Humans Should not Be Obstacles

Anca Dragan

anca@berkeley.edu

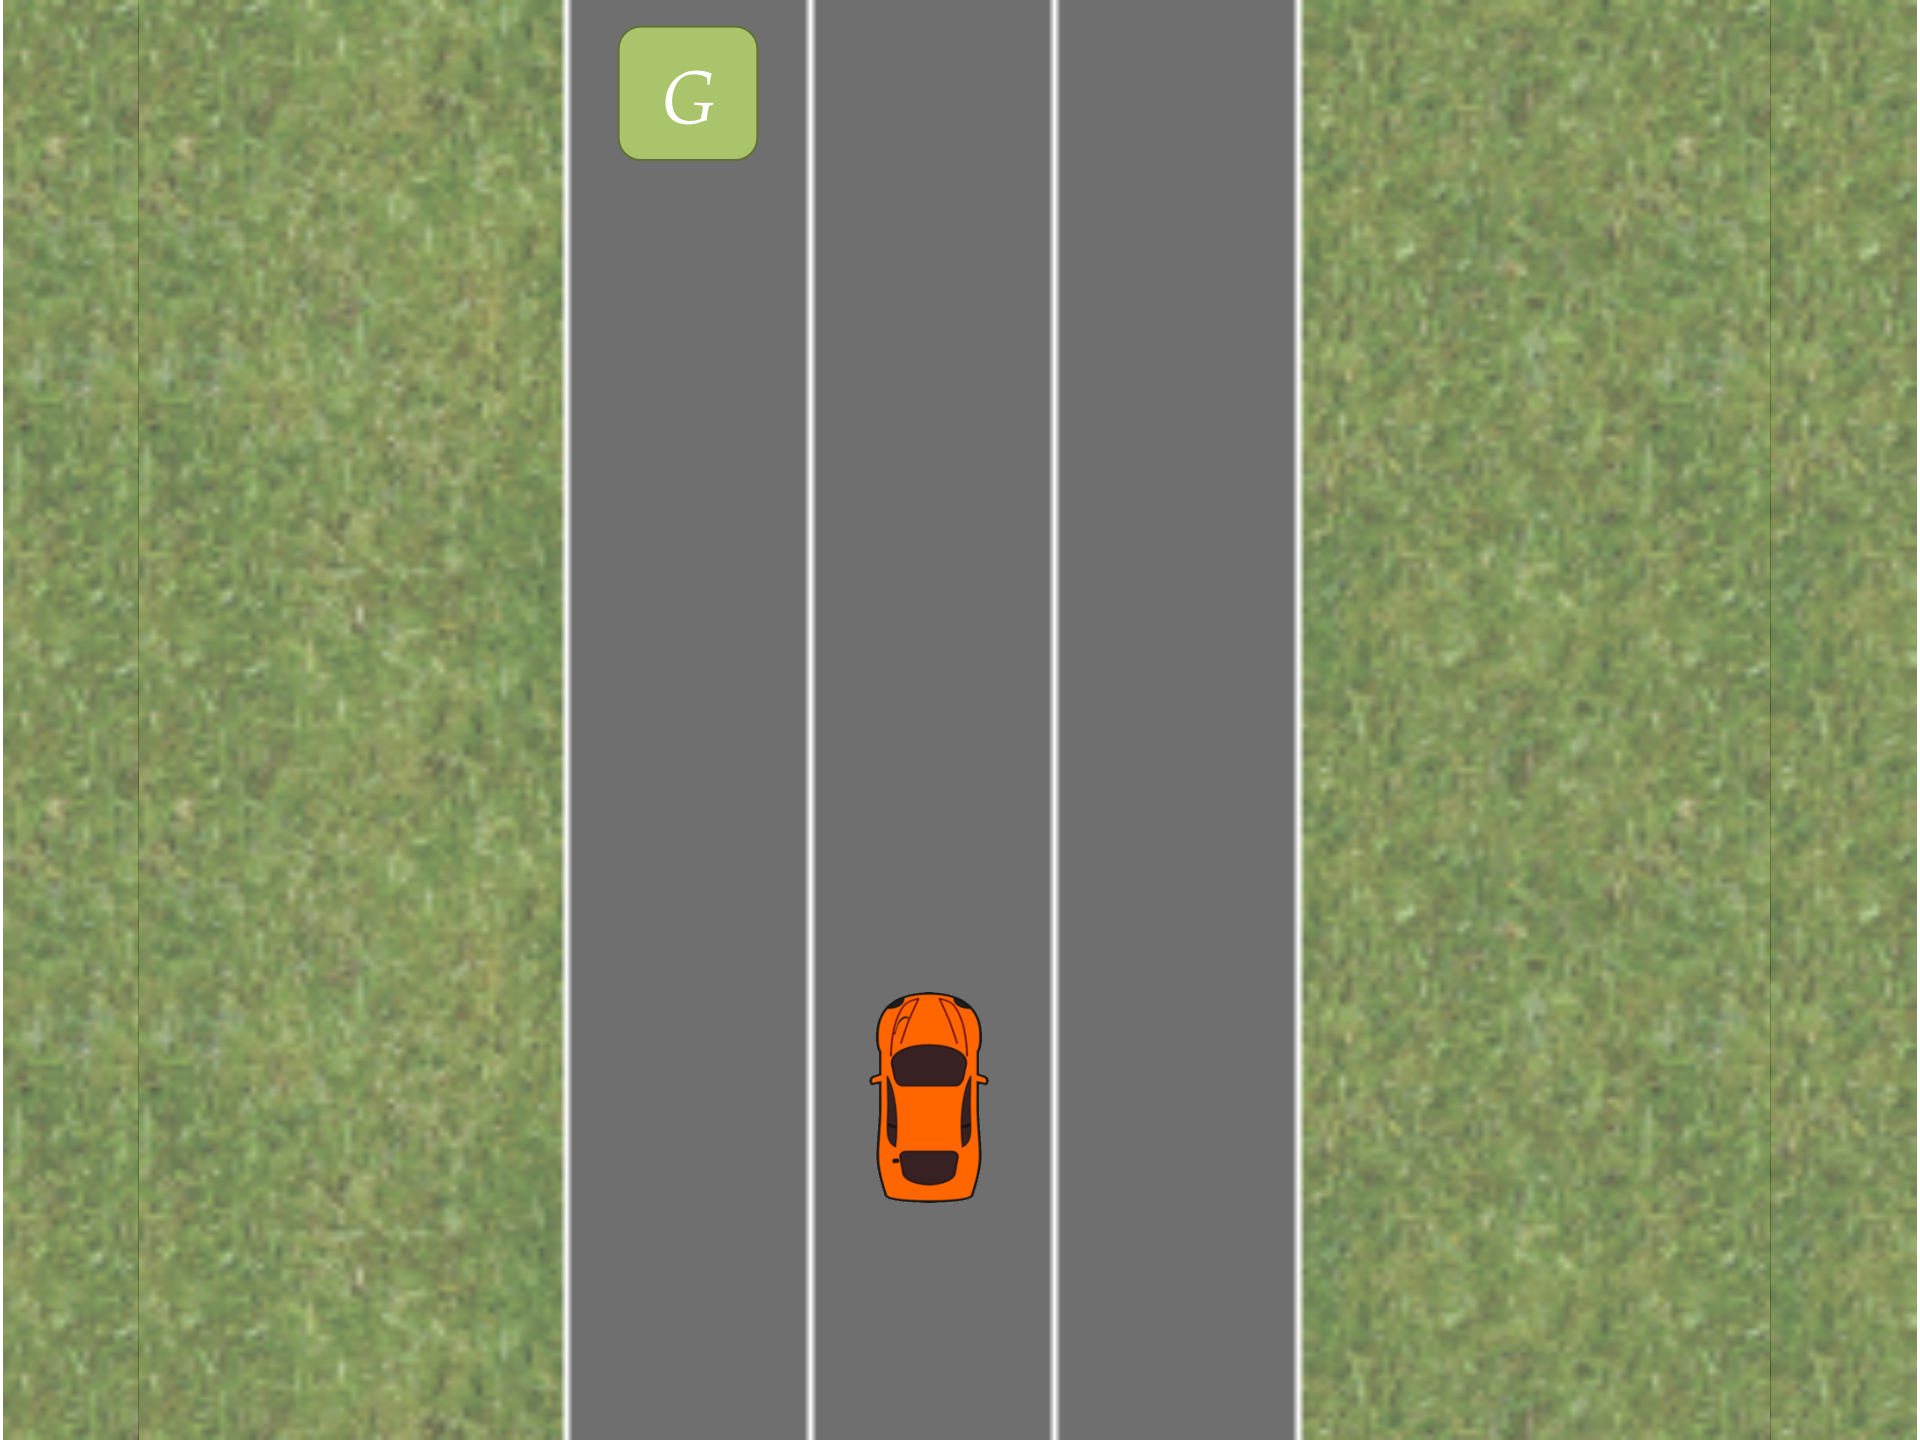
G

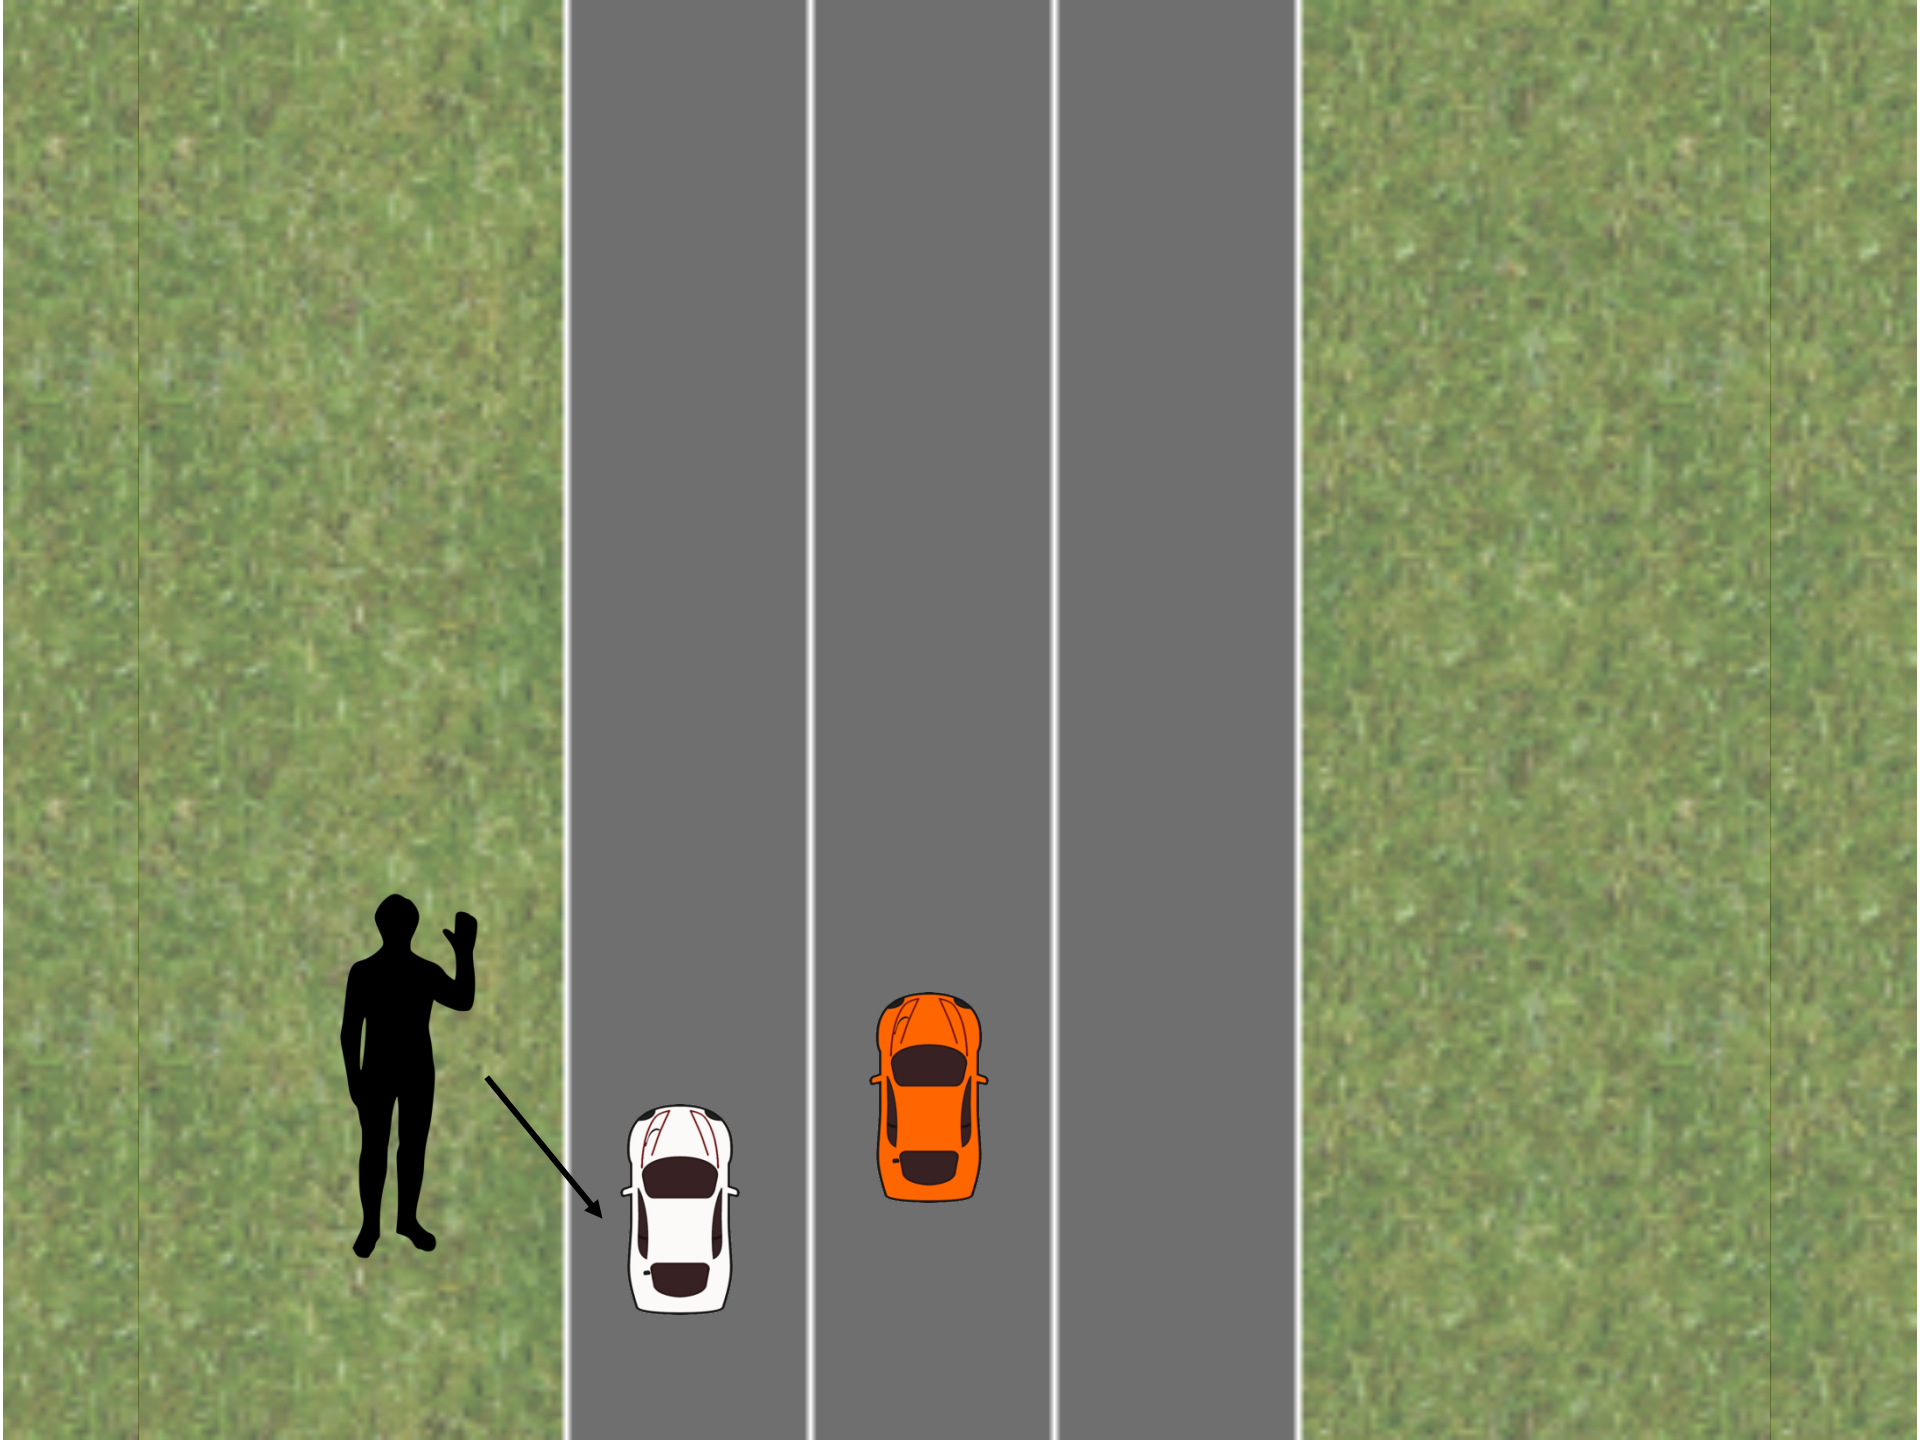

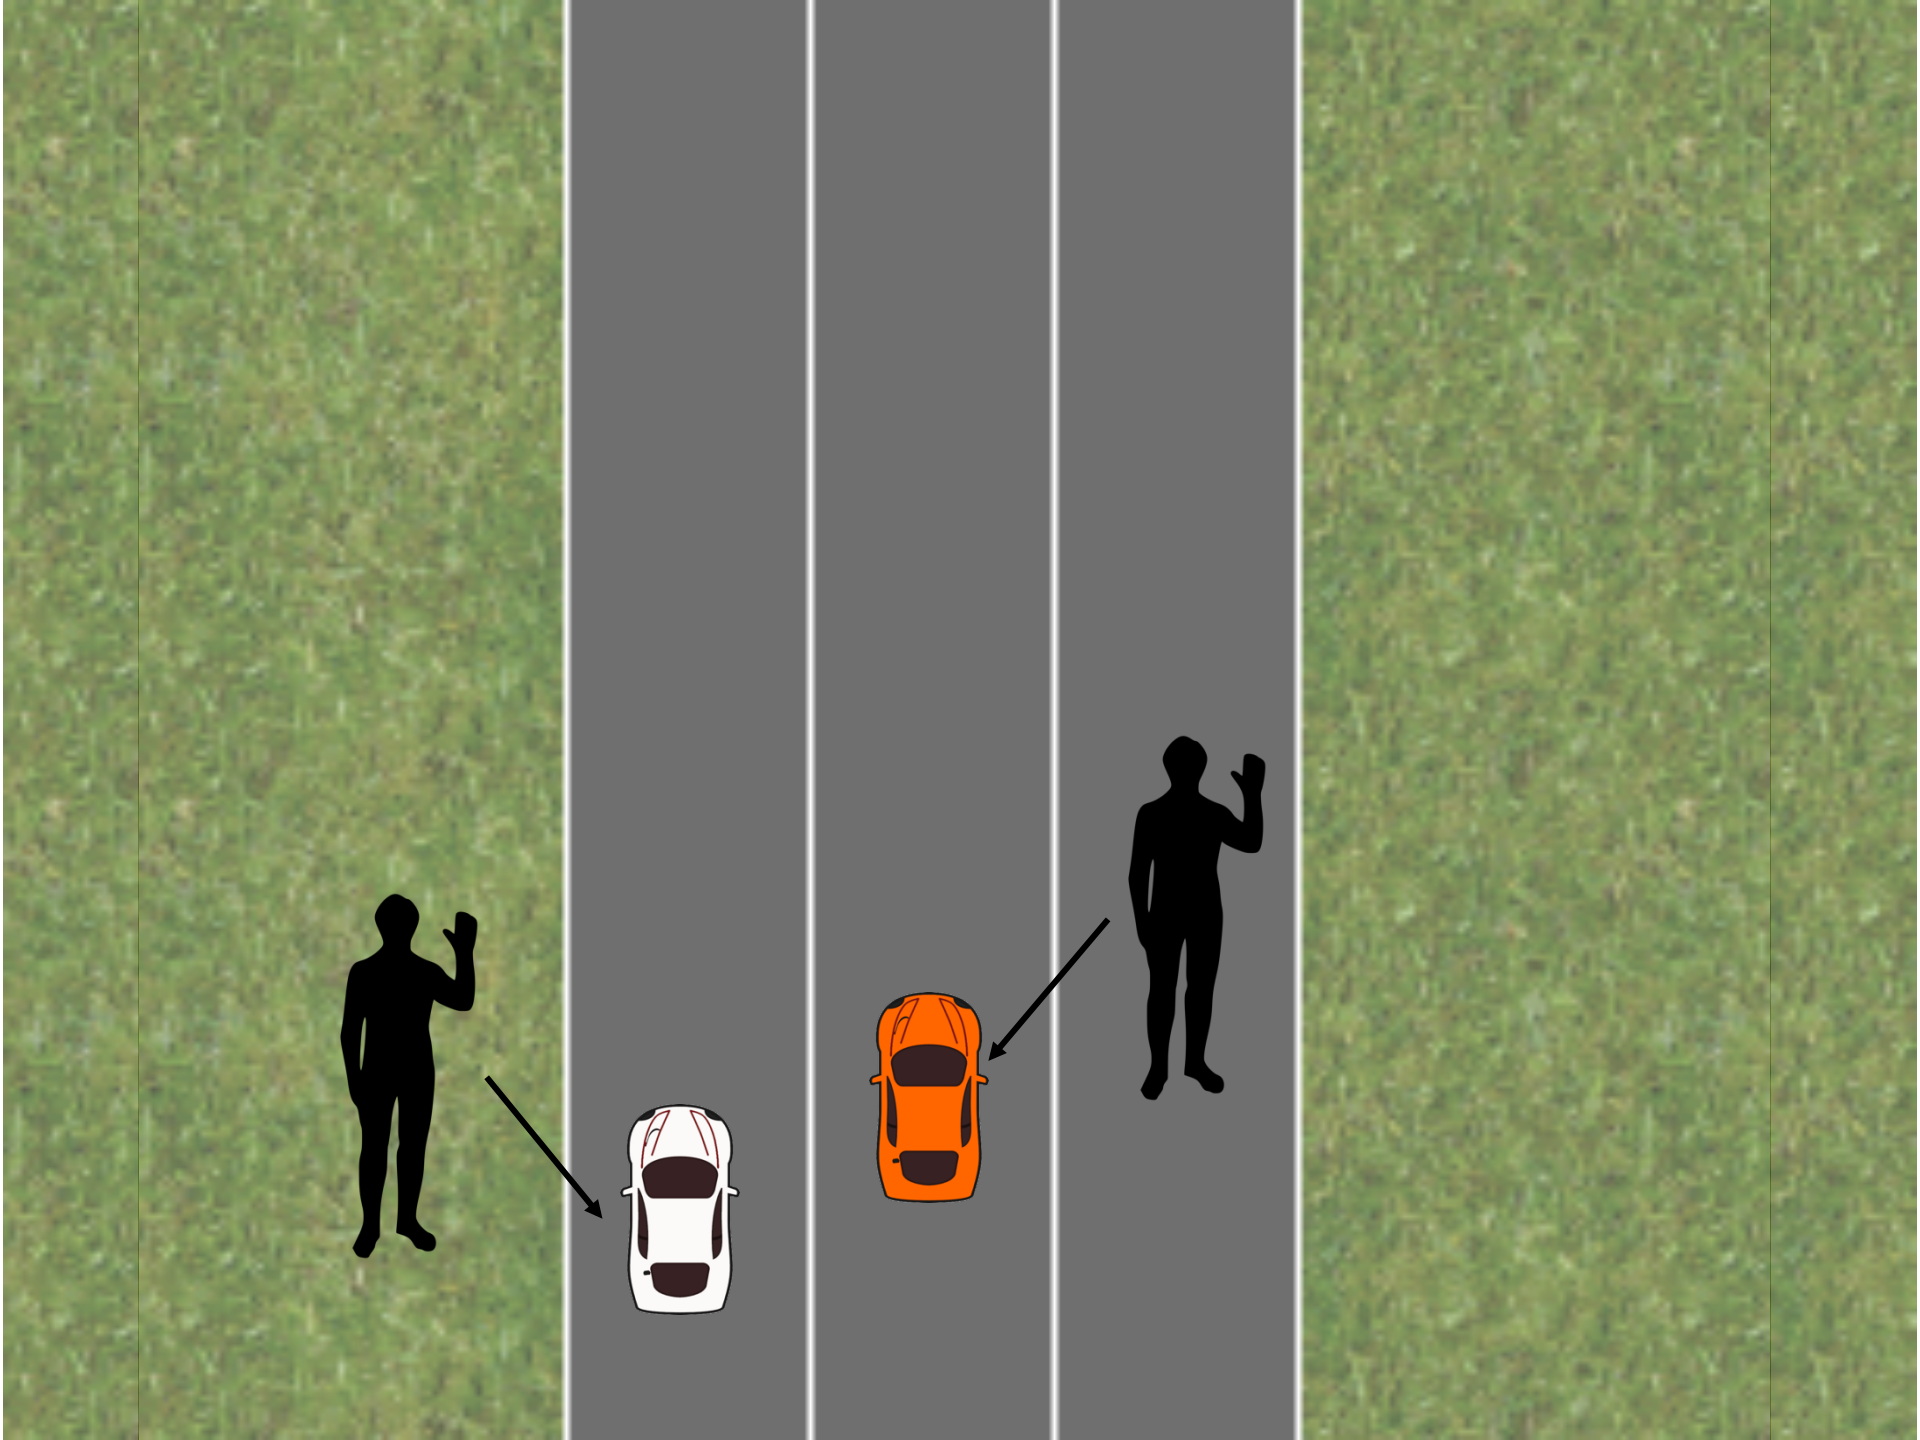

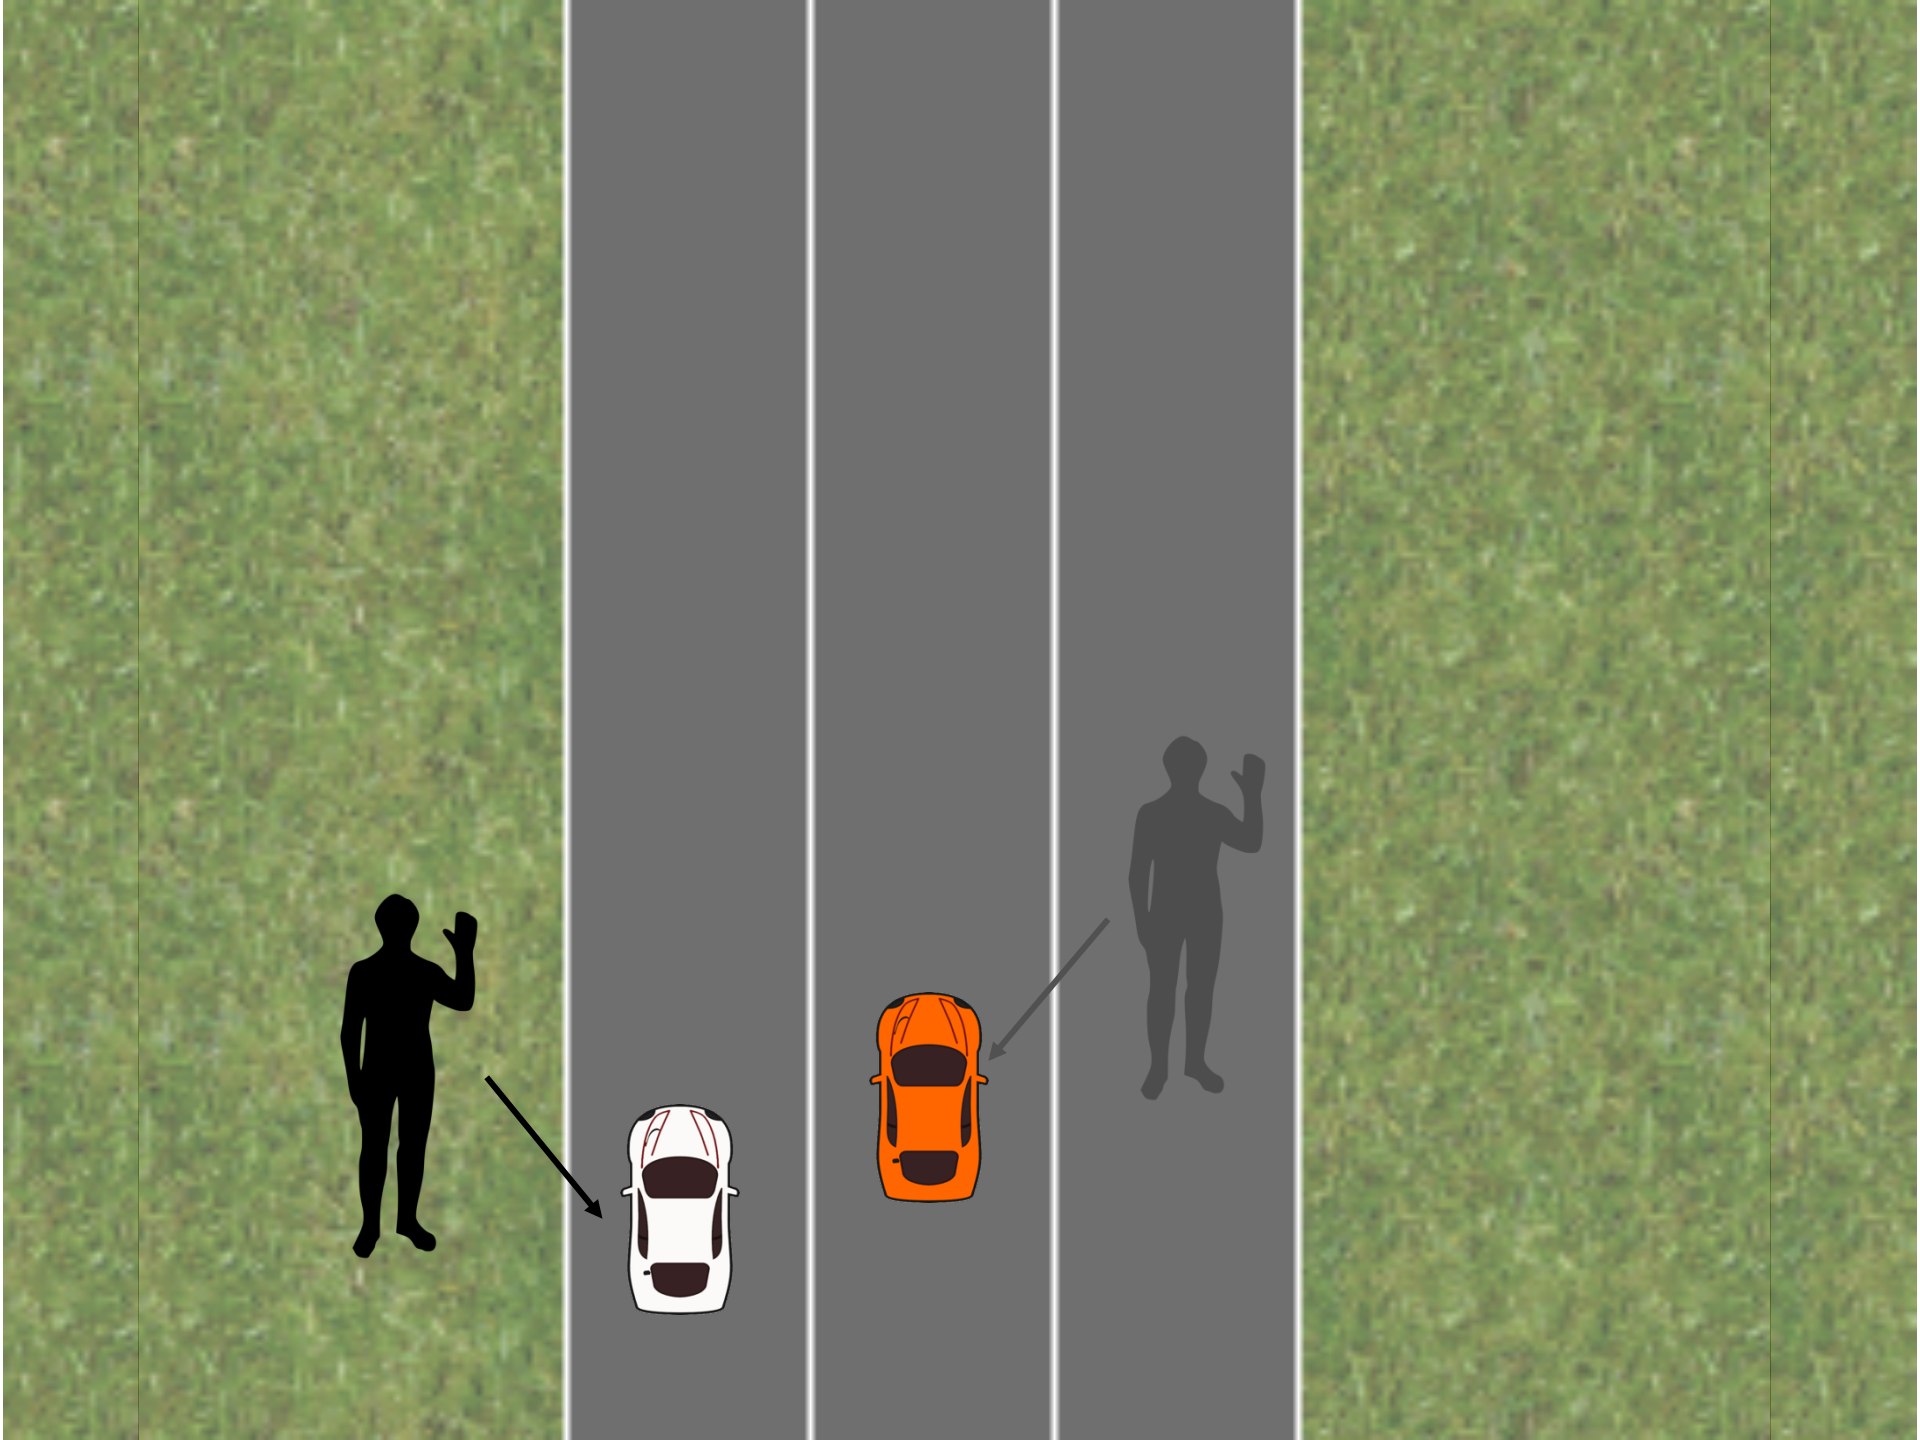

Compatibility = Stay out of the way

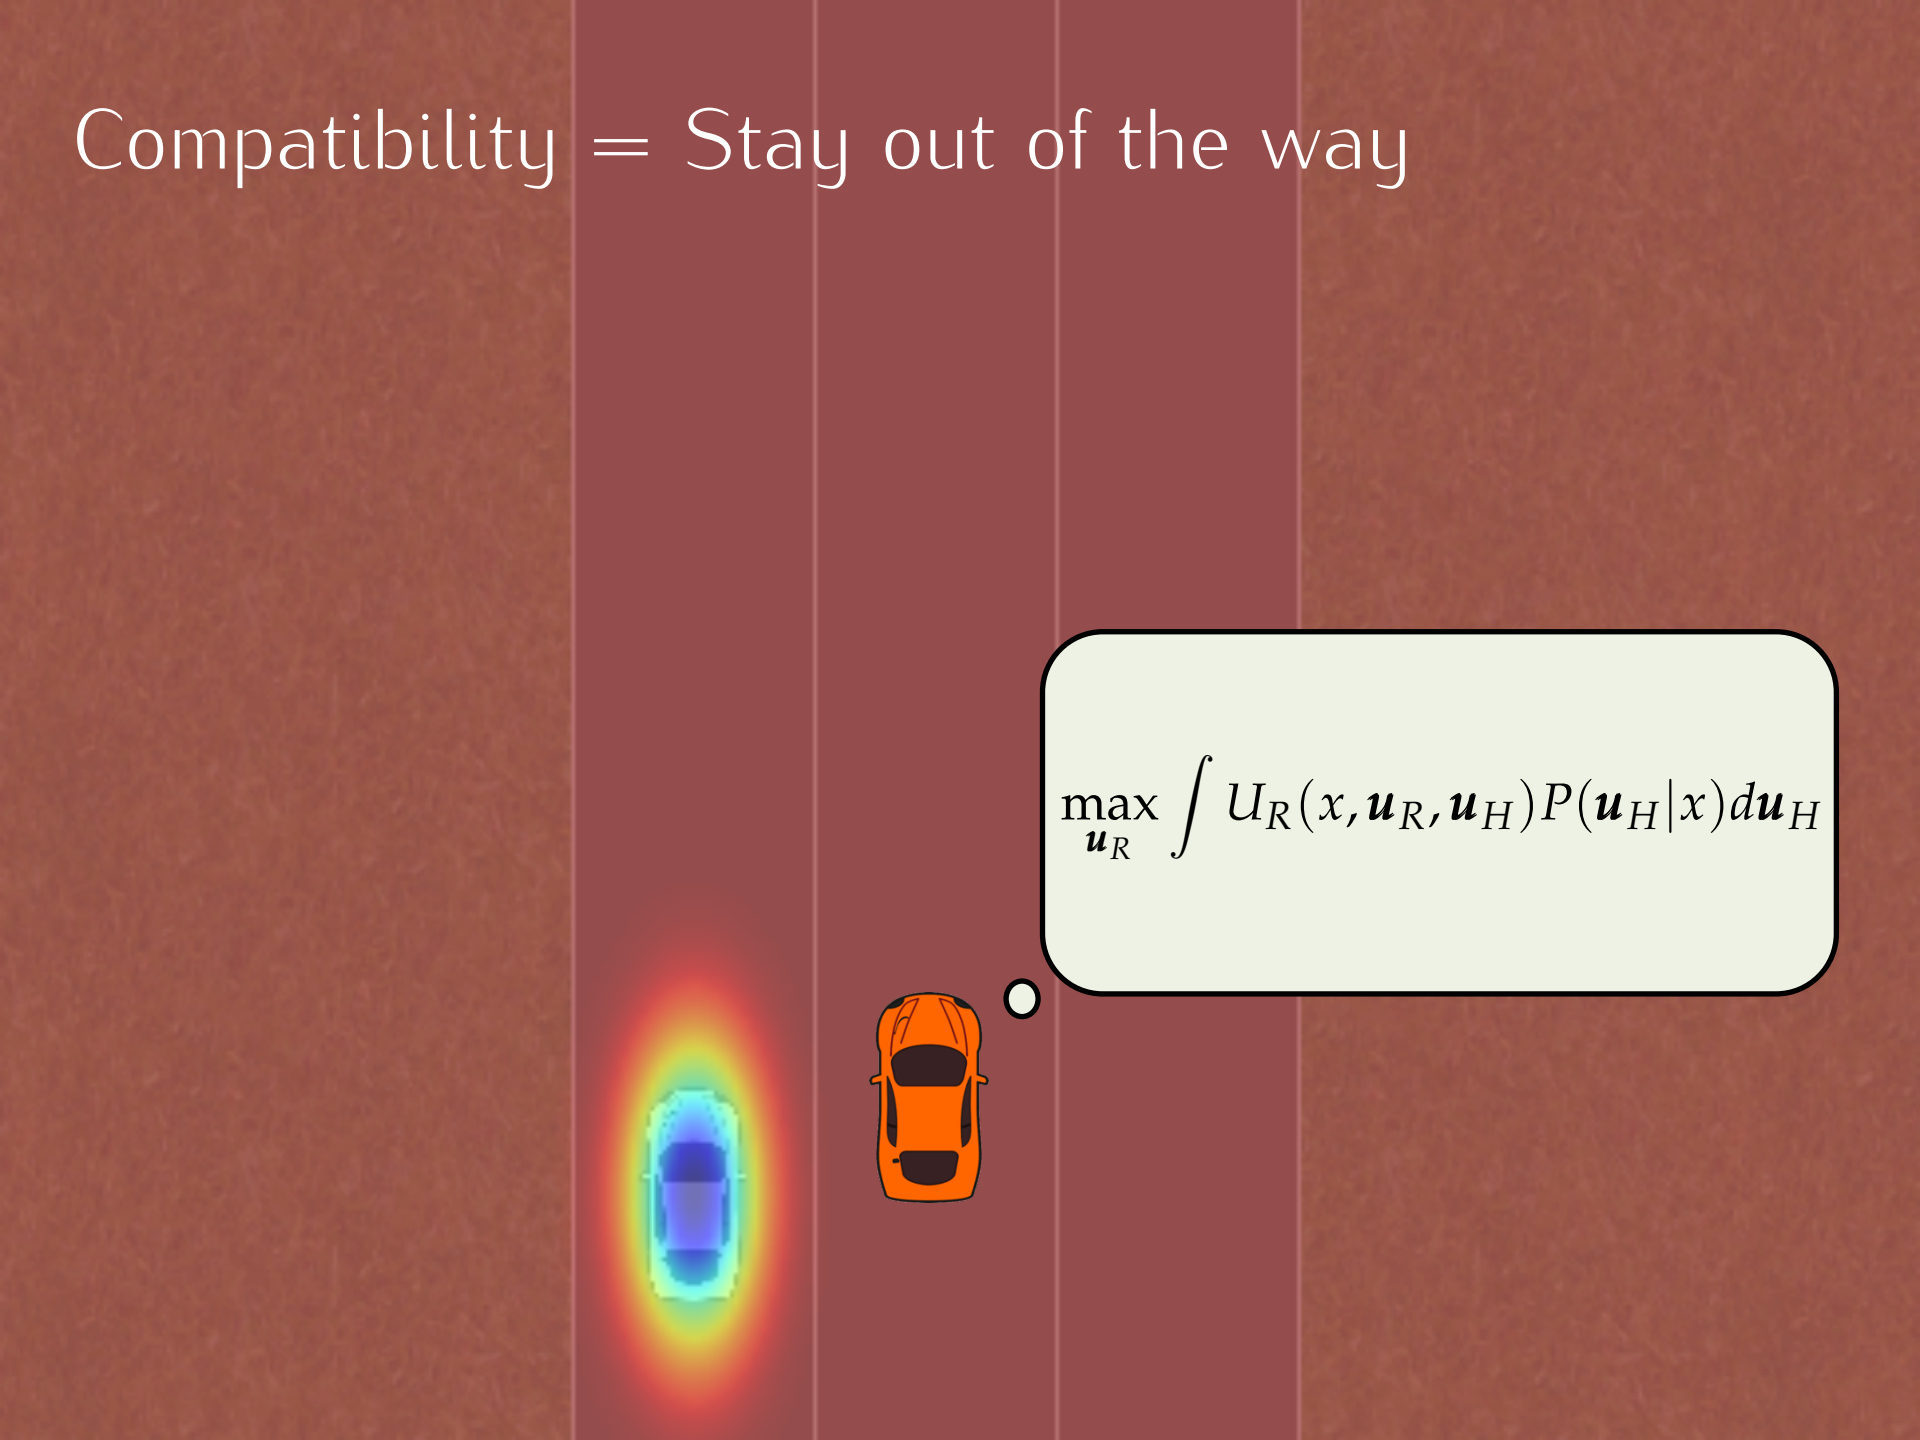


SPEED
SOCIETY
2015/02/06 23:09:54


SPEED
SOCIETY
2015/02/06 23:10:06


civic

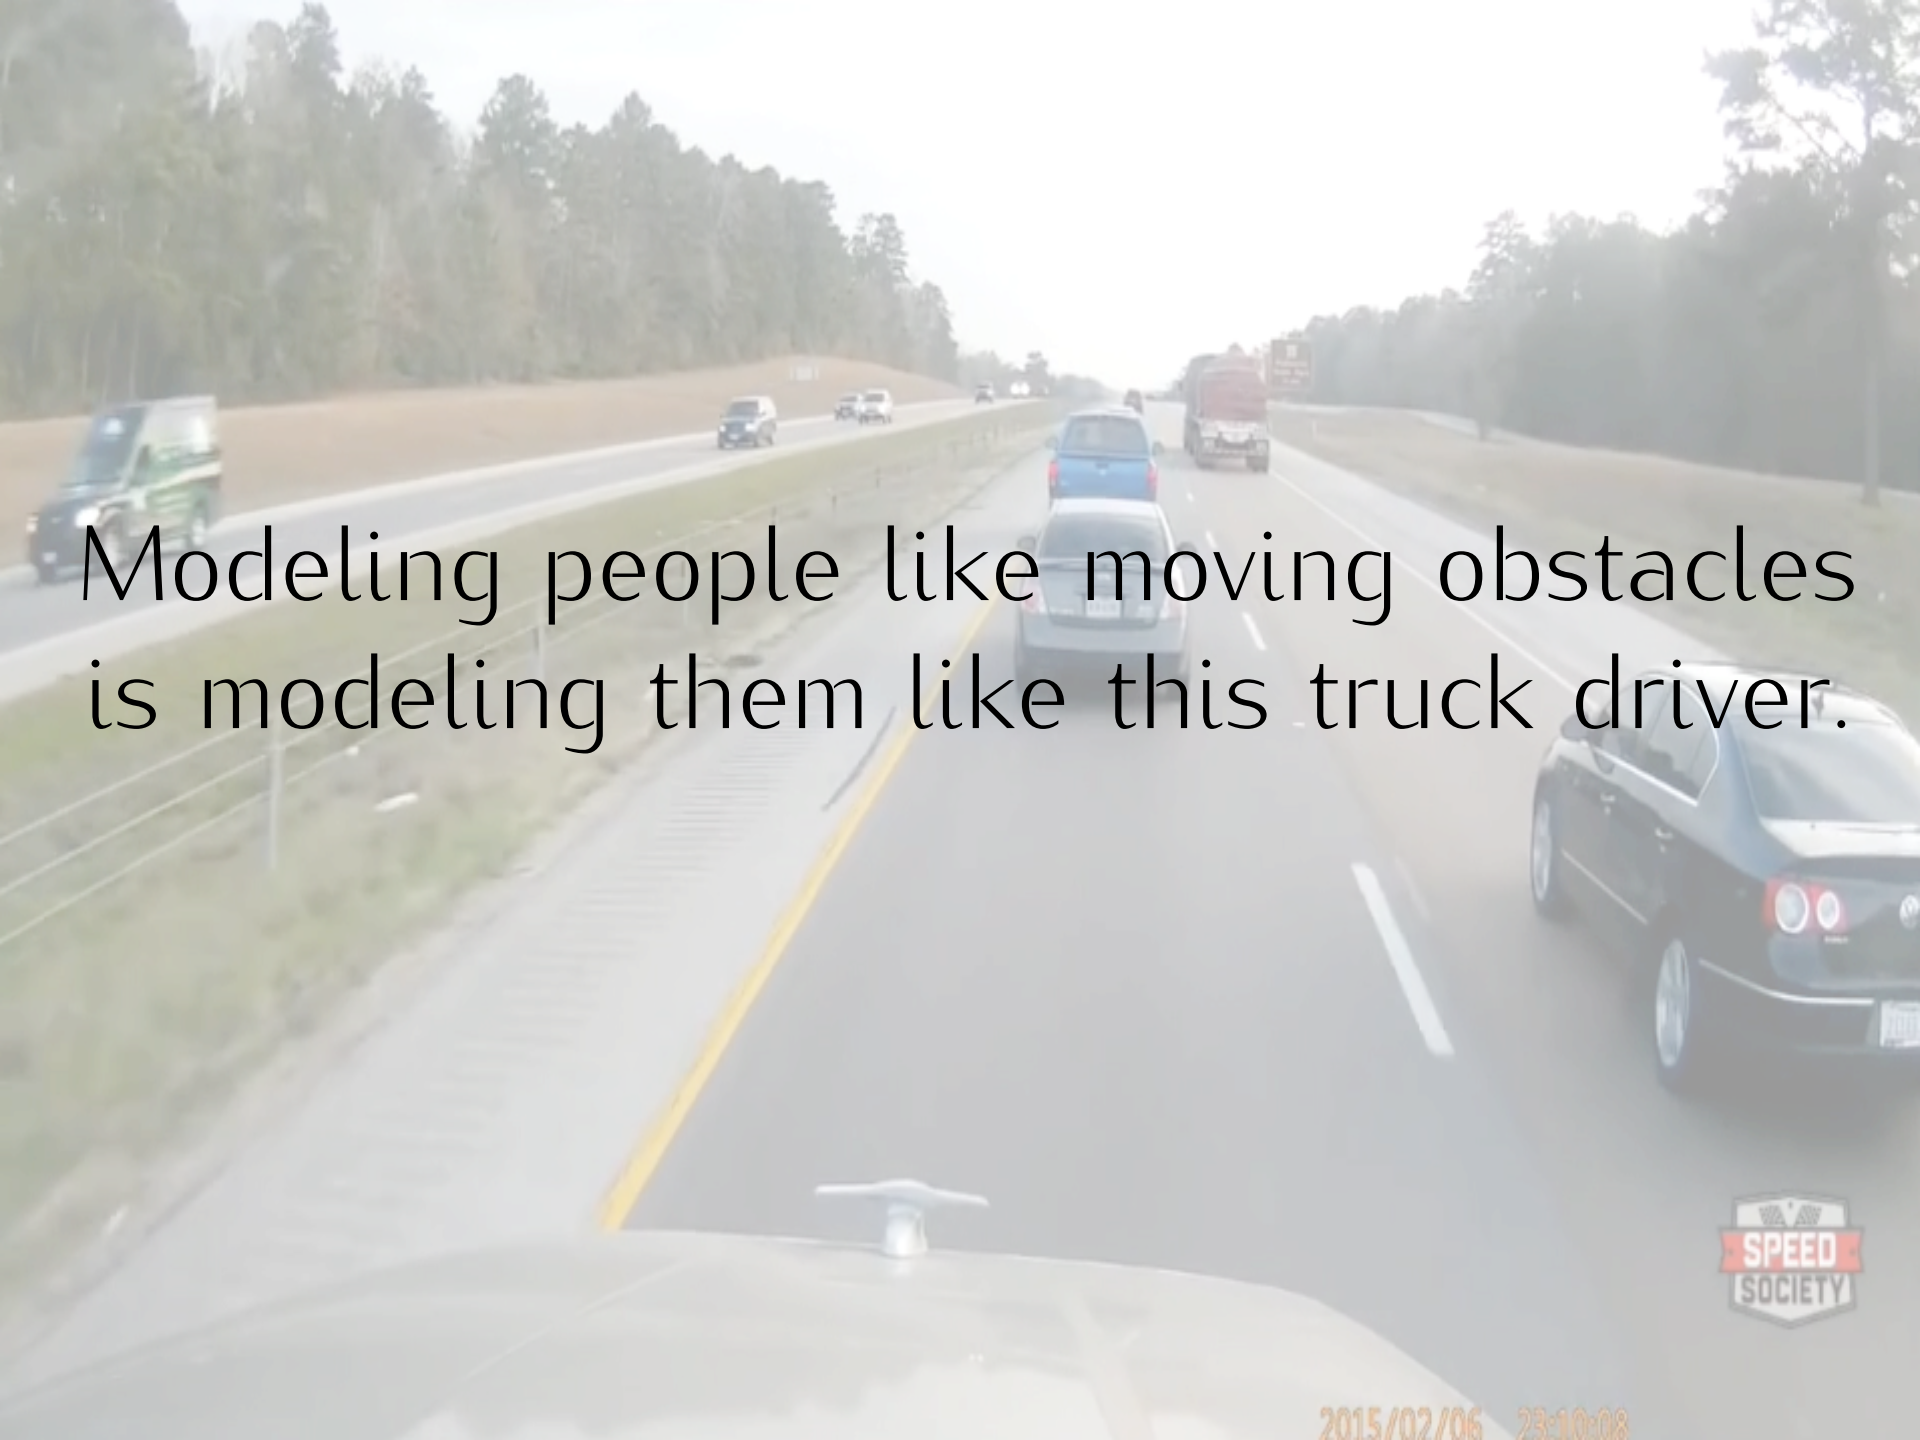

Modeling people like moving obstacles is modeling them like this truck driver.

Google's fleet of autonomous test cars is programmed to follow the letter of the law. But it can be tough to get around if you are a stickler for the rules. One Google car, in a test in 2009, couldn't get through a four-way stop because its sensors kept waiting for other (human) drivers to stop completely and let it go. The human drivers kept inching forward, looking for the advantage — paralyzing Google's robot.

"Google's Driverless Cars Run Into Problem: Cars With Drivers"[Richtel&Dougherty]

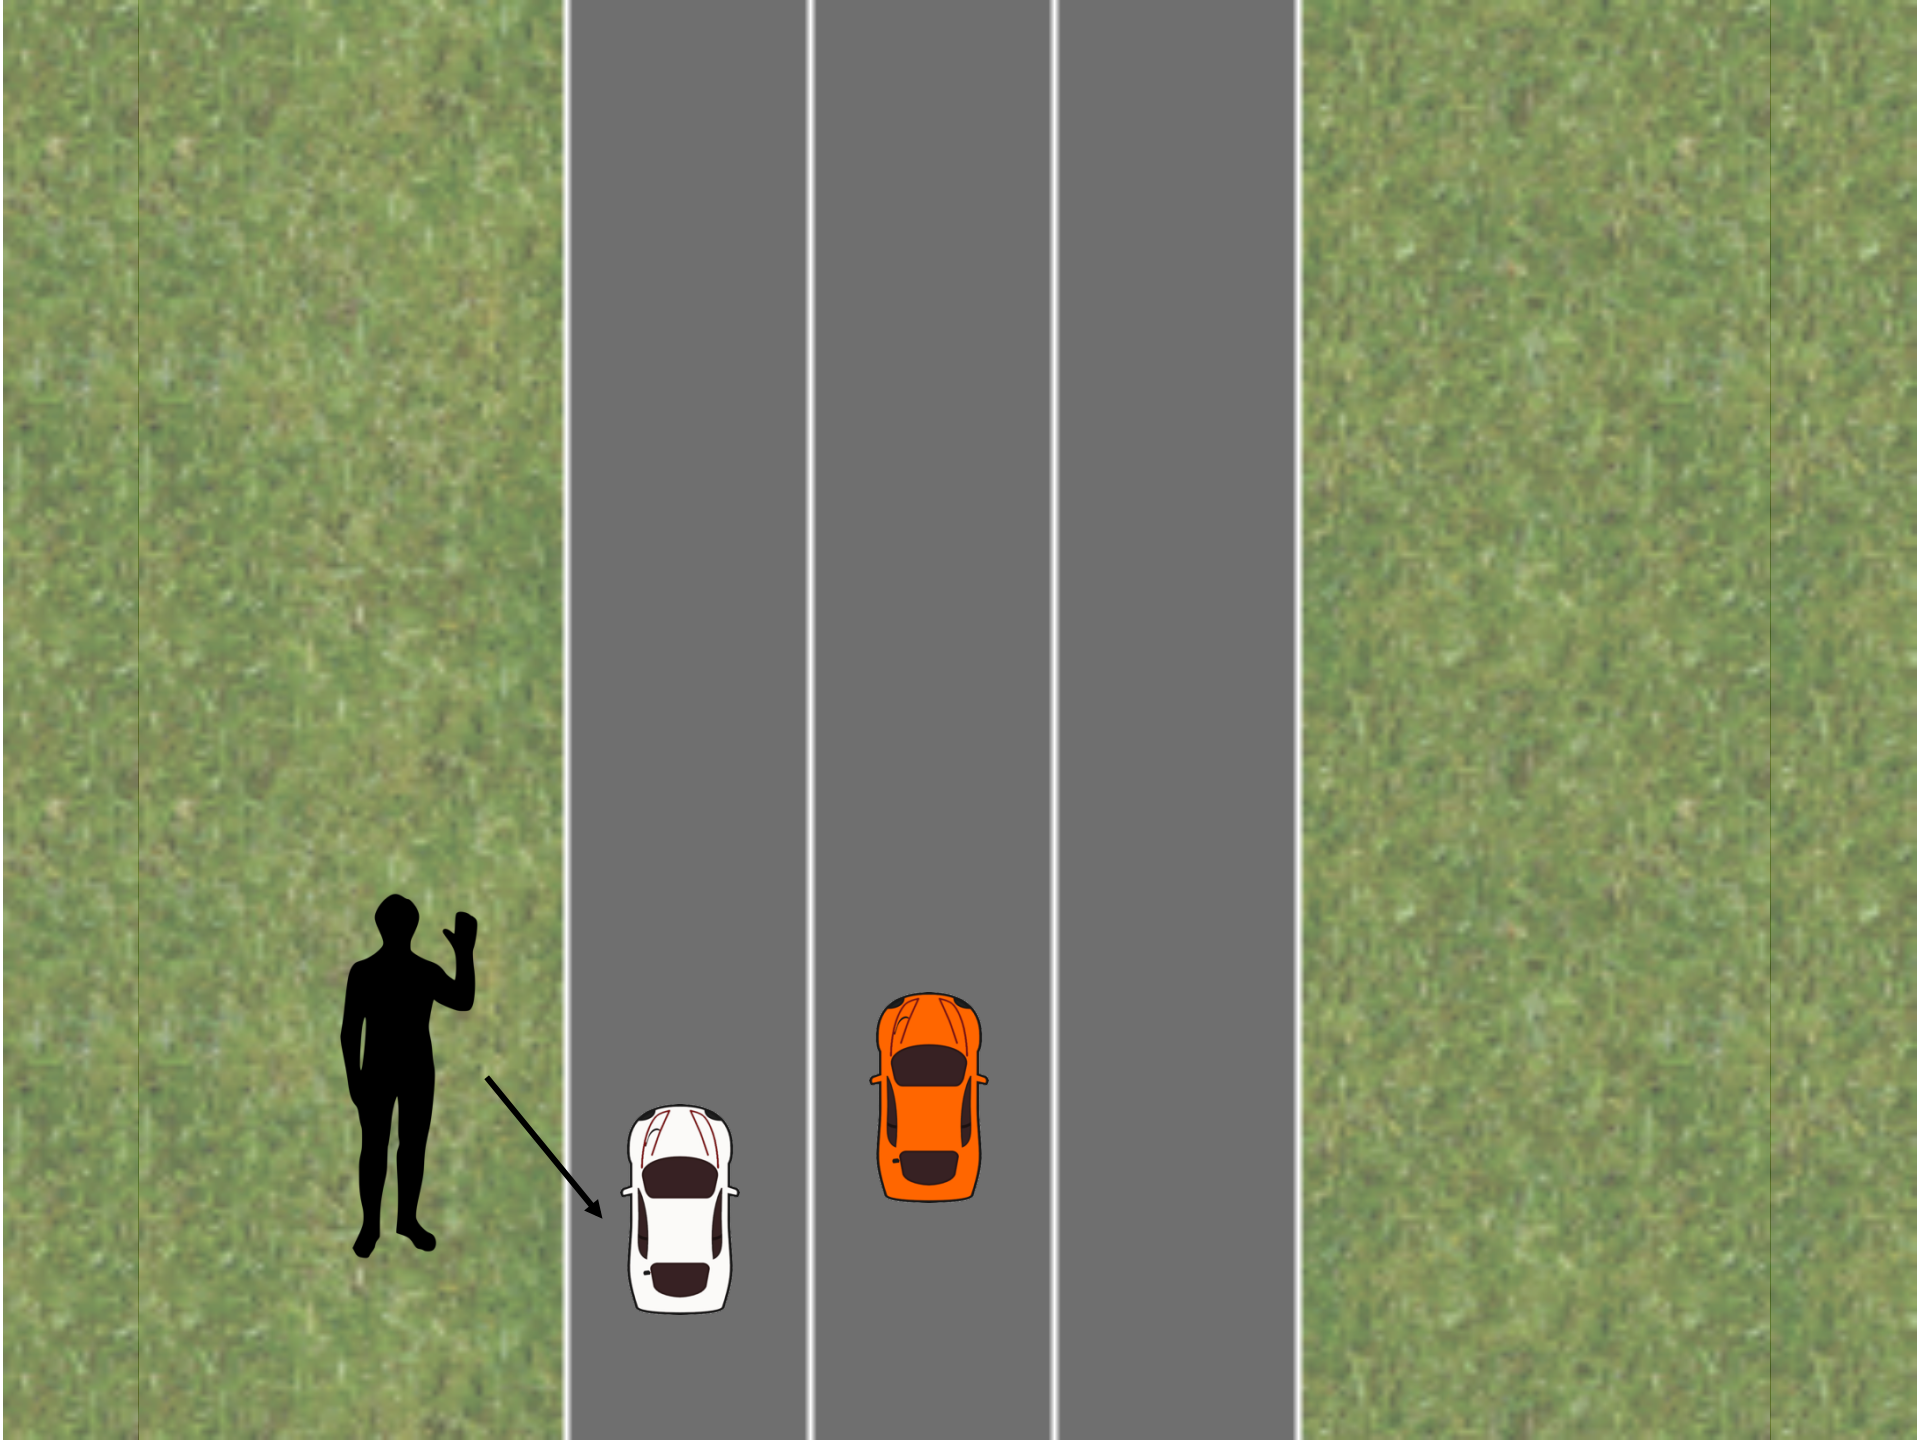

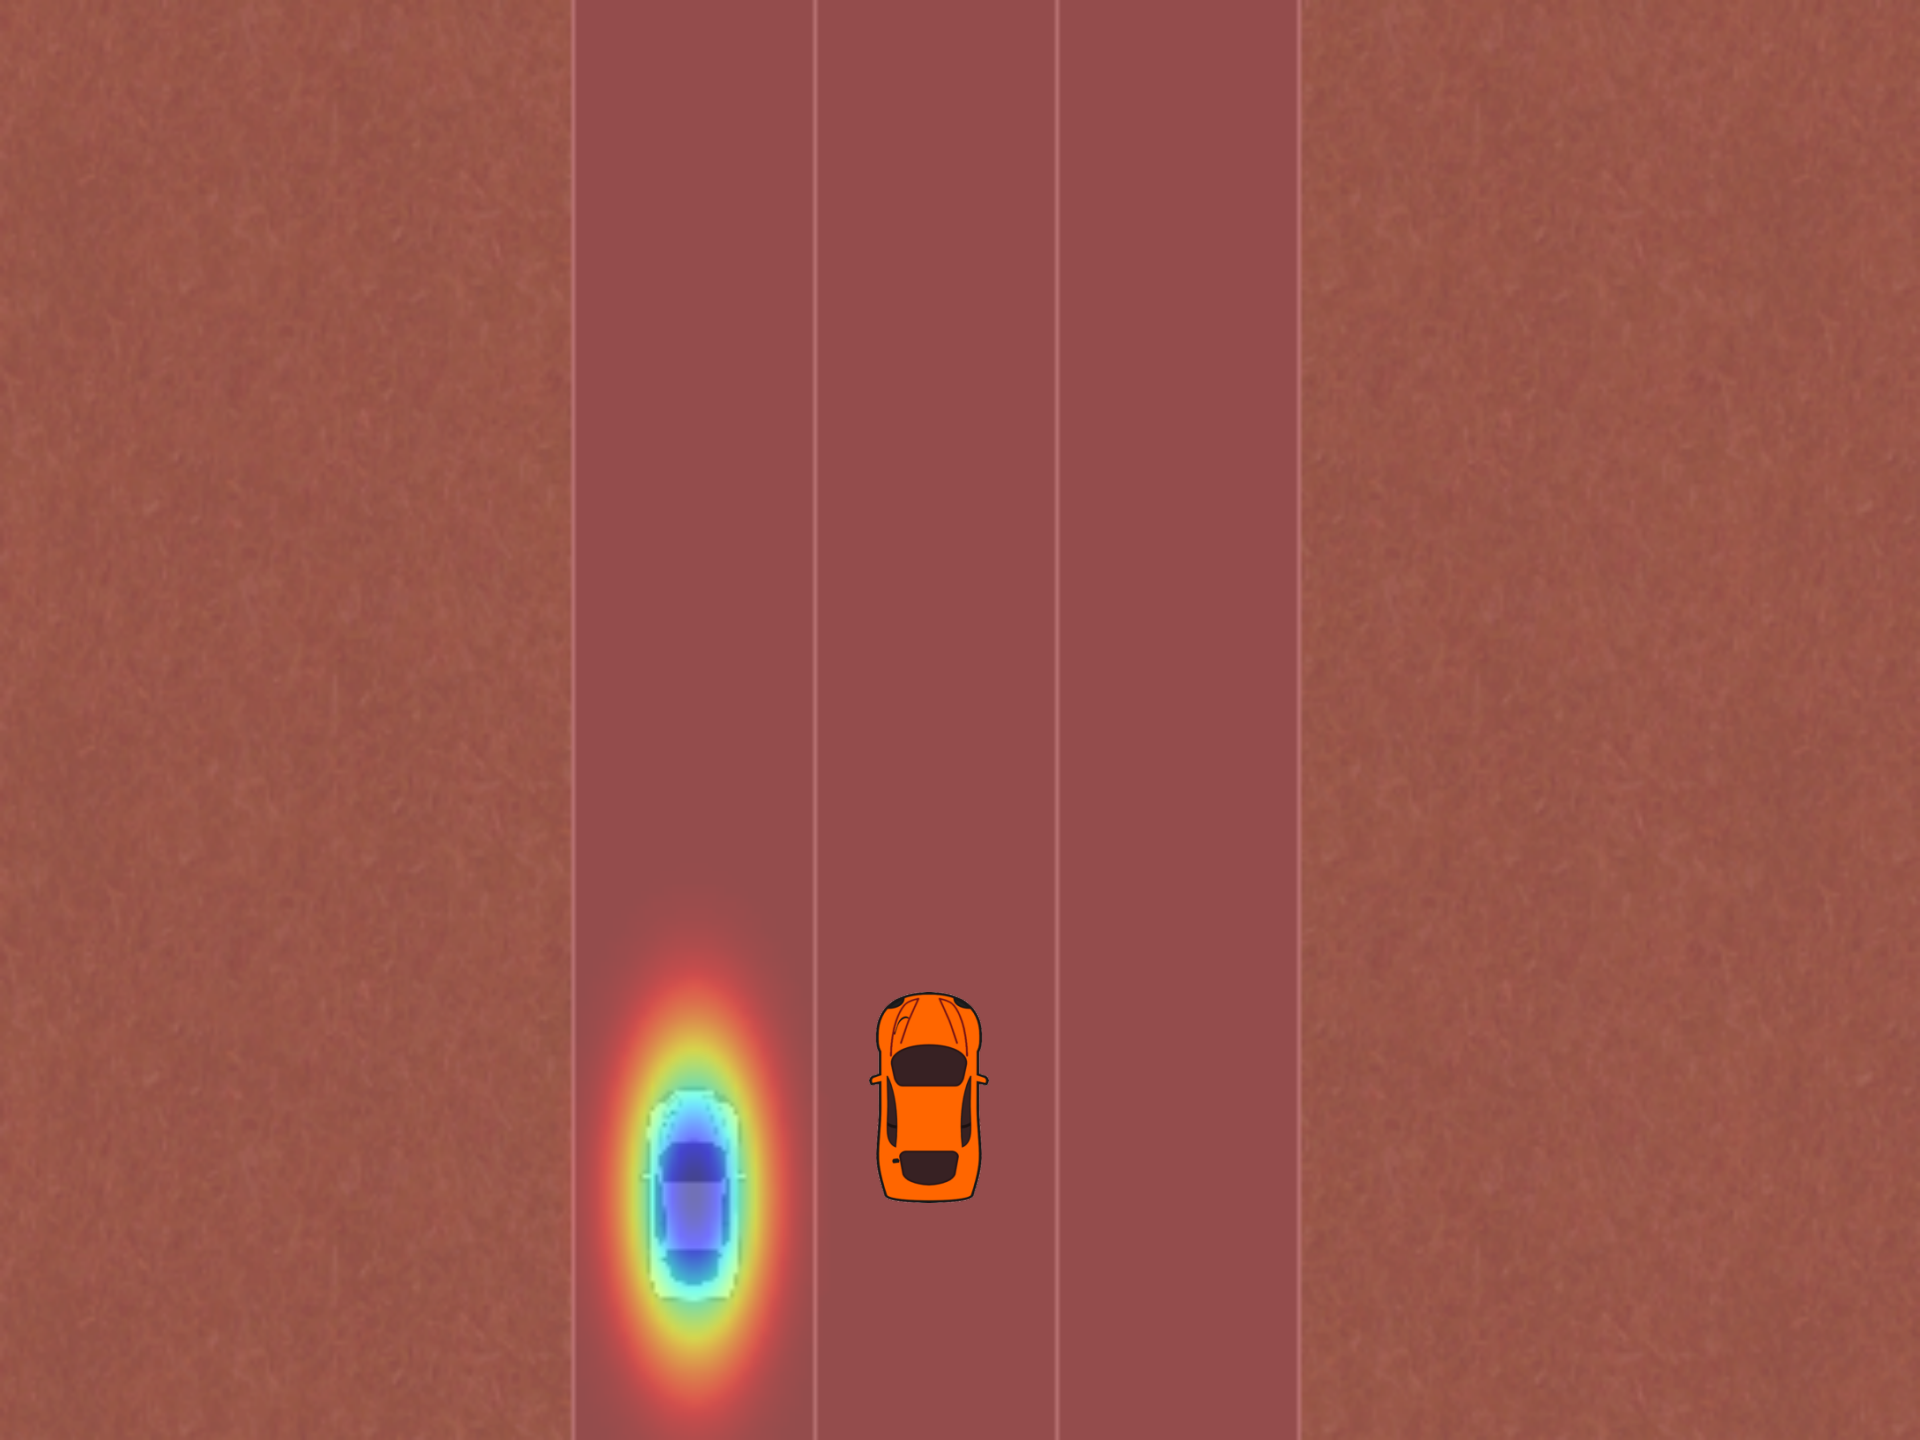

A Human-Robot Game

$$U_H(x, u_R, u_H) \quad U_R(x, u_R, u_H)$$

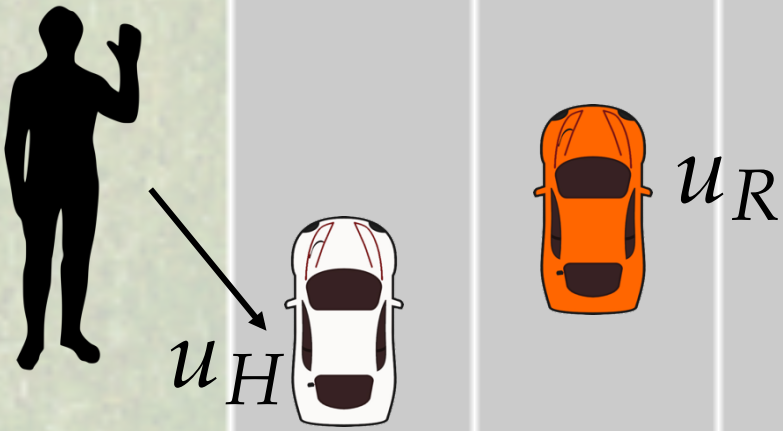

A Partial Information Human-Robot Game

$U_H(x, u_R, u_H; \theta_H)$ $U_R(x, u_R, u_H)$

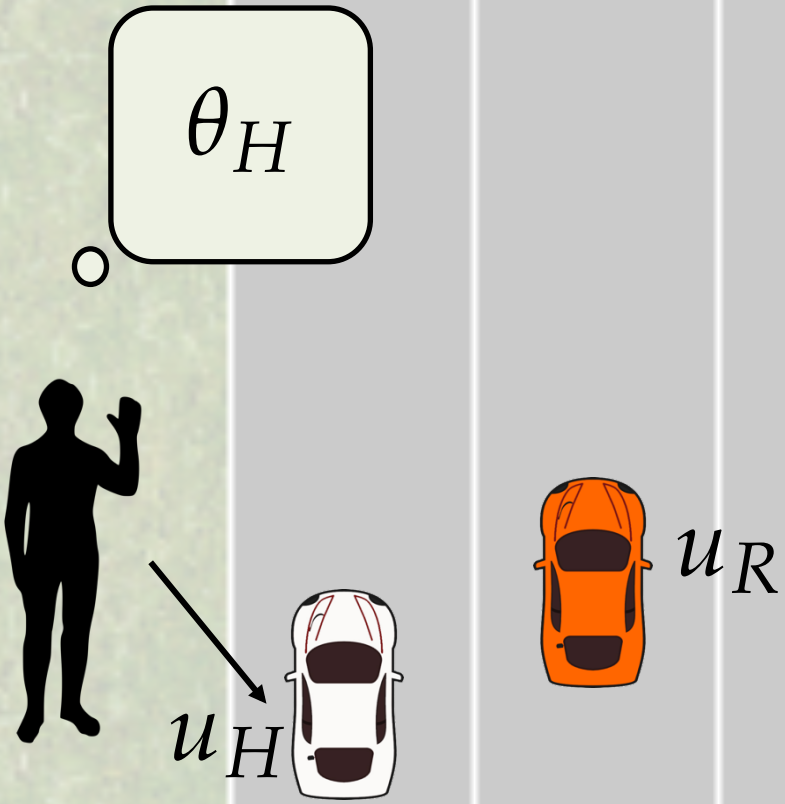

A Partial Information Human-Robot Game

$U_H(x, u_R, u_H; \theta_H)$ $U_R(x, u_R, u_H; \theta_R)$

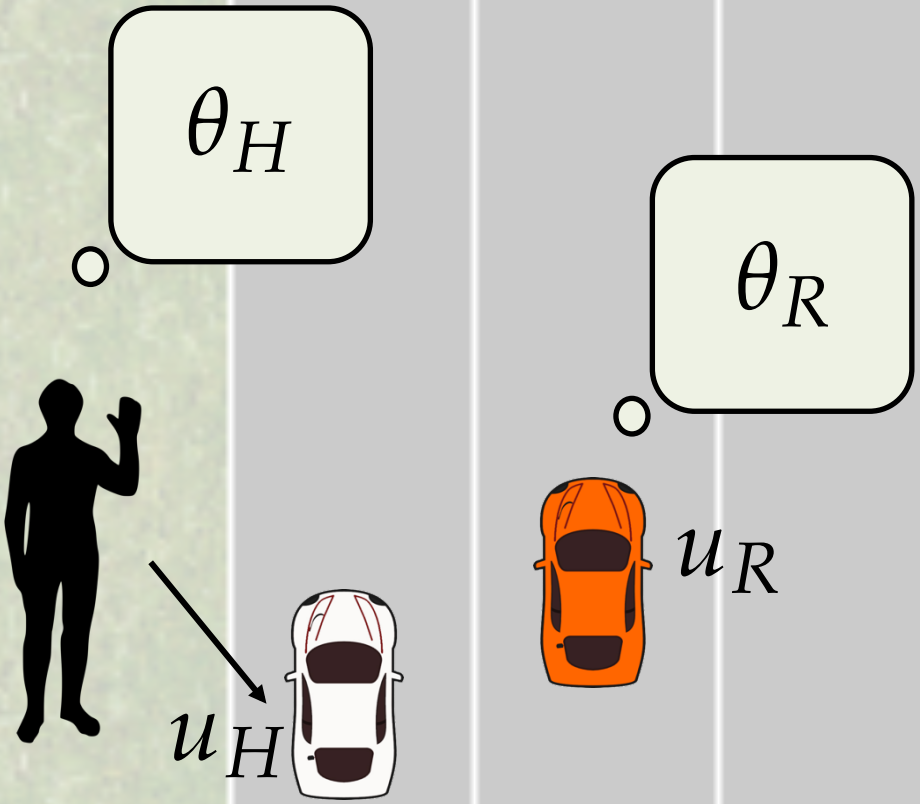

A Partial Information Human-Robot Game

$$U(x, u_R, u_H; \theta_H)$$

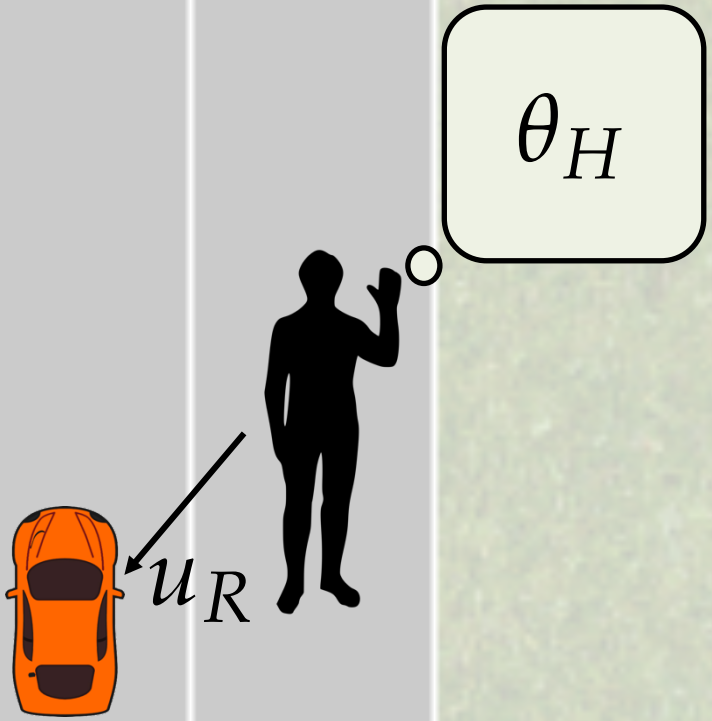

A Partial Information Human-Robot Game

$U_H(x, u_R, u_H; \theta_H)$ $U_R(x, u_R, u_H)$

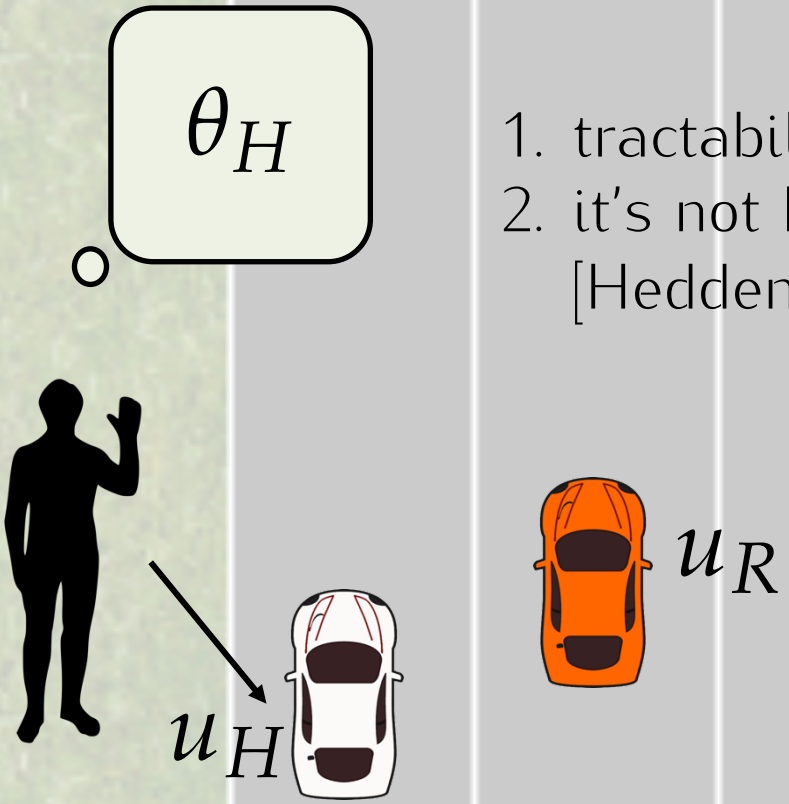

1. tractability, e.g. [Bernstein'02]
2. it's not how people work, e.g. [Hedden'02]

Approximation as Underactuated System

$U_H(x, u_R, u_H; \theta_H)$ $U_R(x, u_R, u_H)$

θ_H $\mathbf{u}_H^* = \arg\max_{\mathbf{u}_H} U_H(x, \hat{\mathbf{u}}_R, \mathbf{u}_H; \theta_H)$

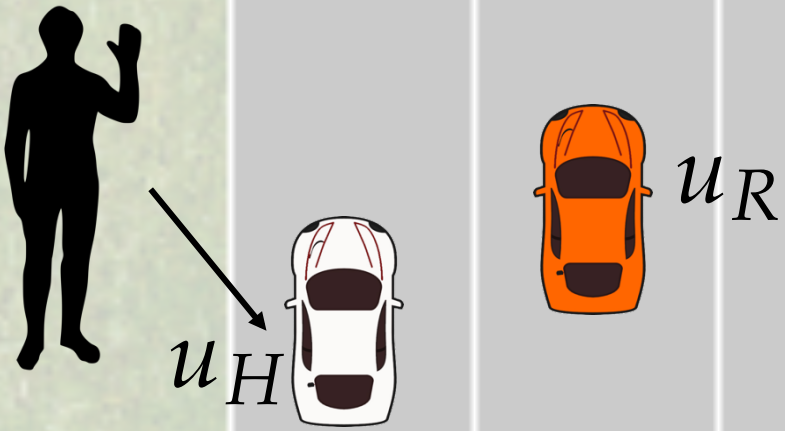

Approximation as Underactuated System

$U_H(x, u_R, u_H; \theta_H)$ $U_R(x, u_R, u_H)$

θ_H $\mathbf{u}_H^* = \arg\max_{\mathbf{u}_H} U_H(x, \mathbf{u}_R, \mathbf{u}_H; \theta_H)$

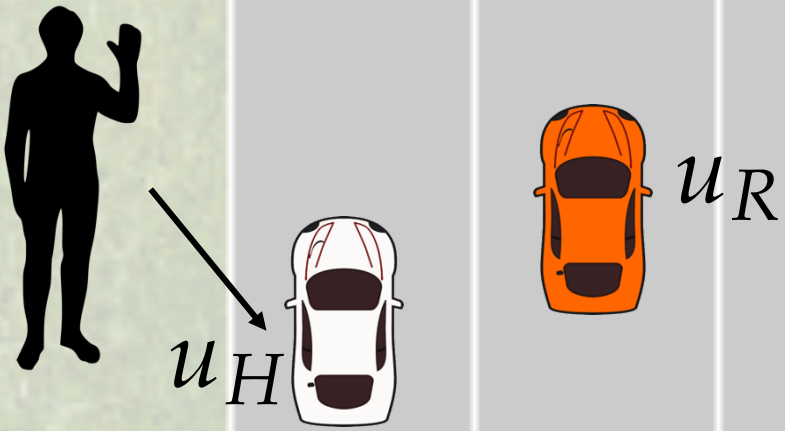

Planning for Autonomous Cars that Leverage Effects on Human Actions [RSS'16]

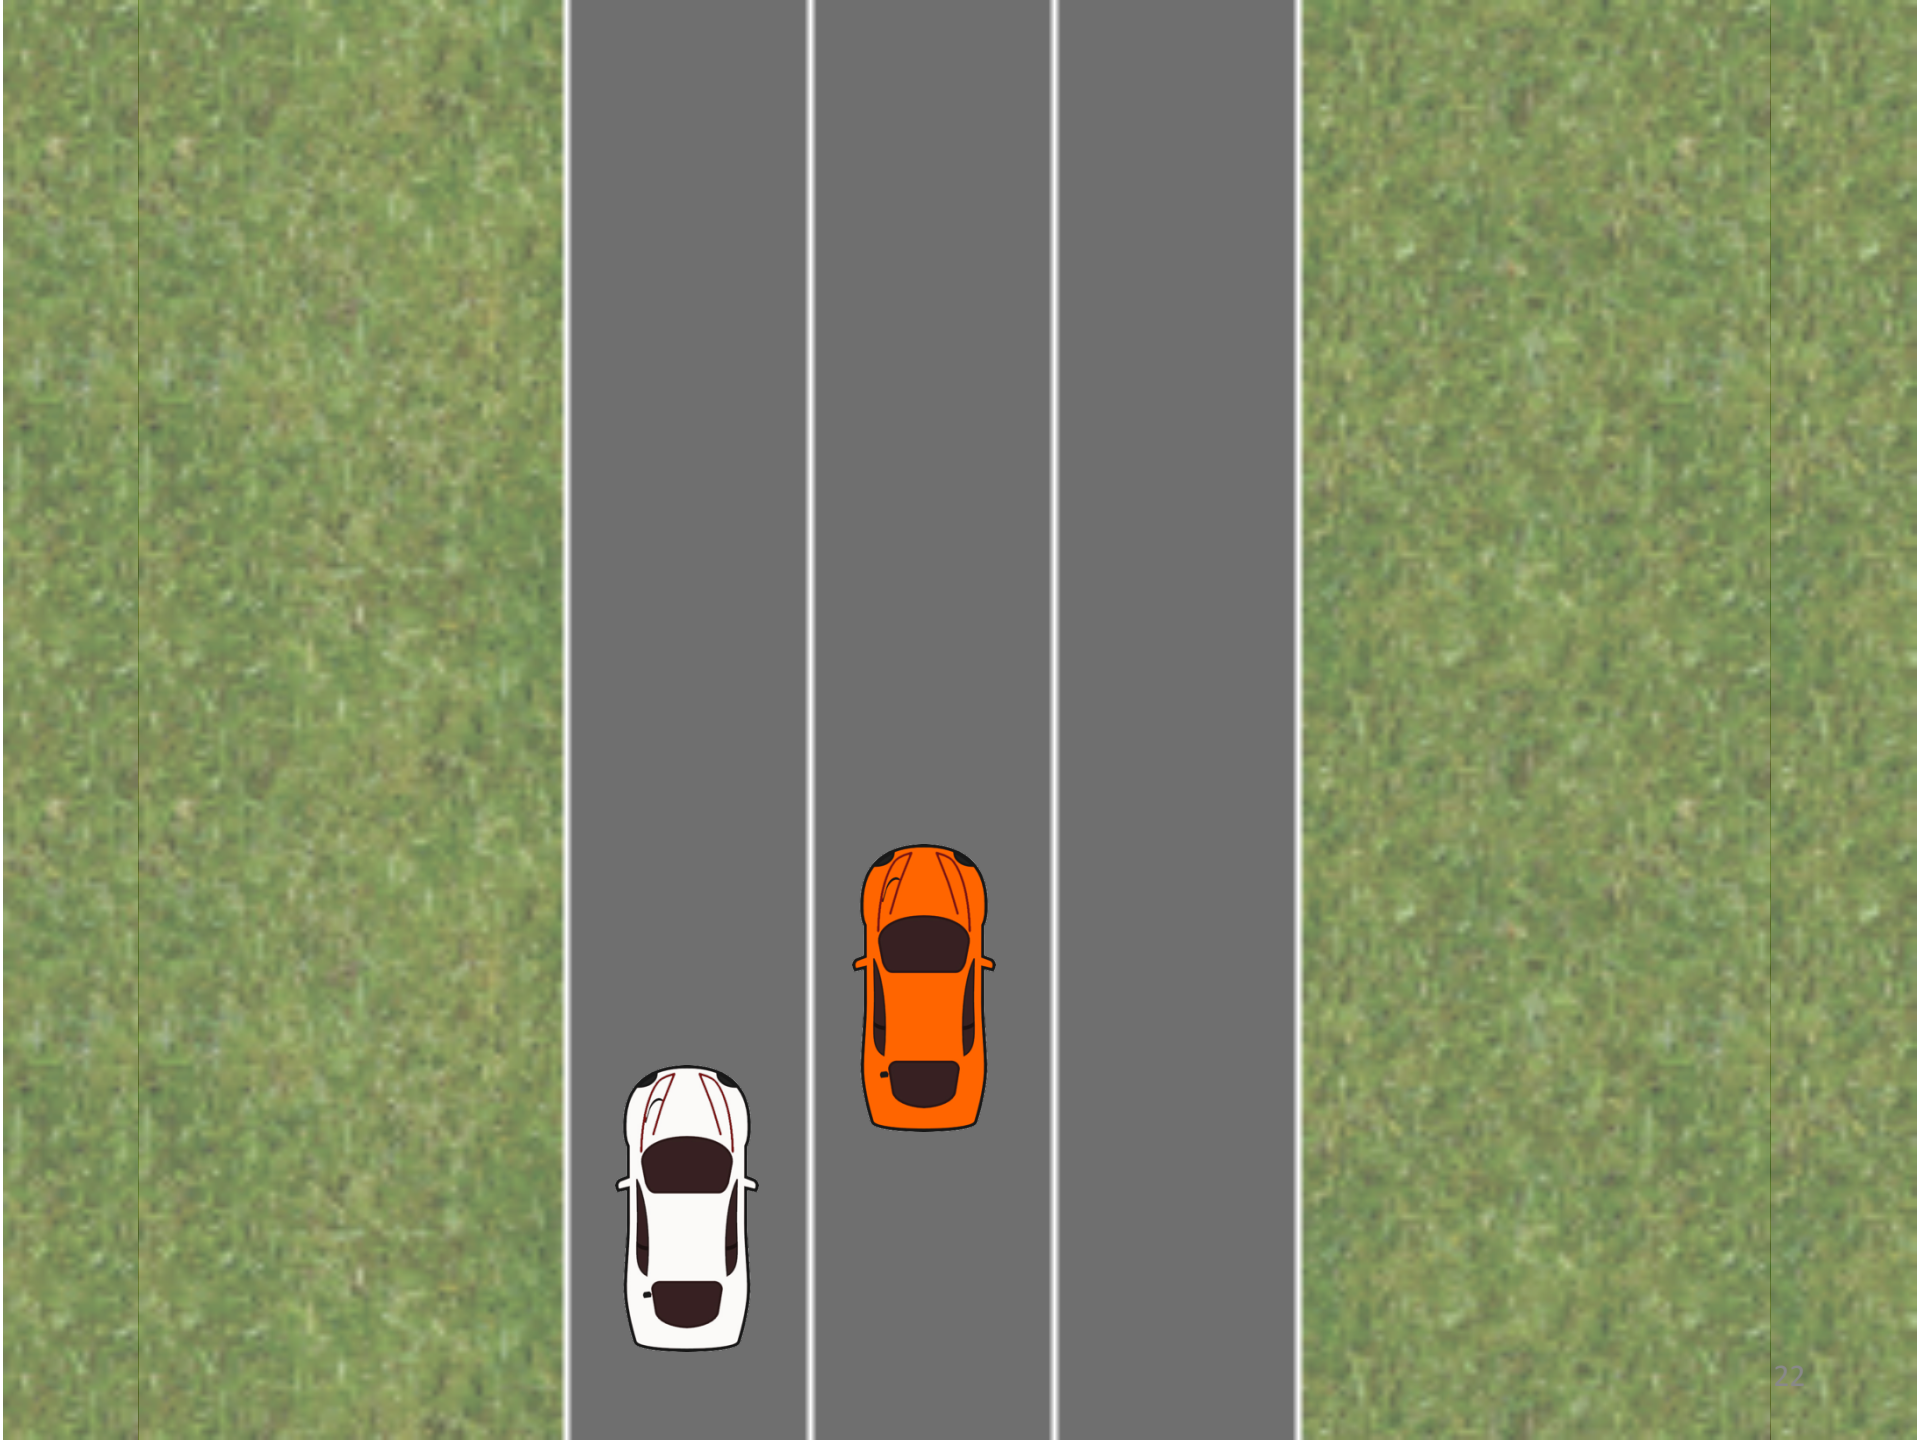

Human as Obstacle

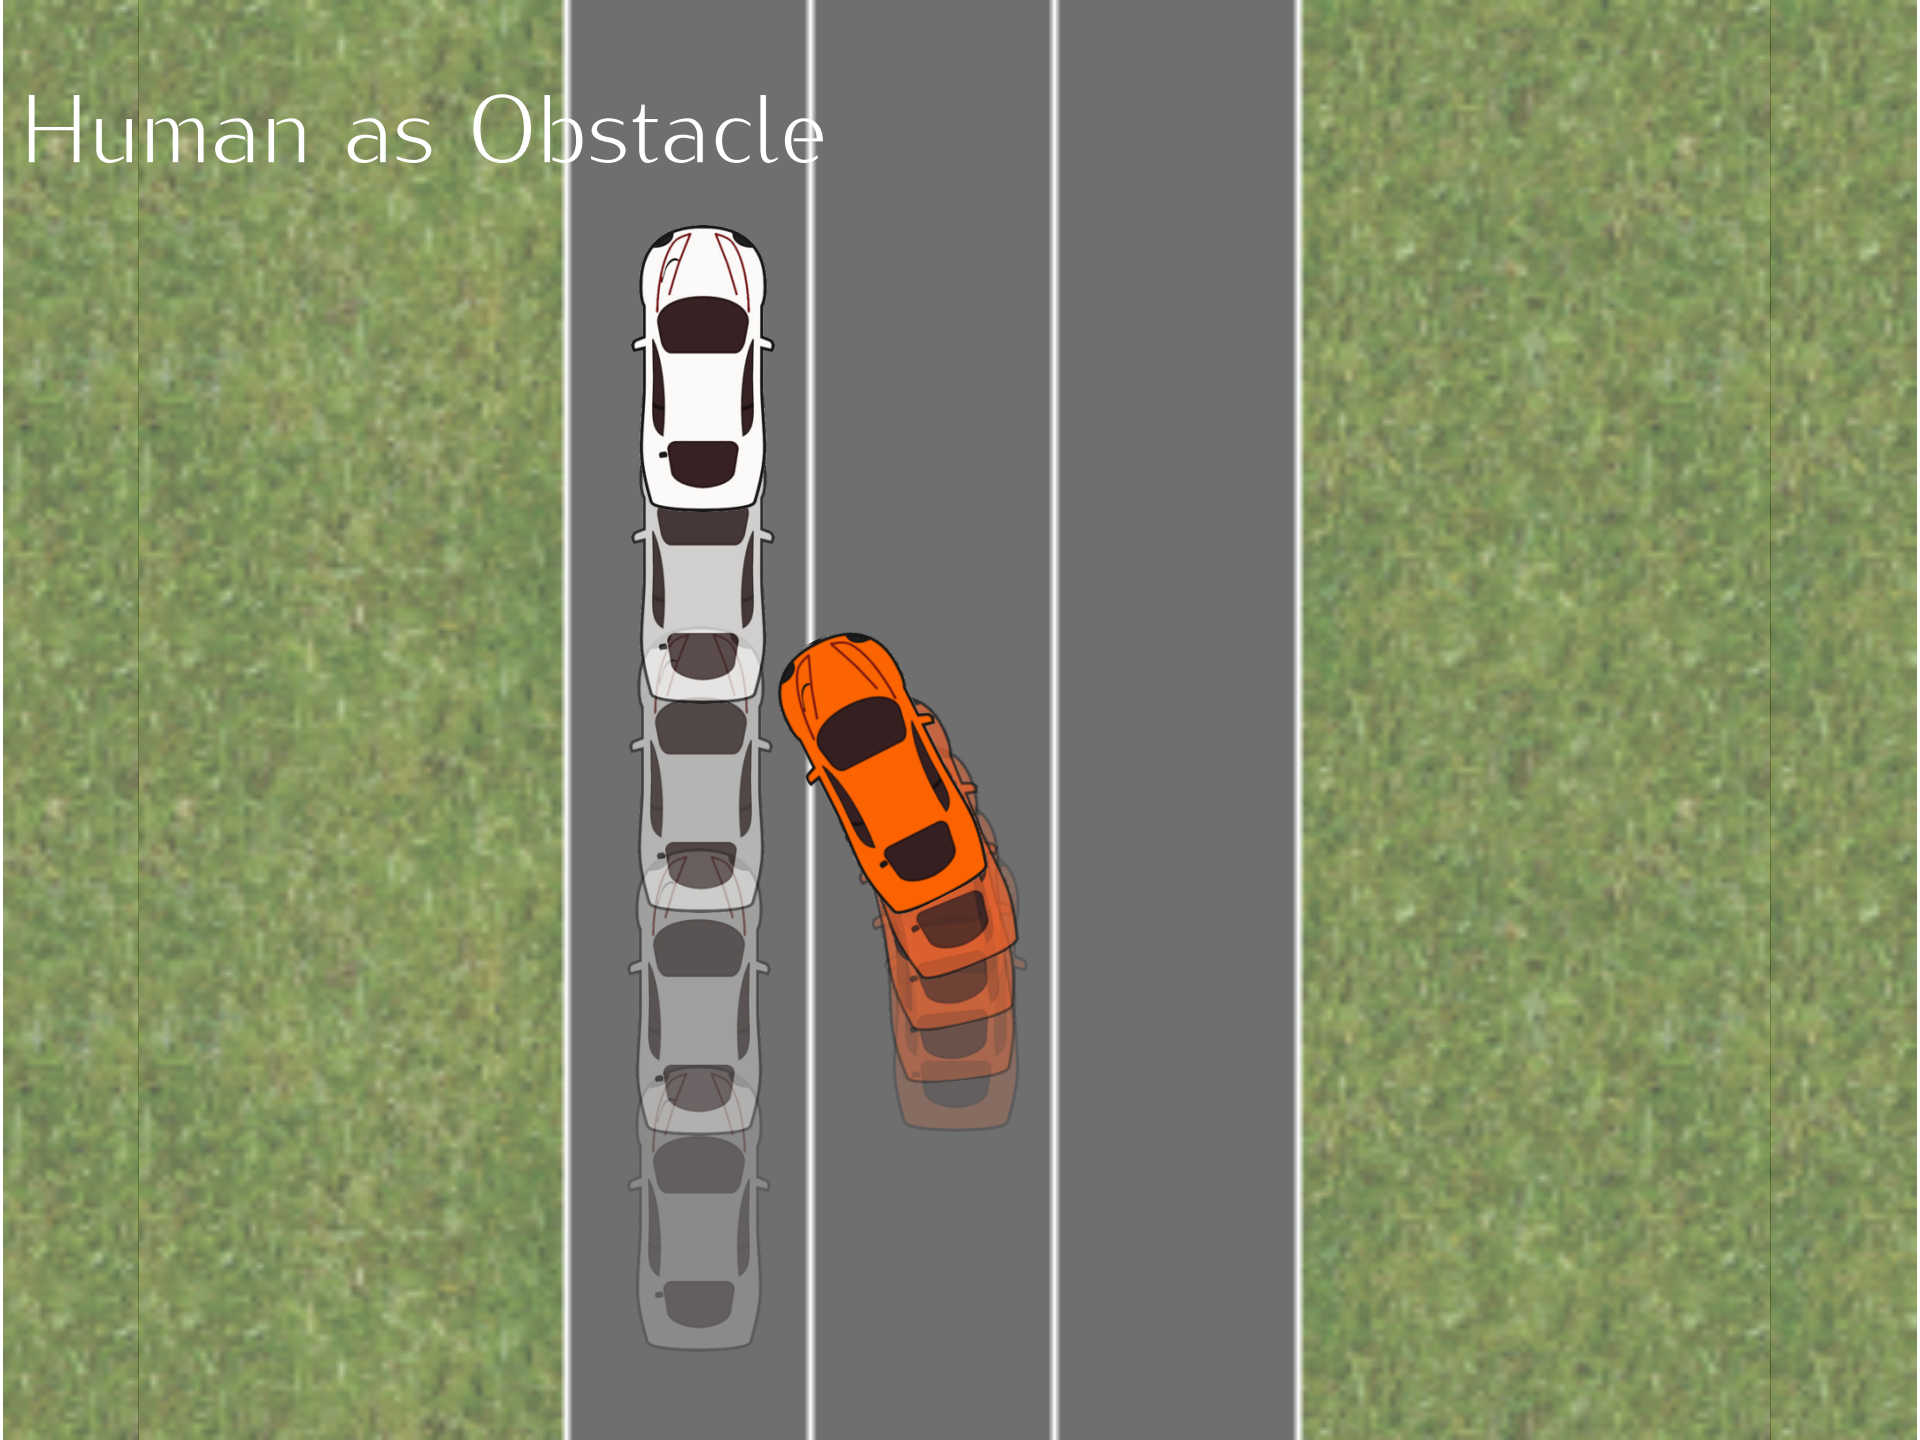

Underactuated System, U_H Learned Offline

Approximation as Underactuated System

$U_H(x, u_R, u_H; \theta_H)$ $U_R(x, u_R, u_H)$

θ_H $\mathbf{u}_H^* = \arg\max_{\mathbf{u}_H} U_H(x, \mathbf{u}_R, \mathbf{u}_H; \theta_H)$

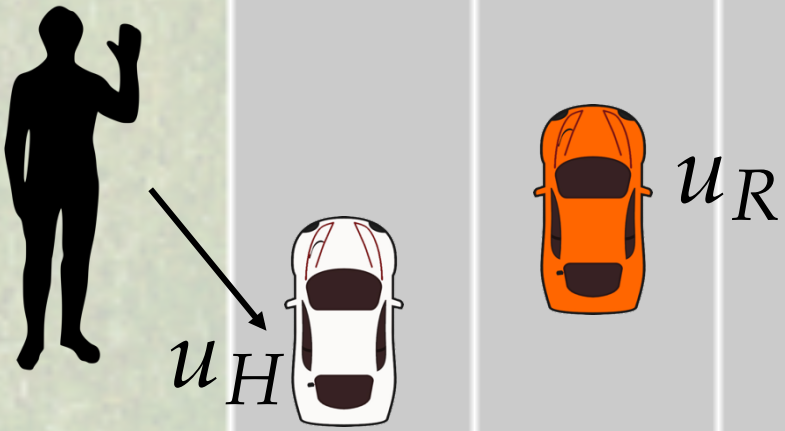

Planning for Autonomous Cars that Leverage Effects on Human Actions [RSS'16]

Also useful in collaboration, when H is myopic!

$$U(x, u_R, u_H; \theta_H)$$

θ_H $u_H^* = \arg\max U_H(x, u_R, u_H; \theta_H)$

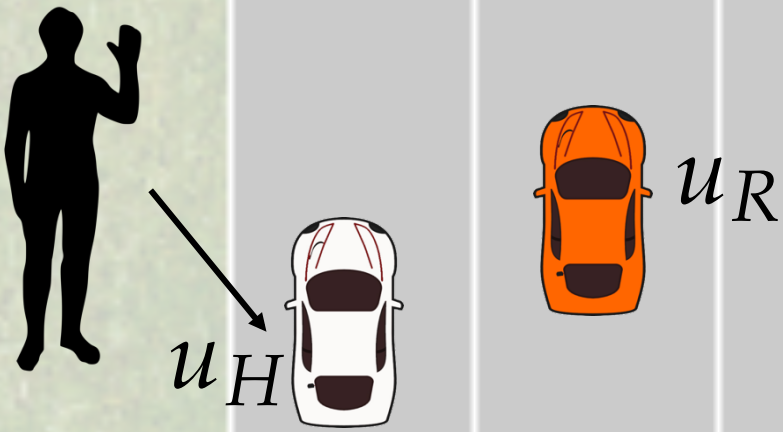

Guide People to Better Plans in Collaboration

citly Assisting Humans to
se Good Grasps [ISER'16]


Big Limitation

Underactuated System, Online U_H

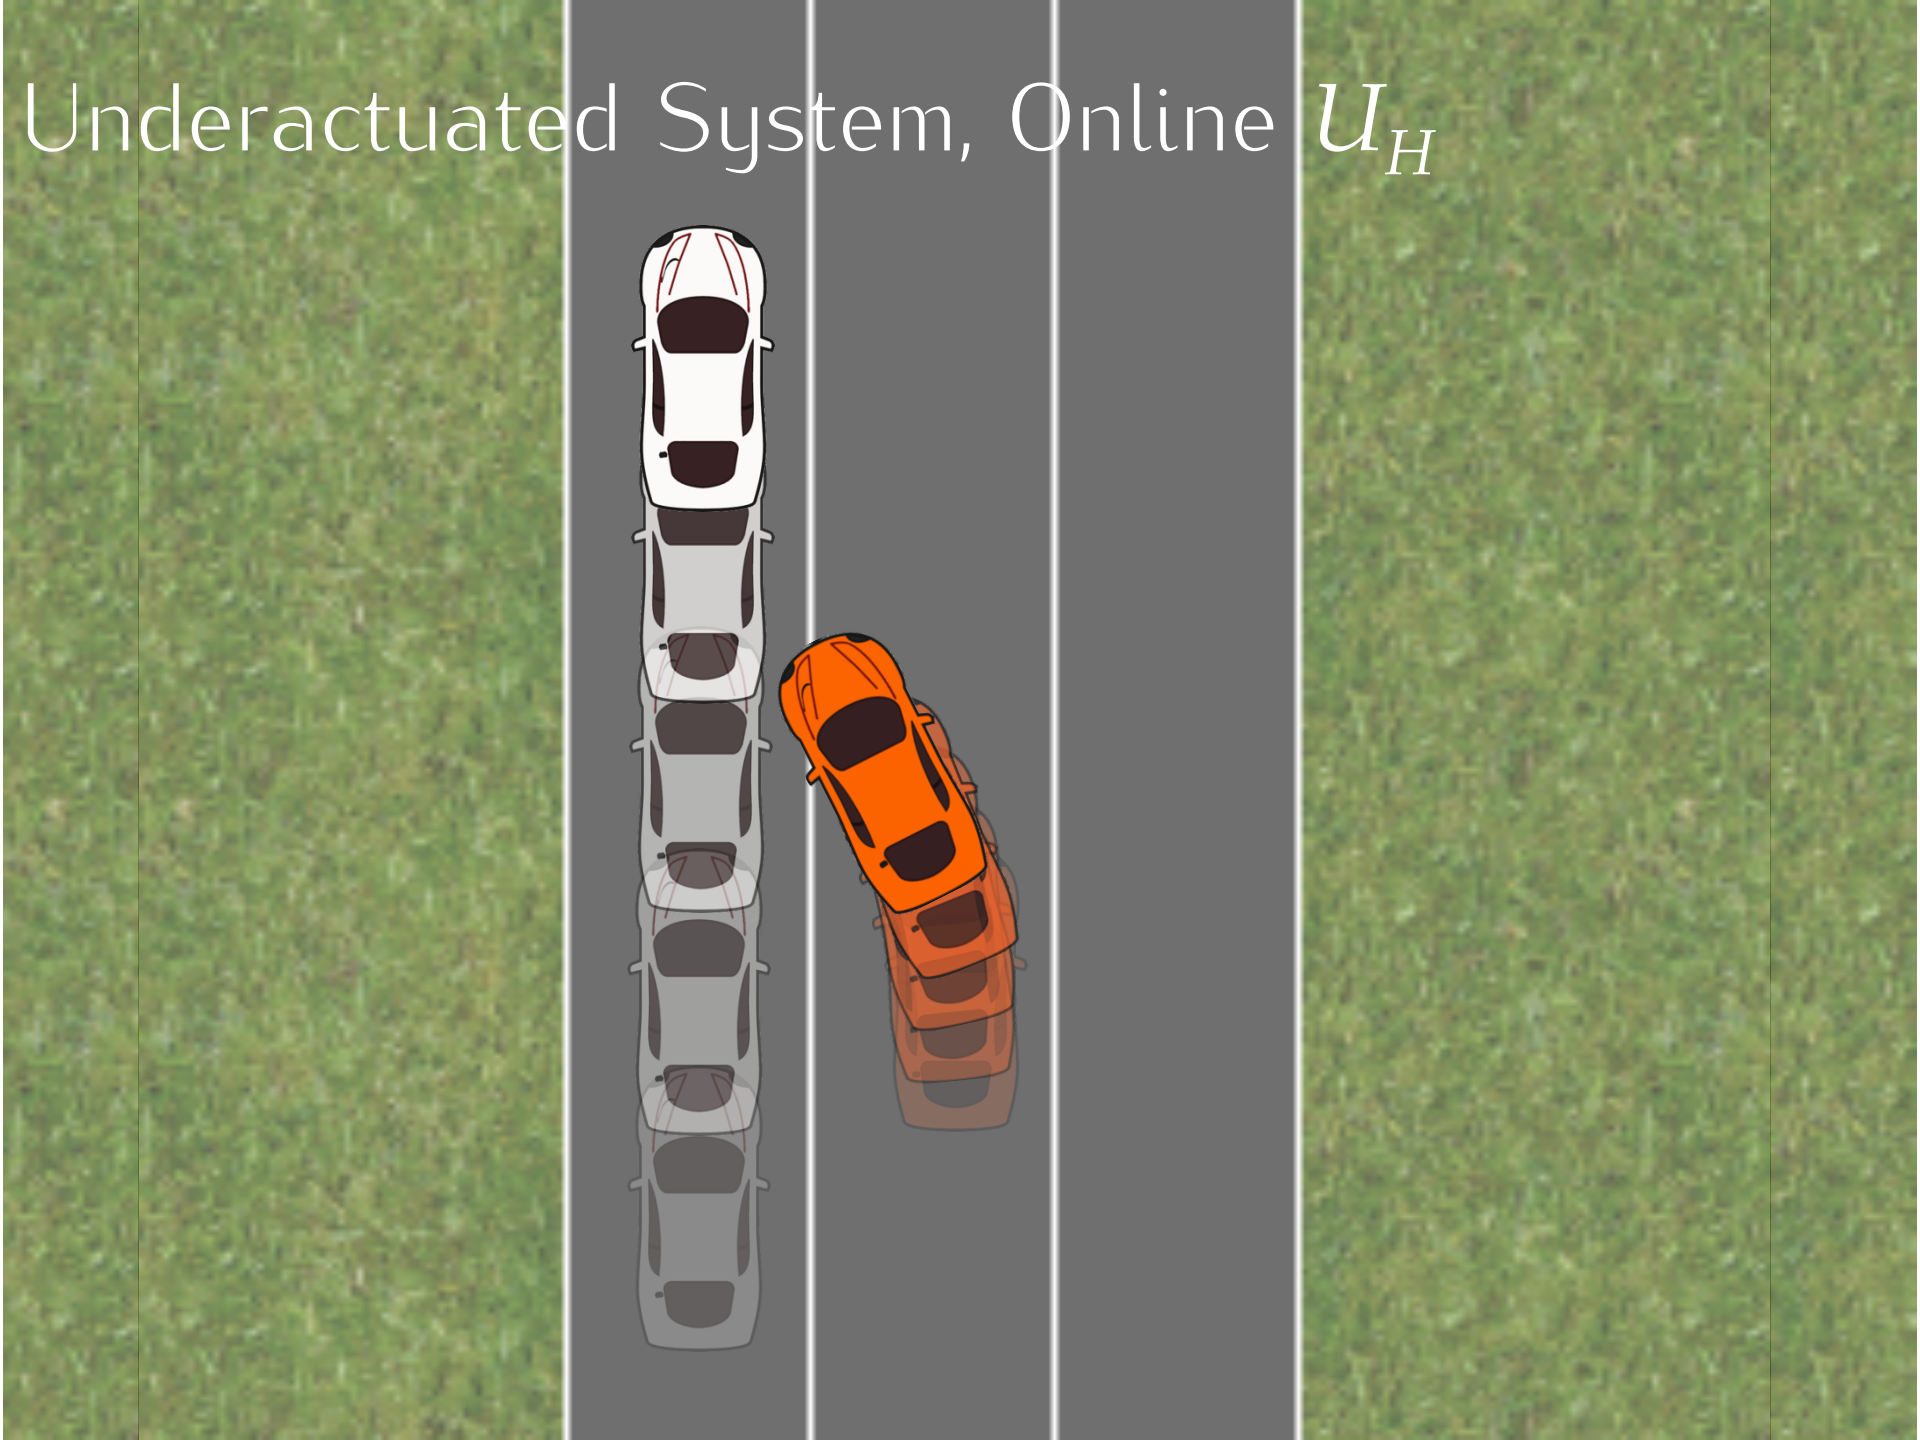

All Users React in Almost the Same Way

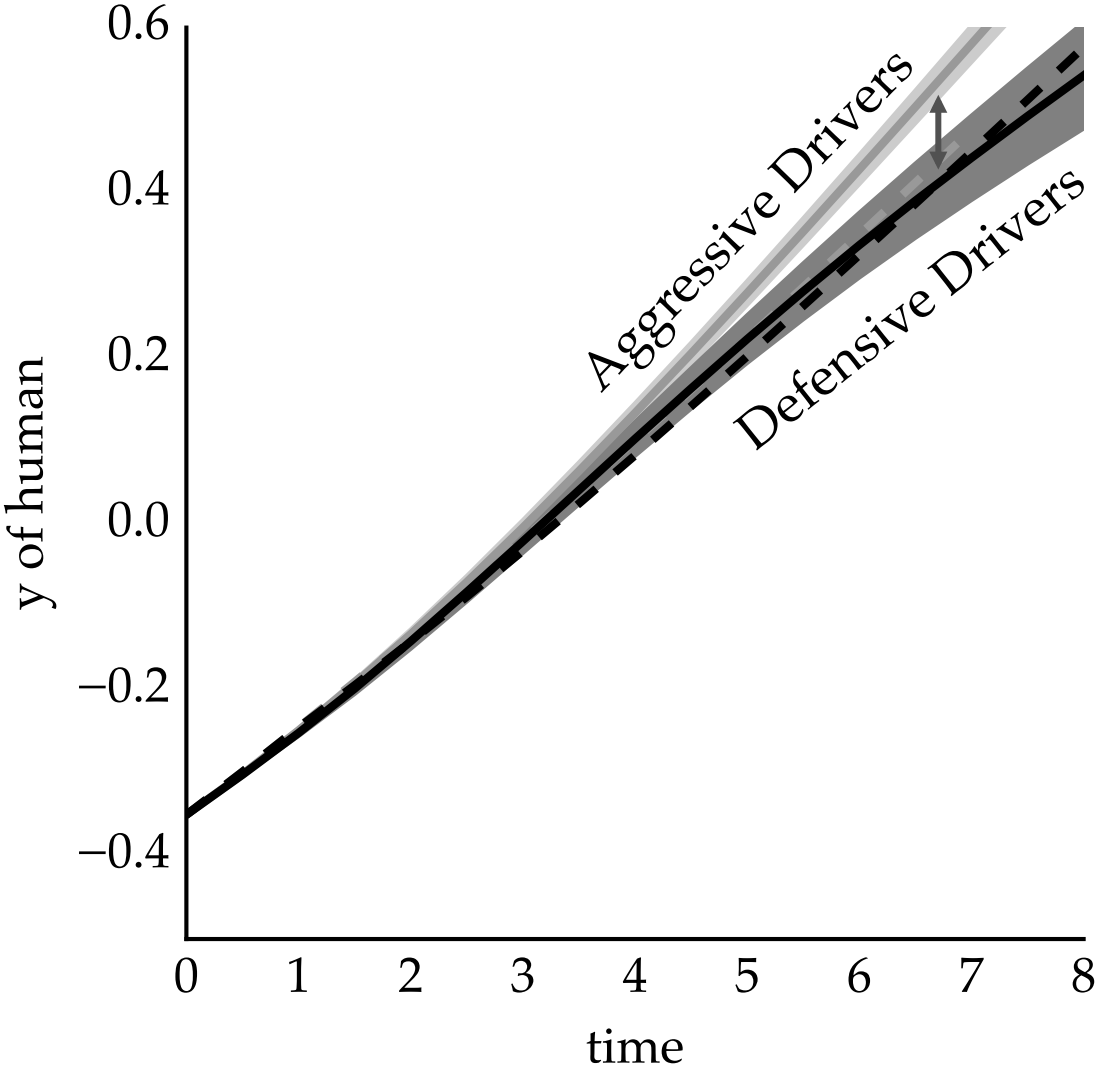

Approximation as Underactuated System

$U_H(x, u_R, u_H; \theta_H)$ $U_R(x, u_R, u_H)$

θ_H $\mathbf{u}_H^* = \arg\max_{\mathbf{u}_H} U_H(x, \mathbf{u}_R, \mathbf{u}_H; \theta_H)$

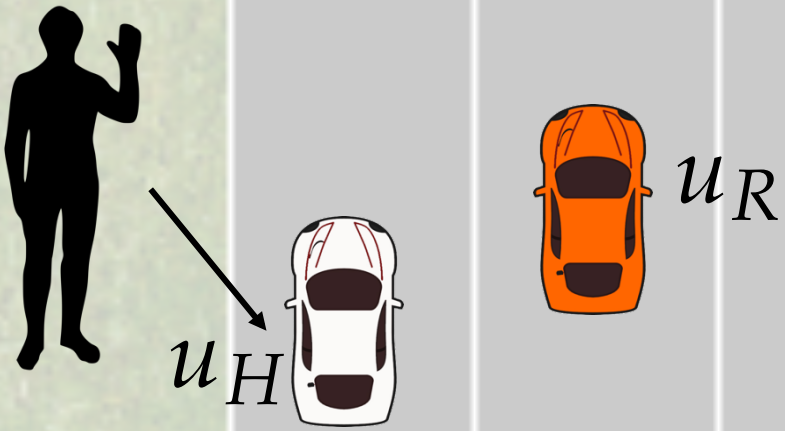

Active Information Gathering over Human Internal State [IROS'16]
[best cognitive robotics paper finalist]

Idea: Leverage the robot's actions!

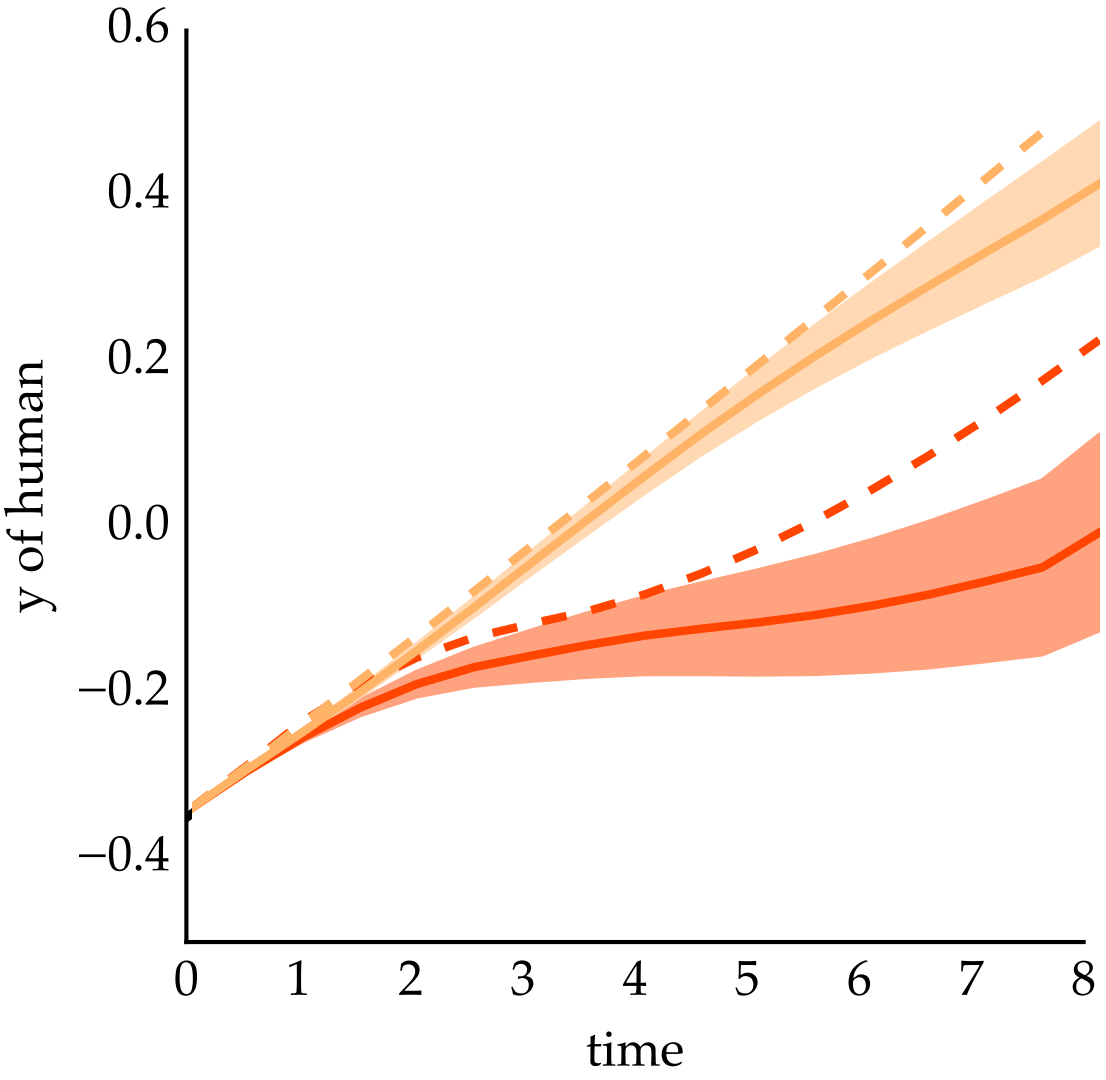

Underactuated System, Active u_H

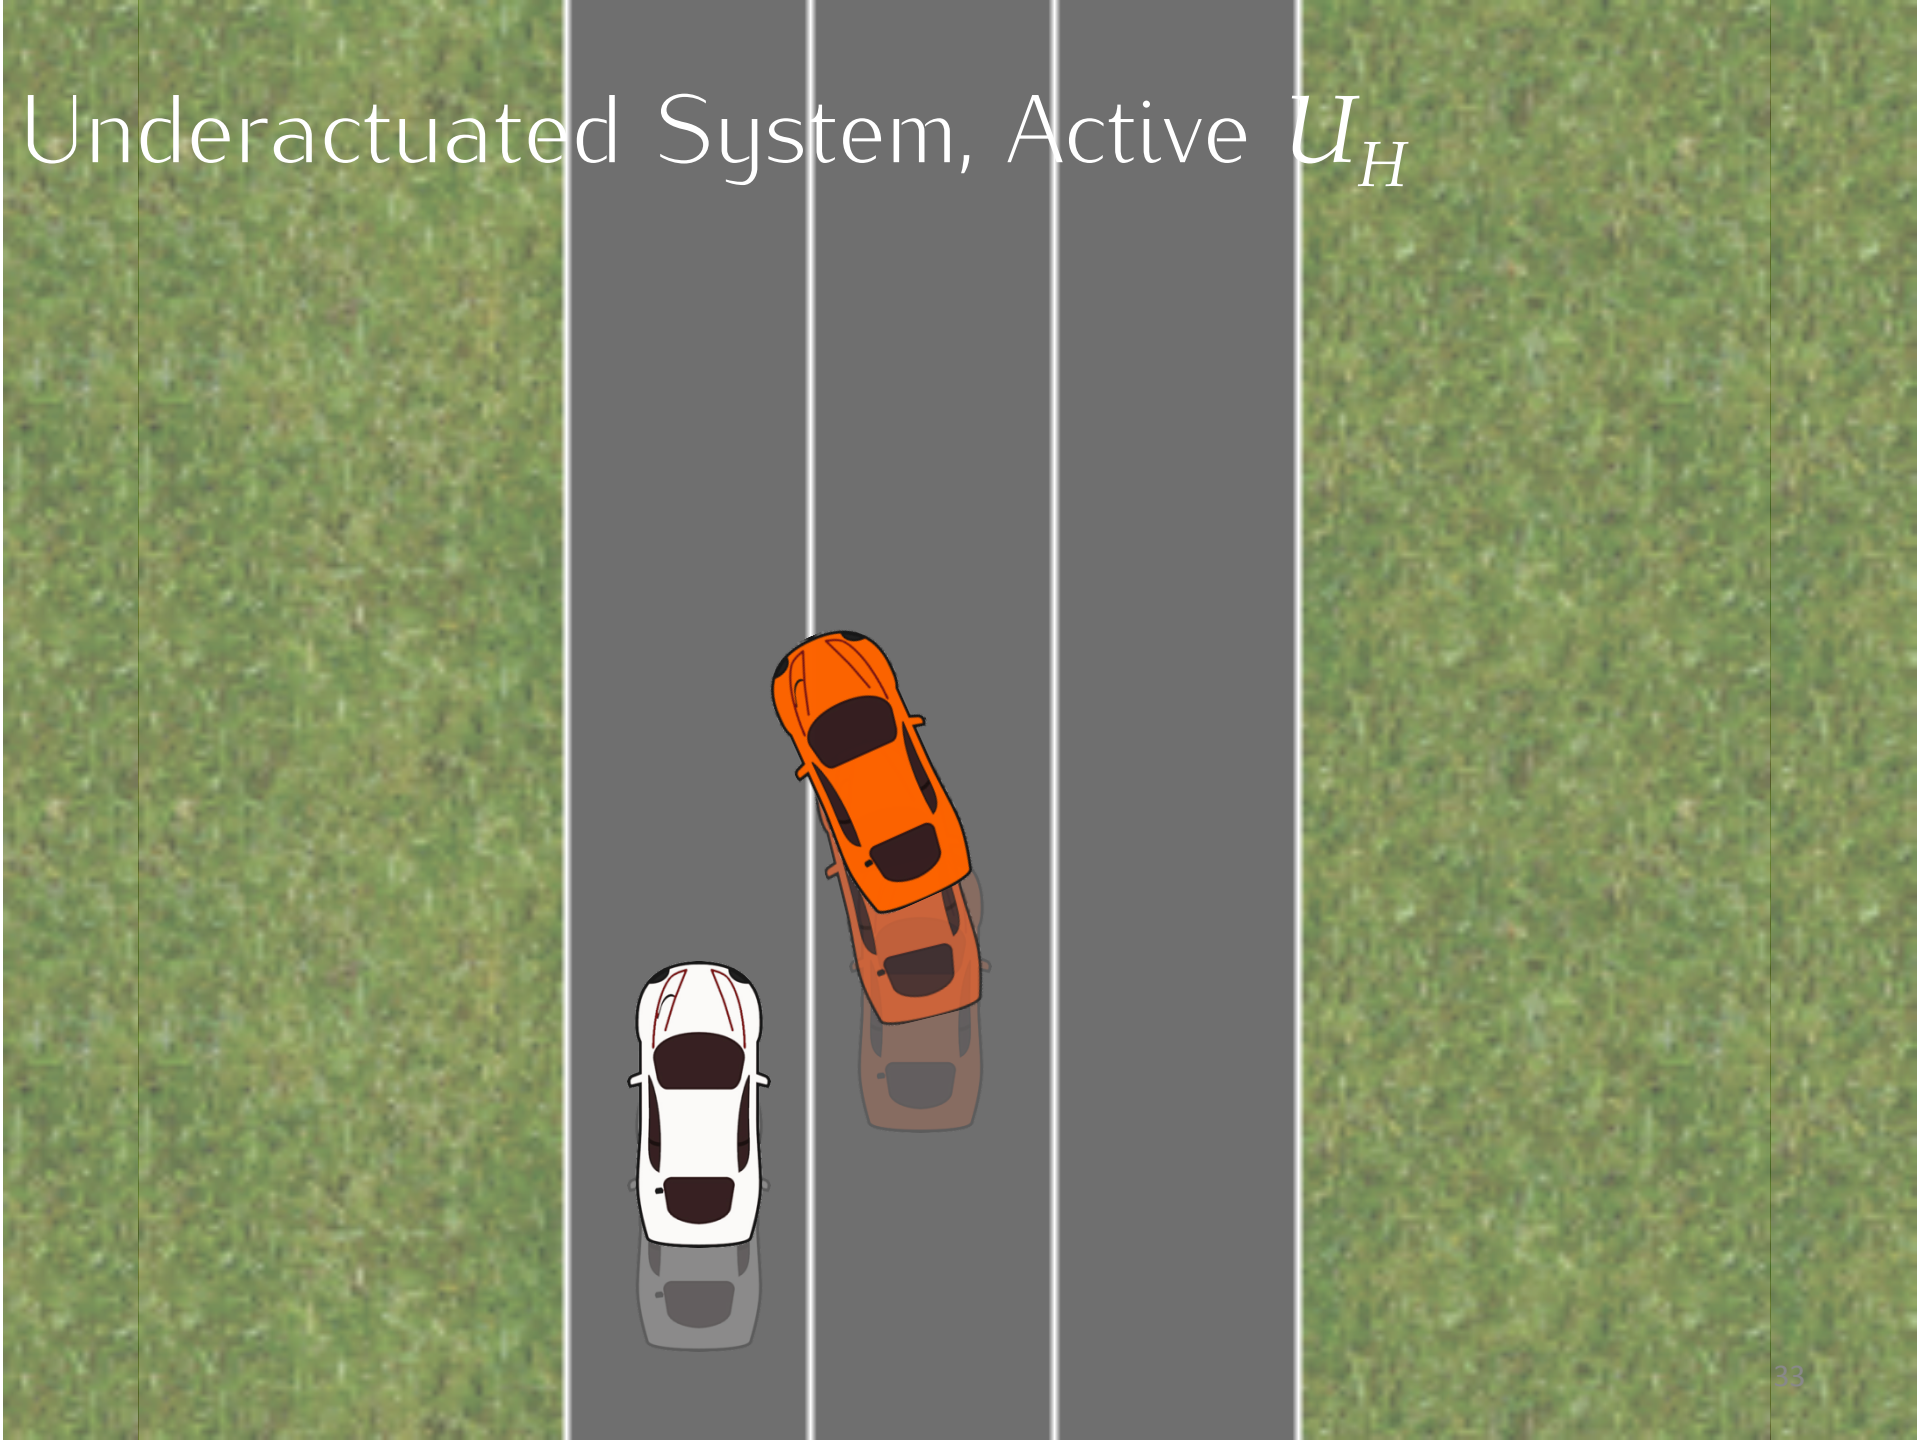

Underactuated System, Active U_H - *Defensive*

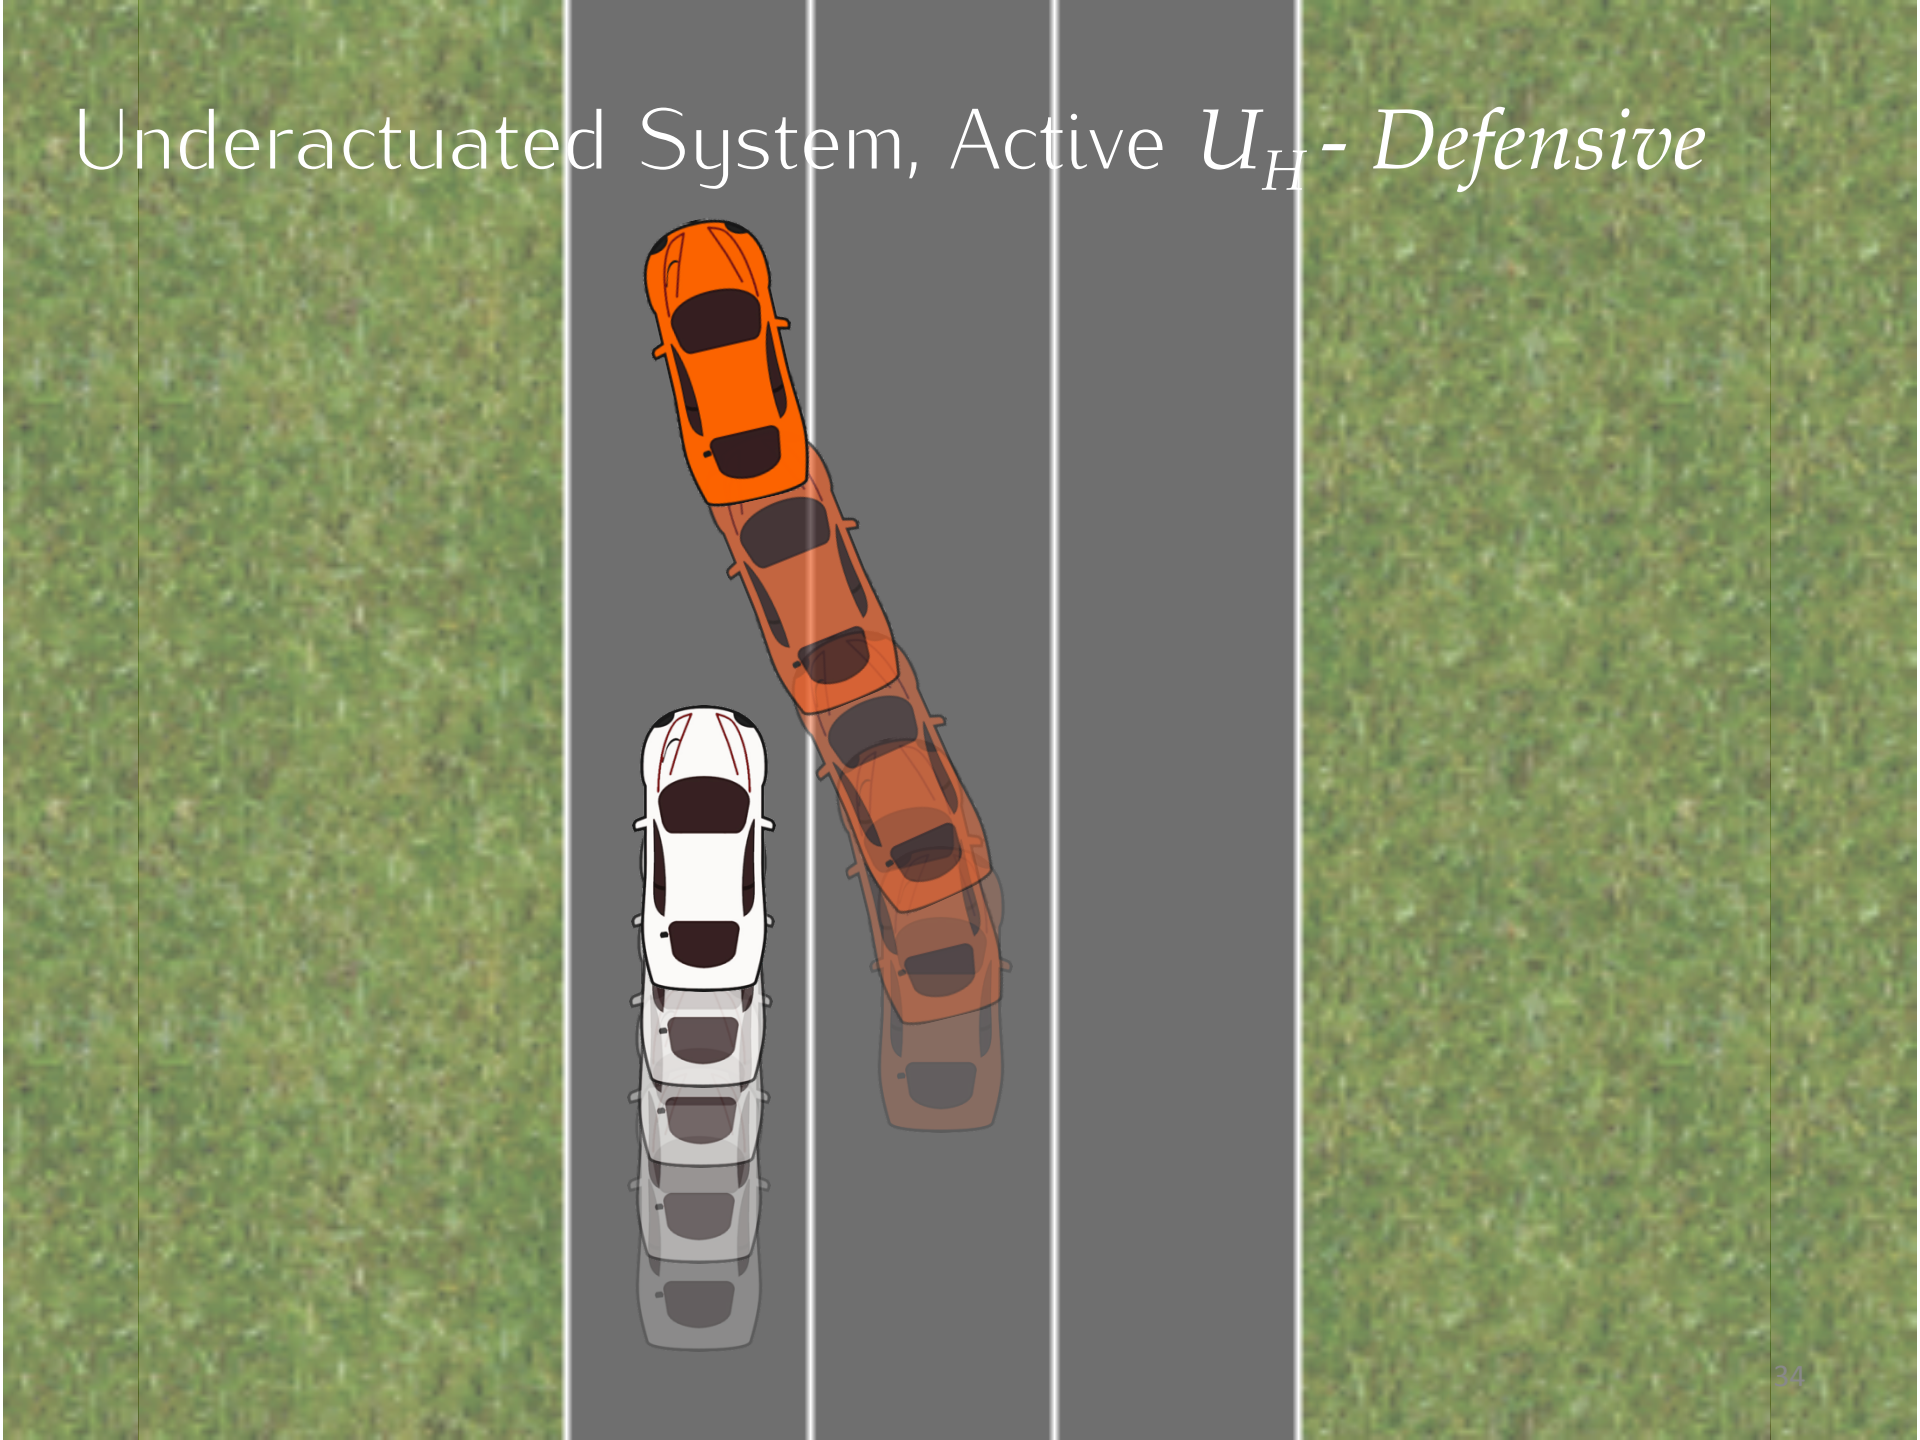

Underactuated System, Active u_H - *Aggressive*

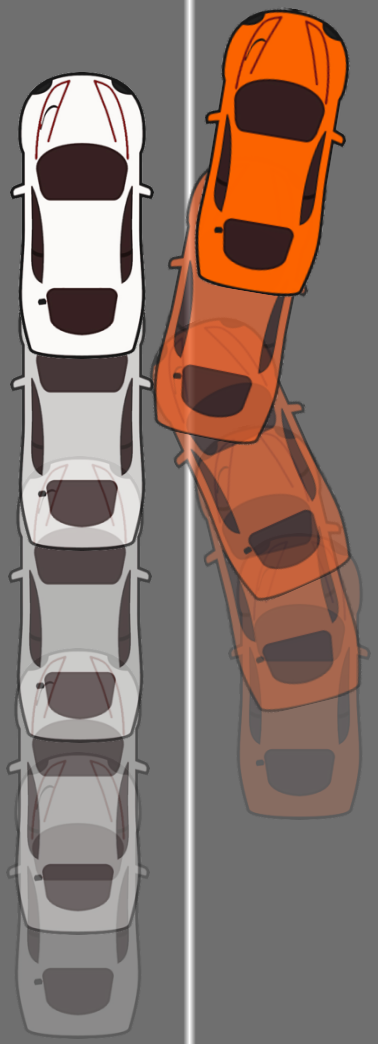

Coordination at 4-Way Stops

Robot Trajectories

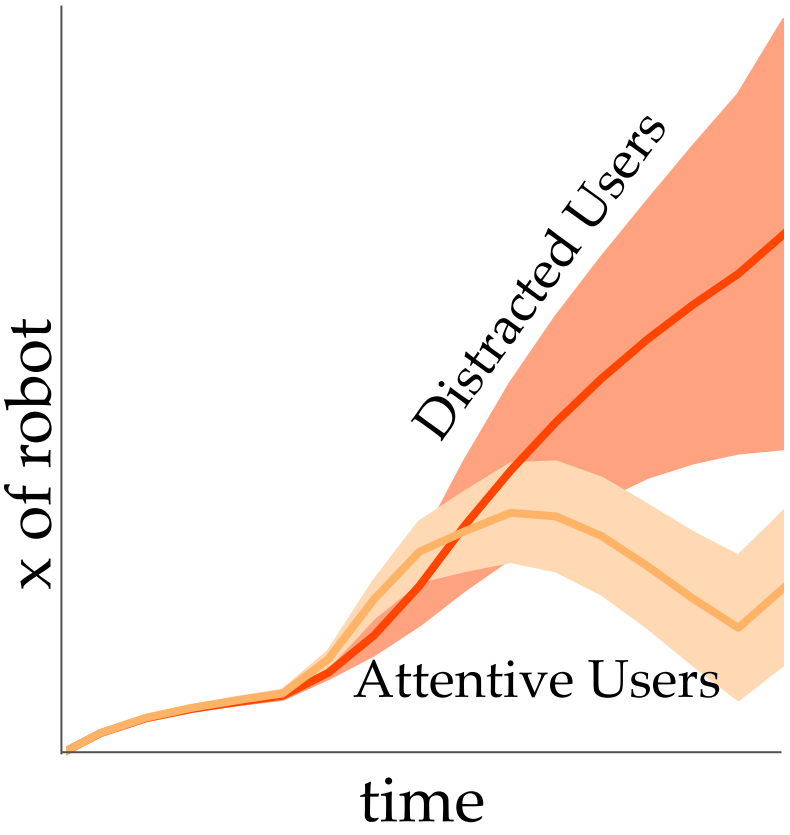

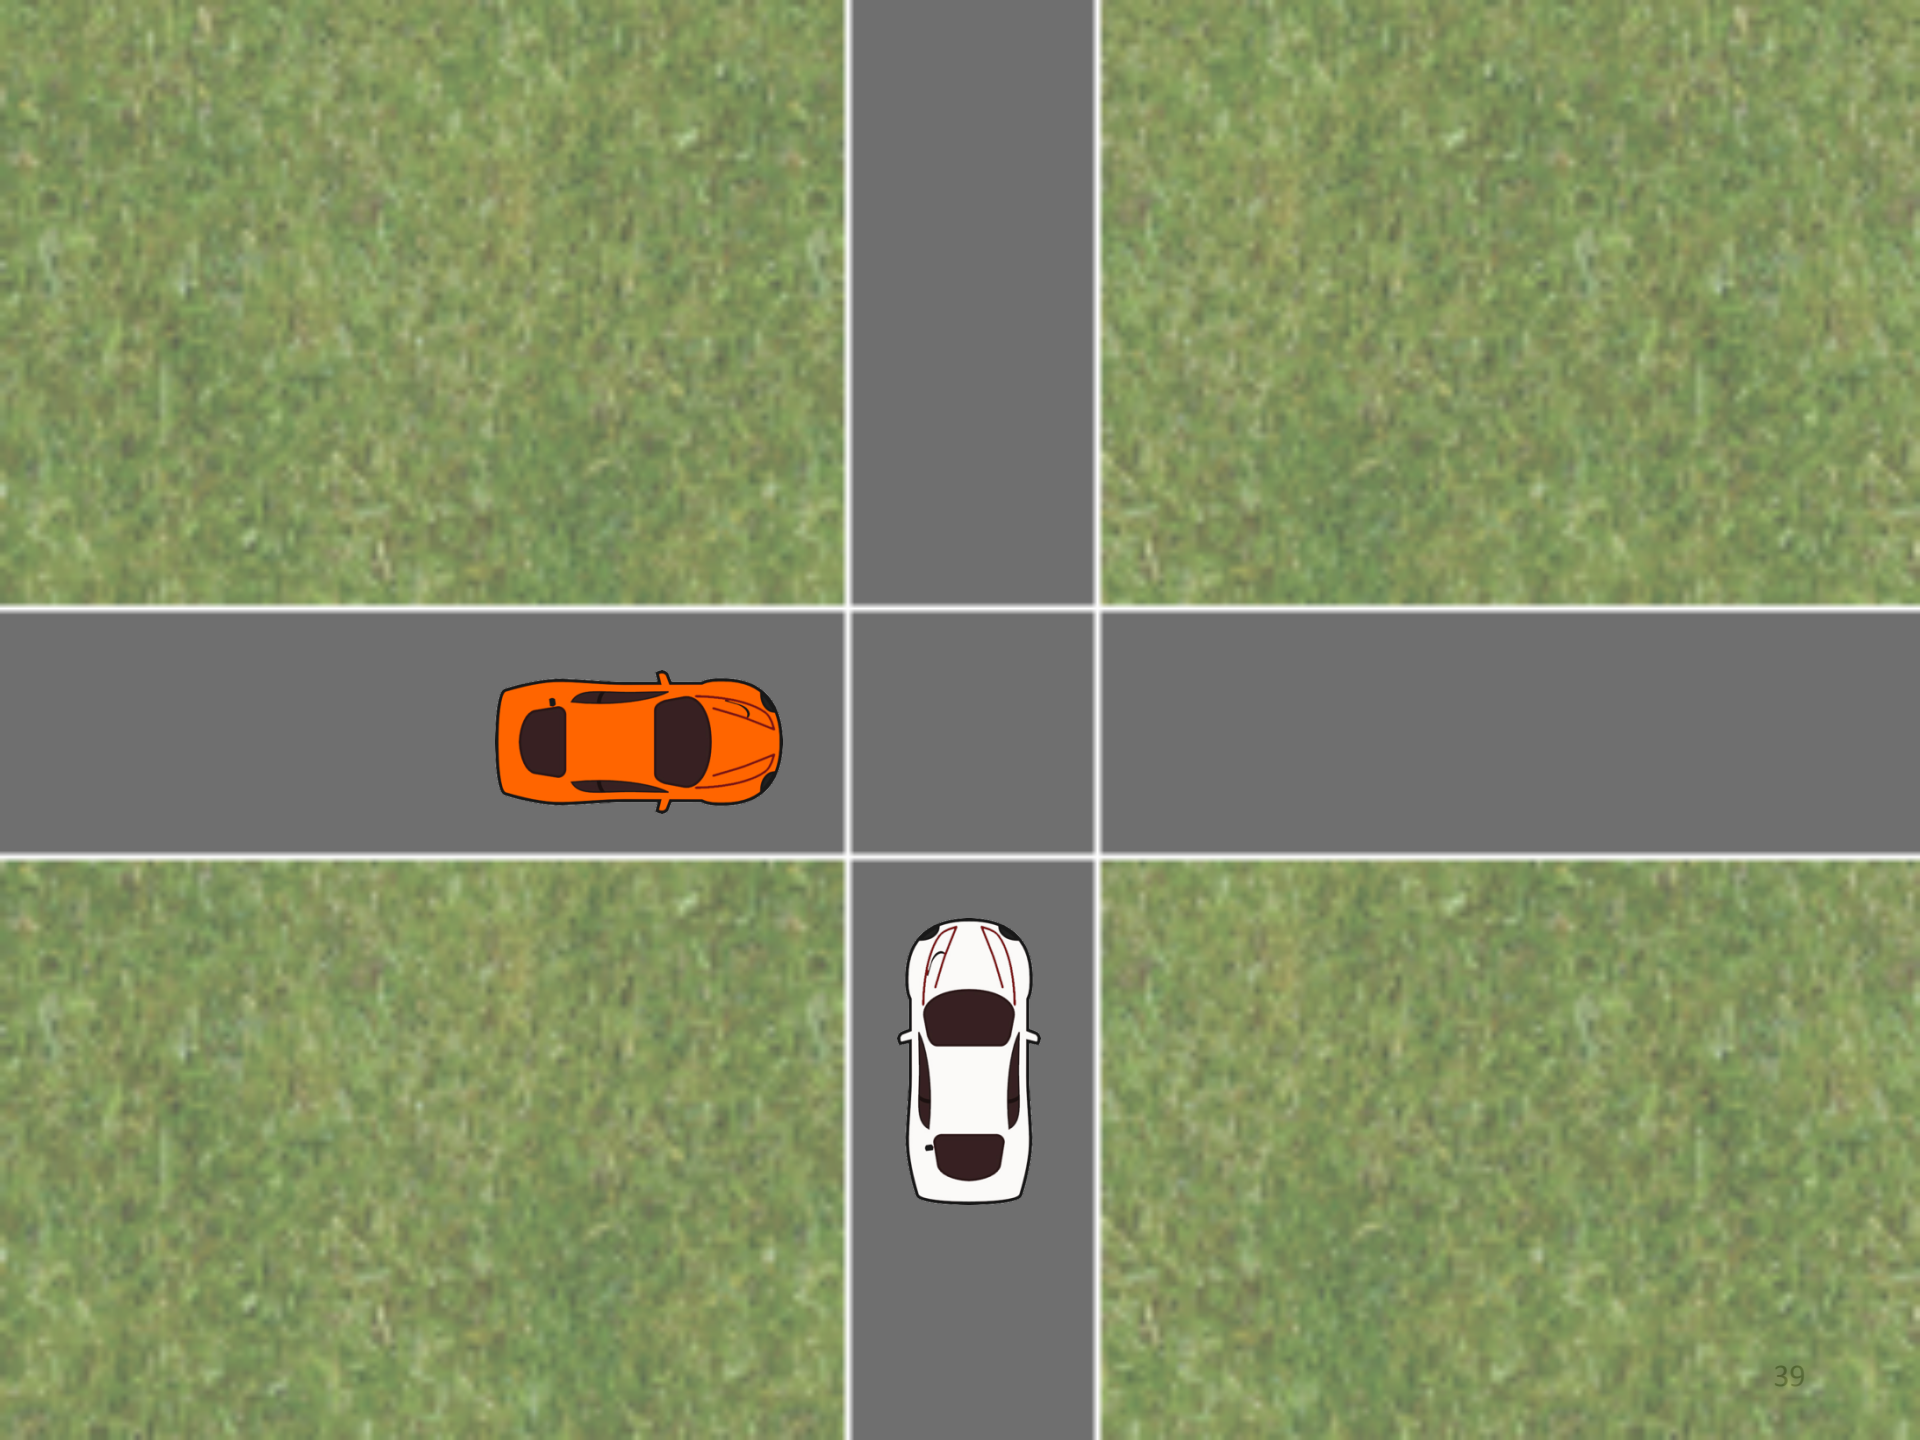

Inch Forward

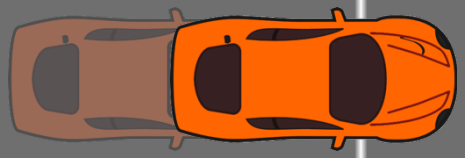

Google's fleet of autonomous test cars is programmed to follow the letter of the law. But it can be tough to get around if you are a stickler for the rules. One Google car, in a test in 2009, couldn't get through a four-way stop because its sensors kept waiting for other (human) drivers to stop completely and let it go. The human drivers kept inching forward, looking for the advantage — paralyzing Google's robot.

"Google's Driverless Cars Run Into Problem: Cars With Drivers"[Richtel&Dougherty]

Info Gathering Improves Estimation

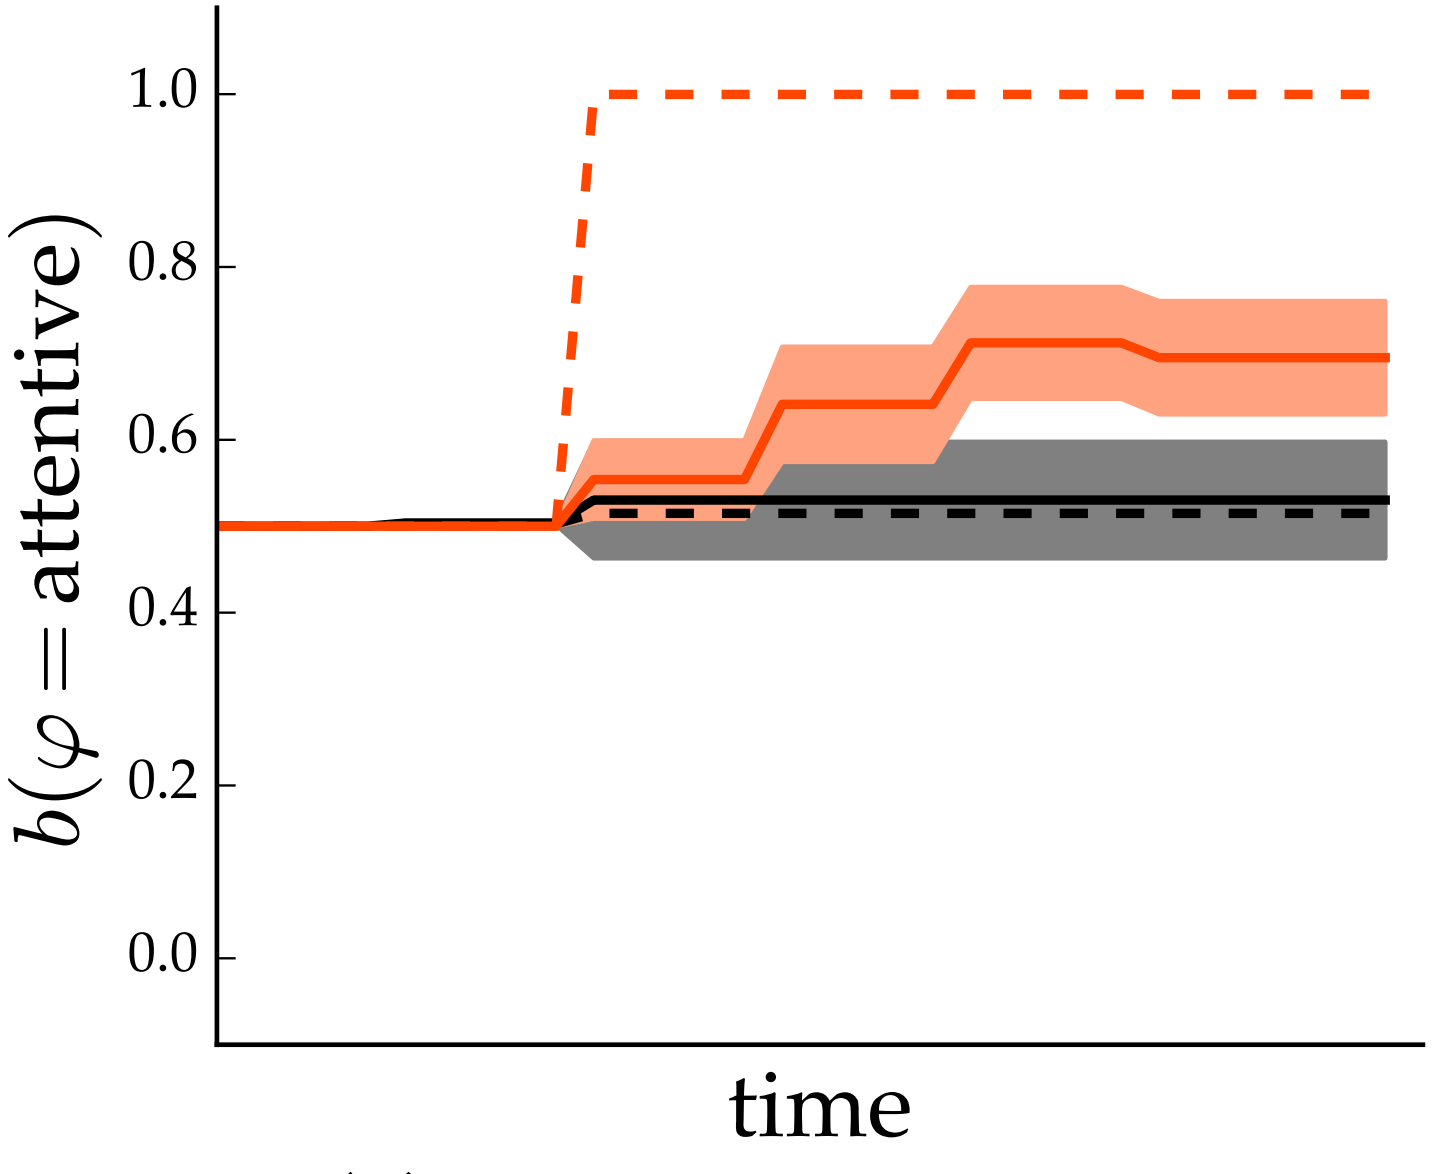

Attentive Users: Continue

Inch Forward

Distracted Users: Go Back

$\mathcal{U}_R$: Human Should Go First

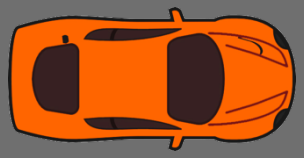

$\mathcal{U}_R$: Human Should Go First

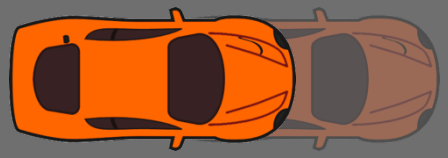


Are you going? Or should I go?
What if I point a lot and flail my arms around?
Wait, maybe you should go.
You go first.
This is confusing.
Let's just sit here and reflect.

Real Effects on Real Users' Actions

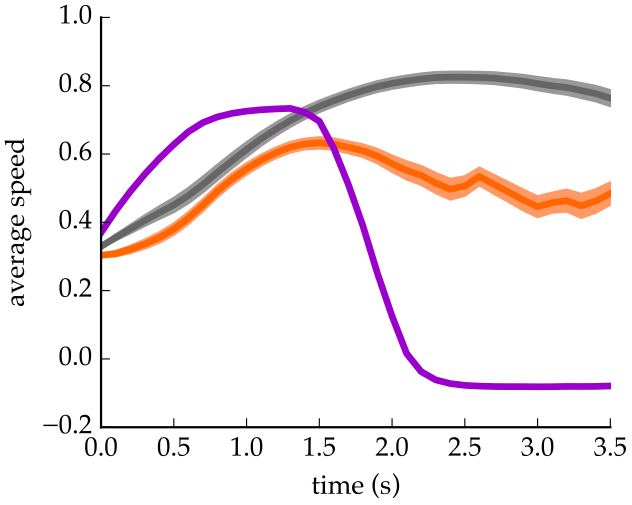

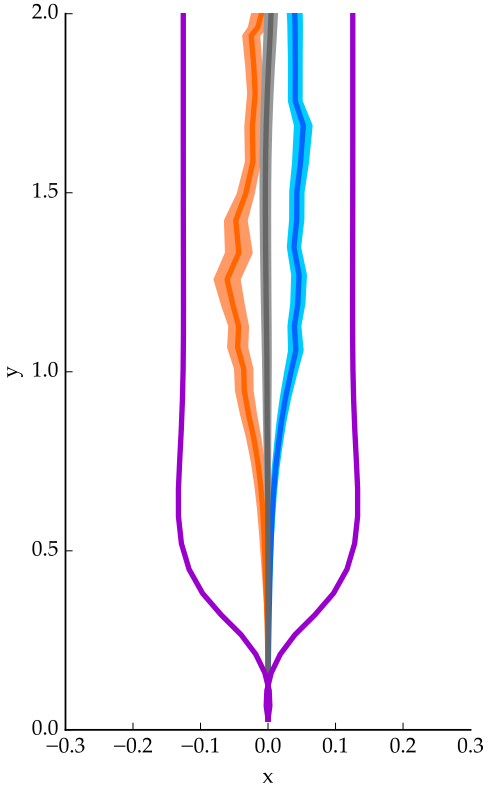

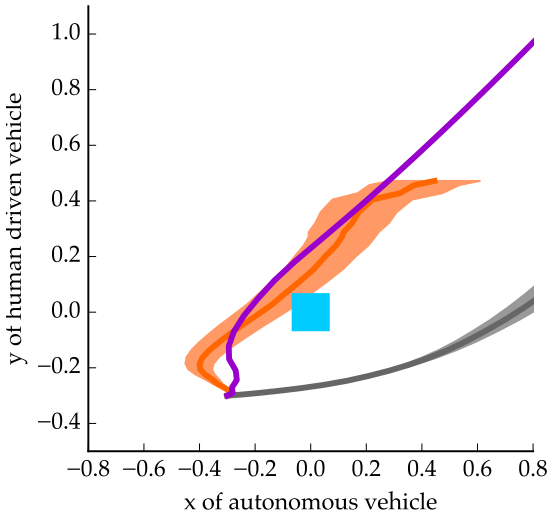

Learned Human Model Avoid Human Affect Human

Google's fleet of autonomous test cars is programmed to follow the letter of the law. But it can be tough to get around if you are a stickler for the rules. One Google car, in a test in 2009, couldn't get through a four-way stop because its sensors kept waiting for other (human) drivers to stop completely and let it go. The human drivers kept inching forward, looking for the advantage — paralyzing Google's robot.

Working interaction strategies emerge out of the optimization.

A Partial Information Human-Robot Game

$$U(x, u_R, u_H; \theta_H)$$

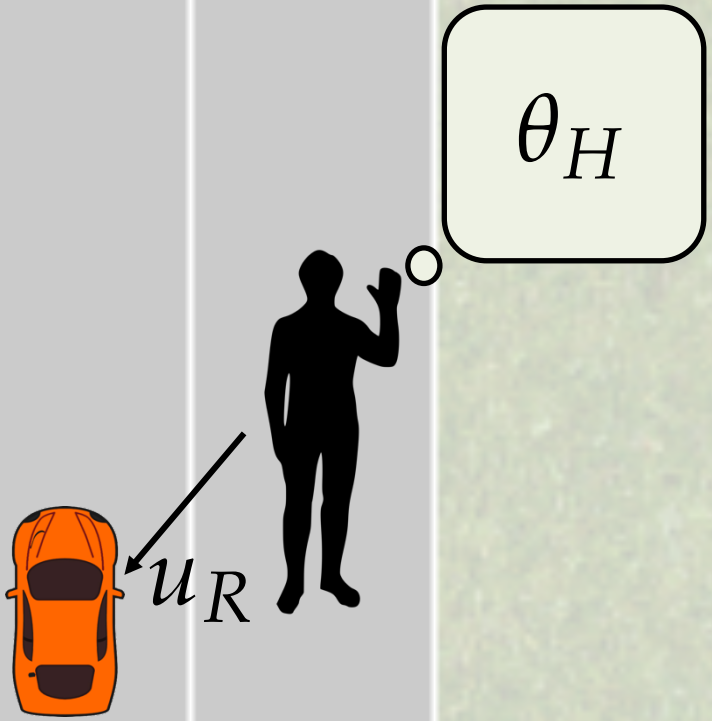

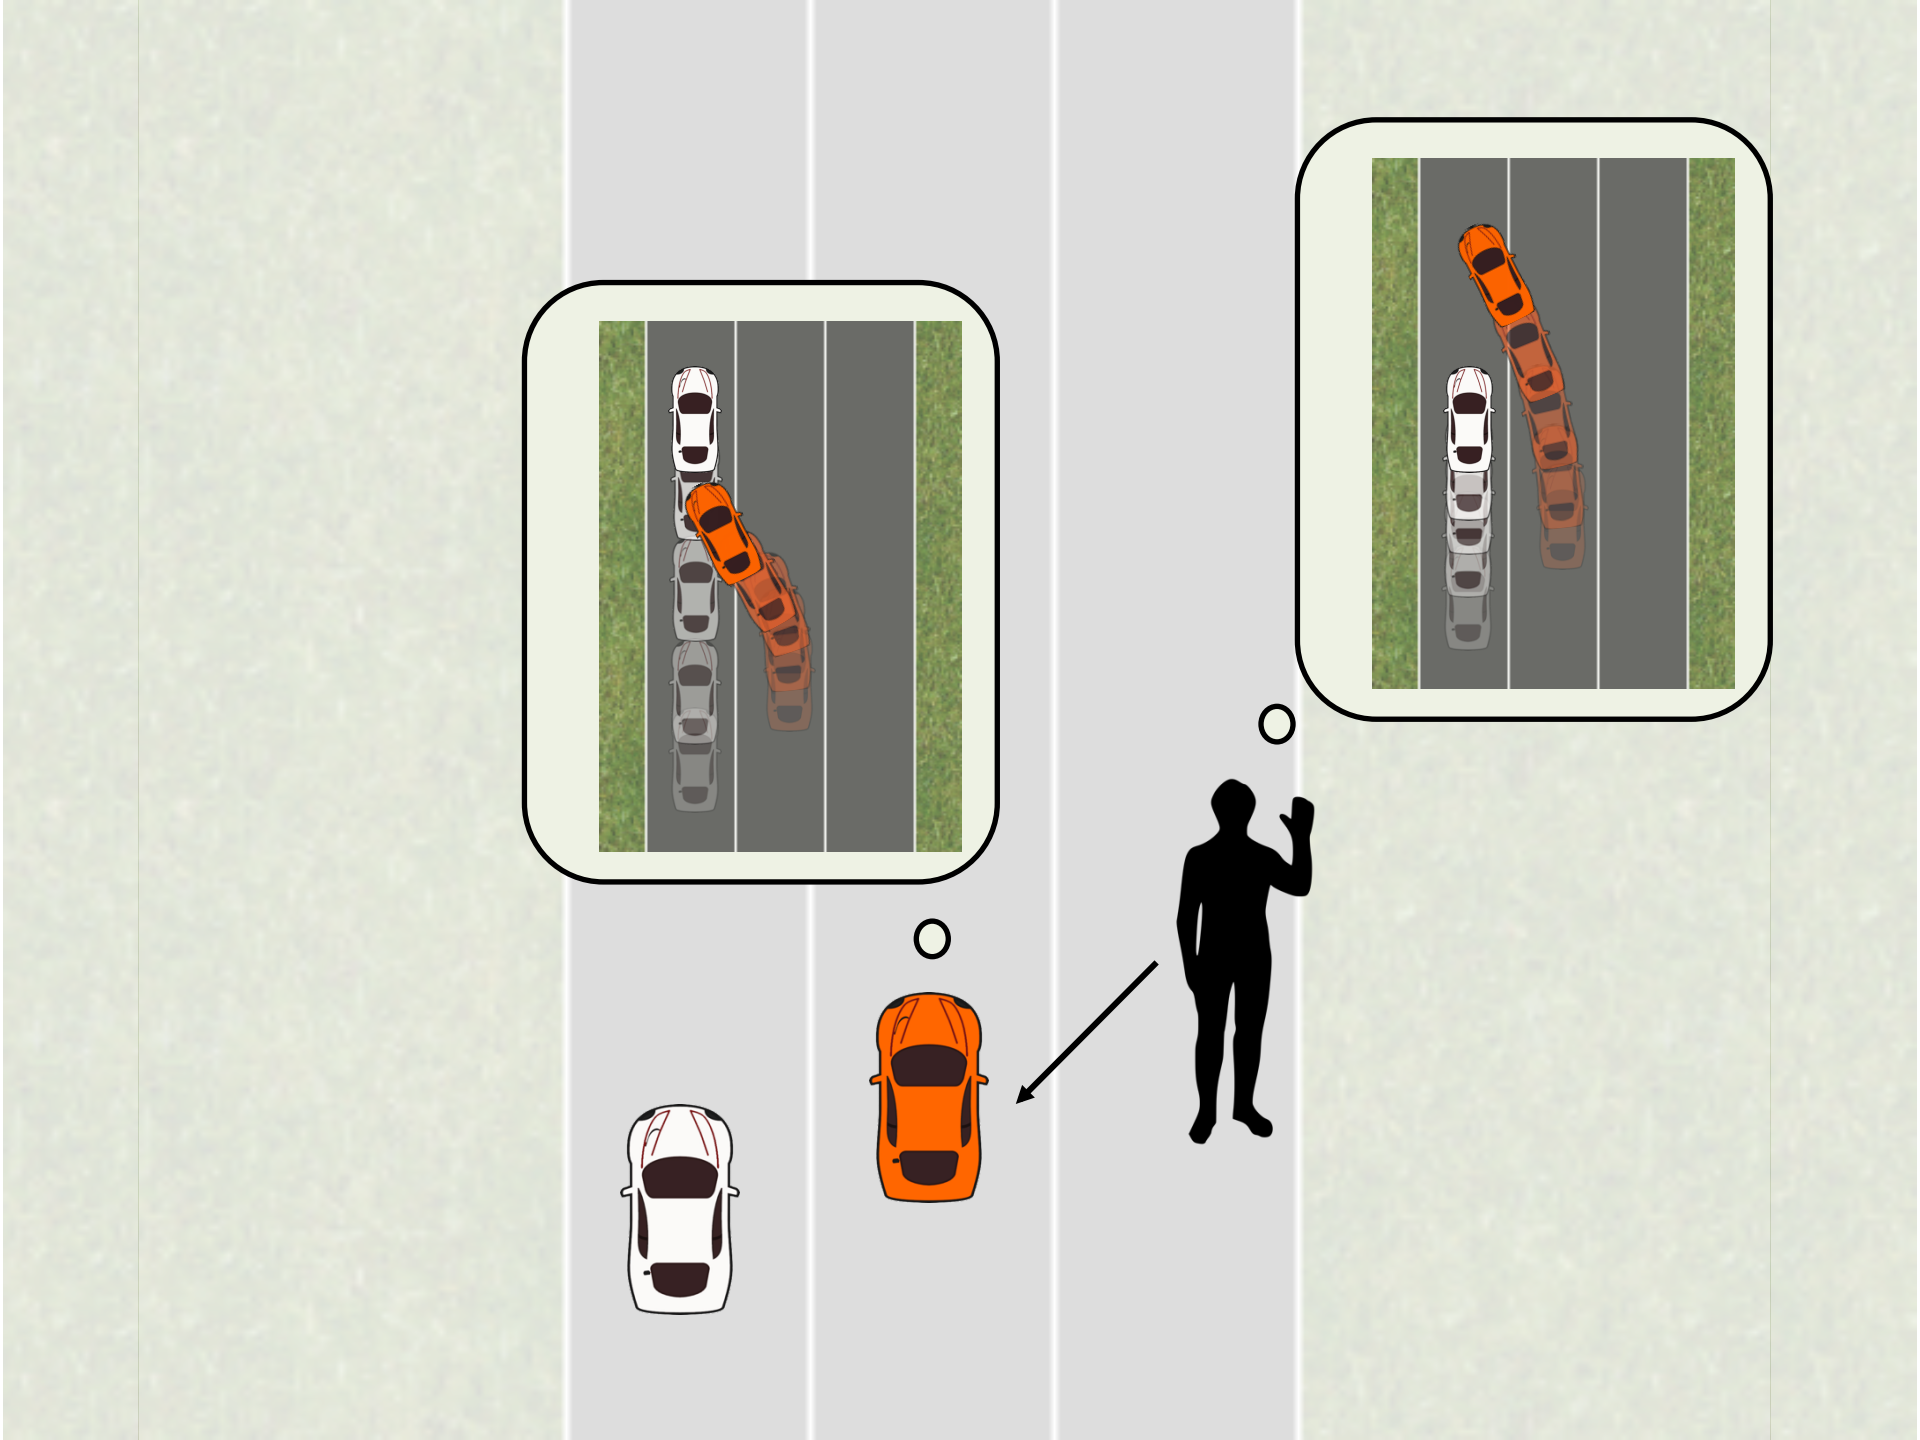

Use human actions as observations!

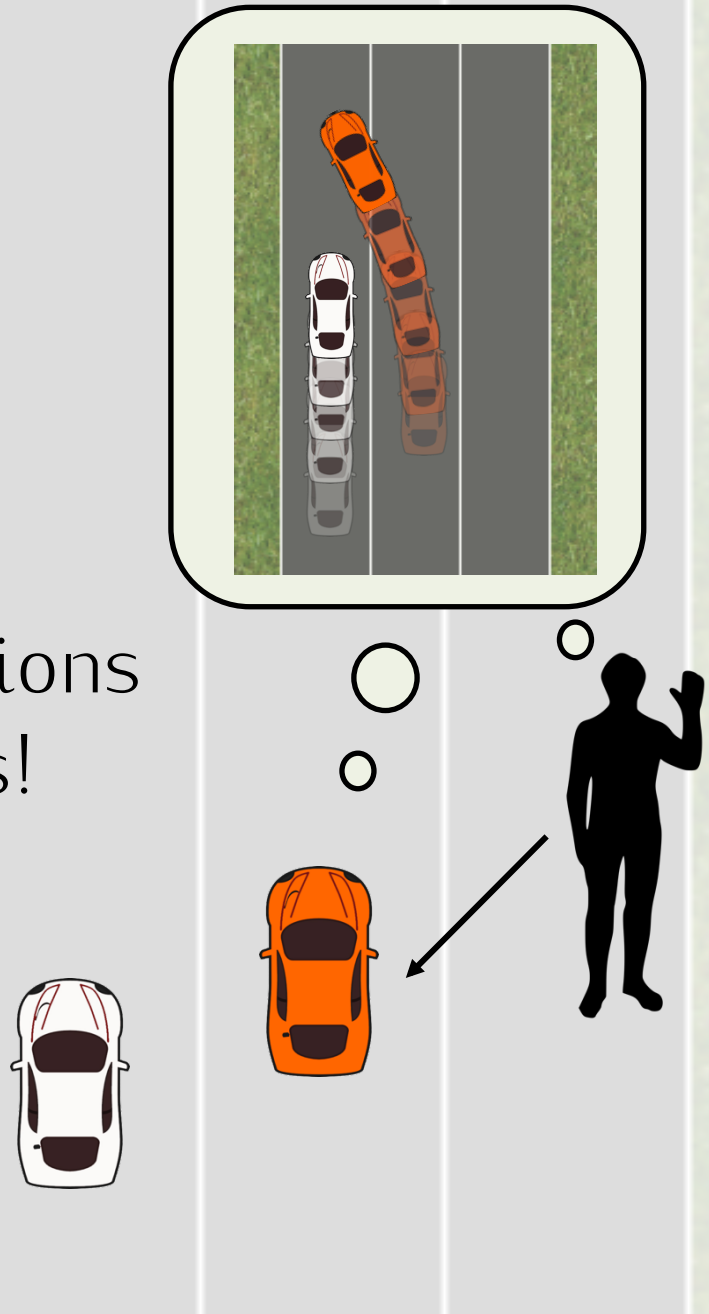

User Study (in Driving Simulator)

Manipulated Factor: Driving Style

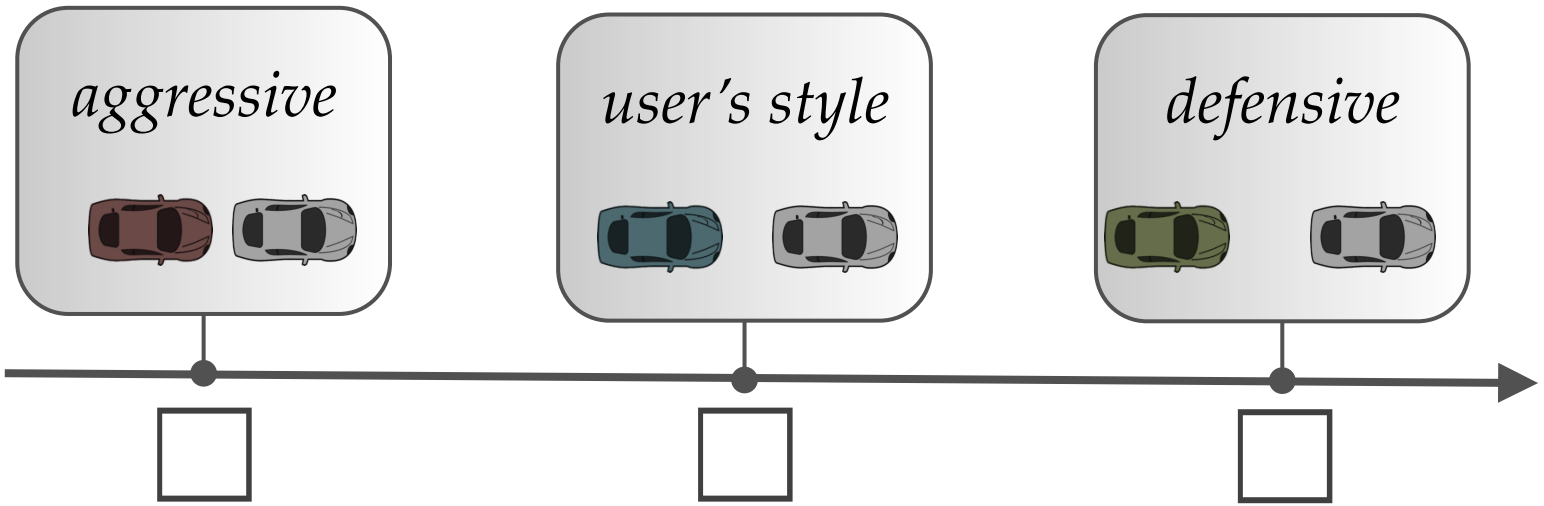

Dependent Measures

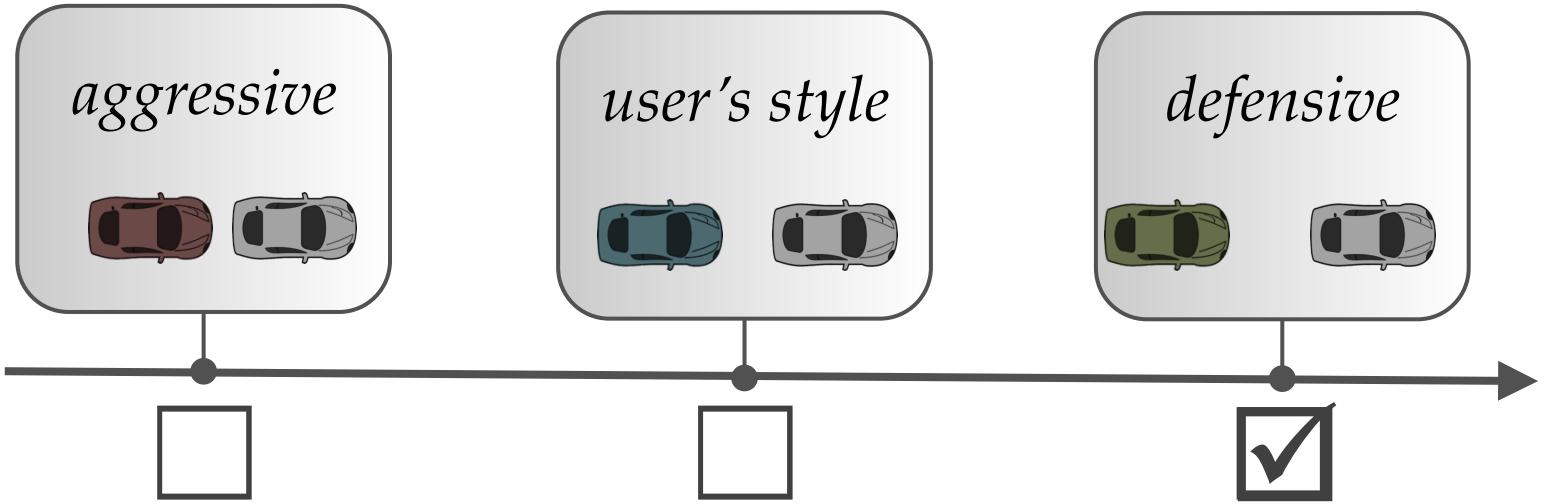

1. Preference for daily use
2. Perceived similarity with own driving style

Own driving style vs. preferred driving style

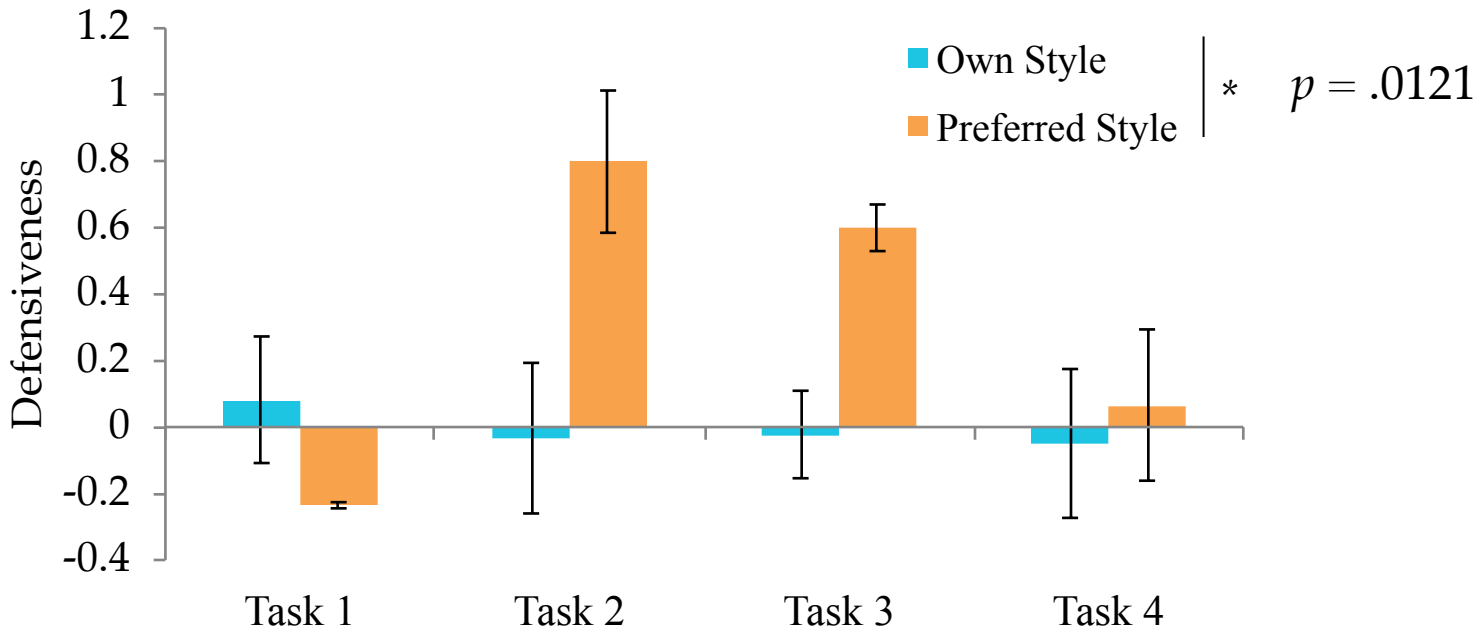

Do you want your autonomous car to drive like you? [HRI'17]

Perceived own style vs preferred driving style!

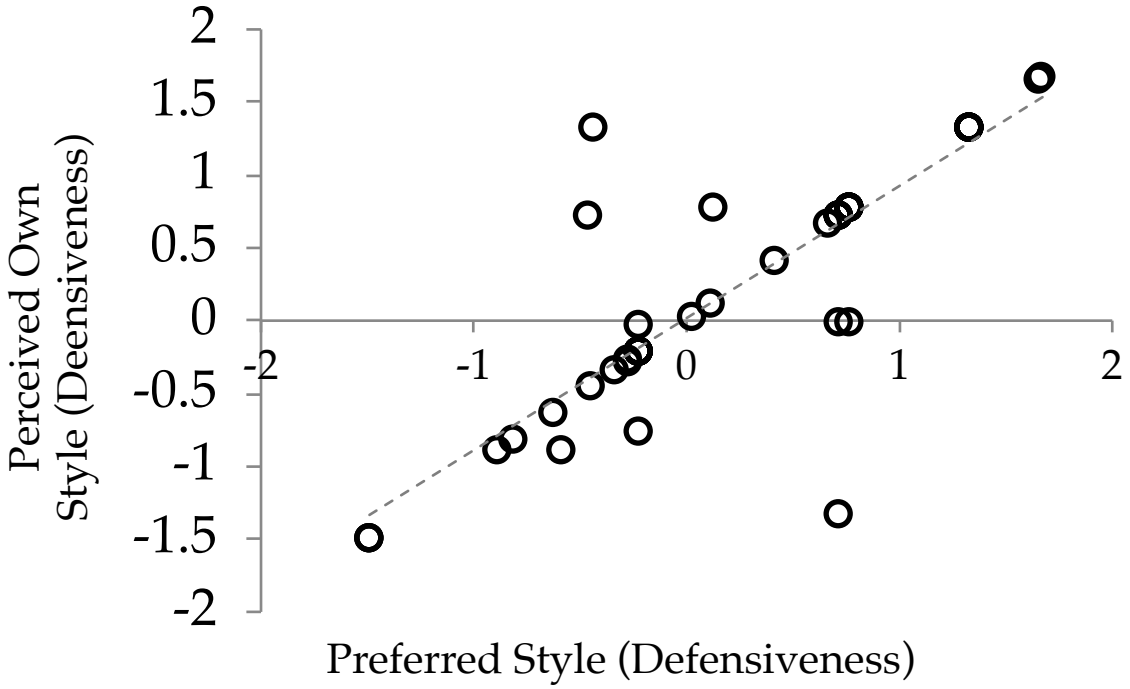

$r(58) = 0.86$

Do you want your autonomous car to drive like you? [HRI'17]

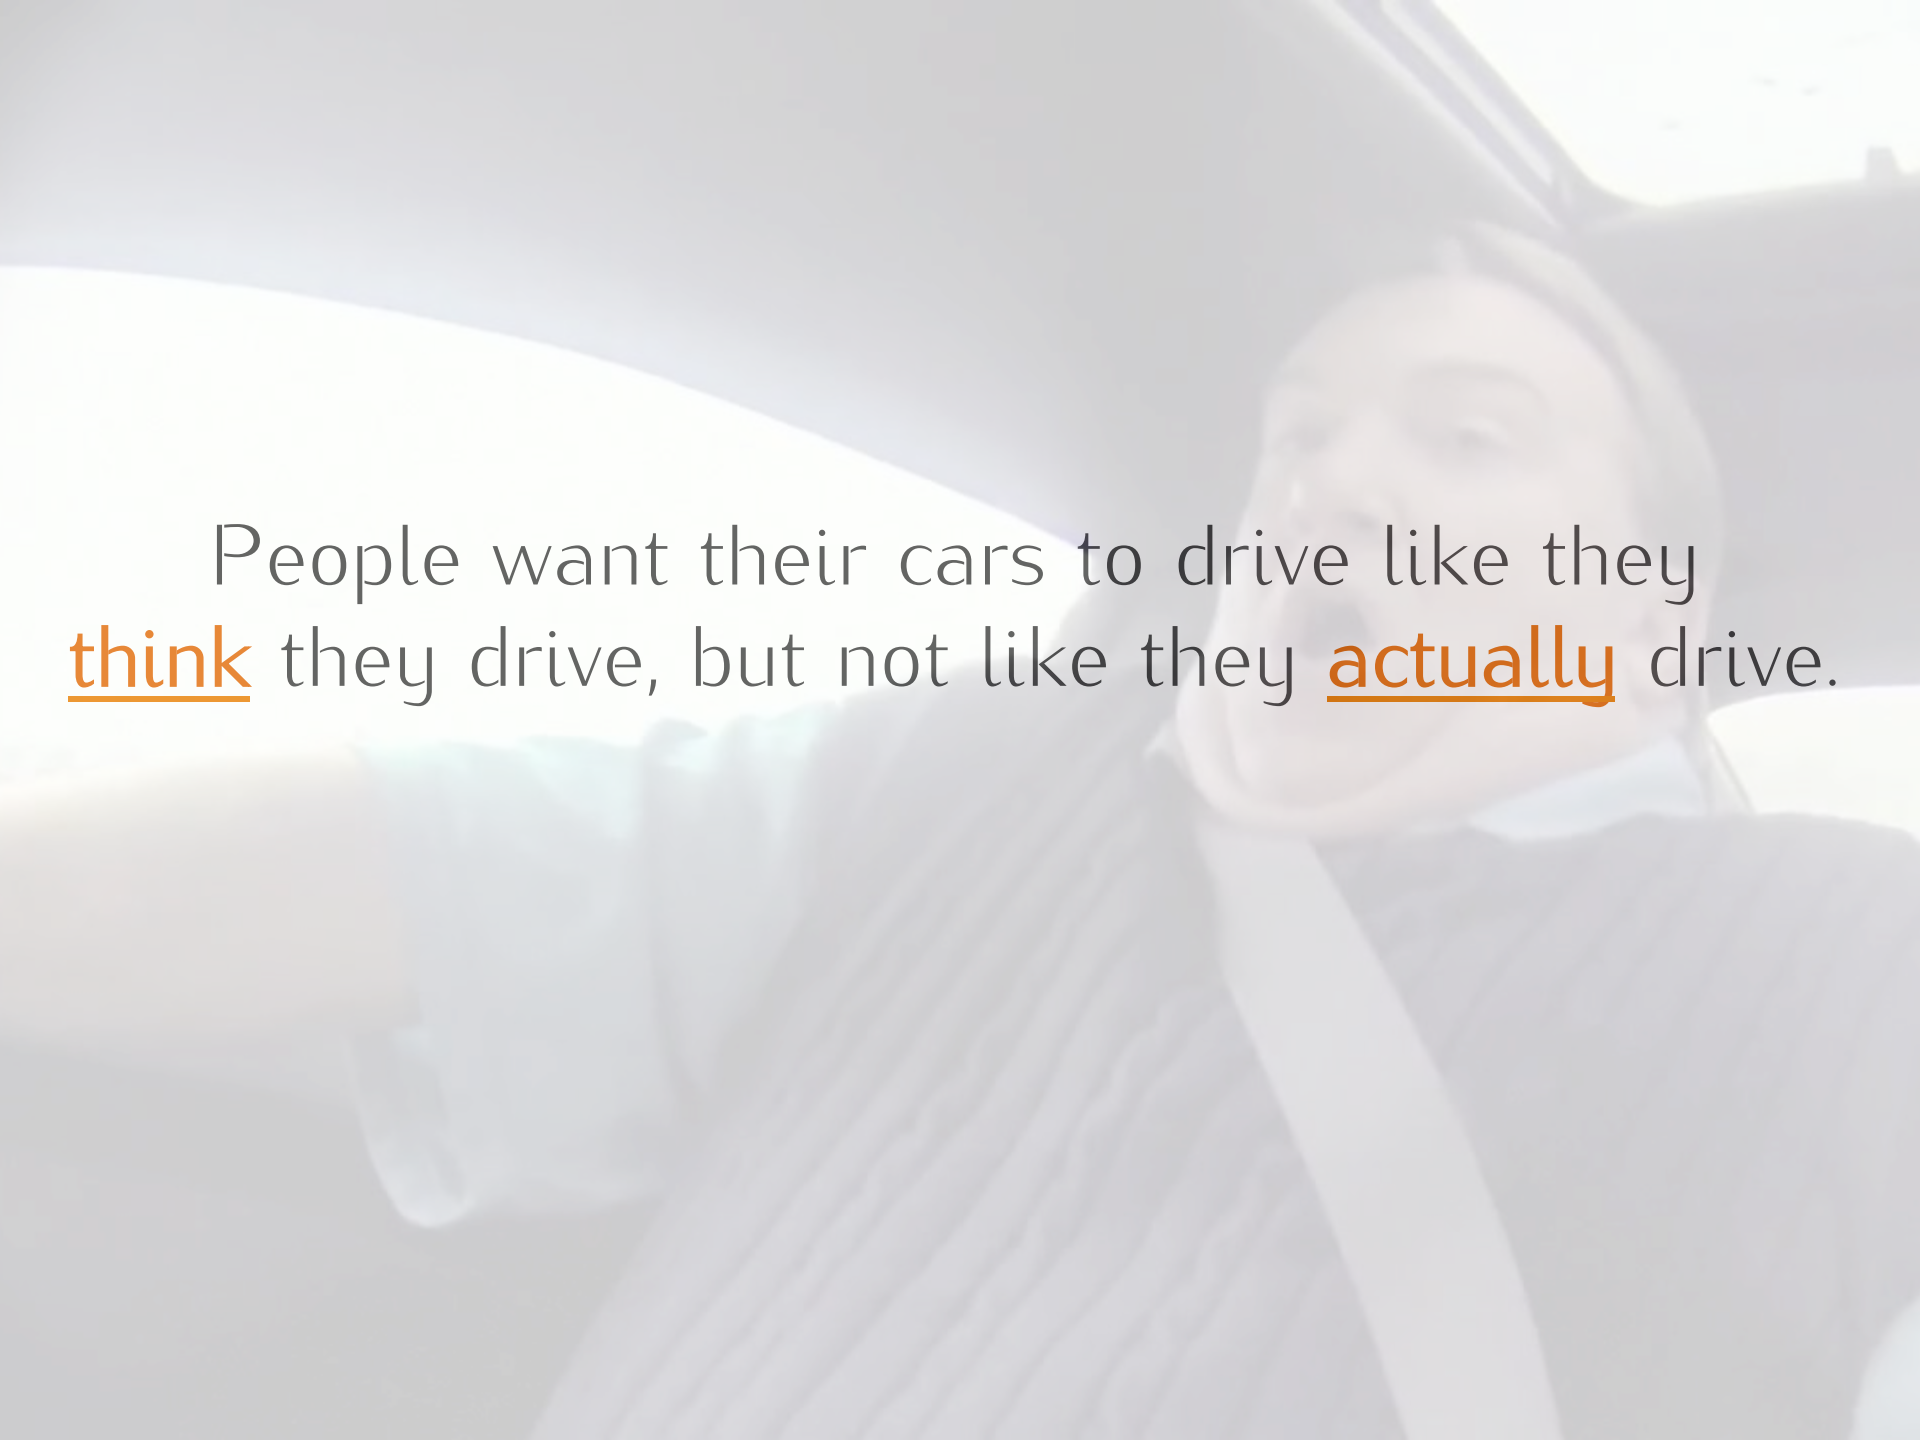
People want their cars to drive like they
think they drive, but not like they actually drive.

Implication: Can't rely on demonstrations!

Comparison-Based Learning

Is ξ_1 or ξ_2 better?

ξ_1

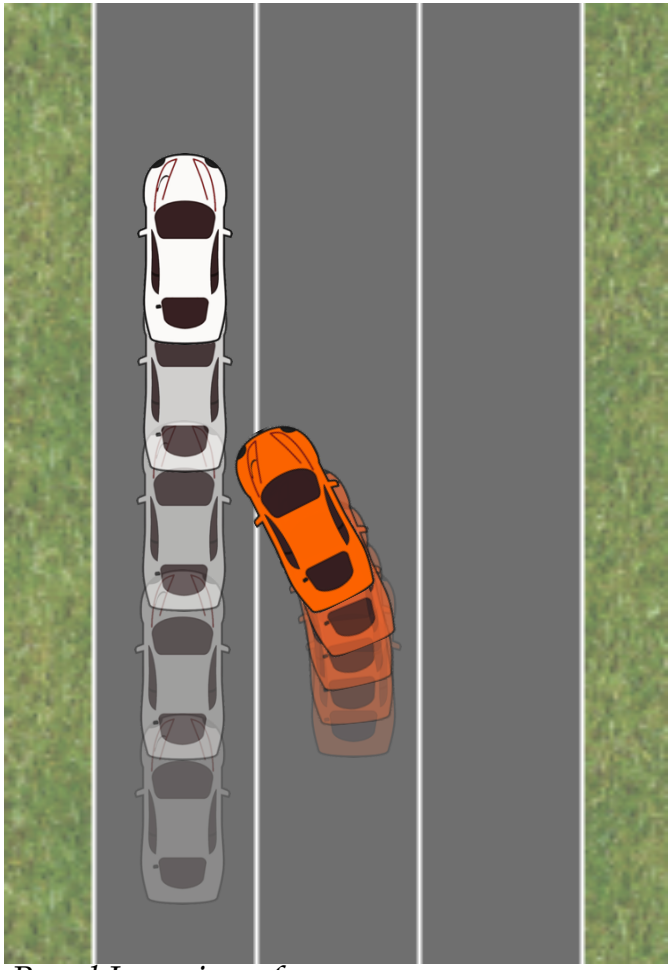

ξ_2

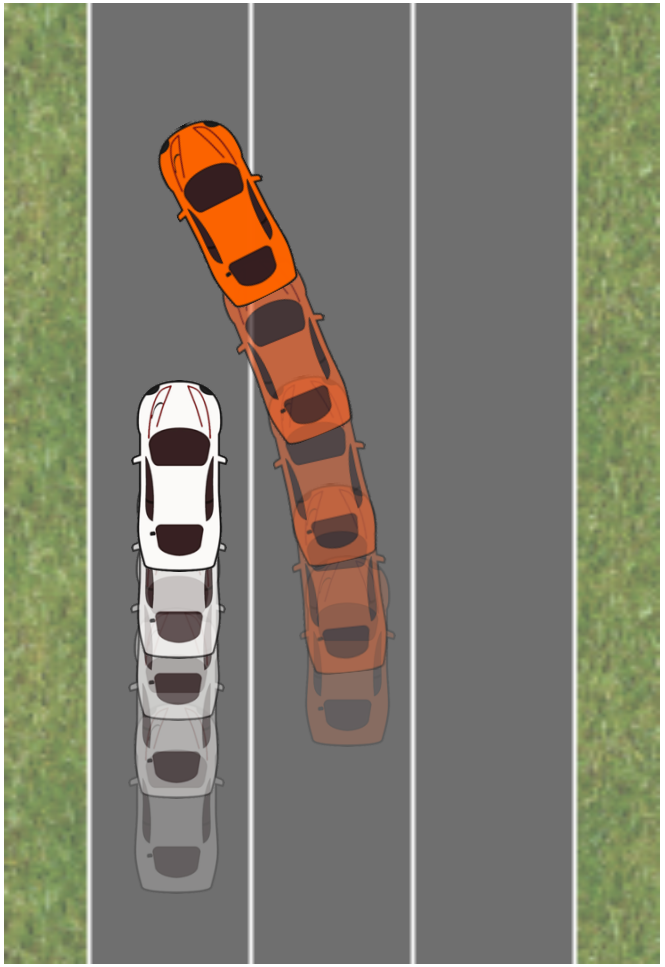

Active Preference-Based Learning of Reward Functions [RSS'17]

Every Answer Updates Belief over Hidden Params

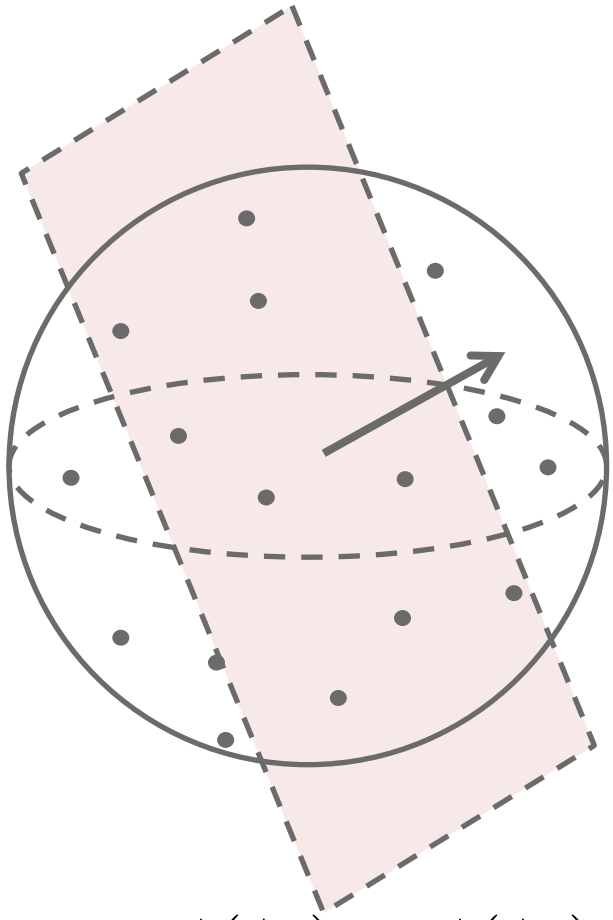

$$\varphi = \phi(\xi_1) - \phi(\xi_2)$$

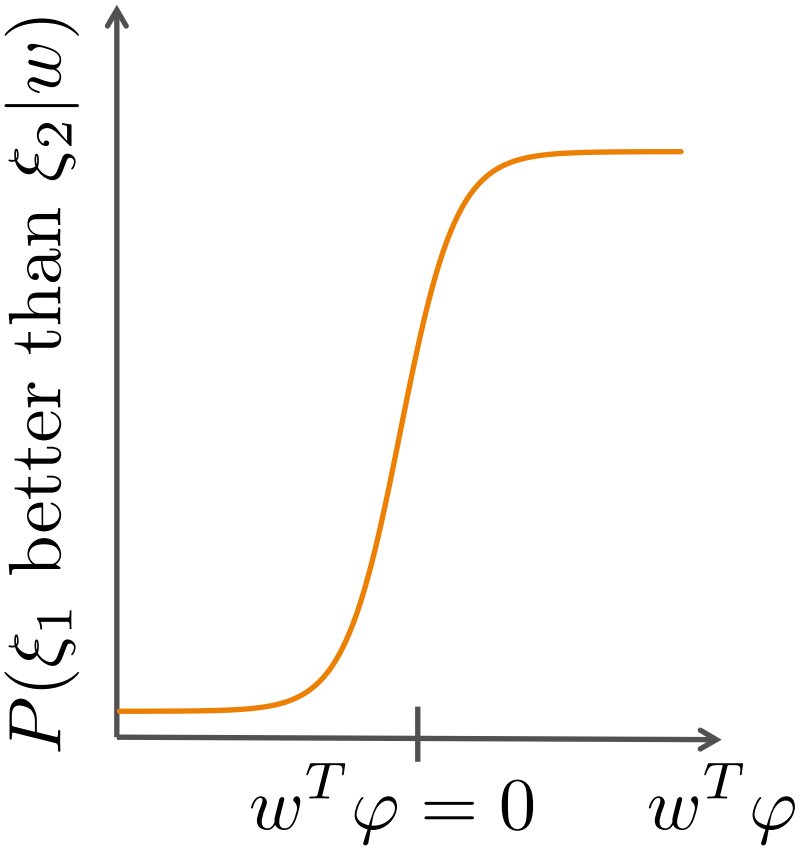

Actively Gather Information

$$\arg \max_{x_0, \mathbf{u}_R^1, \mathbf{u}_R^2} \mathbb{E}_\theta[H(b) - H(b')]$$

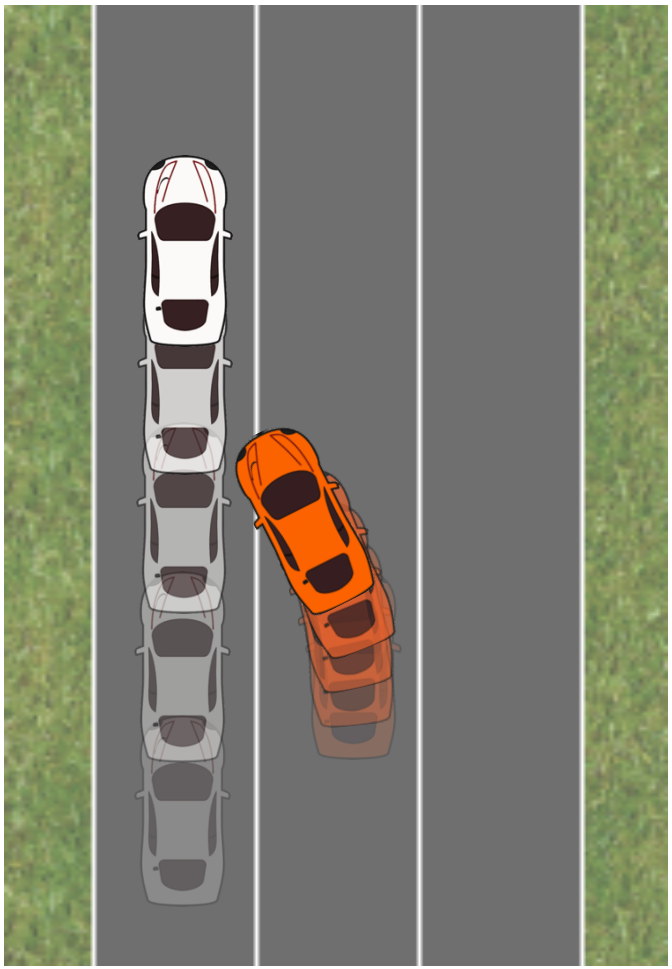

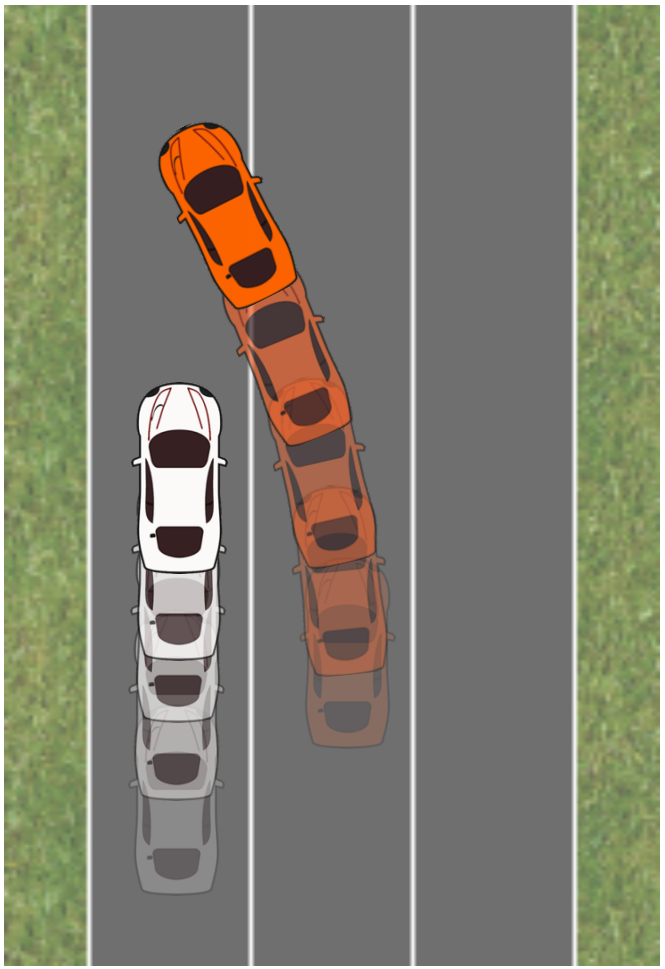

Belief before..

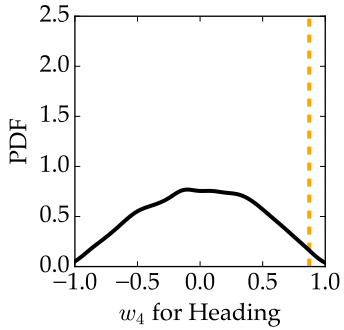

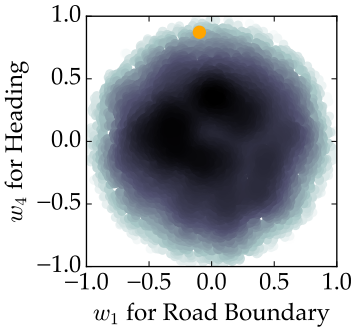

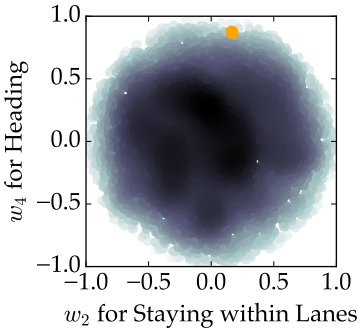

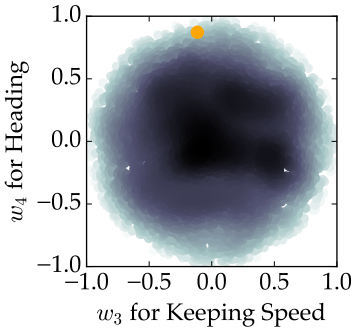

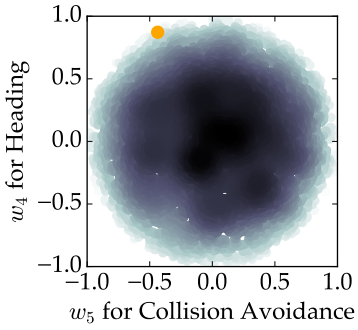

Belief after learning

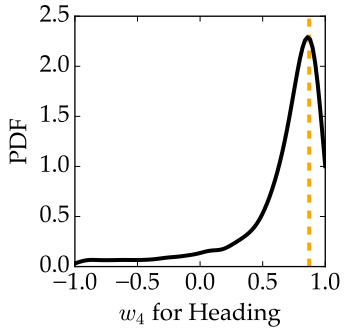

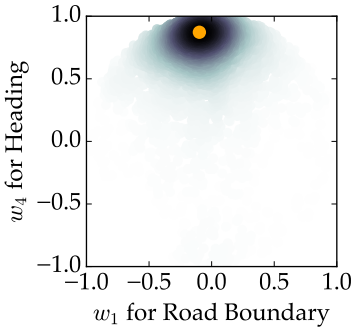

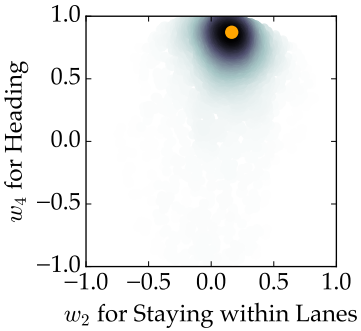

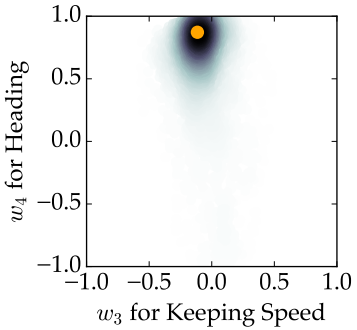

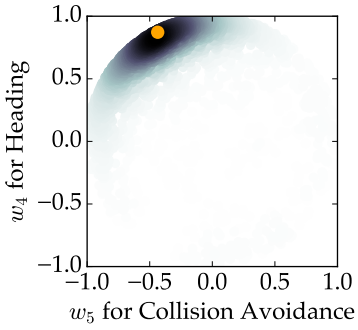

Actively Synthesizing Queries Helps

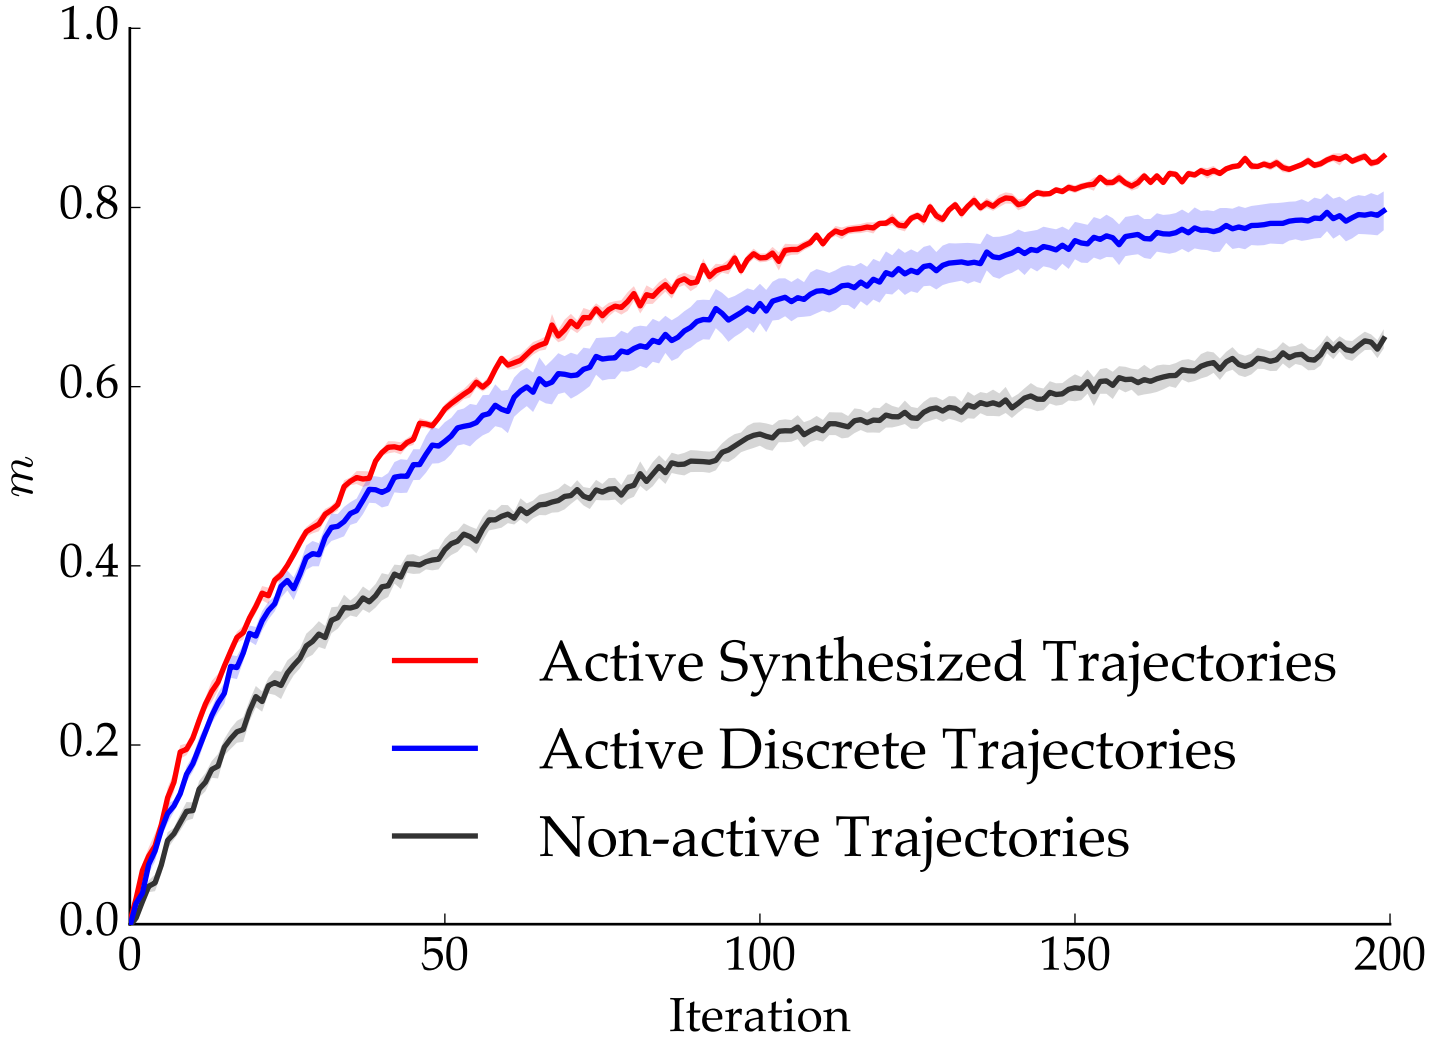

People are not just obstacles to be avoided.

... but they are not perfect game solvers either.

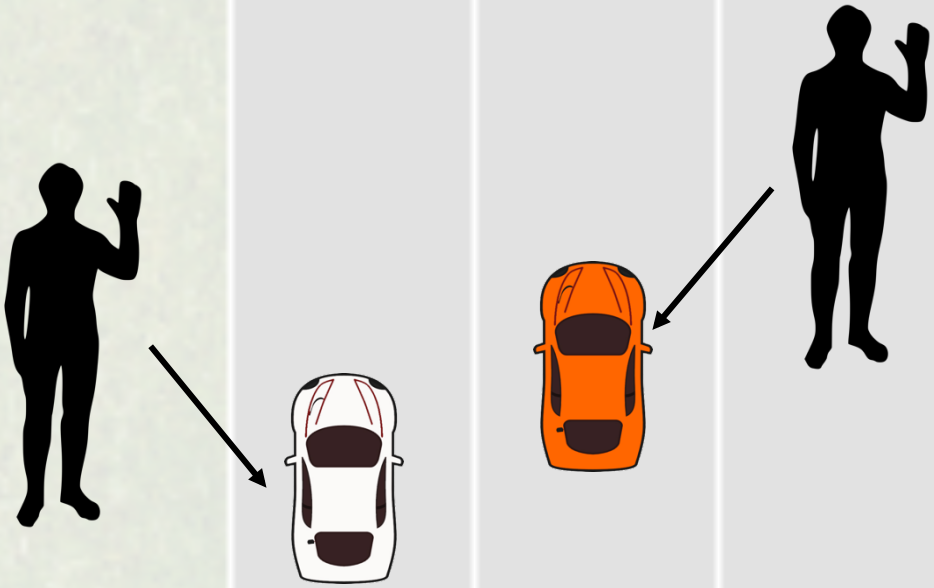


InterACT
Laboratory

Thanks!




NSF

CAREER




BERKELEY DeepDrive

Humans Should not Be Obstacles

Anca Dragan

anca@berkeley.edu