Formal Scenario-Based Testing of Autonomous Vehicles: From Simulation to the Real World

Daniel Fremont (UC Santa Cruz), Sanjit A. Seshia (UC Berkeley)

https://ieeexplore.ieee.org/document/9294368

We present a formal scenario-based testing methodology on the safety of autonomous vehicles, especially those using advanced artificial intelligence-based components, spanning both simulation-based evaluation as well as testing in the real world. Our approach is based on formal methods, combining formal specification of scenarios and safety properties, algorithmic test case generation using formal simulation, test case selection for track testing, executing test cases on the track, and analyzing the resulting data. Experiments with a real autonomous vehicle at an industrial testing facility support our hypotheses that (i) formal simulation can be effective at identifying test cases to run on the track, and (ii) the gap between simulated and real worlds can be systematically evaluated and bridged.

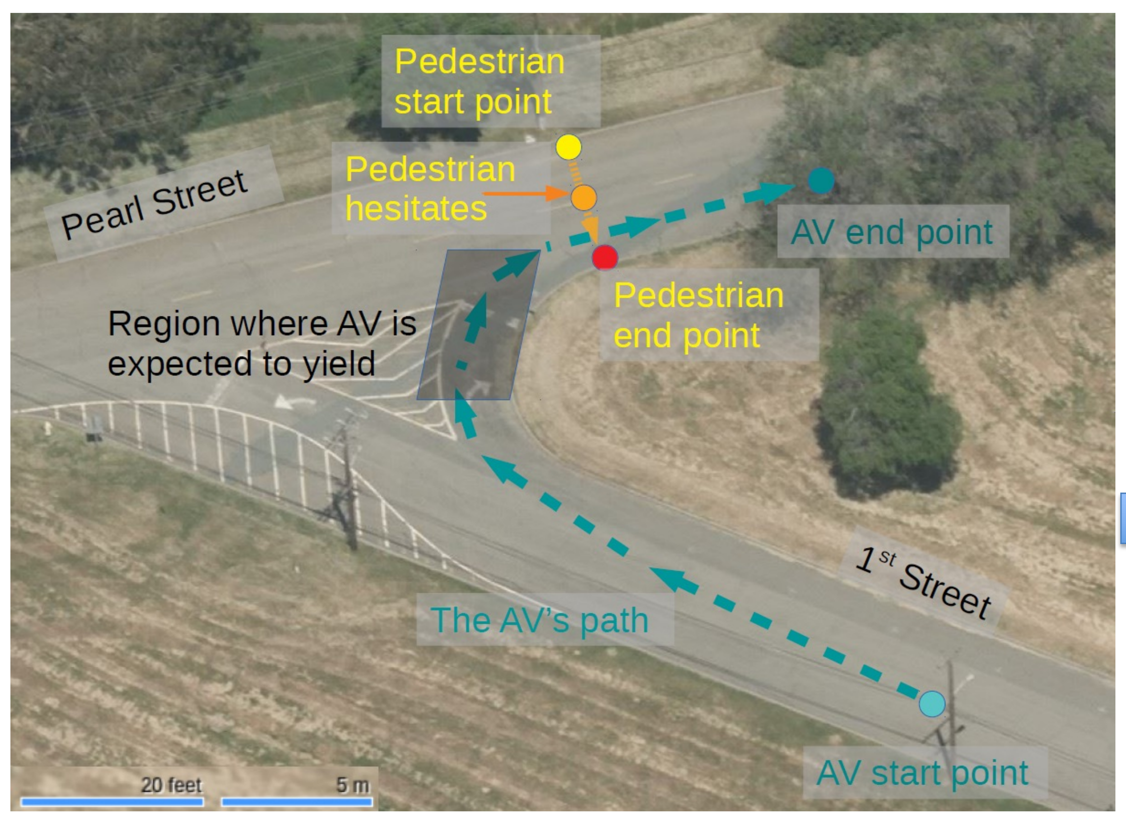

Scenic Program

```
ego = EgoCar at 38.6 @ 183.9,  # LGSVL coordinates
      facing 10 deg relative to roadDirection,
      with behavior DriveTo(40 @ 225.2)
ped = Pedestrian at 19.782 @ 225.680,
      facing 90 deg relative to roadDirection,
      with behavior Hesitate,
      with startDelay (7, 15),   # (a,b) = uniform
      with walkDistance (4, 7),  # distribution on
      with hesitateTime (1, 3)   # that interval
```

ϕ_{safe} = G (dist > 2.5 m)

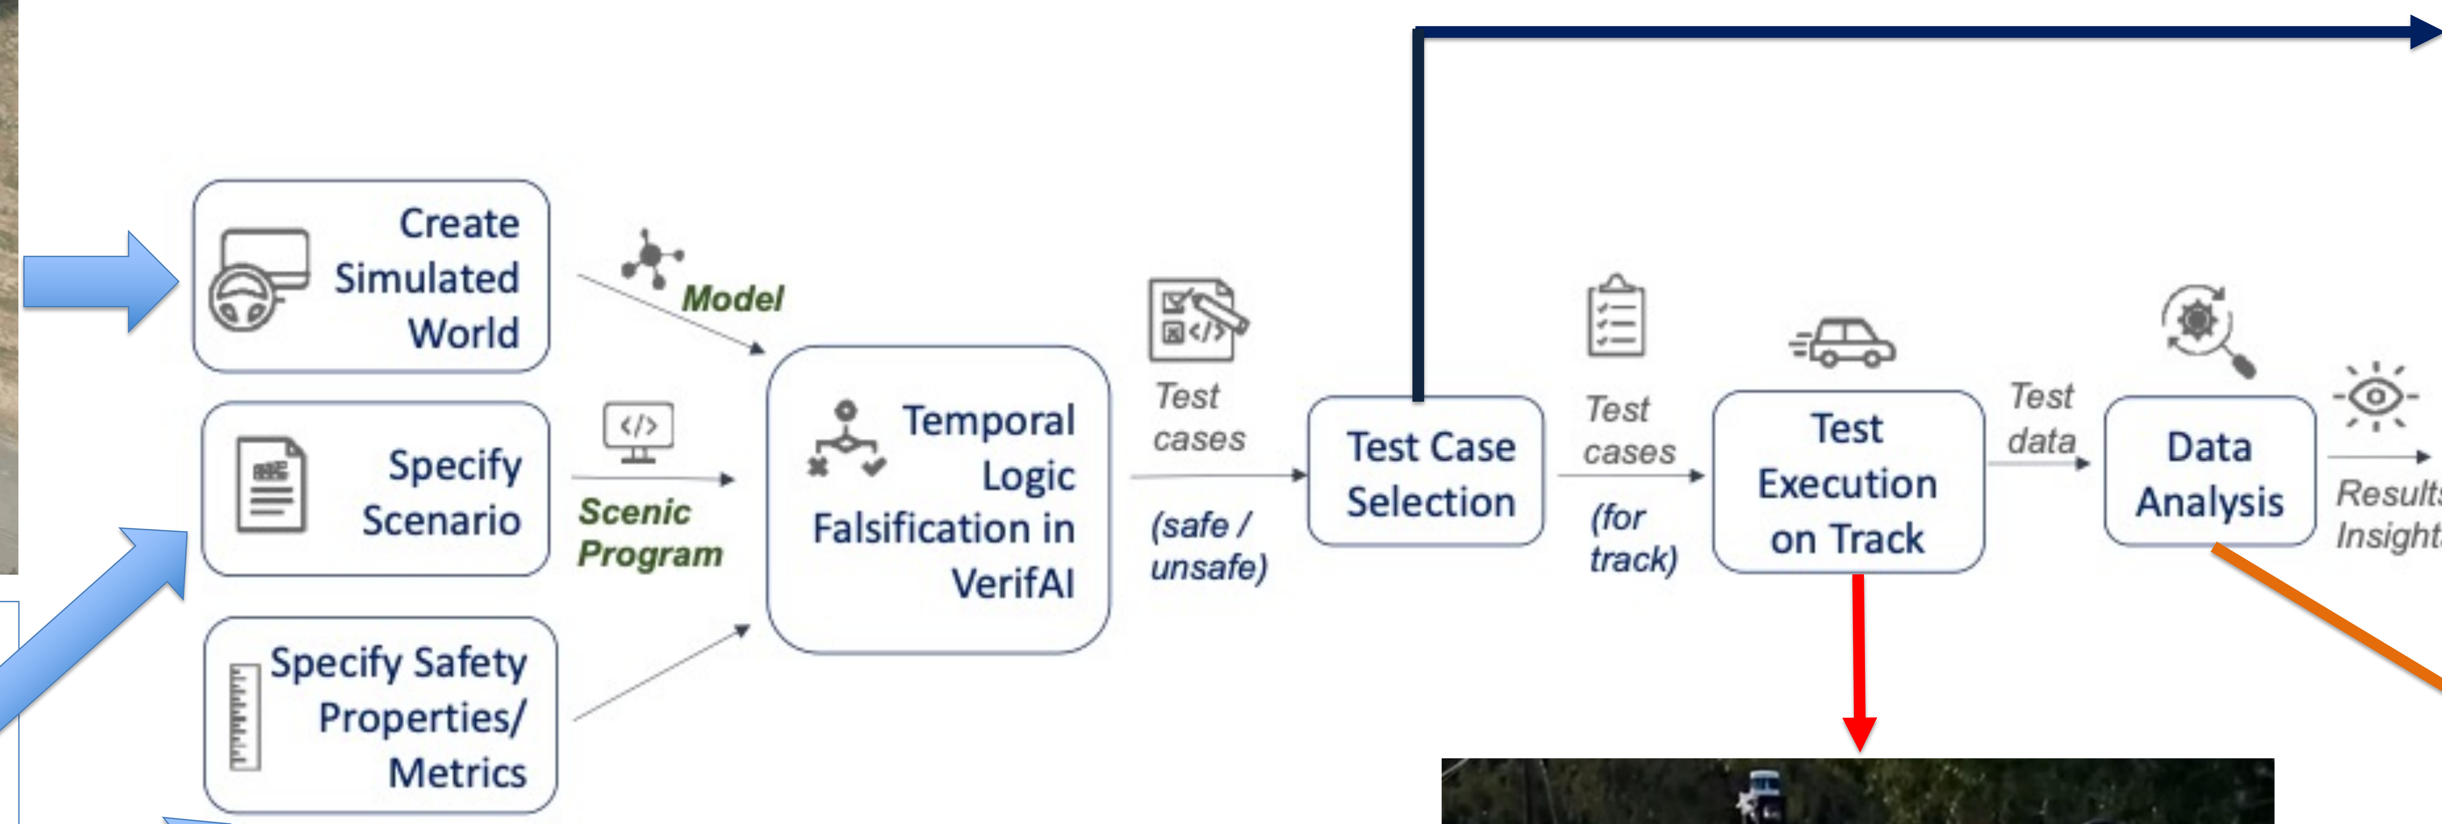

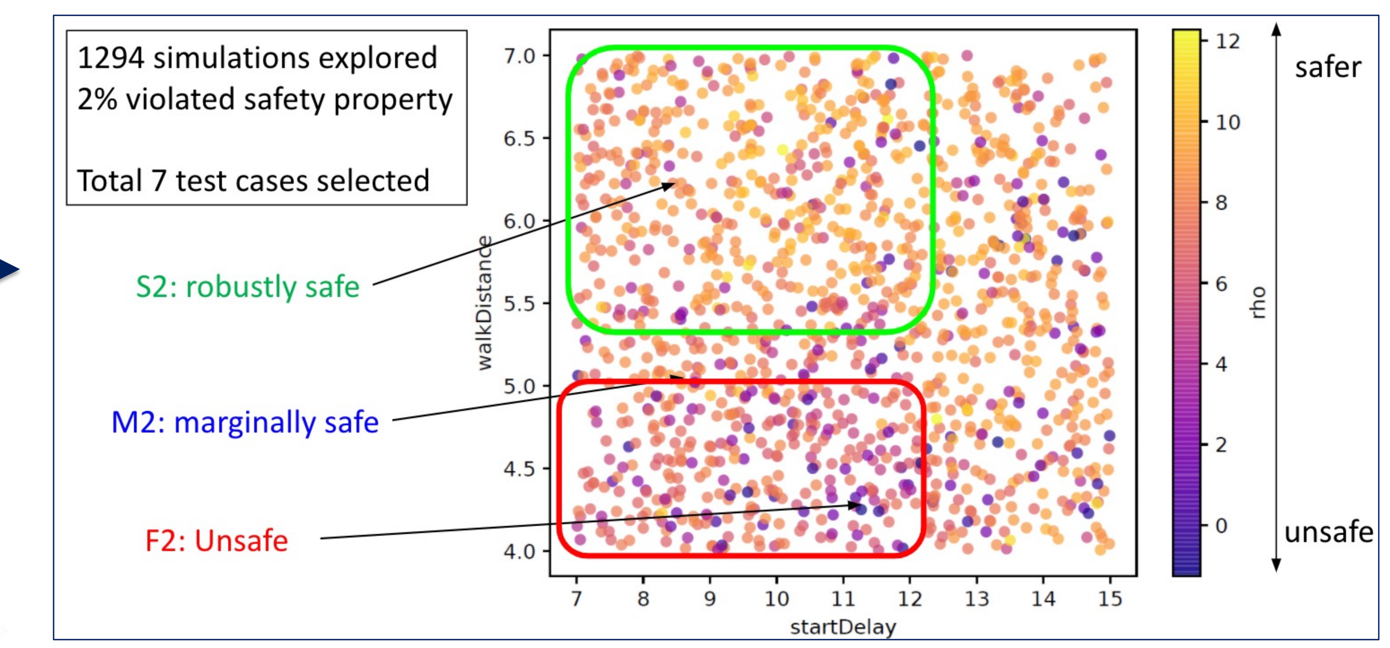

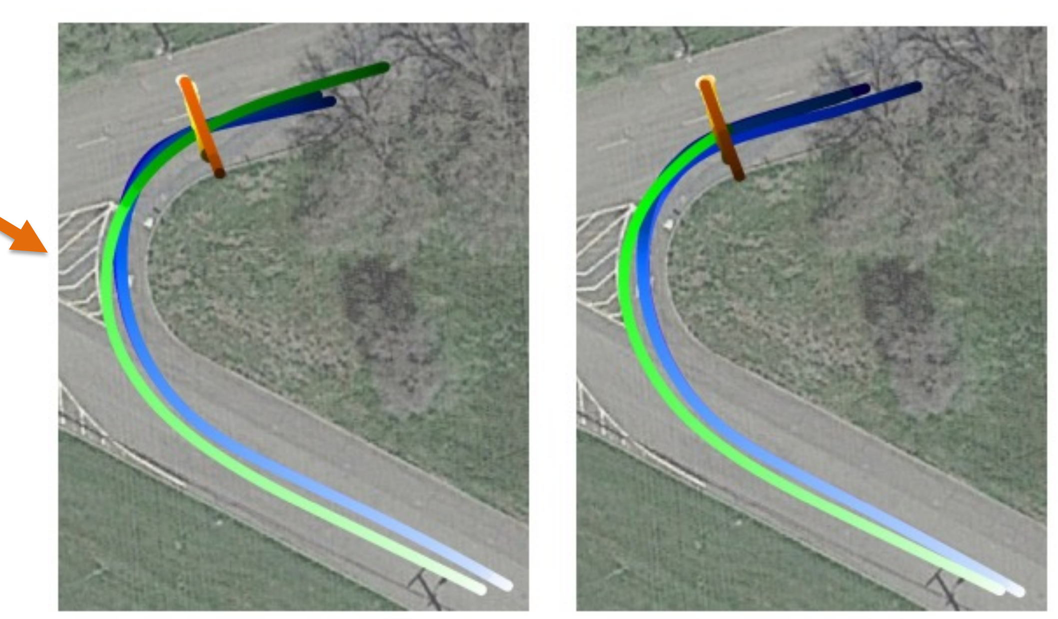

Fig. 6. Trajectories from a track test and 5 resimulations respectively for the AV (green/blue) and the pedestrian (orange/yellow). Color darkness indicates time. Scenarios: S1 Run 2 (left), F1 Run 1 (right).

Unsafe Tests in Simulation → Unsafe in Real World: 62.5%

Safe in Simulation → Safe in Real World: 95%

Broader Impact:

Our methodology is directly applicable to testing self-driving cars at track testing facilities to identify effective test cases, which is crucial for a scalable testing. However, at a larger scope, this methodology is applicable in testing systems which operate in a dynamic, interactive, and multi-agent environment which can be modelled as scenarios.

From education perspective, the outcome of our experiment across simulation and reality signifies the sensor realism issue where autopilot may perform differently on synthetic versus real sensor data.

References

[1] D. Fremont, T. Dreossi, et al, “A language for scenario specification and scene generation,” Programming Language Implementation and Design (PLDI), 2018

[2] Daniel Fremont, Edward Kim, et al. “Scenic: A Language for Scenario Specification and Data Generation,” https://arxiv.org/abs/2010.06580

[3] T. Dreossi, D. Fremont, et al. “VerifAI: A Toolkit for the Formal Design and Analysis of Artificial Intelligence-Based Systems,” *International Conference on Computer Aided Verification (CAV)*, July 2019

2021 NSF Cyber-Physical Systems Principal Investigators' Meeting
June 2-4, 2021

Award ID#: CNS-1545126

What Can Simulation Teach Us About Grasping 3D Deformable Objects?

Isabella Huang, Ruzena Bajcsy, in collaboration with NVIDIA

Motivation

Grasping deformable objects is underexplored in robotics, and can even be unintuitive for humans. We seek to **build intuition for deformable grasping** through simulation of ~4600 grasps

How would you grasp each of these deformable objects? Deformation should be minimized on the cup to avoid dislodging its contents. Stresses should be minimized on the tofu to prevent breakage. On the teddy bear, any grasp works.

Contributions

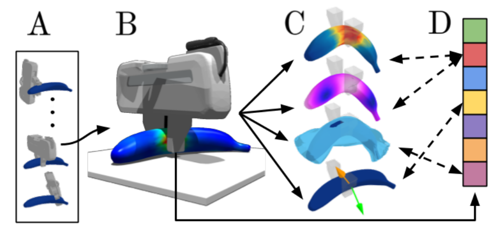

(A) For a broad set of candidate grasps on a deformable objects,
(B) We simulate the object's response with FEM,
(C) Measure 7 performance metrics (e.g., stress, controllability), and
(D) Identify 7 pre-pickup grasp features (e.g. squeezing distance, gripper distance to object center of mass) that are correlated with the metrics.

Prediction on Unseen Objects

Some features are found to be strongly correlated to some metrics. We then use these **correlations to predict the metrics** on unseen objects.

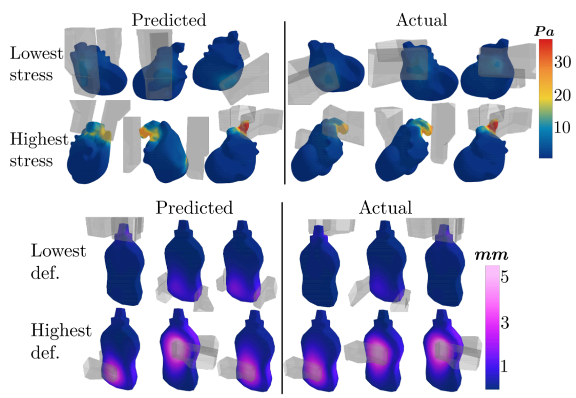

We demonstrate good predictions for the most extreme grasps for stress on a heart (top) as well as the most extreme grasps for deformation on a hollow bottle (bottom)

Sim-to-Real Validation

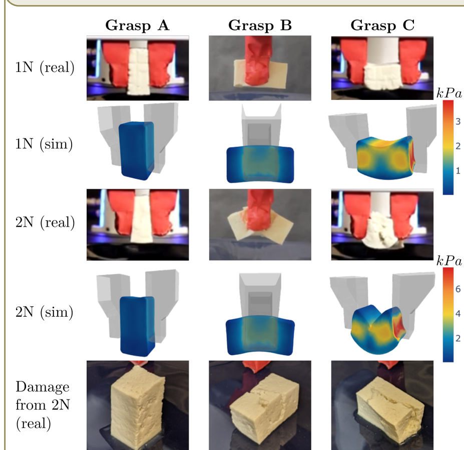

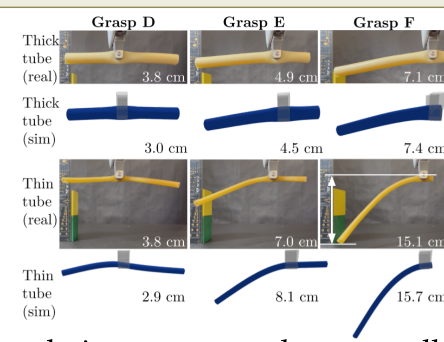

Simulation corresponds very well to the real world, e.g., 3 grasps on tofu (left) and a hollow latex tube (right) without parameter tuning.

2021 NSF Cyber-Physical Systems Principal Investigators' Meeting
June 2-4, 2021

Award ID#: 1545126

DEC-LOS-RRT: Decentralized Path Planning for Multi-robot Systems with Line-of-sight Constrained Communication [To Appear in CCTA 2021]

Victoria Tuck, Yash Vardhan Pant, PIs: Sanjit Seshia, S. Shankar Sastry
https://vehical.org

Project Goal

A decentralized algorithm that given line-of-sight communication between agents (including via multi-hop), has agents

- reach their goal position from a valid starting position
- avoid static obstacles in a known space
- maintain a desired distance from other agents

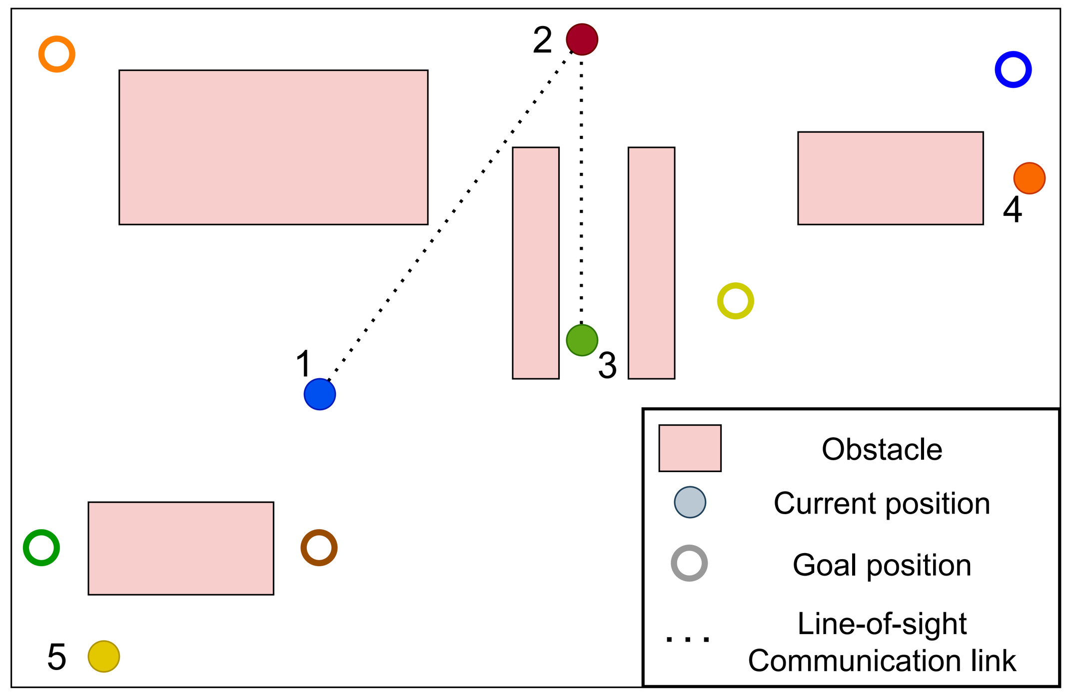

Agents can only communicate with agents in their subgraph. A subgraph is defined by an agent's visible neighbors and any agent in a visible neighbor's subgraph.

DEC-LOS-RRT Algorithm

Algorithm assumes valid starting positions, instantaneous stop, lossless communication with no latency, and single integrator dynamics.

1. Start base RRT-based, safe, decentralized algorithm for each subgraph
2. Update agent waypoints per base decentralized algorithm
3. Stop movement when subgraph changes (e.g., a new agent is seen)
4. Restart base decentralized algorithm for new subgraph of agents
5. Repeat 2-5 until all agents reach their goal or a lock is reached

The algorithm introduces the use of **delta obstacles**. In the right figure, green, solid boxes are obstacles, and blue, dashed boxes are delta obstacles. Avoiding delta obstacles with use of instantaneous stop ensures safety.

11 agents run the DEC-LOS-RRT Algorithm. Safety is assured. Although it is not guaranteed that agents will reach their final positions, most runs resulted in goal attainment.

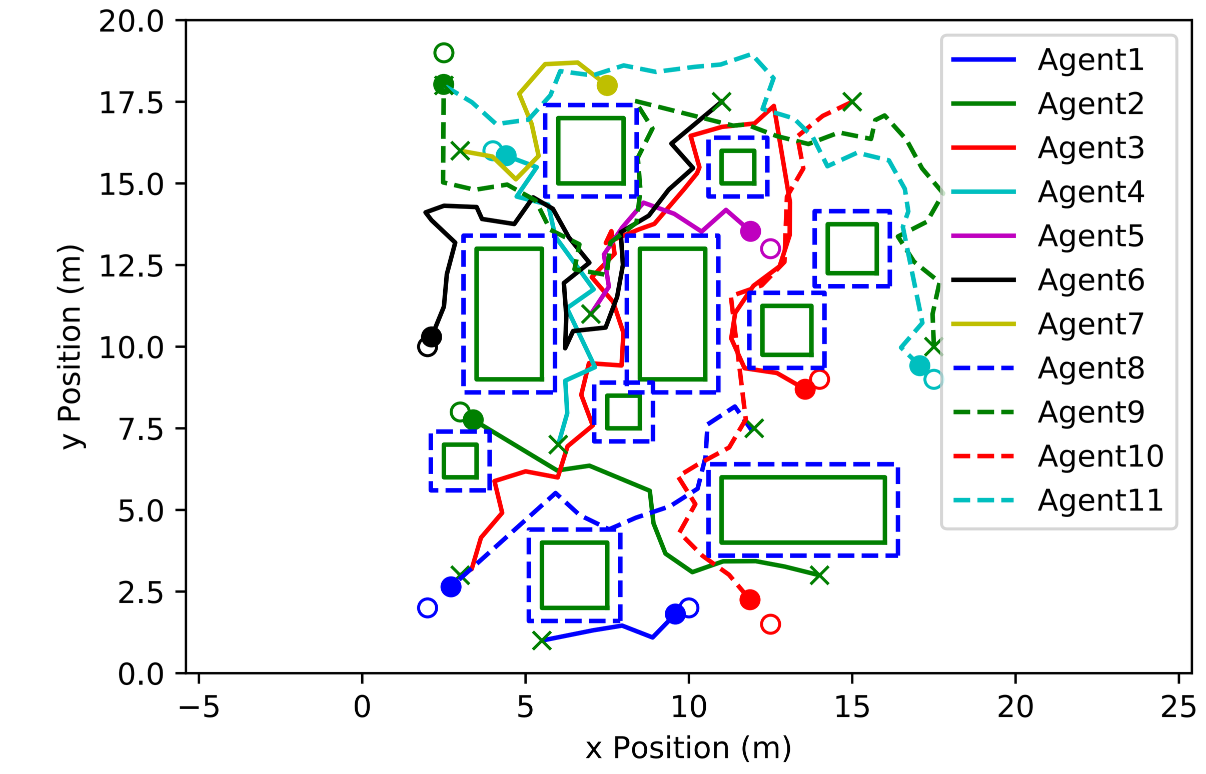

Future Directions: Assumptions such as instantaneous stop and single integrator dynamics limit applicability. In future iterations of this project, we will approach a similar problem for differentially flat systems with more realistic communication and jerk models.

CPS Applications: Low-power communication links that cannot be established through solid obstacles may necessitate an algorithm that accounts for the possibility of an impending crash with an agent that is close but not yet seen. Additionally, such an algorithm would assist autonomous vehicles in avoiding situations where a hidden pedestrian moves into a position that the vehicle cannot avoid.

Broader Impact: In large CPS fleets, a centralized solution to the communication constrained setting will likely not scale, necessitating a decentralized solution that can be trusted in safety-critical societal systems.

Outreach Participation by Authors: Bay Area Scientists in Schools, Girls in Engineering, Be A Scientist

Award ID#:1545126

Model-based Formalization of the Autonomy-to-Human Perception Hand-off

Yash V. Pant, Balasaravanan T. Kumaravel , Ameesh Shah, Erin Kraemer, Marcell Vazquez-Chanlatte, Kshitij Kulkarni, Bjoern Hartmann, Sanjit A. Seshia

https://vehical.org/

Motivation: Autonomous vehicles (AVs) are far from achieving `Full-Self Driving' and need to bring the driver into the decision-making loop in safety-critical situations. This however ensures safety only when the human is attentive and makes a correctly and timely decision. We focus on the *perception hand-off*, where an AV's perception module requires human supervision to interpret the environment. We formalize this Human-Robot Interaction (HRI) to develop an approach for modeling and influencing human attention, even in the presence of a non-driving related task (NDRT), for timely and correct decision making in perception hand-offs.

Challenge problems:

1. How does attention impact human decision-making in safety and time-critical situations?
2. How does attention evolve over time?
3. How can attention be *estimated* and *influenced* via *active information gathering* (AIG)?

Scientific impact:

Human-aware Model-based design for the Operator-Autonomy interface.

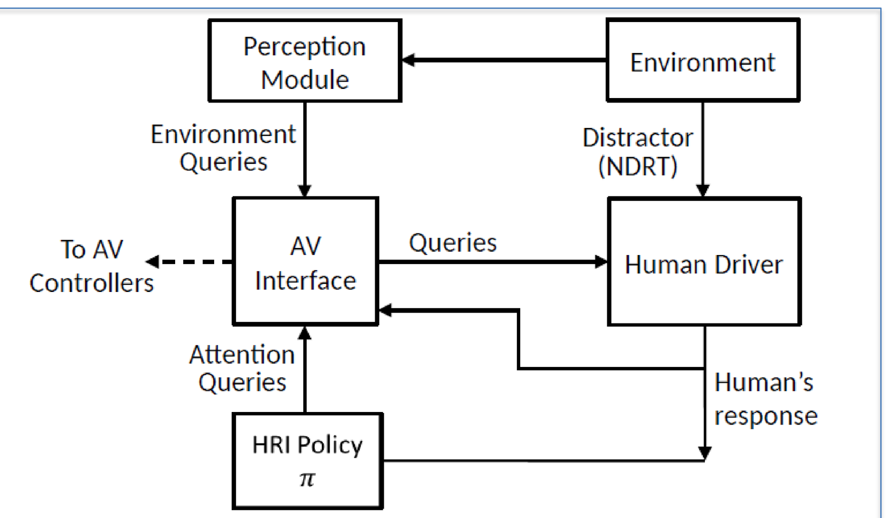

Perception hand-off: The AV queries the human driver when its perception module requires help in decision making, or to influence the human's state of attention. Further influencing the state is a NDRT.

Web-based human study (40 participants):

- ❑ Query-based AIG mechanism
- ❑ Dual-task: Driving related (primary) and simple arithmetic (distractor, or NDRT).

Key insights:

- ❑ Presence of NDRT degrades human response (slower and incorrect).
- ❑ AIG actions can improve human response times (≅ 7%) in the presence of NDRT.

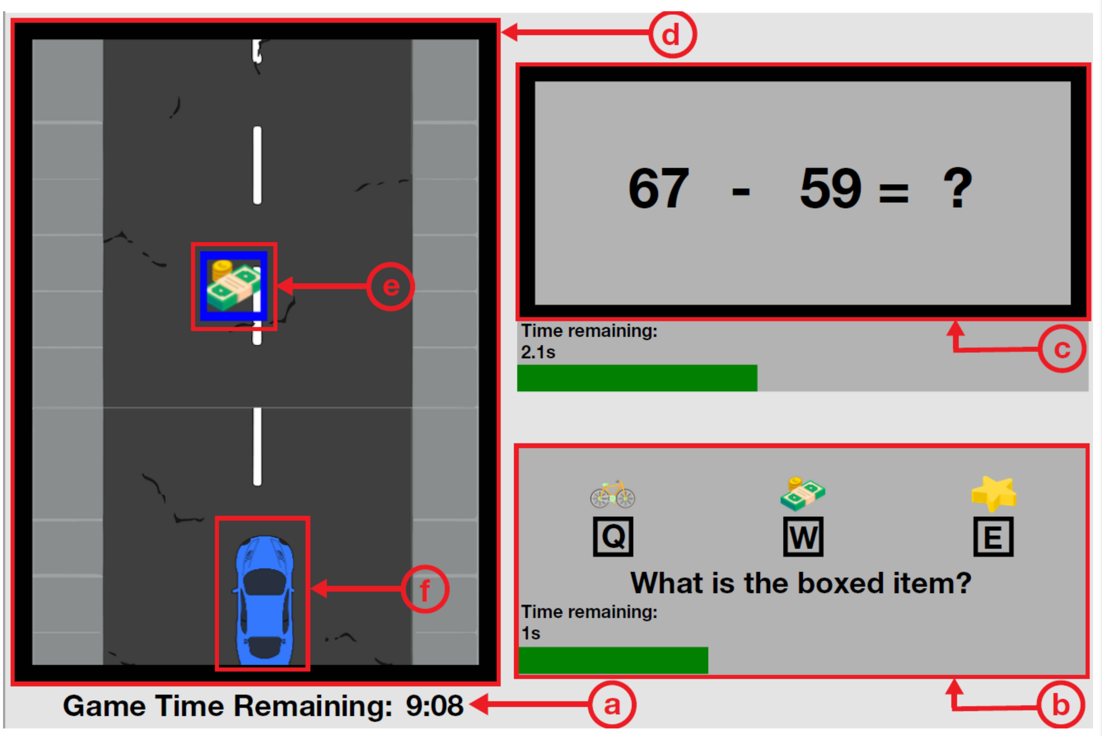

Dual-task experiment: The human identifies objects, while also solving arithmetic problems.

Methodology: Model-based perception hand-off

- ❑ Partially Observable Markov Decision Process (POMDP) model with novel structure.
- ❑ Hidden states are attention levels ($l_1 \prec \ l_2 \prec \cdots \prec l_N$)
- ❑ Actions are queries from the perception module or active information gathering actions
- ❑ Observations are human responses over time
- ❑ Learn parameters from human study data
- ❑ Compute policy for AIG actions

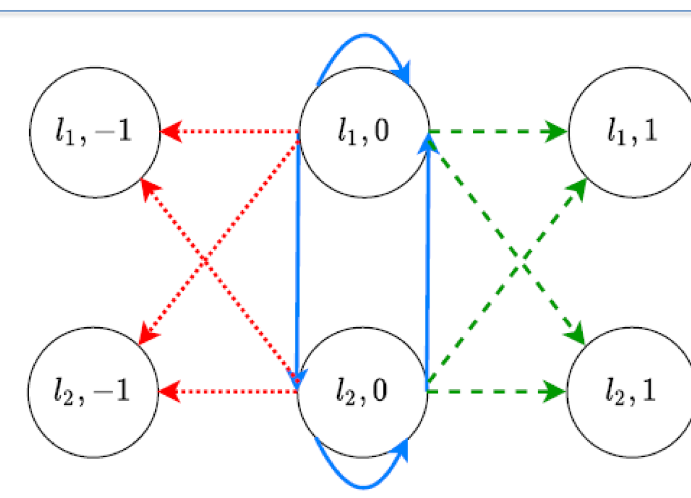

Simplified POMDP state-space: Blue arrows show transitions of attention in absence of queries. Red and green arrows show transitions when queries are active. Two attention levels for simplicity.

Simulation study with learnt POMDP model as a surrogate for the human

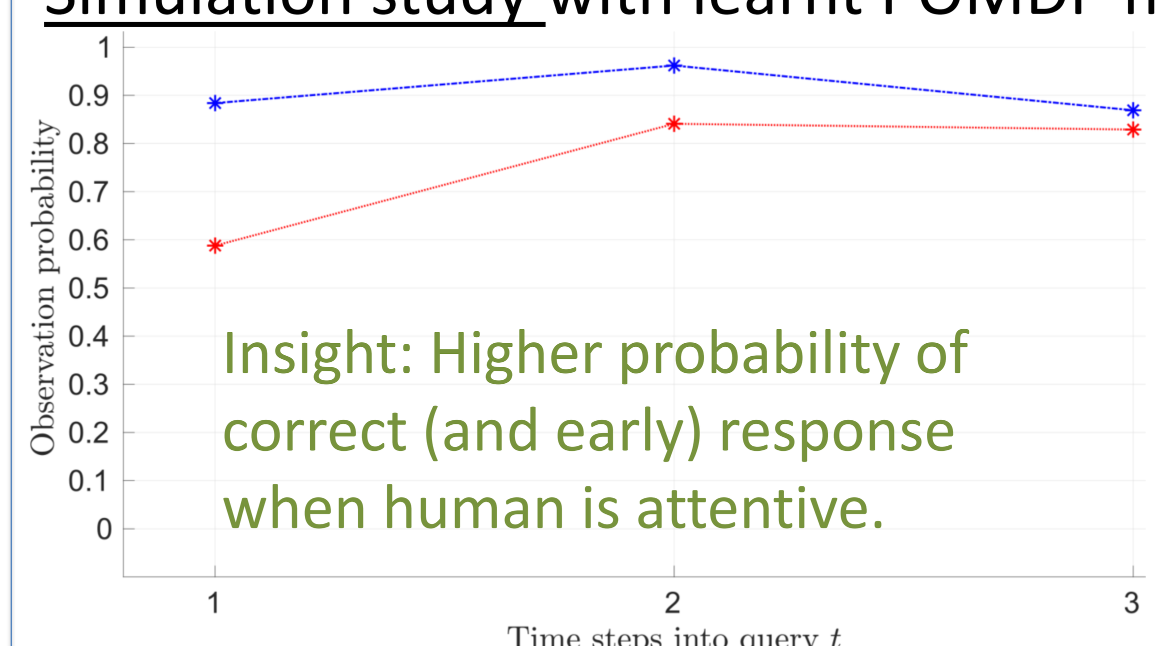

Observation probabilities: Learnt probabilities of correct response vs time steps into query, when latent state is attentive (blue) and inattentive (red).

Policy	Reward (R)	T_{resp}	$f * 100$	$\#a^{AIGA} : \#a^{PER}$
Learned	**15.52** ± 5.27	1.56 ± 0.05	**98.2**± 2.8	**0.83** ± 0.12
No AIGA	11.29 ± 5.55	1.57 ± 0.07	92.8± 3.9	0
Random	11.78± 6.42	1.55 ± 0.04	95.4 ± 3.1	1.48 ± 0.14
Belief	13.83 ± 4.39	**1.54** ± 0.03	95.9 ± 2.8	2.9 ± 0.11

Simulation results: Optimal `Learned' policy for a reward (R) that incentivizes correct and fast responses from the human, verses baseline methods.

- The fraction of perception hand-offs correctly responded to (f) is highest for this policy.
- It also uses fewer AIG queries per each perception hand-off query ($\#a^{\{AIGA\}}$: $\#a^{\{PER\}}$) than other non-trivial baselines.

Conclusion and ongoing work

- ❑ Model-based formulation allows for estimating and influencing human attention for safer perception hand-offs.
- ❑ Immersive human study in development to overcome limitations of current setup, and to add richer signals to improve modeling.

Potential broader impact (societal):

- ❑ Improved safety of autonomous and semi-autonomous vehicles.
- ❑ Principled operator-autonomy interaction beyond AVs.

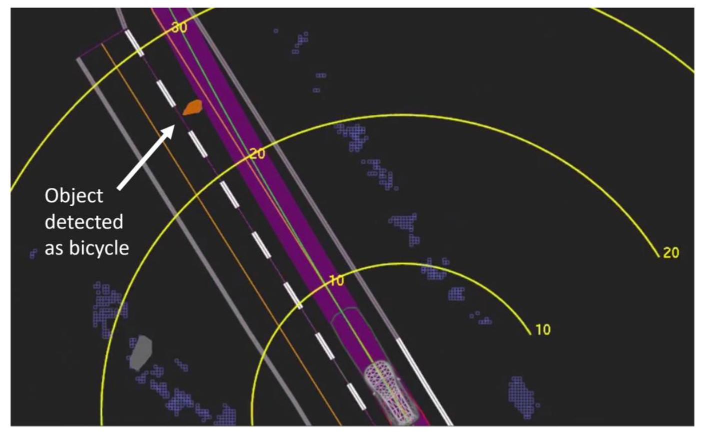

Uber crash in Arizona, 2019. Image from NTSB report.

Education and Outreach:

Researchers from the project served as mentors in the UC Berkeley Girls in Engineering program, serving as technical experts. They also discussed their research with the school-going participants.

2021 NSF Cyber-Physical Systems Principal Investigators' Meeting
June 2-4, 2021

Award ID#: CNS-1545126

Specifications from demonstrations; A Maximum Entropy Approach

Marcell Vazquez-Chanlatte Sanjit A. Seshia

NSF Grant #1545126

What was the agent trying to do?

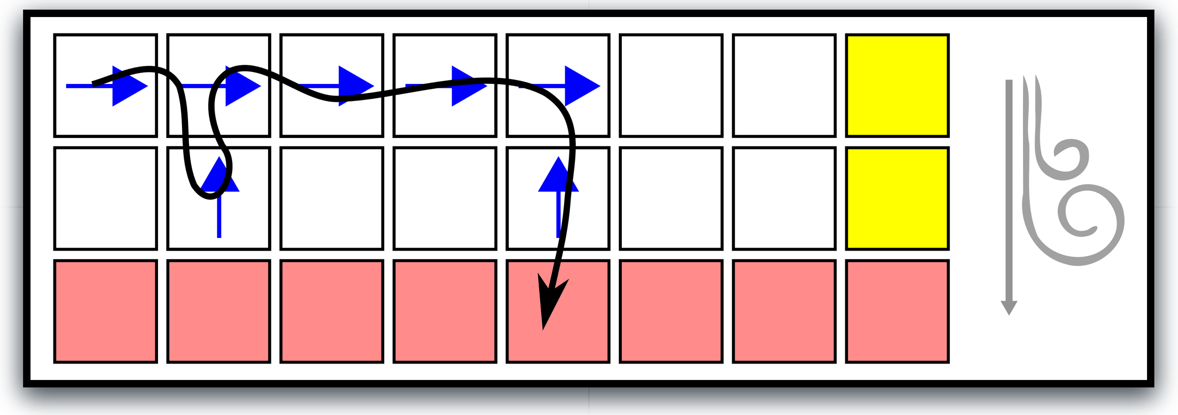

Q: Did the agent intend to touch the red tile?

Problem Statement

Given unlabeled demonstrations, learn a formal specification that "explains" the teachers behavior.

Why not Rewards?

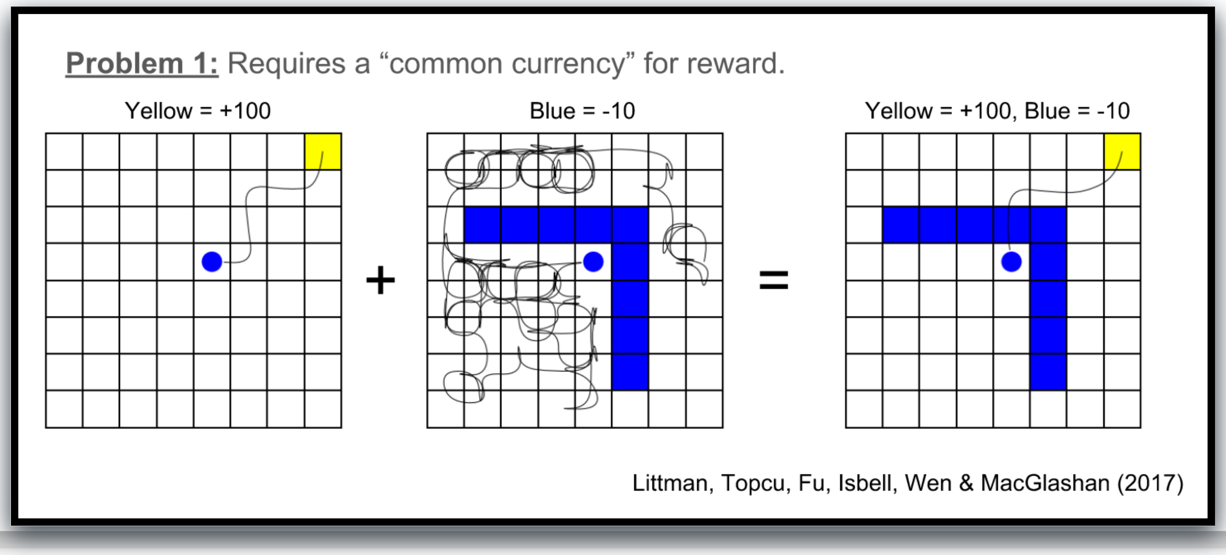

Contributions

1. Robustly learn trace properties from **unlabeled** demonstrations in Markov Decision Processes.
2. Symbolic approach for efficently representing Markov Decision Processes as **Binary Decision Diagrams**.

Symbolic Maximum Causal Entropy Likelihood Estimatation

Key Observation: Can think of soft constraint as binary reward.

$$r_\lambda(\xi) \triangleq \lambda \cdot 1[\xi \in \varphi]$$

- By adding history to state space, can reduce to Maximum Causal Entropy Inverse Reinforcement Learning.
- **Problem:** Potential combinatorial explosion.
- **Solution:** Encode MDP as a Binary Decision Diagram.

1. Write the **composition** of the dynamics and property as a circuit with access to biased coins.

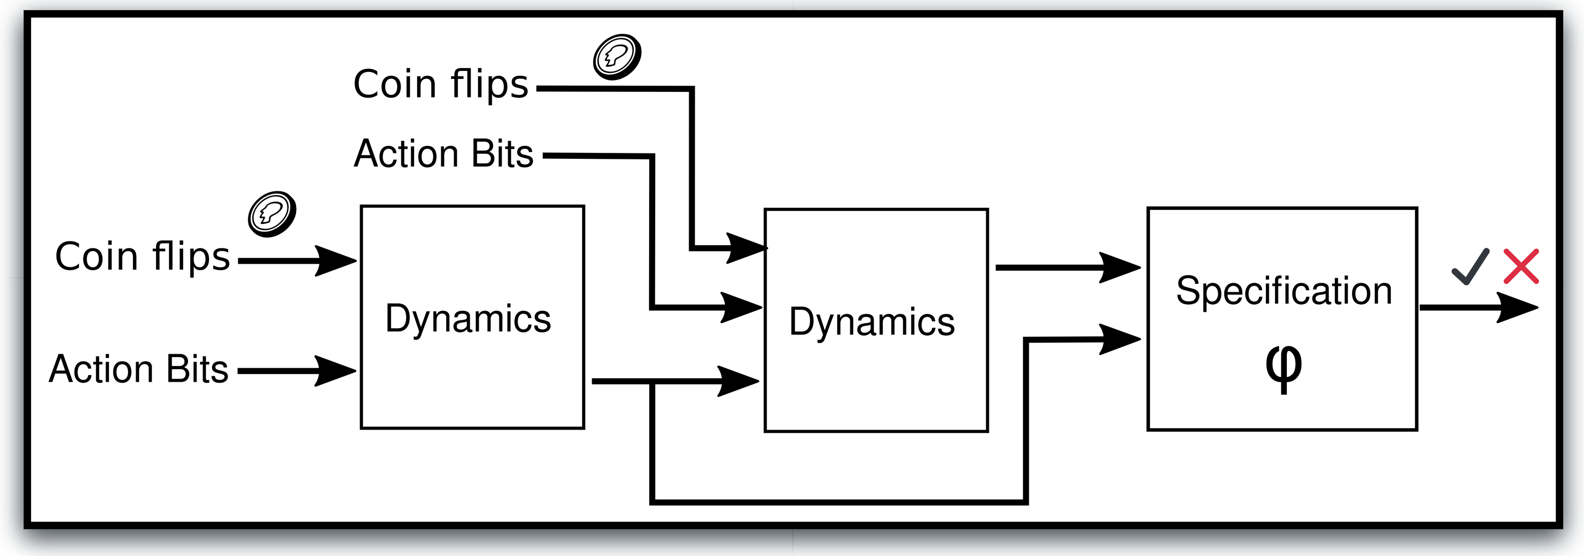

2. **Idea:** Symbolically encode MDP as a Binary Decision Diagram:

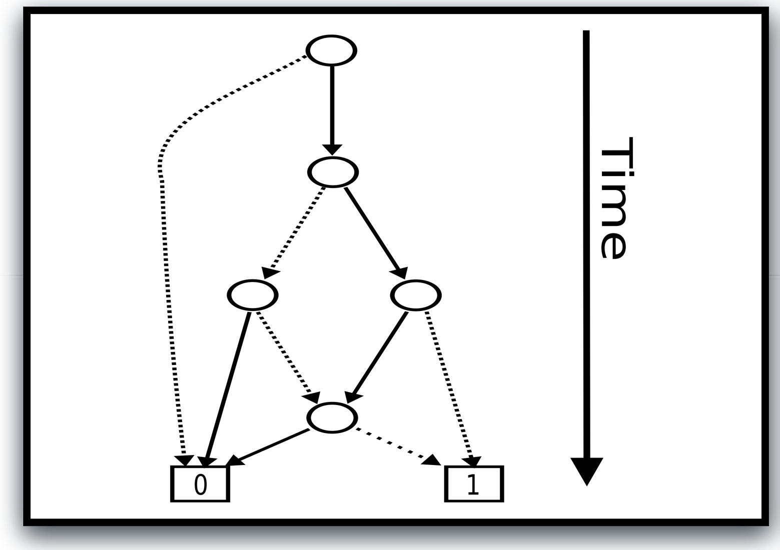

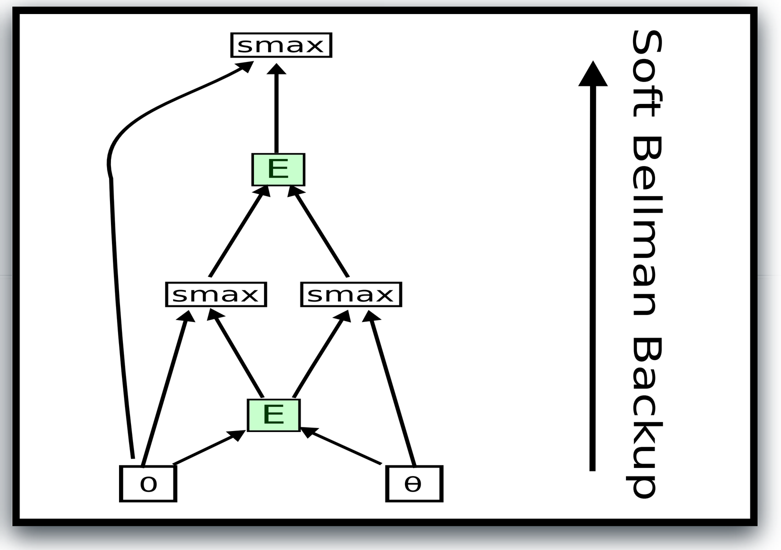

Conservative size bound:

$$O(|\text{horizon}| \cdot |S/\varphi| \cdot |\text{Actions}| \log(|\text{Actions}|))$$

3. We show you can efficiently compute maximum causal entropy policy on compressed MDP.

Application: Used to learn temporal logic constraint from **unlabeled** demonstrations, e.g.,

φ = "Avoid Lava, eventually recharge, and don't recharge while wet."

Experiment: Learn rules given 6 *unlabeled* demos.

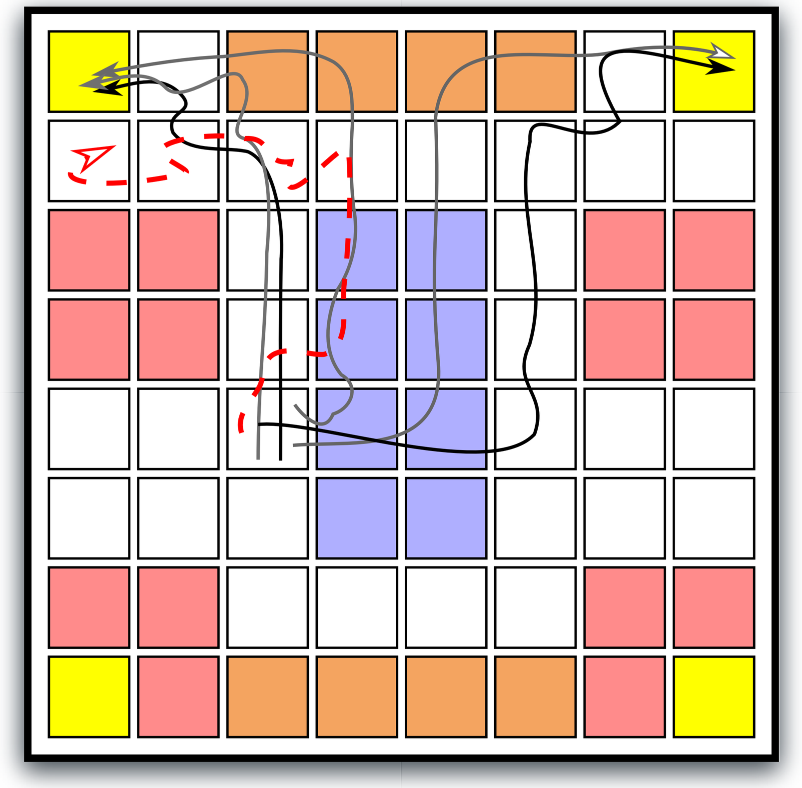

Dynamics

Actions = {↑, ↓, ←, →}.
Probability $\frac{1}{32}$ to slip and move ←.

Rules

1. Go to and stay at the yellow tile.
2. Avoid red tiles.
3. If you enter a blue, touch a brown tile **before** recharging.

Spec	Policy Size (#nodes)	ROBDD build time	Relative Log Likelihood (Compared to True)
true	1	0.48s	0
φ_1 = rule 1	1628	1.2s	5
φ_2 = rule 2	1797	1.5s	-22
φ_3 = rule 3	750	1.6s	-10
$\varphi_4 = \varphi_1 \wedge \varphi_2$	523	1.9s	4
$\varphi_5 = \varphi_1 \wedge \varphi_3$	1913	1.5s	-2
$\varphi_6 = \varphi_2 \wedge \varphi_3$	1842	2s	15
$\varphi_\star = \varphi_1 \wedge \varphi_2 \wedge \varphi_3$	577	1.6	27

Key observation: φ_* more likely than consistent specifications.

Future Work

1. Teaching through demonstrations.
2. Inference in continuous domains.
3. Data driven concept classes - Natural Language Processing, Sampling consistent automata, etc.
4. Estimating Membership Queries: Is a given behavior is ok?

Rules of the Road: Formal Guarantees for Autonomous Vehicles with Behavioral Contract Design

Karena X. Cai*, Tung Phan-Minh*, PIs: Richard M. Murray, Soon-Jo Chung, California Institute of Technology
https://vehical.org

Abstract:

The ability to make formal guarantees on safety and performance for autonomous vehicles in **highly-interactive, dense environments** largely remains unsolved. With a **well-defined behavioral contract**, we can not only provide **formal guarantees** on agent safety and progress, but we also have a mechanism for **assigning blame** when accidents invariably occur. In this paper, we define a behavioral contract for a particular class of agents on a **road network environment** in a **quasi-simultaneous discrete-time game**. We provide **proofs of correctness** of the behavioral contract and **validate our results** in simulation.

Challenge:

How do we design a high-level decision making strategy for autonomous agents in highly-interactive environments to behave 'correctly', i.e. be safe, be lawful, and make progress towards its destination?

Extremely challenging because:

- Robot-freezing problem and unbounded rationality.
- Joint action space grows exponentially.
- Other agents can act to intentionally make safety impossible.
- Can't satisfy all road rules all the time, which to violate?

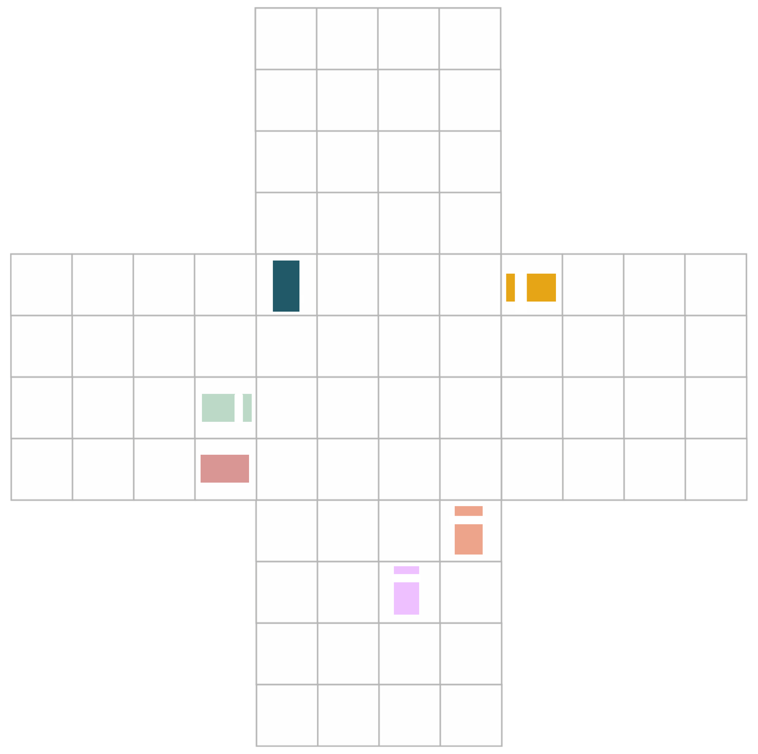

Scientific Impact:

Agent strategy (defined in a discrete-game and in specific road network environments that provides:

Safety guarantee

Safety Theorem

Given that all agents $Ag \in \mathfrak{A}$ in the quasi-simultaneous game $\mathfrak{G}$ select actions in accordance to the agent protocol defined, we can show the safety property:

$P \implies \Box Q$

P assertion that the game is in a state where every agent has a backup plan action that is safe.

Q assertion that agents never occupy the same grid point at the same time.

Performance guarantee

Liveness Theorem

Given the sparsity conditions hold, and that all agents $Ag \in \mathfrak{A}$ in the quasi-simultaneous game $\mathfrak{G}$ select actions in accordance to the agent protocol defined, we can show all agents will eventually reach their respective destinations.

Scalability & Interpretability

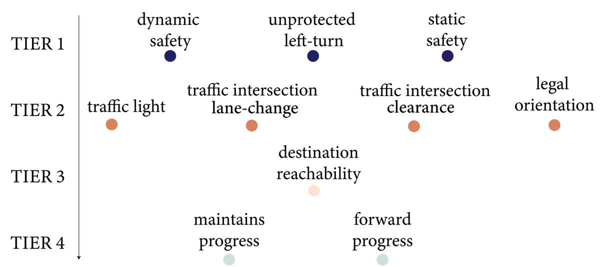

Notion of Blame/Liability

$$C_j = (A_j, G_j)$$
$$\forall j \in \mathcal{J}.\forall i \in \mathcal{J} - j.G_j \subseteq A_i$$

Definition II.2 (Blameworthy action). *A blameworthy action/strategy is one in which an agent violates its guarantees, thereby causing another agent's assumptions not to be satisfied and thus resulting in an unwanted situation where blame must be assigned.*

Solution:

Propose the design of a behavioral protocol agents should use to select actions.

Strategy ensures agents are always entitled to safely execute their backup plan action (i.e. maximal braking)

Pt. 1 Behavioral Profile

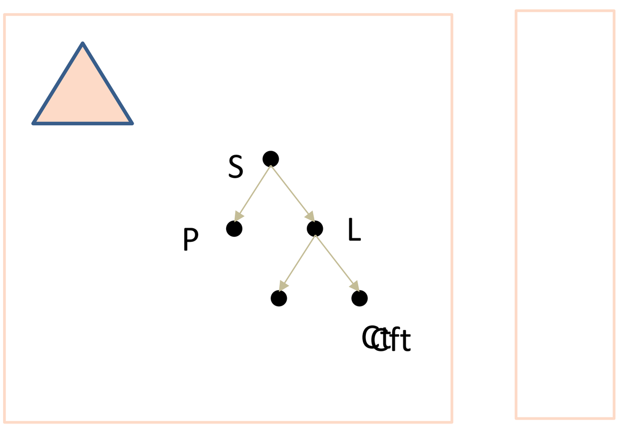

Pt. 2 Conflict Resolution Scheme

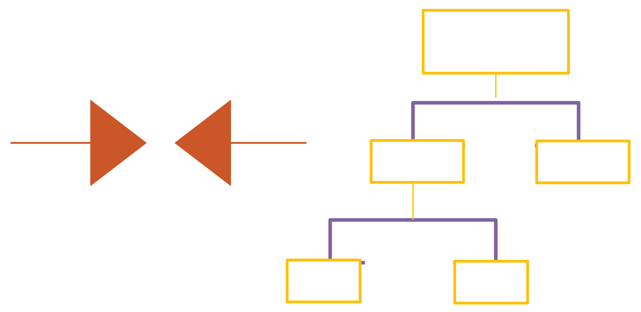

Proofs

1. Safety: no collisions.
2. Performance: agents make progress towards destinations. (under sparsity assumptions)

Simulations

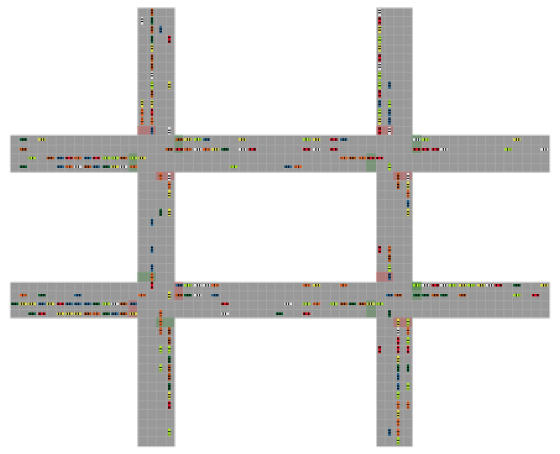

Broader Impact on Society

- Adoption of this type of framework will lead to safer and more interpretable autonomous vehicles on the road..
- Serves as a novel framework for designing vehicle behavior with the collective in mind (instead of the individual).
- Could be integrated alongside data-driven/machine learning approaches.

Broader Impact: Education and Outreach

Designed and hosted workshop on 'Building Effective Research Collaborations' to teach grad students communication and conflict prevention/management skills. Resources can be found: http://healthycollab.caltech.edu/

Quantifying Broader Impact:

- Potential to design autonomous vehicle algorithms that reduce number of collisions on the road.
- Also could help inform design of autonomous vehicle road rules and regulations.

2021 NSF Cyber-Physical Systems Principal Investigators' Meeting
June 2-4, 2021

Award ID#: CNS-1545126