# Integrating Autonomy into Urban Systems

## A Reinforcement Learning Perspective


**Cathy Wu** | Postdoc, MSR AI

Assistant Professor (Fall19--)


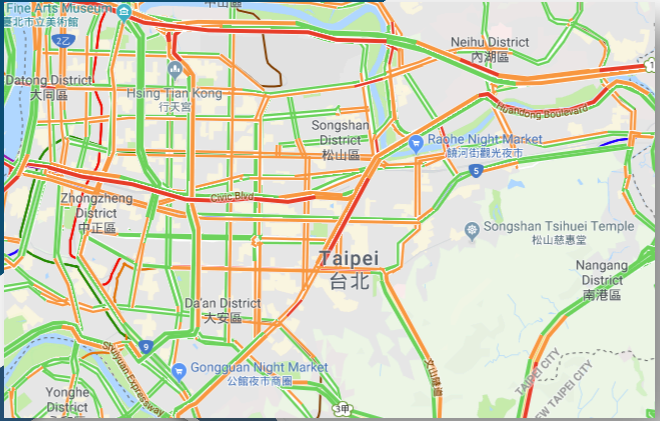



Microsoft

# Year 2050: How could self-driving cars change urban systems?


- **Traffic accidents**:
  - 37,000 fatalities
  - 41% deaths of young adults (ages 15-24)
  - 94% of serious crashes caused by human error
- **Greenhouse gas emissions**:
  - 28% from transportation
- **Congestion**:
  - 6.9 billion hours wasted
  - 3.1 billion gallons of fuel wasted (160$B)
- **Access to mobility**: 30% of population
  - 20% youth or elderly
  - 10% disabled (ages 18-64)

U.S. Energy Information Administration, 2017; U.S. Census Bureau, 2017.

# Year 2050: Current expert opinion on impact of AVs

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

Transportation today:
**31% US energy consumption**

100% self-driving cars:
**-40%** to **+100%** energy

Impact on safety? Access? Congestion? Environment?

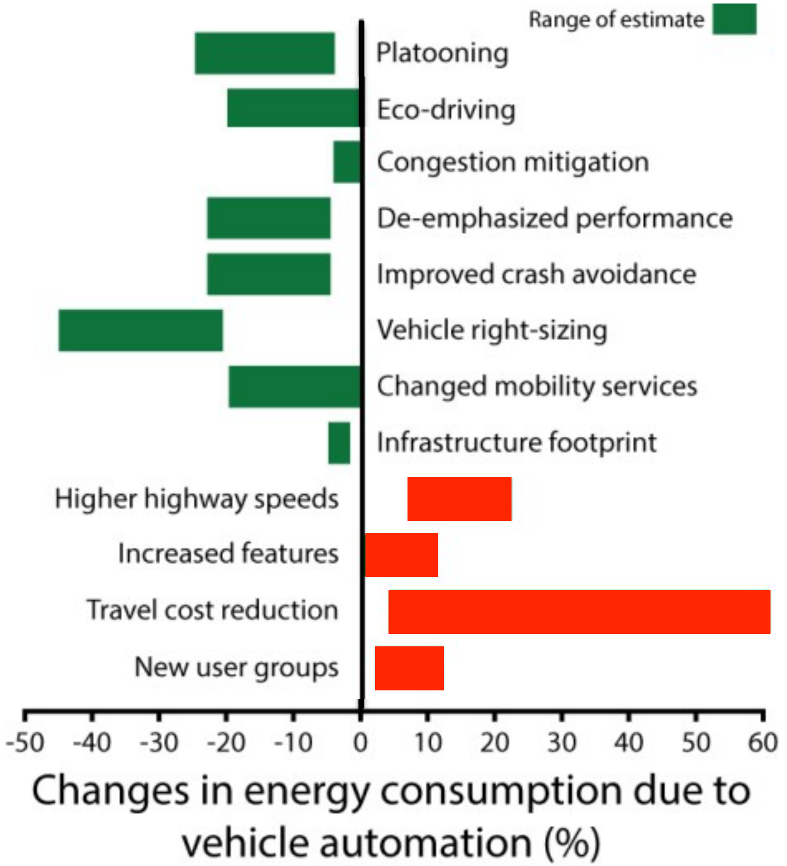


Wadud, et al. 2016. TR-A; U.S. Energy Information Administration, 2017.

# Integrating autonomy

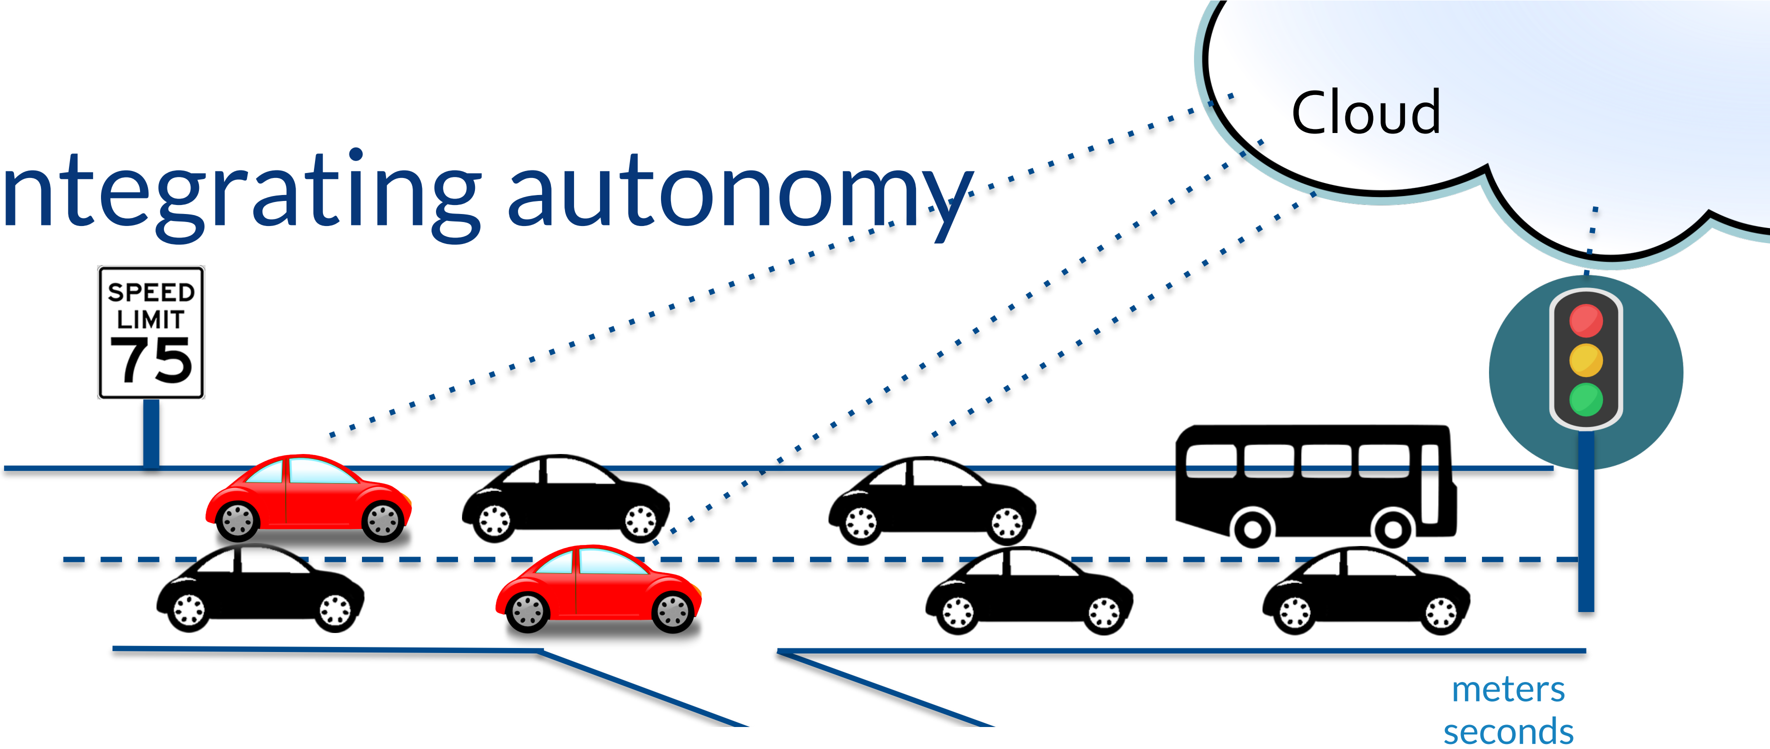


How can we gain understanding for integrating autonomy into complex systems?

In particular: traffic congestion.

# Long-standing challenges

(Deep reinforcement learning)

↑

(Policy optimization)

↑

- Highly complex non-linear delayed dynamics → Search possibilities
- Severe data limitations → Simulation
- Human behavior modeling → Leverage mature models
- Large-scale, heterogeneity → Seek insights in small settings
- Computational cost → Seek insights in small settings
- Limited benchmarks → Create some

Solutions:

- Search possibilities
- Simulation
- Leverage mature models
- Seek insights in small settings
- Create some

# Flow Research Team

FLOW

Microsoft

aws

riselab UC Berkeley

Berkeley DeepDrive

BAIR
BERKELEY ARTIFICIAL INTELLIGENCE RESEARCH

CALIFORNIA PATH

Eugene Vinitsky
UCB MechE, PhD

Aboudy Kreidieh
UCB CEE, PhD

Kathy Jang
UCB EECS, RA

Kanaad Parvate
EECS, ugrad

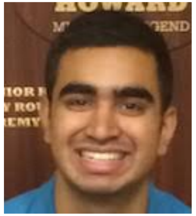

Nishant Kheterpal
EECS, ugrad

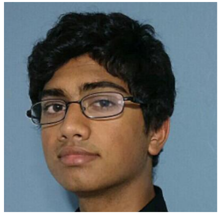

Ananth Kuchibhotla
EECS, ugrad

Leah Dickstein
EECS, ugrad

Nathan Mandi
EECS, ugrad

Alexandre Bayen
Principal Investigator
UCB EECS/CEE

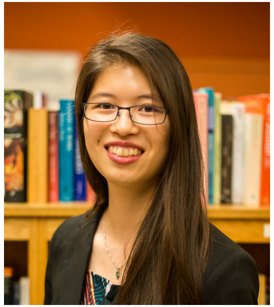

Cathy Wu
Founder & Advisor
MSR AI, MIT CEE/IDSS

# Deep reinforcement learning (RL) is a decision making framework

*Decisions in urban systems:*

Vehicle accelerations
Tactical maneuvers
Transit schedules
Traffic lights
Land use
Parking
Tolling
…

Recent successes of RL

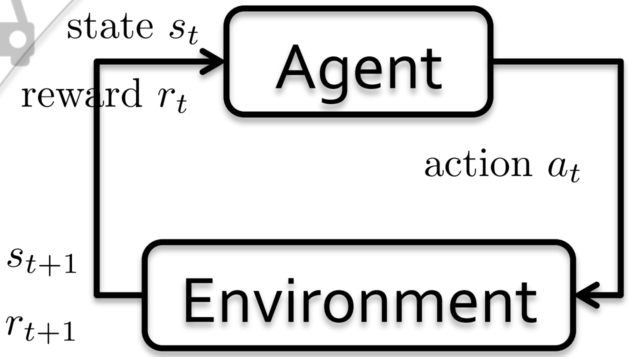


Video games

DQN (2015)

Locomotion

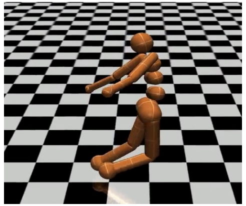

TRPO (2015)

Go

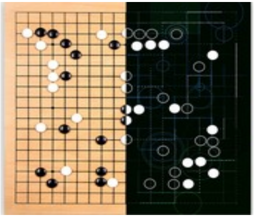

AlphaGo (2016)

# Reinforcement learning

- Automated
- Observed
- Unobserved

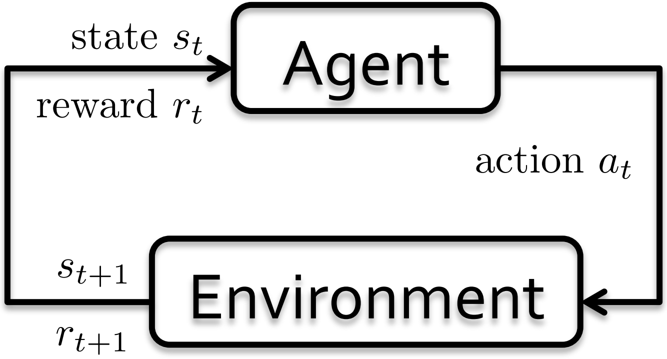


**Reward** -100 for every car crash; +1 otherwise

10 m

**State** ?? 8 m/s 8 m/s ?? ??

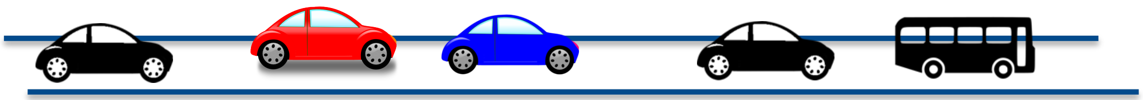

**Action** 5 m/s$^2$

# Reinforcement learning

Automated
Observed
Unobserved

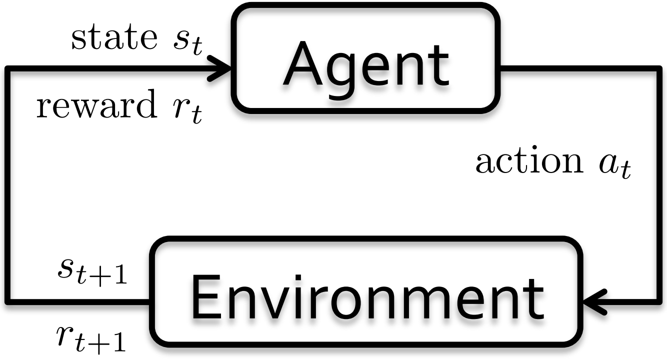


## Goal:

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

**Reward** -100 for every car crash; +1 otherwise

**State**

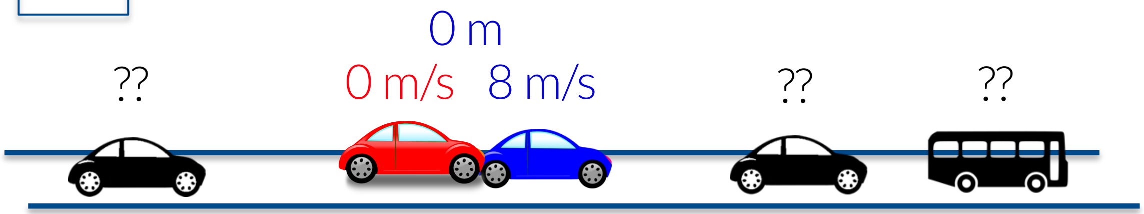

# Reinforcement learning

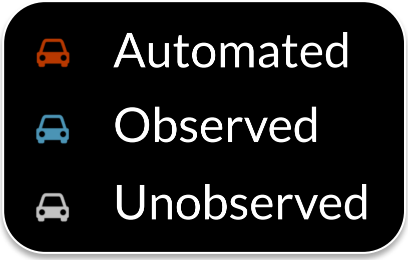


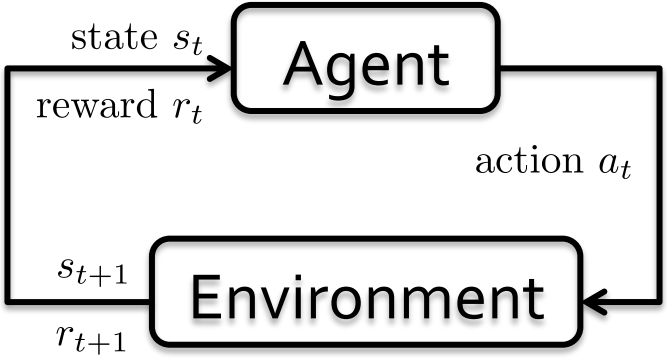


## Goal:

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

**Global** rewards

- average velocity
- energy consumption
- travel time
- safety
- comfort

# Deep reinforcement learning

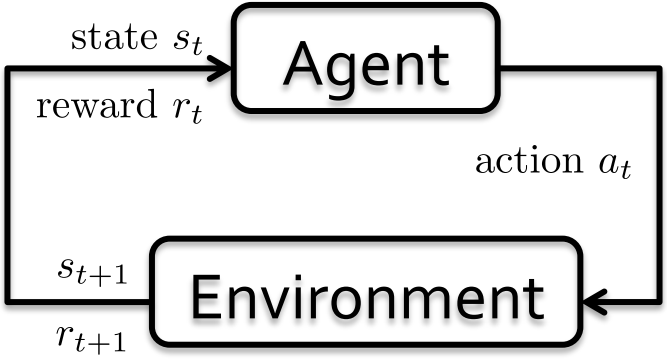


**Goal:**

learn policy $\pi : S \to A$

to maximize reward

Deep neural networks

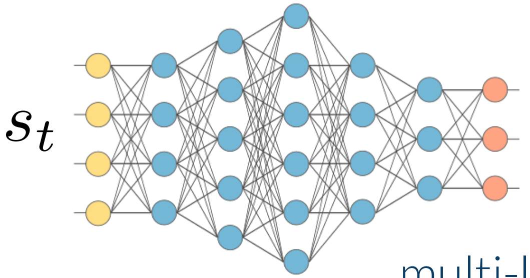


multi-layer perceptron (MLP)

Example Deep RL algorithms

- Deep Q Networks (DQN)
- Policy gradient

# Flow: full networks

## OpenStreetMaps

**Setting**: ~2000 vehicles

**Dynamics**:

- cascaded nonlinear systems
- bottlenecks
- multi-lane merges
- toll plaza dynamics

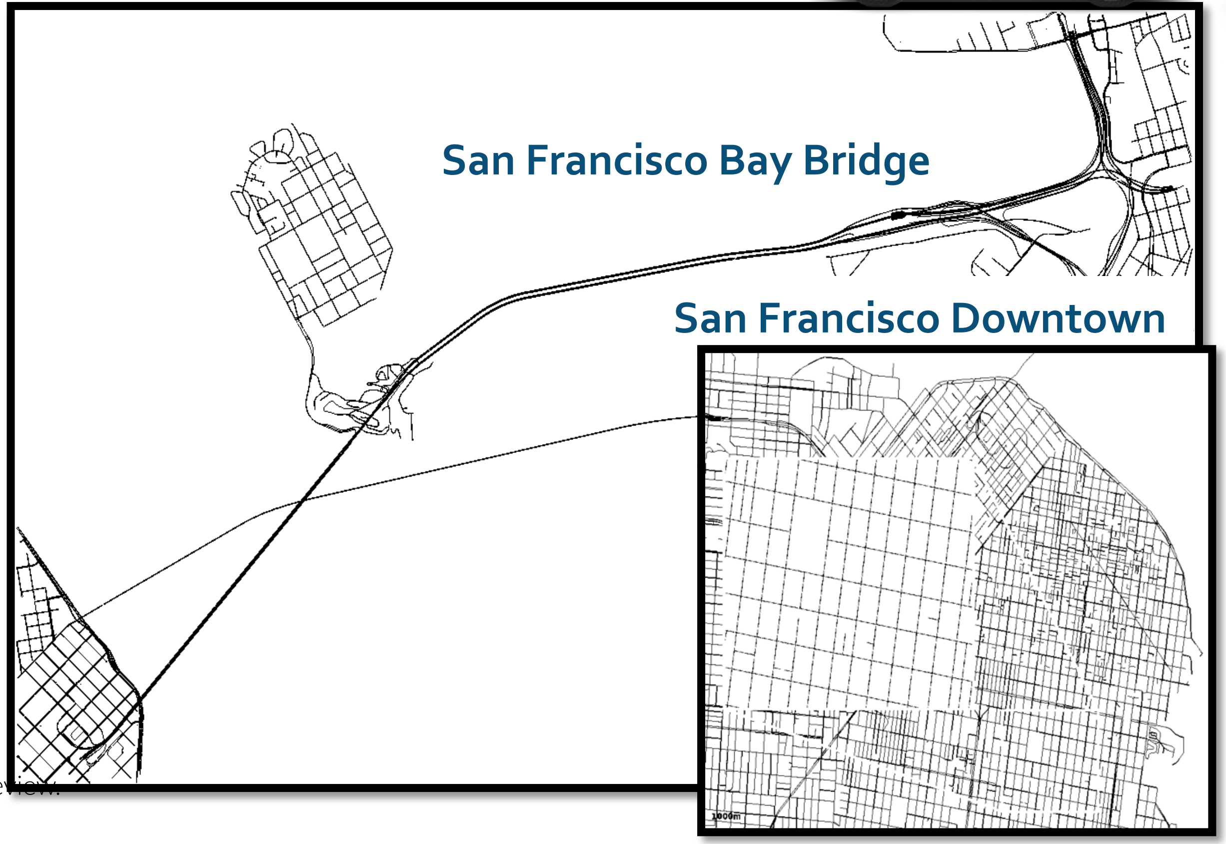


**Wu**, et al. IEEE ITSC, 2017; **Wu**, et al. IEEE T-RO, in review.

# Flow: traffic LEGO blocks

## Benchmarks for autonomy in transportation

Single-lane

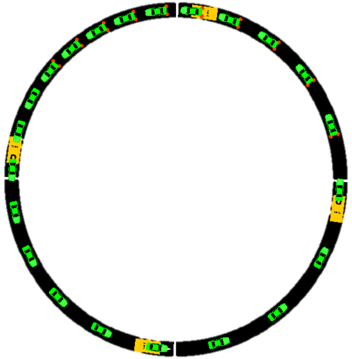

Multi-lane

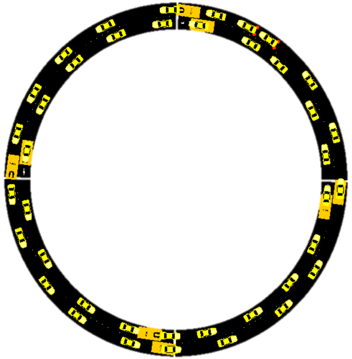

Intersection

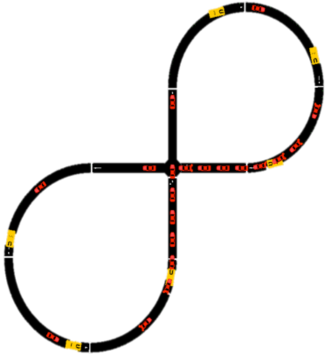

On/off-ramp

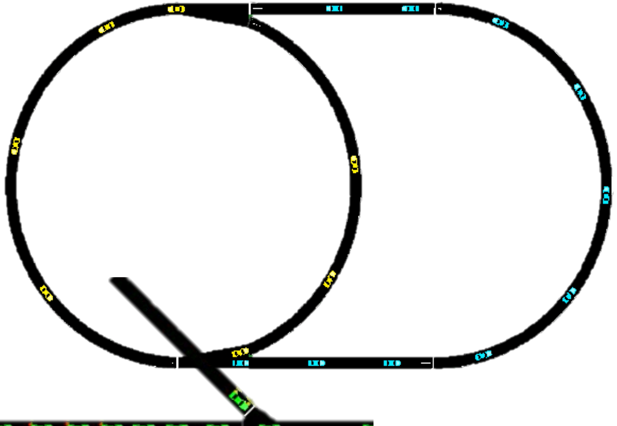

Straight highway

Bottleneck

Grid network

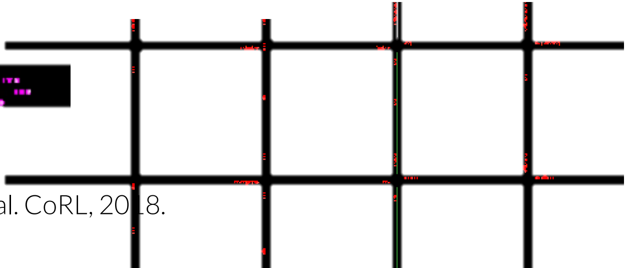

Signalized intersection

**Wu**, et al. IEEE ITSC, 2017; **Wu**, et al. IEEE T-RO, in review; Vinitsky, **Wu**, et al. CoRL, 2018.

# Traffic jams

**Sugiyama**, et al.

1955 — 900 papers on PDEs for traffic — **2008** — 2019

Partial differential equations (PDE)

**Setting**: 22 human drivers

**Instructions**: drive at 19 mph.

No traffic lights, stop signs, lane changes.

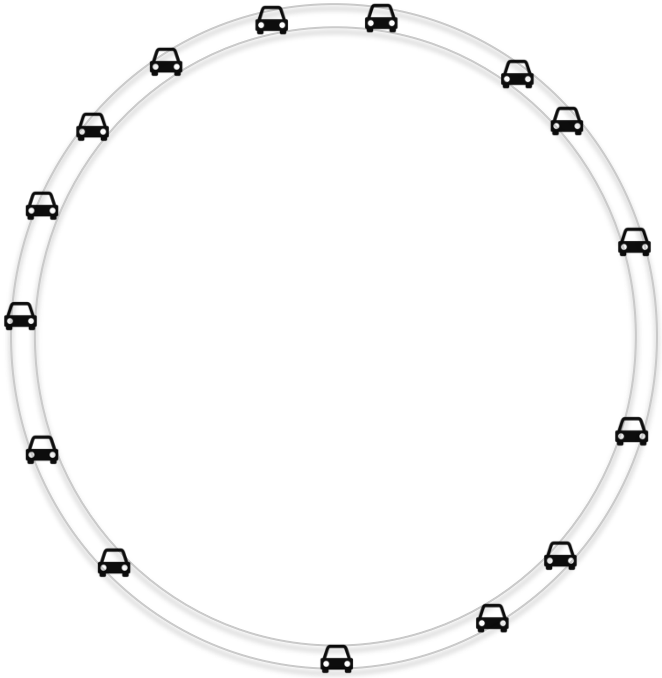

# Traffic jams

**Sugiyama**, et al.

1955 — 900 papers on PDEs for traffic — 2008 — 2019

Partial differential equations (PDE)

**Setting**: 22 human drivers

**Instructions**: drive at 19 mph.

No traffic lights, stop signs, lane changes.

**Traffic jams still form.**

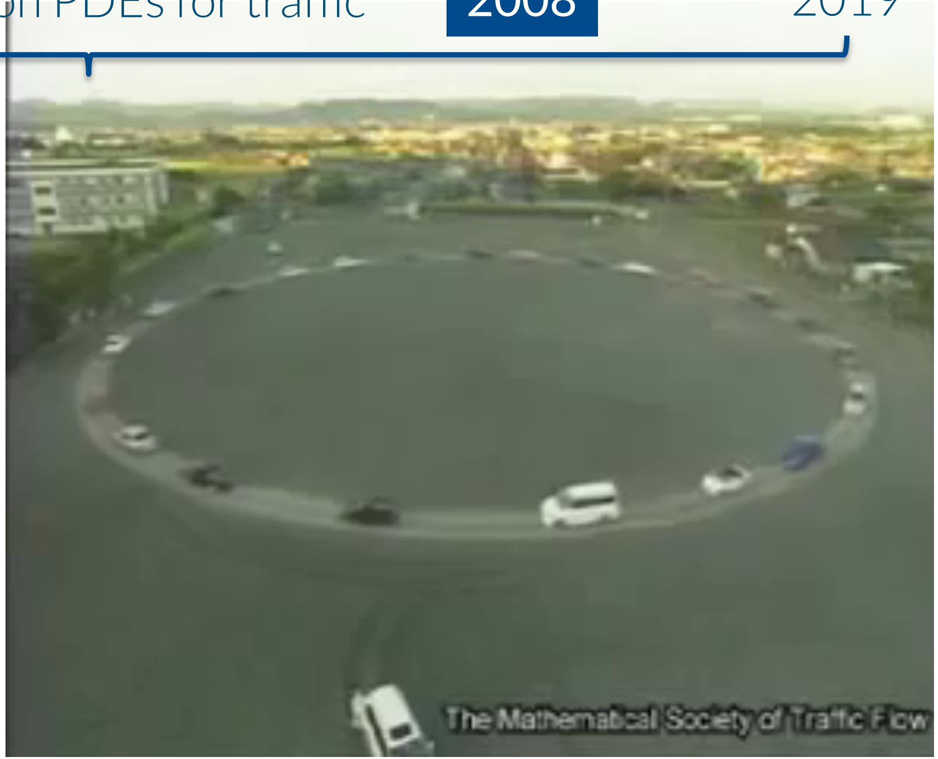


Video credits: NewScientist.com

# Single-lane traffic

Wu, et al. 2017

Sugiyama, et al.

1955 — 2008 — 2019

**Setting**: 1 AV, 21 human

**Experiment**

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

**Results**

- 1 AV: **+49%** average velocity
- Uniform flow at **near-optimal velocity**
- **Generalizes** to out-of-distr. densities

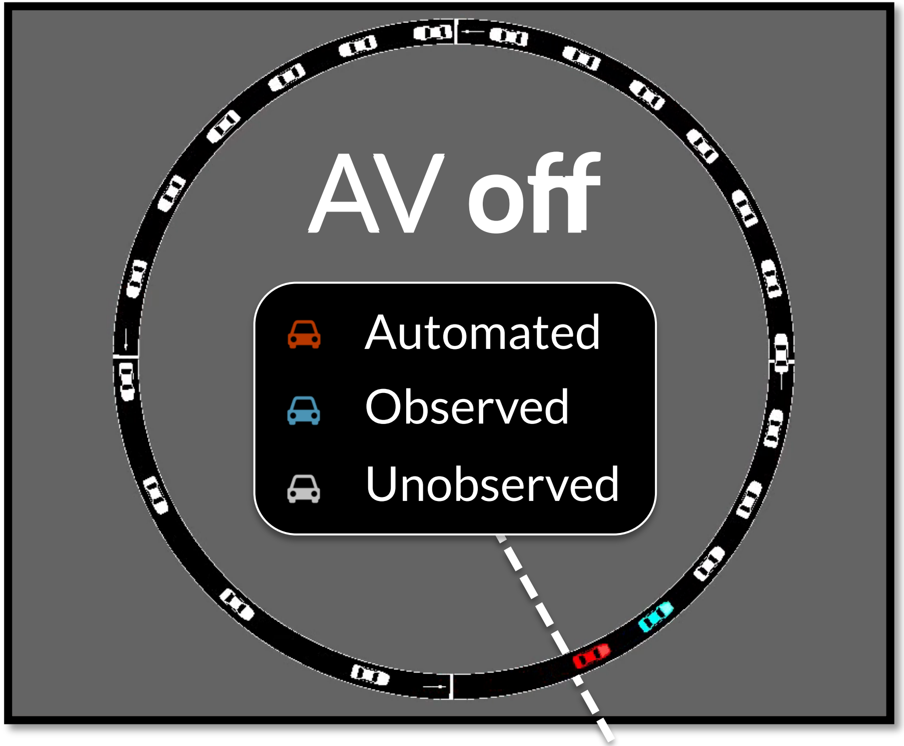


**Wu**, et al. CoRL, 2017; **Wu**, et al. IEEE T-RO, 2018

# San Francisco Bay Bridge

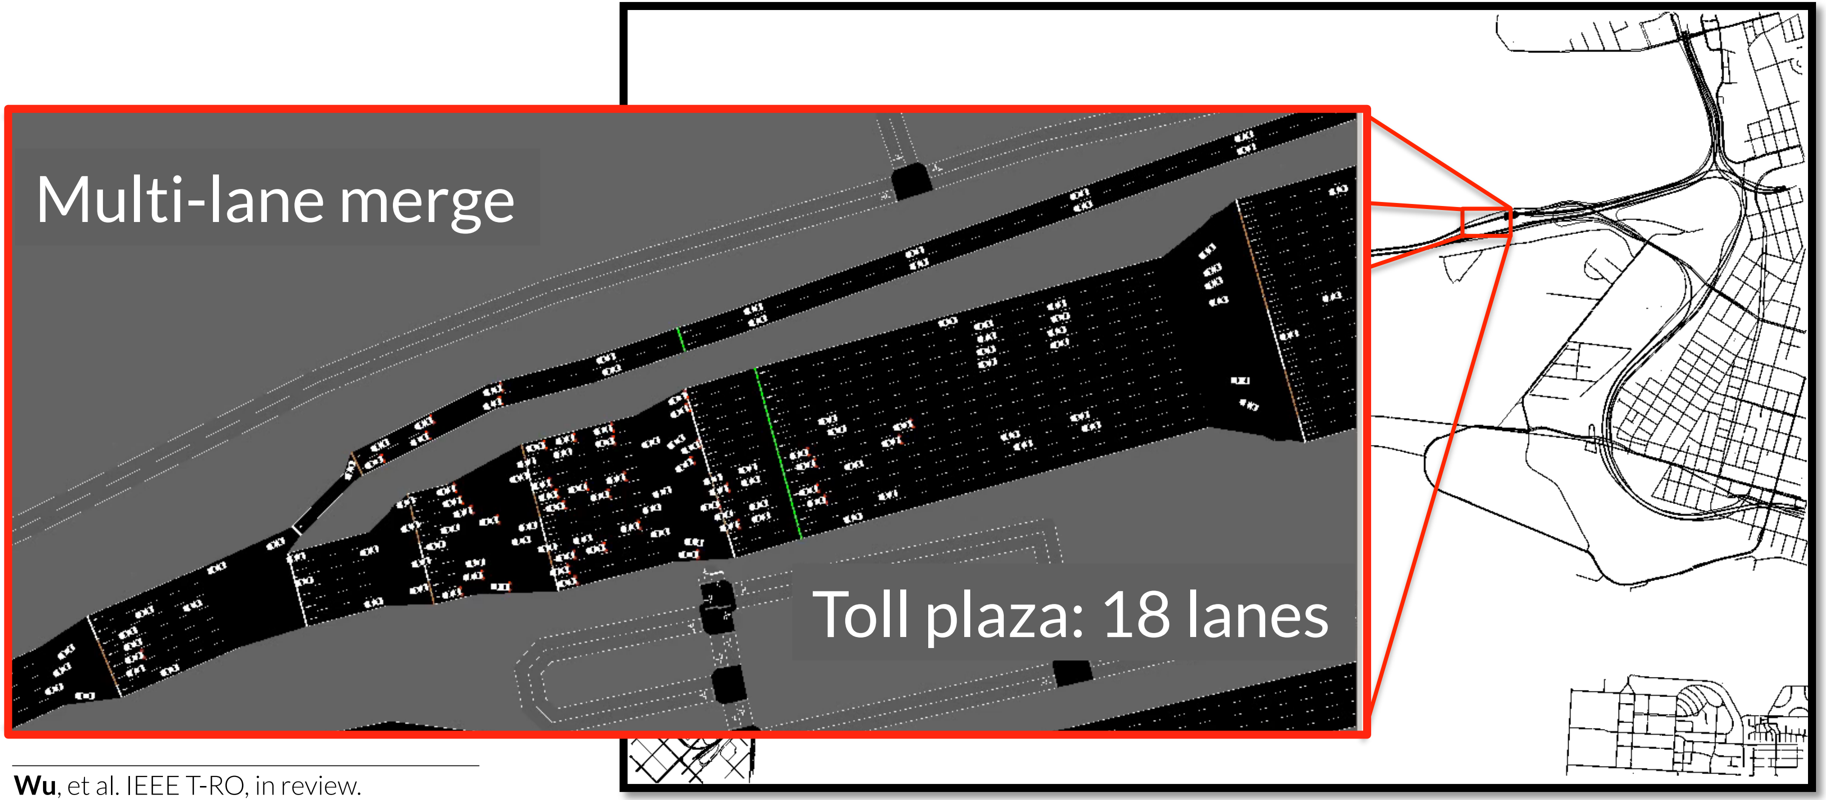


**Wu**, et al. IEEE T-RO, in review.

# Core problem: traffic bottleneck


Eugene Vinitsky

Setting: No AVs

720 veh/hr

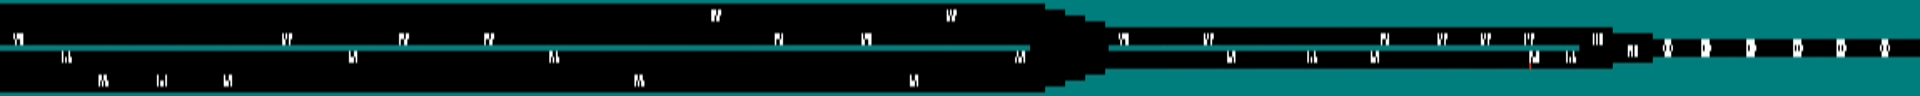

**Phenomenon: capacity drop**

Setting: 10% AVs

1020 veh/hr

**40% improvement**
**Avoids capacity drop**

Dynamics:

- Four lanes → Two lanes → One
- Cascaded nonlinear systems with right-of-way dynamics model, **merge conflicts**, and **excessive, fluctuating inflow**

Vinitsky, Parvate, Kreidieh, **Wu**, Bayen. IEEE ITSC, 2018

# FLOW Flow: platform for RL + urban decisions

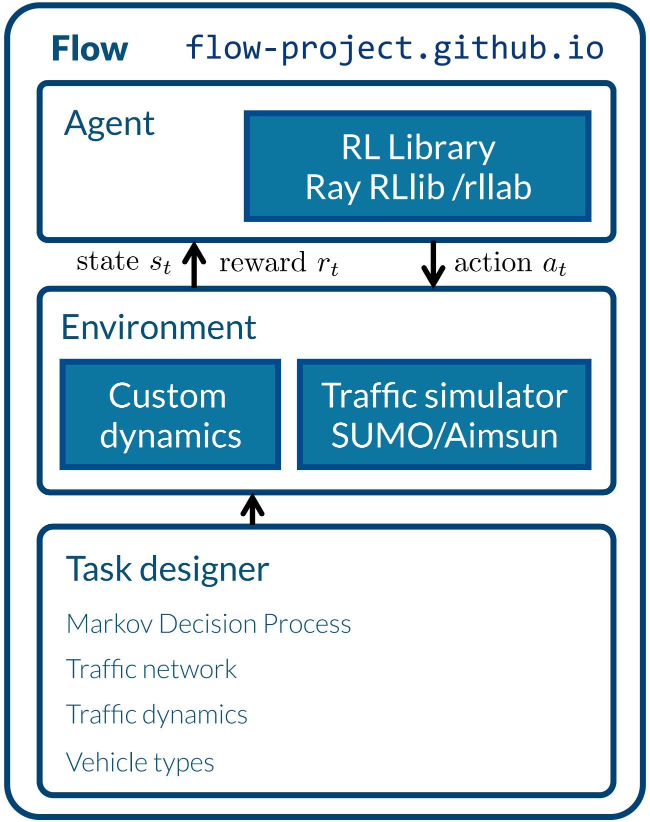


**Control signals**

Longitudinal, lateral control

Traffic light control, ramp meters

**Large-scale reinforcement learning**

Hierarchical policy

Multi-agent environments

Distributed simulation and sampling

**Scenarios and networks**

Parameterized python scenario creation

A variety of open and closed networks

OSM network import

**Libraries**

Rich models via SUMO/Aimsun

OpenAI gym interface

Supports rllab and Rllib

AWS Machine Learning Research Award

Flow is *open-source*. Check it out!

Team: `flow-dev@googlegroups.com`

Docs: `flow.readthedocs.io`

Website: `flow-project.github.io`

**Wu**, et al. arXiv, 2019; **Wu**, et al. ITSC, 2017; Krajzewicz, et al. IJASM, 2012; Duan, et al. 2016. arXiv:1604.06778; Liang, et al. 2017. arXiv:1712.09381.

# Exercise 02: Running RLlib Experiments

This tutorial walks you through the process of running traffic simulations in Flow with trainable RLlib-powered agents. Autonomous

# Exercise 03: Running rllab Experiments

This tutorial walks you through the process of running traffic simulations in Flow with trainable rllab-powered agents. Autonomous

# Exercise 04: Visualizing Experiment Results

This tutorial describes the process of visualizing and replaying the results of Flow experiments run using RL. The process of

# Exercise 05: Creating Custom Scenarios

This tutorial walks you through the process of generating custom scenarios. Scenarios define the network geometry of the problem, as well as the constituents of the network, e.g. vehicles, traffic lights, etc... Various scenarios are available in Flow, depicting a diverse set of open and closed traffic networks such as ring roads, intersections/grids, straight highway merges, and more.

In this exercise, we will recreate the ring road network, seen in the figure below.

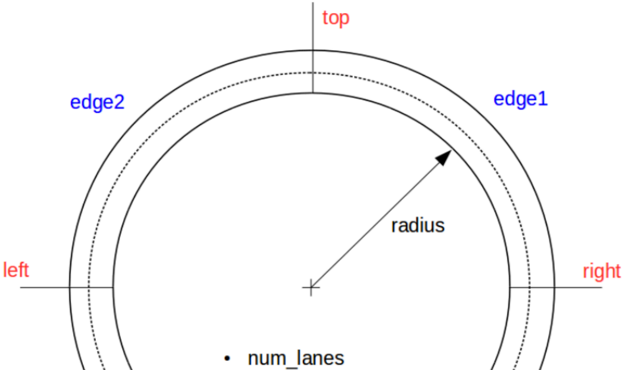


A. Kreidieh

FLOW Tutorials

flow-project.github.io

# Towards: reliable decisions in urban systems

Controller design,
traffic control for AVs

Understanding
adversarial driving

Scalable RL for
networked systems

KACST

System
verification

Behavior modeling
& distribution shift

Urban decision
support systems

# Integrating autonomy into urban systems

**Deep reinforcement learning** provides **understanding** for integrating autonomy into urban systems.


**Small % of AVs** greatly affect **traffic dynamics**, which in turn, affects all parts of the urban system.

**5-10%** AVs

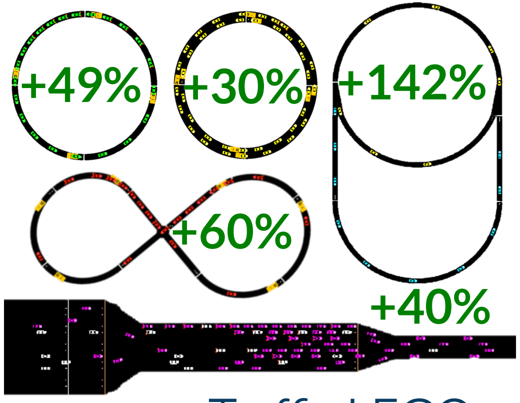


Traffic LEGOs

**Flow**: open source project to enable RL for traffic control

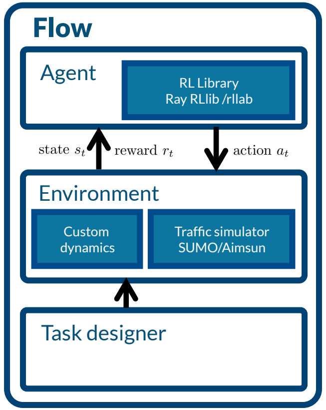


FLOW

flow-project.github.io


Cathy Wu
cathywu@mit.edu