Solving Traffic Problems using Autonomous Vehicles

[redacted]

Alexandre Bayen
Director, Institute of Transportation Studies
Professor, EECS & CEE
Faculty Scientist, LBNL

High level motivation

Planning the future of mobility: mixed autonomy

Neom, KSA

Jurong district, Singapore

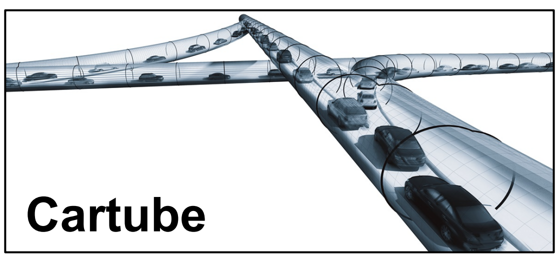

Cartube

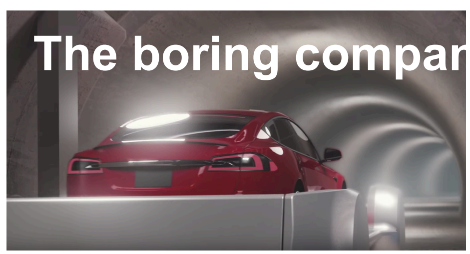

The boring company

Hyperloop

E-taxi

Example

Impact of automation on the energy footprint of mobility

Short answer:
it is **highly uncertain**.

Transportation today:
28% US energy consumption

100% self-driving cars:
-60% to **+200%** energy

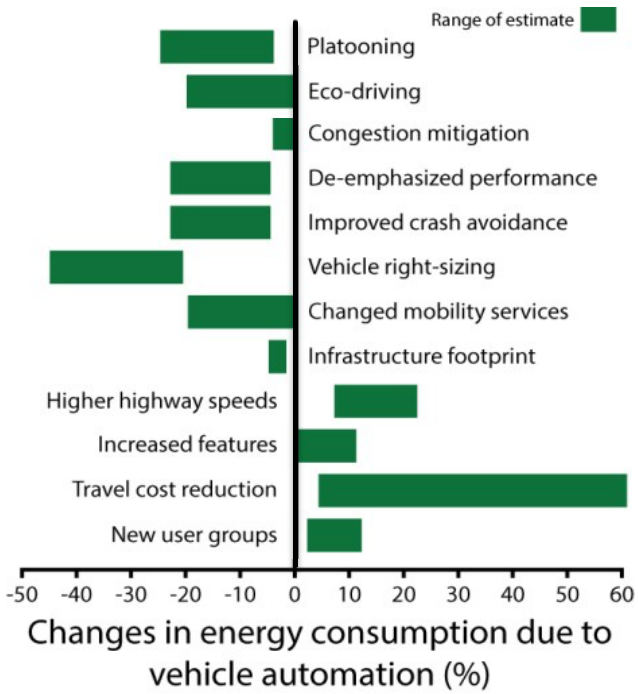

Z. Wadud, D. MacKenzie, and P. Leiby, “Help or hindrance? the travel, energy and carbon impacts of highly automated vehicles,” *Transportation Research Part A: Policy and Practice*, vol. 86, pp. 1 – 18, 2016.

Traffic control, traffic management

Forward simulation models

Variety of tools historically developed at different scales

- Energy-based vehicular models
- Microscopic models
- Mesoscopic models
- Macroscopic models
- Agent based models
- Excel accounting models

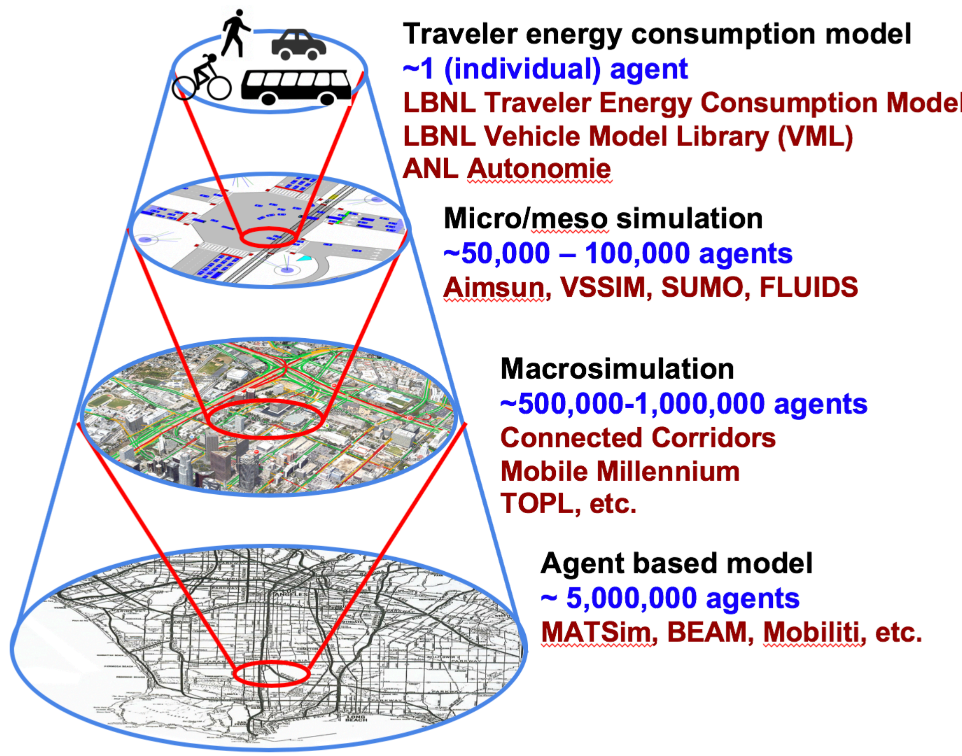

The state of the art microsim today

Microscopic simulation models: simulating 100,000s of vehicles

Example

- App "problem"
- Thru-traffic
- 20% app users
- nextGen DTA

Aimsun ONLINE

PTV VISSIM

the mind of movement

SUMO
Simulation of Urban MObility

Always in motion the future is

The next battlefield (2-5years)

Data:

~~Floating car data (GPS, cell tower, CAV data)~~
~~Asset data (signal timing, metering etc.)~~
~~Event data (closures, games, special events)~~
~~Maps assets (#lanes, speed limits etc.)~~

Calibration:

~~Estimation: vehicle-based~~
~~ID~~
~~Missing data inference~~
Demand
Routing

Model computation:

Computational time for model forecast
~~Distribution of the model on AWS EC2~~

Control / deep-RL:

Model free deep-RL
End2end learning / pixel learning
Sample inefficiency
Curse of dimensionality in the action space
Multi-agent learning training

A call to action

Can we over the next few (2) years demonstrate ML-based microsim?

Framework:
State of the art microsim: SUMO, Aimsun or other
RLLAB, RLLIB, TensorFlow, Caffe, etc.
All in AWS EC2 or similar cloud

(Mixed)-autonomy traffic control:
Every controllable asset (infrastructure or CAV) modeled
Most common scenarios solved: merge, intersections, freeways, arterials, roundabouts, tools, metered bridges, etc.

Can we demonstrate the following within 2 years?
Benchmarks for all, winner algorithms for all
Actual migration on assets (static and vehicles)

Acceleration of history

Deep-RL is about to leapfrog 80 years of model-based research

1935:

First aggregate model of congestion

1955:

First PDE model of traffic

Acceleration of history

Deep-RL is about to leapfrog 80 years of model-based research

1935:
First aggregate model of congestion

1955:
First PDE model of traffic

2008:
First experiment showing instability

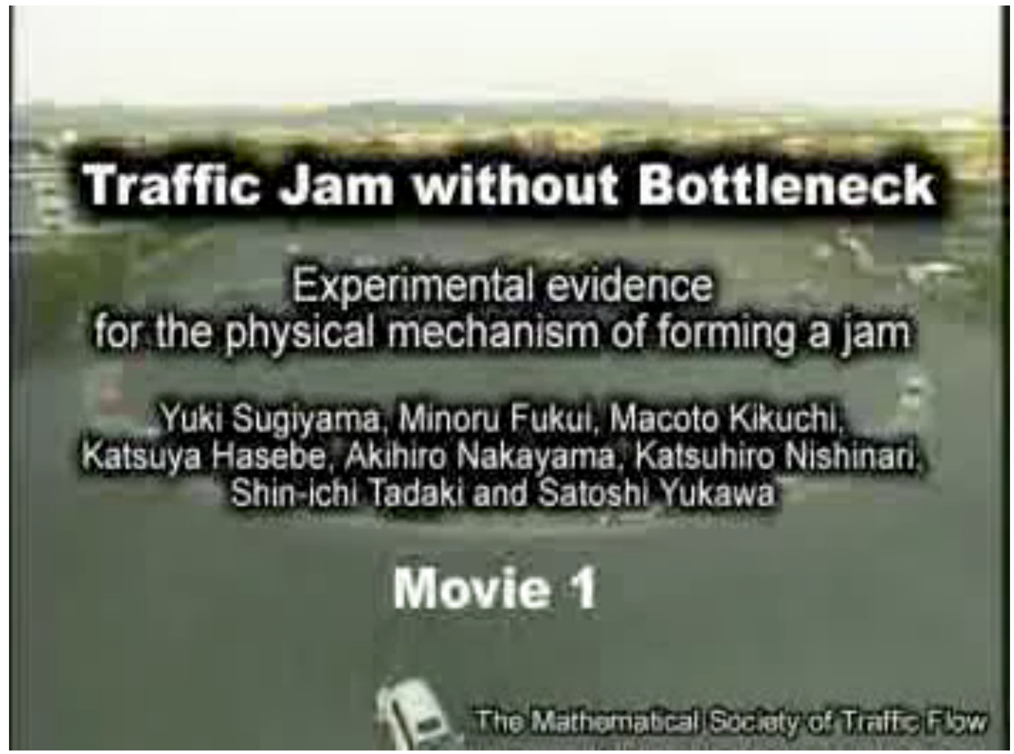

Acceleration of history

Deep-RL is about to leapfrog 80 years of model-based research

1935:
First aggregate model of congestion

1955:
First PDE model of traffic

2008:
First experiment showing instability

2017:
First controller implemented

Dissipation of stop-and-go traffic waves via control of a single autonomous vehicle

ILLINOIS UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN · RUTGERS · TEMPLE UNIVERSITY · THE UNIVERSITY OF ARIZONA




Prof. Daniel Work, Vanderbilt
Prof. Benedetto Piccoli, Rutgers
Prof. Benjamin Seibold, Temple
Prof. Jonathan Sprinkle, UoA

Acceleration of history

Deep-RL solution to the same problem

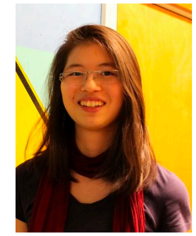

Prof. Cathy Wu, MIT

1935:
First aggregate model of congestion

1955:
First PDE model of traffic

2008:
First experiment showing instability

2017:
First controller implemented

2018:
Better result with deep-RL

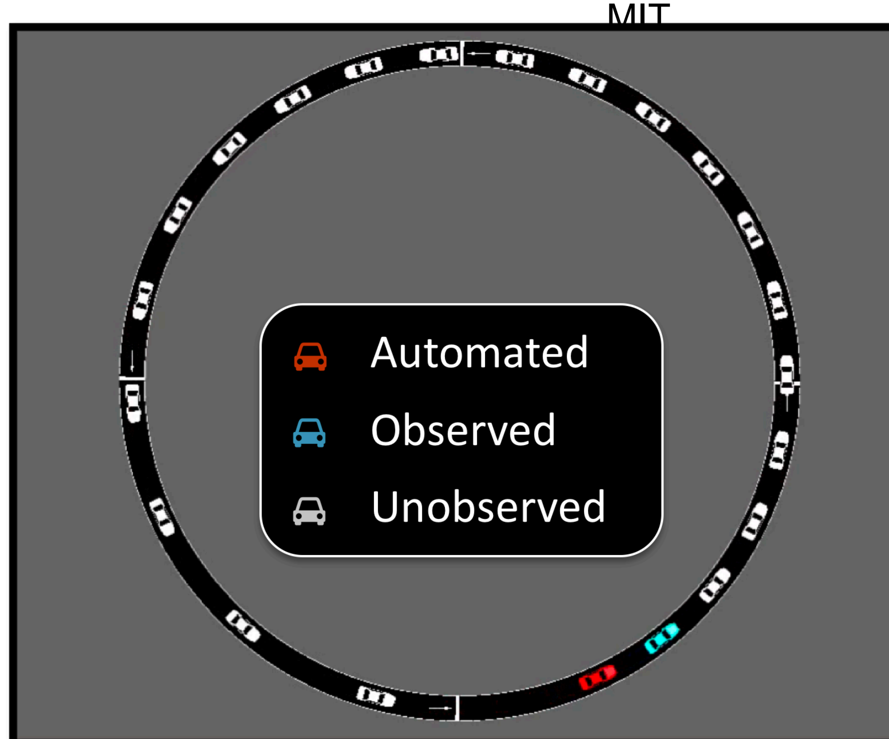

Acceleration of history

Deep-RL solution to the same problem

Prof. Cathy Wu, MIT

1935:
First aggregate model of congestion

1955:
First PDE model of traffic

2008:
First experiment showing instability

2017:
First controller implemented

2018:
Better result with deep-RL (Cathy Wu's PhD)

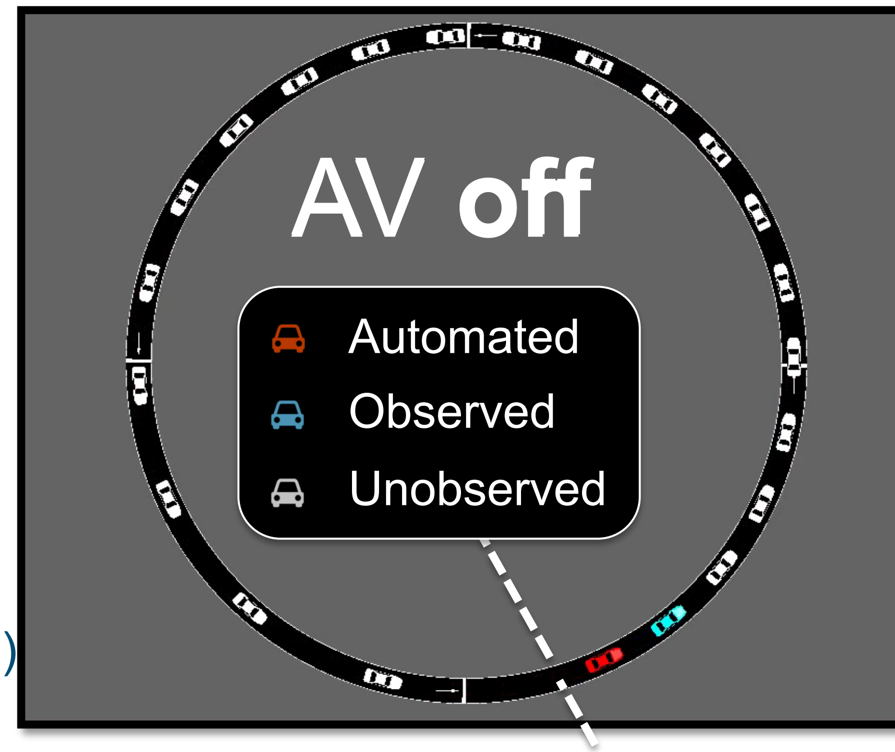

Acceleration of history

Deep-RL solution to the same problem

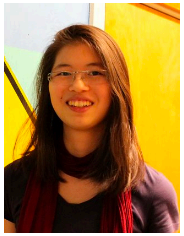
Prof. Cathy Wu, MIT

Setting: 1 AV, 21 human

Experiment:

- **Goal**: maximize average velocity
- **Observation**: relative velocity & headway
- **Action**: acceleration
- **Policy**: multi-layer perceptron (MLP)
- **Learning algorithm**: policy gradient

Results:

- 1 AVs: **+49%** average velocity
- **Stabilization at near-optimal velocity**

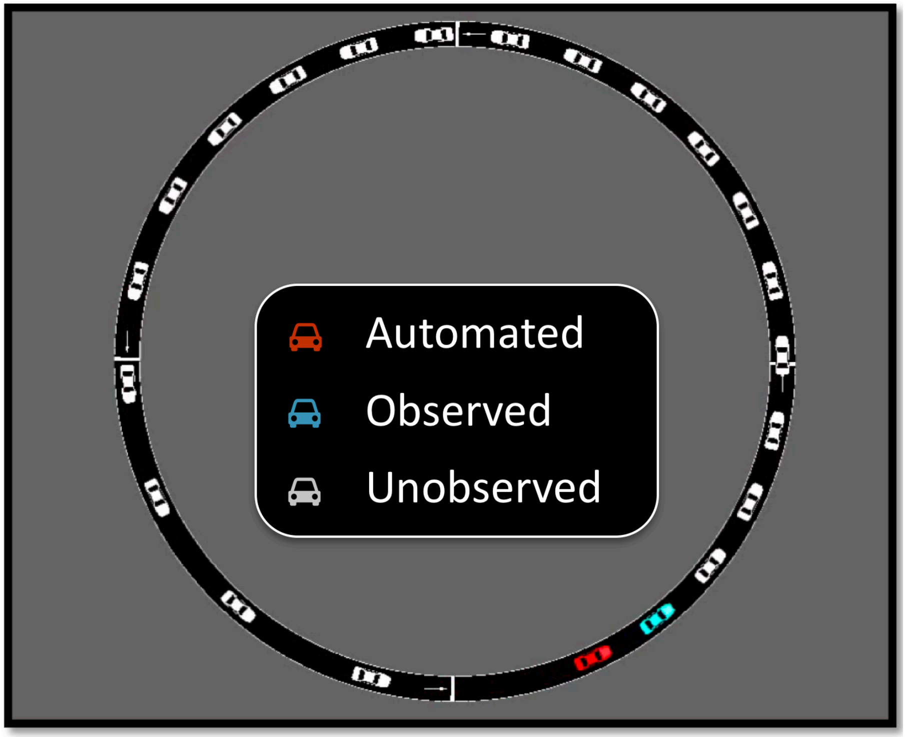

Acceleration of history

Will deep-RL leapfrog 80 yrs. of model-based research?

Prof. Cathy Wu, MIT

1935:
First aggregate model of congestion

~10,000 articles

1955:
First PDE model of traffic

2008:
First experiment showing instability

1,000 articles

2017:
First controller implemented

1 article

2018:
Better result with deep-RL (Cathy Wu's PhD

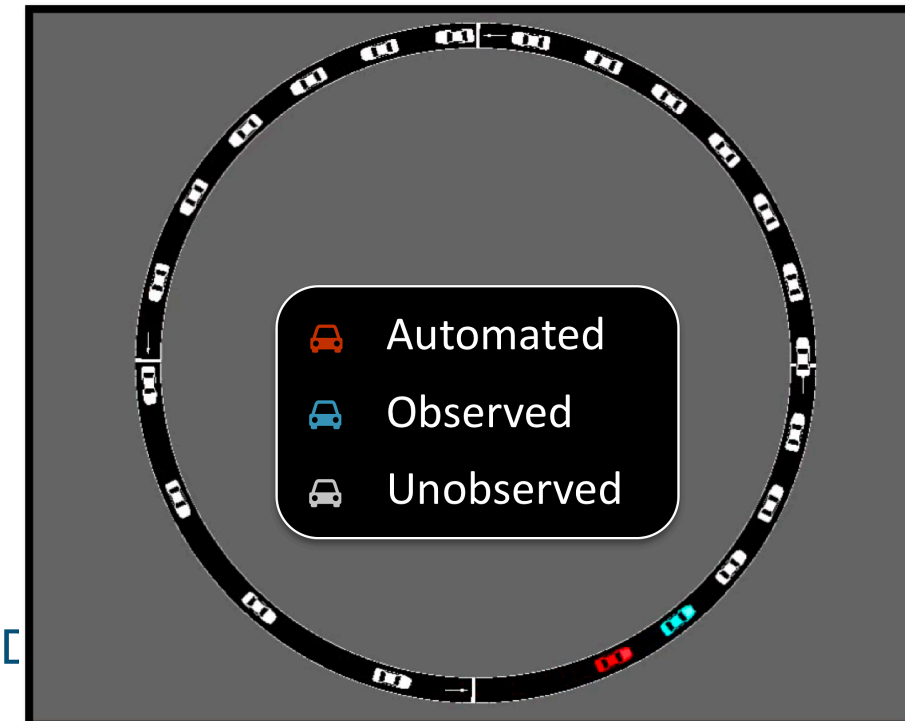

Problem statement

Traffic flow control by CAV (and static assets if needed)

CAV: 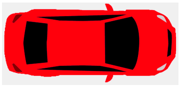
Autonomous,
Onboard policy (learned)
Connected to other CAVs

Sensed vehicle 
Sensed by CAV proximity
Or other [C]onnected vehicle

Other vehicle 
Following human dynamics
(car following model)

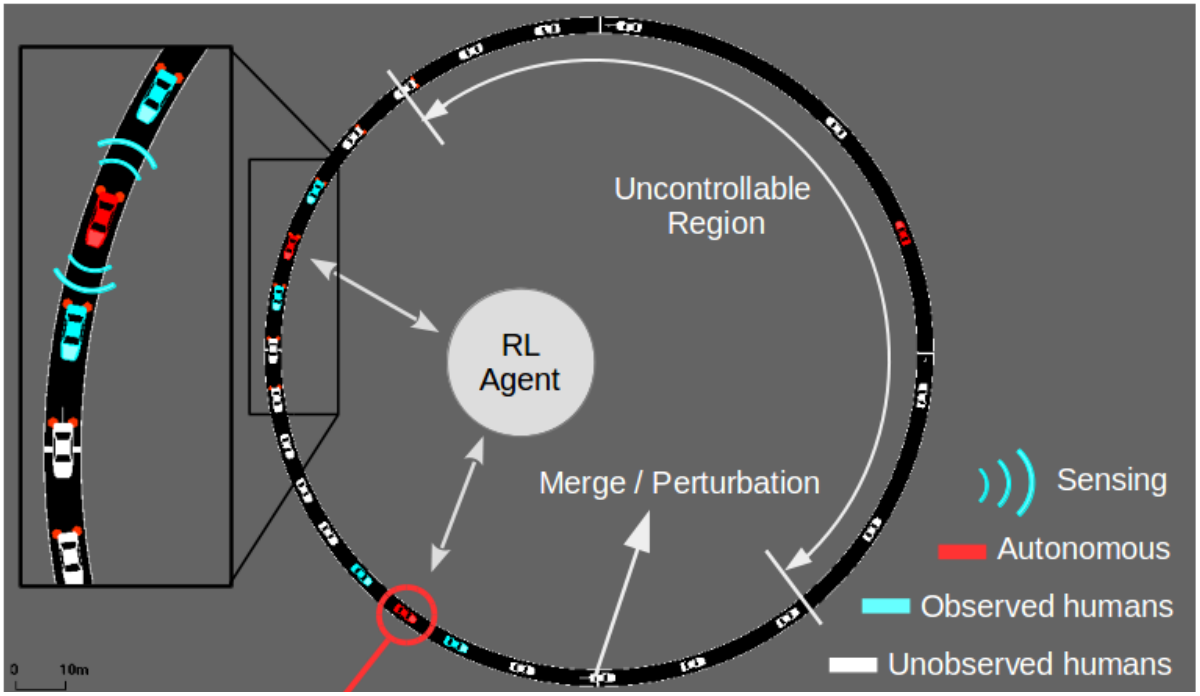

Building lego blocks

Deep-RL lego blocks = science fiction of model based approaches

Setup: 1 AV, 41 human

Experiment

- **Goal**: Maximize average velocity
- **Observation**: following headways, velocity
- **Action**: acceleration and lane change

Results

- **Insight**: A single AV can stabilize multiple lanes of traffic
- **Emergent traffic break**

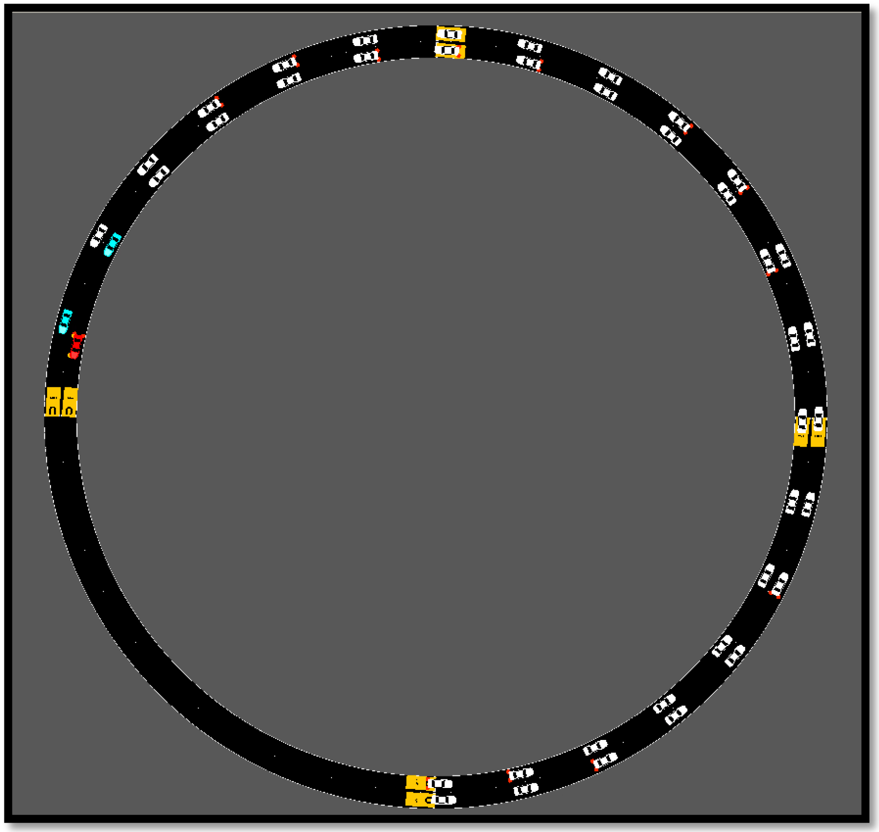

Building lego blocks

Deep-RL lego blocks = science fiction of model based approaches

Setup: 1 AV, 41 human

Experiment

- **Goal**: Maximize average velocity
- **Observation**: following headways, velocity
- **Action**: acceleration and lane change

Results

- **Insight**: A single AV can stabilize multiple lanes of traffic
- **Emergent traffic break**

Intersection control

Moving towards automated intersections

Queuing theory
Reservation systems
Model predictive control

A multiagent approach to autonomous intersection management.
Dresner, Stone. JAIR, 2008.

Polling-systems-based control of high-performance provably-safe autonomous intersections.
Miculescu, Karaman. CDC, 2014.

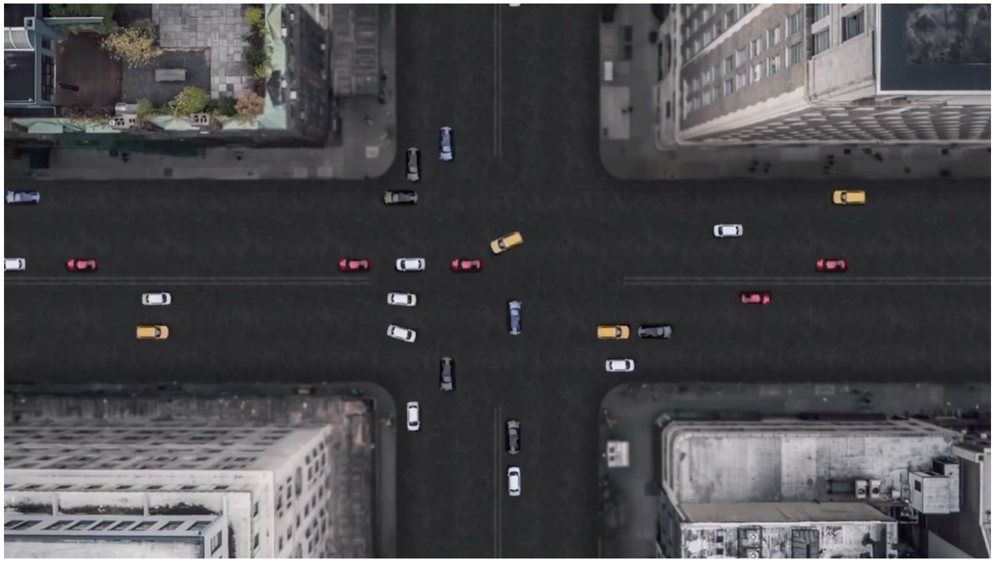

What if even one of these vehicles is not automated?

Intersection control

Moving towards automated intersections

Setting: 0 AV, 14 human

Dynamics: cascaded nonlinear systems with **right-of-way dynamics model**

No autonomy

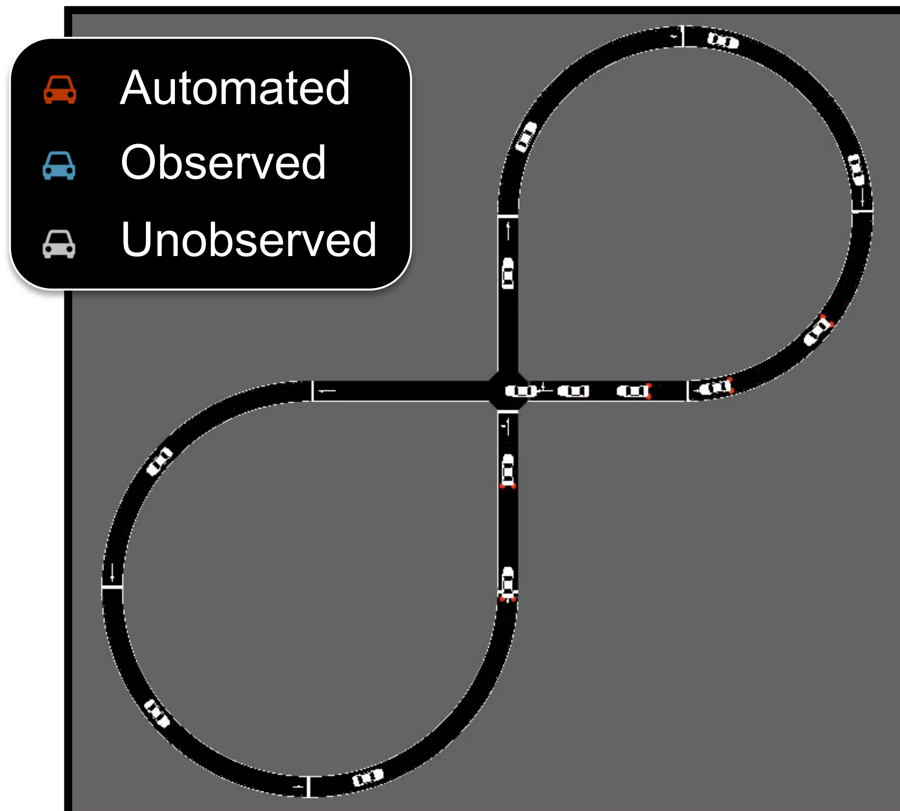

Intersection control

Moving towards automated intersections

Setting: 1 AV, 13 human

Experiment:

- Goal: maximize average velocity
- Observation: fully observed
- Action: acceleration

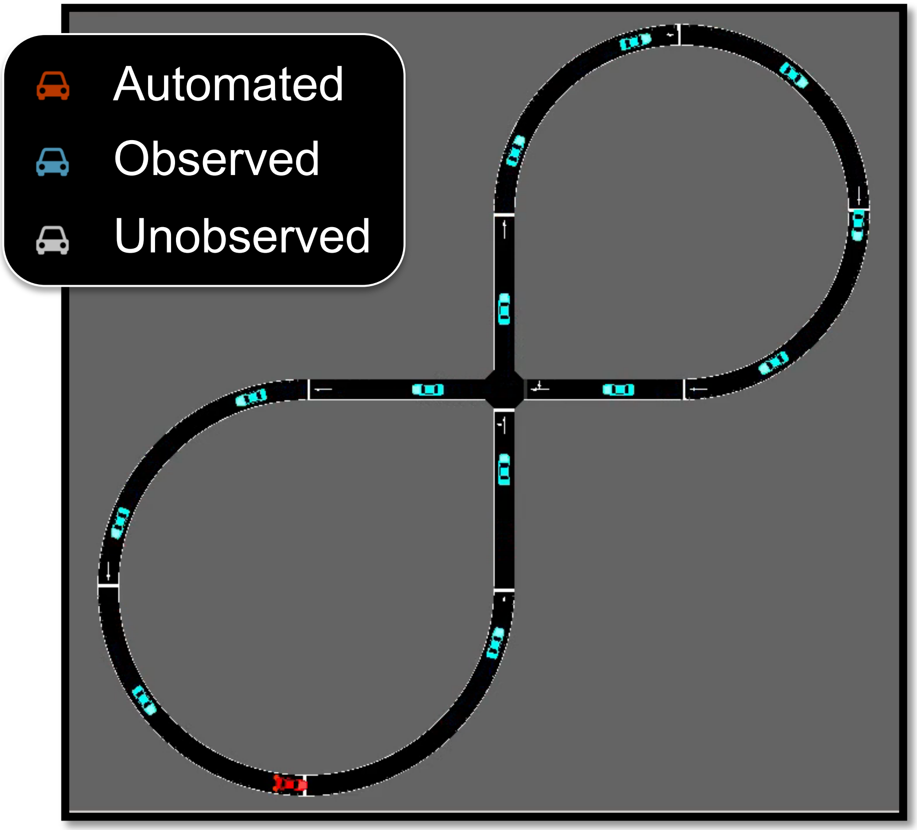

Results

- **Emergent mixed-autonomy platoon**
- **Insight**: A single AV can slow or stop ALL vehicles behind it
- **1 AV: +60% average velocity**

Intersection control

For time space diagram afficiona

Setting: 1 AV, 13 human

Experiment:

- Goal: maximize average velocity
- Observation: fully observed
- Action: acceleration

Results

- **Emergent mixed-autonomy platoon**
- **Insight**: A single AV can slow or stop ALL vehicles behind it
- **1 AV: +60% average velocity**
- **14 AVs: +170% average velocity**

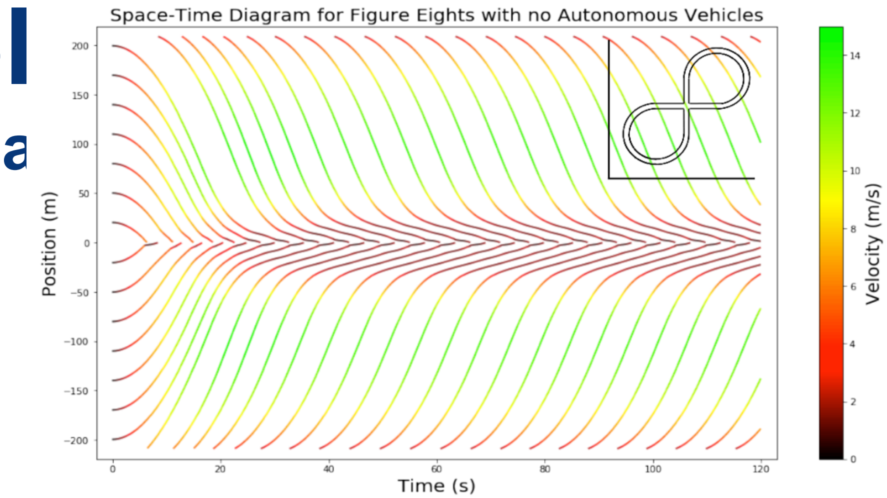

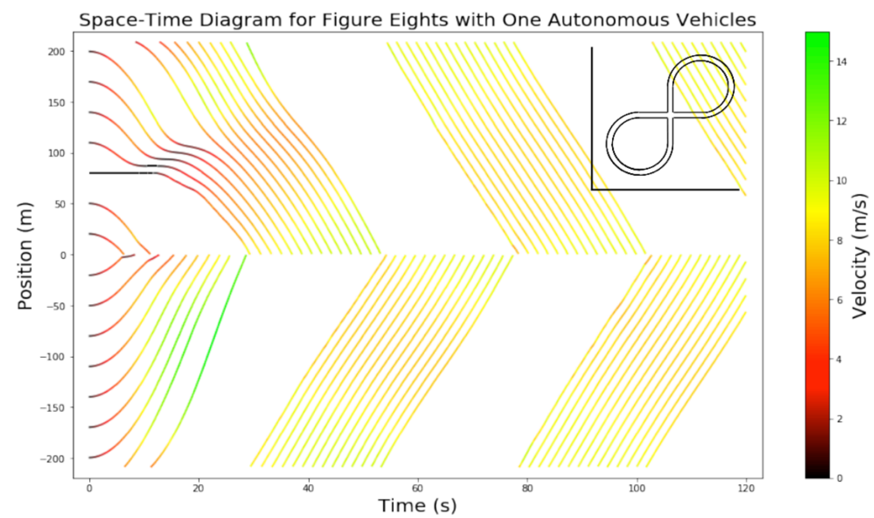

Merge control

Moving towards automated merges

Impacts: 40% of highway congestion

Setting: 0 AV, 17 human

Dynamics:

cascaded nonlinear systems with right-of-way dynamics model

Longitudinal control algorithm for automated vehicle merging.
Lu, Hedrick. IJC, 2003.

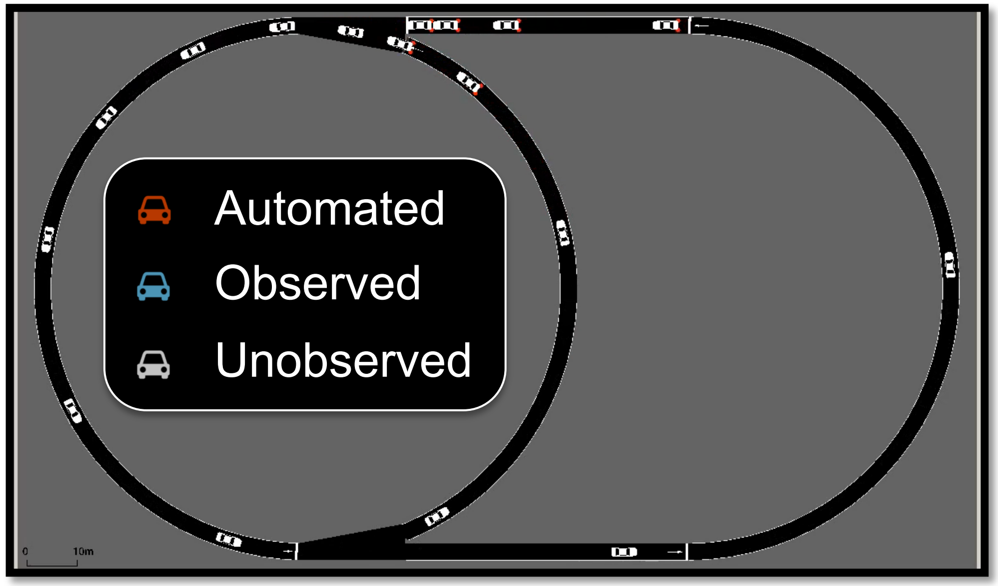

The impacts of a communication based merging assistant on traffic flows of manual and equipped vehicles at an on-ramp using traffic flow simulation.
Pueboobpaphan, et al. IEEE ITSC, 2010.

Merge control

Moving towards automated merges

Setting: 1 AV, 16 human

Experiment:

- **Goal**: maximize average velocity
- **Observation:** Local and merging vehicles, statistics, e.g. queue length
- **Action:** acceleration

Results

- **Emergent mixed-autonomy cooperative merge**
- 1 AV: **+142%** average velocity (6.3 m/s)
- 0 AV: 2.6 m/s

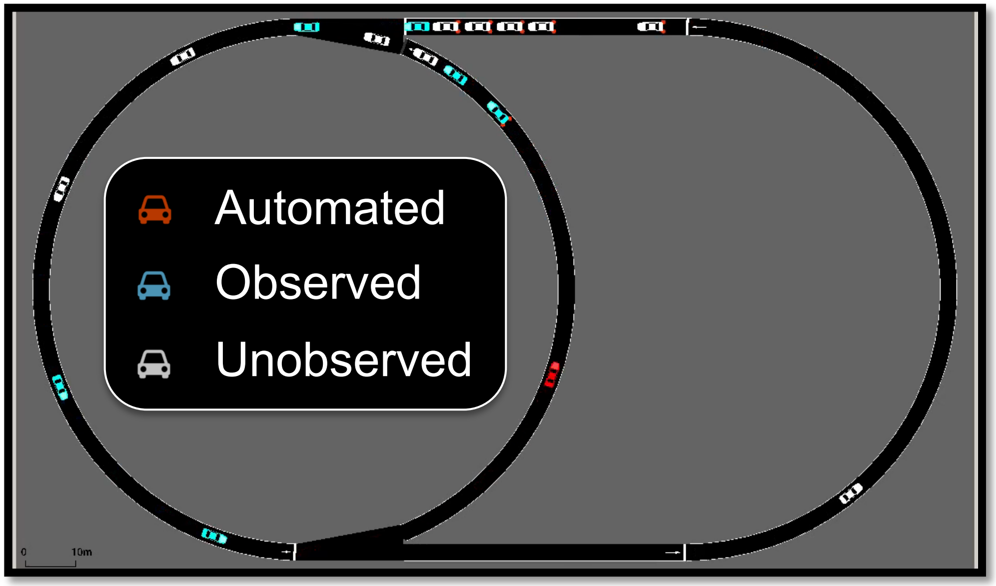

Transfer learning

Moving towards automated merges

Setting: p% CAV penetration

Experiment:

- **Goal**: maximize average velocity
- **Observations:** 1 vehicle ahead/behind
- **Actions:** acceleration

Transfer Learning

- Initial training on **ring road** with periodically **induced perturbations**
- Resultant policy extracted and tested on straight highway with merge.

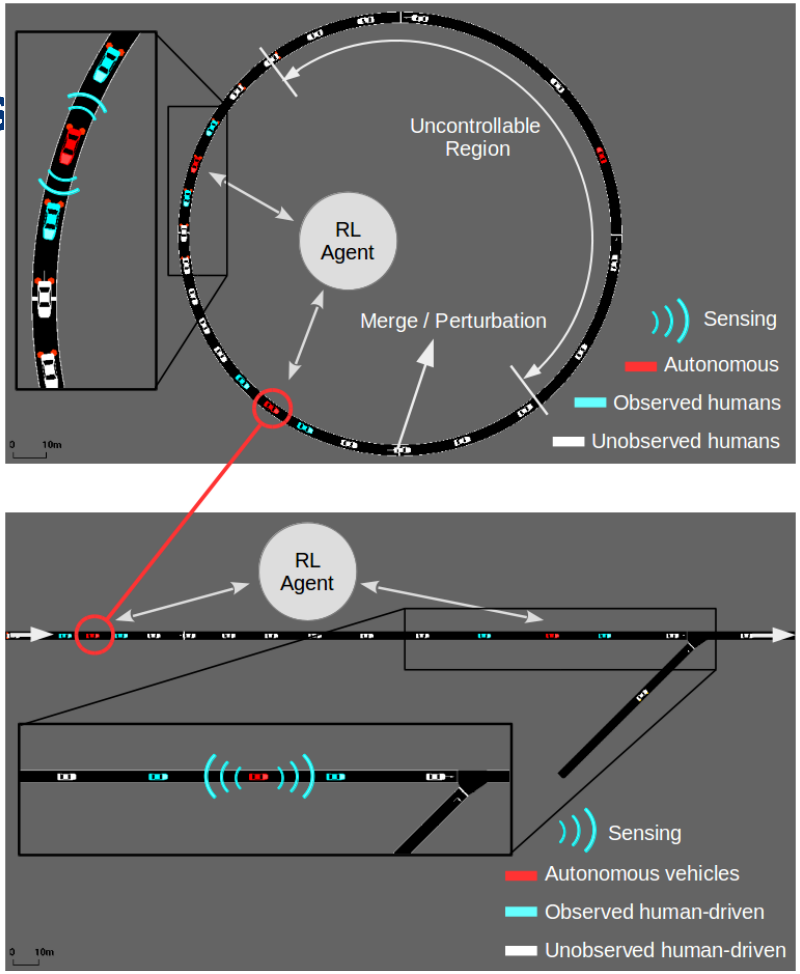

Transfer learning

Moving towards automated merges

Transfer Learning Results:

- Ring road policy initially outperforms human-driven dynamics
- **Significance**: Control strategies derived from simplified closed network geometries are somewhat transferable to open network problems.

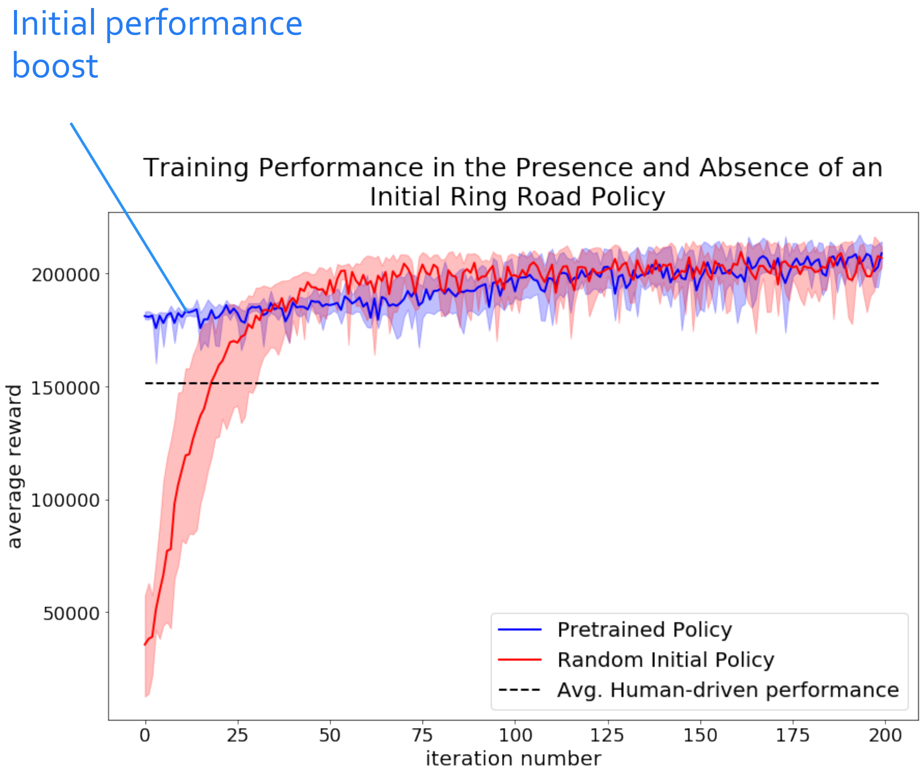

Transfer learning

Moving towards automated merges

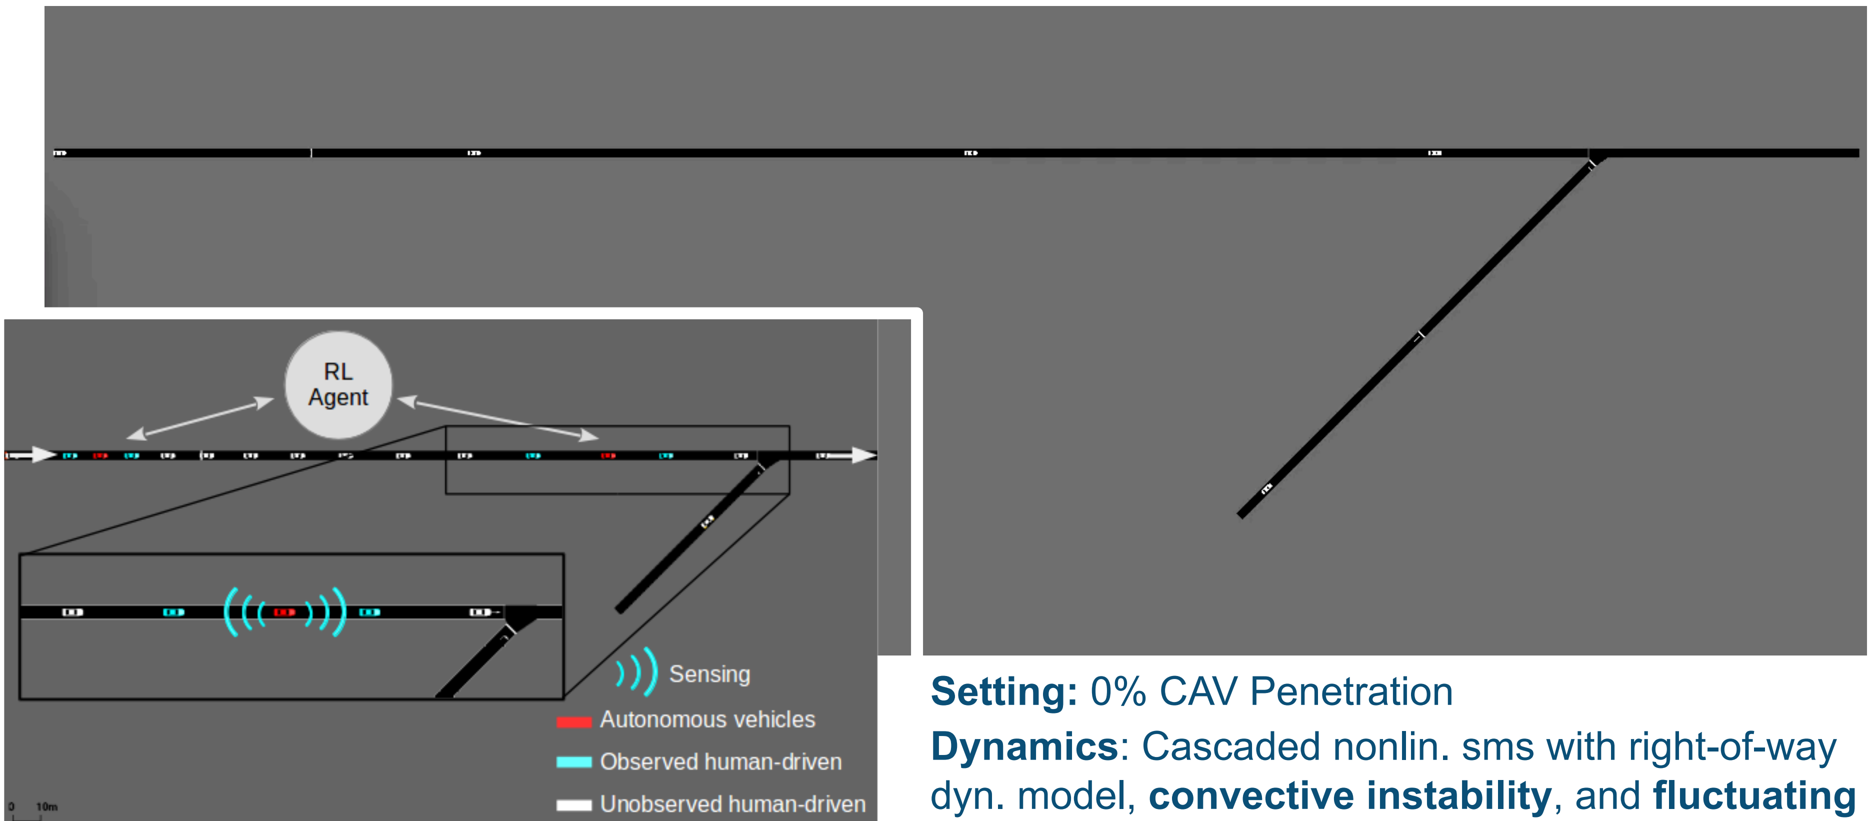

Setting: 0% CAV Penetration

Dynamics: Cascaded nonlin. sms with right-of-way dyn. model, **convective instability**, and **fluctuating densities**

Transfer learning

Moving towards automated merges

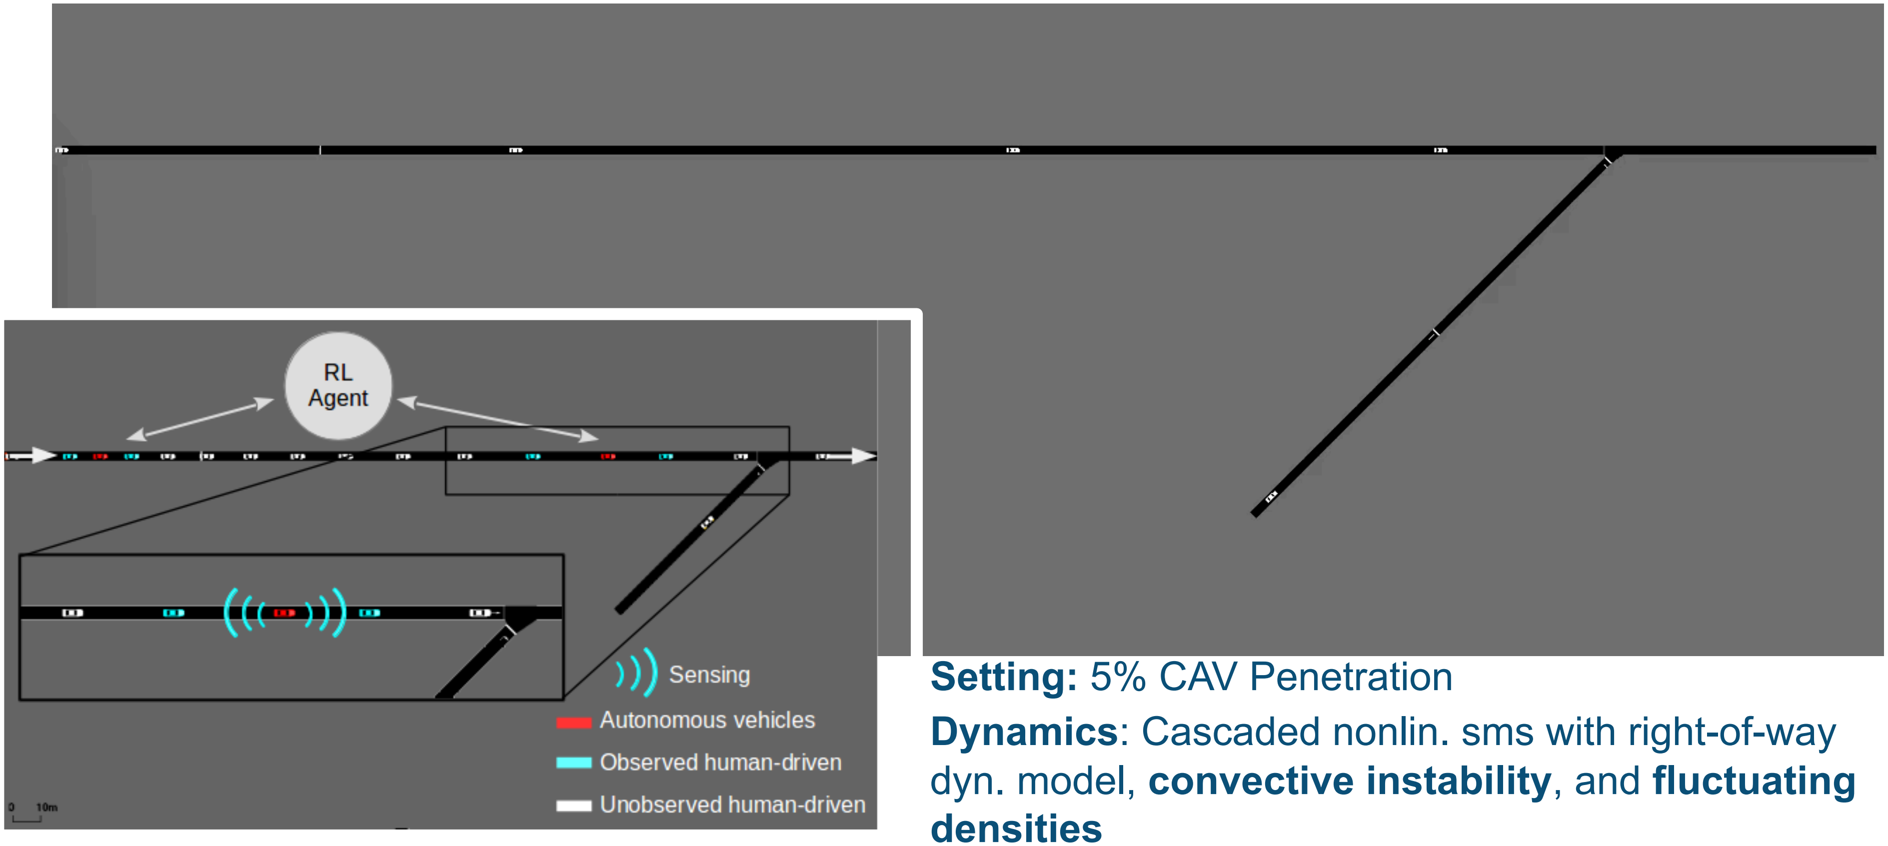

Setting: 5% CAV Penetration

Dynamics: Cascaded nonlin. sms with right-of-way dyn. model, **convective instability**, and **fluctuating densities**

Transfer learning

This slide is also dedicated to the time-space diagram afficionados

6.7% improvement in throughput

Short story: deep-RL just learned to create gaps with forward waves

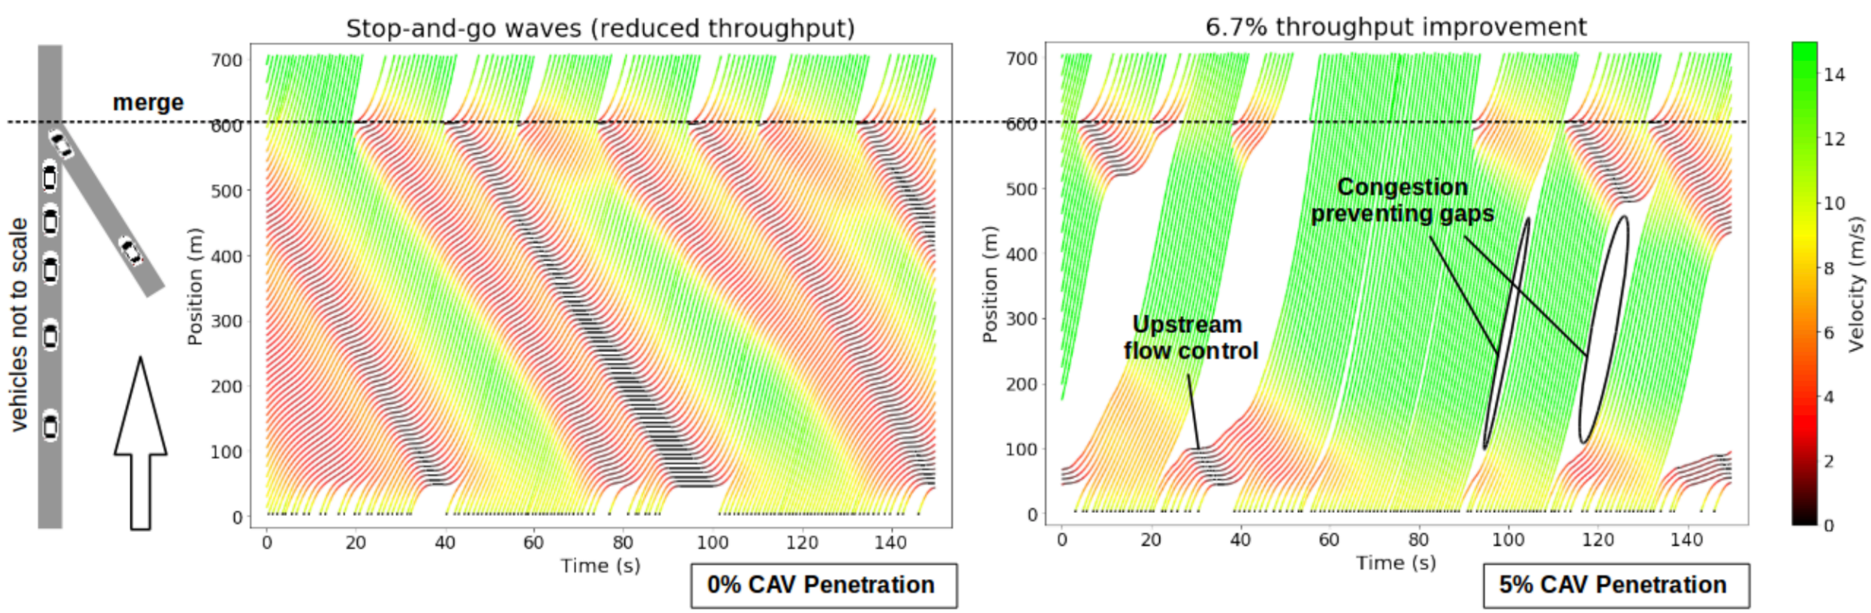

Penetration studies

How many CAVs does it take to smooth traffic?

A. Kreidieh

E. Vinitsky

0% AV (baseline)

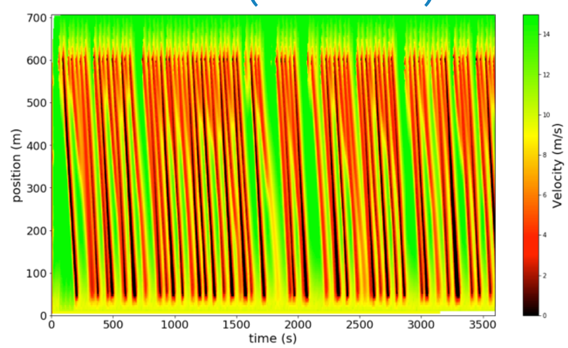

2.5% AV (0% throughput)

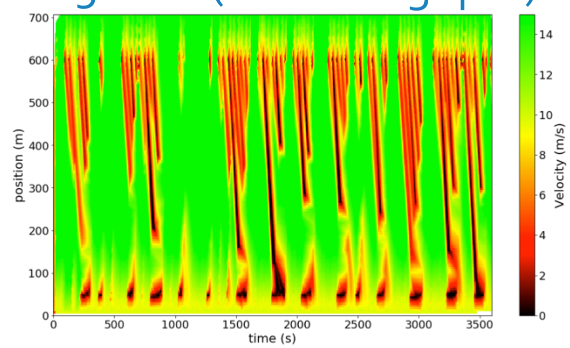

5% AV (+6.7% throughput)

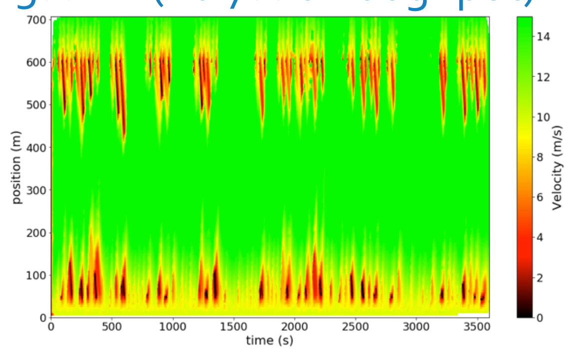

10% AV (+13% throughput)

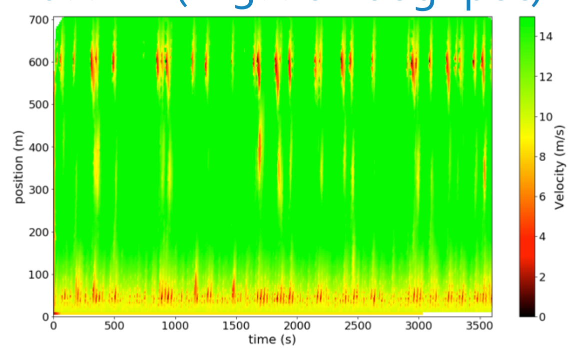

TRAFFIC

Just a few self driving cars on the road could improve traffic for everyone

Jennifer Olney

Friday, November 16, 2018 10:45PM

Science

Watch just a few self-driving cars stop traffic jams

By Matthew Hutson | Nov. 16, 2018 , 3:55 PM

Anyone can start a traffic jam—all it takes is tapping on your brakes. The driver behind you will brake, as will the next driver, starting a shock wave of stop-and-go reactions that can travel backward for kilometers. Now, scientists have shown that a few self-driving cars can prevent such jams—and in some cases double the average speed of surrounding vehicles.

The researchers used a video game–style interface to control simulated cars on made-up roadways. In one scenario, the cars drove around a figure eight with a central intersection. In others, one or several lanes of traffic merged, or the cars traversed a Manhattan-like city grid with traffic lights at each crossing. The team looked at various ratios of self-driving cars mixed with regular cars that simulated typical human driving.

In each scenario, the researchers tested four algorithms that used reinforcement learning—a type of artificial intelligence (AI) that learns skills through trial and error. In the figure eight and merging scenarios, a central algorithm controlled all self-driving cars, experimenting by changing their patterns of acceleration and braking. In the Manhattan scenario, the AI-controlled traffic lights instead of cars.

In real life

Oscillations exist naturally in highway traffic

Flow smoothing on I210

How many CAVs does it take to smooth traffic?

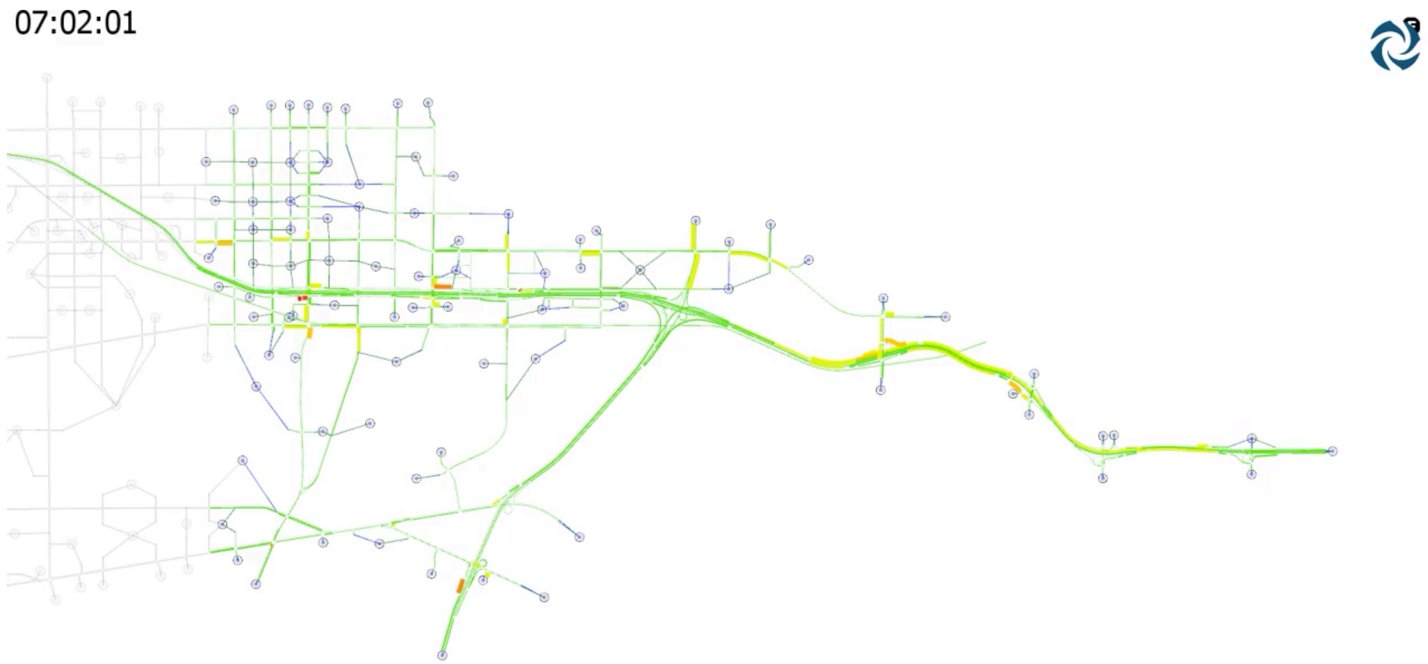

Flow smoothing on I210

How many CAVs does it take to smooth traffic?

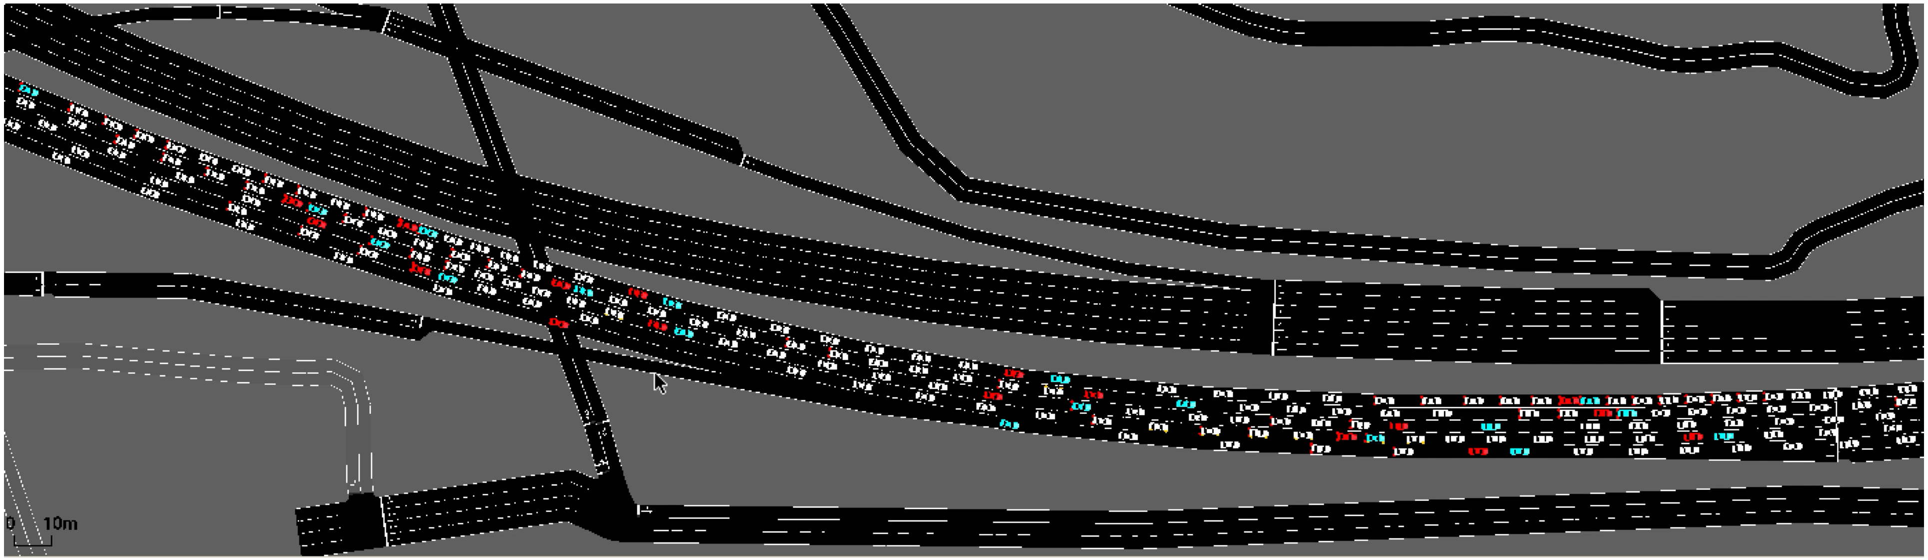

Flow smoothing on I210

How many CAVs does it take to smooth traffic?

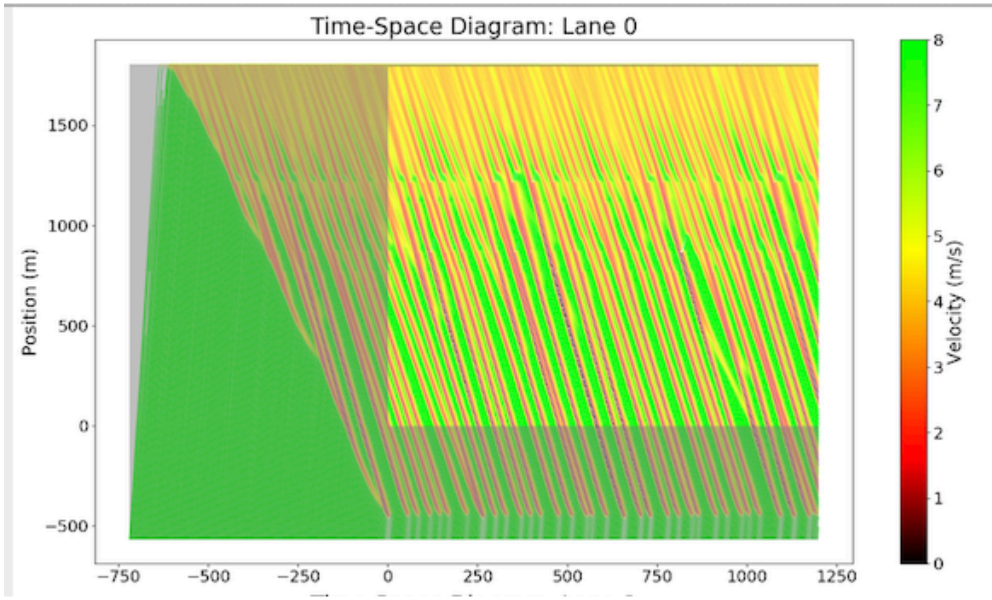

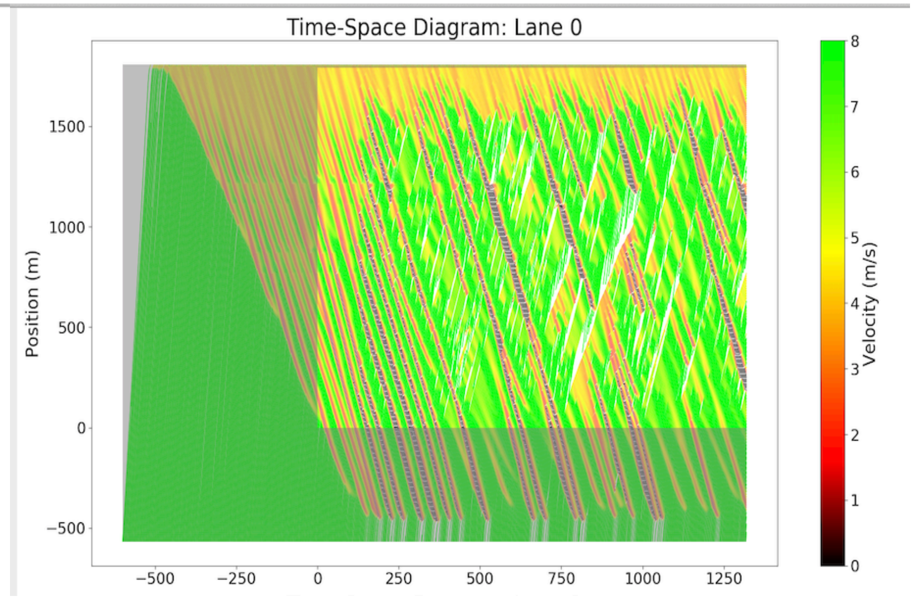

Flow smoothing on I210

How many CAVs does it take to smooth traffic?

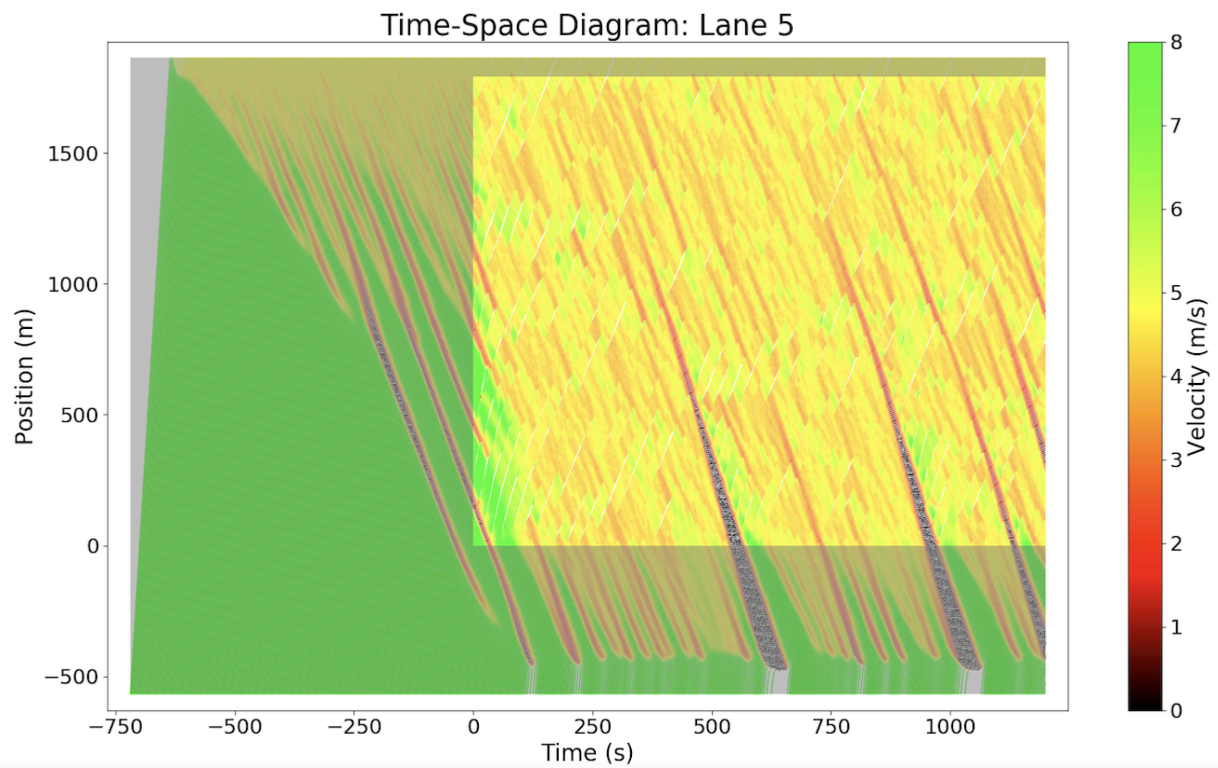

Flow smoothing on I210

How many CAVs does it take to smooth traffic?

$$\text{reward } = -\text{energy consumed}$$

Problem
The agent can optimize that reward by not moving

Need an incentive to move forward:

- Penalty for staying too long in the network
- Additional state (time since entered)

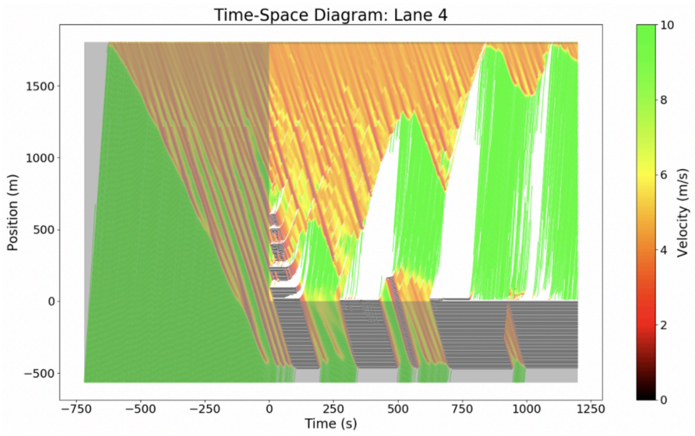

Bridge metering

Lagrangian metering

NextGen infrastructure

Can we remove the metering light and can we replace it with CAVs?

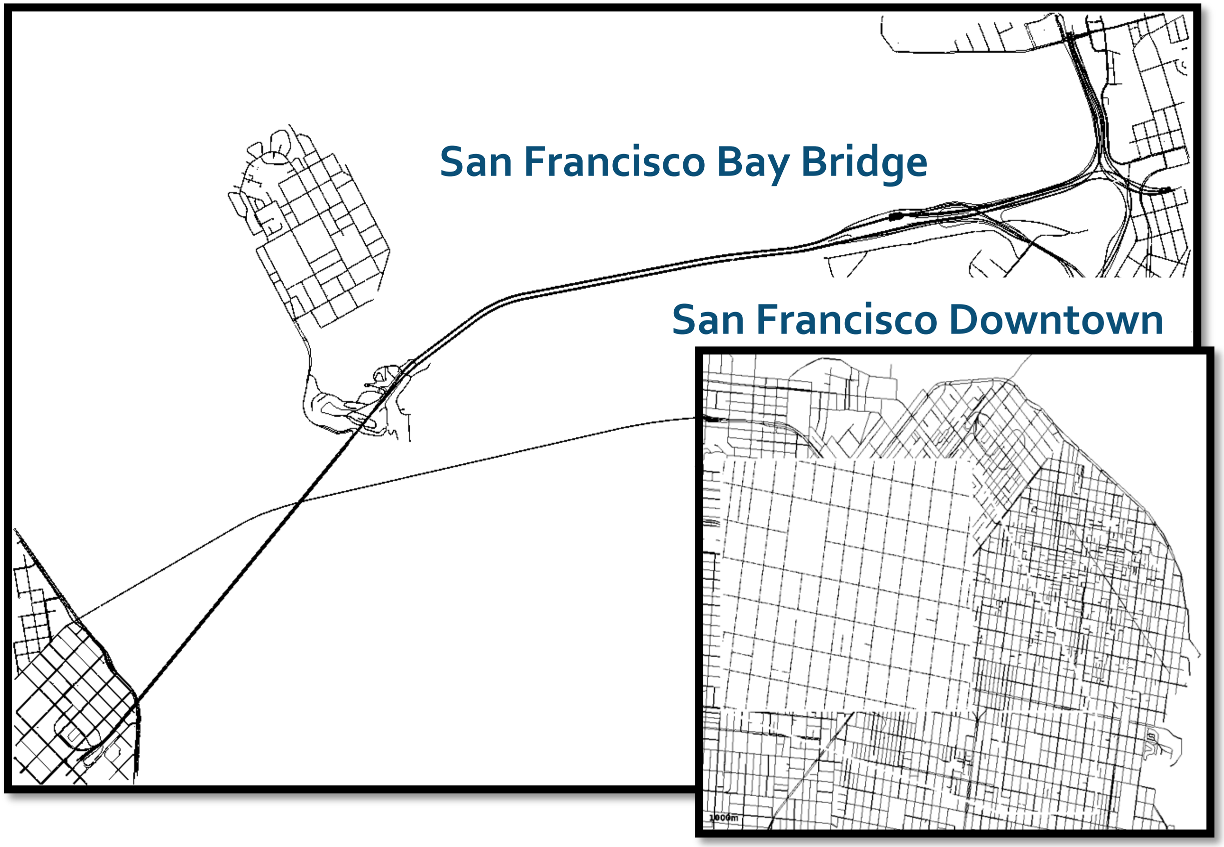

Bridge metering ([not] waiting for Godot)

Lagrangian metering: 33% improvement (throughput)

Setting:

- 10% CAV Penetration
- Four lanes -> Two lanes -> One

Dynamics:

Cascaded nonlinear systems with right-of-way dynamics model, **merge conflicts**, and **excessive, fluctuating inflow**

Bridge metering

Lagrangian metering

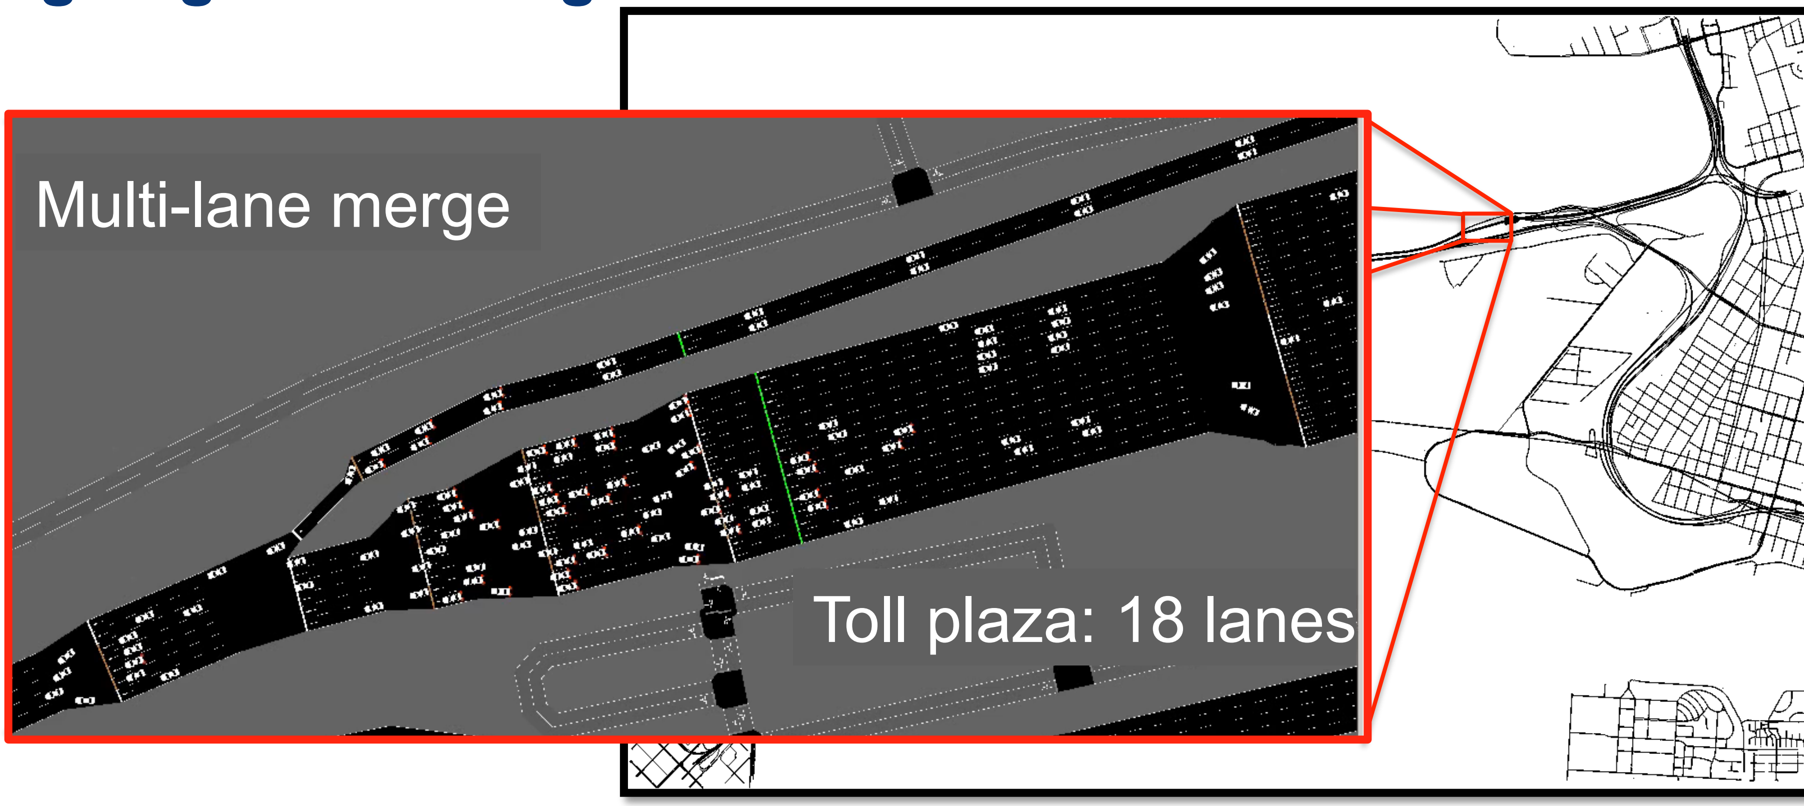

Policy transfer

Left: baseline scenario; right: flow maximization


K. Jang


E. Vinitsky

Policy transfer

Go to our demo booth tomorrow!

F. Wu

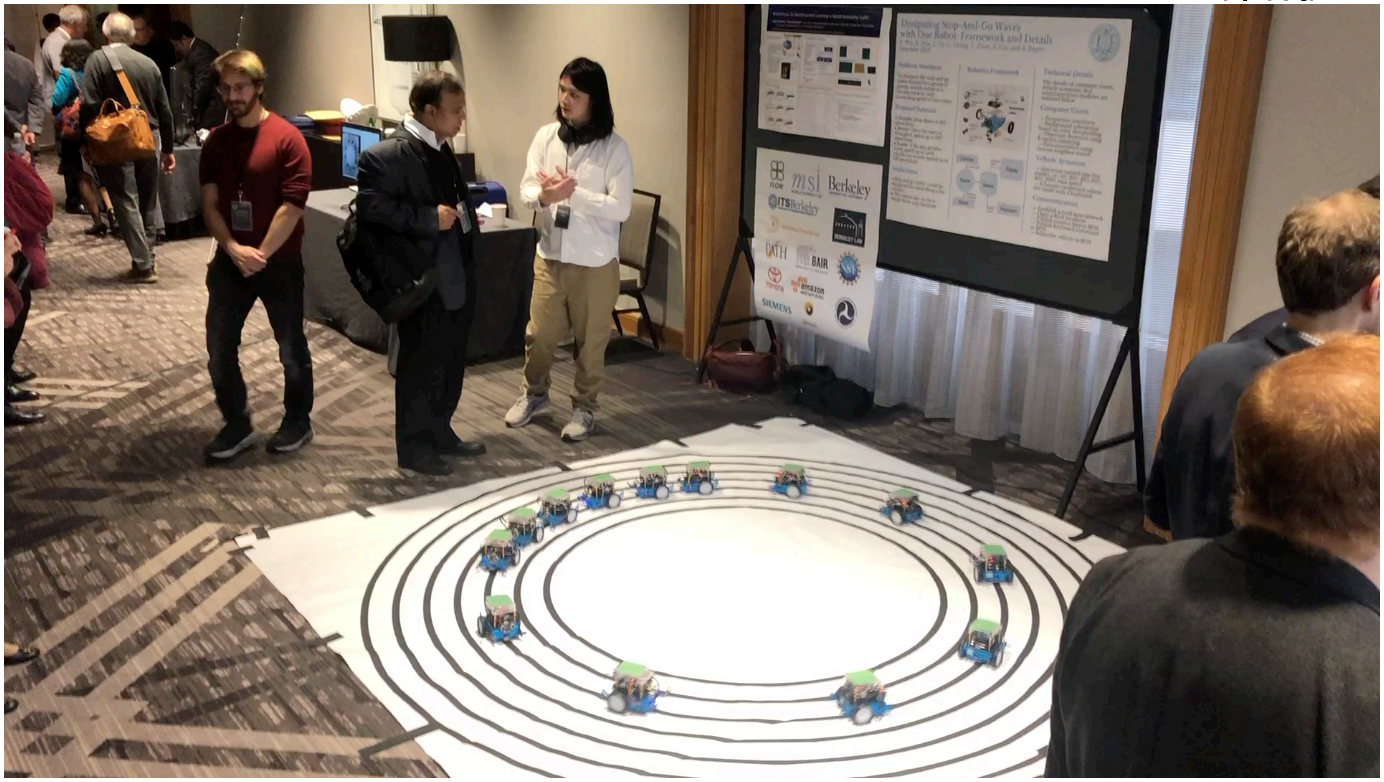

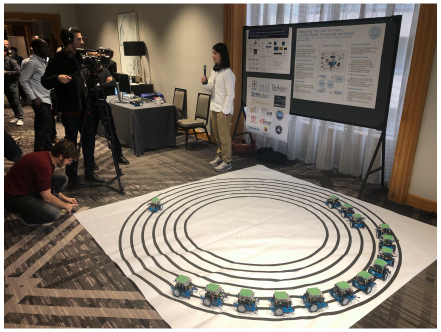

End-to-end pixel learning

Deep RL as a Markov decision process

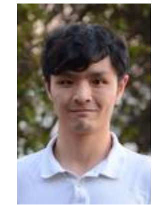

F. Wu

Markov decision process:

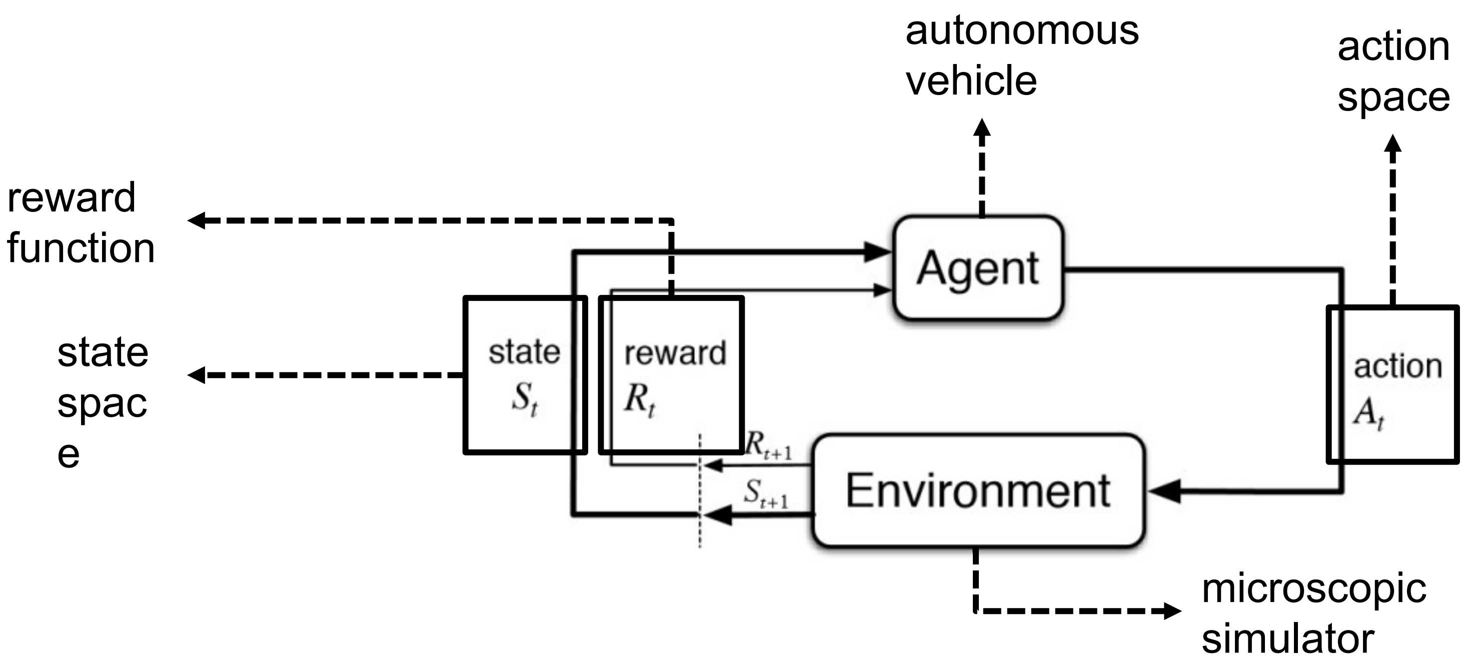

End-to-end pixel learning

State space design

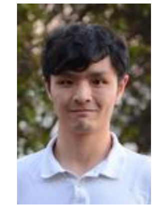

F. Wu

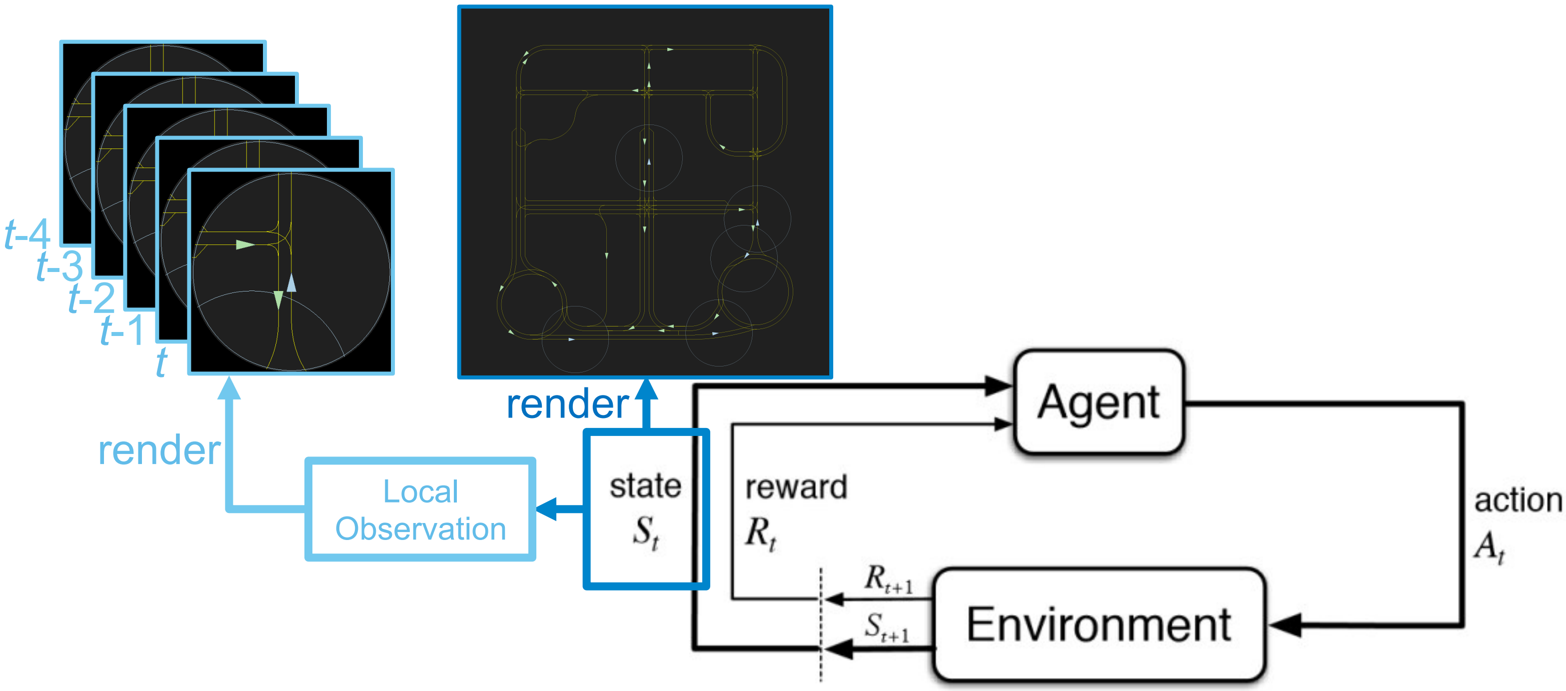

End-to-end pixel learning

Action space design

F. Wu

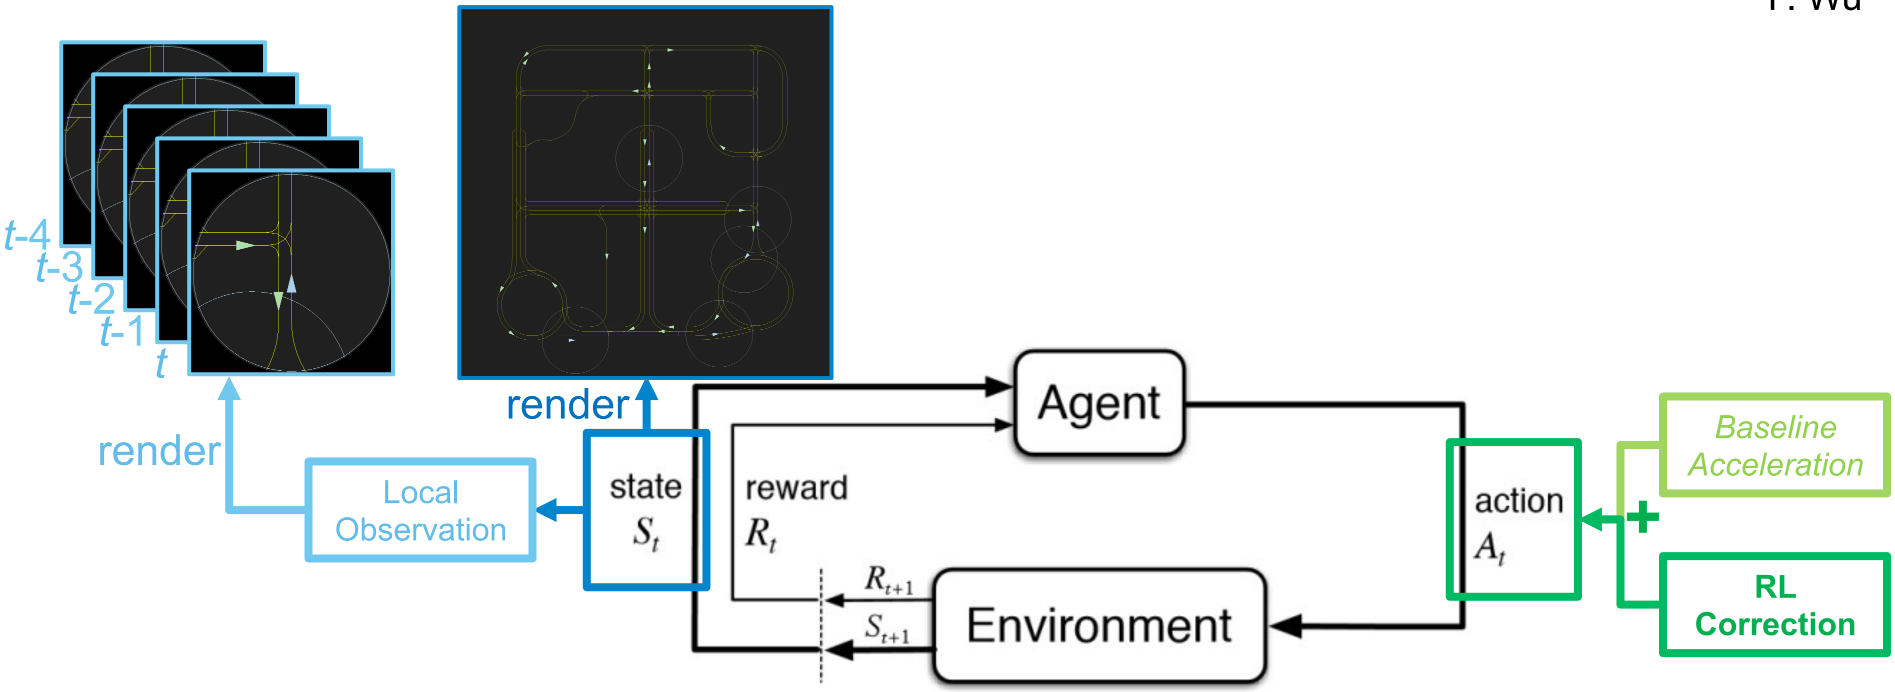

End-to-end pixel learning

Reward function design

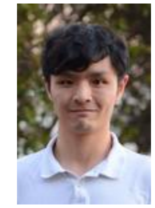
F. Wu

$$r = \frac{1}{v_{\max}}\left(\beta \boxed{v_{\mathrm{avg}}} - (1-\beta)\boxed{v_{\mathrm{std}}}\right)$$

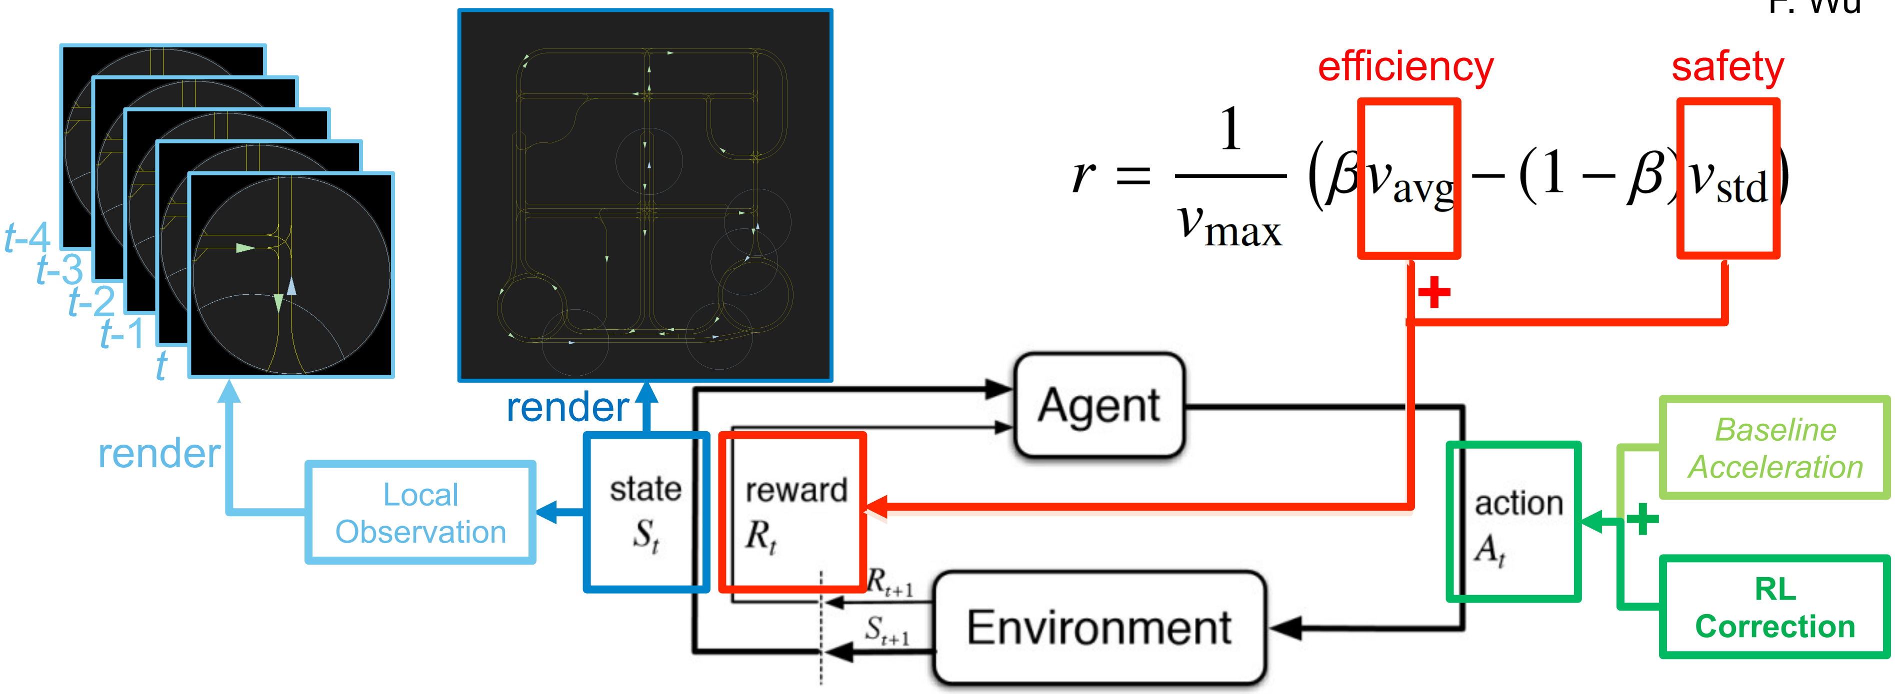

End-to-end pixel learning

Learning on simple environments

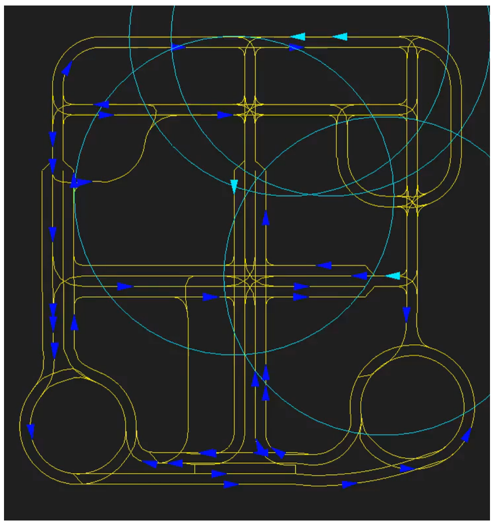

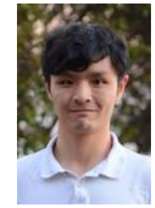

F. Wu

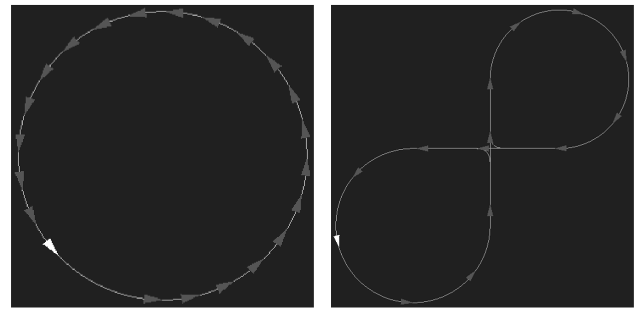

(a) *Circle* environment. (b) *Cross* environment.

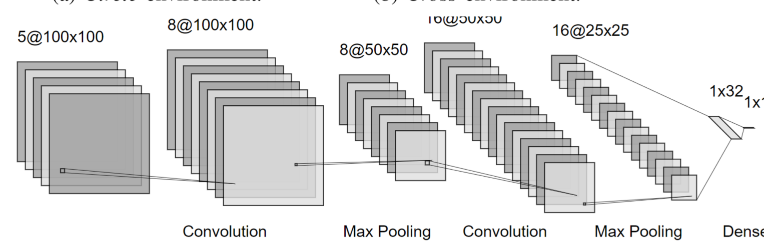

RL augmentation improves human drivers' skills to the level of an optimized AI.

	0%	10%	20%	30%	40%	50%	60%	70%	80%	90%	100%
Circle	210	210	324	336	334	340	342	**343**	340	336	342
Cross	348	336	353	350	510	548	545	**562**	420	447	436

humans humans+AI AI

End-to-end pixel learning

What if one could remove the need for state space access?

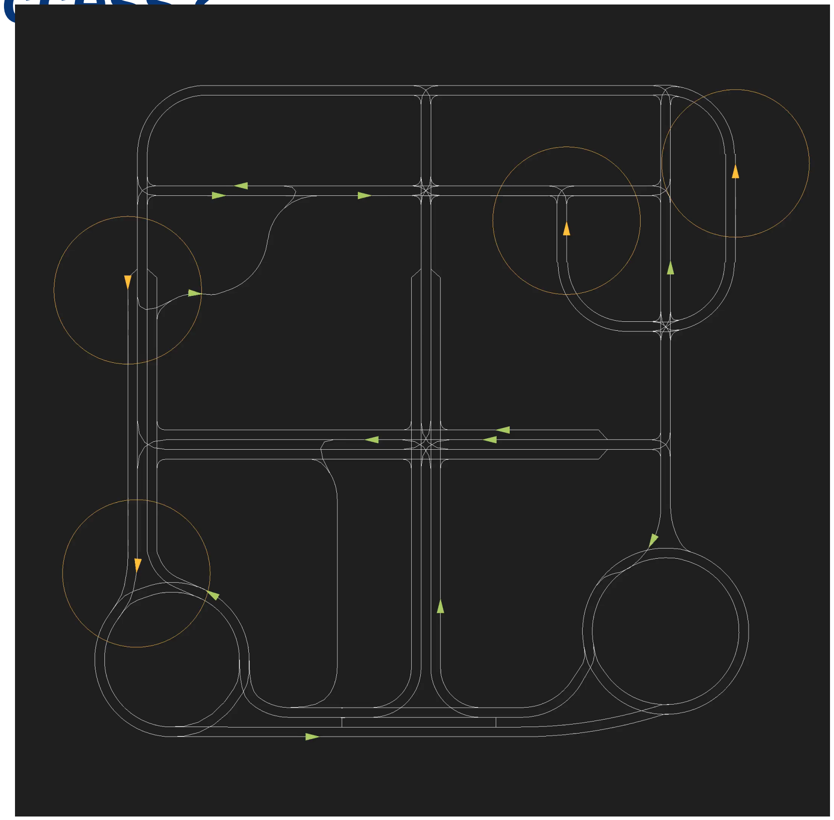

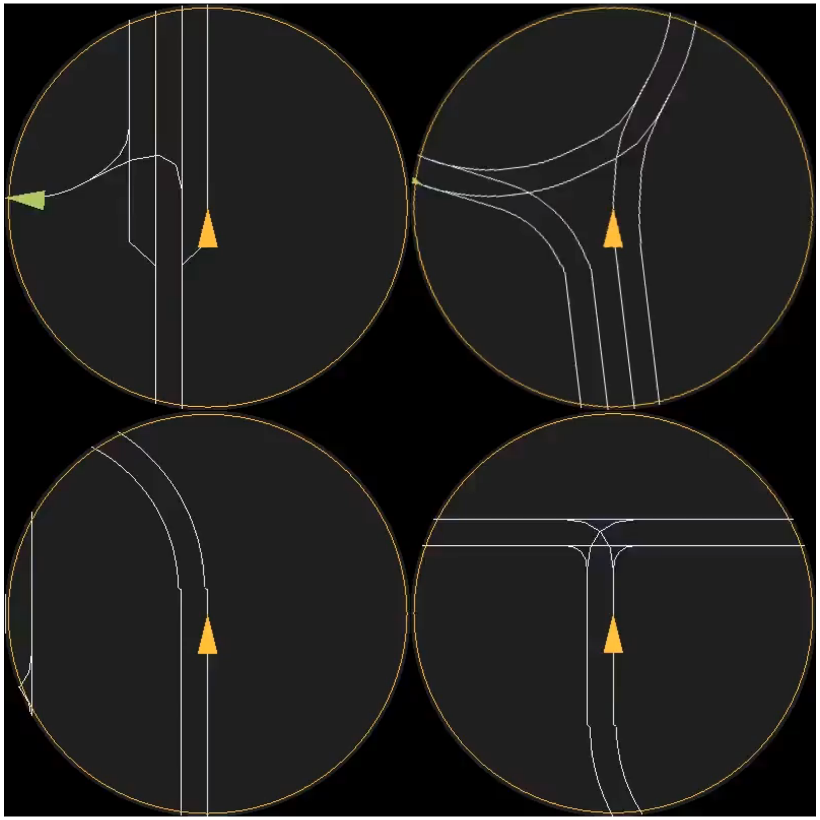

The vision

Eventually linked to dashcam data

F. Wu

Fully centralized:

Pixel learning

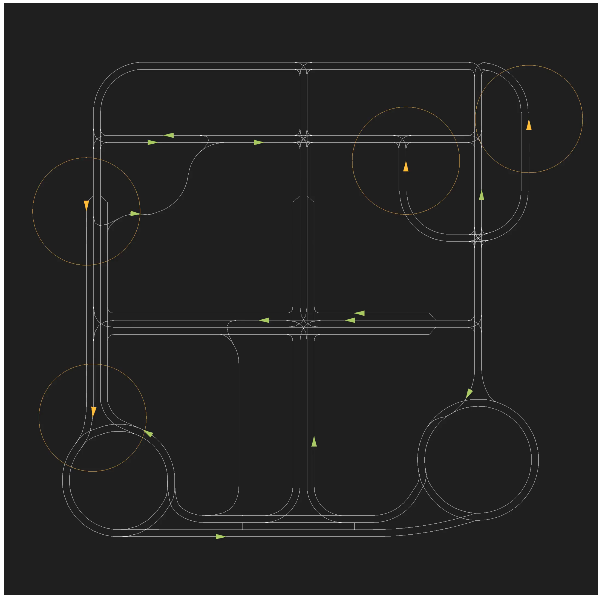

Decentralized:

Pixel learning, multi-agent

Decentralized:

Dashcam segmented data

Who is next?

Traffic management

[redacted]

[redacted]

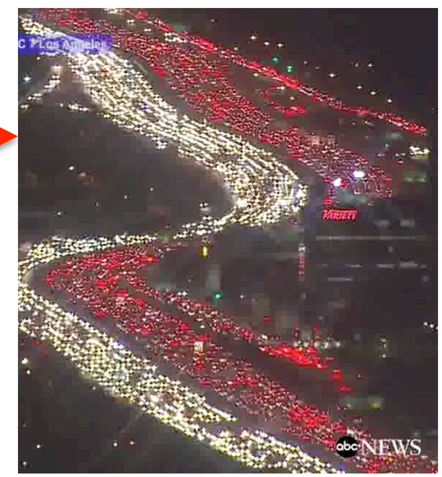

Deep-RL

Basic optimization framework

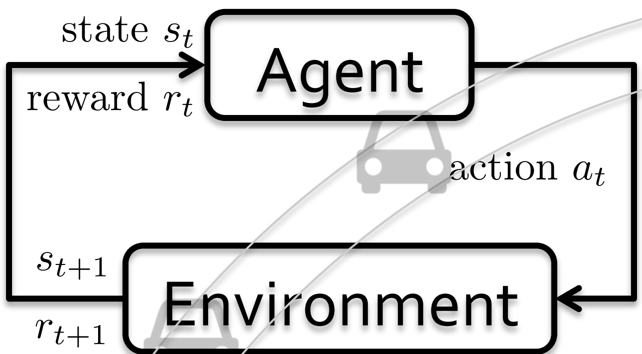

Global rewards

- Average velocity
- Energy consumption
- Travel time
- Safety, comfort

Goal:
learn policy $\pi : S \rightarrow A$
to maximize reward

$$\max_{\theta} \mathbb{E}\left[\sum_{t=0}^{H} r(s_t, a_t) \middle| \pi_\theta\right]$$

Cumulative rewards, returns

Policy parameters (deep neural network)

Flow

Brief presentation of FLOW

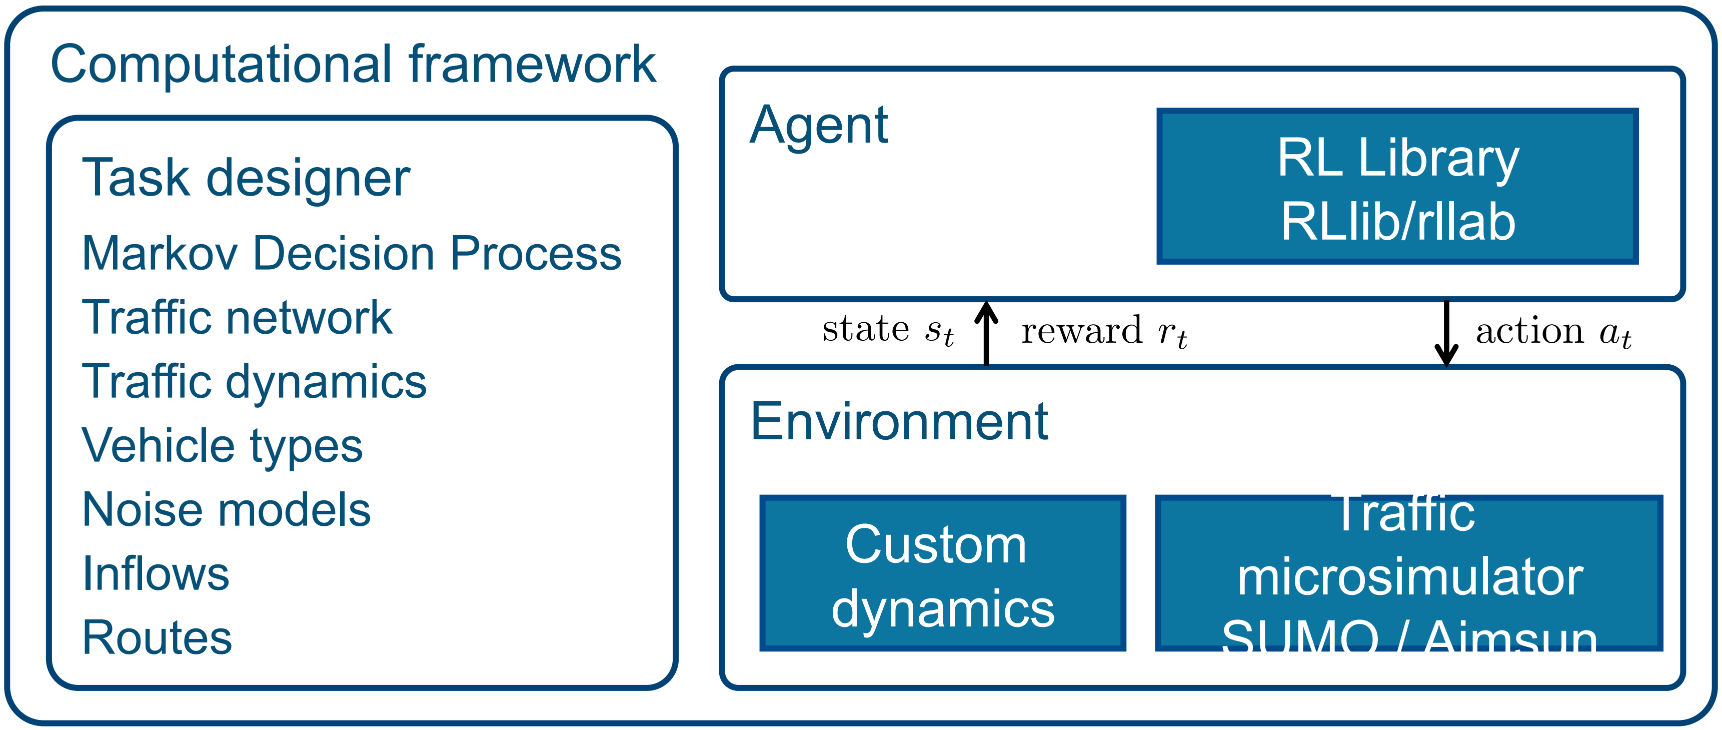

Lego-blocks

Building a library

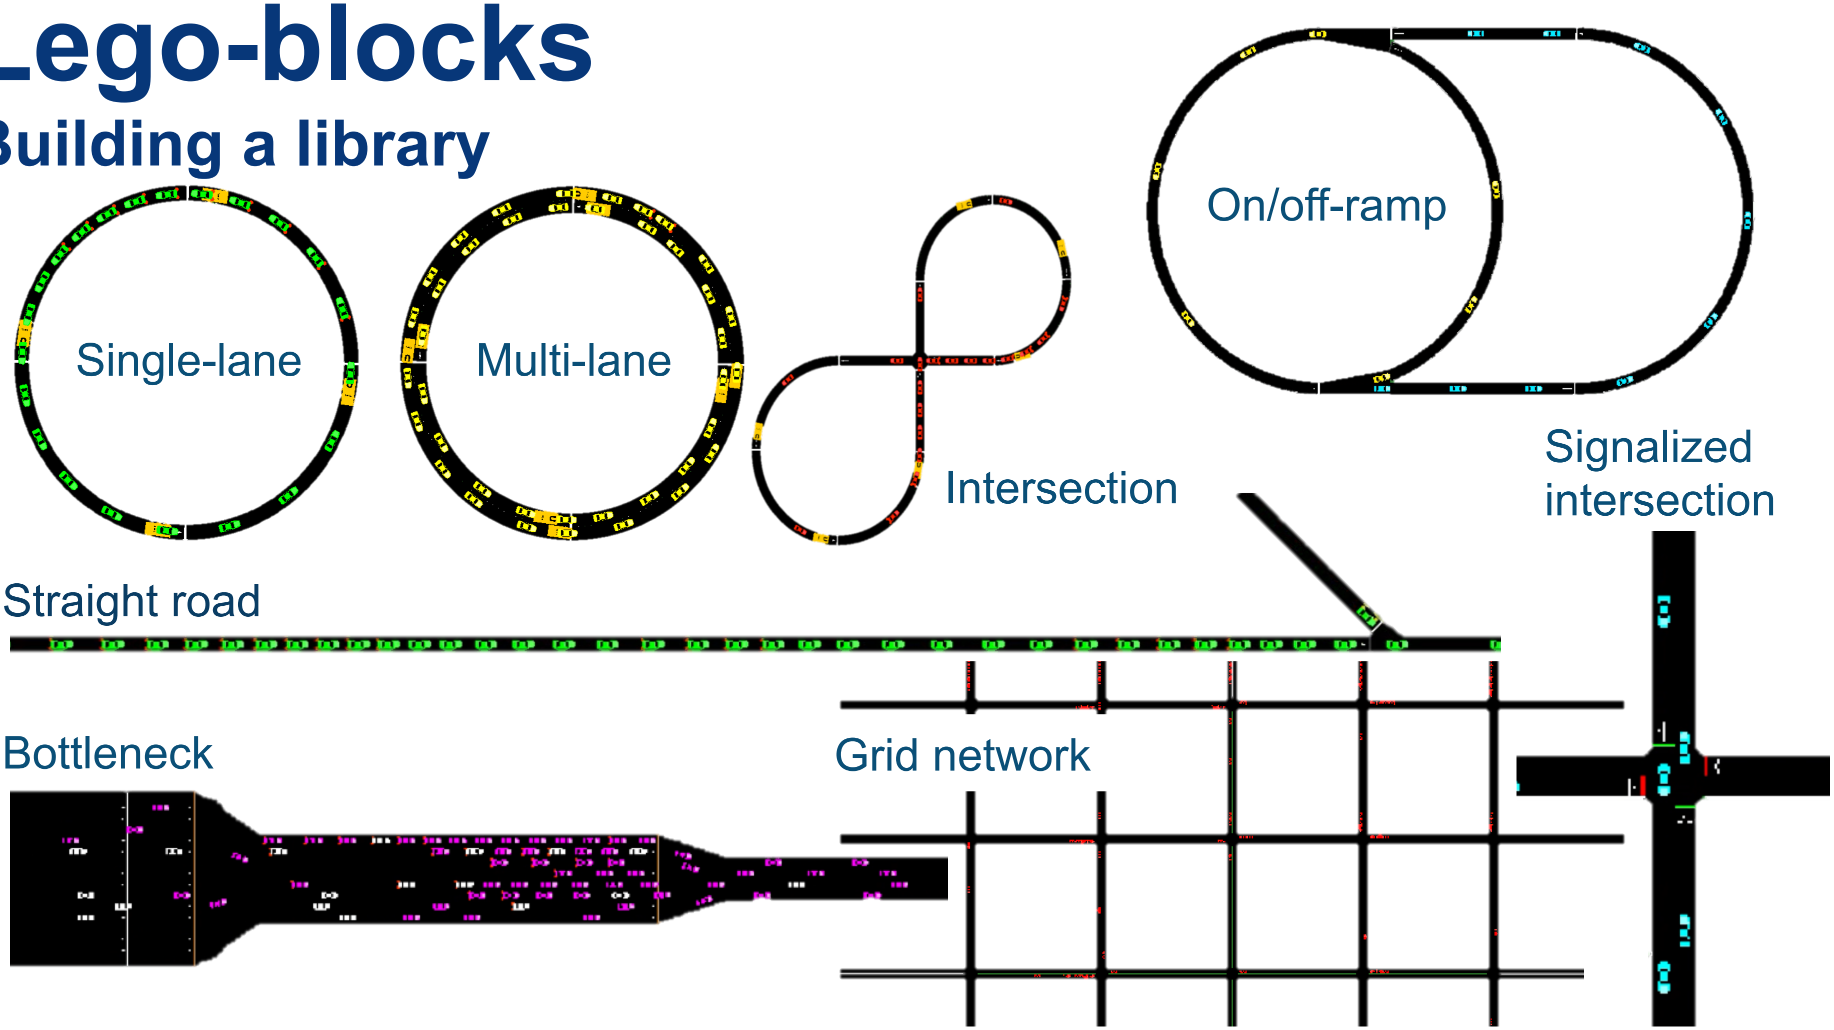

Benchmarks, launched 2018, CORL

https://flow-project.github.io/

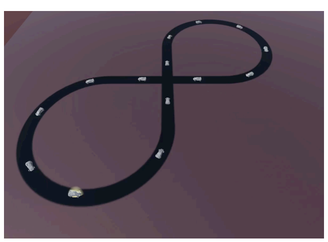

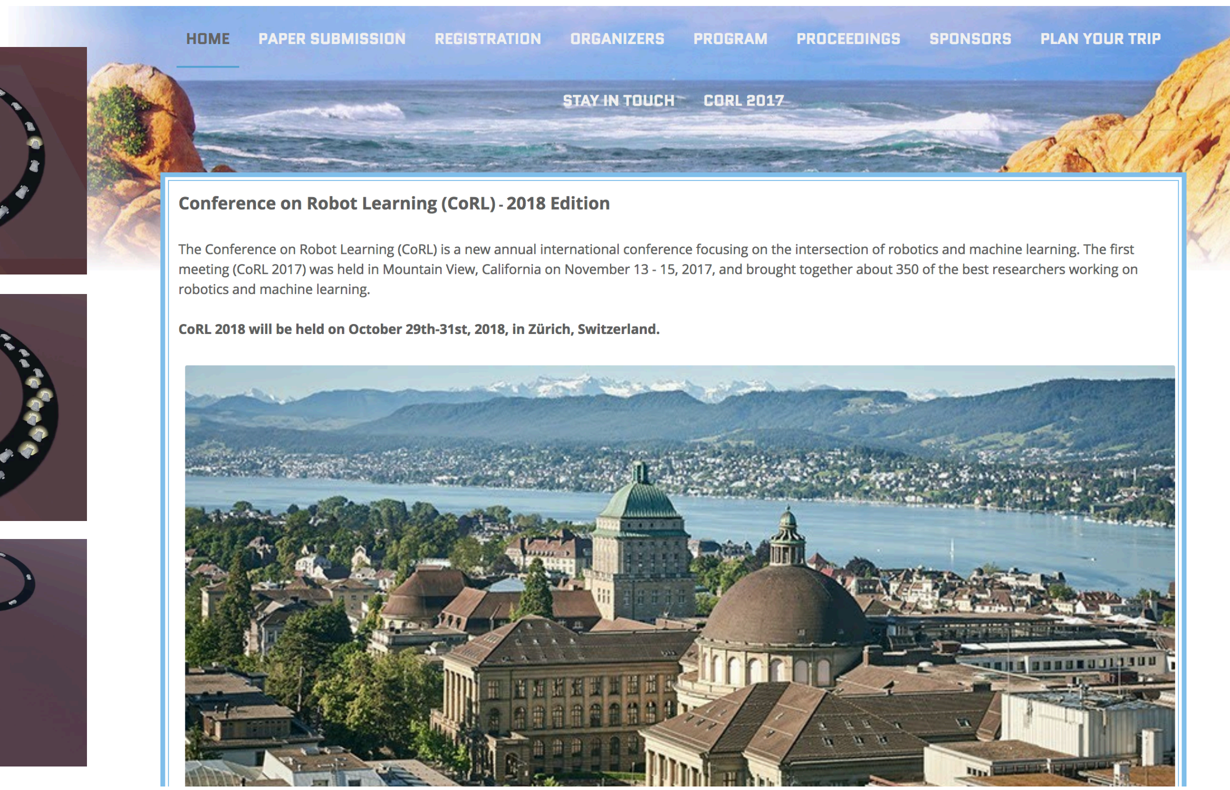

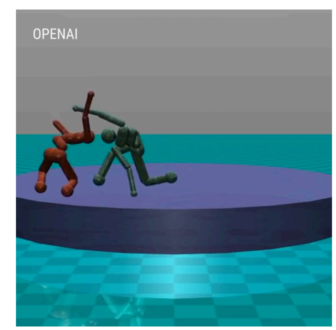

Dashboard CIRCLES

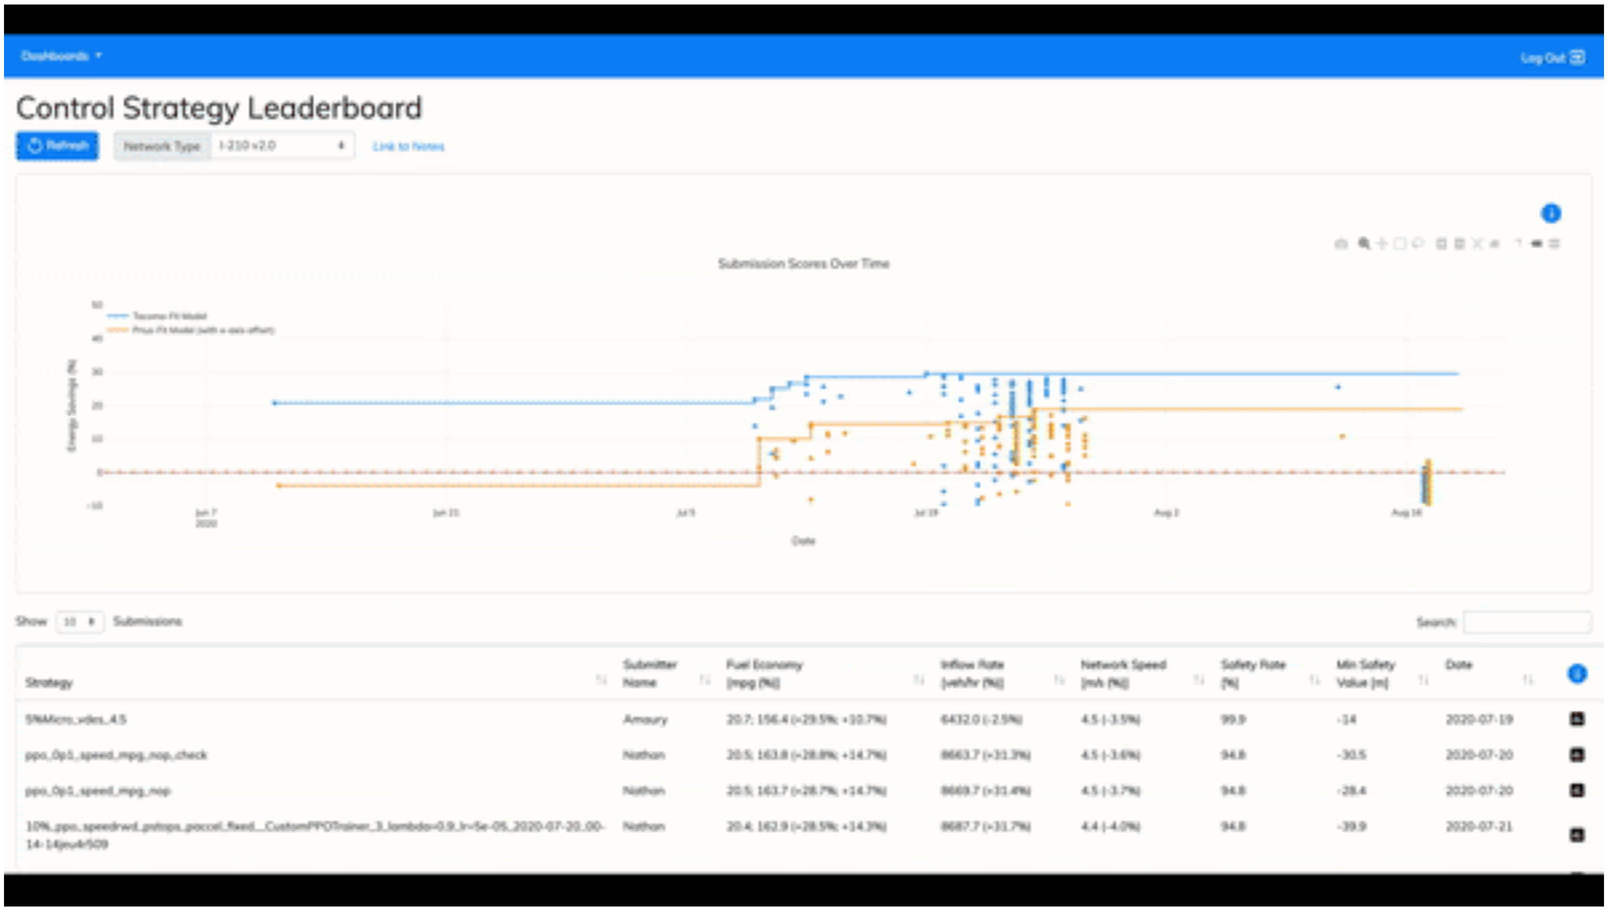

Dashboard CIRCLES

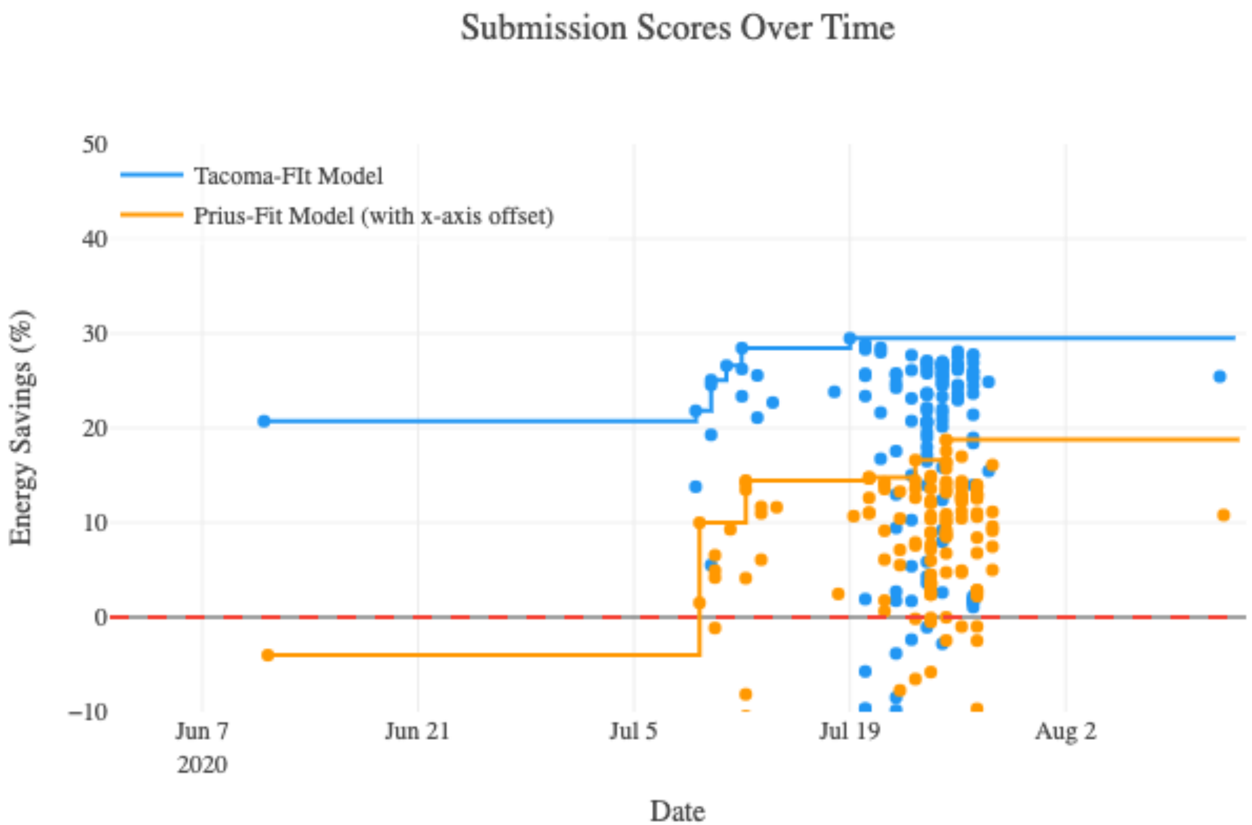

User community

Classes, workshops, tutorials, events, users…

IEEE CDC, Nice, France

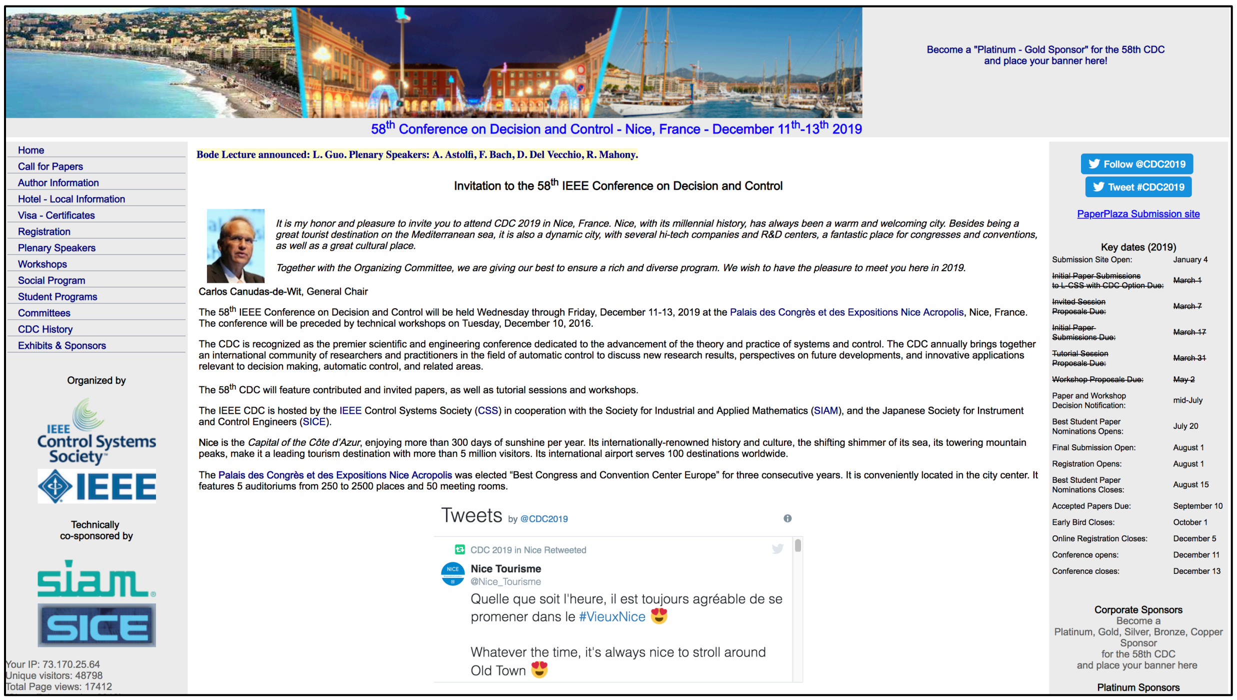

Become a "Platinum - Gold Sponsor" for the 58th CDC
and place your banner here!

58th Conference on Decision and Control - Nice, France - December 11th-13th 2019

- Home
- Call for Papers
- Author Information
- Hotel - Local Information
- Visa - Certificates
- Registration
- Plenary Speakers
- Workshops
- Social Program
- Student Programs
- Committees
- CDC History
- Exhibits & Sponsors

Organized by

IEEE Control Systems Society

IEEE

Technically co-sponsored by

siam

SICE

Your IP: 73.170.25.64
Unique visitors: 48798
Total Page views: 17412

Bode Lecture announced: L. Guo. Plenary Speakers: A. Astolfi, F. Bach, D. Del Vecchio, R. Mahony.

Invitation to the 58th IEEE Conference on Decision and Control

It is my honor and pleasure to invite you to attend CDC 2019 in Nice, France. Nice, with its millennial history, has always been a warm and welcoming city. Besides being a great tourist destination on the Mediterranean sea, it is also a dynamic city, with several hi-tech companies and R&D centers, a fantastic place for congresses and conventions, as well as a great cultural place.

Together with the Organizing Committee, we are giving our best to ensure a rich and diverse program. We wish to have the pleasure to meet you here in 2019.

Carlos Canudas-de-Wit, General Chair

The 58th IEEE Conference on Decision and Control will be held Wednesday through Friday, December 11-13, 2019 at the Palais des Congrès et des Expositions Nice Acropolis, Nice, France. The conference will be preceded by technical workshops on Tuesday, December 10, 2016.

The CDC is recognized as the premier scientific and engineering conference dedicated to the advancement of the theory and practice of systems and control. The CDC annually brings together an international community of researchers and practitioners in the field of automatic control to discuss new research results, perspectives on future developments, and innovative applications relevant to decision making, automatic control, and related areas.

The 58th CDC will feature contributed and invited papers, as well as tutorial sessions and workshops.

The IEEE CDC is hosted by the IEEE Control Systems Society (CSS) in cooperation with the Society for Industrial and Applied Mathematics (SIAM), and the Japanese Society for Instrument and Control Engineers (SICE).

Nice is the *Capital of the Côte d'Azur*, enjoying more than 300 days of sunshine per year. Its internationally-renowned history and culture, the shifting shimmer of its sea, its towering mountain peaks, make it a leading tourism destination with more than 5 million visitors. Its international airport serves 100 destinations worldwide.

The Palais des Congrès et des Expositions Nice Acropolis was elected "Best Congress and Convention Center Europe" for three consecutive years. It is conveniently located in the city center. It features 5 auditoriums from 250 to 2500 places and 50 meeting rooms.

Tweets by @CDC2019

CDC 2019 in Nice Retweeted

Nice Tourisme
@Nice_Tourisme

Quelle que soit l'heure, il est toujours agréable de se promener dans le #VieuxNice 😍

Whatever the time, it's always nice to stroll around Old Town 😍

Follow @CDC2019

Tweet #CDC2019

PaperPlaza Submission site

Key dates (2019)

Submission Site Open:	January 4
~~Initial Paper Submissions to L-CSS with CDC Option Due:~~	~~March 1~~
~~Invited Session Proposals Due:~~	~~March 7~~
~~Initial Paper Submissions Due:~~	~~March 17~~
~~Tutorial Session Proposals Due:~~	~~March 31~~
~~Workshop Proposals Due:~~	~~May 2~~
Paper and Workshop Decision Notification:	mid-July
Best Student Paper Nominations Opens:	July 20
Final Submission Open:	August 1
Registration Opens:	August 1
Best Student Paper Nominations Closes:	August 15
Accepted Papers Due:	September 10
Early Bird Closes:	October 1
Online Registration Closes:	December 5
Conference opens:	December 11
Conference closes:	December 13

Corporate Sponsors
Become a
Platinum, Gold, Silver, Bronze, Copper Sponsor
for the 58th CDC
and place your banner here

Platinum Sponsors

Flow open source library

https://flow-project.github.io/

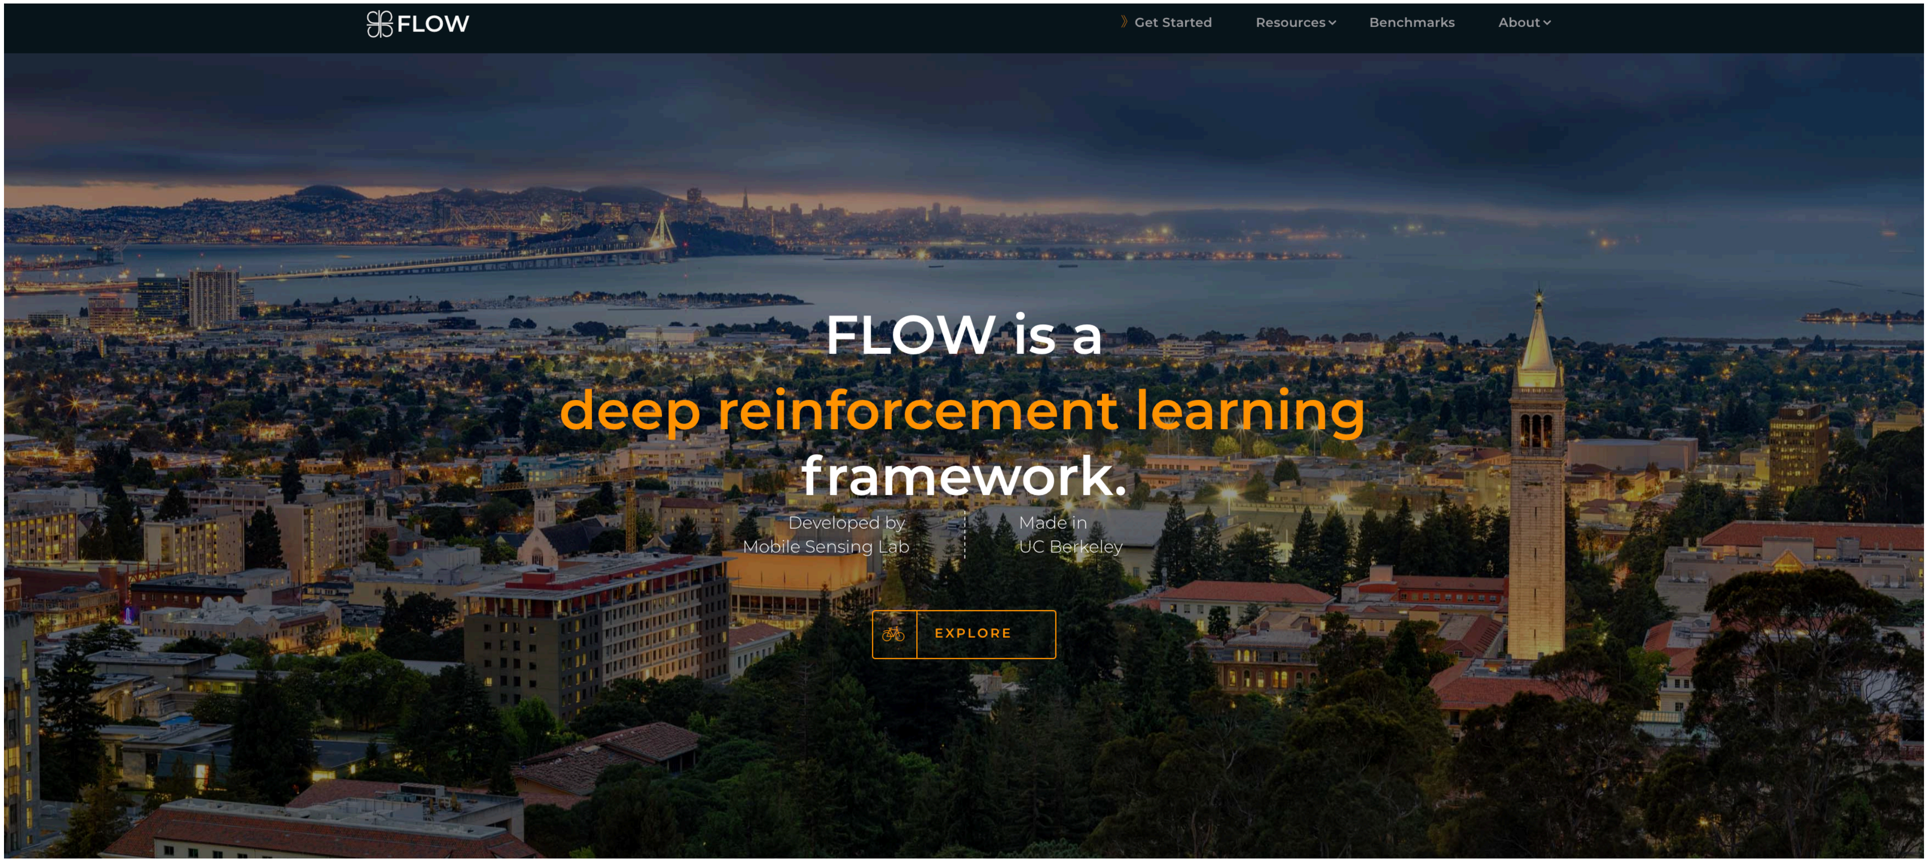

Solving Traffic Problems using Autonomous Vehicles

[redacted]

Alexandre Bayen
Director, Institute of Transportation Studies
Professor, EECS & CEE
Faculty Scientist, LBNL

Solving Traffic Problems using Autonomous Vehicles

amazon RE:MARS

Alexandre Bayen
Director, Institute of Transportation Studies
Professor, EECS & CEE
Faculty Scientist, LBNL