Game-Theoretic Modeling of Driver/Vehicle Interactions in Traffic

Anouck Girard
Aerospace Engineering
University of Michigan, Ann Arbor



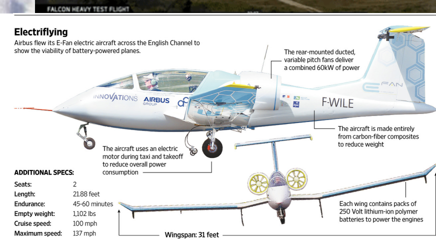


Collaborators/Funding/Thanks

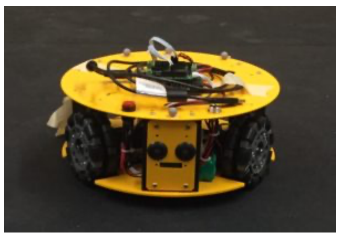

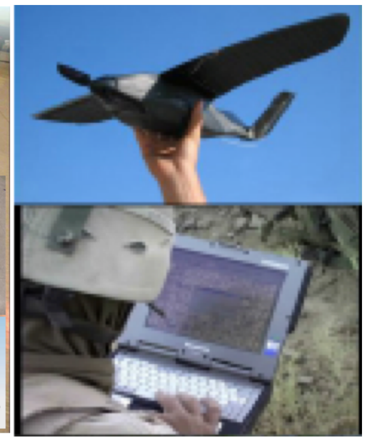

- **Vehicle Optimization, Dynamics, Control and Autonomy Laboratory:** Dynamics, control and optimization of advanced, increasingly autonomous vehicles operating in the space, air, ground or marine domains.

- **Collaborators:** **I. Kolmanovsky,** A. Berning, W. Dunham, **N. Li,** R. Sutherland, R. Tian

- **Funding Sources:** AFRL, AFOSR, NASA, NSF, ONR, TARDEC, Boeing, Luna Rossa, Oracle, and the automotive industry.

- **Website:** vodca.engin.umich.edu

Robots are increasingly capable and prevalent

Robots, including collaborative robots, are becoming more capable, better accepted, and more commonplace.

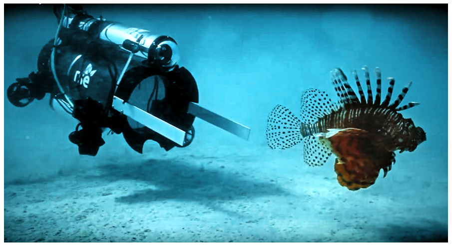

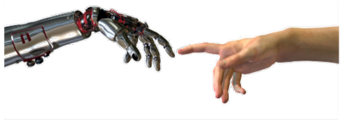

Unmanned Air Vehicles are here

- ❑ Estimated number of UAVs in US by 2020: 7 million
- ❑ Estimated number of commercial UAVs in the air by 2018: 600,000
- ❑ Number of UAVs registered with the FAA: 770,000
- ❑ Estimated value of UAV industry: $ 3.3 billion
- ❑ Projected value of UAV industry in 2025: $ 90 billion
- ❑ Percentage of Americans who own a “drone:” 8%

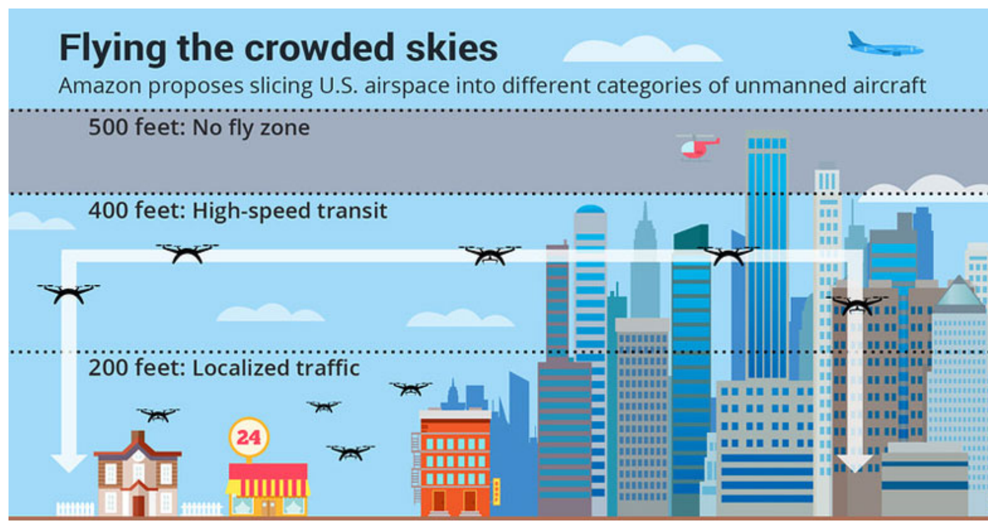

Self-driving cars are being tested

MOLLY MCHUGH GEAR 10.14.15 05:19 PM

TESLA'S CARS NOW DRIVE THEMSELVES, KINDA

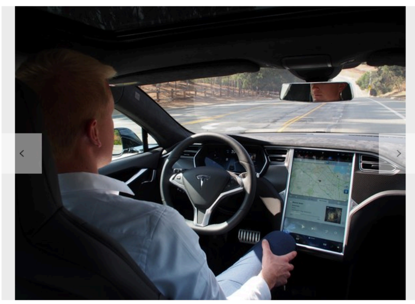

1 / 7 Technically, it's advised to keep hands resting on the wheel—but you can go hands-free.
MOLLY MCHUGH/WIRED

Self-driving cars: from 2020 you will become a permanent backseat driver

Driverless cars will revolutionise motoring, claim the manufacturers. But is the greatest danger that they will be too safe?

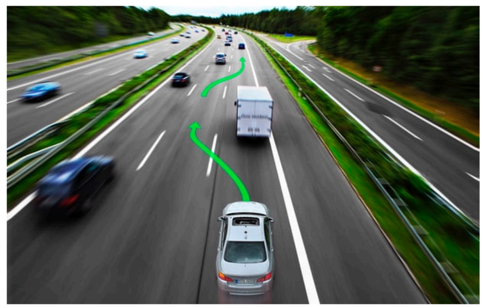

▲ A BMW 'highly automated' prototype on the German autobahn. Photograph: PR

Self-driving cars: it's only a matter of time until they take over

Nicholas Tucker, Staff Reporter
April 11, 2018
Filed under Opinion

Self-driving trucks are coming—and this law just made things even worse for truckers

Safety/validation concerns are relevant

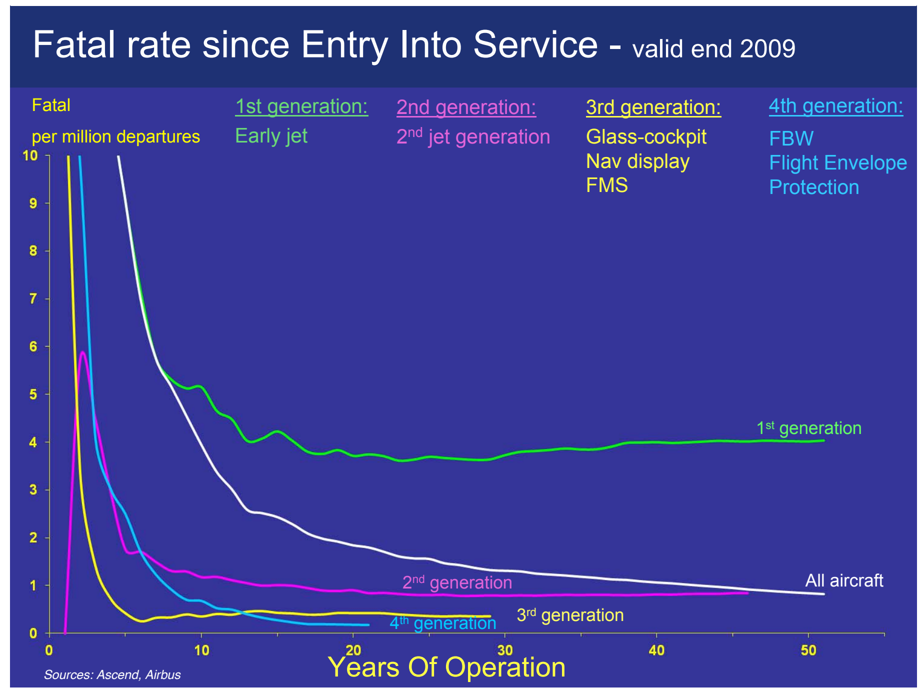

D. Chatrenet, Air transport safety technology and training, ETP 2010

My research: control/autonomy for multi-vehicle systems

- Research at the **intersection of computer science, systems theory and control systems** for **aerospace and autonomous vehicle applications**
- **Emphasis on interacting systems** (multiple vehicles, actors, or subsystems)
- **Mixed-initiative operations** (manned and unmanned vehicles, adjustable autonomy):
  - Human pilot/driver modeling (game theory),
  - Integration of unmanned vehicles into the airspace,
  - Task scheduling, allocation and planning with adjustable autonomy,
  - Collaborative/adversarial environments.

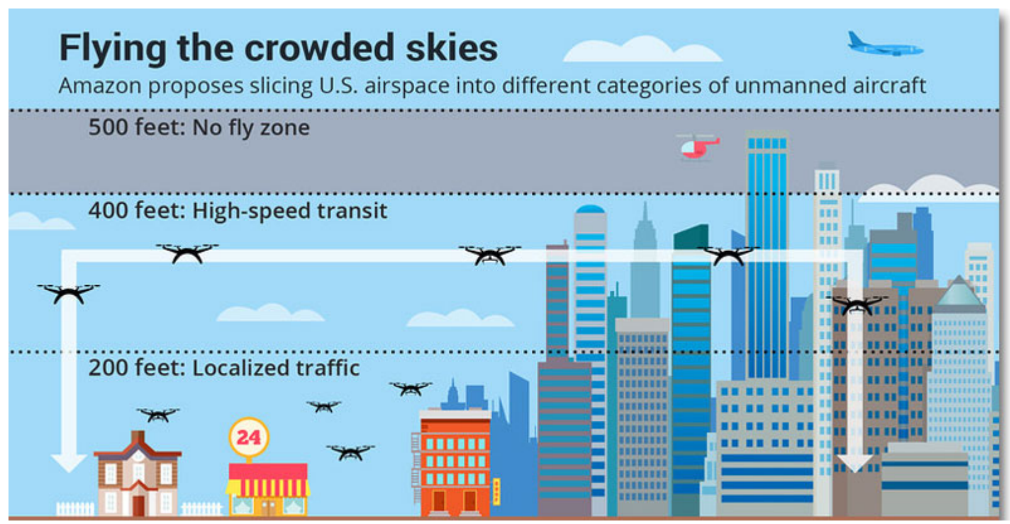

... and integrated control of multiple vehicle subsystems

- **Integrated vehicle control,** including coupled control of the propulsion and power systems for aircraft, with the inclusion of heat management constraints, in light of more or all electric aircraft, design considerations for the 6th generation fighter, and the increasing use of small UAVs.

- **Predictive and nonlinear control,** as made possible by increasingly capable flight hardware.

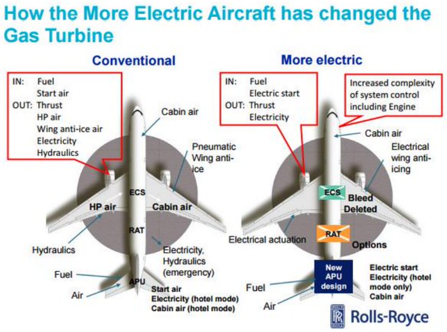

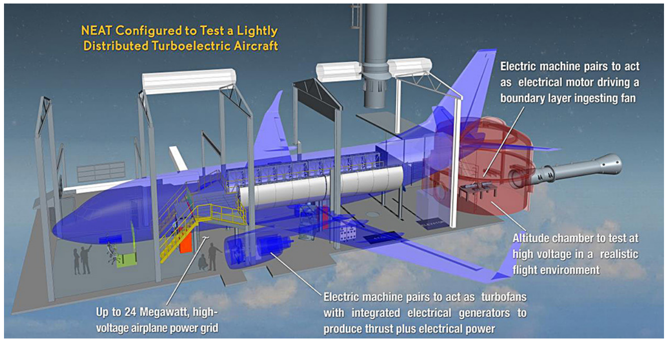

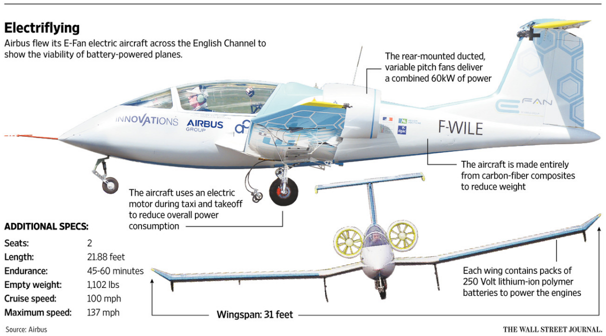

Research in distributed control of AV

Autonomy and Decision Making	PE games (UAV) 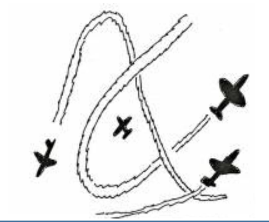	TSP stability regions (UAVs) 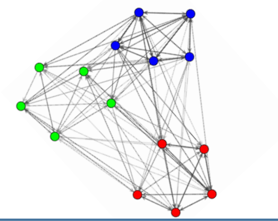	Game theoretic traffic simulation 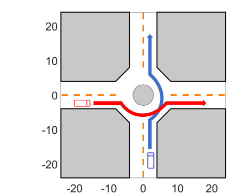
Path Planning, Trajectory Opt.	UAV integration in airspace 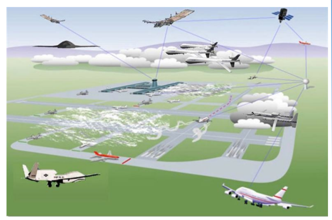	Information collection 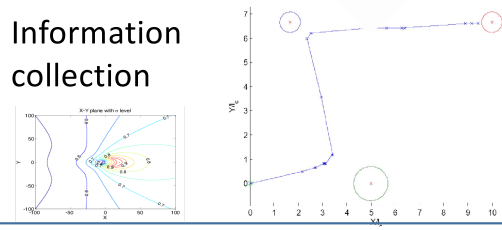	Formation flight 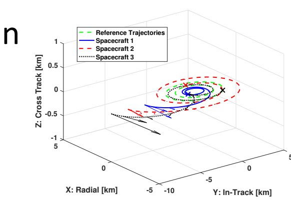
Energy Management	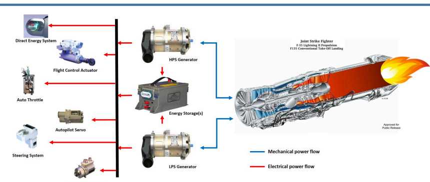 MEA/AEA/ fighters	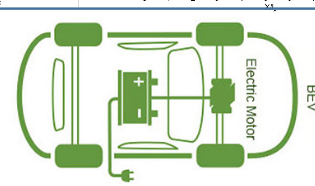 Automotive	Disjunctive sensing/ control 
Unusual Configurations		Gyrocart 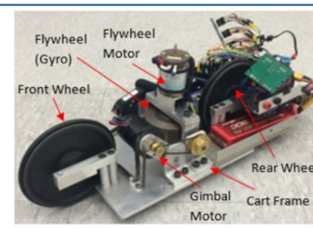	AC 

Today’s Topic

Autonomy and Decision Making	PE games (UAV) 	TSP stability regions (UAVs) 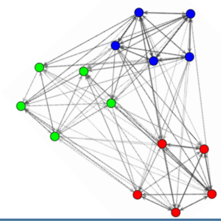	Game theoretic traffic simulation 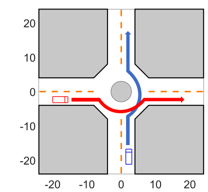
Path Planning, Trajectory Opt.	UAV integration in airspace 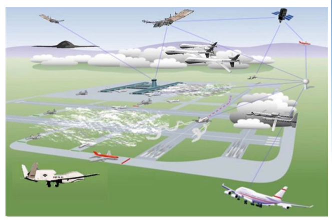	Information collection 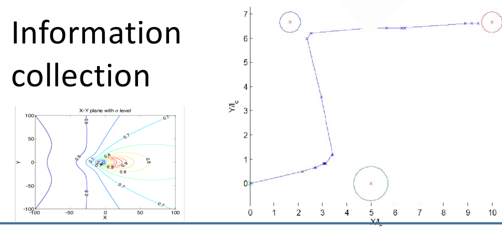	Formation flight 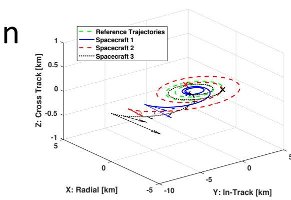
Energy Management	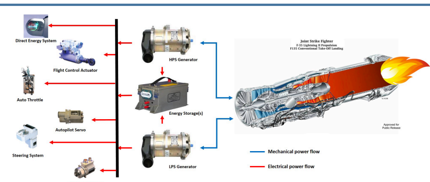 MEA/AEA/ fighters	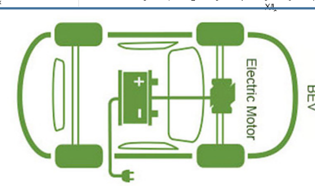 Automotive	Disjunctive sensing/ control 
Unusual Configurations		Gyrocart 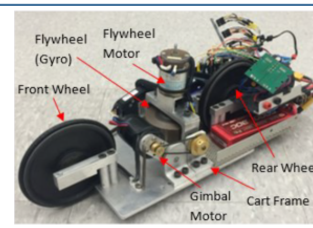	AC 

1995-2005s: Mostly ignore the humans

2005-15: Control theory for humans

Intelligence, Surveillance, and Reconnaissance missions

- Objects of interests
  - threat or friend
- Unmanned aerial vehicles (UAVs)
  - carry a suite of sensors and a communication device
- Human operators
  - direct the UAVs
  - inspect data and make classification decisions
- Need high quality classification decisions in the presence of uncertainties

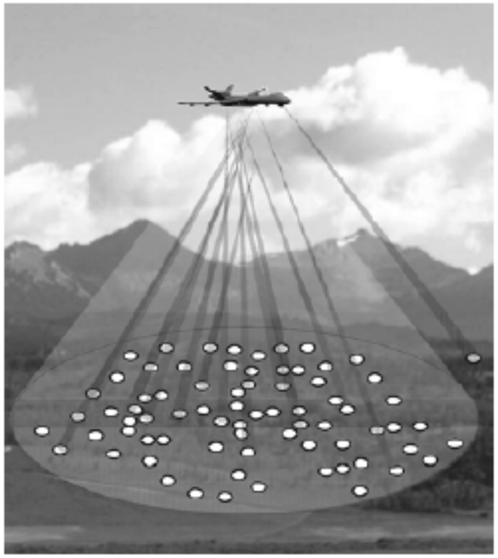

"Inverting the ratio"

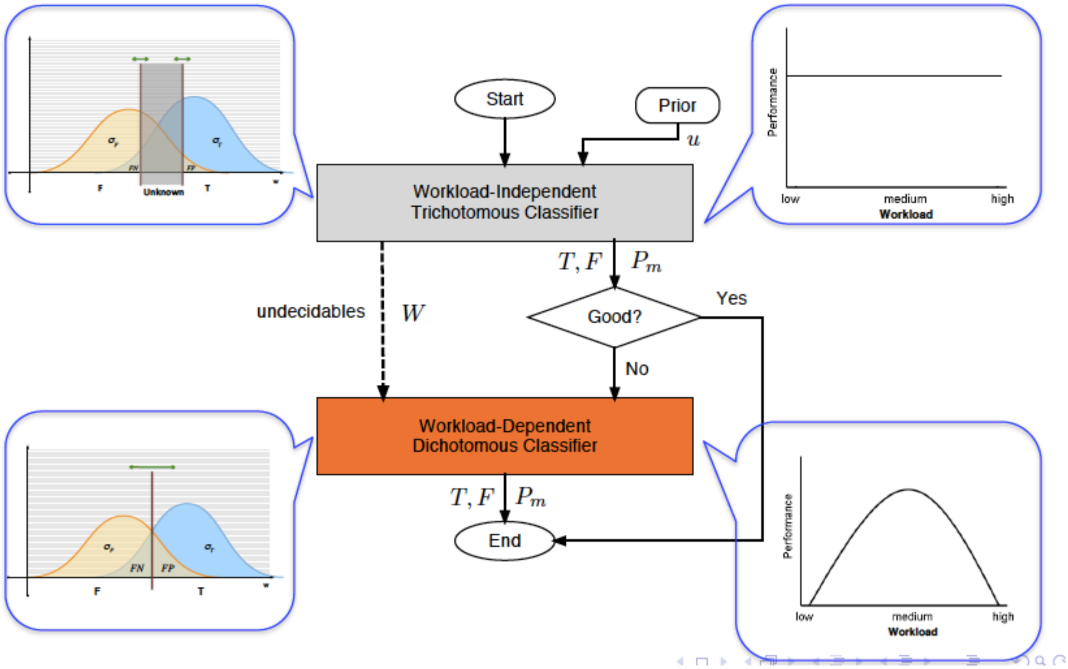

Some results:

- More information (in the sense of Shannon) does not guarantee better classification performance.
- Knowing what the information is used for (e.g., classification) allows one to collect more useful information.
- Team theory (how to combine classifiers, how to give feedback to human operators)
- Main lesson: Humans exhibit a lot of variability, and are difficult to model

AV Challenges

- In the near to medium term, AV penetration will be small
- AVs will operate in traffic together with human-driven vehicles (HVs)
- AV control algorithms must account for the interactions between AVs and HVs
- This is a harder problem than in an “all-AVs” world

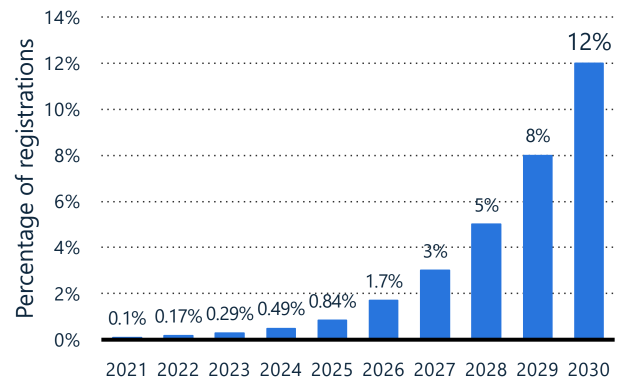

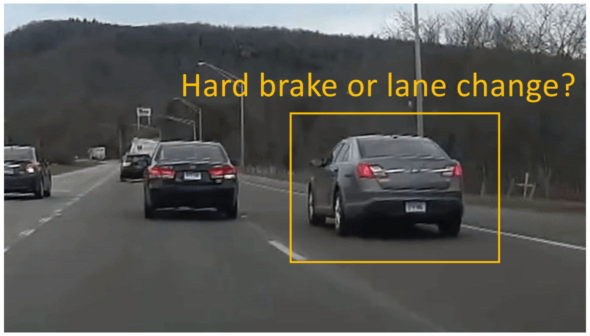

Why do we need traffic simulation?

- AV control algorithms must account for the interactions between AVs and HVs
- Estimated that one must drive **>100 million kilometers** to validate autonomous driving software and assure against "faults" in the algorithms, e.g., conditions under which the algorithm may cause unsafe driving behavior
- Verification and validation (V&V) in virtual world (i.e., using simulation tools) is appealing. However, simulators must be able to represent realistic vehicle interactions in traffic
- This leads to the question: How to model and account for the interactions between AVs and HVs?

RAND CORPORATION

Driving to Safety

How Many Miles of Driving Would It Take to Demonstrate Autonomous Vehicle Reliability?

Nidhi Kalra, Susan M. Paddock

Table 1. Examples of Miles and Years Needed to Demonstrate Autonomous Vehicle Reliability

		Benchmark Failure Rate		
	How many miles (years[a]) would autonomous vehicles have to be driven…	**(A) 1.09 fatalities per 100 million miles?**	**(B) 77 reported injuries per 100 million miles?**	**(C) 190 reported crashes per 100 million miles?**
Statistical Question	(1) without failure to demonstrate with 95% confidence that their failure rate is at most…	275 million miles (12.5 years)	3.9 million miles (2 months)	1.6 million miles (1 month)
	(2) to demonstrate with 95% confidence their failure rate to within 20% of the true rate of…	8.8 billion miles (400 years)	125 million miles (5.7 years)	51 million miles (2.3 years)
	(3) to demonstrate with 95% confidence and 80% power that their failure rate is 20% better than the human driver failure rate of…	11 billion miles (500 years)	161 million miles (7.3 years)	65 million miles (3 years)

[a] We assess the time it would take to compete the requisite miles with a fleet of 100 autonomous vehicles (larger than any known existing fleet) driving 24 hours a day, 365 days a year, at an average speed of 25 miles per hour.

Game Theoretic Traffic Simulation

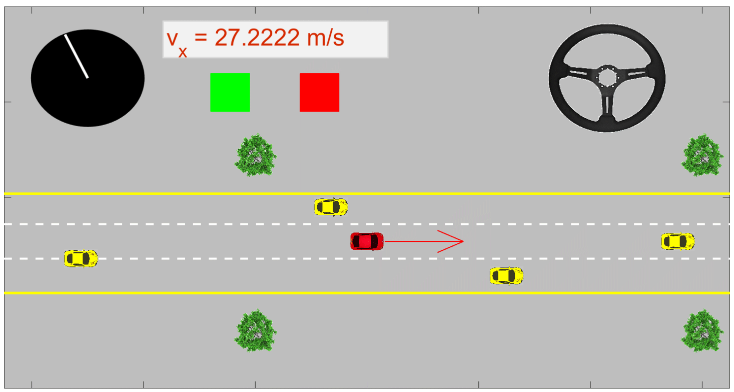

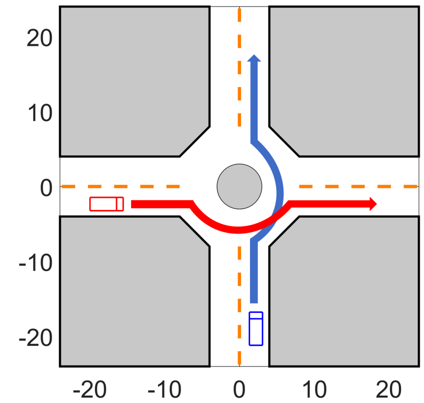

Main Collaborators: Ilya Kolmanovsky, Yildiray Yildiz, Nan Li, Ran Tian.

Game theory

Game theory: The study of strategic decision-making
The payoff (reward/cost) received by each participant depends on her own action and the actions of the other participants of the game.

Game	⟷	Traffic Scenario
Participants	⟷	Drivers/Vehicles
Payoffs	⟷	Driving Objectives (Safety, Performance, Comfort, …)

Conventional tool: Nash equilibrium
Problems:
1) Outcomes of experimental subjects systematically violate Nash-equilibrium predictions
2) Nash equilibrium is difficult to compute for multi-player games (traffic scenarios with multiple interacting vehicles)

Game theoretic approaches for traffic simulation

Hierarchical reasoning/Level-k approach [1,2]

- Players have bounded rationality and are categorized by the "depth" of their strategic reasoning (called **level**).
- A level-*k* player assumes all other players are level-(*k*-1) and makes her own decision as the optimal response to their level-(*k*-1) decisions.
- Use different *k* to model different driving styles (aggressive/conservative) and driving intentions (proceeding/yielding)

[1] Li, Oyler, Zhang, Yildiz, Kolmanovsky and Girard (TCST 2018).
[2] Tian, Li, Kolmanovsky, Yildiz and Girard (TITS 2020).

Leader-follower/Stackelberg-like approach [3]

- Players have asymmetric roles: Leader vs Follower.
- The leader moves before the follower and has the **first-mover advantage**: The leader can influence the follower's decision through her own action choice.
- Use the leader-follower asymmetry to represent "right of way" rules.

[3] Li, Yao, Kolmanovsky, Atkins and Girard (TITS 2020).

Hierarchical reasoning game theory

- **Modeling humans precisely is difficult** (Variability, small data sets).
- **Hierarchical reasoning game theory (level-k game theory)** attempts to describe human thought processes in strategic games. Assumes that players base their decisions on their predictions of the likely actions of other players.
- Players in strategic games can be categorized by the “depth” of their strategic thought. Players have bounded rationality.
- Level-k theory **predicts human decisions better than equilibrium-based models** in a range of games.

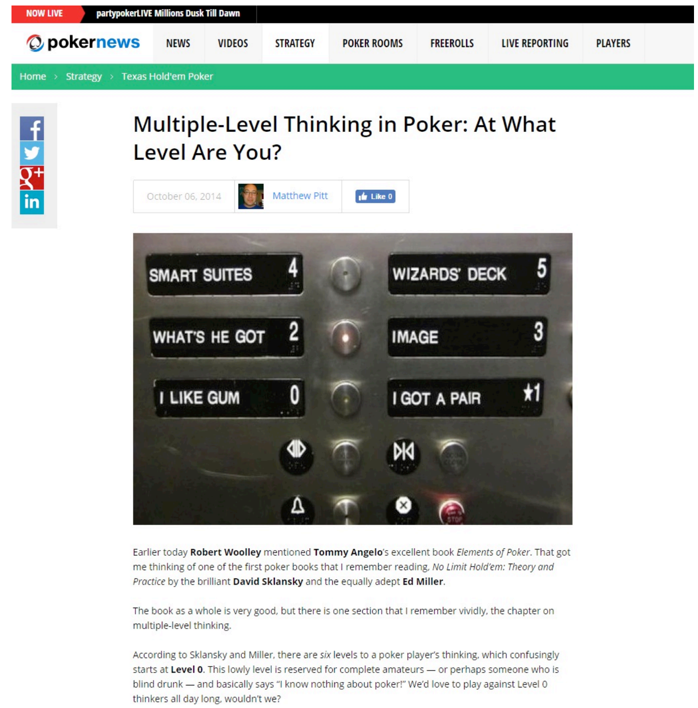

NOW LIVE partypokerLIVE Millions Dusk Till Dawn

pokernews NEWS VIDEOS STRATEGY POKER ROOMS FREEROLLS LIVE REPORTING PLAYERS

Home › Strategy › Texas Hold'em Poker

Multiple-Level Thinking in Poker: At What Level Are You?

October 06, 2014 Matthew Pitt Like 0

Earlier today **Robert Woolley** mentioned **Tommy Angelo**’s excellent book *Elements of Poker*. That got me thinking of one of the first poker books that I remember reading, *No Limit Hold’em: Theory and Practice* by the brilliant **David Sklansky** and the equally adept **Ed Miller**.

The book as a whole is very good, but there is one section that I remember vividly, the chapter on multiple-level thinking.

According to Sklansky and Miller, there are *six* levels to a poker player’s thinking, which confusingly starts at **Level 0**. This lowly level is reserved for complete amateurs — or perhaps someone who is blind drunk — and basically says “I know nothing about poker!” We’d love to play against Level 0 thinkers all day long, wouldn’t we?

An example: the Keynesian “beauty” contest

- **Keynesian beauty contest:** Pick a number between 0 and 100. Winner is the person whose number is the closest to half of the average of all the (many) participants' guesses.

- **Level zero players:** choose a number non-strategically (at random, 42 for Doug Adams, birthday, etc.)

- **Level one players:** choose their number consistent with the belief that all other players are level zero. If all other players in the game are level zero, the average of those guesses would be about 50. Therefore, a level one player chooses 25.

- **Level two players:** choose their number consistent with the belief that all other players are level one. Since a level one player will choose 25, a level two player should choose 12 or 13. This process repeats for higher-level players.

- Studies from other domains show **humans are usually level 0, 1 or 2, rarely 3**.

Level-k reasoning and reinforcement learning

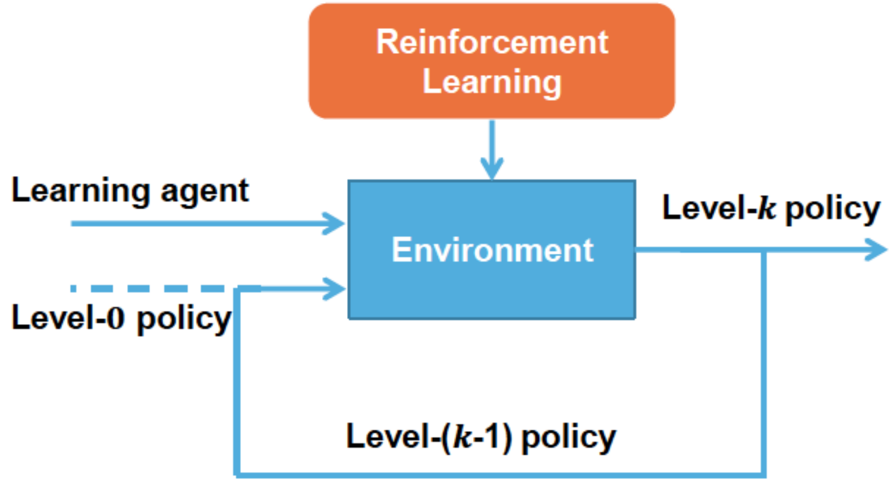

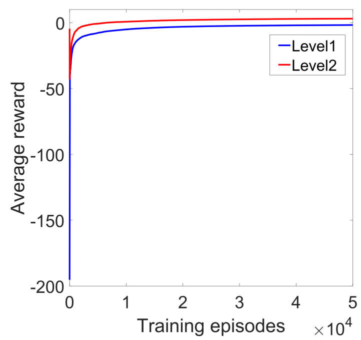

- **Level-0** agent makes instinctive decisions and does not take into account the interactions.
- **Level-1** agent assumes all of the other agents are **Level-0** and makes optimal responses based on this assumption.

⋮

- **Level-k** agent assumes all of the other agents are **Level-(k-1)** and makes optimal responses based on this assumption.

- To obtain the level-k policy, we put a learner in traffic consisting of level-(k-1) drivers, and use reinforcement learning to train the learner. In other domains, humans have been shown to usually be levels 0, 1 or 2 decision makers.

Markov Decision Processes

- **Markov Decision Process (MDP):** A stochastic transition system + a reward function
  - A set of states, $\mathcal{S}$
  - A set of actions, $\mathcal{A}$
  - A transition probability function, T: $\mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R}^+$
  - An initial state, $s_0 \in \mathcal{S}$
  - A reward function, R: $\mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R}^+$

- If action a is applied from state s, then a transition from state s to state s' occurs with probability T(s, a, s'). If this transition is made, a reward R(s, a, s') is collected.

- Both transition probabilities and rewards only depend on the present state, not on the history of the state (future states and rewards are independent of the past, given the present).

- The state is known exactly (only the transitions are stochastic). If the state is not known exactly: **Partially Observable Markov Decision Process (POMDP)**.

MDP for highway driving

❑ Control Hierarchy. Focus: Higher-level Controller

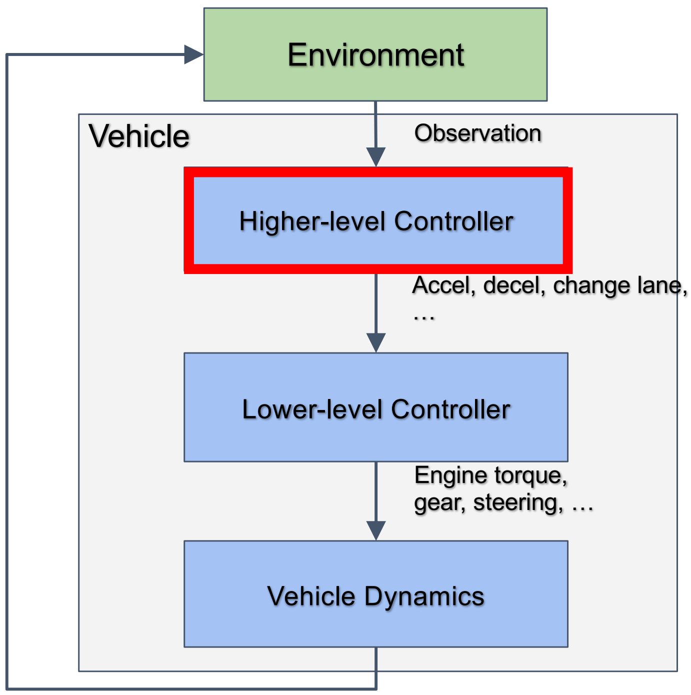

- **Set of states:**
  - 3 states per vehicle, longitudinal position and velocity, lateral position
  - 30-50 vehicles in simulation

- **Set of actions: 7 dimensional**
  - Maintain
  - Change lane left/right
  - Accelerate/decelerate (regular maneuver)
  - Emergency Accelerate/decelerate

- **Reward/Penalty:** (safe distance) constraint violation, vehicle speed, maneuver effort, headway

$$R = w_1 c + w_2 v + w_3 e + w_4 h$$

- **Observation space (POMDP): 18-dimensional**
  - Observe states of 5 neighboring vehicles + self

Q Learning (Reinforcement Learning)

- In some cases, the exact details of the MDP (the transition probabilities) are not known. **Q-learning** (a reinforcement learning technique) **estimates the total collected reward for state-action pairs.**

- The **Q-factor, Q(s, a),** is an estimate of the total collected reward collected by
  i. starting at state s,
  ii. applying action a at the first time step,
  iii. acting optimally for all future times.

- The Q-factor update law, based on an observed transition from s to s' under action a, is

$$Q(s,a) \leftarrow (1-\alpha_t)Q(s,a) + \alpha_t(R(s,a,s') + \gamma \max_{a'} Q(s',a'))$$

- **Exploration/Exploitation trade-off created.**

- γ is a discount factor (Bellman equation); α_t should decay over time for convergence (for example, as 1/t).
- May use approximation techniques to deal with large state spaces. This may cause instability.

- Use Jaakkola version for POMDP.

Application to highway driving

- **Policy:** a rule that determines a decision given the available information at state s.

$$\pi^*(s) = \arg\max_a Q^*(s,a) \; \forall s \in \mathcal{S}$$

- **In practice:**
  - Discretize states (message elements).
  - Create simulation environment, initialize with n cars.
  - Assign the level-0 policy to n-1 cars. This “replaces” the need for a transition probability model for those cars.
    - **Level-0:** Drive at constant speed unless vehicle in front. If vehicle in front, maintain headway.

  - Use Jaakkola RL to obtain Q(m,a) for the ego car. The message m replaces the state s in POMDP problems. Simulation time: approximately 2 days.

  - Obtain **stochastic policy** from Q(m,a).
    - A **deterministic policy** always returns the same action (that with the highest expected Q value).
    - A **stochastic policy** models a distribution over actions and draws an action according to this distribution.

Game theoretic highway driving

- **A dynamic game of two players**

$$\langle \mathbf{P}, \mathbf{X}, \mathbf{U}, \mathbf{T}, \mathbf{R} \rangle$$

- Two players
  $\mathbf{P} = \{1, 2\}$
- State of the game
  $\mathbf{x} = (x^1, x^2) \in \mathbf{X}$
- Actions of the players
  $\mathbf{u} = (u^1, u^2) \in \mathbf{U} = U^1 \times U^2$
- Dynamics of the game, can be stochastic
  $\mathbf{T} : (\mathbf{x}, \mathbf{u}) \to \mathbf{x}^+$ $\quad \mathbb{P}\big(\mathbf{x}^+ | \mathbf{x}, \mathbf{u}\big) = \mathbf{T}(\mathbf{x}, \mathbf{u}, \mathbf{x}^+) \in [0, 1]$
- Reward functions of the players
  $\mathbf{R} = \{R^1, R^2\}$ $\quad R^i : \mathbf{X} \times \mathbf{U} \to \mathbb{R}$

- **A (stochastic) policy of player *i*** $\quad \pi^i : \mathbf{x} \to u^i \quad \mathbb{P}\big(u^i | \mathbf{x}\big) = \pi^i(\mathbf{x}, u^i) \in [0, 1]$

- **A level-*k* policy of player *i* optimally responds to level-(*k*-1) policy of the other**

$$\pi^{1,k} \in \arg\max_{\pi^1} \mathbb{E}\left\{ \sum_{t=0}^{N} R^1(\mathbf{x}_t, u_t^1, u_t^2) \,\middle|\, u_t^1 \sim \pi^1(\mathbf{x}_t, \cdot), u_t^2 \sim \pi^{2,k-1}(\mathbf{x}_t, \cdot) \right\}$$

Game theoretic traffic modeling results

Level-1 vs Level-0

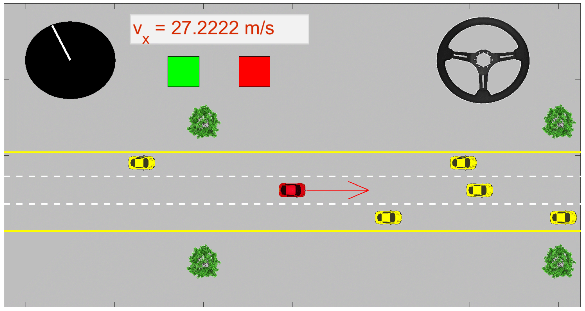

Level-2 vs Level-1

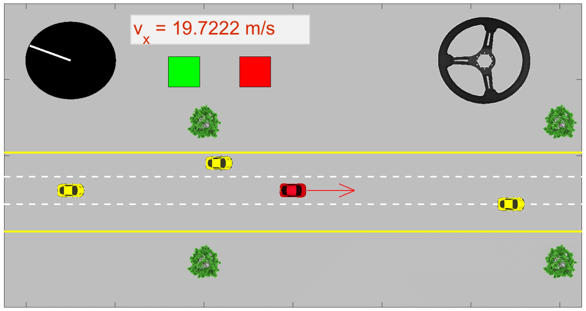

Mixed traffic (10% Level-0, 60% Level-1 and 30% Level-2)

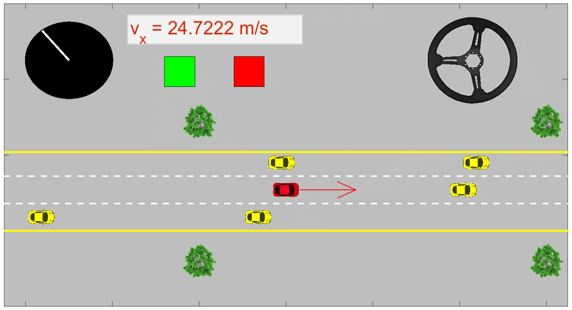

Is level-k a good model for humans?

- Level-k theory has been shown by numerous experimental results from cognitive and behavioral science to predict human decisions better than conventional analytic methods (such as equilibrium-based models) in a range of games.
- We have also developed our own experiment to validate its effectiveness [4].

2020 American Control Conference
Denver, CO, USA, July 1-3, 2020

Beating humans in a penny-matching game by leveraging cognitive hierarchy theory and Bayesian learning

Ran Tian, Nan Li, Ilya Kolmanovsky, and Anouck Girard

Multiple-Level Thinking in Poker: At What Level Are You?

October 06, 2014 | Matthew Pitt

The levels of thought in poker.

Working out what your opponent may be holding is not easy, and it re
"Multiple Level Thinking" is a concept that was brought forward by Da
defines the different levels of thought that a poker player can occupy:

- **Level 0:** No thinking.
- **Level 1:** What do I have?
- **Level 2:** What do they have?
- **Level 3:** What do they think I have?
- **Level 4:** What do they think I think they have?
- **Level 5:** What do they think I think they think I have?

I think I should probably leave it there now because of the fact that it
addition to the fact that it is getting pretty difficult for me to even write
you can see, you can think on different levels whilst playing poker, wi
the more advanced players of the game playing at 4 or above.

Is level-k a good model for humans?

Level-k Theory Application to a Penny-Matching Game [4]

- An AI algorithm was developed to play a repeated penny-matching game.
- The AI player models the human player as a level-*k* decision-maker with online identification of her *k* value.
- Based on the identification result, the AI player maximizes its own chance of winning.

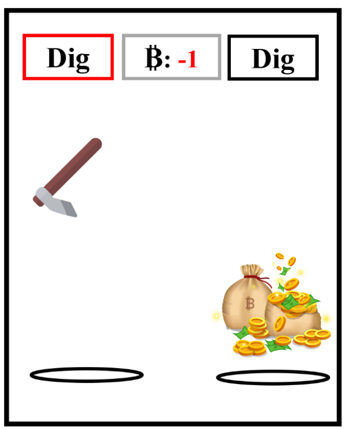

AI won, Human lost

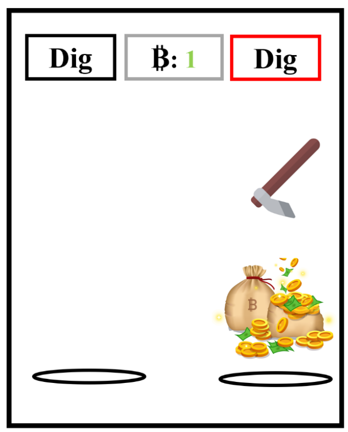

AI lost, Human won

AI vs real humans

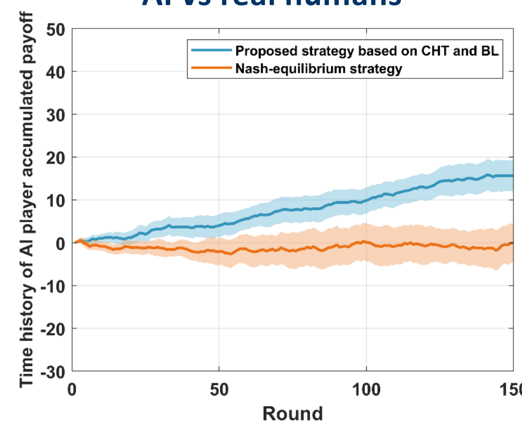

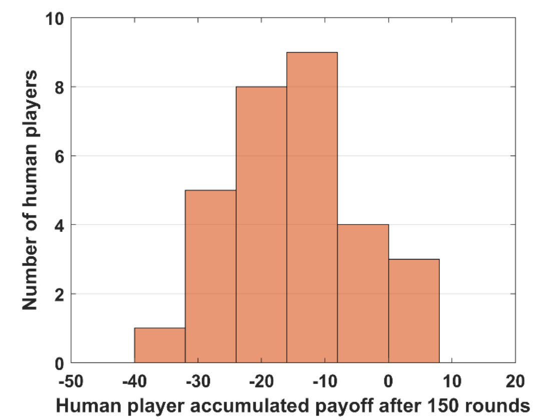

Is level-k a good model for driving?

Naturalistic data analysis – Highlights (Collaborators -- Shan Bao, Huayi Li)

- We studied naturalistic data from the Safety Pilot Model Deployment (SPMD) program led by UMTRI. Valid driving data come from 79 vehicles equipped with Data Acquisition System and Mobileye.
- Our analysis shows that, out of more than 0.3 million miles of highway driving episodes in car-following scenarios, **all vehicles have exhibited level-0, level-1 and level-2 characteristics** based on previous definitions adopted by the traffic model, with a ratio of 11%, 46%, 43% respectively, which are **close to our traffic model assumptions**.

Is level-k a good model for driving?

Naturalistic data analysis – Level-0

- Level-0 characteristic filter: In each event, the driver travels in the **same lane** with the speed between 55mph and 75mph (i.e., no speeding) and with traffic around.
- **All** 79 valid SPMD vehicle drivers exhibited Level-0 characteristics. Average level-0 ratio is 0.11. This means about **11%** of the highway driving events show level-0 characteristic, which is close to the assumption of the simulator.
- Distribution of *number of events vs. event duration*:

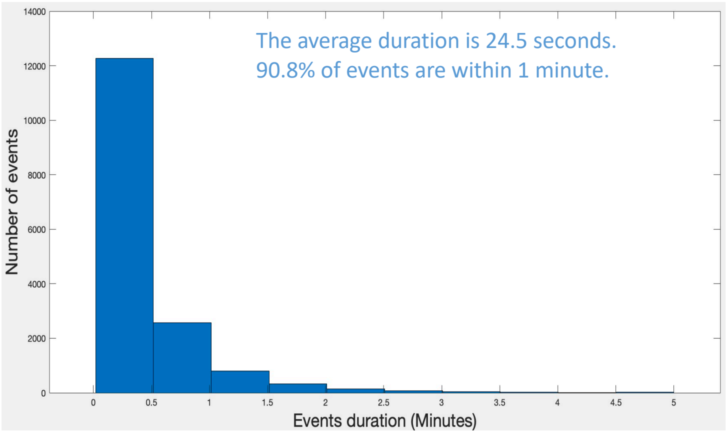

Is level-k a good model for driving?

Naturalistic data analysis – Level-0

- An example of level-0 driving:
  - Driver remains in one lane without lane changing or speeding for approximately 19.2 minutes with an average speed of 71 mph, which is the longest event we have observed.

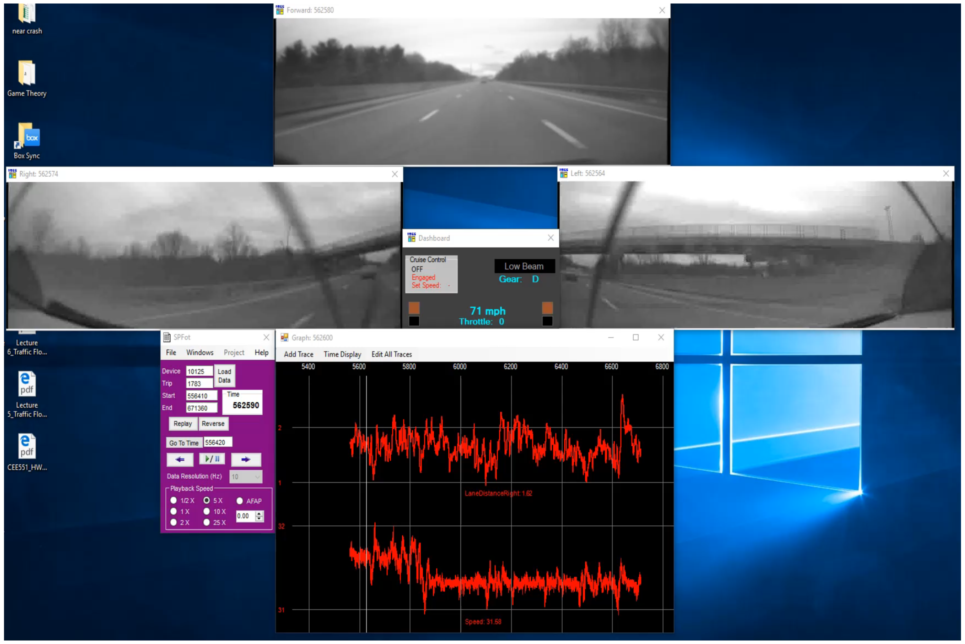

Is level-k a good model for driving?

Naturalistic data analysis – Level-1 and level-2

– Level-1 and level-2 characteristic filter: the driver travels with a minimum speed of 55mph with traffic around, which is similar to the level-0 filter except that **lane change and speeding behavior were observed**.

– We then apply the **unsupervised learning** algorithm, **K-means ++ clustering**, to categorize the events into two groups, and match them to level-1 and level-2 behaviors based on traffic model assumptions.

– *Clustering by the K-means ++ method →*

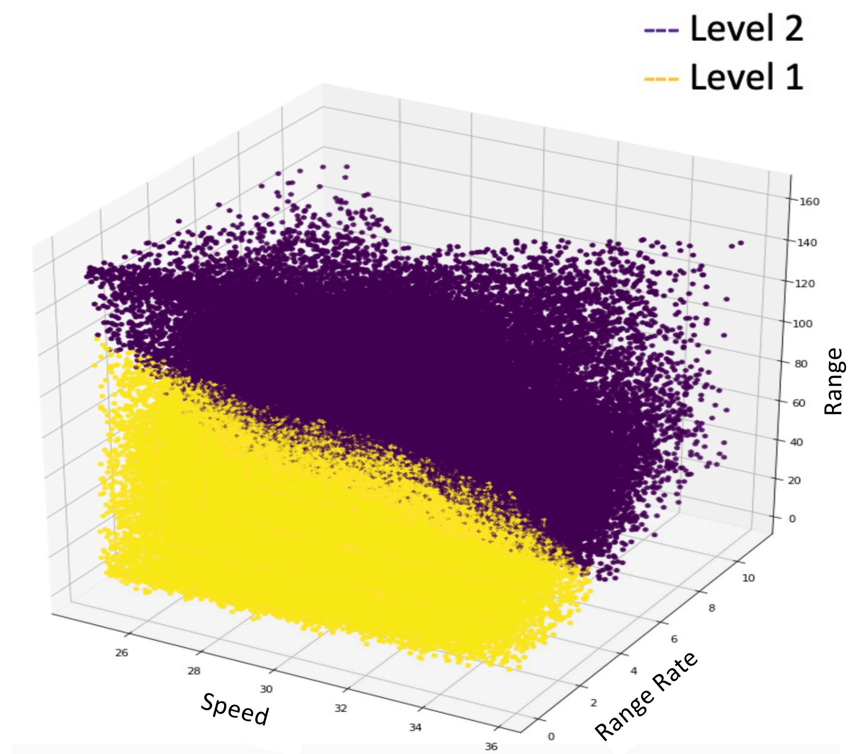

Is level-k a good model for driving?

Naturalistic data analysis – Summary

- The average ratios of level-0, level-1 and level-2 are **11%, 46% and 43%** based on total distance traveled on the highway.
- Examples of behavioral differences:
  - Level-0 is the most defensive, keeping the largest distance from the vehicle in front without speeding or lane changes.
  - Level-1 is the most aggressive, keeping the smallest distance from the vehicle in front and the largest ratio of speeding.
  - The characteristic of level-2 is between level-1 and level-0. Compared with level-1, it has lower number of lane changes and ratio of speeding.

	Event Numbers	Avg. Duration (sec.)	Avg. Range (m)	Avg. Range Rate (m/s)	Avg. Speed (m/s)	Avg. Acceleration (m/s^2)	Speeding Ratio	Lane Changes per 100 Miles
Level-0	91546	24.52	71.72	0.15	29.83	-0.18	0	0
Level-1	61034	143.40	49.65	-0.56	32.96	-0.16	0.45	143.88
Level-2	144708	60.22	60.74	-0.48	31.62	-0.17	0.29	111.26

Interactive traffic simulator

Conventional simulator:

- Ego AV is the only intelligent decision-maker.
- Motions of the other vehicles prescribed as functions of time or responding to traffic state using simple rules.

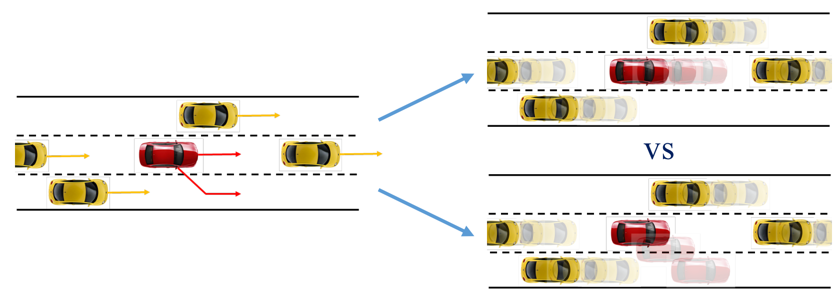

Interactive simulator:

- Ego and other vehicles are all intelligent decision-makers and interact with each other.
- When ego AV takes different actions, the other vehicles will respond differently.

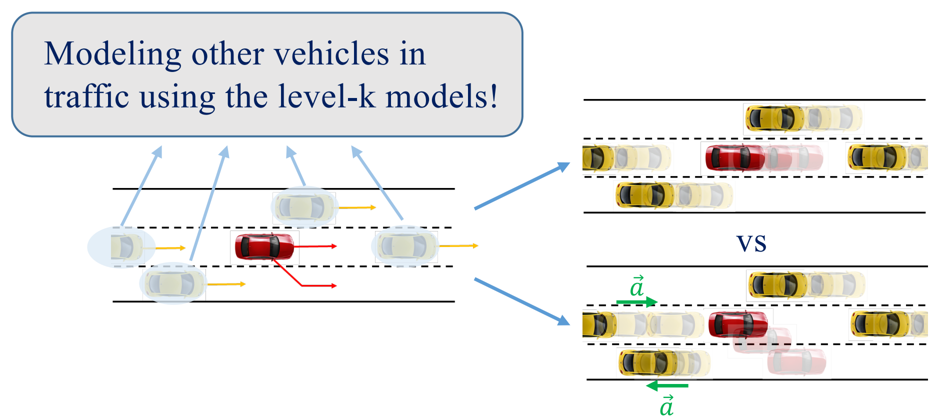

Integration with high-fidelity simulator (TORCS)

Interactive Traffic Simulator [5]

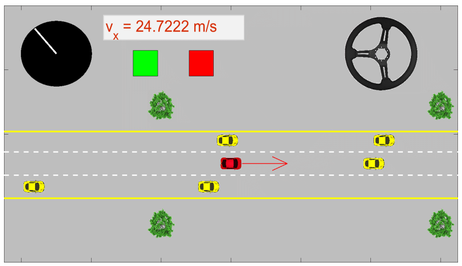

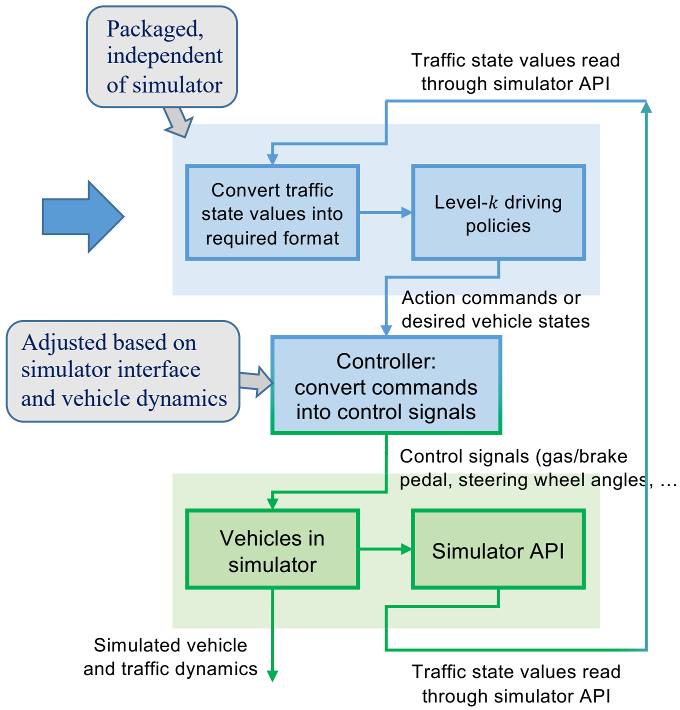

Integration of level-k models and high-fidelity driving simulator

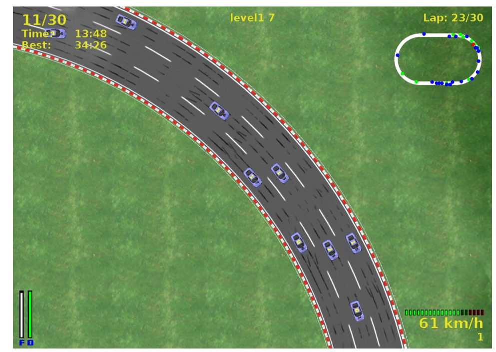

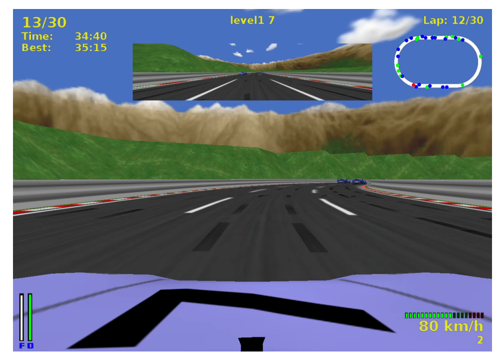

[5] Su, Li, Yildiz, Girard and Kolmanovsky (CPHS 2019).

Integration with high-fidelity simulator (TORCS)

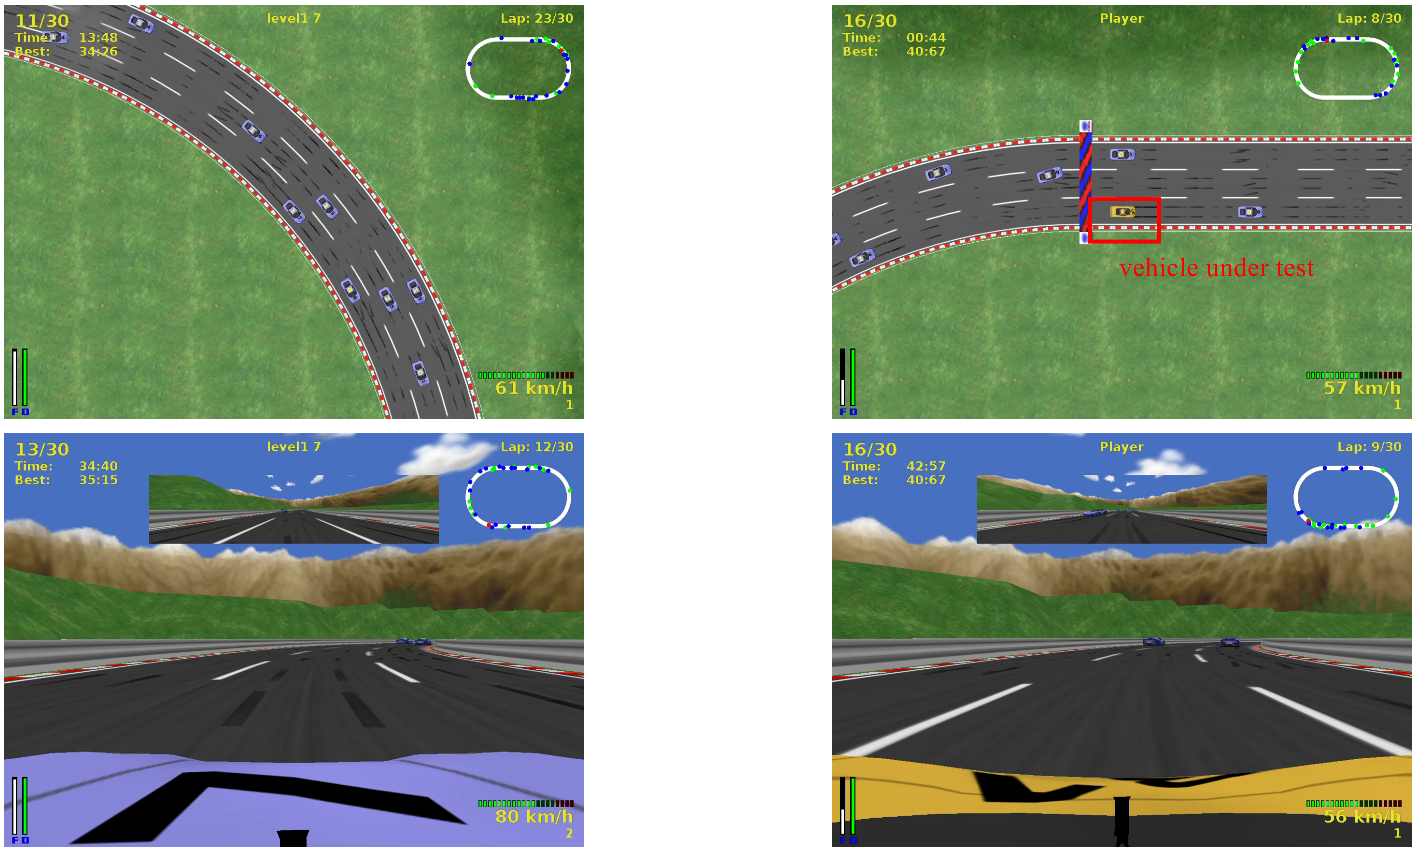

https://cps-vo.org/group/verification_tools/tool/GTTS

Un-signalized Intersections

Four-way intersection

Roundabout

- Represent roughly **50% of intersection crashes** in the US
- Require a balance between overly aggressive actions (that do not take into account behavior of others) and overly conservative actions (deadlock)
- Challenges: predict other drivers' actions (interactive), take optimal decision corresponding to the prediction

Un-signalized intersections: Scenarios

2-vehicle interactions at:

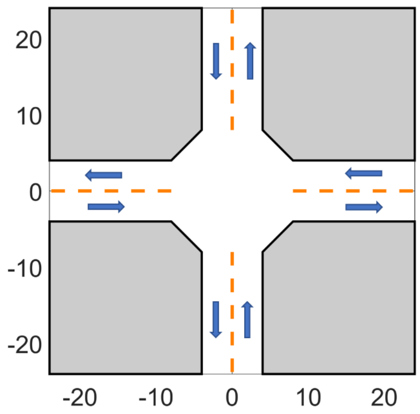

Unsignalized four-way intersection

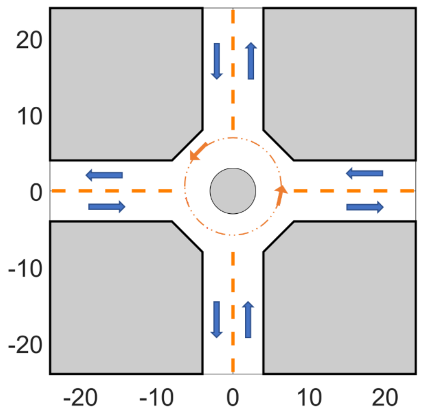

Roundabout

MDP for intersections

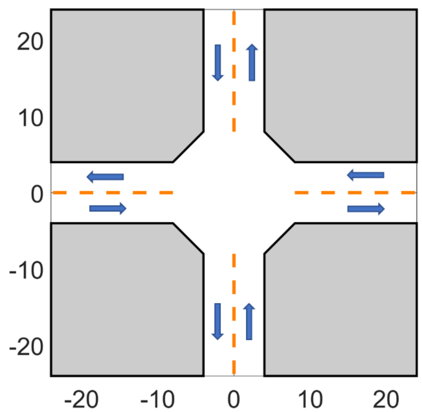

- **Set of states (MDP, not POMDP):**
  - 4 states per vehicle, longitudinal and lateral positions, speed, yaw angle
  - 2-vehicle interactions at first
- **Set of actions: 6 possible combinations**
  - Acceleration
  - Yaw rate
- **Reward/Penalty:** (safe distance) constraint violation, vehicle speed, maneuver effort, headway

$$R = w_1 R_1 + w_2 R_2 + w_3 R_3 + w_4 R_4 + w_5 R_5$$

R_1 penalty for collision, R_2: penalty for being too close, R_3: penalty for driving off road,

R_4 penalty for driving in wrong lane, R_5 reward for approaching objective lane

Un-signalized intersections: Results[2]

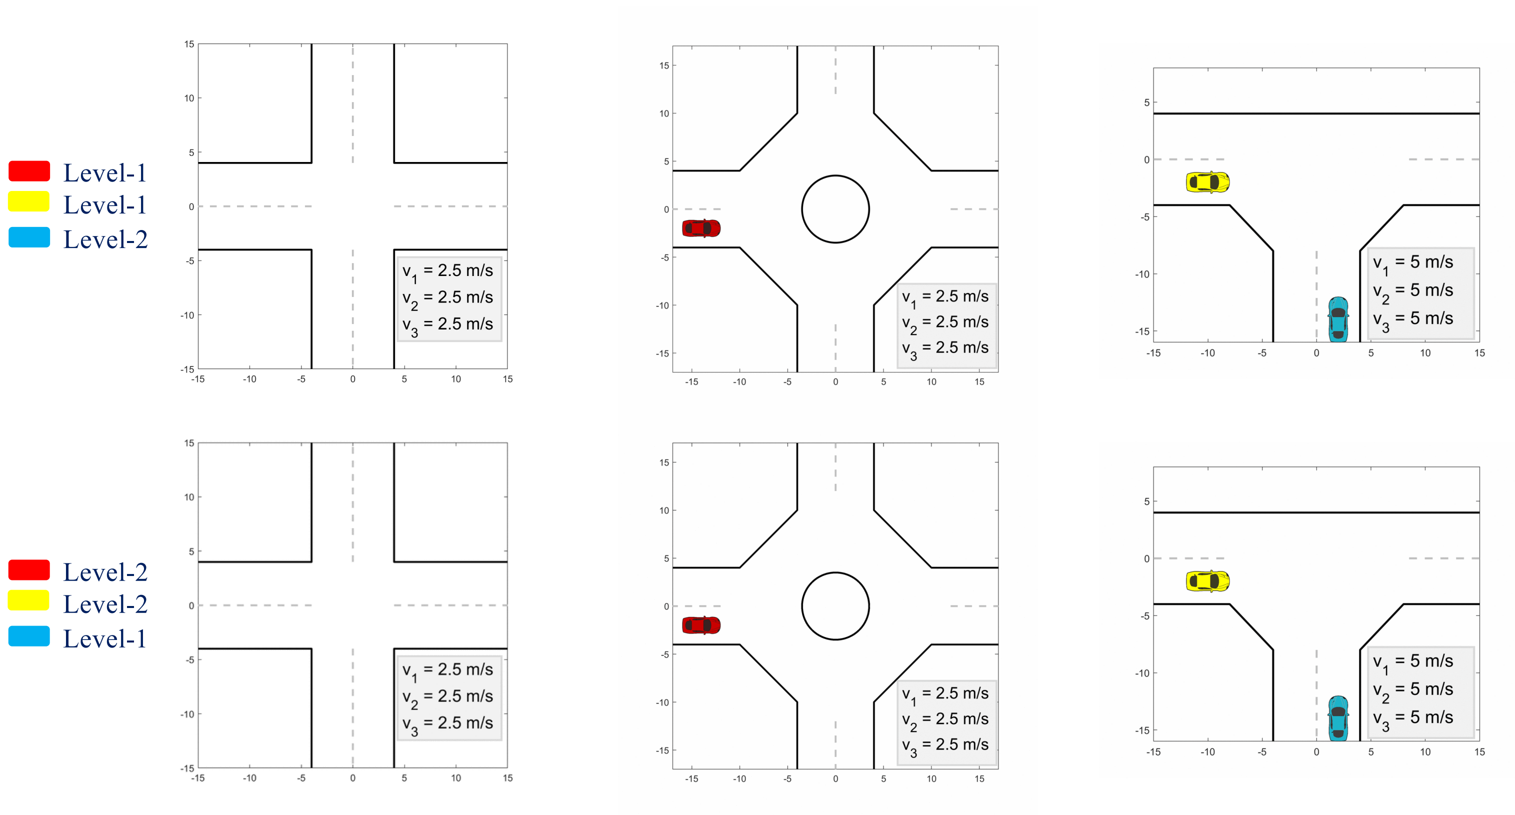

L0 driver treats the other car as a stationary obstacle and does not take her potential actions into account → aggressive

Un-signalized intersections: Larger network[2]

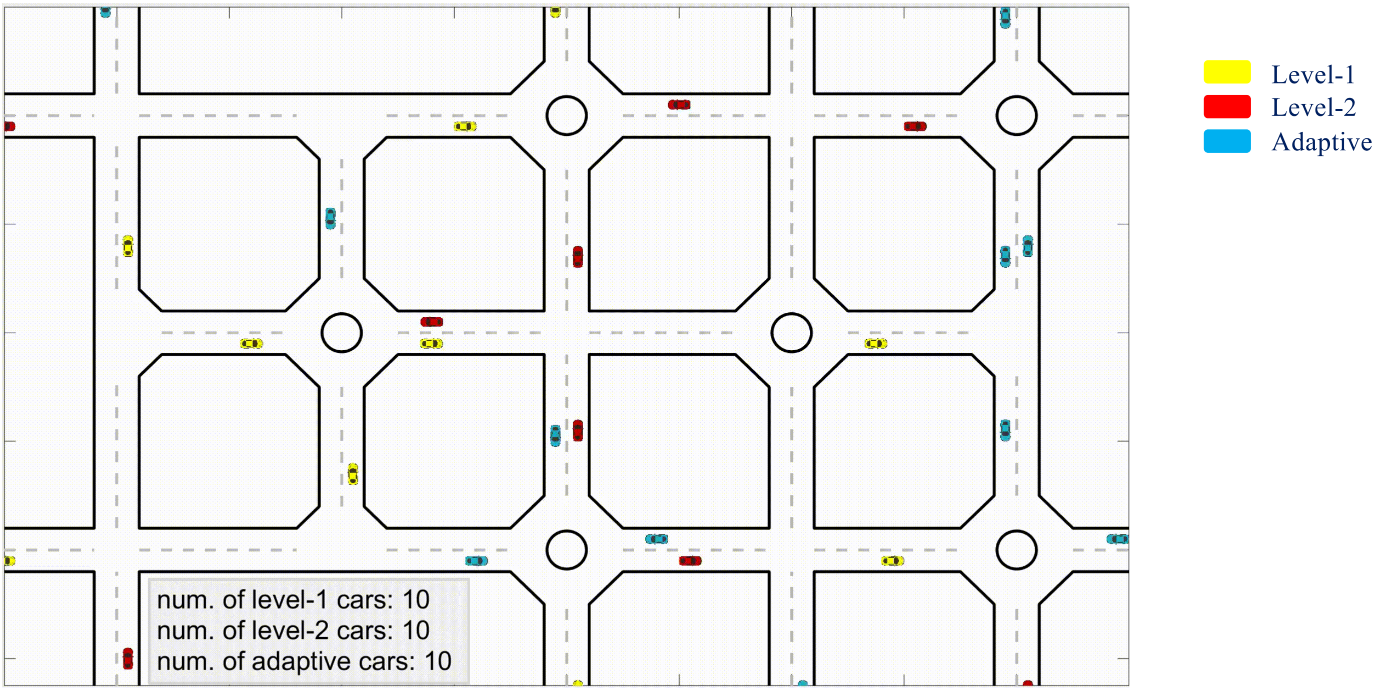

Un-signalized intersections: Validation[2]

Validation against Traffic Data [2]

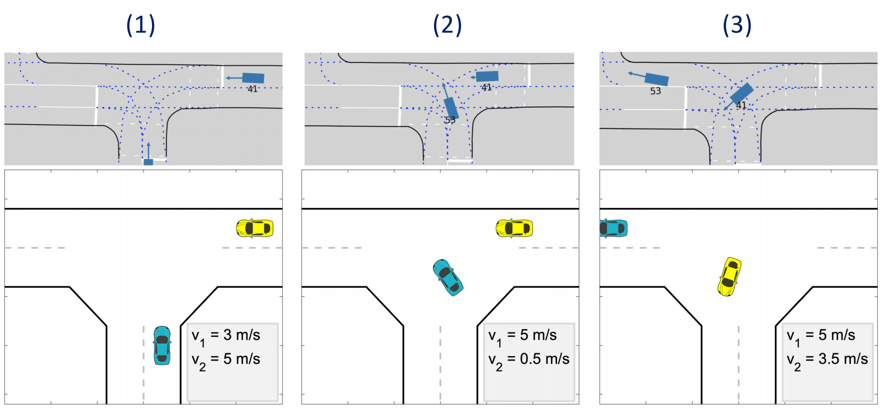

Top: A traffic scenario extracted from the INTERACTION dataset
Bottom: Simulation snapshots of a level-2 vehicle (blue) interacting with a level-1 vehicle (yellow)

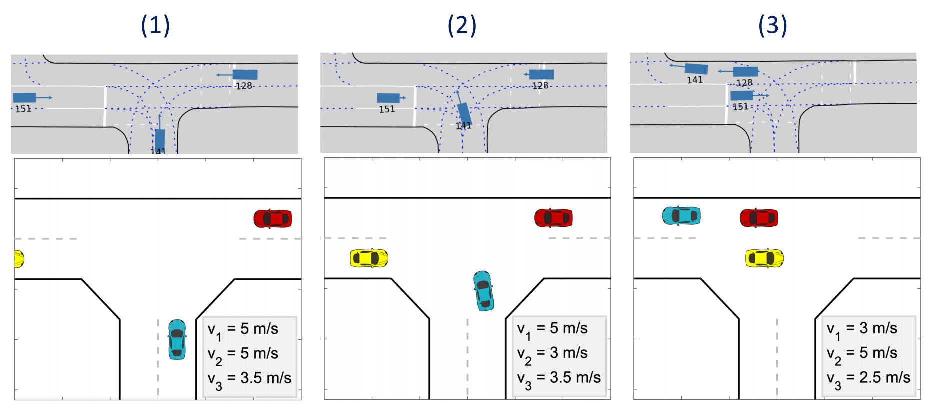

Top: A traffic scenario extracted from the INTERACTION dataset
Bottom: Simulation snapshots of a level-2 vehicle (blue) interacting with two level-1 vehicles (yellow and red)

Generalization to n–vehicle scenarios

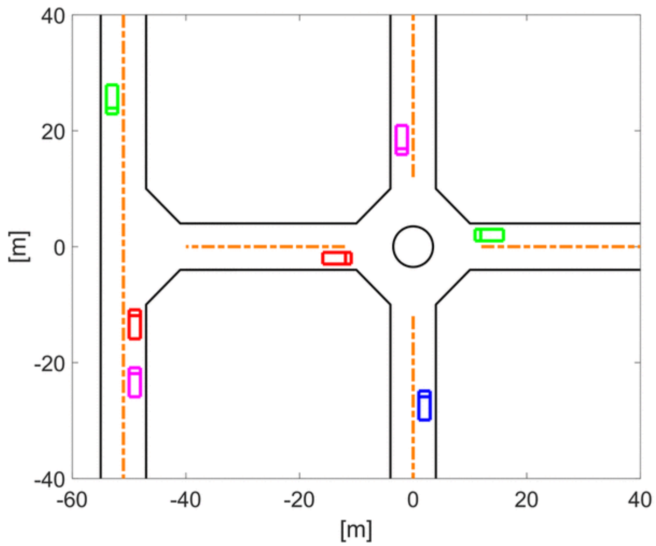

- **POMDP:**
  - The ego car considers its interactions with its two nearest neighbors

Generalization to arbitrary road geometries

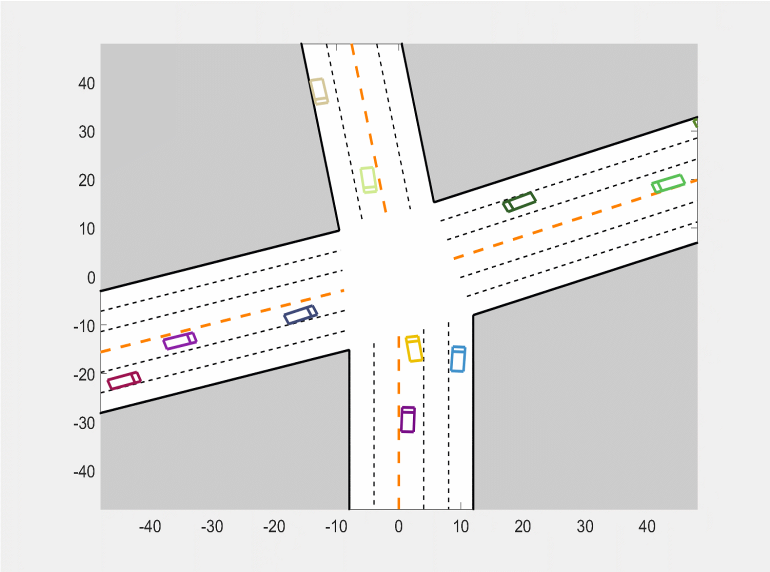

Using the simulator for control design

Interaction-Aware Decision-Making for AVs [6]

Game Theory
Generate a set of interactive driver behavior models, corresponding to different driving styles (aggressive/conservative) and driving intentions (proceeding/yielding).

Bayesian Inference
Identify the model(s) that most accurately characterizes the present interaction scenario based on observed vehicle trajectories, with corresponding identification confidence.

Stochastic Model Predictive Control
Make decisions for ego AV, accounting for driving objectives (crossing, merging, …) and providing probabilistic collision-avoidance guarantee.

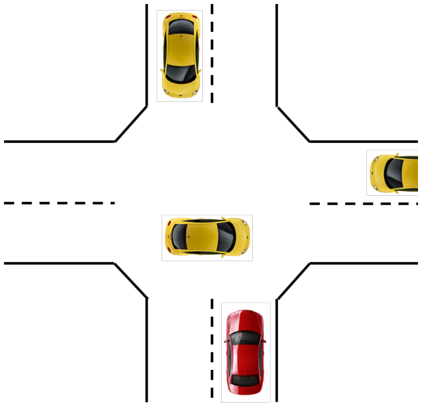

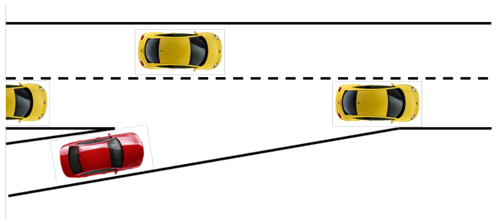

[6] Li, Li, Girard and Kolmanovsky (CDC 2019).

Using the simulator for control design

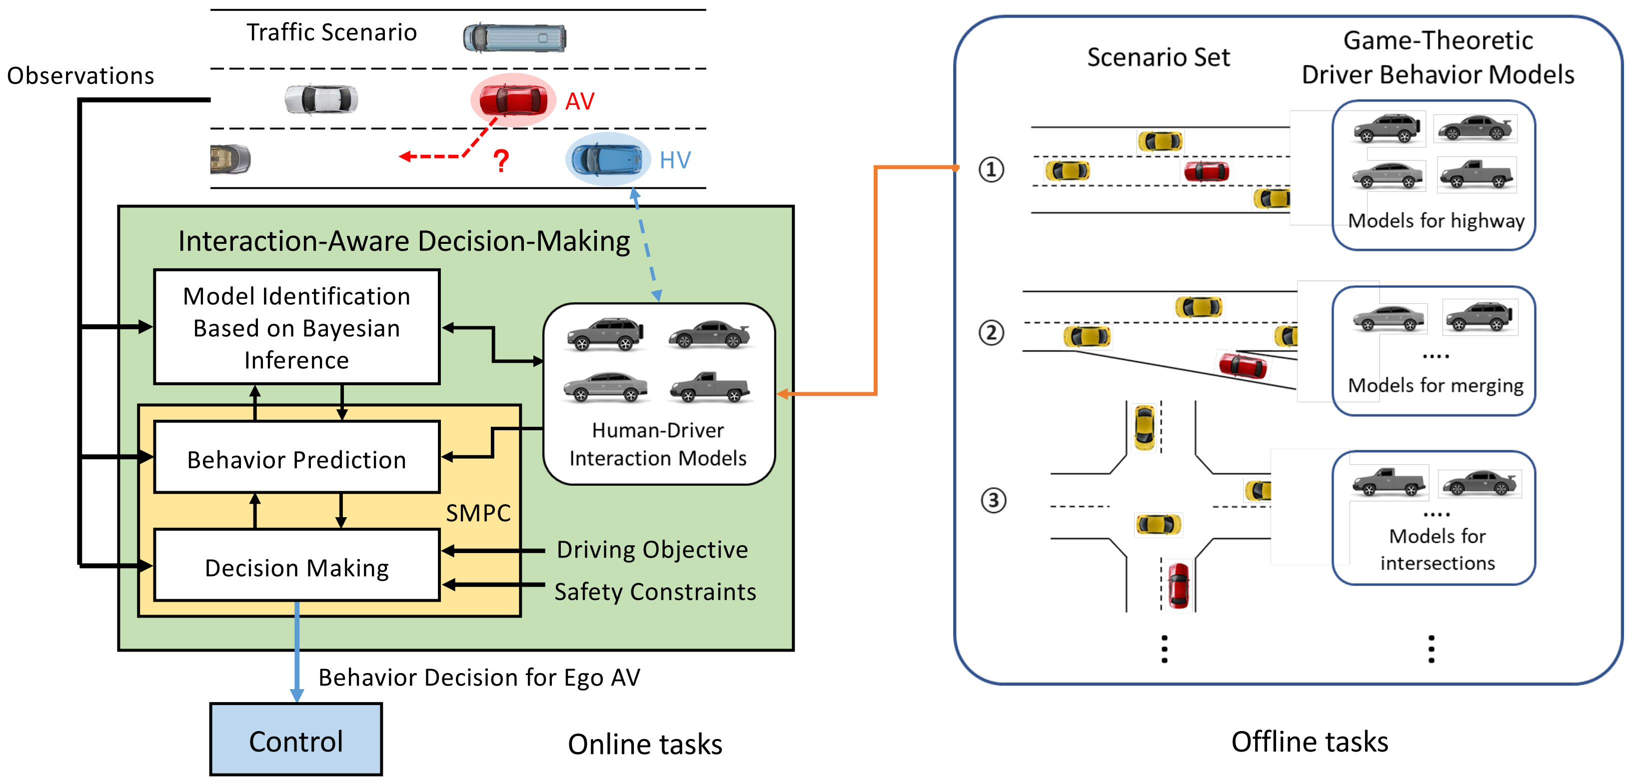

Using the simulator for control design

Interaction-Aware Decision-Making for AVs [6]

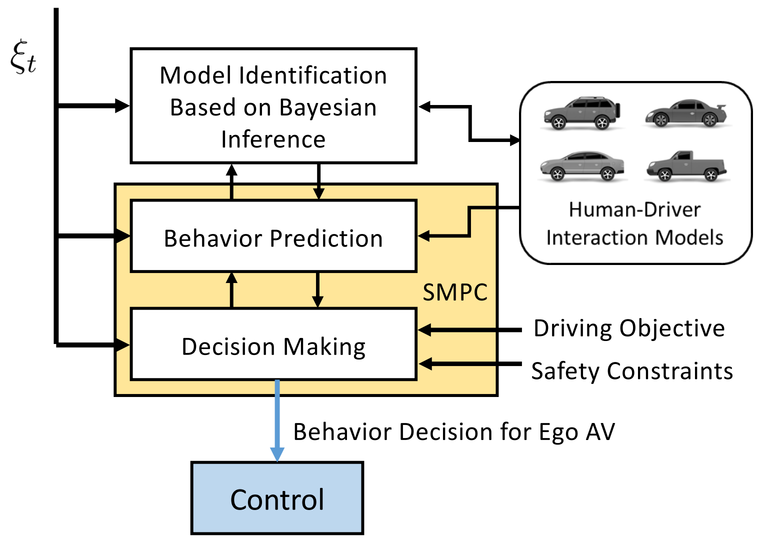

Augmented state $\overline{\mathbf{x}} = (\mathbf{x}, \sigma)$

Physical state (observable) Model Id (hidden)

Model predictions $\mathbb{P}\big(u_t^2 = u | \mathbf{x}_t, \sigma\big) = \pi^{2,\sigma}(\mathbf{x}_t, u)$

Observations $\xi_t = \big\{\mathbf{x}_0, \dots, \mathbf{x}_t, u_0^2, \dots, u_{t-1}^2\big\}$

Bayesian inference formula

Identification confidence (belief) $\mathbf{b}_t : \mathbb{P}\big(\sigma | \xi_t\big)$

SMPC

$$\max_{u^1_{\tau|t} \in U^1,\, \tau=0,\dots,N} \quad \mathbb{E}\left\{ \sum_{\tau=0}^{N} R^1(\mathbf{x}_{\tau|t}, u^1_{\tau|t}, u^2_{\tau|t}) \,\Big|\, \mathbf{b}_t \right\}$$

$$\text{s.t.} \quad \mathbb{P}\big(\mathbf{x}_{\tau+1|t} \in \mathbf{X}_{\text{safe}},\ \tau = 0, \dots, N \,\big|\, \mathbf{b}_t\big) \geq 1 - \varepsilon$$

Interaction-Aware Decision Making for AVs

- The AV does not know the driving styles/intentions of the two interacting vehicles a priori.
- Same initial condition, different interaction behavior!

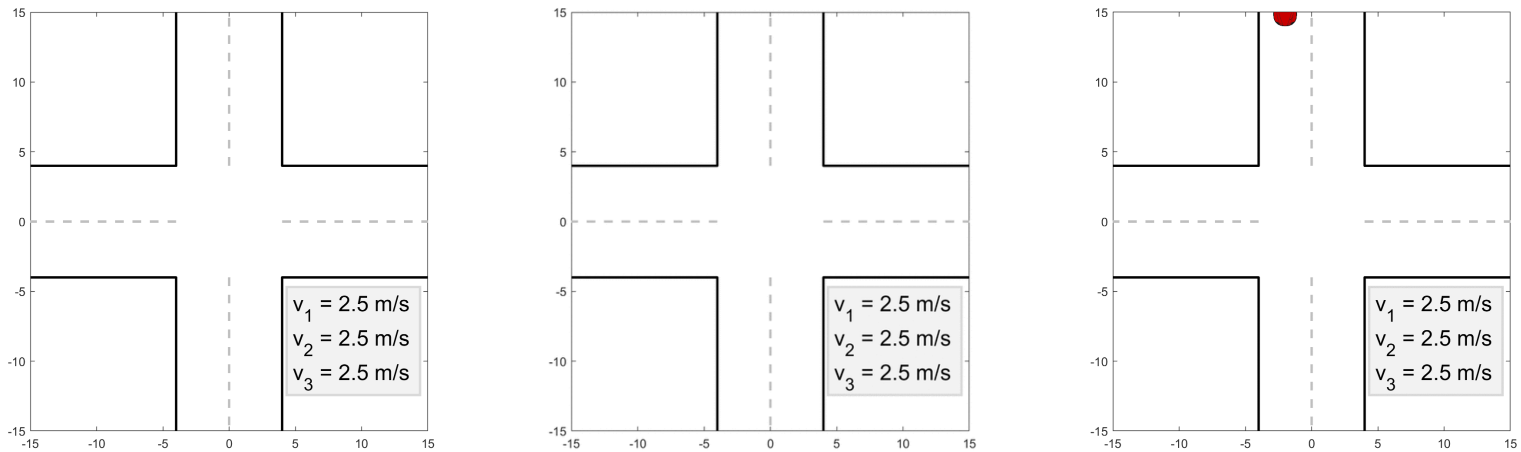

Level-1 (conservative)
Level-1 (conservative)
AV

Level-2 (aggressive)
Level-2 (aggressive)
AV

Level-1 (conservative)
Level-2 (aggressive)
AV

Interaction-Aware Decision Making for AVs

- The AV does not know the driving styles/intentions of the two interacting vehicles a priori.
- Same initial condition, different interaction behavior!

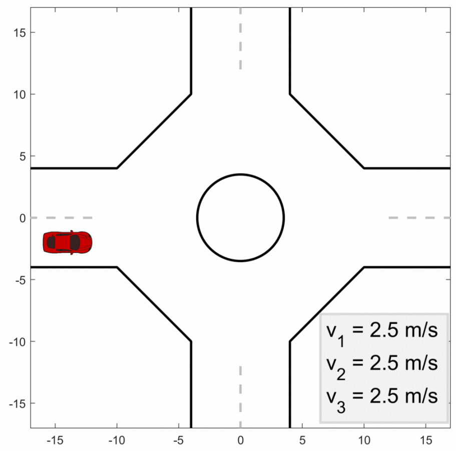

Level-1 (conservative)
Level-1 (conservative)
AV

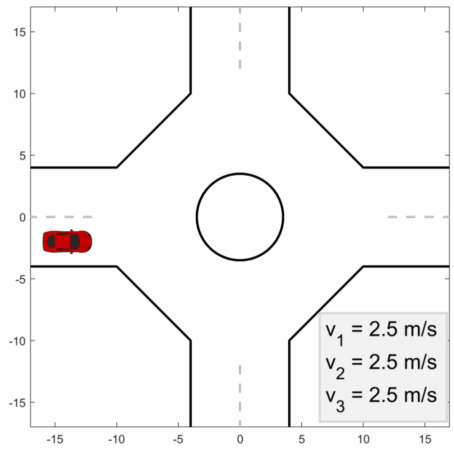

Level-2 (aggressive)
Level-2 (aggressive)
AV

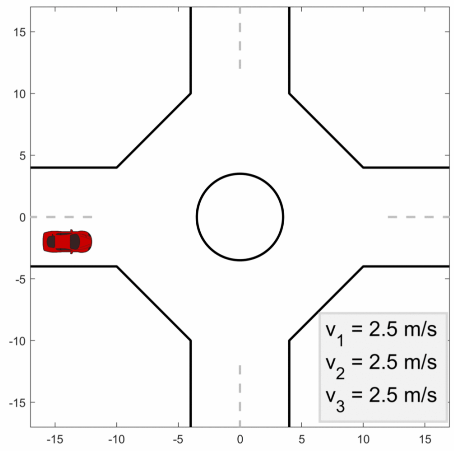

Level-1 (conservative)
Level-2 (aggressive)
AV

Interaction-Aware Decision Making for AVs

- The AV does not know the driving styles/intentions of the two interacting vehicles a priori.
- Same initial condition, different interaction behavior!

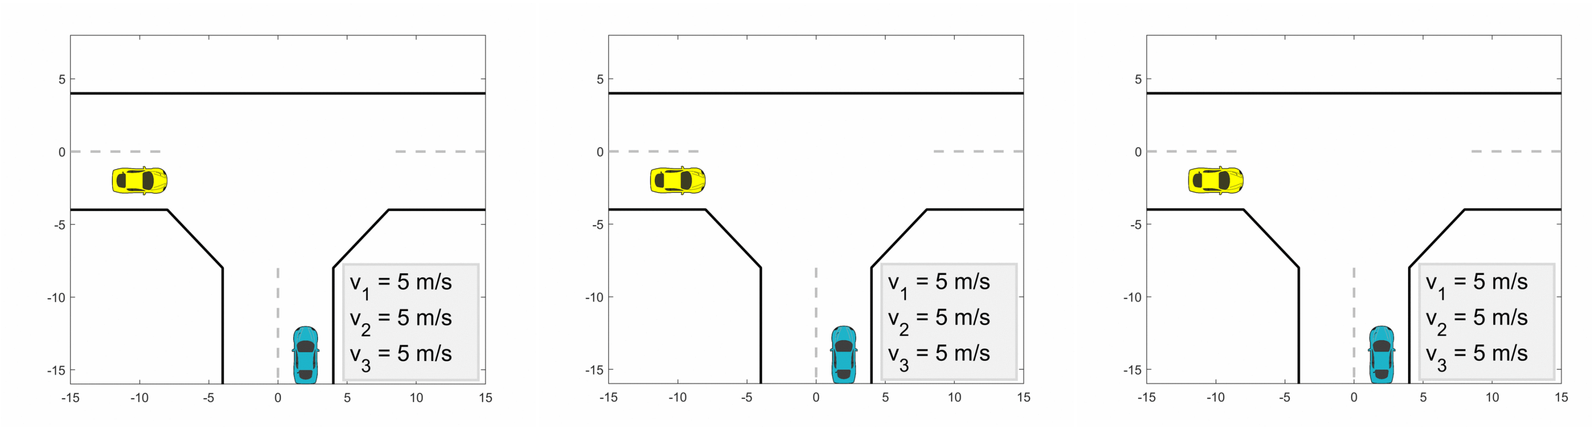

Level-1 (conservative)
Level-1 (conservative)
AV

Level-2 (aggressive)
Level-2 (aggressive)
AV

Level-1 (conservative)
Level-2 (aggressive)
AV

Interaction-Aware Decision Making for AVs

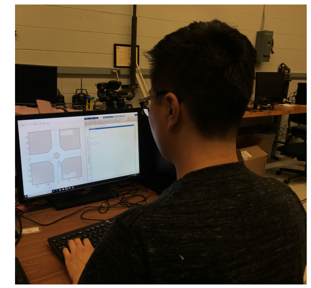

- The AV does not know the driving style/intention of the interacting vehicle a priori.
- Same initial condition, different interaction behavior!

Roundabout scenario [7]:

We let a human operator control the red car using the keyboard.

AV vs Level-1

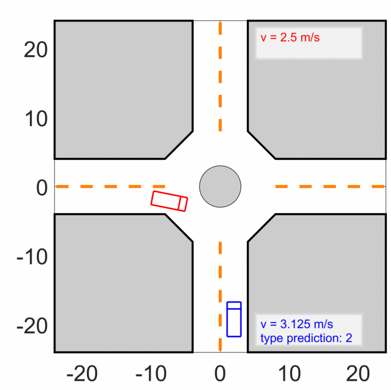

AV vs Level-2

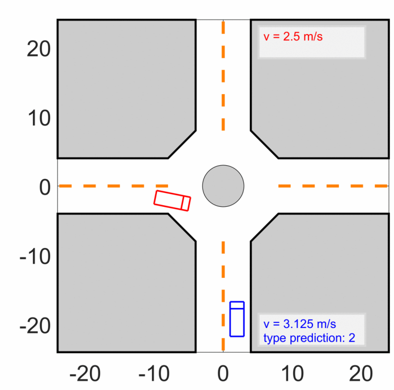

AV vs Human 1

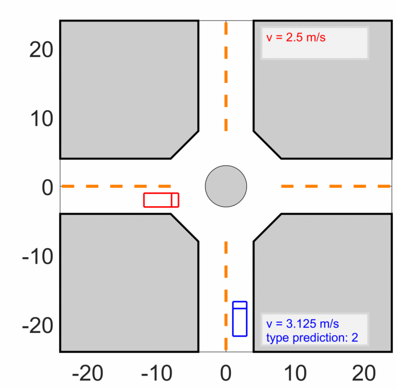

AV vs Human 2

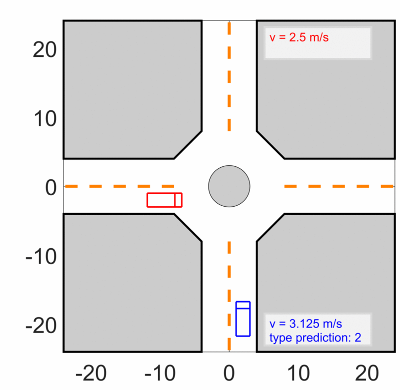

[7] Tian, Li, Li, Kolmanovsky, Girard and Yildiz (CDC 2018).

Concluding Remarks

1. Autonomous vehicles (AVs) will likely operate in traffic together with human-driven vehicles (HVs). Their control algorithms must account for the interactions between AVs and HVs.
2. Game theory is a useful tool for modeling such interactions.
3. The generated interaction models have various applications, e.g., for AV control system testing and for interaction-aware AV control.
4. It is possible to apply the proposed game-theoretic approaches to modeling interactions of other traffic participants, e.g., pedestrians, cyclists, etc.

Backup

Traffic Rules for Intersections

Incorporating "Right of Way" Rules through a Leader-follower/Stackelberg-like Formulation

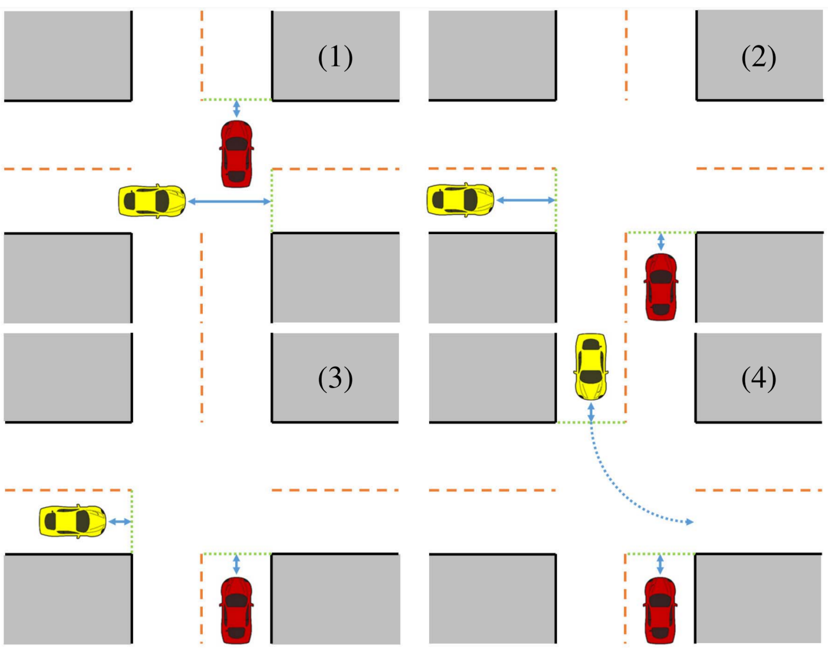

(1) If vehicles i, j have both entered the intersection, the vehicle with a strictly smaller signed distance to the exit of the intersection is the leader.

(2) If at most one of vehicles i, j has entered the intersection, the vehicle with a strictly smaller signed distance to the entrance of the intersection is the leader.

(3) If no leader has been assigned to (i, j) according to (1)-(2), then the vehicle on the right is the leader when the two vehicles are coming from adjacent road arms.

(4) If no leader has been assigned to (i, j) according to (1)-(3), then the vehicle going straight is the leader when the other vehicle is making a turn.

Algorithm 1 Leader-follower role assignment

Input : an ordered pair of vehicles (i, j) and their states $(s_i(t), s_j(t))$

Output: whether i is the leader of j

if $(\Delta\rho_i^{en}(t) \leq 0$ *and* $\Delta\rho_j^{en}(t) \leq 0)$ *and* $\Delta\rho_i^{ex}(t) < \Delta\rho_j^{ex}(t) - \delta$ **then** $i \prec j$;

else if $(\Delta\rho_i^{en}(t) > 0$ *or* $\Delta\rho_j^{en}(t) > 0)$ *and* $\Delta\rho_i^{en}(t) < \Delta\rho_j^{en}(t) - \delta$ **then** $i \prec j$;

else if *i and j are coming from adjacent ways and i's way is on the right of j's way* **then** $i \prec j$;

else if *i is going straight and j is making a turn* **then** $i \prec j$;

else $i \succeq j$.

Affordance indicators/Automated driving FSM

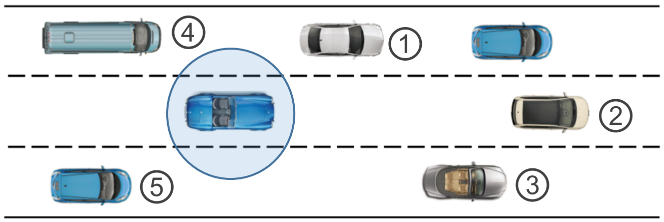

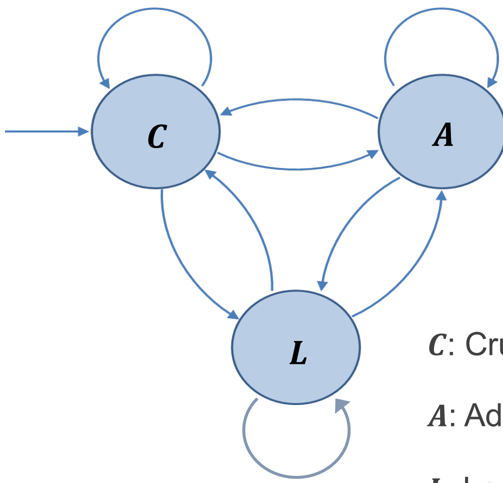

C: Cruise control mode

A: Adaptive cruise control mode

L: Lane change mode

Affordance indicators:

The range to the preceding car in the left lane, and the corresponding range rate,
The range to the preceding car in the current lane, and the corresponding range rate,
The range to the preceding car in the right lane, and the corresponding range rate,
The range to the car following behind in the left lane, and the corresponding range rate,
The range to the car following behind in the right lane, and the corresponding range rate.

Testing and calibration in UM traffic simulator

Red car = FSM controller, Yellow cars = level-k game theory

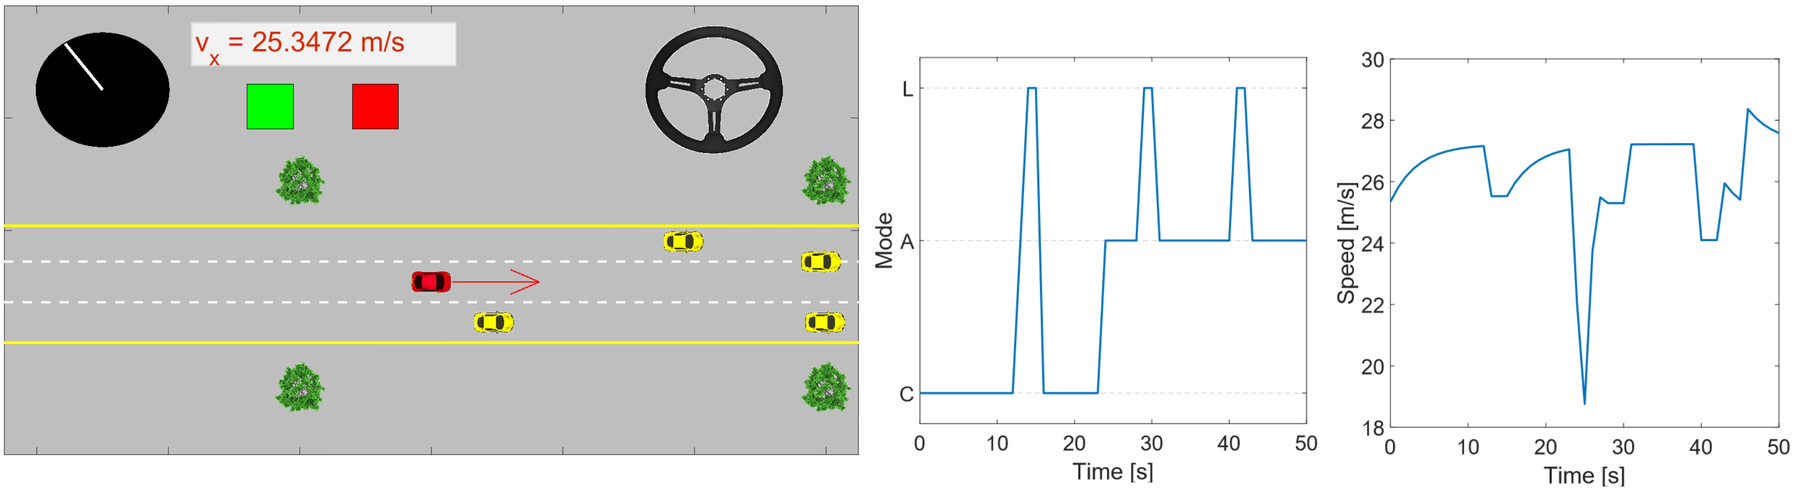

Testing and calibration in UM traffic simulator

Finite State Machine (rule-based)

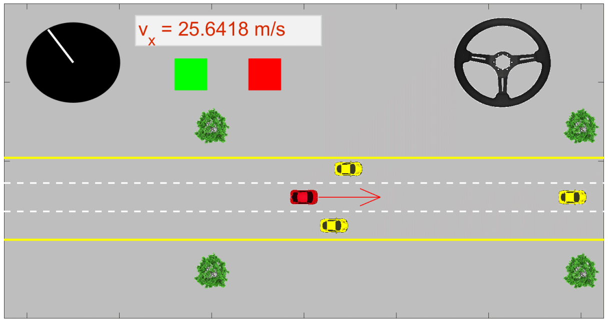

Human Driver

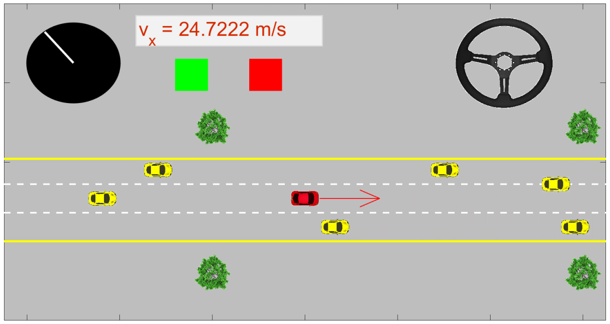

Zhang, M., Li, N., Girard, A., & Kolmanovsky, I. (2017, October). A Finite State Machine Based Automated Driving Controller and its Stochastic Optimization. In *ASME 2017 Dynamic Systems and Control Conference* (pp. V002T07A002-V002T07A002). American Society of Mechanical Engineers.

Stochastic optimization of FSM parameters

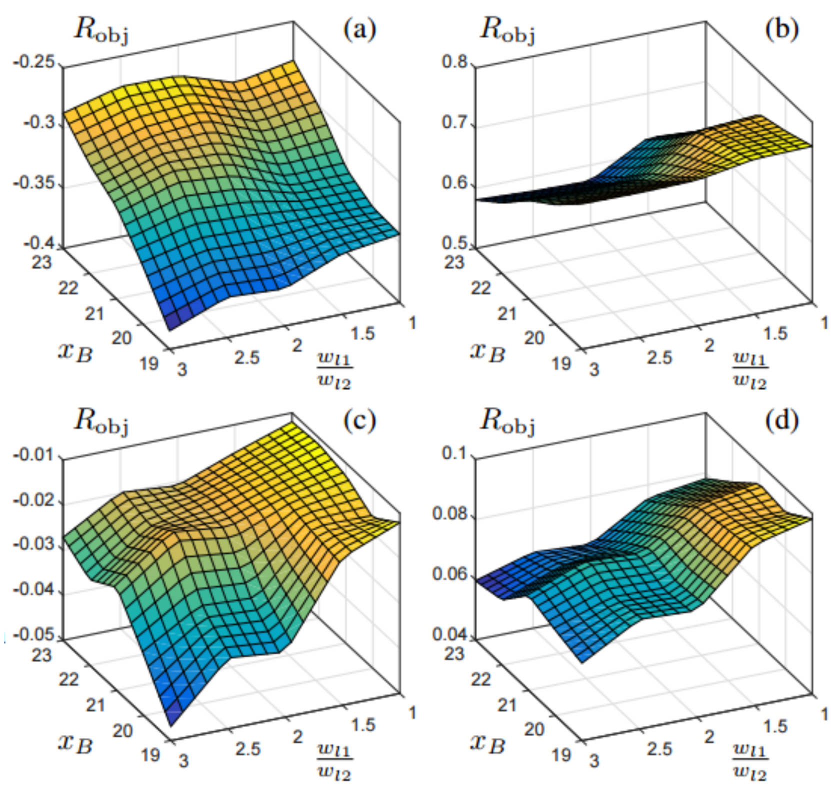

$$R = \gamma_1(-c) + \gamma_2 \frac{\bar{v}_x - v_{min}}{v_{max} - v_{min}}$$

c is the safety distance constraint violation rate, $\bar{v}_x$ is the average speed.

Different combinations of weights:
(1) $\gamma_1 = 1, \gamma_2 = 0;$
(2) $\gamma_1 = 0, \gamma_2 = 1;$
(3) $\gamma_1 = 0.7, \gamma_2 = 0.3;$
(4) $\gamma_1 = 0.6, \gamma_2 = 0.4.$

The user can pick the best pair of parameters based on their objective function design.

- To obtain the best action sequence, a reward is evaluated after each action. A weighted sum of the 2 levels' rewards is computed to represent how good that action sequence is. $\frac{w_{l1}}{w_{l2}}$ is the weight ratio for the 2 levels.
- x_B is the size of a safe zone in the longitudinal direction. If this safe zone is violated, the host car decelerates.

Bridging Past, Present and Future

- CONSISTENT
- SAFETY CRITICAL
- COMMERCIAL
- COMPETITIVE
- CERTIFIED
- CONGESTED
- CONTESTED

- AUTONOMOUS
- COLLABORATIVE
- INTELLIGENT
- LOW-COST
- RECONFIGURABLE
- MULTI-AGENT
- HUMAN-CENTERED

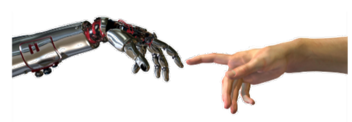

Opportunities for Engineering

- ❑ New players (Space X, Blue Origin, UAV companies, self-driving car technology) create demand for a new engineering curriculum

- ❑ Computing (new forms of computation, parallel, neuromorphic, quantum), AI and data literacy will have to be mainstream knowledge like maths and physics

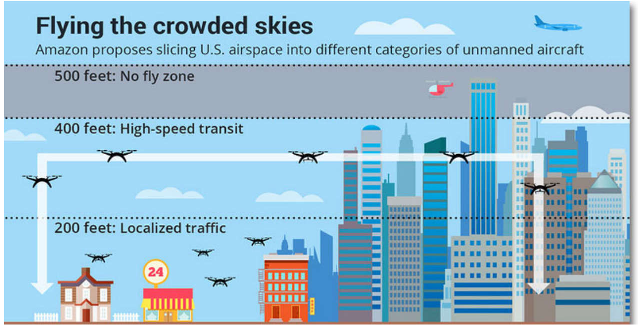

A 4th Industrial Revolution?

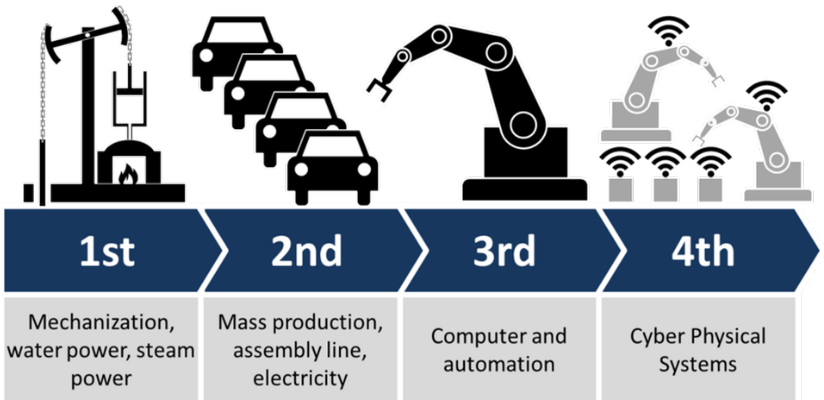

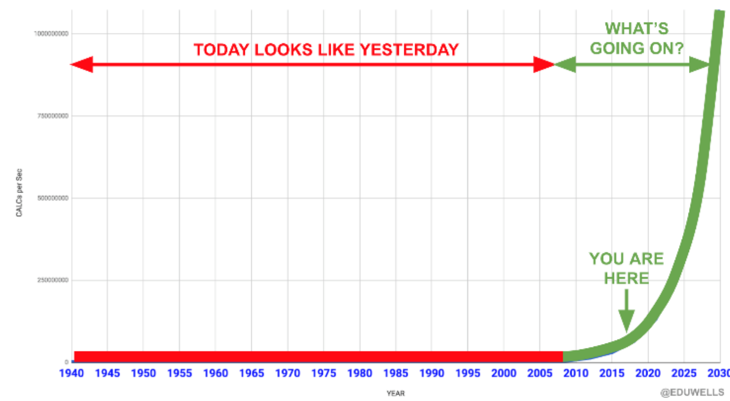

Mathematical rigor vs experimental validation
Deterministic vs stochastic
Physics vs data based

Gap between "cutting edge" and "rigor" in an era of exponential technologies

- Leverage the abundance of computation power
- Adapt/replace **standards** for the future
- Discuss **certification** of new technologies
- Prepare for tradeoffs and new paradigms

About the speaker

Anouck R. Girard received the Ph.D. degree in ocean engineering from the University of California, Berkeley, CA, USA, in 2002. She has been with the University of Michigan, Ann Arbor, MI, USA, since 2006, where she is currently an Associate Professor of Aerospace Engineering.

She has co-authored the book Fundamentals of Aerospace Navigation and Guidance (Cambridge University Press, 2014). Her current research interests include flight dynamics and control systems. Dr. Girard was a recipient of the Silver Shaft Teaching Award from the University of Michigan and a Best Student Paper Award from the American Society of Mechanical Engineers.

Background

- 2006 – Current: Assistant/Associate Professor, Aerospace Engineering, University of Michigan
- 2004 – 2006: Assistant Professor, Mechanical Engineering, Columbia University
- 2002 – 2004: Post-Doctoral Researcher/Lecturer, UC Berkeley
- 1998 – 2002: PhD, UC Berkeley, Karl Hedrick's group

- **Research Interests:** Controlling advanced and increasingly autonomous vehicles and vehicle systems operating in the air, space, ground and marine domains. These vehicles / systems exhibit complex nonlinear dynamics, and must function in uncertain environments with limited resources, while satisfying stringent constraints and counteracting the effects of disturbances.

- **Publications:** 1 textbook (navigation and guidance), 46 journal papers, 130+ conference papers

- **Center Experience:** PI for the MACCCS Center, 2007-2016

- **Teaching Experience:** Undergraduate: Aircraft Performance, Spacecraft Dynamics, Aircraft Dynamics and Control, Classical Control; Graduate: Linear Systems, Nonlinear Systems, Navigation and Guidance, Spacecraft Dynamics and Control, Predictive and Nonlinear Control, Autonomy.

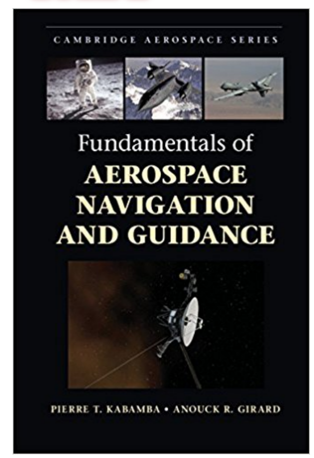

Current/Recent Research

Autonomy and Decision Making	PE games (UAV) 	TSP stability regions (UAVs) 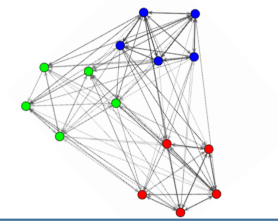	Game theoretic traffic simulation 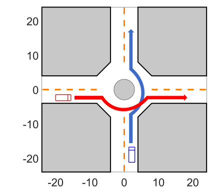
Path Planning, Trajectory Opt.	UAV integration in airspace 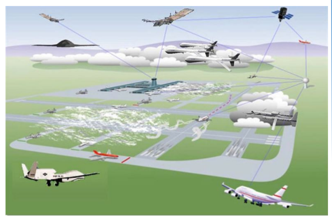	Information collection 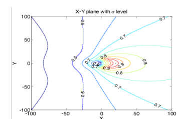 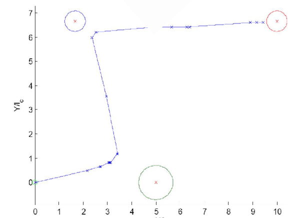	Formation flight 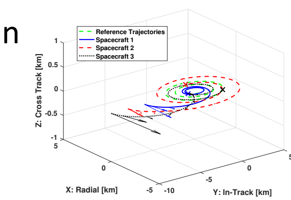
Energy Management	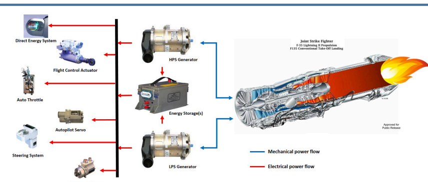 MEA/AEA/ fighters	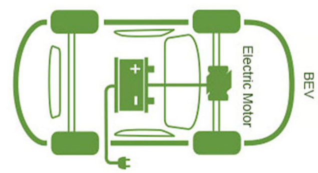 Automotive	Disjunctive sensing/ control 
Unusual Configurations		Gyrocart 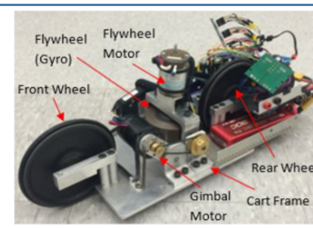	AC