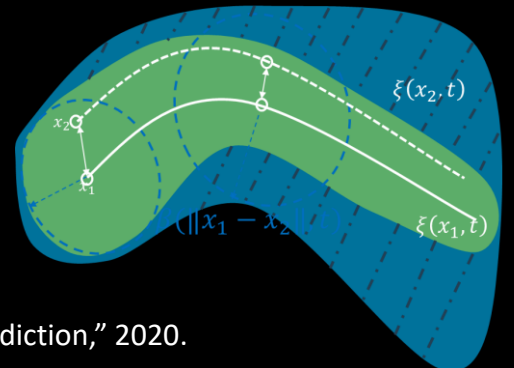


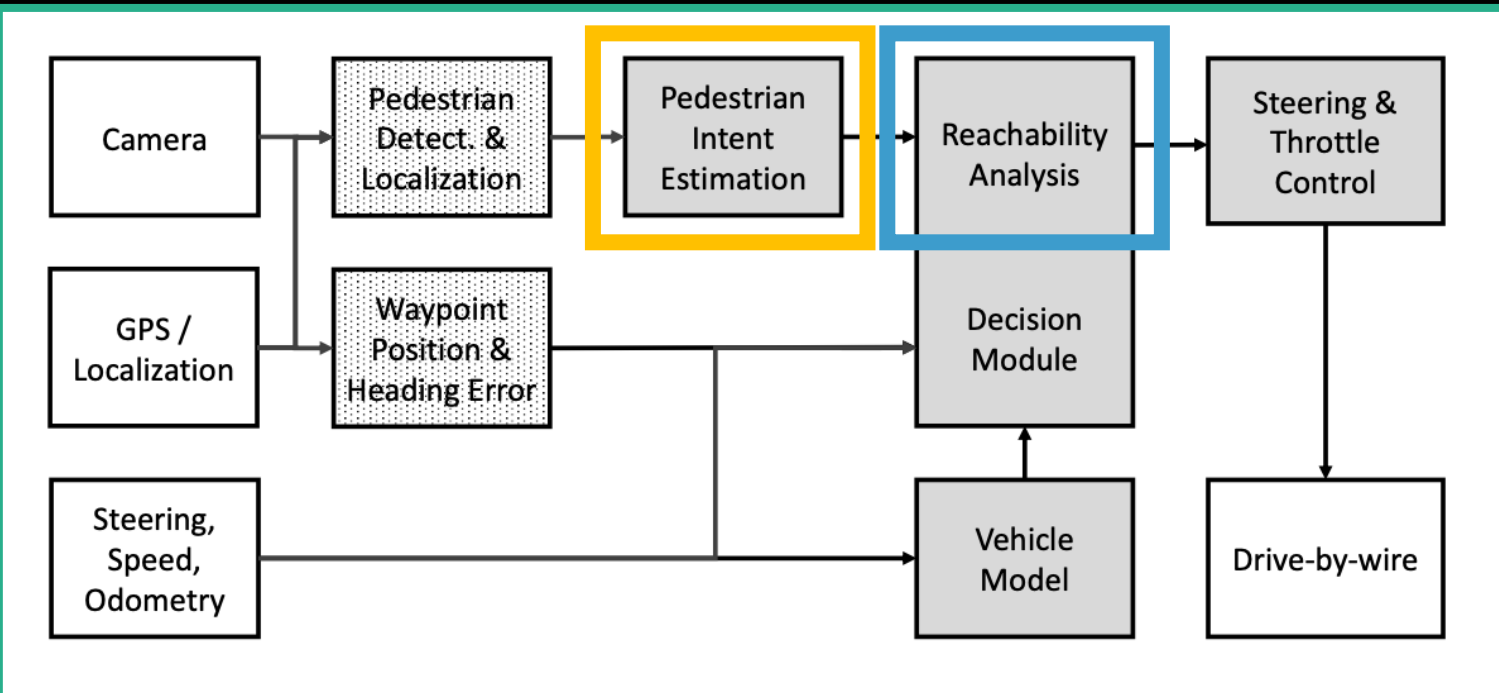
Method	AOE	FOE
SF	3.124	3.909
LSTM	2.132	3.005
SLSTM	1.524	2.510
ALSTM	0.986	1.311
SGAN	1.042	2.088
MIF-WLSTM ₁	0.665	1.236
MIF-WLSTM ₄	0.636	1.179



Online Monitoring for Vehicle-Pedestrian Interaction




Online Monitoring for Vehicle-Pedestrian Interaction



Z. Huang, A. Hasan, and K. Driggs-Campbell. "Intention-aware Residual Bidirectional LSTM for Long-term Pedestrian Trajectory Prediction," 2020.
 S. Duggirala, S. Mitra, and M. Viswanathan, "Verification of annotated models from executions," EMSOFT 2013.
 C. Fan, B. Qi, and S. Mitra, "Data-driven formal reasoning and their applications in safety analysis of vehicle autonomy features," IEEE Design & Test, 2018.

Online Reachability with DryVR

- Compute the reachable set $Reach_A(\Theta, p, T)$, given system A with state space X , mode $p \in P$, set of initial states Θ , and look-ahead time T
- DryVR uses a sensitivity function β to bound distance between trajectories
- Θ is partitioned into regions from which a numerical simulation $\xi(x_i, t)$ is computed for T time
- An over-approximation of the reachable set $Reach_A(\Theta, p, T_{Look})$ is obtained by bloating $\xi(x_i, t)$ with β and then taking a union

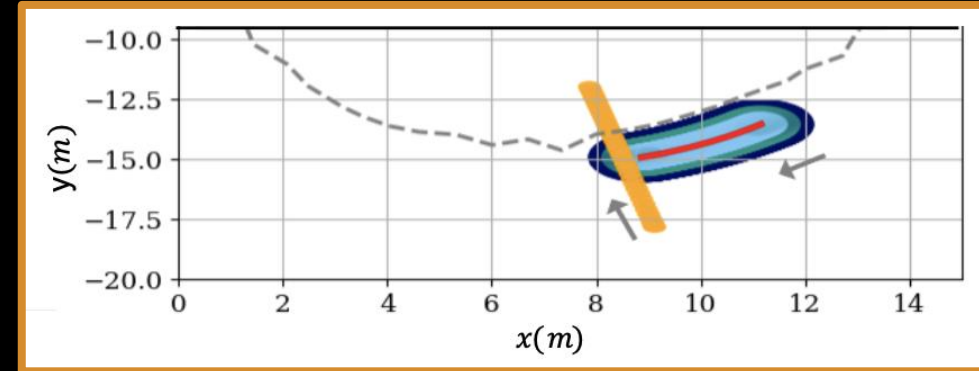



$\dot{x} = v \cos \theta$, $\dot{y} = v \sin \theta$, $\dot{\phi} = u$, $\dot{v} = a$, $\dot{\theta} = \frac{v}{L} \tan \phi$
 where (x, y) is the position, v is the speed, θ is the steering angle, ϕ is the heading angle, and L is the length of the vehicle

T_{Look} (s)	3.0	3.5	4.0	4.5	5.0
Compute Time (s)	0.096	0.103	0.129	0.136	0.163


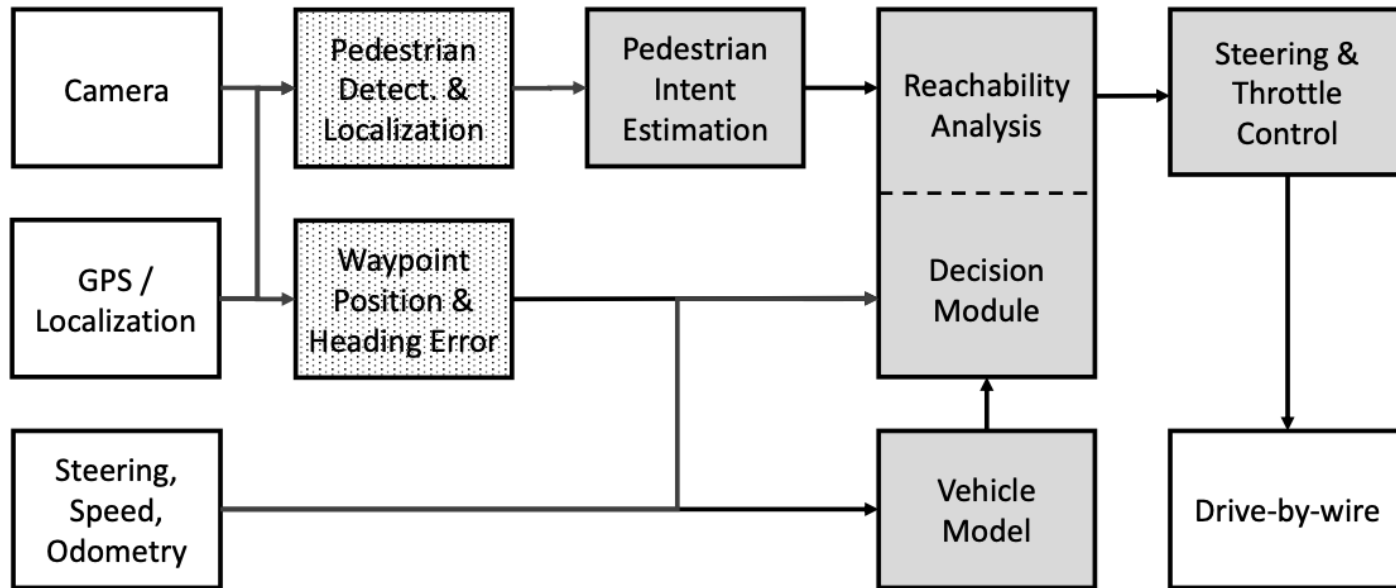
T_{Look} (s)	3.0	3.5	4.0	4.5	5.0
High Conf. (%)	98.905	98.338	98.617	97.825	96.851
Med Conf. (%)	98.161	96.629	97.138	95.917	94.874
Low Conf. (%)	95.825	94.078	94.672	93.416	92.078

Online Monitoring for Vehicle-Pedestrian Interaction



- Challenges**
1. Perception for non-experts remains a challenge
 2. Reachability (and often decision-making and control) assumes single (most likely) trajectory prediction
 3. Scaling to multiple agents adds complexity

Online Monitoring for Vehicle-Pedestrian Interaction



- Challenges**
1. Perception for non-experts remains a challenge
 2. Reachability (and often decision-making and control) assumes single (most likely) trajectory prediction
 3. Scaling to multiple agents adds complexity

Today's Destinations

Effective Pedestrian Prediction Methods for Safe Vehicle-Pedestrian Interactions



Inferring Occluded Pedestrians



www.wash.com.ua



965

2015/09/31

10:46:27




Can we make improved inferences about the state of the **environment** by observing other agents' behaviors?

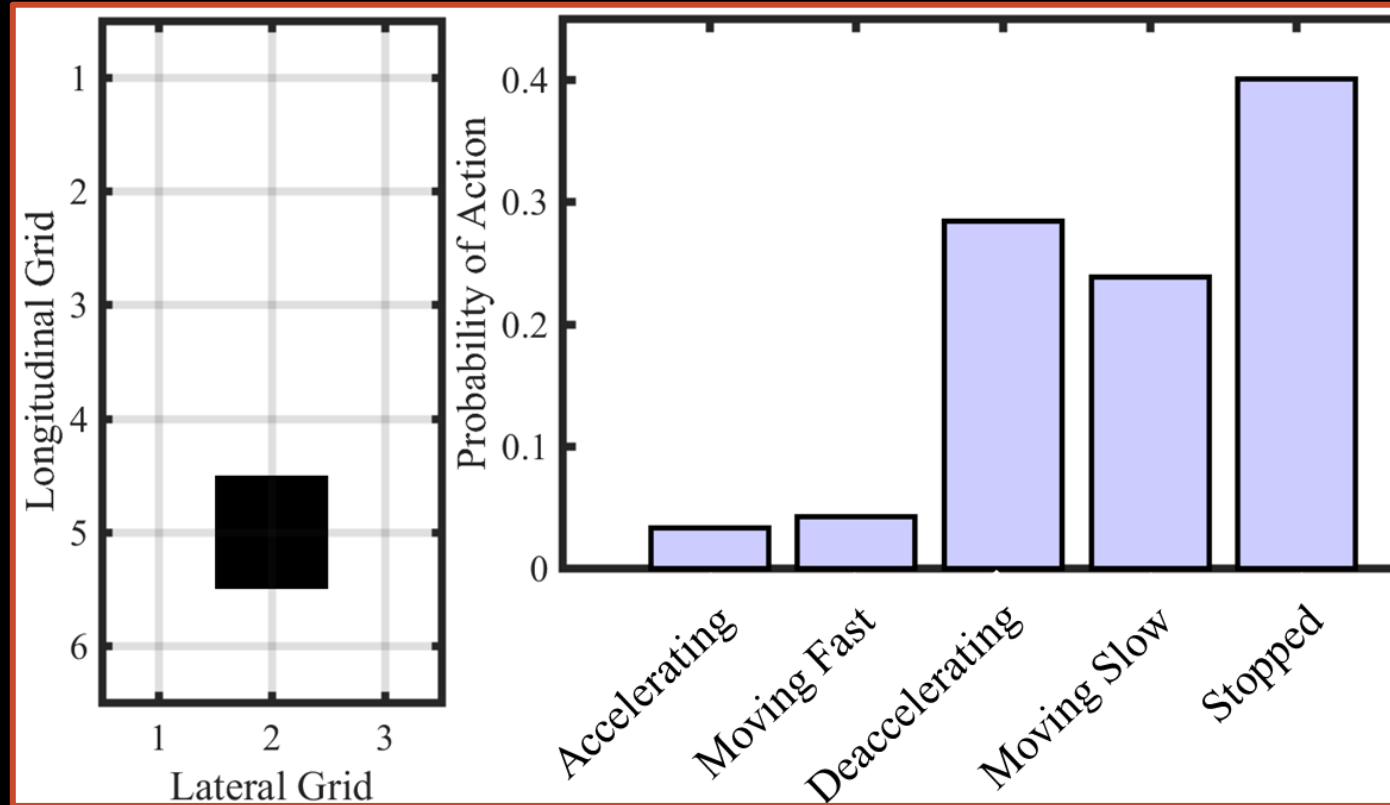
Imputing Maps from Human Actions

Consider the occluded region as a map we must estimate.


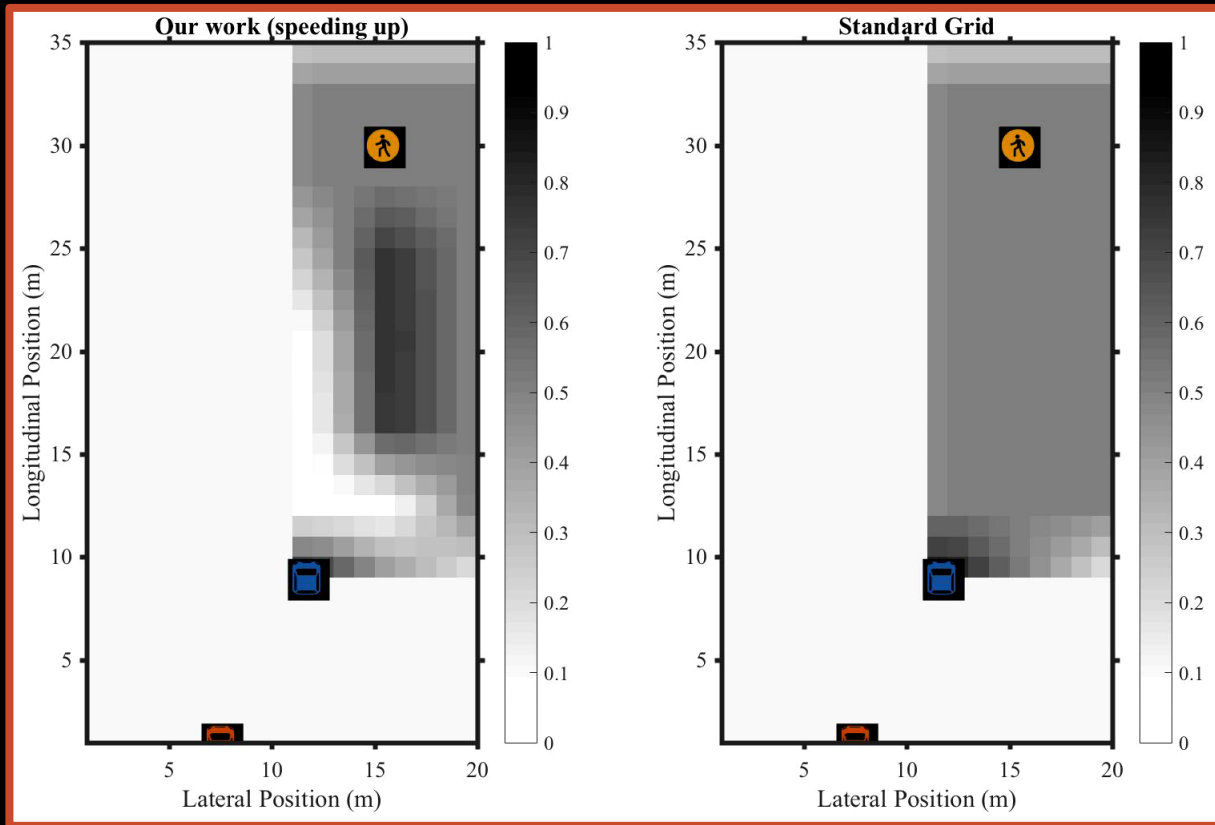
Occupancy Grid Approach: $p_t(m|x_{1:t}, z_{1:t}) = \prod_{i=1}^n p_t(\mathbf{m}_i|x_t, z_t)$





People as Sensors





Imputing Maps from Human Actions



Real-World Tests and Additional Complexities



Acknowledgements



Today's Destinations

Effective Pedestrian Prediction Methods for Safe Vehicle-Pedestrian Interactions



Inferring Occluded Pedestrians

