VeHICaL: Verified Human Interfaces, Control, and Learning for Semi-Autonomous Systems

S. A. Seshia, R. Bajcsy, T. Griffiths, B. Hartmann, S. S. Sastry, C. Tomlin (Berkeley), R. Murray (Caltech), C. Sturton (UNC Chapel Hill)

http://vehical.org

Annual Meeting
September 19, 2017

Human Cyber-Physical Systems (h-CPS)

CPS that operate in concert with humans

Driver Assistance in Cars

Fly-by-wire Cockpit Interfaces

UAVs with Human Operators

Robotic Surgery & Medicine

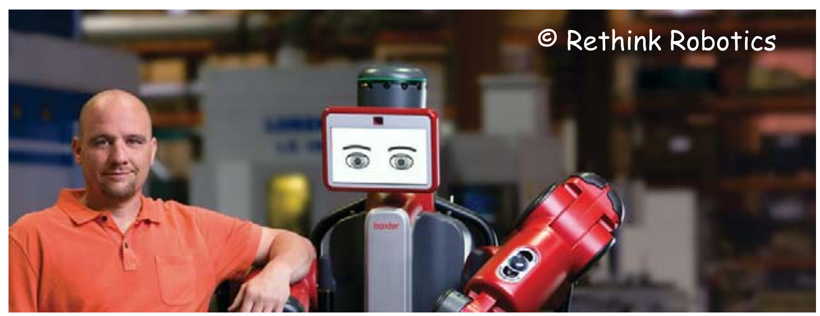

Semi-Autonomous Manufacturing

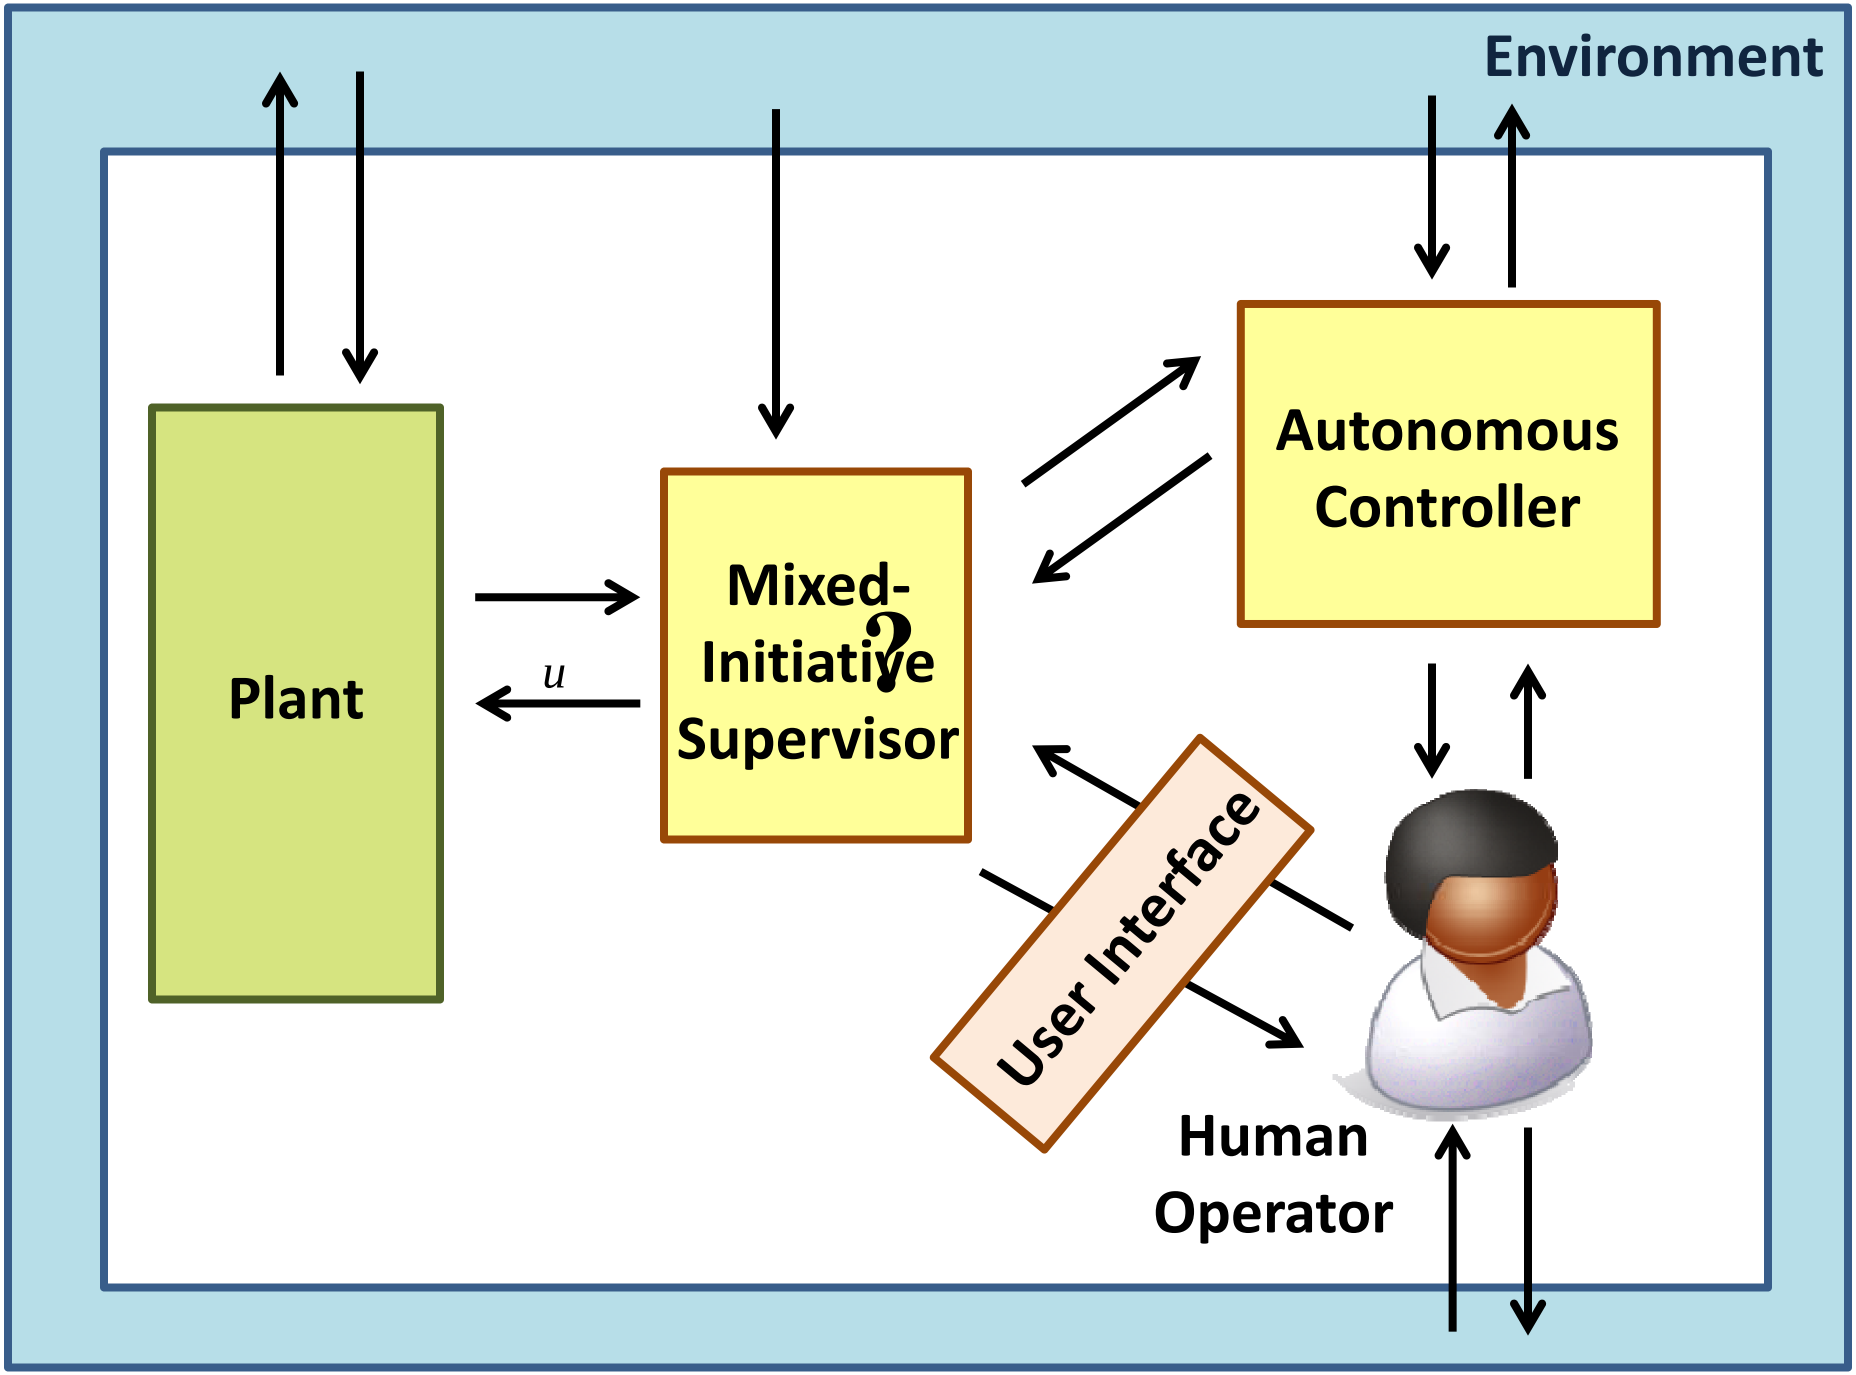
Environment
Plant
u
Mixed-Initiative Supervisor
?
Autonomous Controller
User Interface
Human Operator

Overall Project Objective of VeHICaL

To develop a ***science of verified co-design***
of controllers for semi-autonomous
cyber-physical systems and
interfaces between humans and
cyber-physical components

Motivating Applications

- **Semi-Autonomous Automobiles**

TECHNOLOGY — The New York Times

The 15-Point Federal Checklist for Self-Driving Cars

By CECILIA KANG SEPT. 20, 2016

A Lesson of Tesla Crashes? Computer Vision Can't Do It All Yet

- **Semi-Autonomous UAVs**

FAA Expects 600,000 Commercial Drones In The Air Within A Year

August 29, 2016 · 3:10 PM ET npr

AARIAN MARSHALL TRANSPORTATION 08.30.16 1:28 PM WIRED

HOW TO ACE THE FAA'S NEW TEST AND BECOME A PRO DRONE PILOT

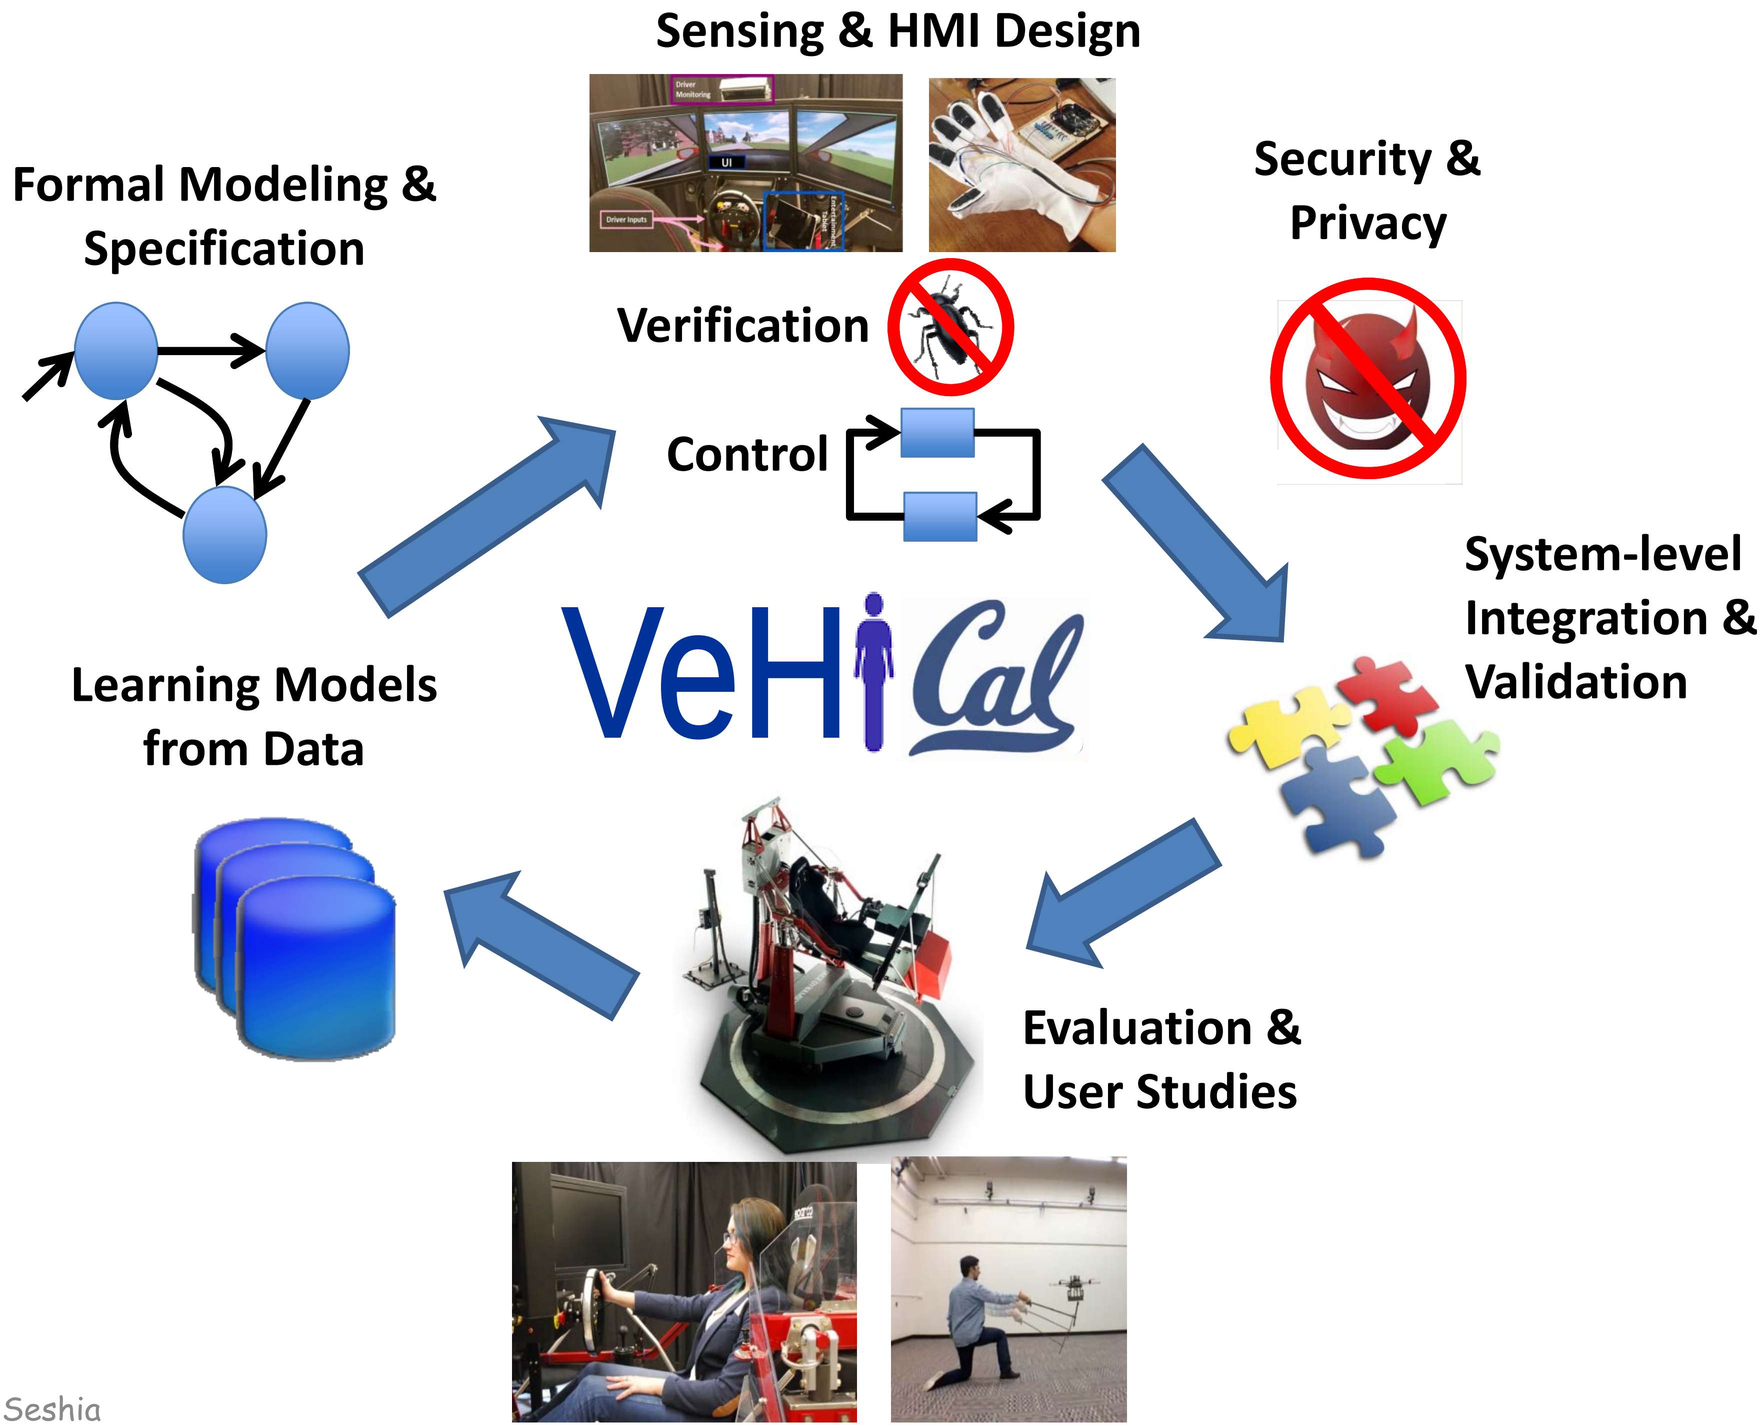
Sensing & HMI Design
Driver Monitoring
UI
Driver Inputs
Formal Modeling &
Specification
Security &
Privacy
Verification
Control
System-level
Integration &
Validation
Learning Models
from Data
VeHiCal
Evaluation &
User Studies

Four Project Thrusts

1. Specification and Modeling

2. Learning, Verification, and Control

3. Human-Machine Interface Design & Verification

4. Testbed and Evaluation

Quick Overview of Year 1 Results

- Characterizing Behavior of Bounded Agents [Griffiths]
- Interaction-Aware Control [Sastry, Seshia]
- Computing Approximately-Optimal Controls for Stochastic Systems [Tomlin]
- FaSTrack: Fast and Safe Tracking for High Dimensional Systems [Tomlin]
- aDOBO: Dynamics Optimization via Bayesian Optimization [Tomlin]
- Verification of Learning-based CPS [Seshia]
- User Interfaces that Convey Internal and External Awareness [Bajcsy, Sastry, Seshia]
- Optimizing the Information-Performance Tradeoff between Humans and Autonomy [Bajcsy]
- Privacy Preserving Drowsiness Detection [Sturton]

Recent Graduates and New Additions

- Katherine Driggs-Campbell
  - Ph.D. UC Berkeley 2017
  - Currently Postdoctoral Researcher, Stanford

- Dorsa Sadigh
  - Ph.D. UC Berkeley 2017
  - Currently Asst. Professor, Stanford

- Mark Ho
  - Joining as VeHICaL postdoctoral researcher in Jan. 2018
  - Currently finishing Ph.D. at Brown University

Education and Outreach

- New course on "Reimagining Mobility" at the Jacobs Insitute of Design Innovation (Bjoern Hartmann)
- New undergraduate course on "Robotic Manipulation and Interaction" (Ruzena Bajcsy)
  - first undergraduate robotics course that emphasizes design and human-robot interactions, and applications
- Girls in Engineering summer outreach program (Claire Tomlin)
- Two Industry Workshops (Nov. 9, 2016 and May 9, 2017)
- NHTSA Engagement (ongoing) on defining a testing methodology for autonomous vehicles

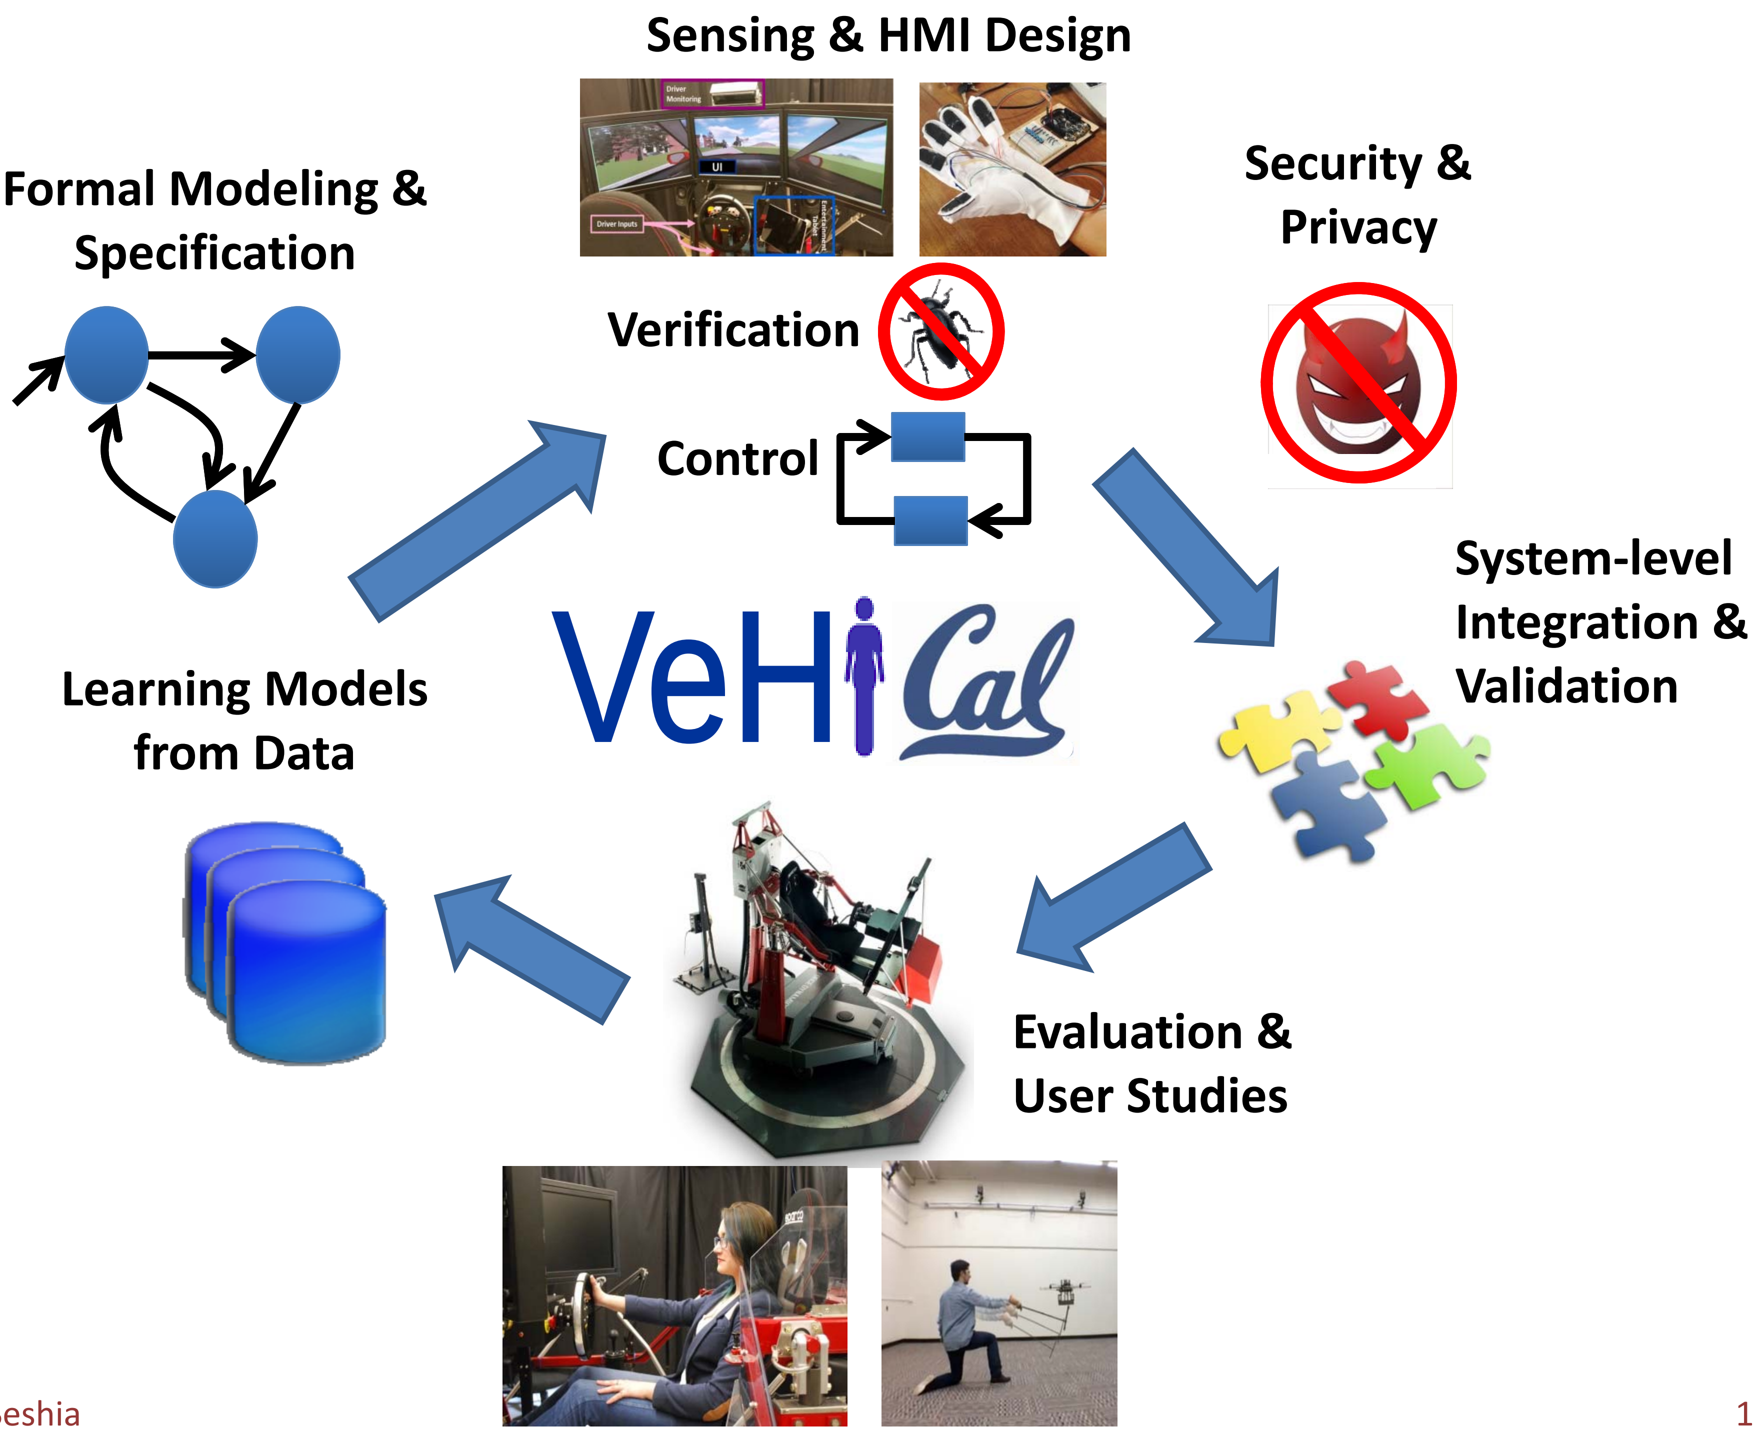
Sensing & HMI Design
Formal Modeling &
Specification
Security &
Privacy
Verification
Control
System-level
Integration &
Validation
Learning Models
from Data
VeHiCal
Evaluation &
User Studies

Interaction-Aware Control

- D. Sadigh, S. Sastry, S. Seshia, A. Dragan. Information Gathering Actions over Internal Human State. In IROS, 2016.
- D. Sadigh, S. Sastry, S. Seshia, A. Dragan. Planning for Autonomous Cars that Leverages Effects on Human Actions. In RSS, 2016.

Lane Change on a Highway

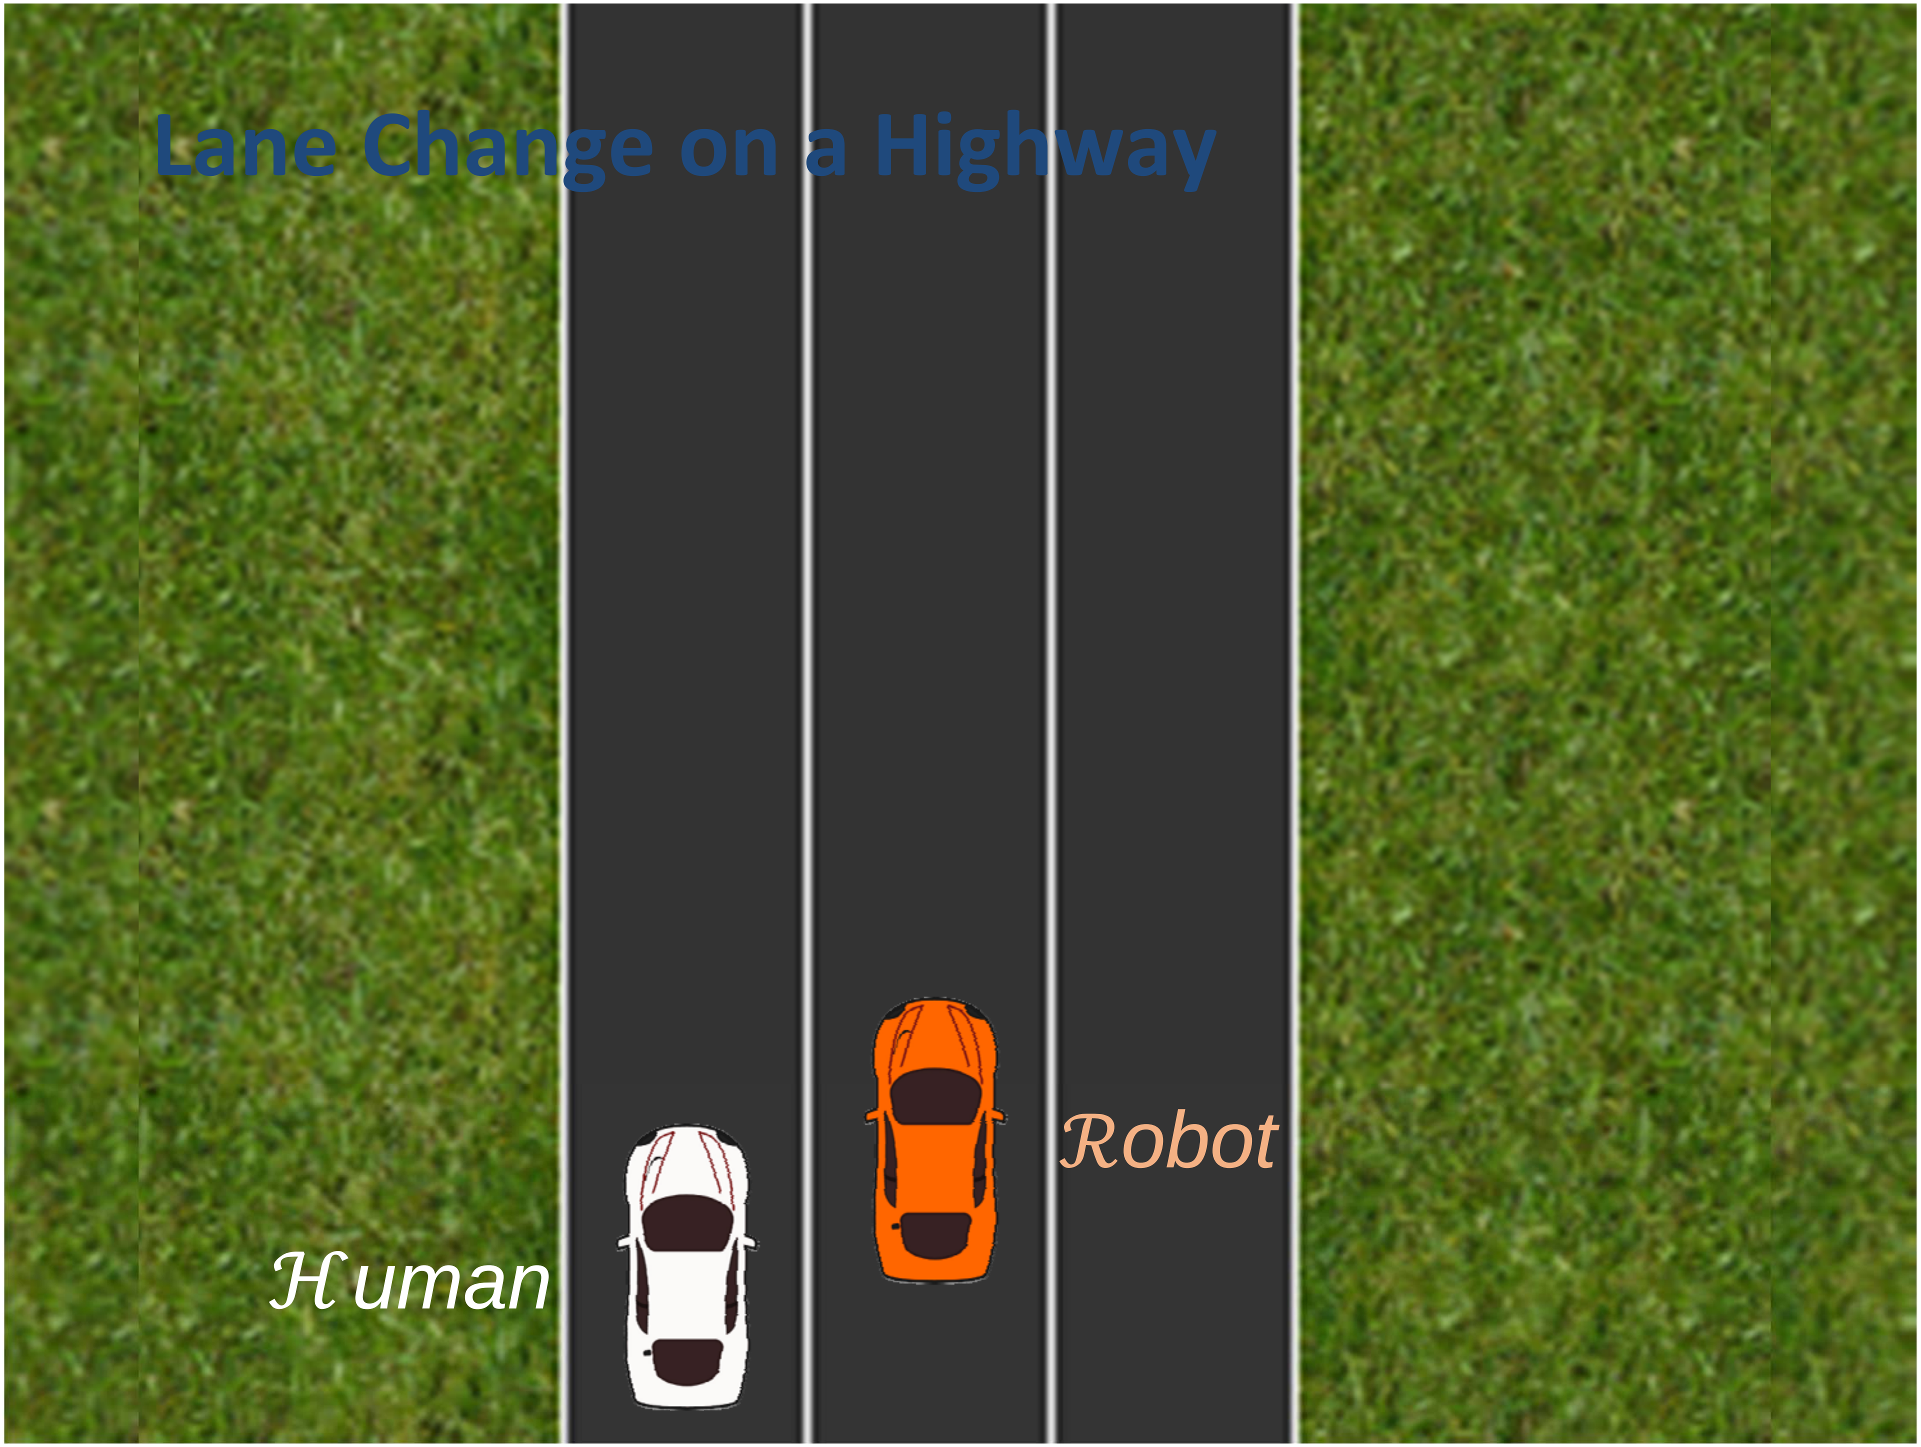

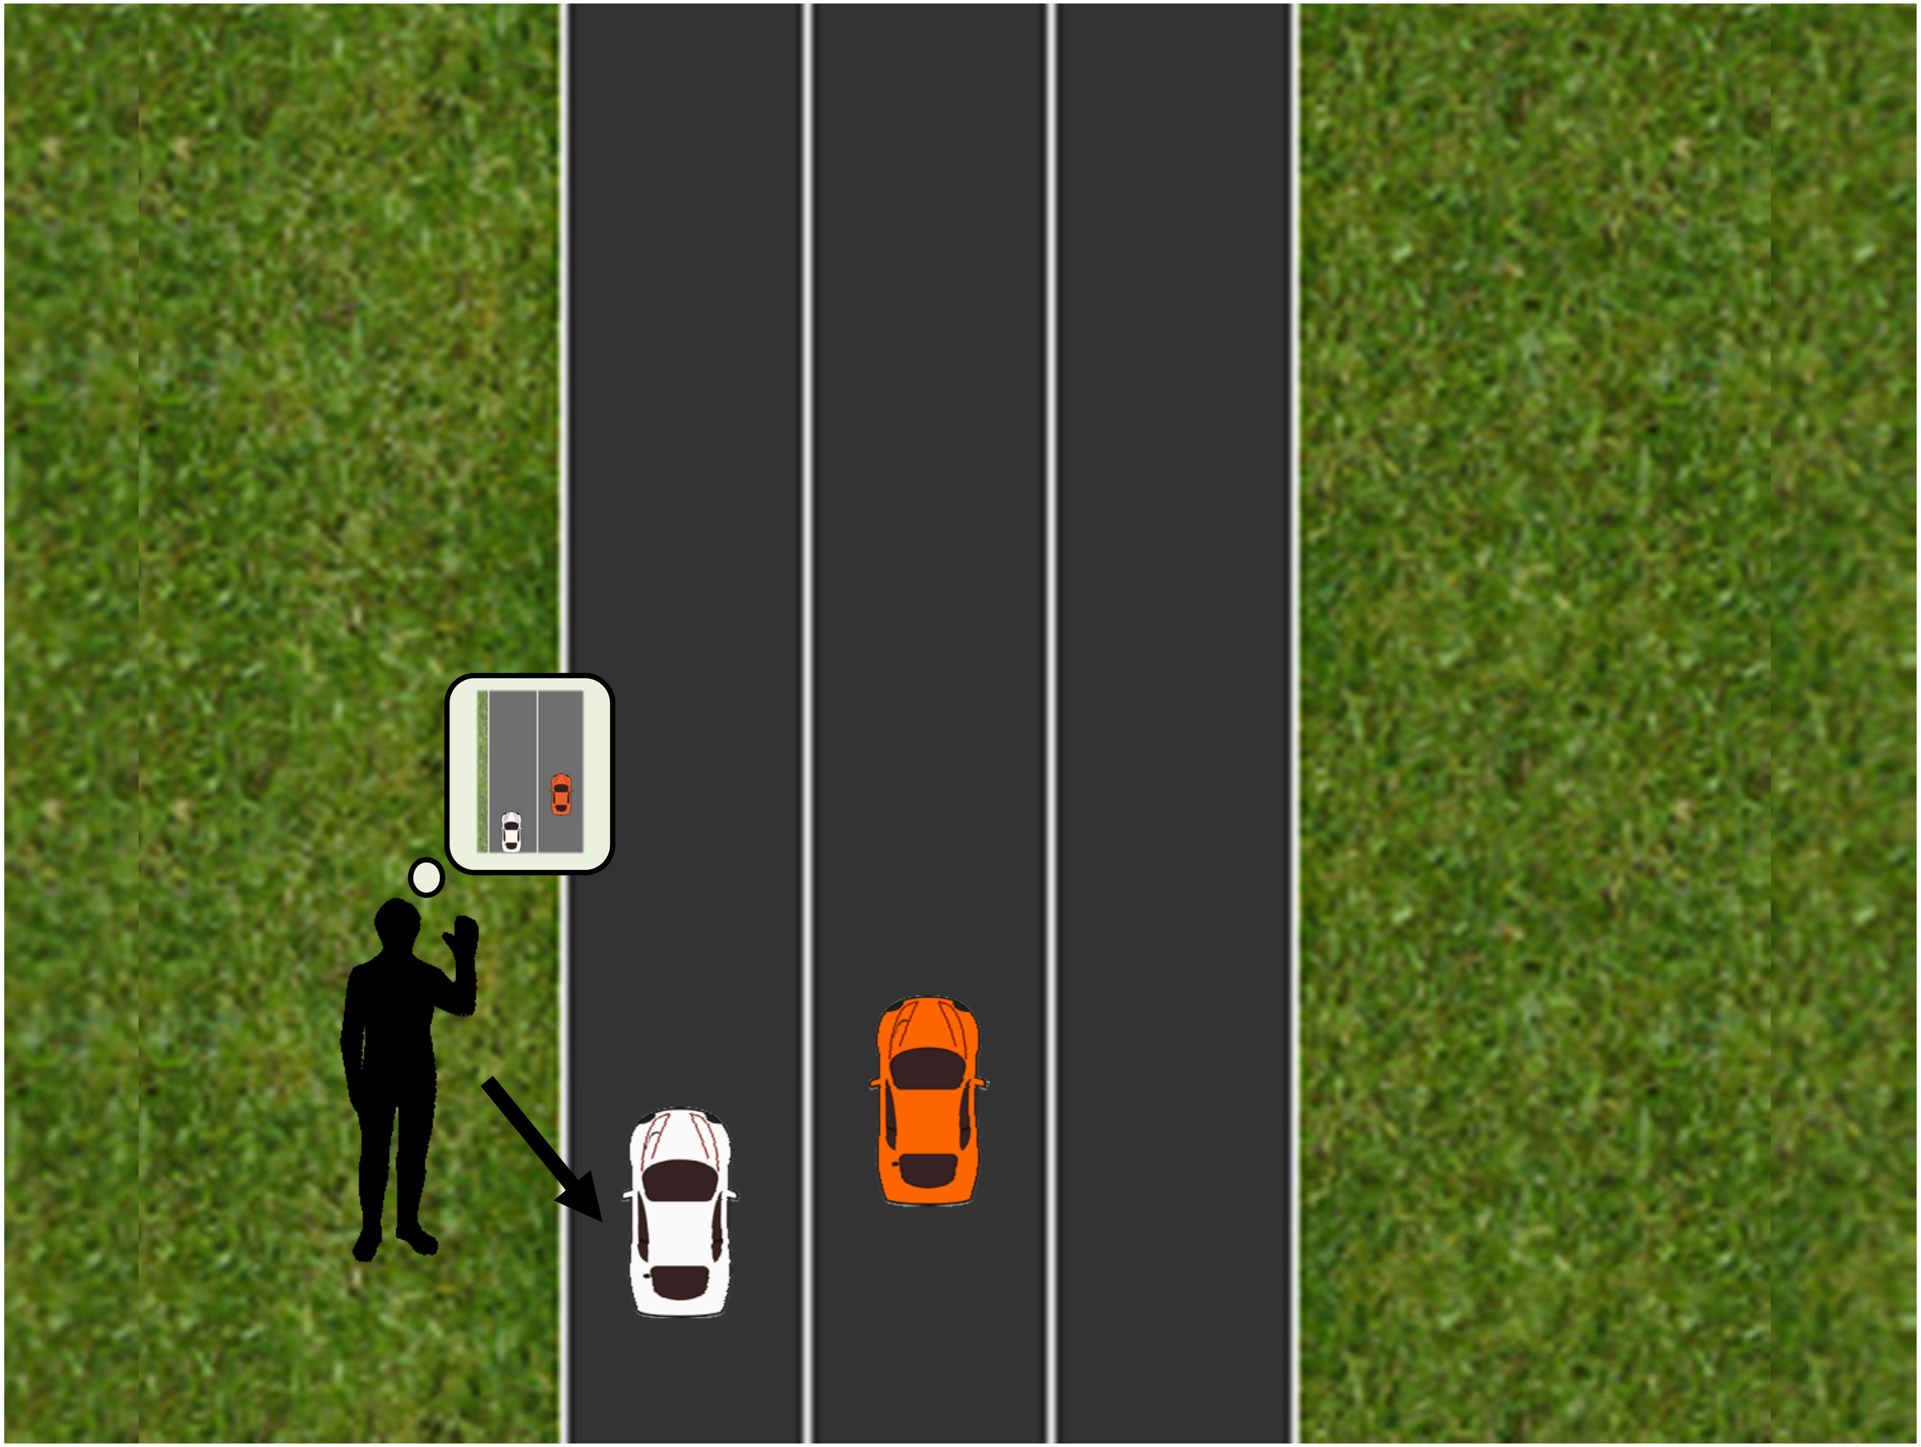

Interaction as a Dynamical System

$$x^{t+1} = f_{\mathcal{H}}(f_{\mathcal{R}}(x^t, u^t_{\mathcal{R}}), u^t_{\mathcal{H}})$$

Robot actions $\boldsymbol{u}_R$

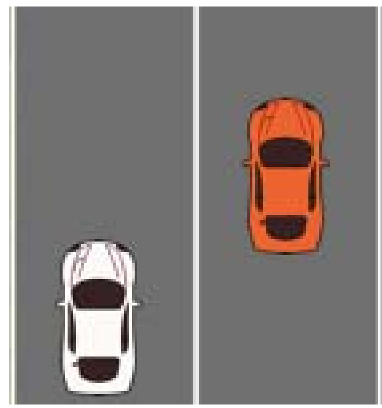

Human actions $\boldsymbol{u}_H$

Model the problem as a *Stackelberg (turn-based) Game*. Robot moves first.

Assumptions/Simplifications

Model Predictive (Receding Horizon) Control:

Optimize over short time horizon N, replan at every step t.

$$R_{\mathcal{R}}(x, u_{\mathcal{R}}, u_{\mathcal{H}}) = \sum_{t=1}^{N} r_{\mathcal{R}}(x^t, u_{\mathcal{R}}^t, u_{\mathcal{H}}^t) \qquad R_{\mathcal{H}}(x, u_{\mathcal{R}}, u_{\mathcal{H}}) = \sum_{t=1}^{N} r_{\mathcal{H}}(x^t, u_{\mathcal{R}}^t, u_{\mathcal{H}}^t)$$

Assume *deterministic "rational"* human model, human optimizes reward function which is a linear combination of "features".

Human has full access to $\boldsymbol{u}_{\mathcal{R}}$ for the short time horizon.

$$\boldsymbol{u}_H^*(x_0, \boldsymbol{u}_R) = \underset{\boldsymbol{u}_H}{\operatorname{argmax}}\, R_H(x_0, \boldsymbol{u}_R, \boldsymbol{u}_H)$$

Learning (Human) Driver Models

Learn Human's reward function based on Inverse Reinforcement Learning [Ziebart et al, AAAI'08; Levine & Koltun, 2012].

Assume structure of human reward function:

$$r_H(x^t, u_R^t, u_H^t) = w^\top \phi(x^t, u_R^t, u_H^t)$$

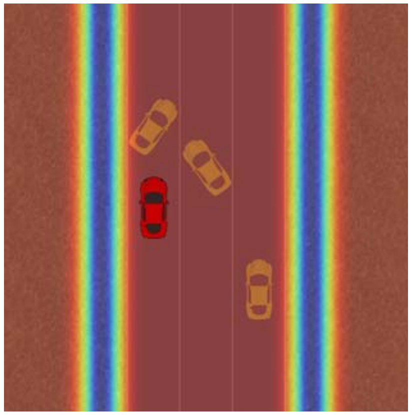

(a) Features for the boundaries of the road

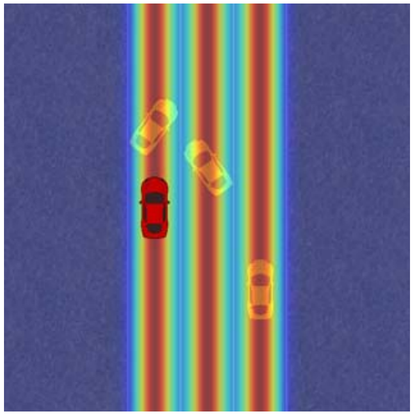

(b) Feature for staying inside the lanes.

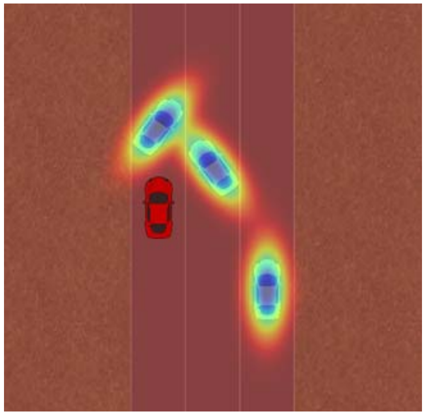

(c) Features for avoiding other vehicles.

B. Ziebart, A. Maas, J. A. Bagnell, and A. K. Dey. Maximum entropy inverse reinforcement learning. In AAAI, 2008.
S. Levine, V. Koltun. Continuous inverse optimal control with locally optimal examples. arXiv , 2012.

Interaction as a Dynamical System

$$\boldsymbol{u}_R^* = \underset{\boldsymbol{u}_R}{\operatorname{argmax}}\, R_R(x_0, \boldsymbol{u}_R, \boldsymbol{u}_H^*(x_0, \boldsymbol{u}_R))$$

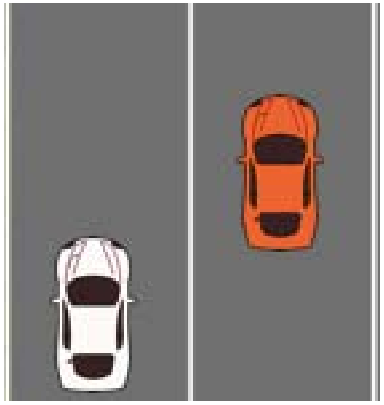

Find optimal actions for the autonomous vehicle while accounting for the human response $\boldsymbol{u}_H^*$.

Model $\boldsymbol{u}_H^*$ as optimizing the human reward function R_H.

$$\boldsymbol{u}_H^*(x_0, \boldsymbol{u}_R) = \underset{\boldsymbol{u}_H}{\operatorname{argmax}}\, R_H(x_0, \boldsymbol{u}_R, \boldsymbol{u}_H)$$

Implication 1: Efficiency

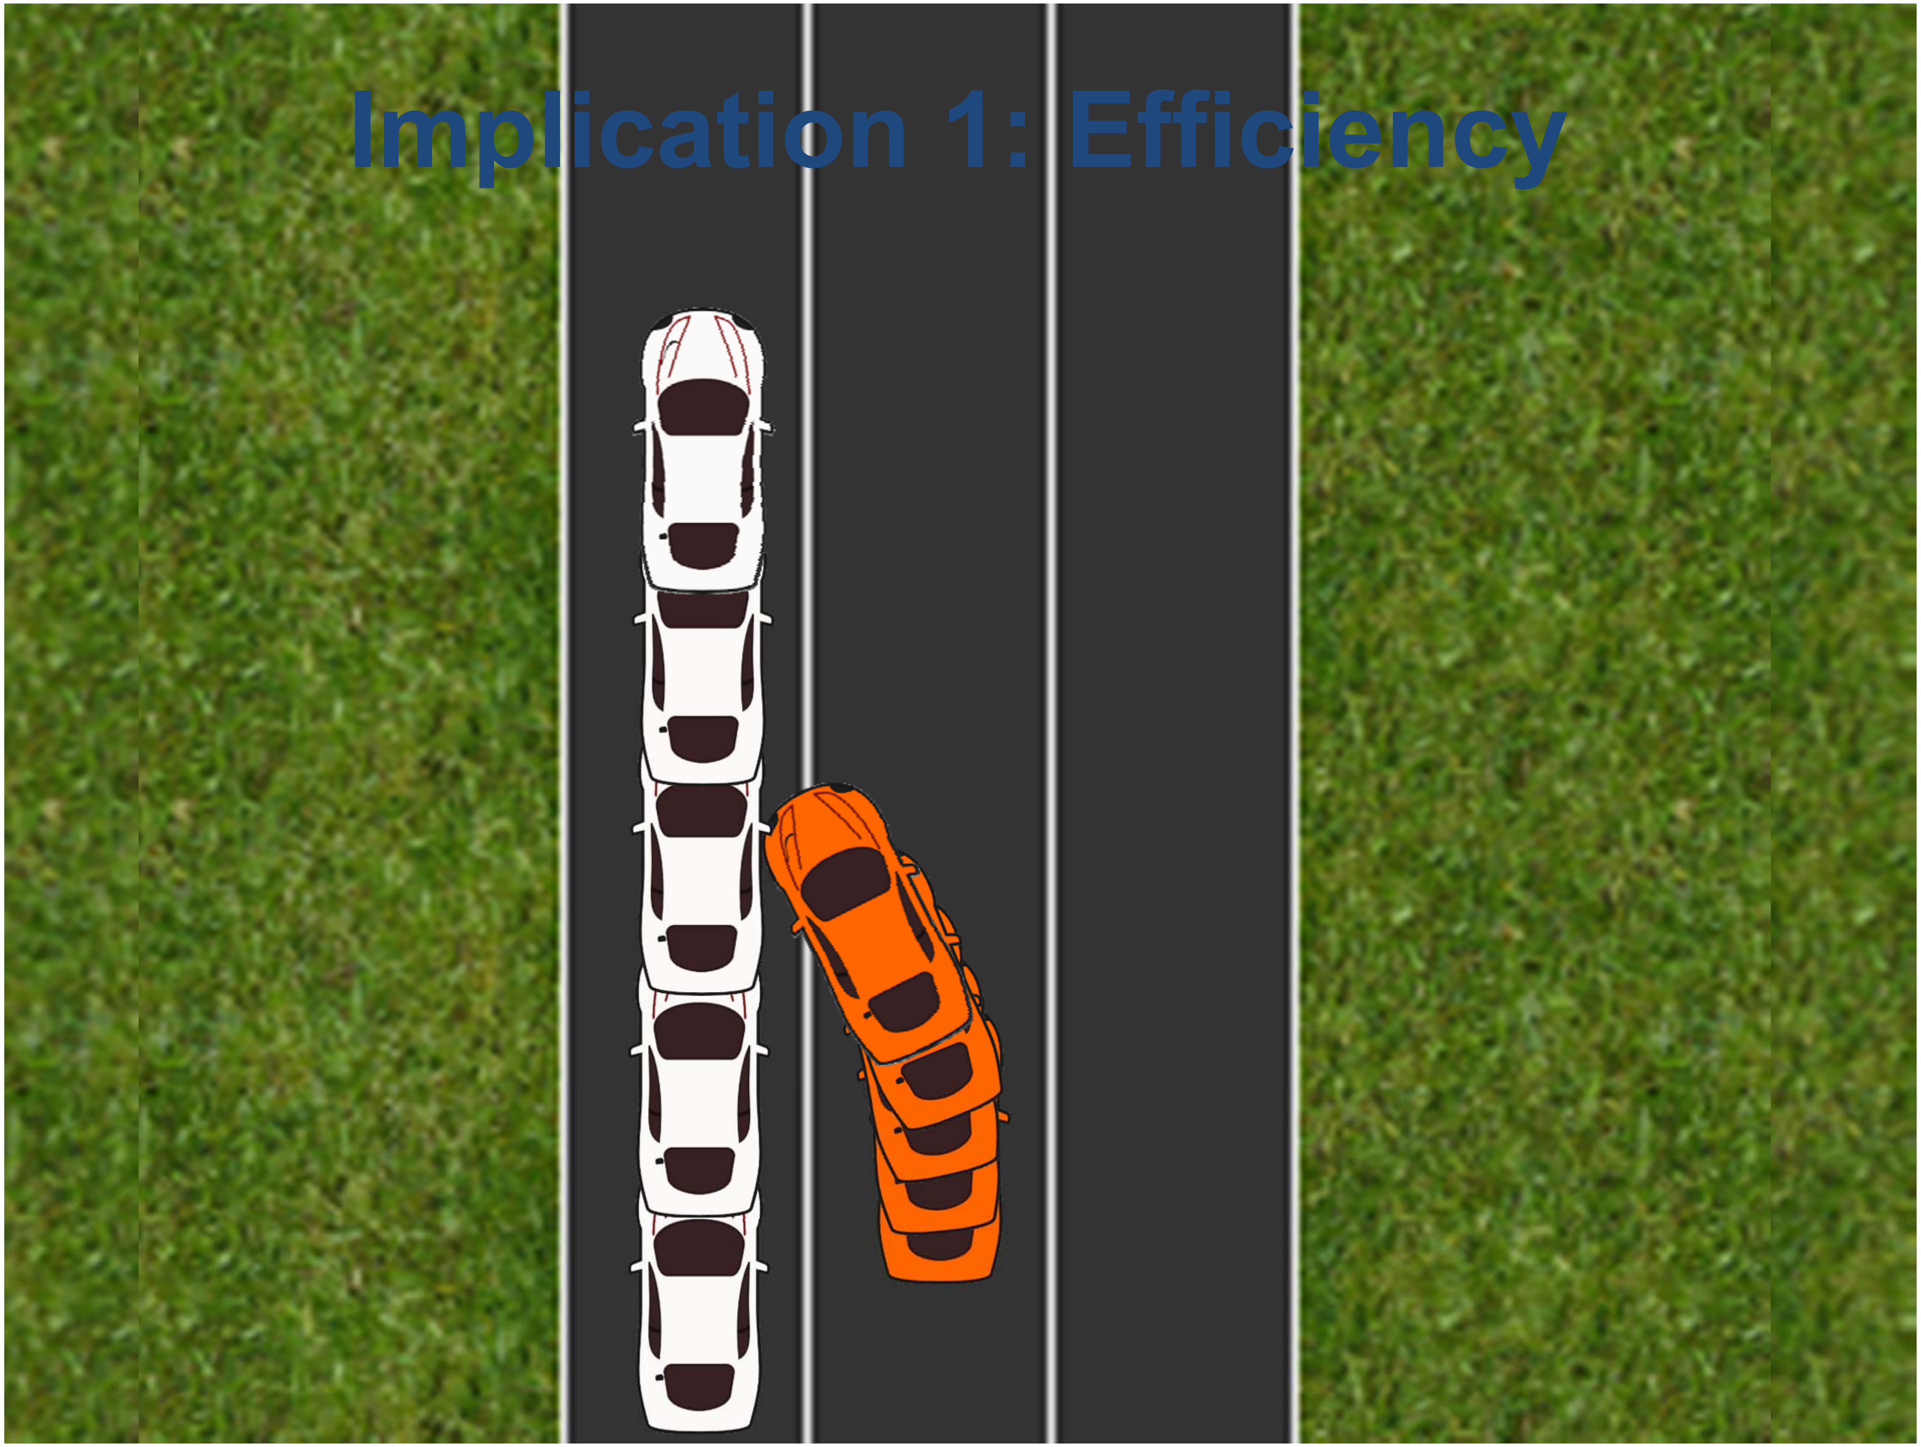

Implication 1: Efficiency

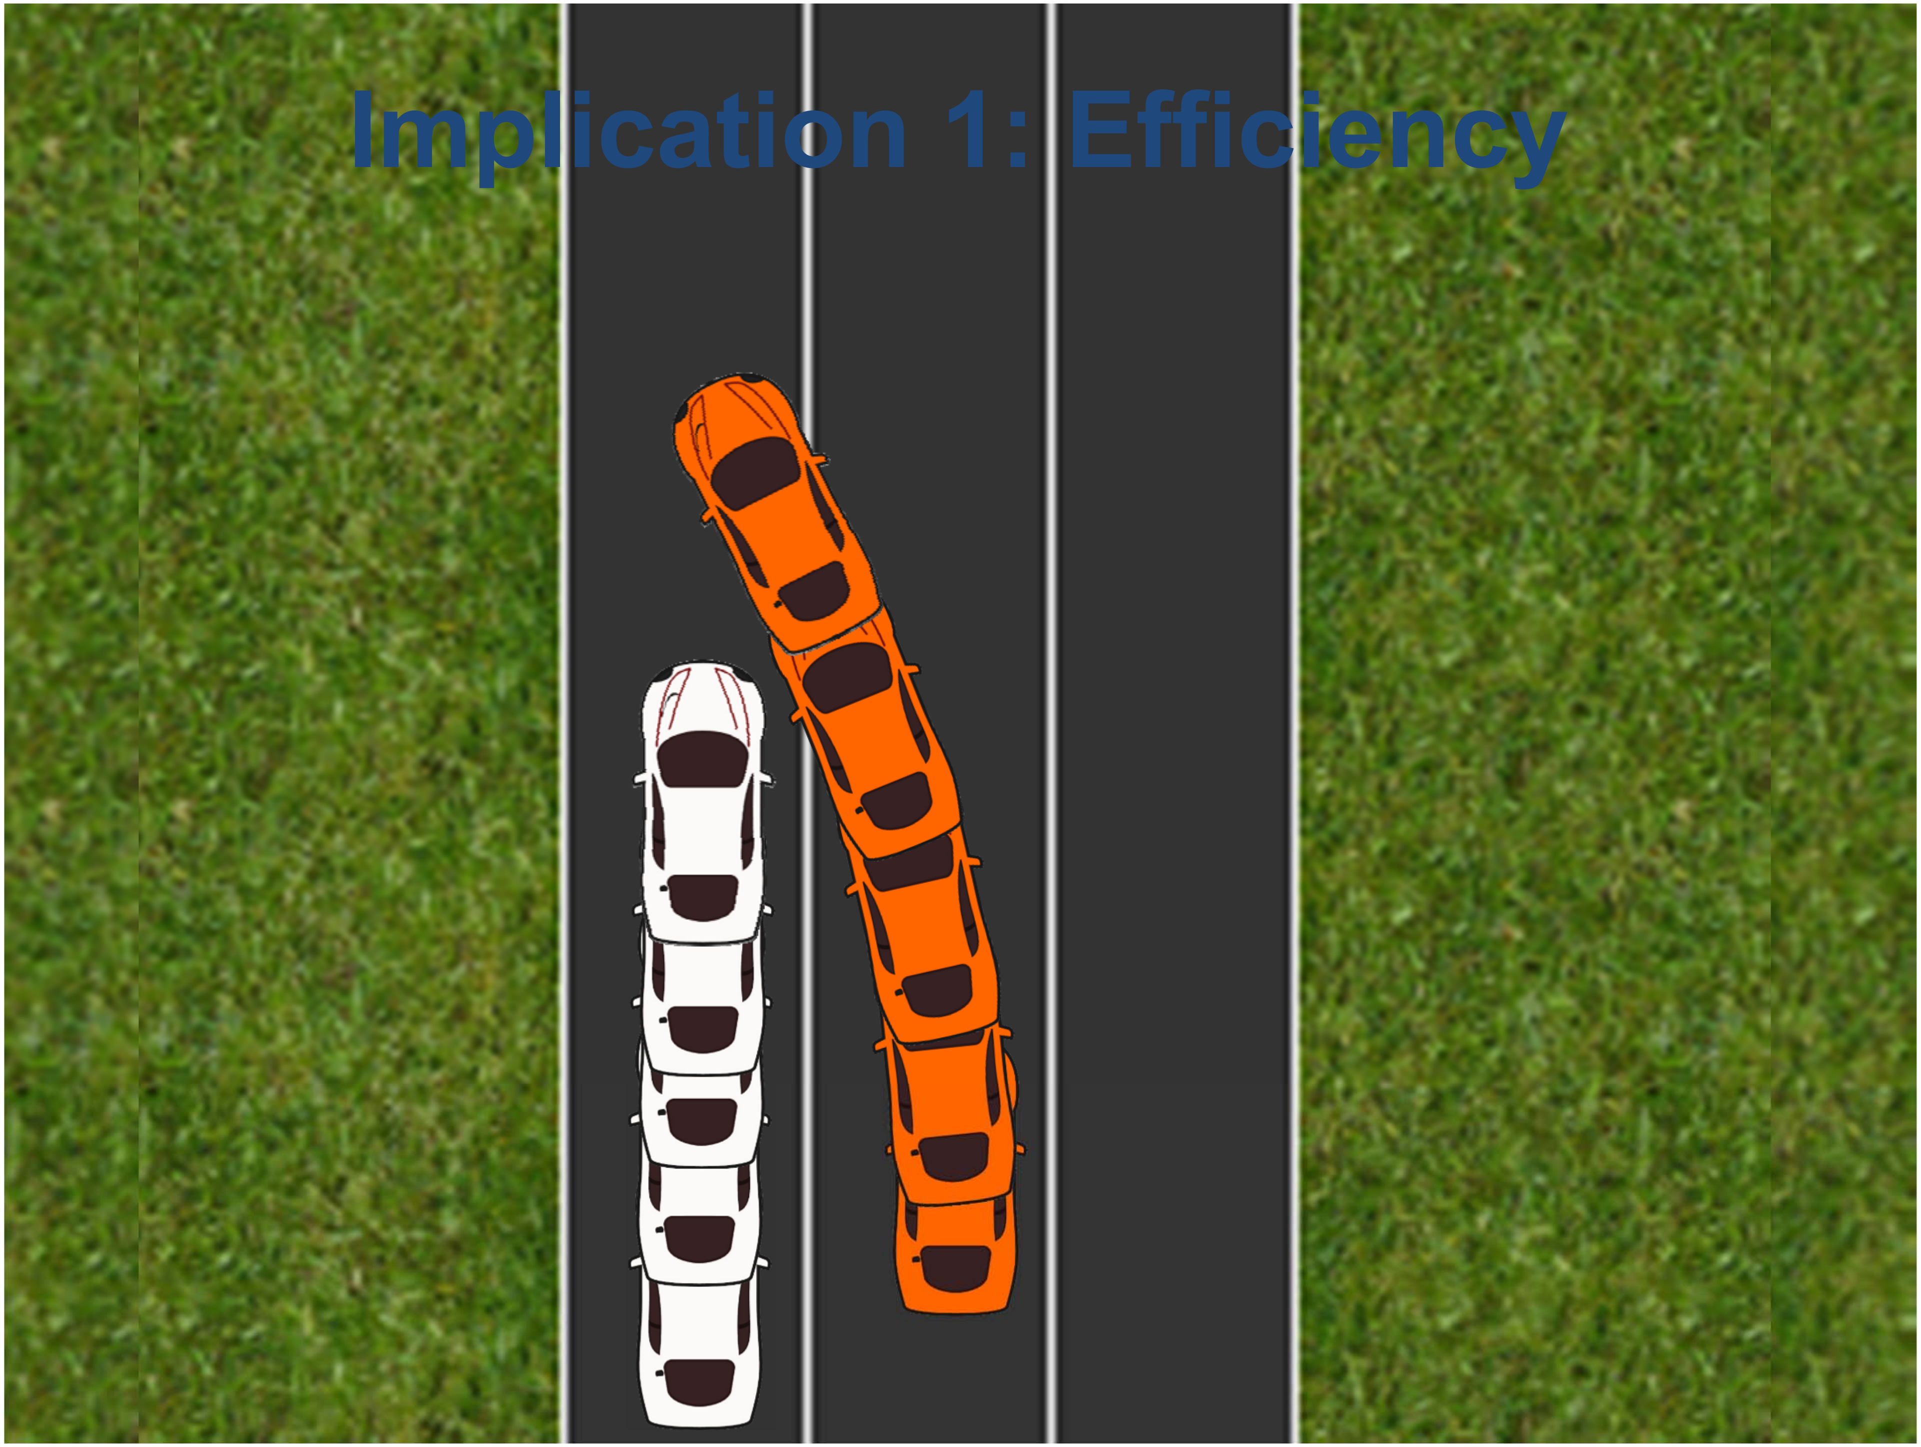

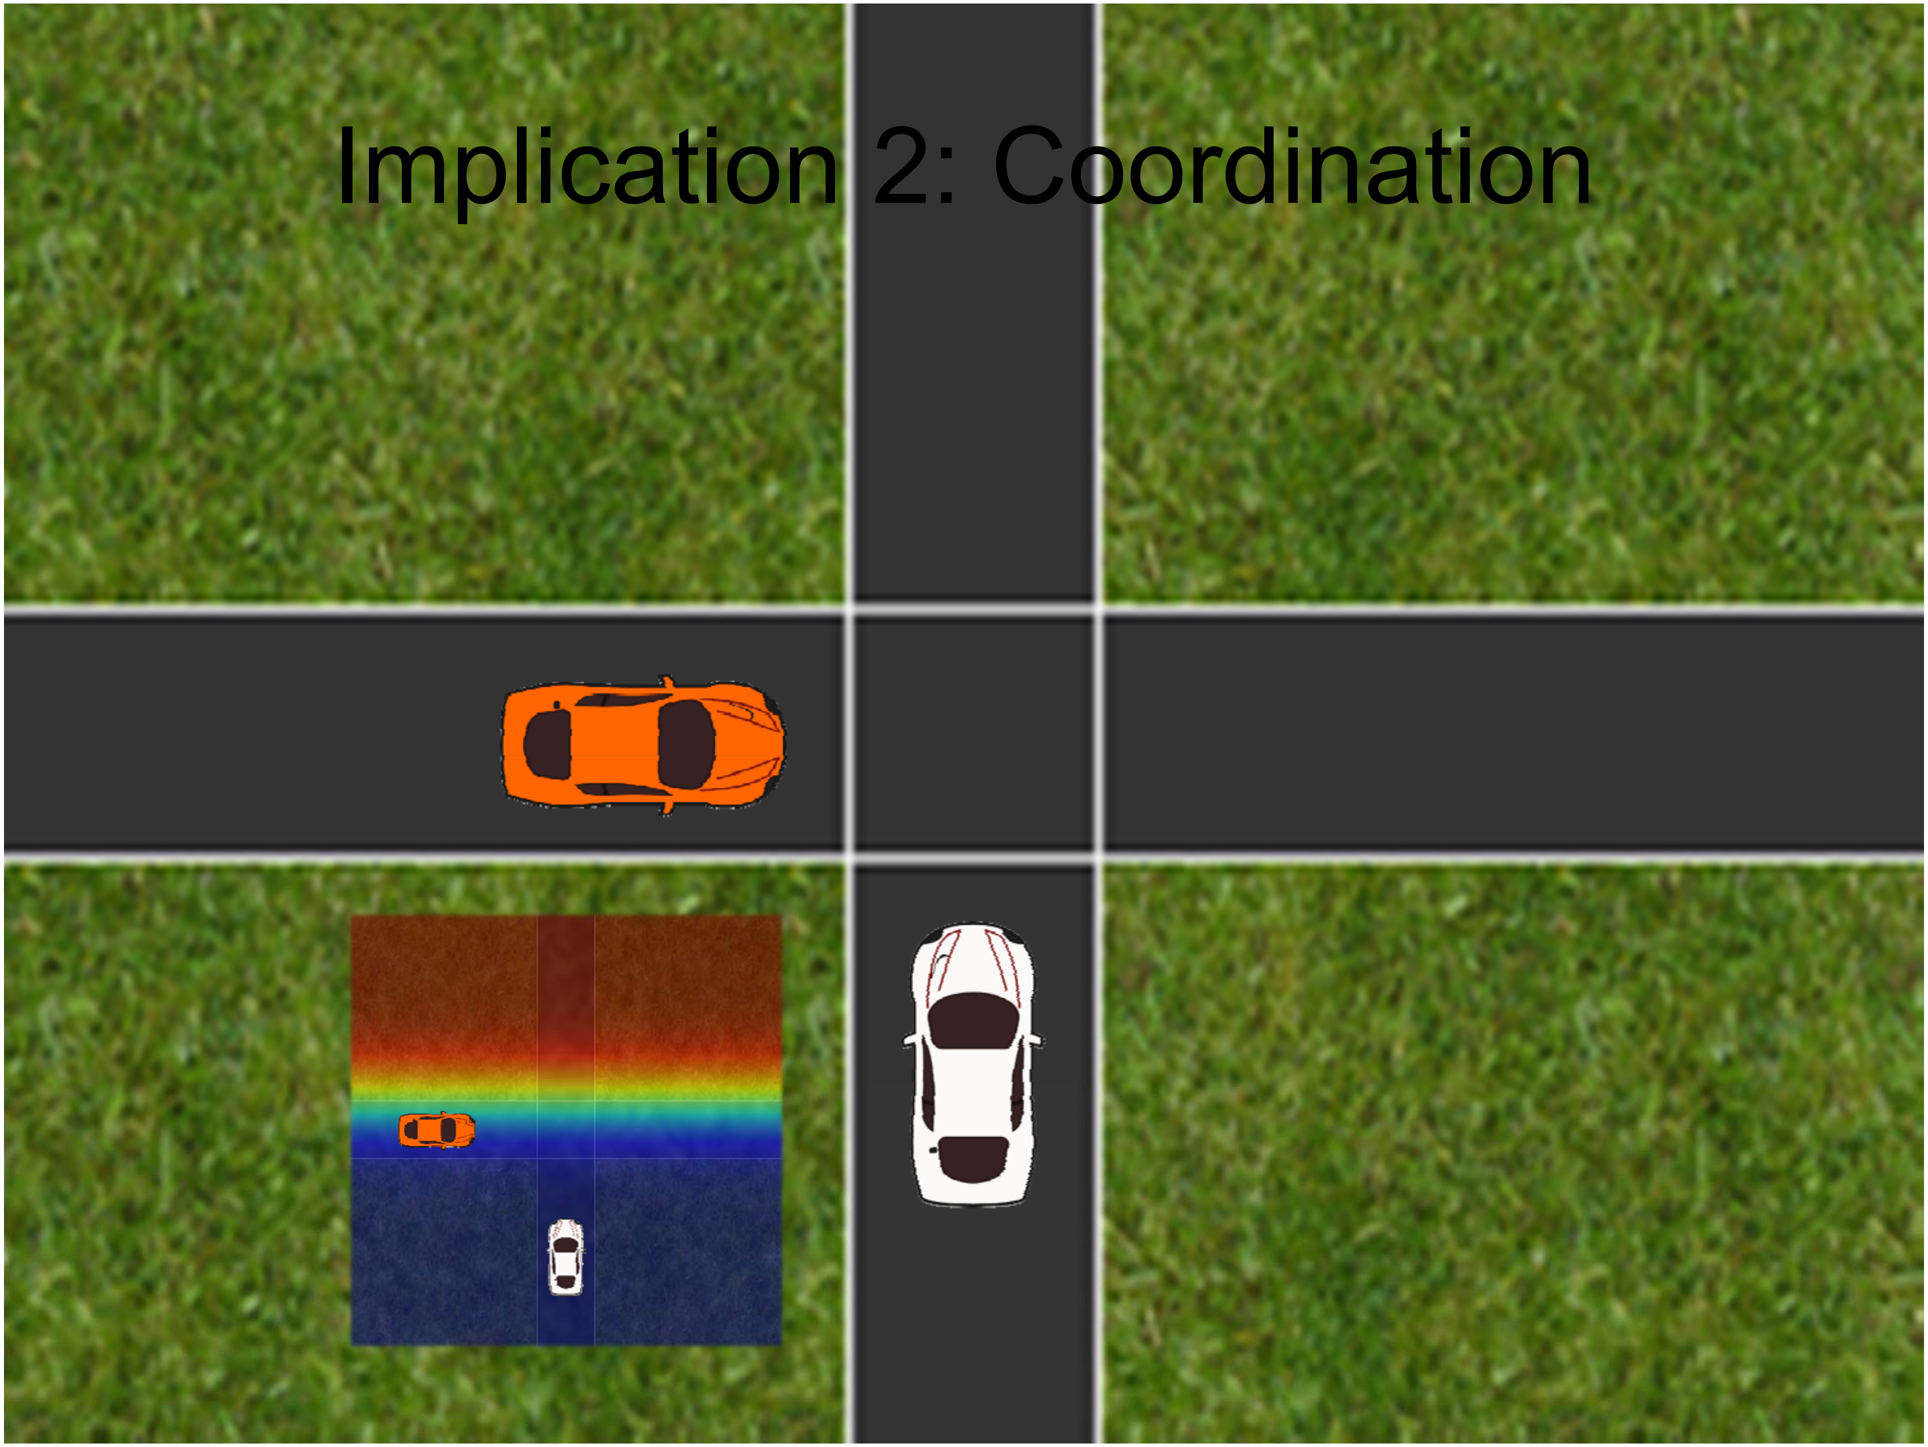
Implication 2: Coordination

Implication 2: Coordination

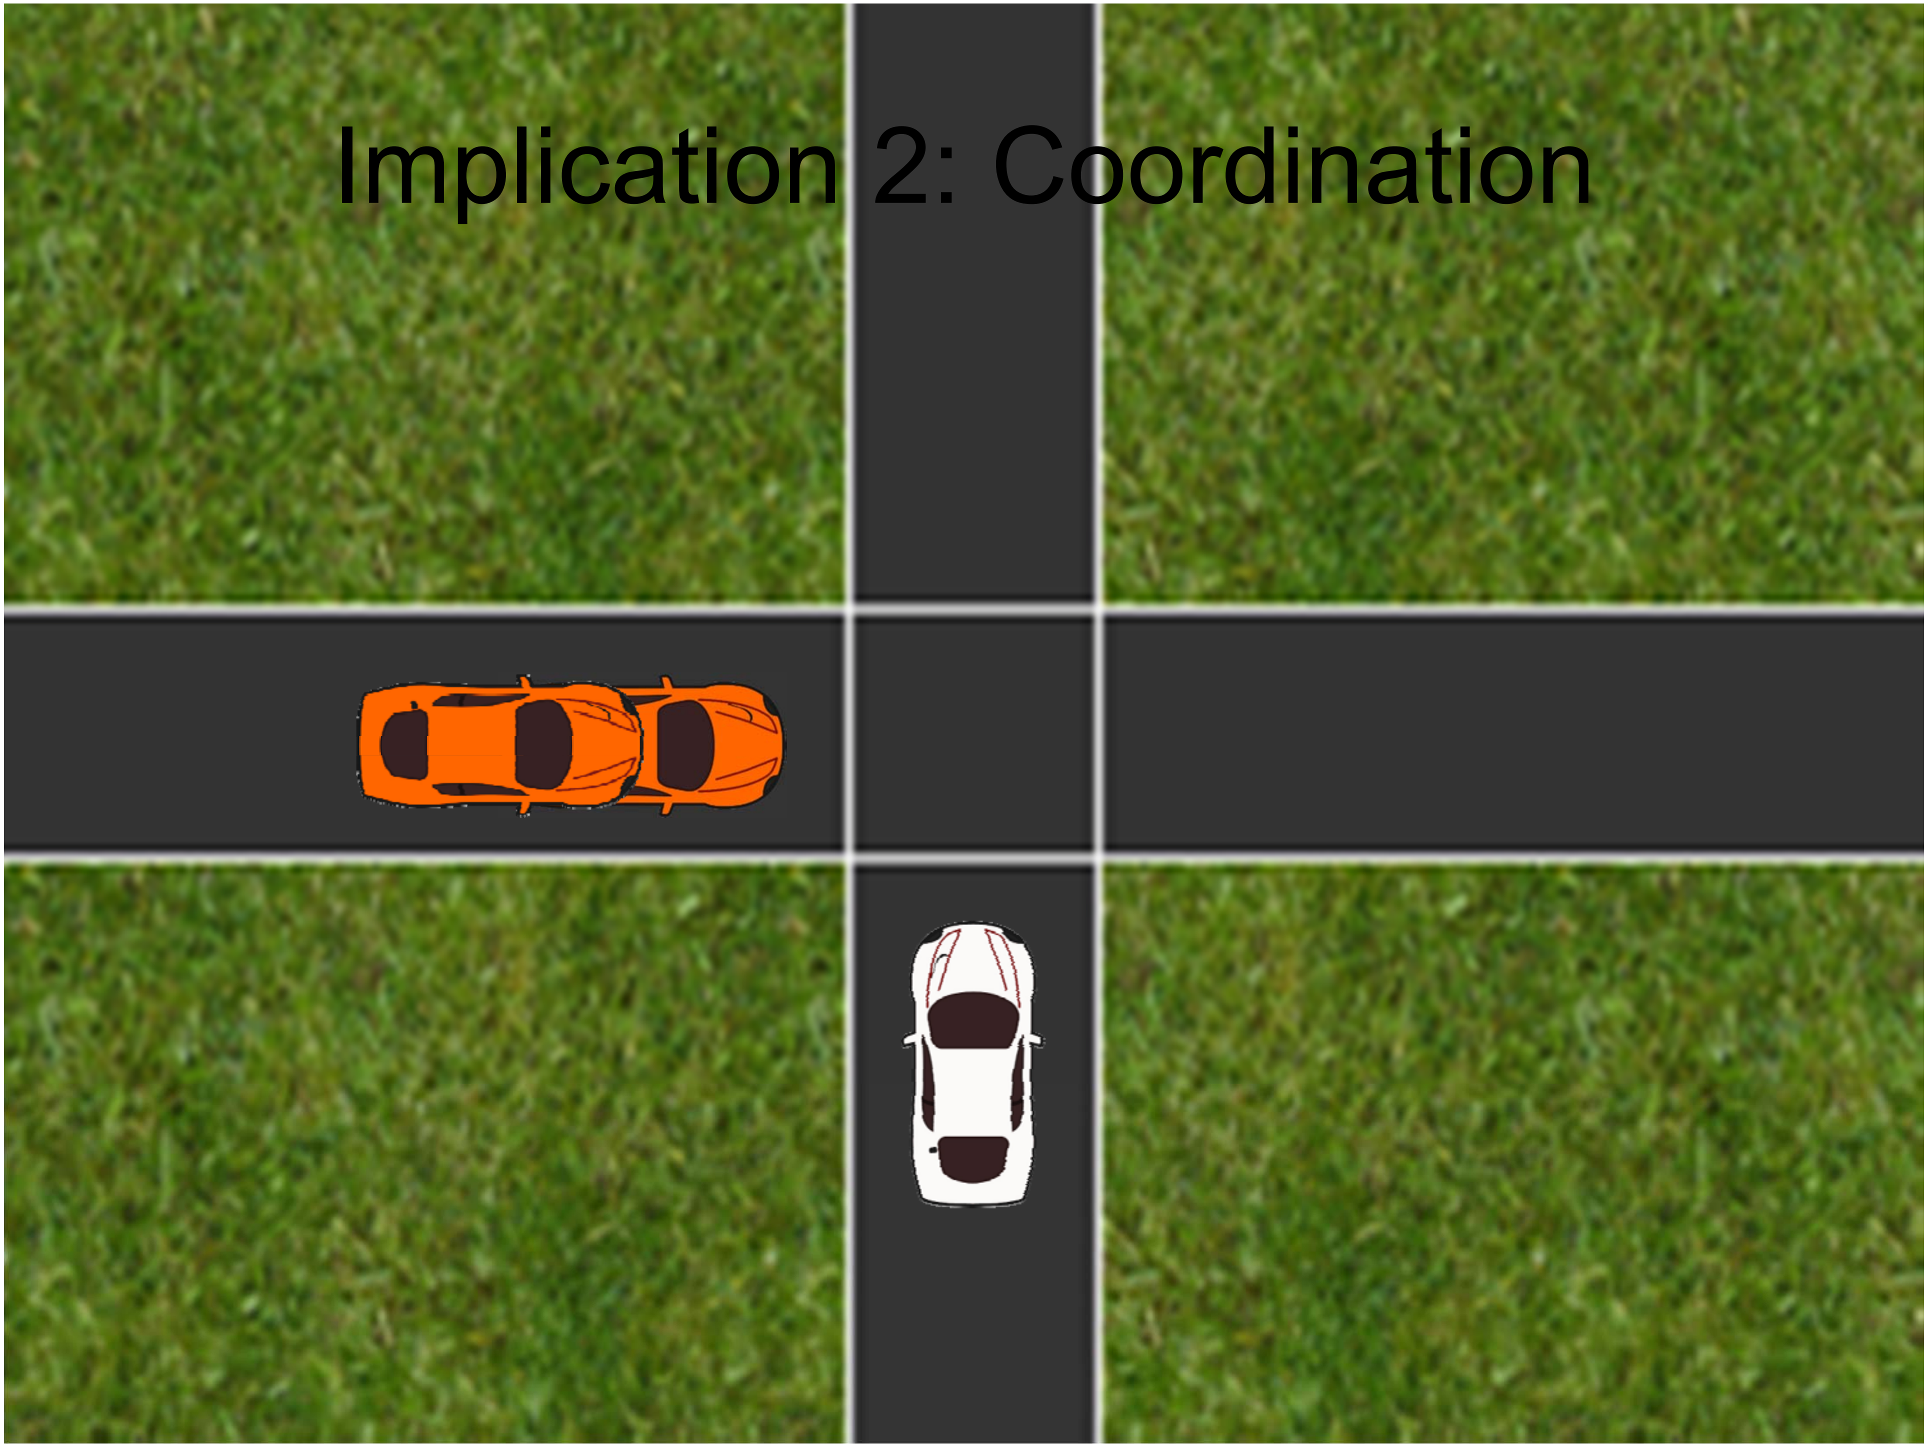


acer

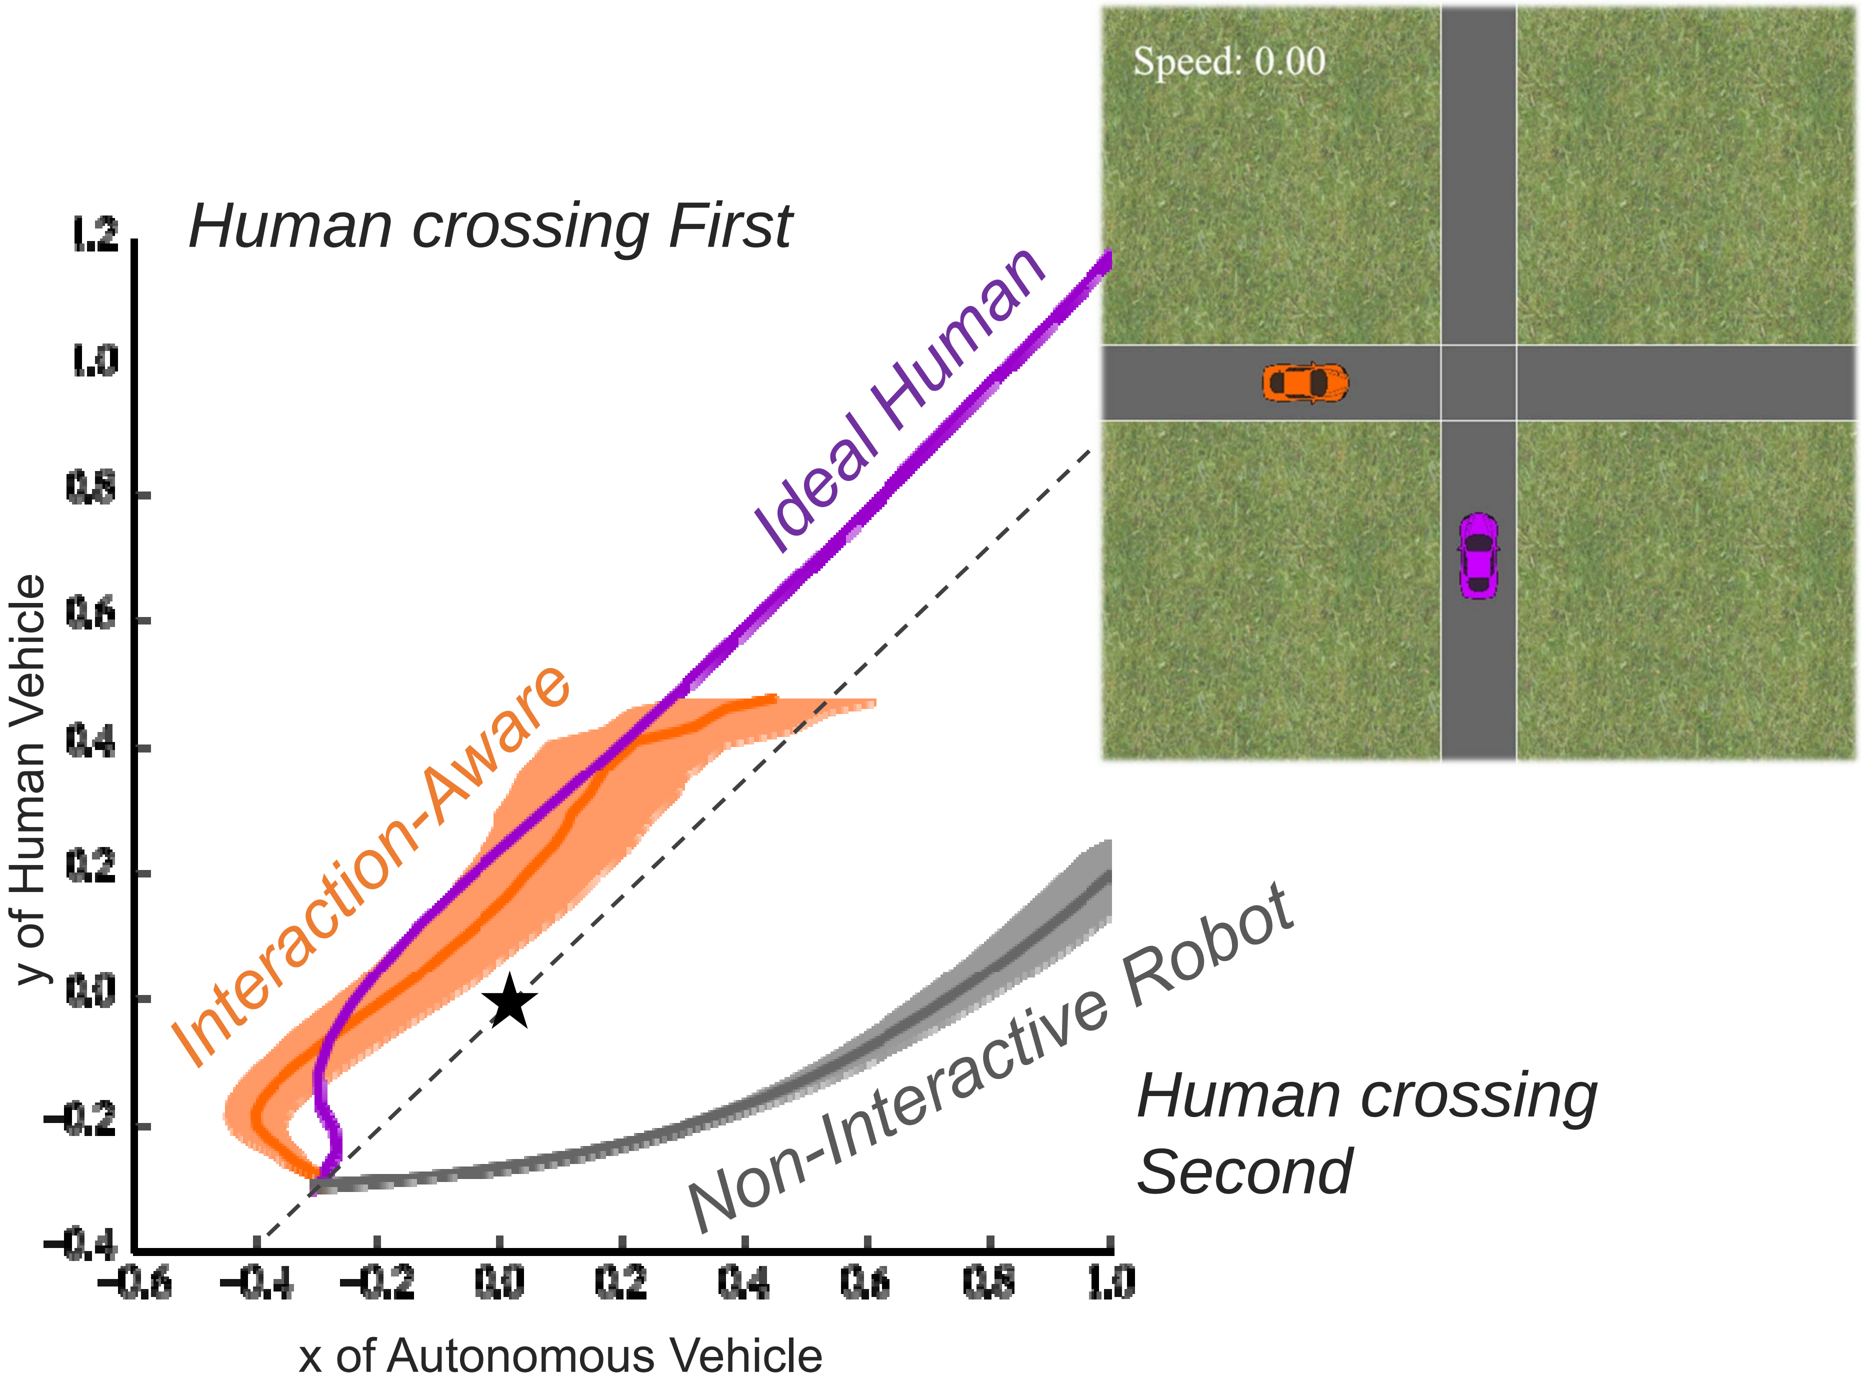

Human crossing First
Ideal Human
Interaction-Aware
Non-Interactive Robot
Human crossing Second
Speed: 0.00
y of Human Vehicle
x of Autonomous Vehicle
1.2
1.0
0.8
0.6
0.4
0.2
0.0
−0.2
−0.4
−0.6
−0.4
−0.2
0.0
0.2
0.4
0.6
0.8
1.0

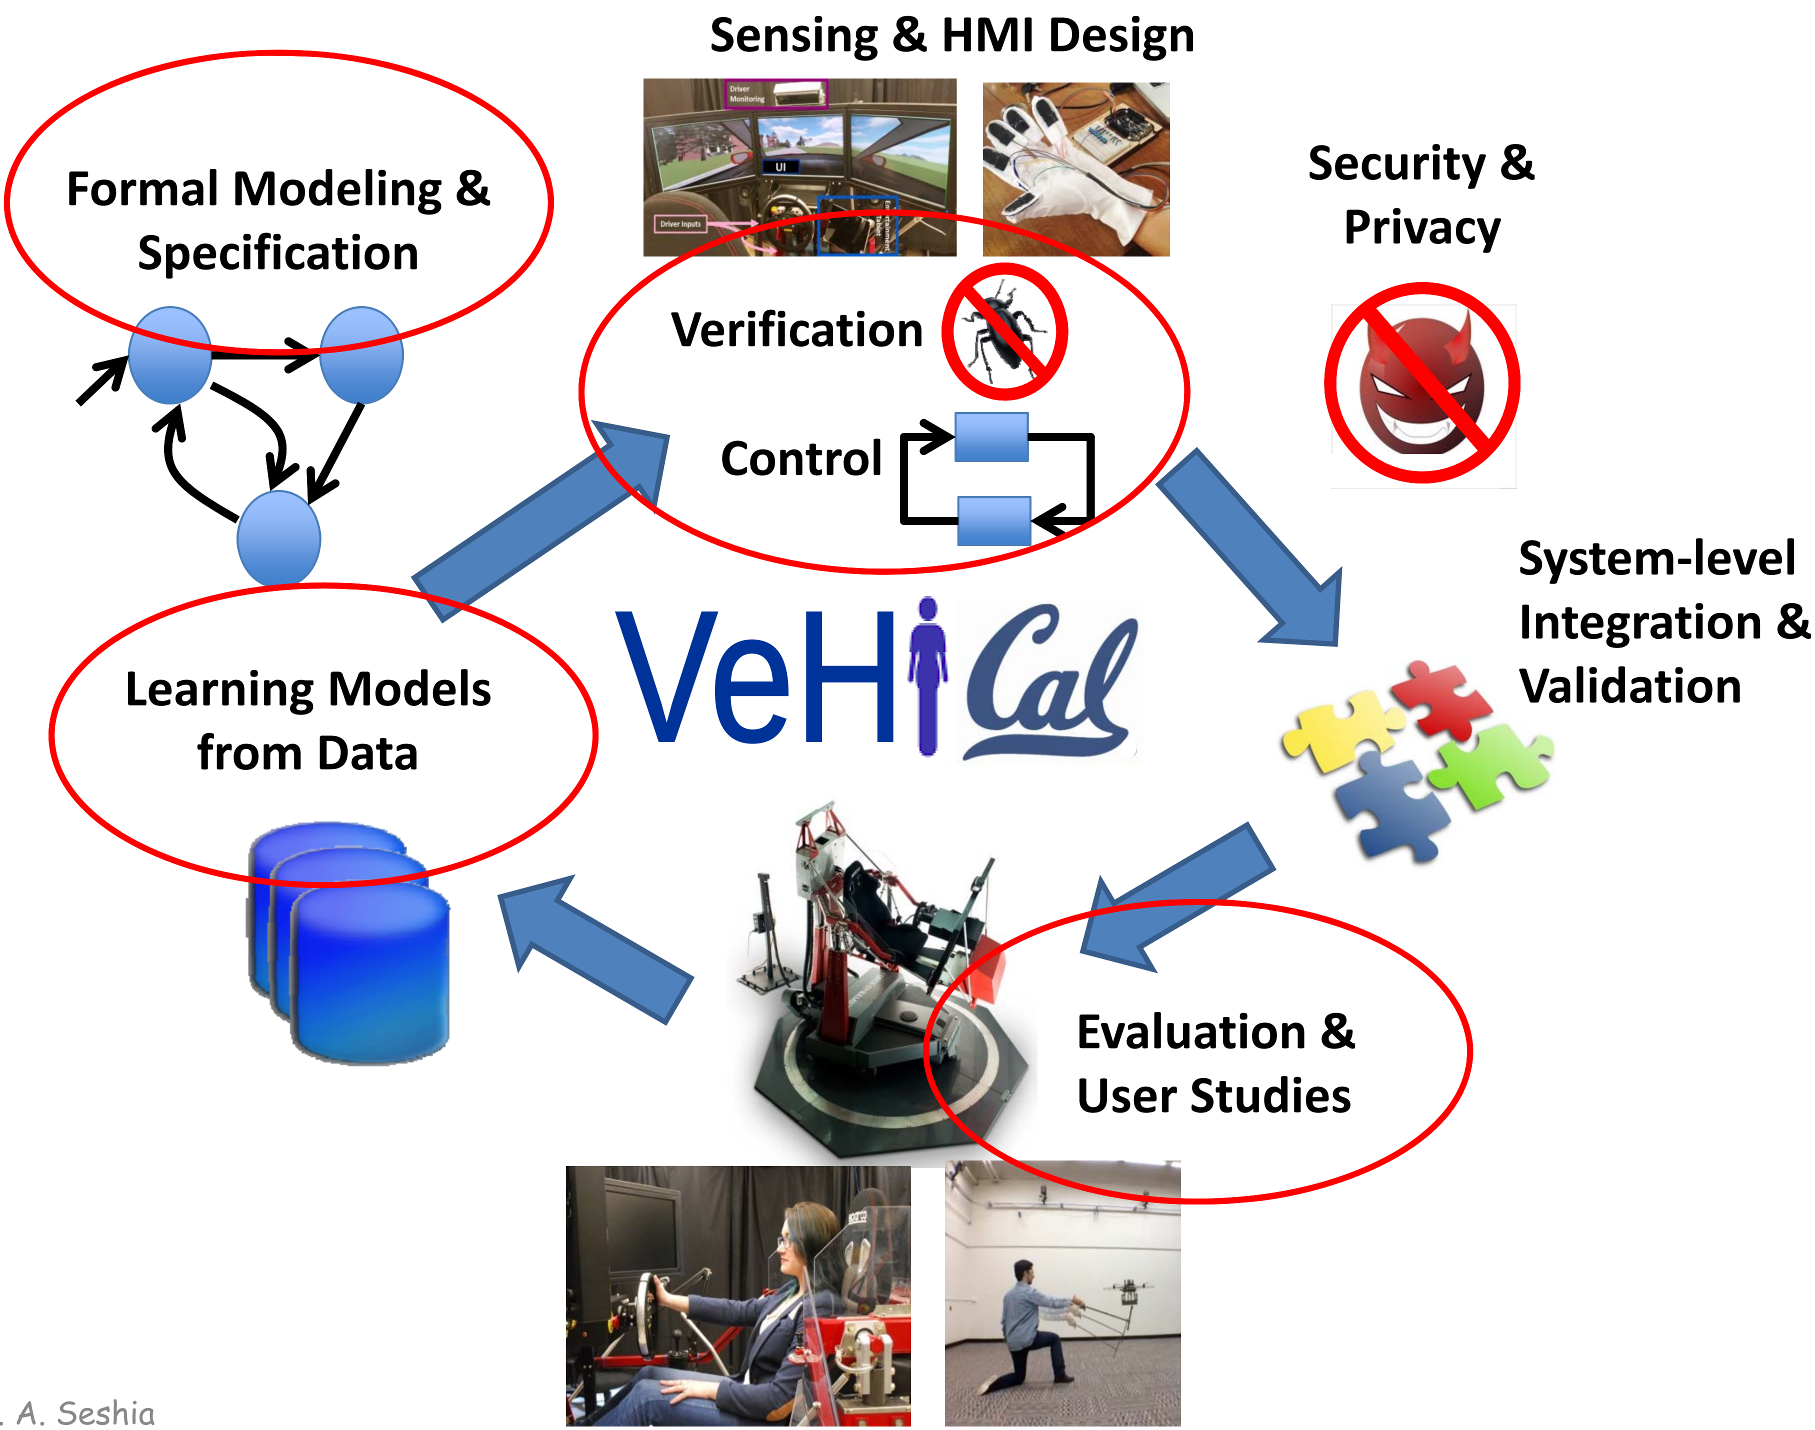
Sensing & HMI Design
Formal Modeling & Specification
Security & Privacy
Verification
Control
System-level Integration & Validation
Learning Models from Data
VeHi Cal
Evaluation & User Studies

Questions for the Meeting

- Human Modeling:
  - Multiple approaches/formalisms: what are the pros and cons of each?
- Learning Systems:
  - What are the unique *specification, verification, and control* problems that arise?
- Interfaces:
  - How do we best integrate control design with interface design?
- Challenge Problems:
  - What are the best integrative challenge problems to tackle (from our target domains)?