# CS 188: Artificial Intelligence

## Advanced Applications: Games and Robotics*

Instructors: Pieter Abbeel & Dan Klein --- University of California, Berkeley

These slides were created by Dan Klein, Pieter Abbeel and Anca Dragan for CS188 Intro to AI at UC Berkeley.

All CS188 materials are available at http://ai.berkeley.edu

# So Far: Foundational Methods

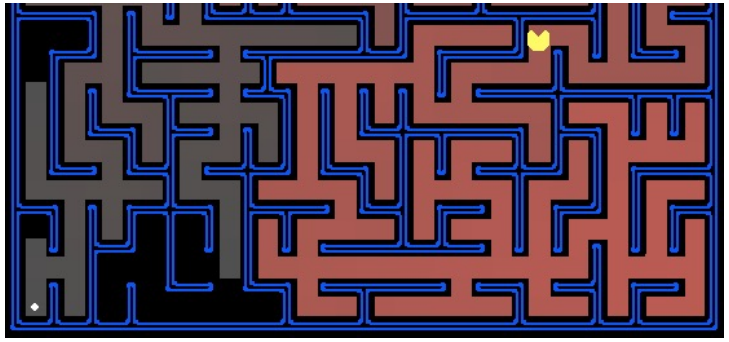

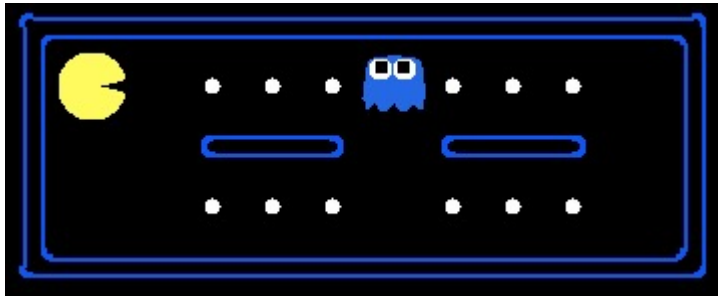

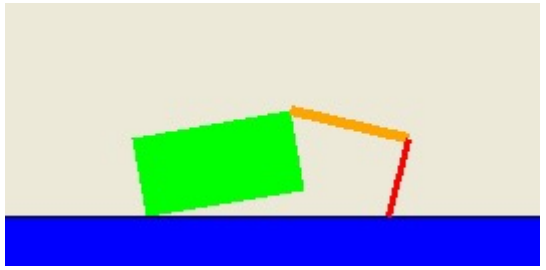

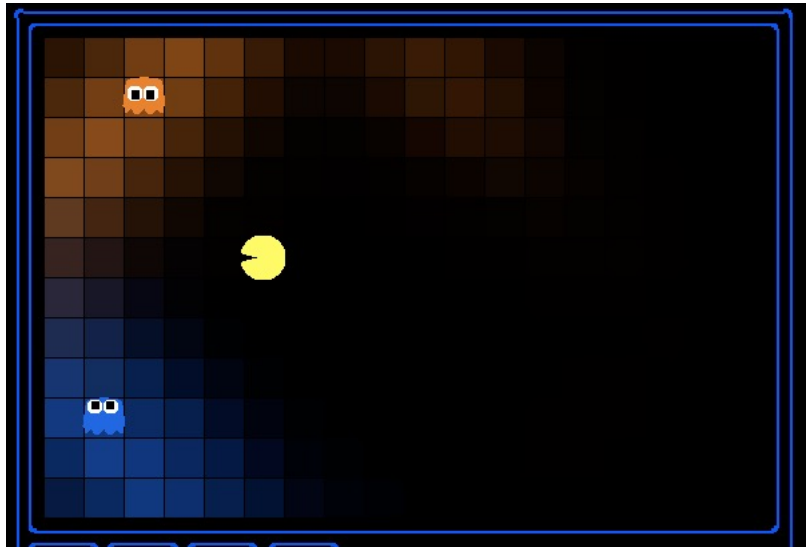

# Now: Advanced Applications


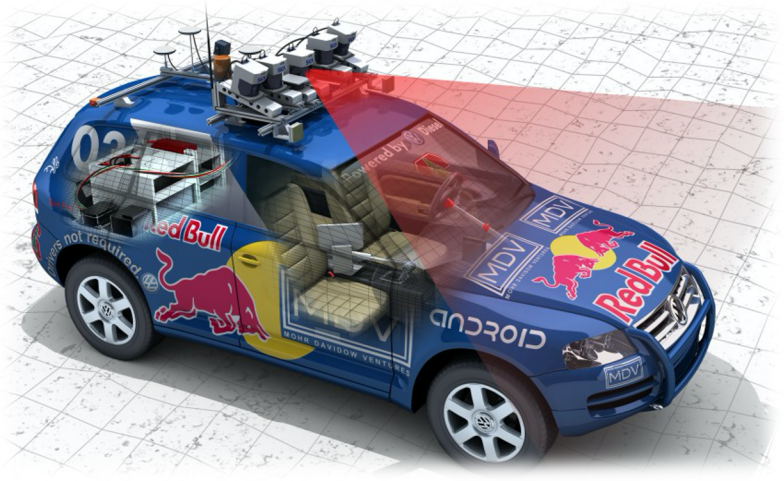


# AlphaGo


[Silver et al (Deepmind), 2015]

# Next level games?

- Dota2 – OpenAI Five

  https://openai.com/five/


- Starcraft – Deepmind's AlphaStar

  https://deepmind.com/blog/article/alphastar-mastering-real-time-strategy-game-starcraft-ii

# Why Dota?

- Problem (seemingly) out of reach of existing Deep RL

- However, surprising finding:
  - Scaling up existing Deep RL + getting the details right got the job done!


70
60
50
40
30
Trueskill
Finals
Buyback
7.21
7.20
4096 LSTM Size
Benchmark
TI
1 Courier
$10^1$
$10^2$
$10^3$
Training PF/s Days

# Deep RL

- Proximal Policy Optimization PPO
- Discount Factor gamma:
  - annealed from 0.998 (half-life of 46 seconds) to 0.9997 half-life of five minutes)
- NN architecture:
  - 1024 unit LSTM
- Self-play training
  - 80% against current self
  - 20% against past selves
- Randomization with ”lane assignment” till a certain time
- Hardcode item and skill builds, and choose which of the builds to use at random

# RL agent defeats in-house OpenAI team at fairly restricted 5v5.

Mirror match of 5 fixed heroes utilizing 5 invulnerable couriers. No neutrals, runes, shrines, wards, invisibility, summons, illusions, or Scan. No Divine Rapier, Bottle, Quelling Blade, Boots of Travel, Tome of Knowledge, or Infused Raindrop.

> READ "OPENAI FIVE"
>
> ▶ WATCH VIDEO

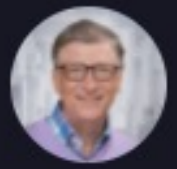

**Bill Gates**
@BillGates

#AI bots just beat humans at the video game Dota 2. That's a big deal, because their victory required teamwork and collaboration – a huge milestone in advancing artificial intelligence.

via Twitter

# OpenAI Five Defeats Dota 2 World Champions

OpenAI Five is the first AI to beat the world champions in an esports game, having won two back-to-back games versus the world champion Dota 2 team, OG, at Finals this weekend. Both OpenAI Five and DeepMind's AlphaStar had previously beaten good pros privately but lost their live pro matches, making this also the first time an AI has beaten esports pros on livestream.

# Starcraft

# What is Starcraft?

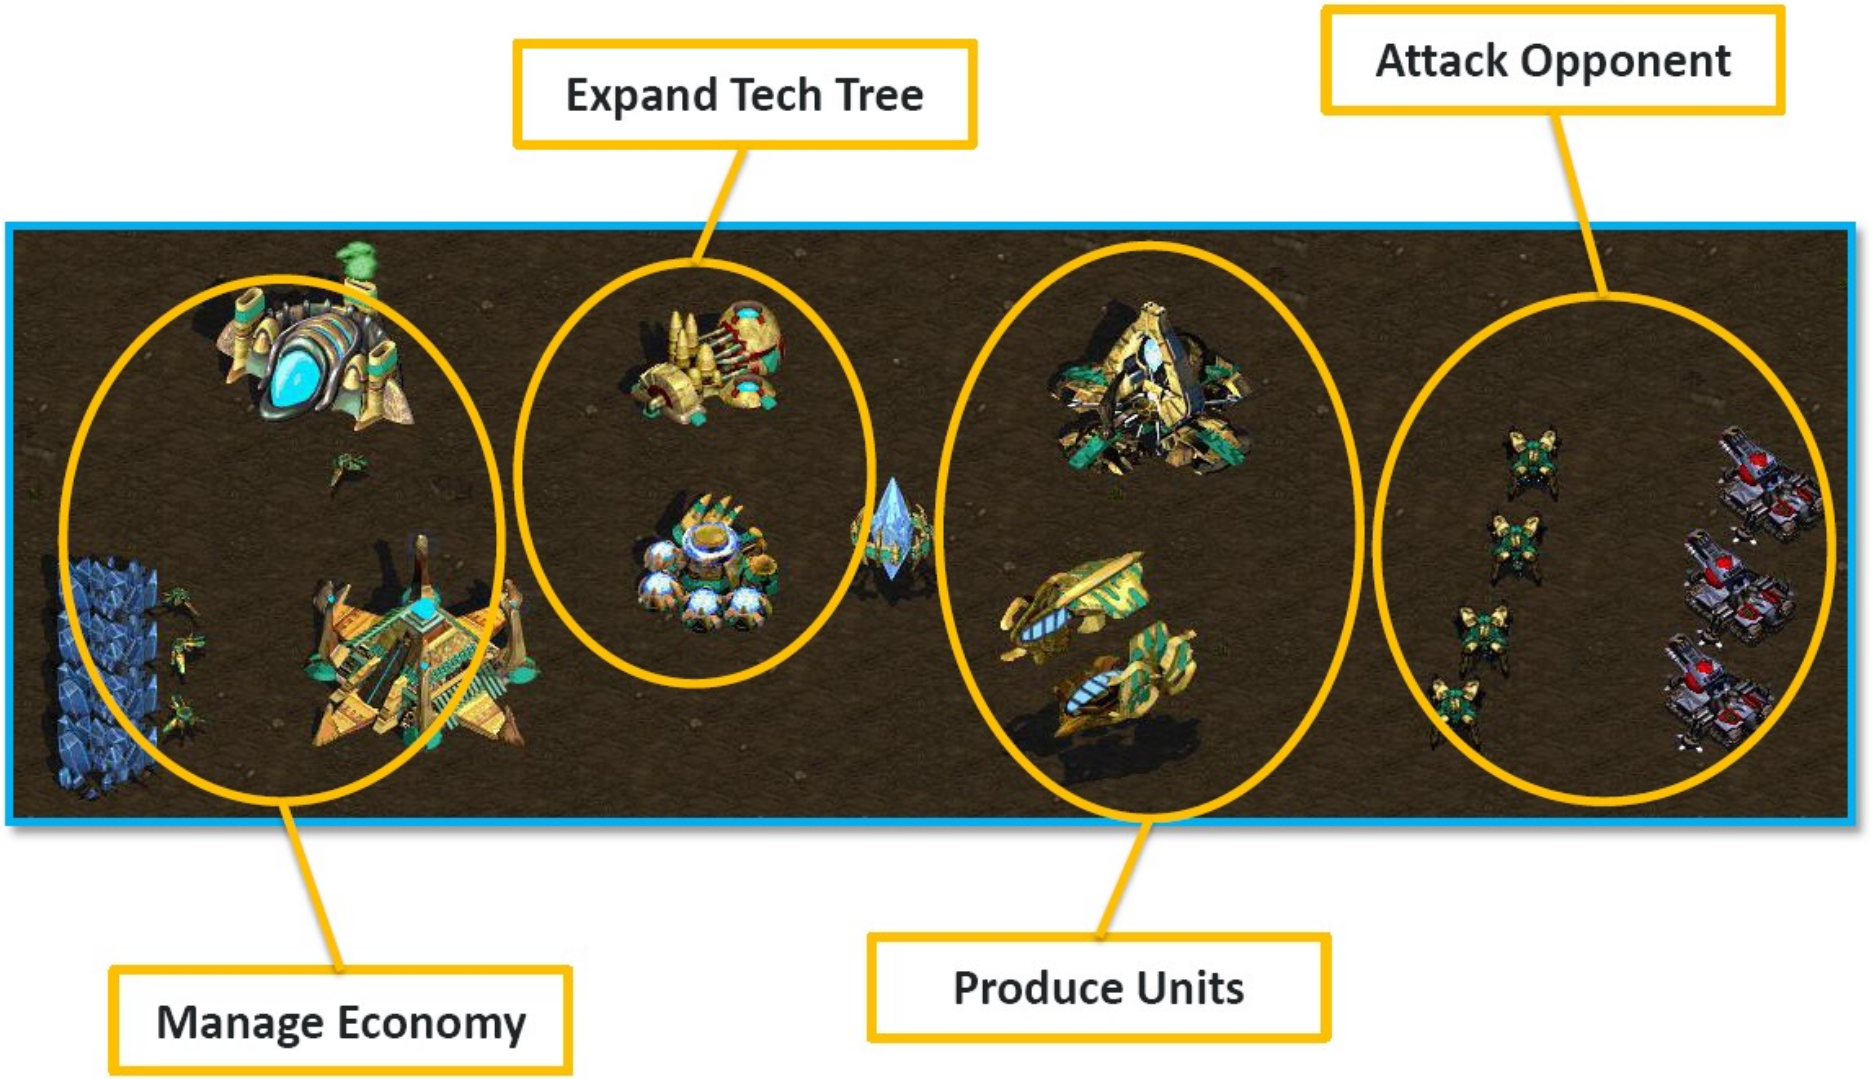


Image from Ben Weber

# Why is Starcraft Hard?

- The game of Starcraft is:
  - Adversarial
  - Long Horizon
  - Partially Observable
  - Realtime
  - Huge branching factor
  - Concurrent
  - Resource-rich
  - …
- No single al… (e.g. minimax) will sol… the-shelf!

Or might it?


# AlphaStar

- Large NN trained:
  - Phase 1: supervised learning to imitate (strong) human players (why?)
  - Phase 2: reinforcement learning

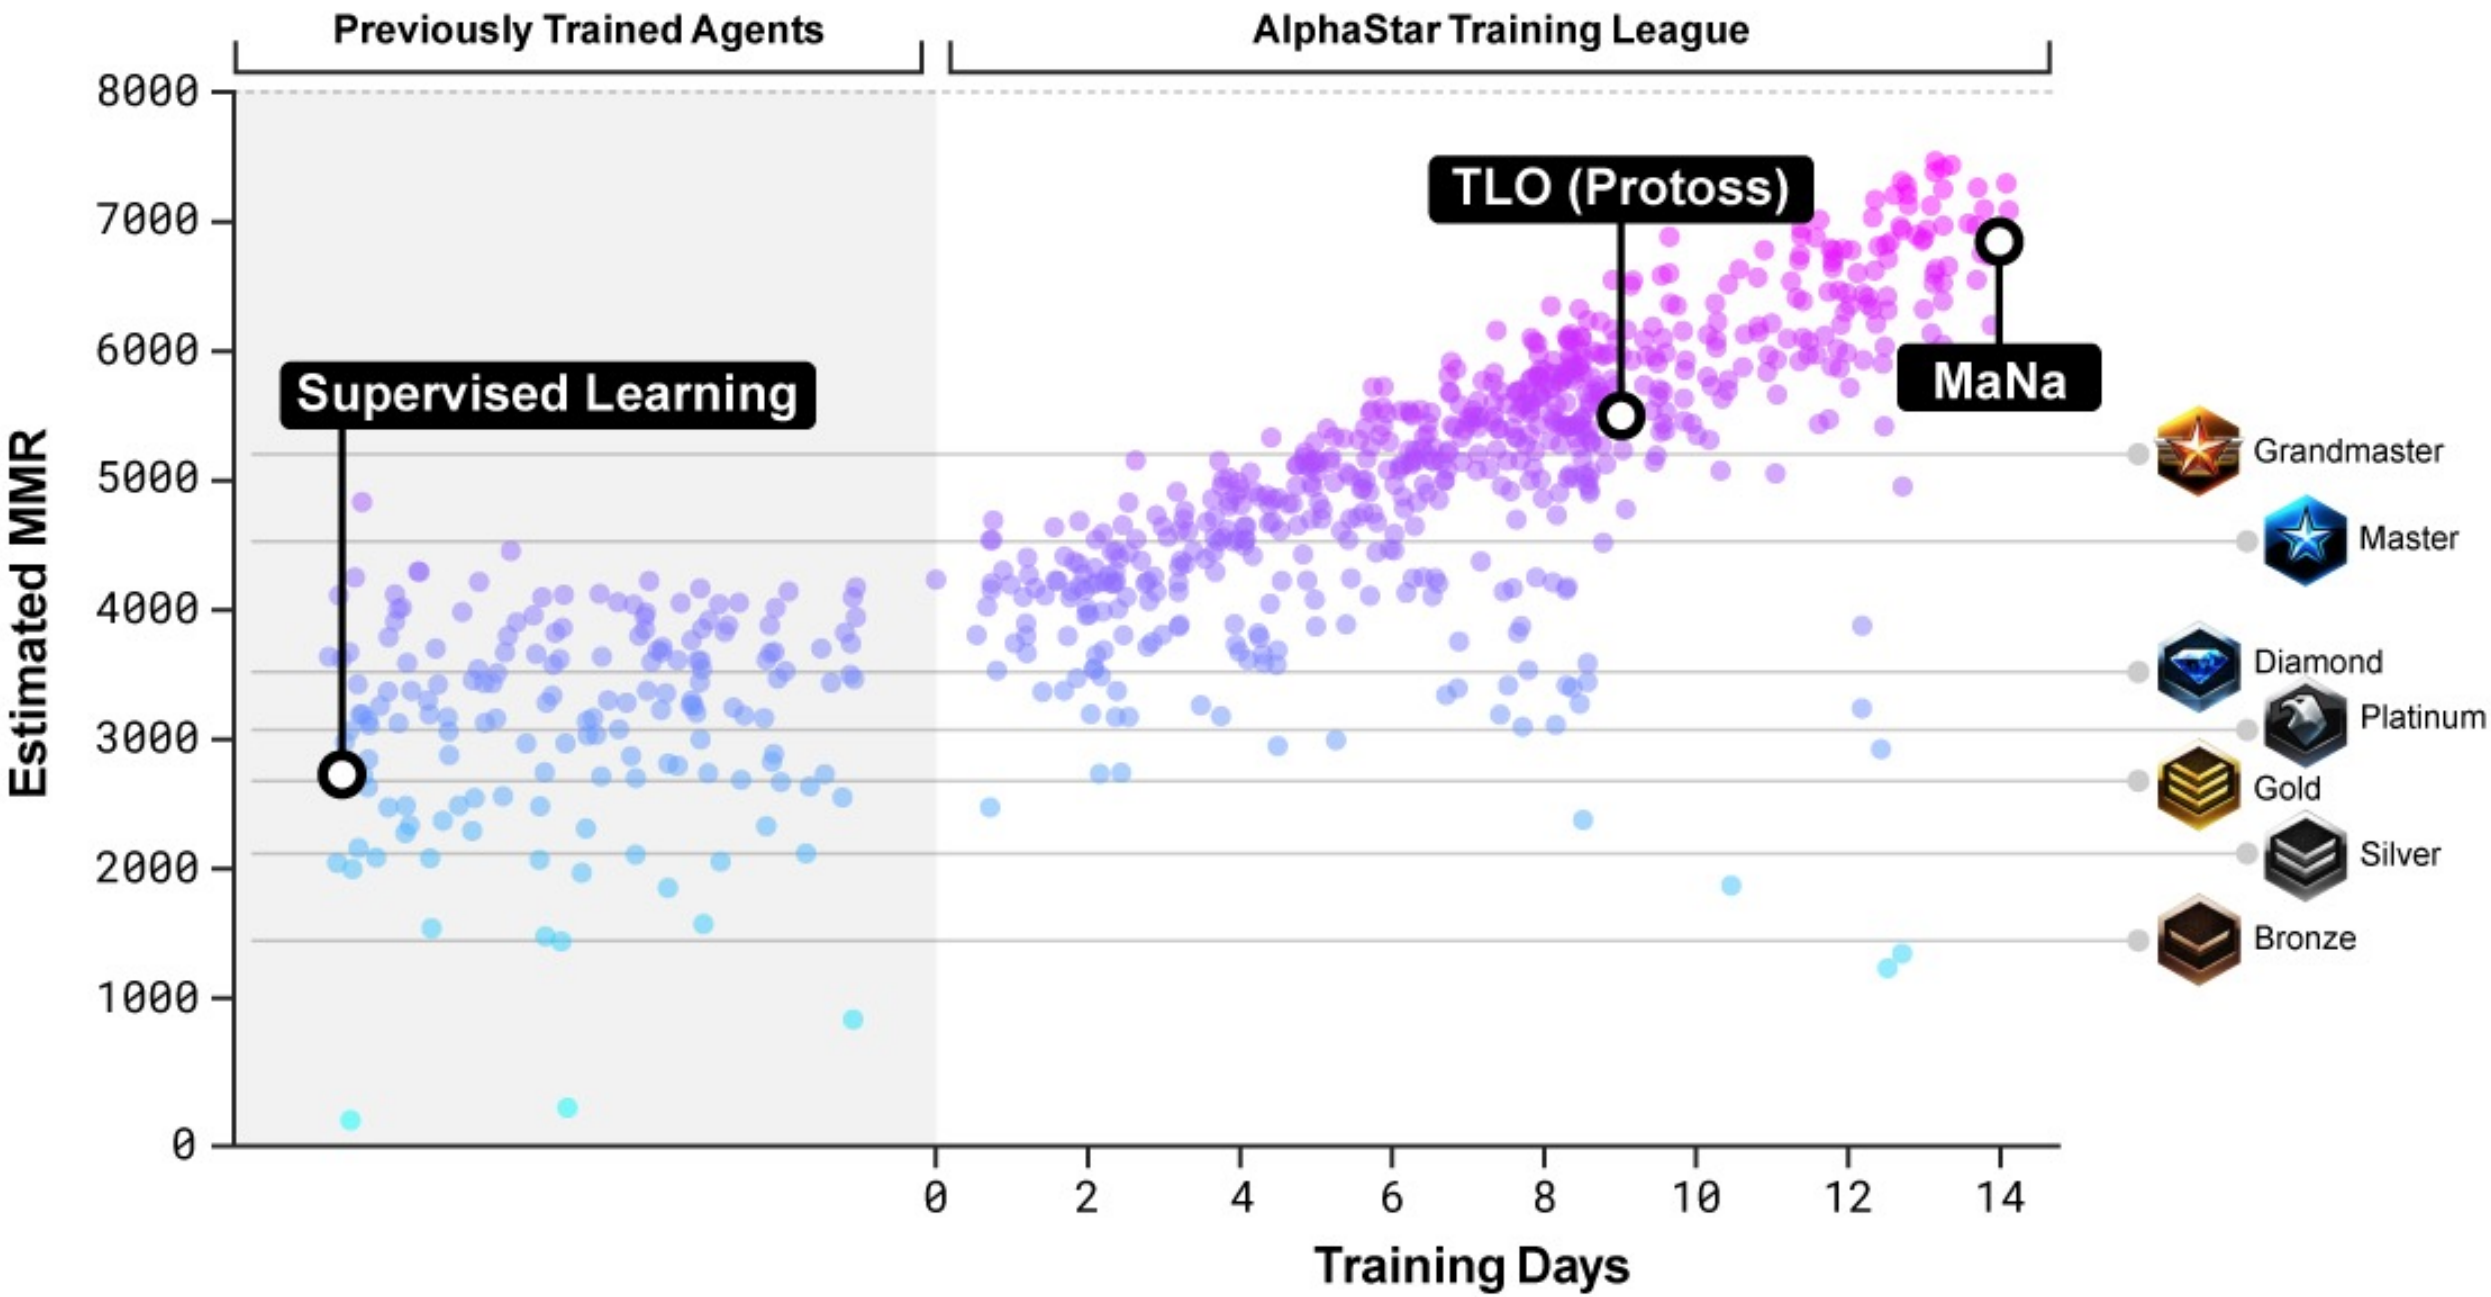
Previously Trained Agents
AlphaStar Training League
8000
7000
6000
5000
4000
3000
2000
1000
0
Estimated MMR
Supervised Learning
TLO (Protoss)
MaNa
Grandmaster
Master
Diamond
Platinum
Gold
Silver
Bronze
0
2
4
6
8
10
12
14
Training Days

# How strong is AlphaStar?

- Played against MaNa, one of the world's strongest StarCraftII players, among the 10 strongest Protoss players
- AlphaStar won 5-0

Article | Published: 30 October 2019

# Grandmaster level in StarCraft II using multi-agent reinforcement learning

Oriol Vinyals, Igor Babuschkin, Wojciech M. Czarnecki, Michaël Mathieu, Andrew Dudzik, Junyoung Chung, David H. Choi, Richard Powell, Timo Ewalds, Petko Georgiev, Junhyuk Oh, Dan Horgan, Manuel Kroiss, Ivo Danihelka, Aja Huang, Laurent Sifre, Trevor Cai, John P. Agapiou, Max Jaderberg, Alexander S. Vezhnevets, Rémi Leblond, Tobias Pohlen, Valentin Dalibard, David Budden, Yury Sulsky, James Molloy, Tom L. Paine, Caglar Gulcehre, Ziyu Wang, Tobias Pfaff, Yuhuai Wu, Roman Ring, Dani Yogatama, Dario Wünsch, Katrina McKinney, Oliver Smith, Tom Schaul, Timothy Lillicrap, Koray Kavukcuoglu, Demis Hassabis, Chris Apps & David Silver — Show fewer authors

*Nature* **575**, 350–354 (2019) | Cite this article

# Robotic Helicopters

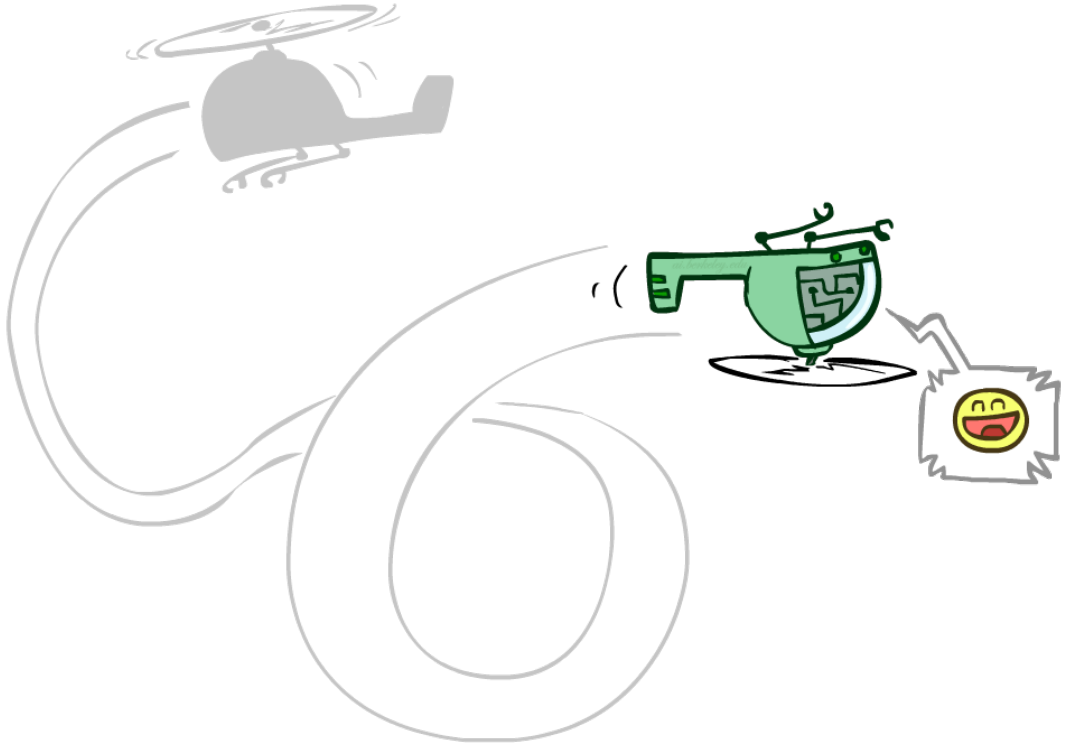

# Motivating Example

- How do we execute a task like this?

[VIDEO: tictoc_results.wmv]

# Autonomous Helicopter Flight


- Key challenges:
  - Track helicopter position and orientation during flight
  - Decide on control inputs to send to helicopter

# Autonomous Helicopter Setup

Position

On-board inertial measurement unit (IMU)

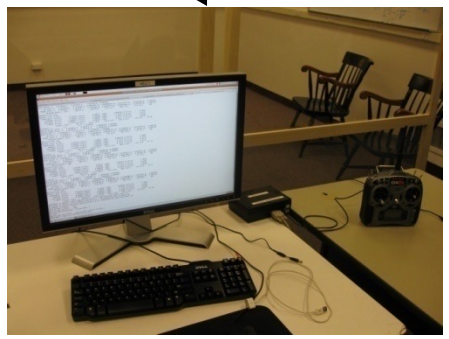

Send out controls to helicopter

# HMM for Tracking the Helicopter

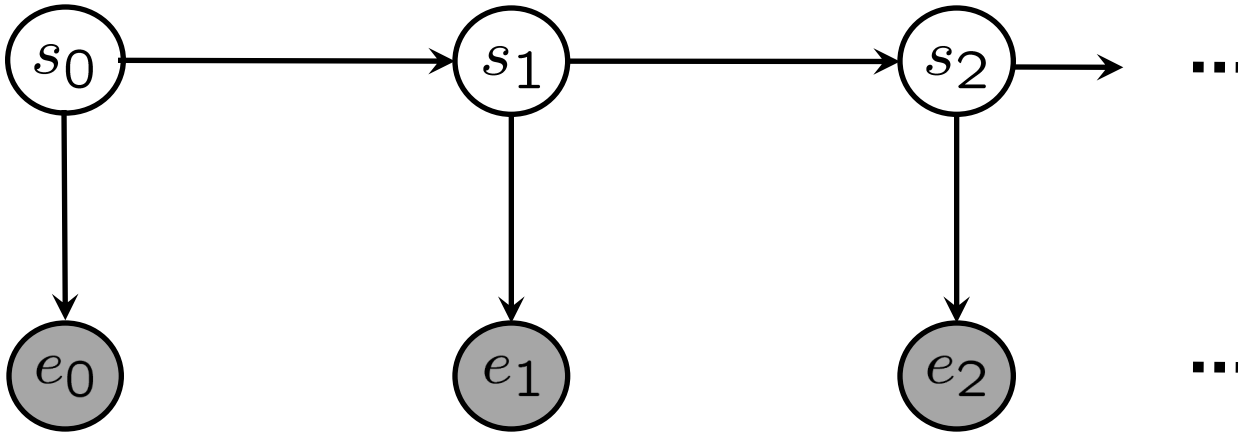


- State: $s = (x, y, z, \phi, \theta, \psi, \dot{x}, \dot{y}, \dot{z}, \dot{\phi}, \dot{\theta}, \dot{\psi})$
- Measurements: [observation update]
  - 3-D coordinates from vision, 3-axis magnetometer, 3-axis gyro, 3-axis accelerometer
- Transitions (dynamics): [time elapse update]
  - $s_{t+1} = f(s_t, a_t) + w_t$ f: encodes helicopter dynamics, w: noise

# Helicopter MDP

- State: $s = (x, y, z, \phi, \theta, \psi, \dot{x}, \dot{y}, \dot{z}, \dot{\phi}, \dot{\theta}, \dot{\psi})$

- Actions (control inputs):
  - $a_{lon}$ : Main rotor longitudinal cyclic pitch control (affects pitch rate)
  - $a_{lat}$ : Main rotor latitudinal cyclic pitch control (affects roll rate)
  - $a_{coll}$ : Main rotor collective pitch (affects main rotor thrust)
  - $a_{rud}$ : Tail rotor collective pitch (affects tail rotor thrust)


- Transitions (dynamics):
  - $s_{t+1} = f(s_t, a_t) + w_t$

  [f encodes helicopter dynamics]

  [w is a probabilistic noise model]

- Can we solve the MDP yet?

# Problem: What's the Reward?

- Reward for hovering:

$$
\begin{aligned}
R(s) \quad = \quad & -\alpha_x (x - x^*)^2 \\
& -\alpha_y (y - y^*)^2 \\
& -\alpha_z (z - z^*)^2 \\
& -\alpha_{\dot{x}} \dot{x}^2 \\
& -\alpha_{\dot{y}} \dot{y}^2 \\
& -\alpha_{\dot{z}} \dot{z}^2
\end{aligned}
$$

# RL: Helicopter Flight

[Andrew Ng]

[Video: HELICOPTER]

# Problem for More General Case: What's the Reward?

- Rewards for "Flip"?
  - Problem: what's the target trajectory?
  - Just write it down by hand?

# Flips (?)

[VIDEO: 20061204---bad.wmv]

# Helicopter Apprenticeship?

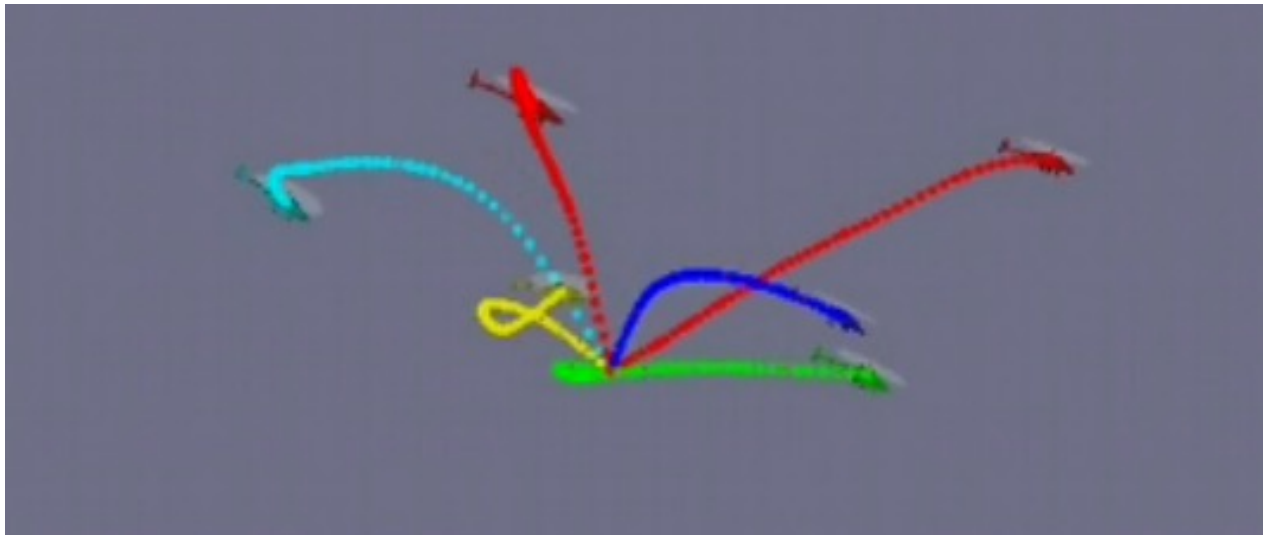

[VIDEO: airshow_unaligned.wmv]

# Demonstrations

# Learning a Trajectory

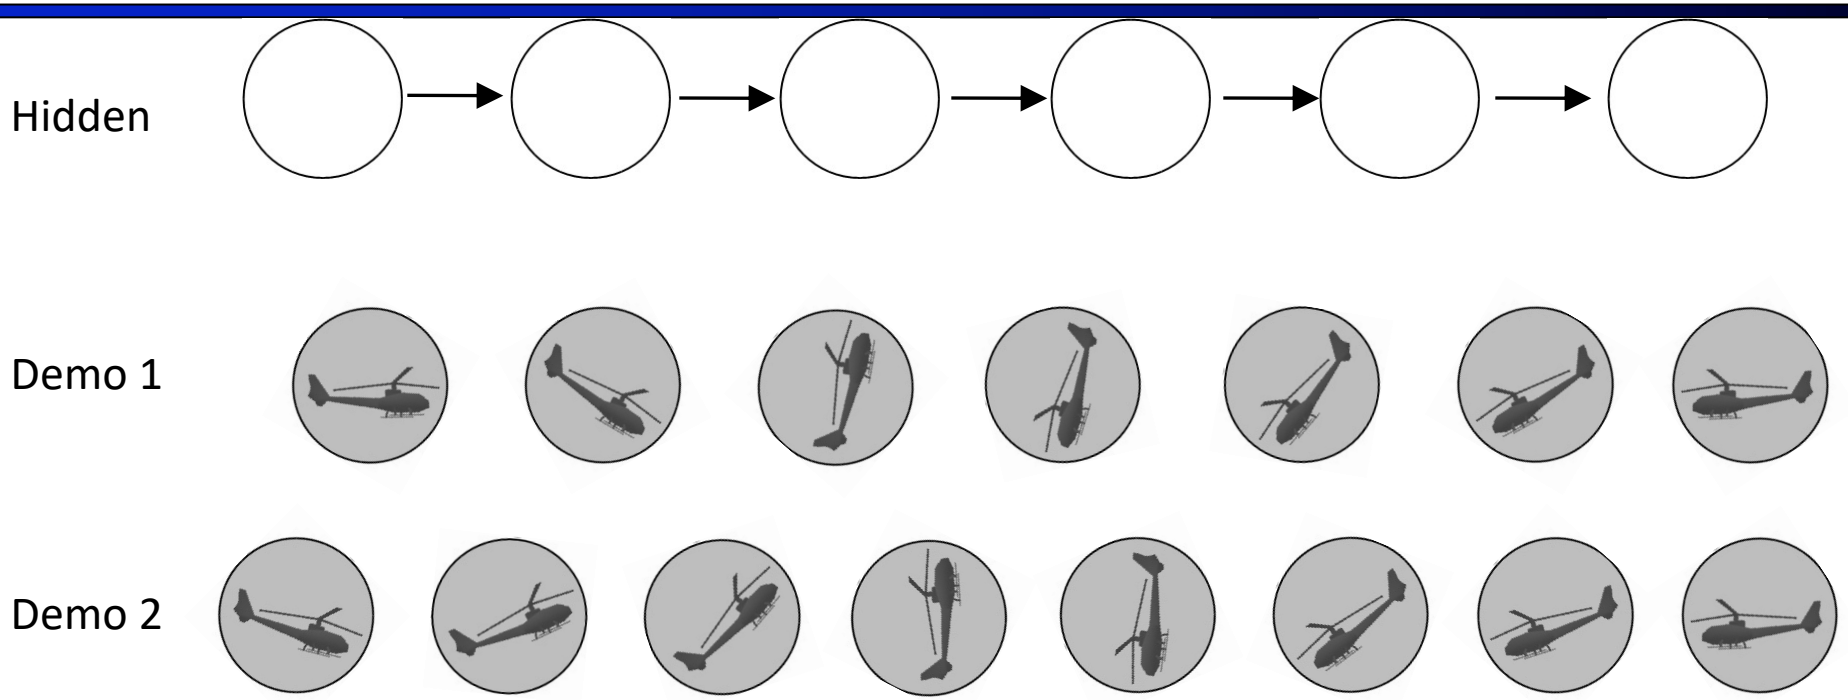


- HMM-like generative model
  - Dynamics model used as HMM transition model
  - Demos are observations of hidden trajectory
- Problem: how do we align observations to hidden trajectory?

Abbeel, Coates, Ng, IJRR 2010

# Probabilistic Alignment using a Bayes' Net

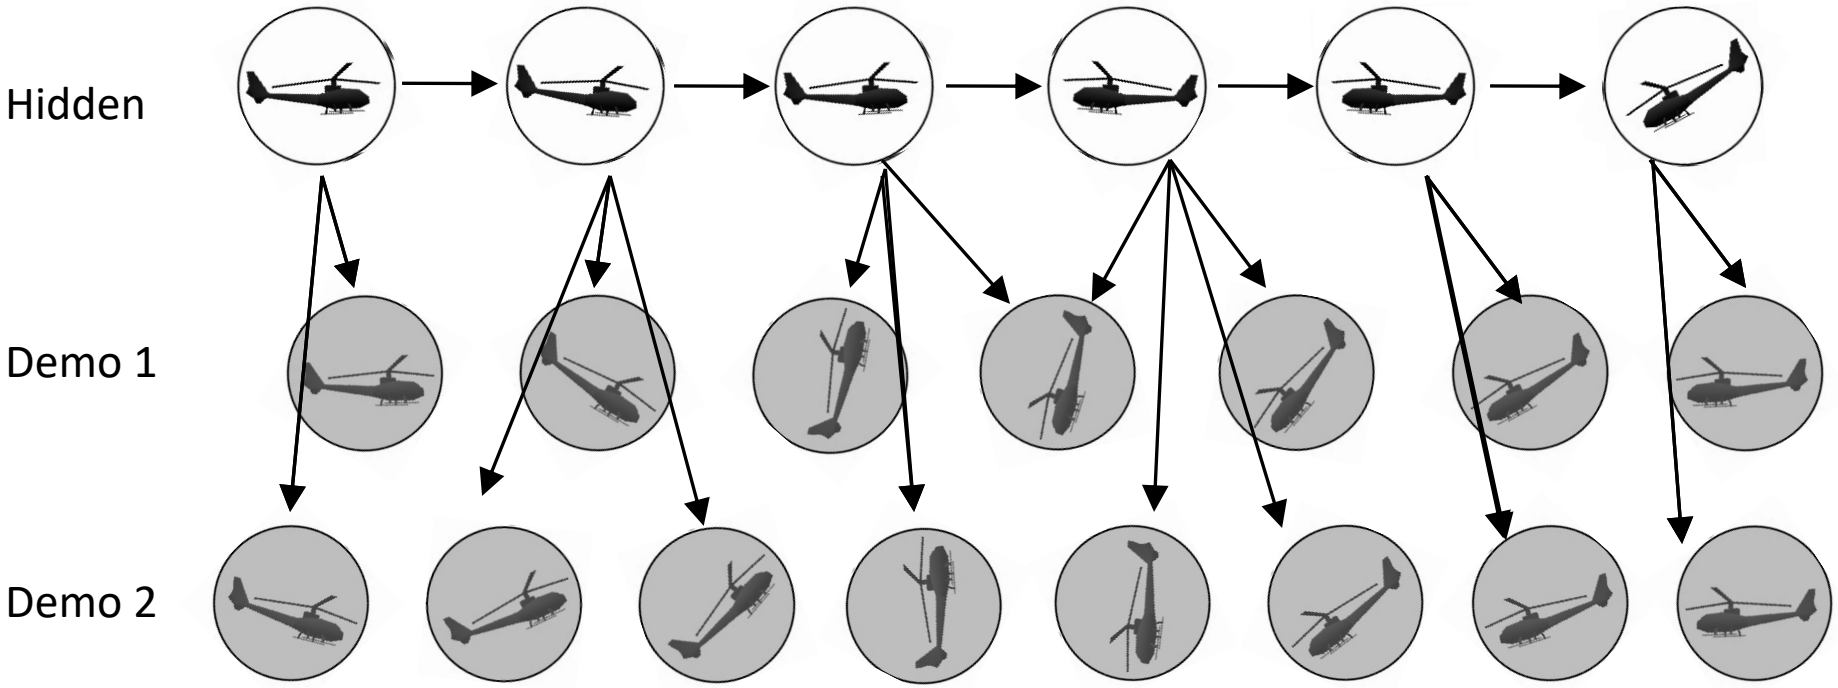


- Dynamic Time Warping

  (Needleman&Wunsch 1970, Sakoe&Chiba, 1978)

- Extended Kalman filter / smoother (=continuous HMM)

Abbeel, Coates, Ng, IJRR 2010

[VIDEO: airshow_unaligned.wmv]

# Aligned Demonstrations

# Alignment of Samples

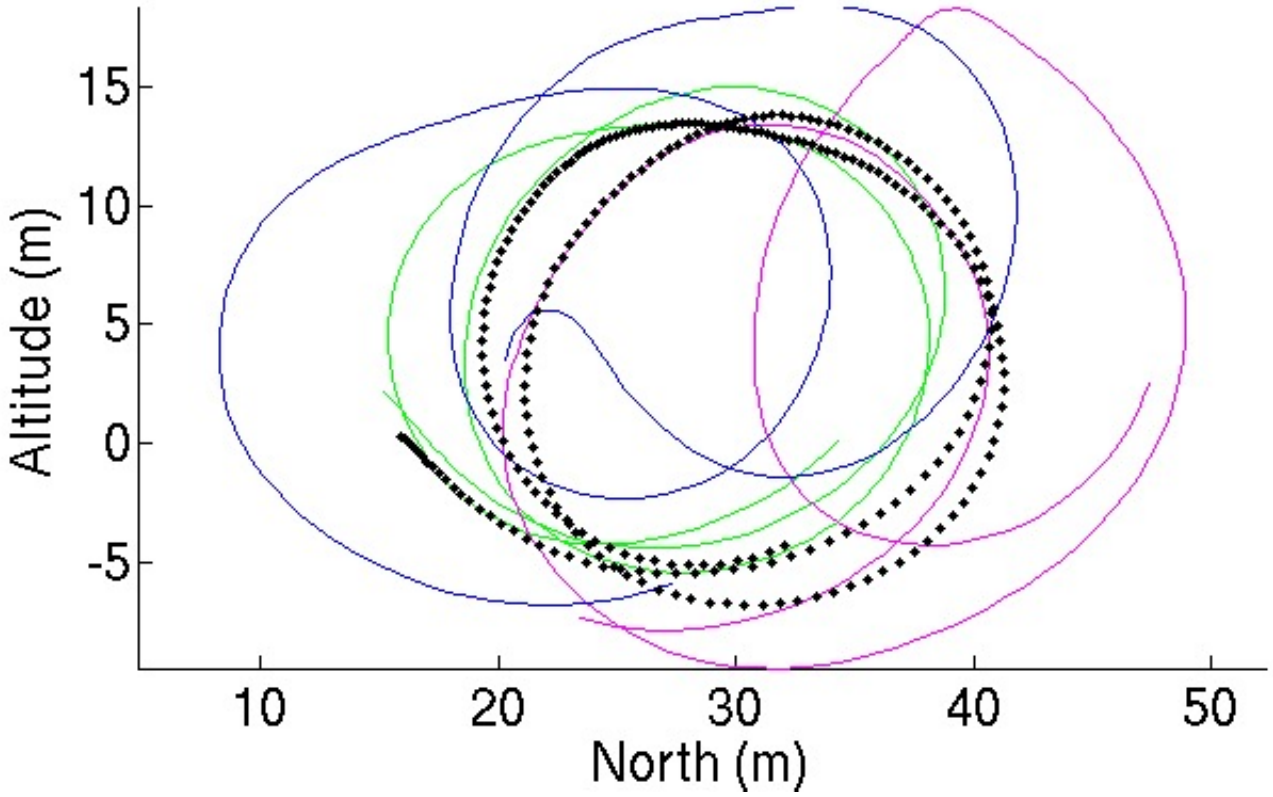


- Result: inferred sequence is much cleaner!

# Learned Behavior

[VIDEO: airshow_trimmed.wmv]

[Abbeel, Coates, Quigley, Ng, 2010]

# Legged Locomotion

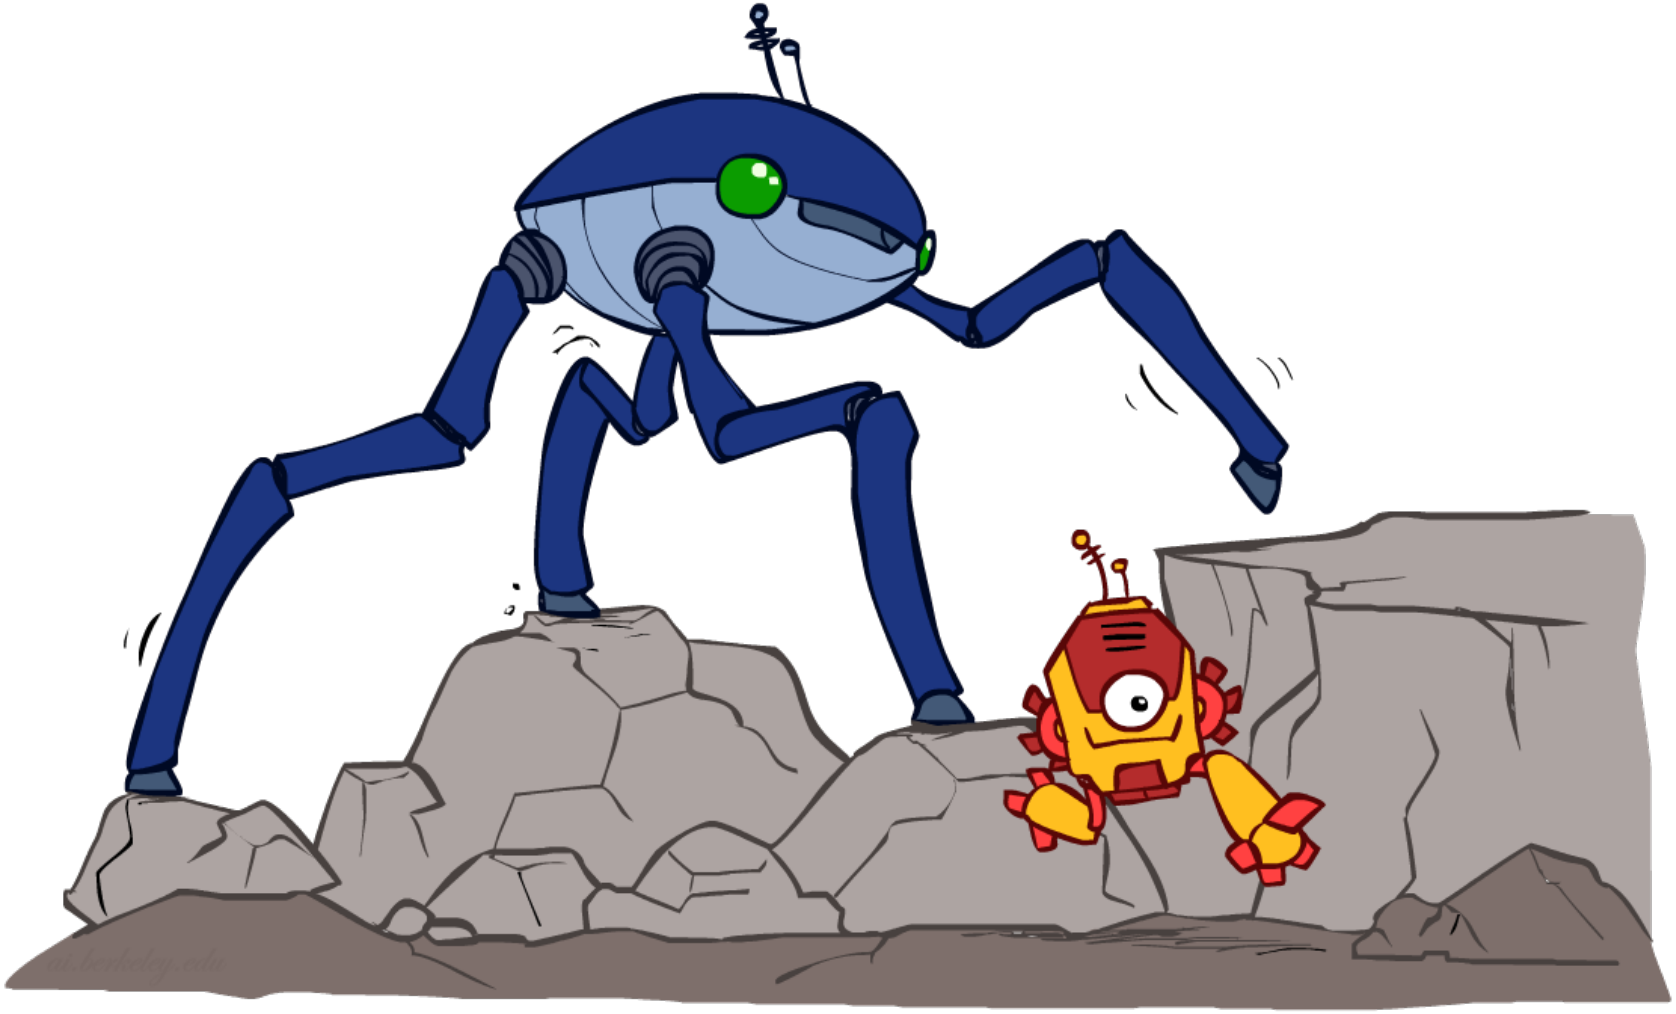

# For Perspective: Darpa Robotics Challenge (2015)

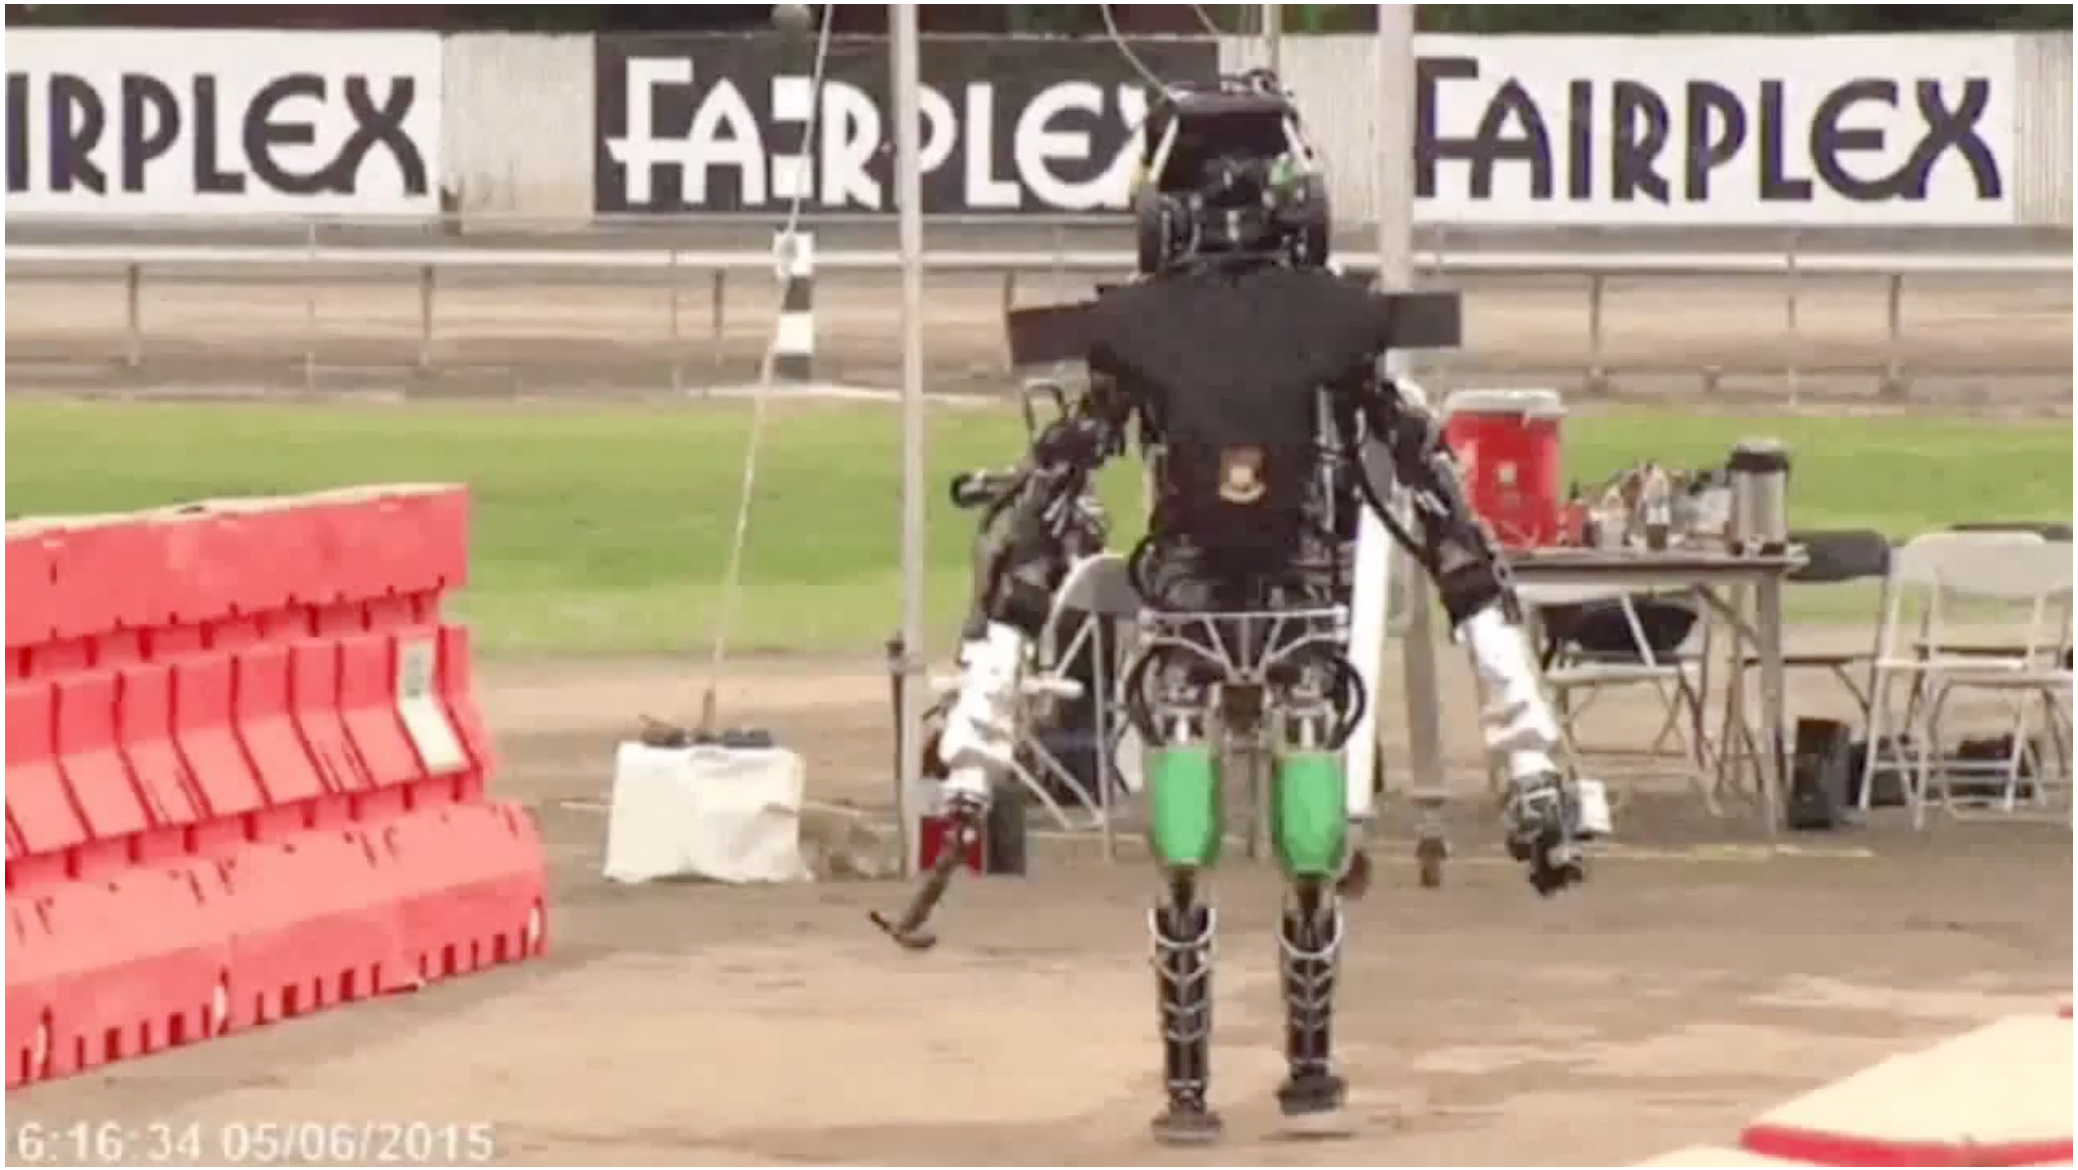

# Boston Dynamics

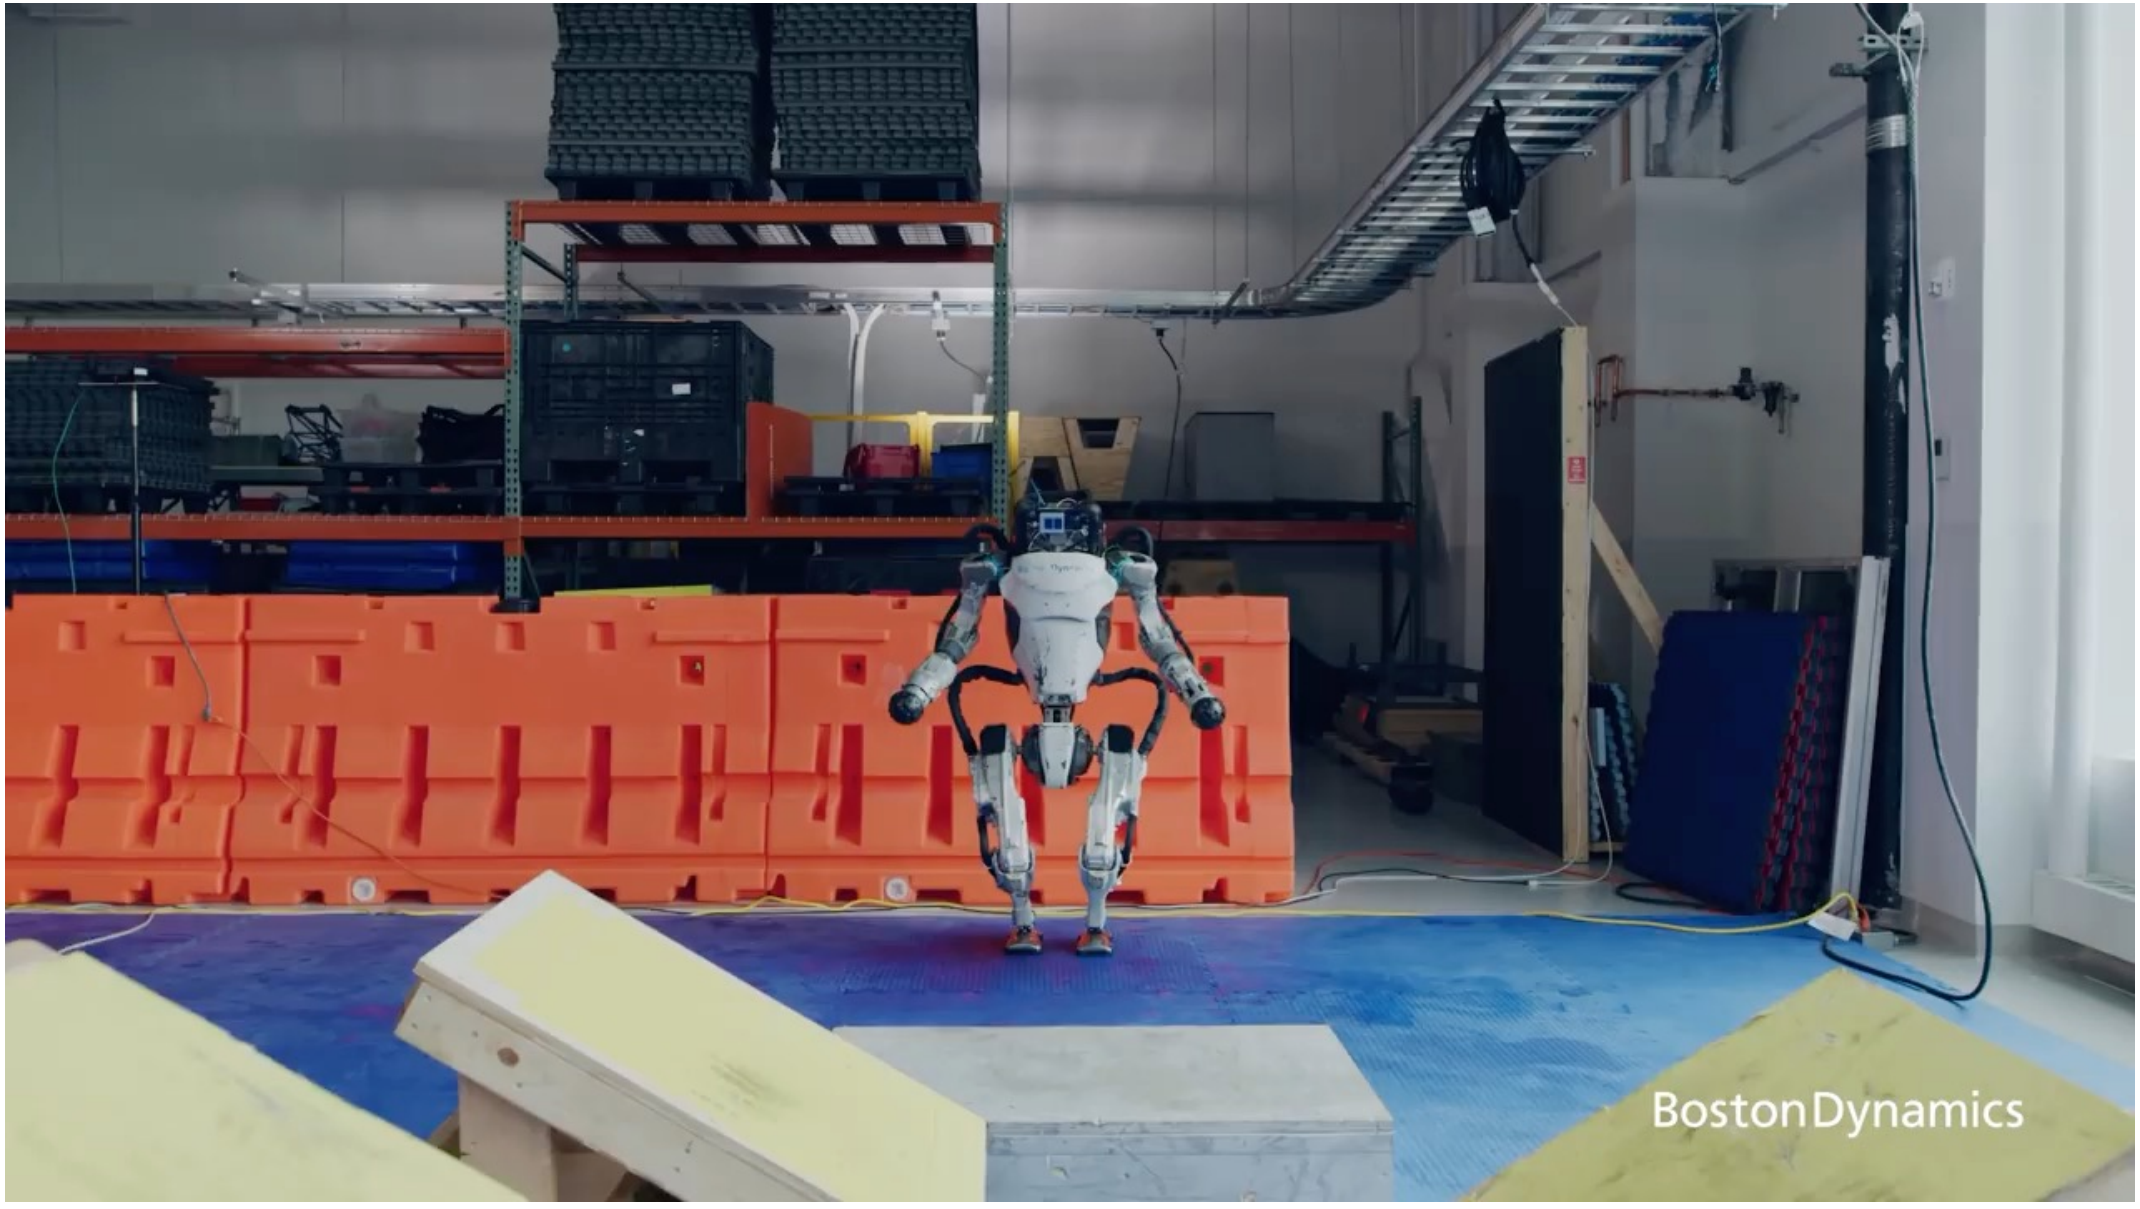

# Reinforcement Learning Locomotion

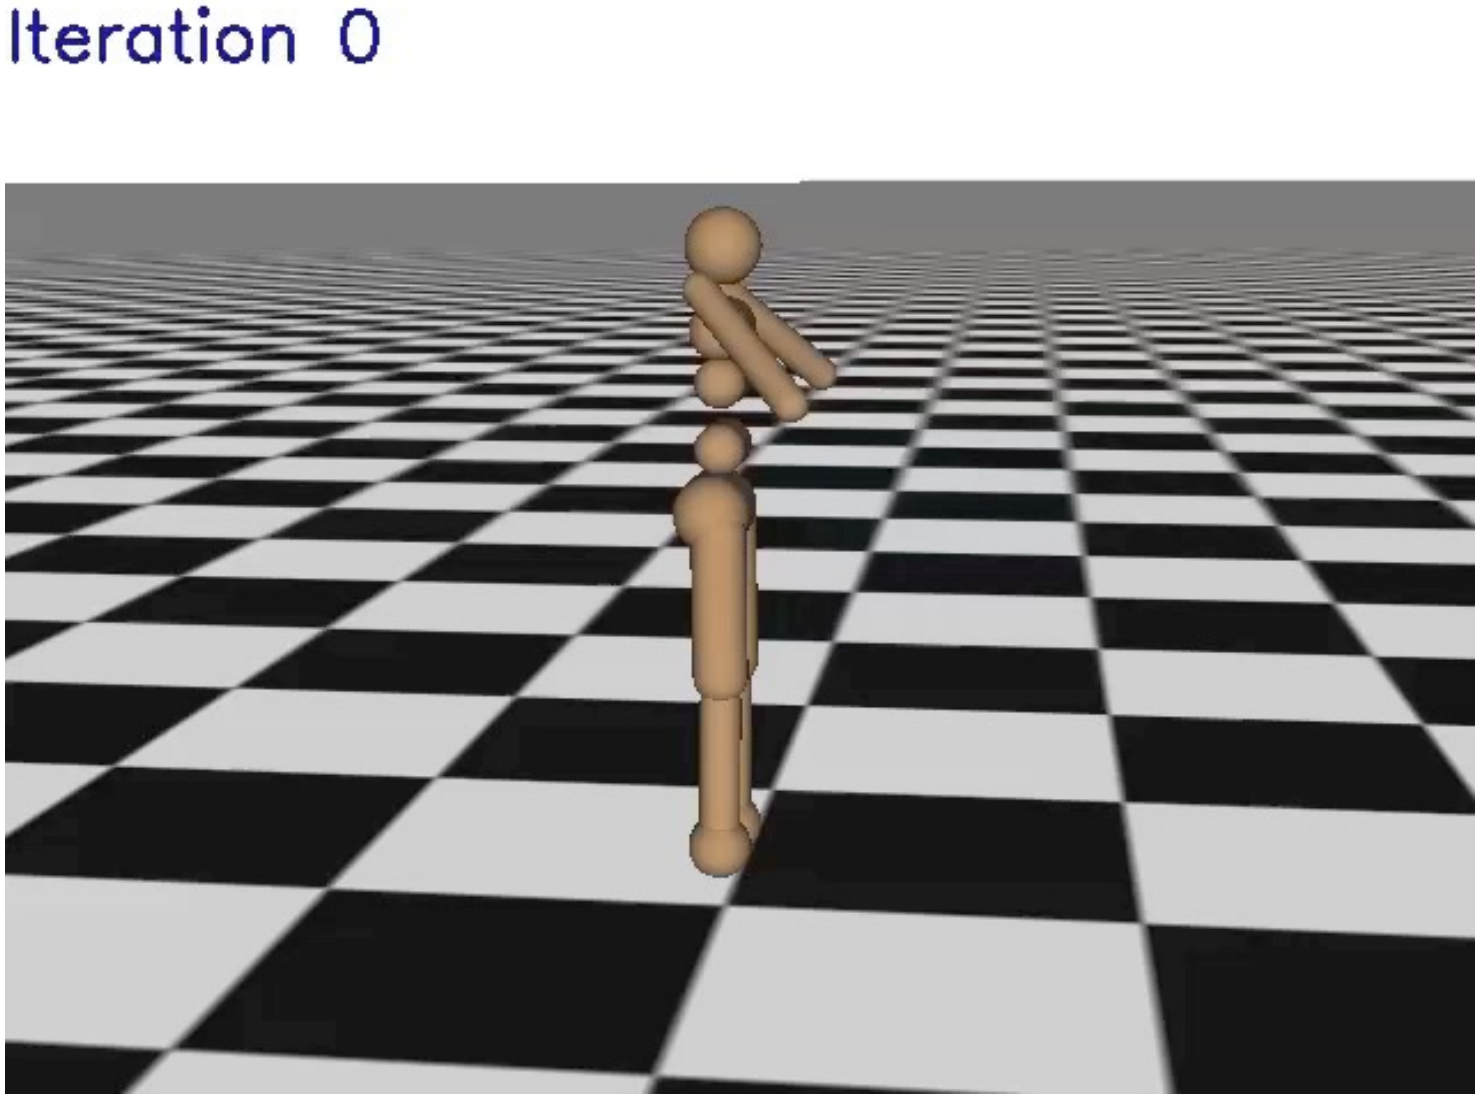


[Schulman, Moritz, Levine, Jordan, Abbeel, 2015]

# Reinforcement Learning a Wide Range of Skills


[Peng, Abbeel, Levine, van de Panne, 2018]

# Rapid Motor Adaptation for Legged Robots

Ashish Kumar
UC Berkeley

Zipeng Fu
CMU

Deepak Pathak
CMU

Jitendra Malik
UC Berkeley/FAIR

Robotics: Science and Systems 2021

# Autonomous Driving

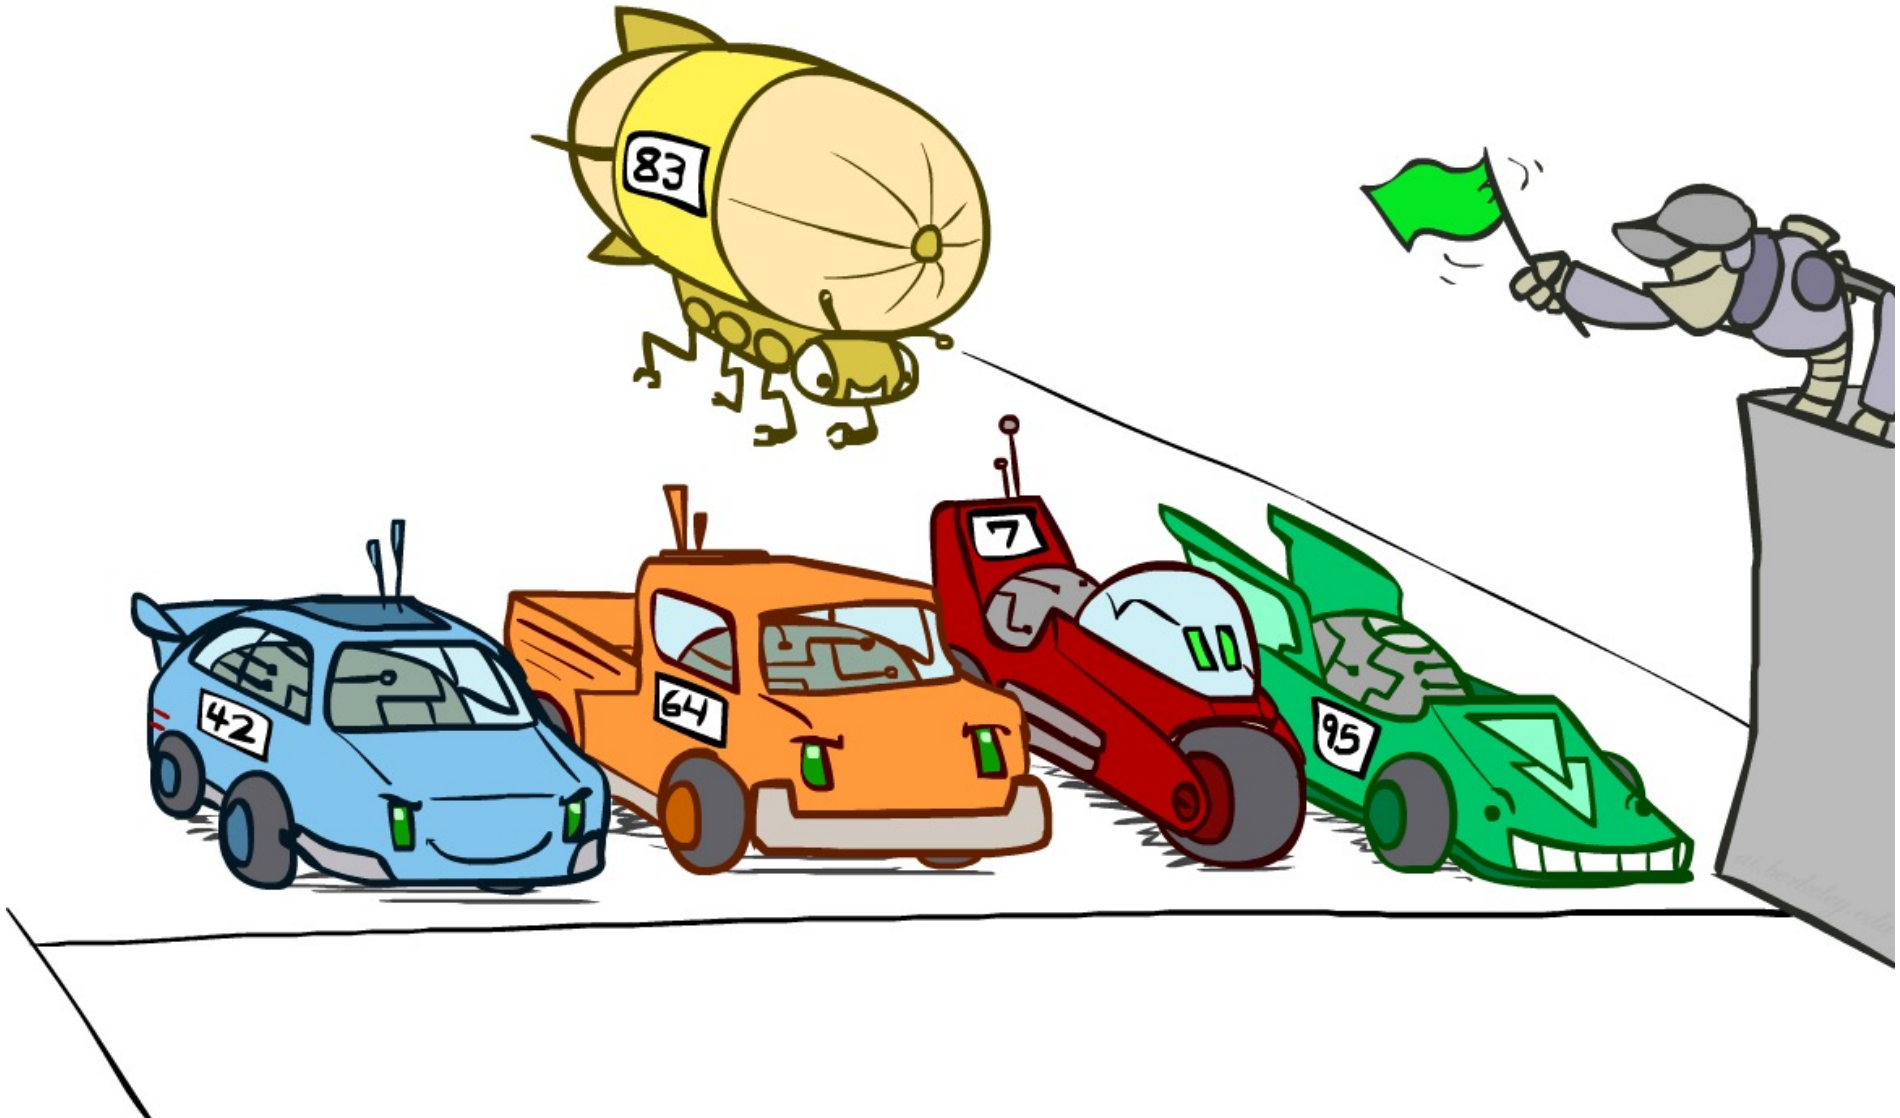

# Grand Challenge 2005: Barstow, CA, to Primm, NV




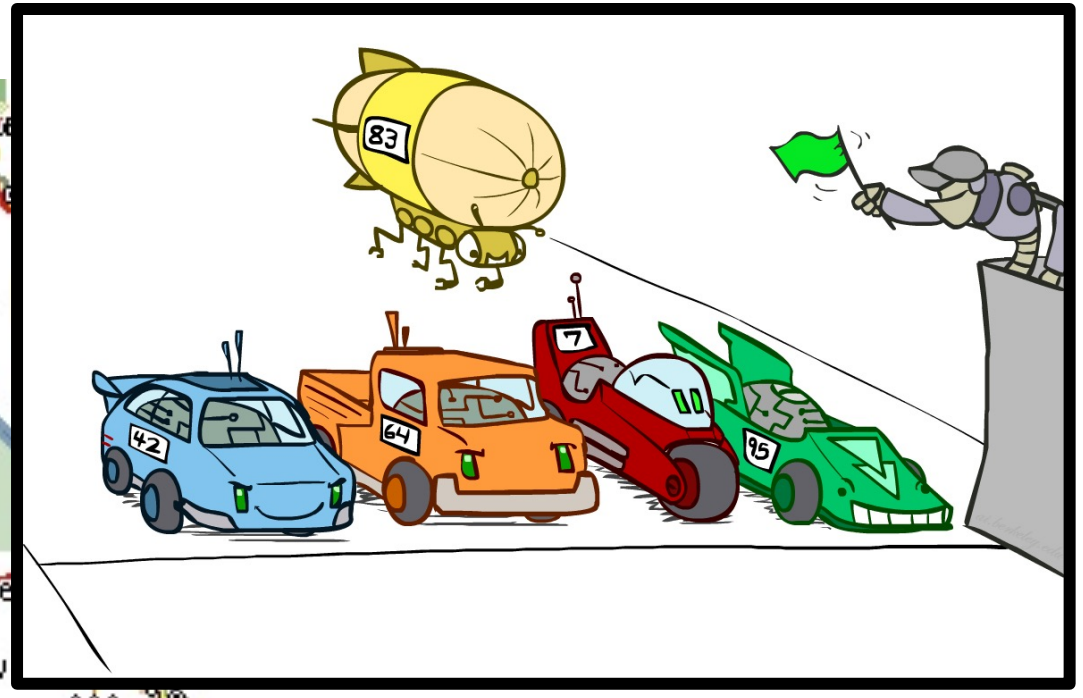


- 150 mile off-road robot race across the Mojave desert
- Natural and manmade hazards
- No driver, no remote control
- No dynamic passing

# Autonomous Vehicles

Autonomous vehicle slides adapted from Sebastian Thrun

# Grand Challenge 2005 Nova Video


[VIDEO: nova-race-supershort.mp4]

# Grand Challenge 2005 – Bad

[VIDEO: grand challenge – bad.wmv]

# An Autonomous Car

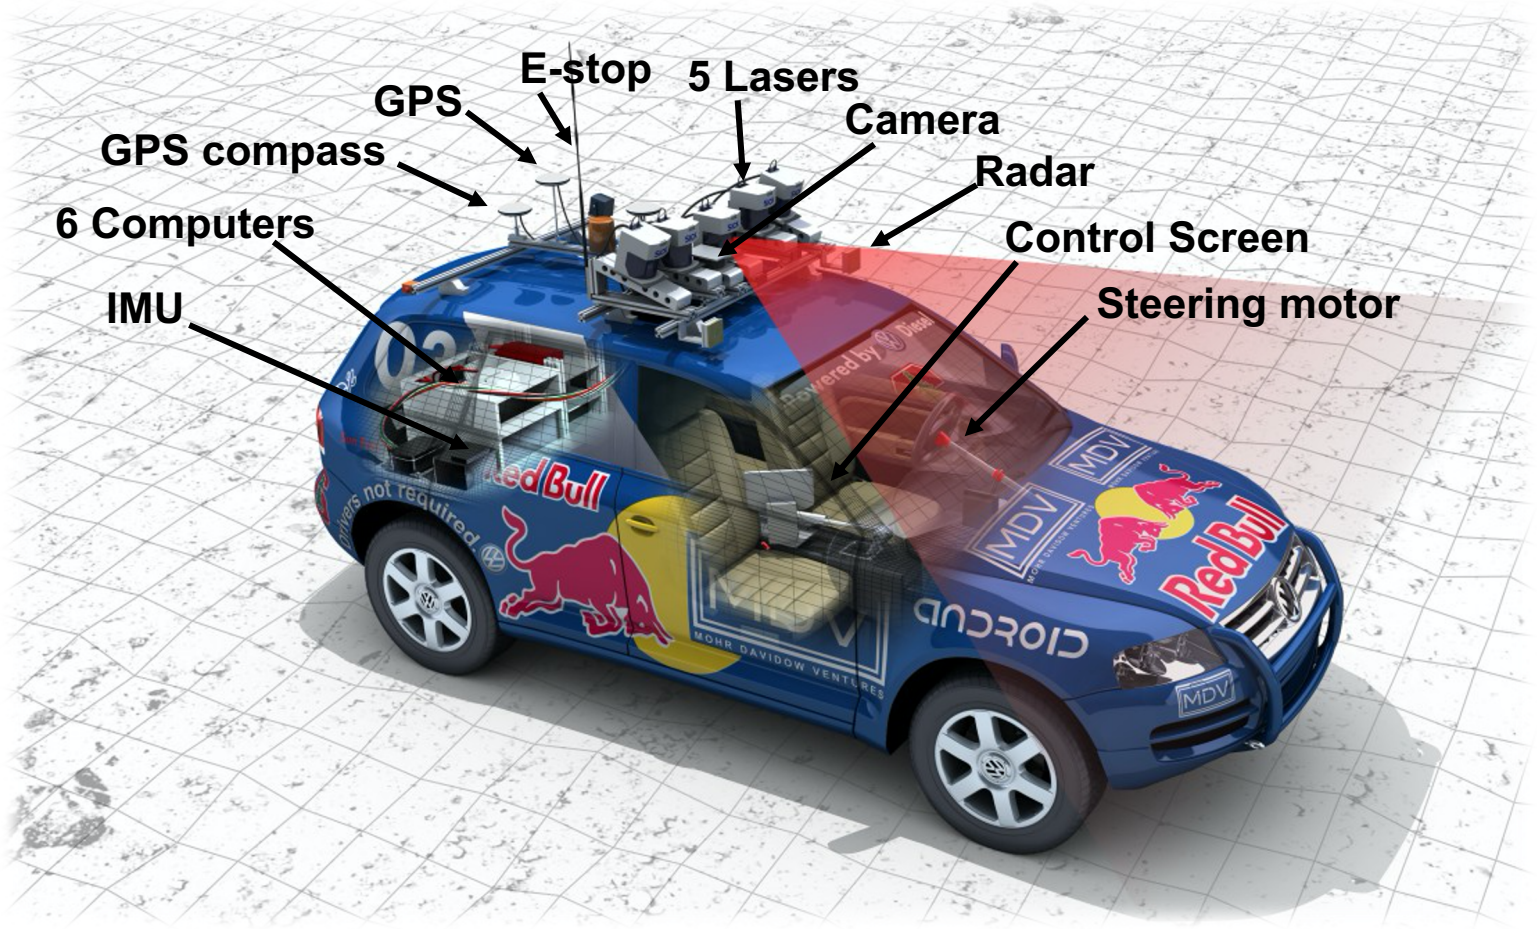

# Actions: Steering Control

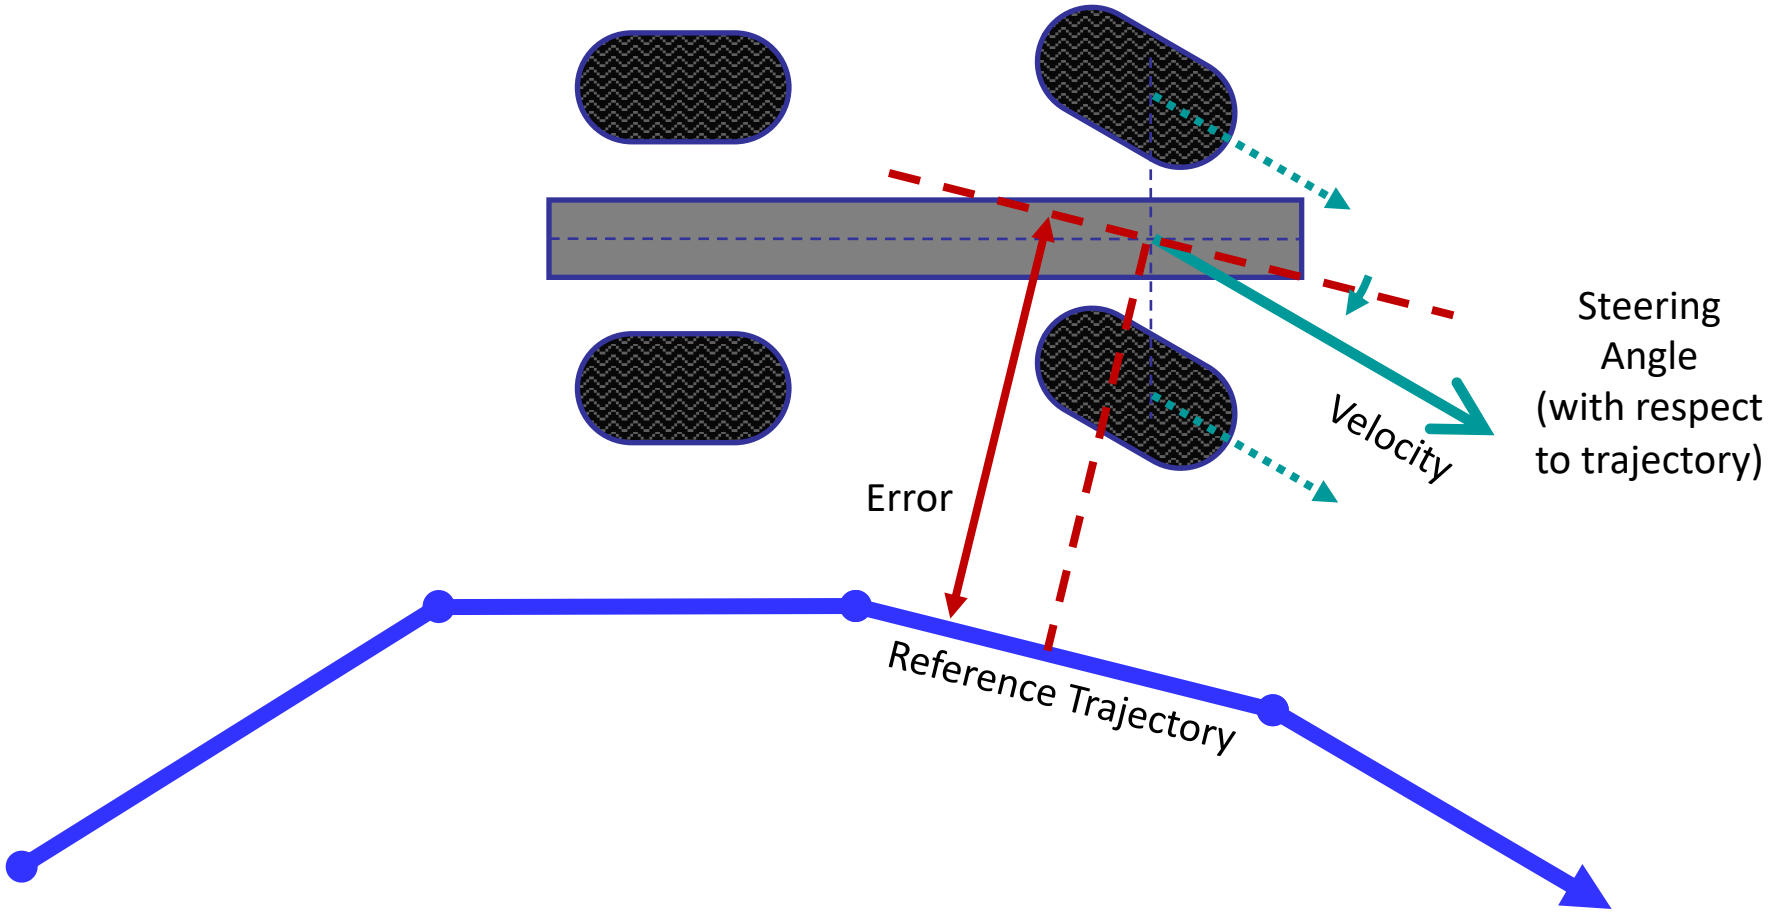

# Laser Readings for Flat / Empty Road

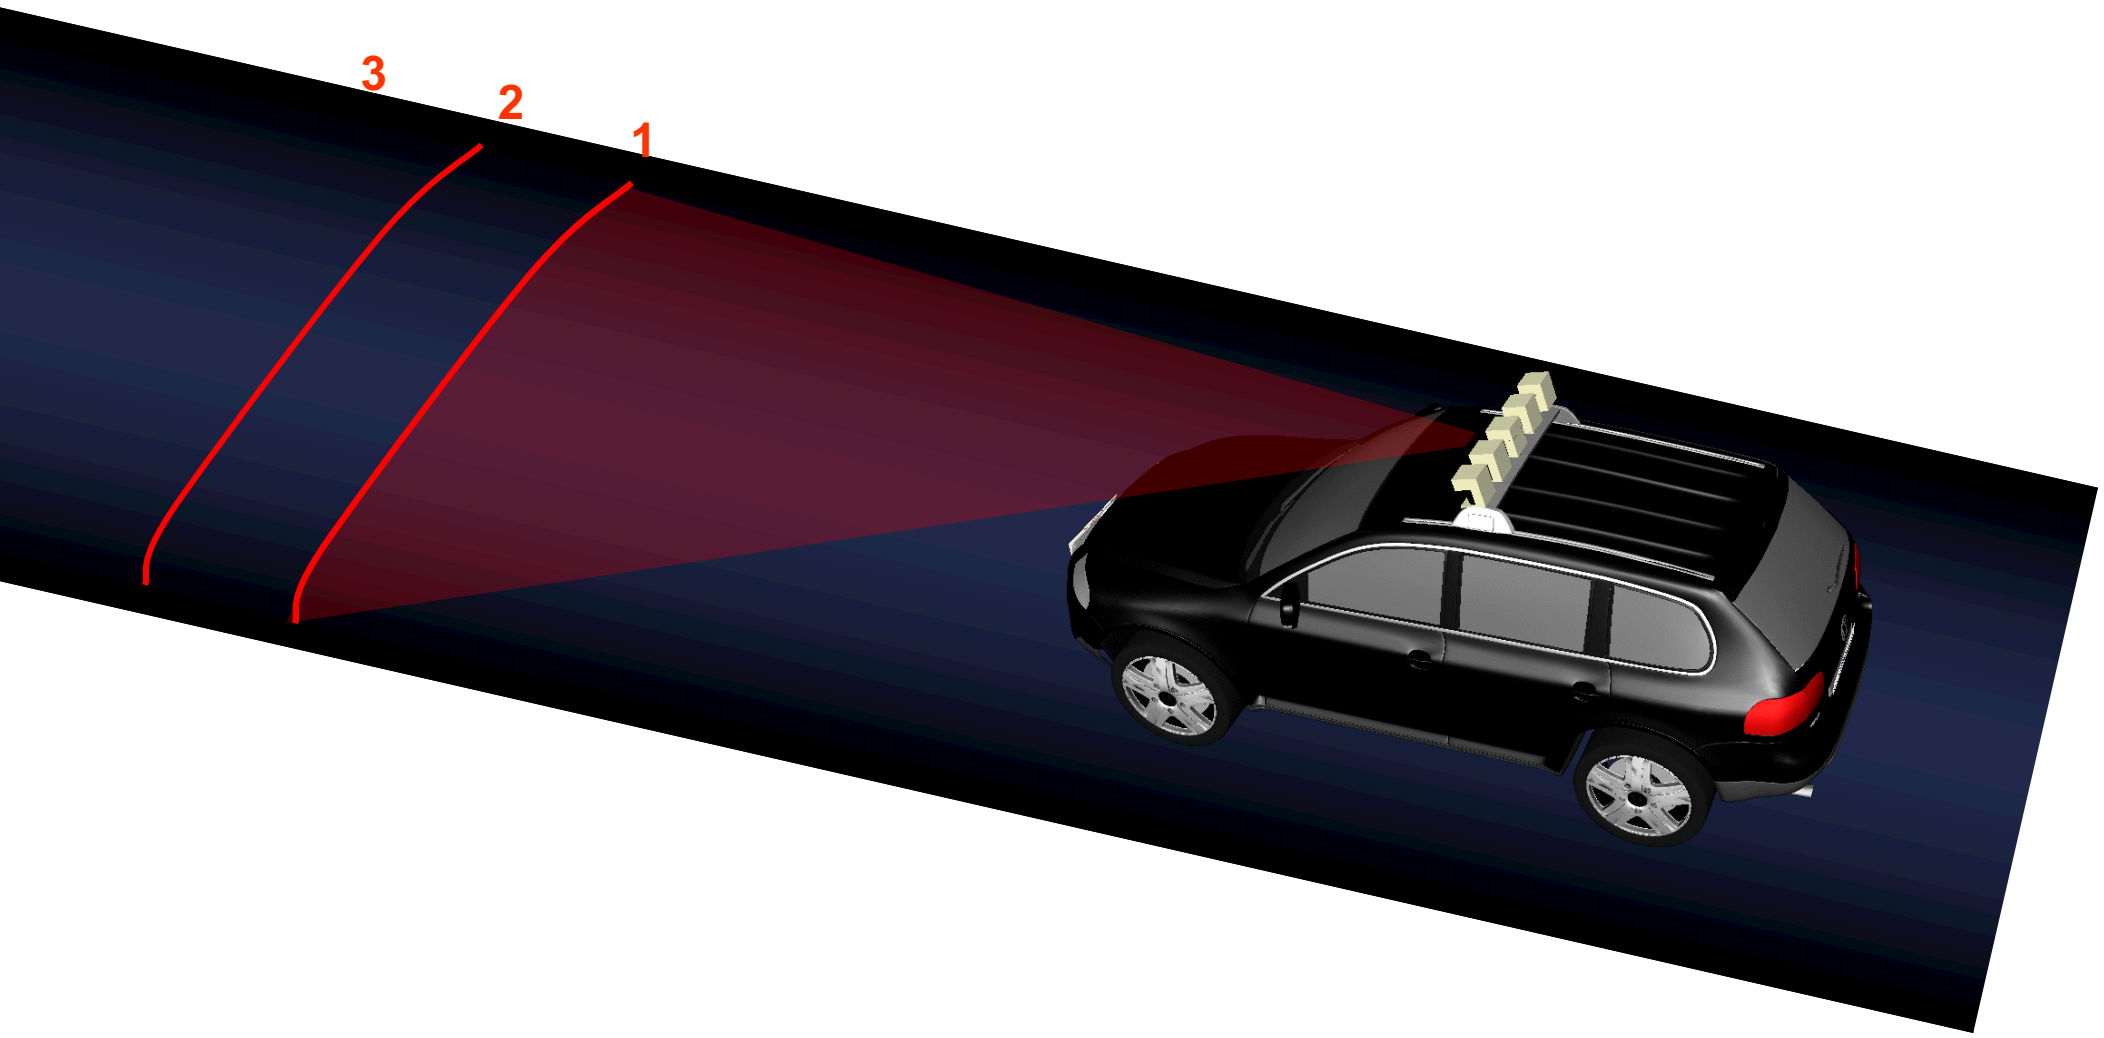

# Laser Readings for Road with Obstacle

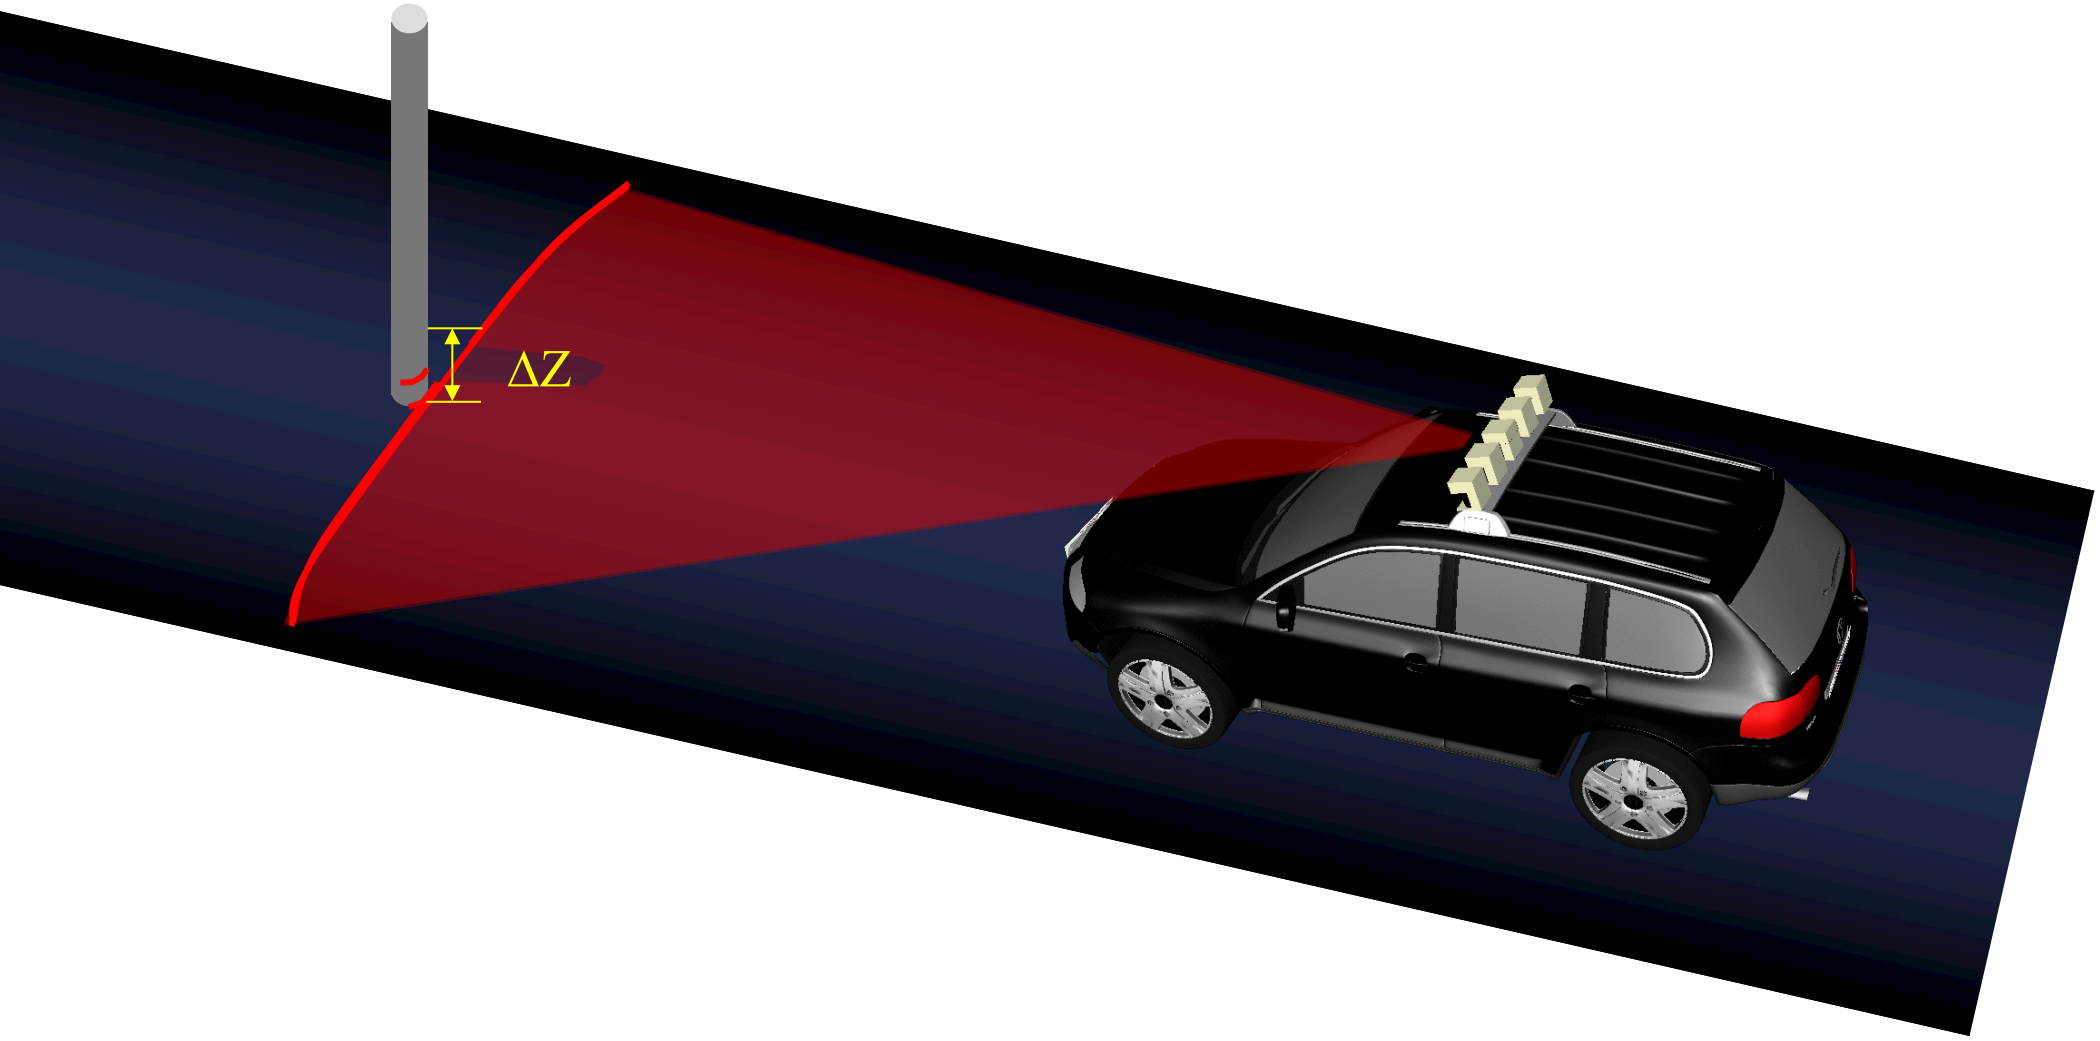

# Urban Environments

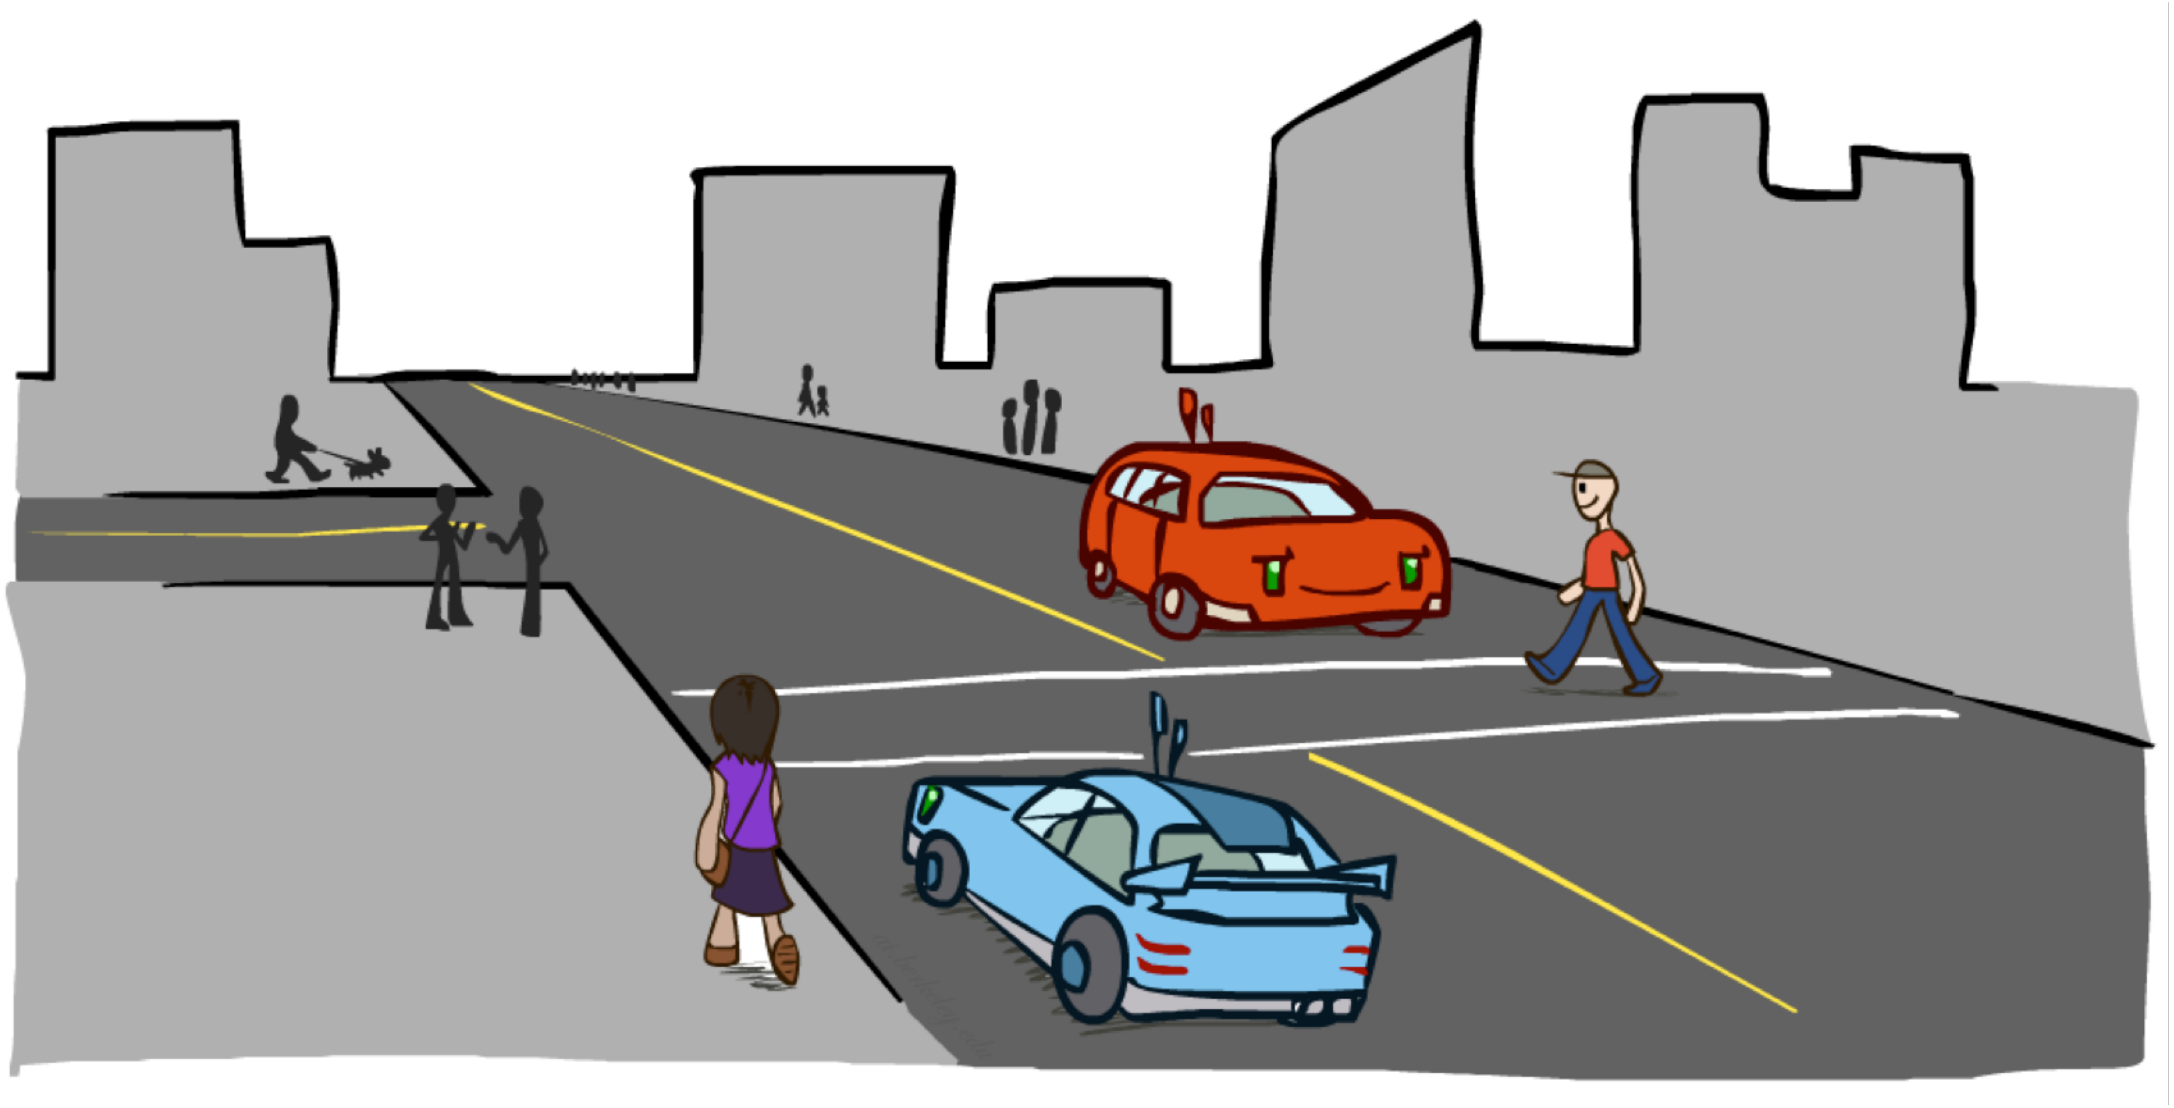

[VIDEO: ROBOTICS – gcar.m4v]

# Google Self-Driving Car (2013)

(mostly lidar)

# Recent Progress: NN Semantic Scene Segmentation

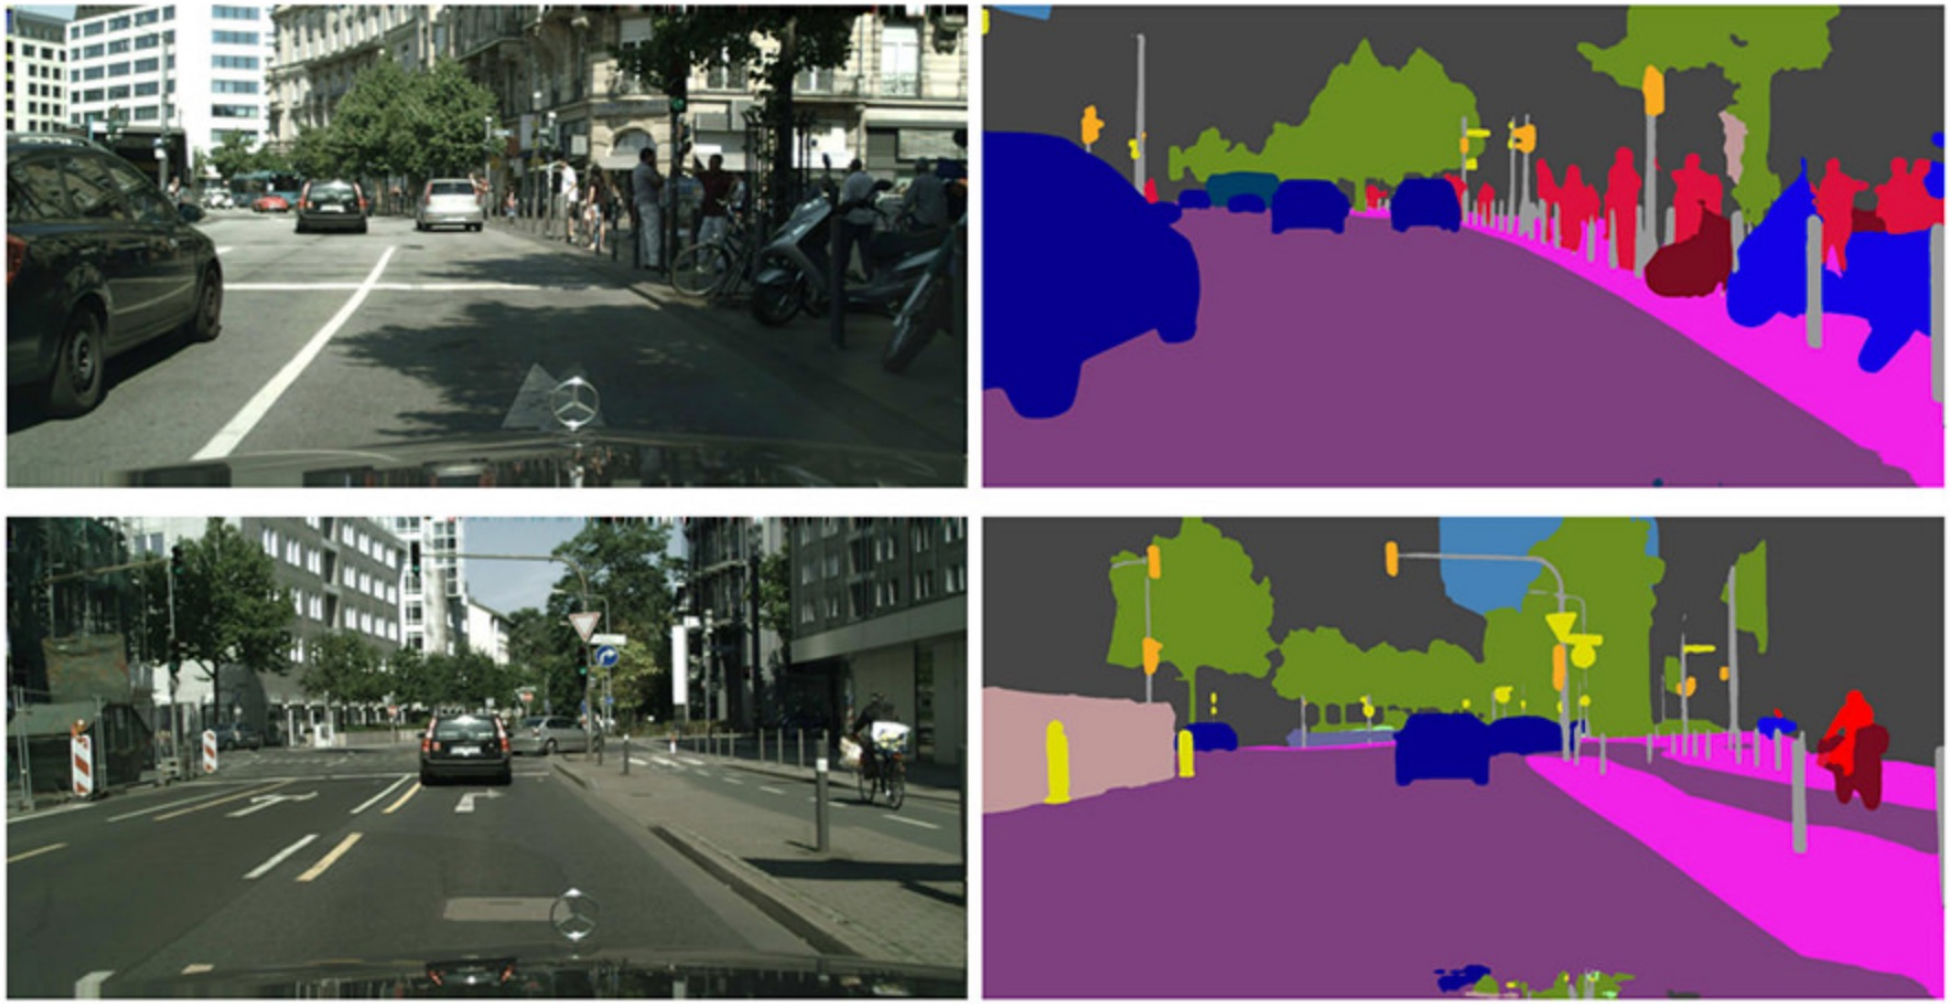

~ neural net classifies every pixel

PSPNet50

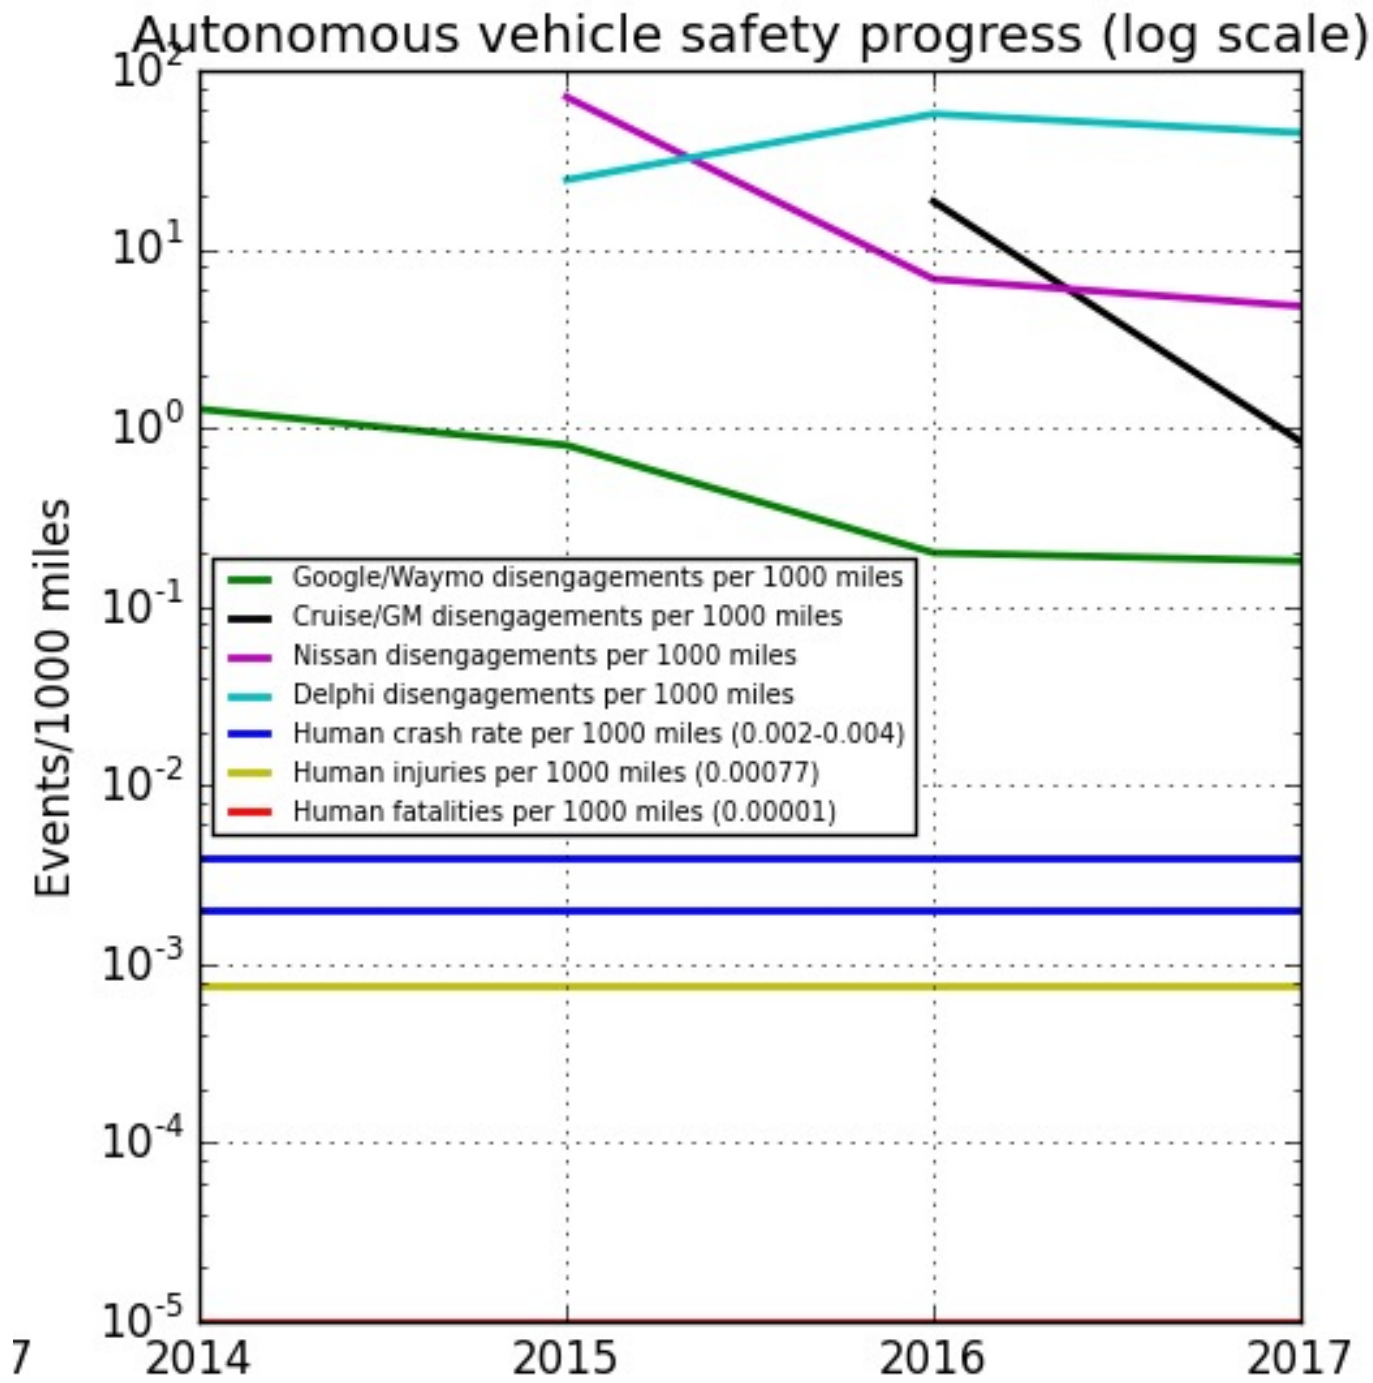
Autonomous vehicle safety progress (log scale)
Events/1000 miles
$10^{2}$
$10^{1}$
$10^{0}$
$10^{-1}$
$10^{-2}$
$10^{-3}$
$10^{-4}$
$10^{-5}$
2014
2015
2016
2017
Google/Waymo disengagements per 1000 miles
Cruise/GM disengagements per 1000 miles
Nissan disengagements per 1000 miles
Delphi disengagements per 1000 miles
Human crash rate per 1000 miles (0.002-0.004)
Human injuries per 1000 miles (0.00077)
Human fatalities per 1000 miles (0.00001)

# THE FUTURE OF TRANSPORTATION STACK

COMET LABS

## SERVICES

### ROUTE PLANNING
SPATIAL, CITIViZ, waze, CommuteMaster, X, moovit, Citymapper

### PARKING
Waypals, Cloud Parc, Aware, VoicePark, SpotHero, PEARL, seevider, Parko, FenSens, LUXE, anxsys

### CAR HAILING + POOLING
TESLOOP, lyft, RideCell, 滴滴

### OTHER: AFTERMARKET, REPAIR, RENTAL
CARVUE, upshift, WRENCH PATROL, zipcar, Getaround, URGENTLY, Tantalum, bitkar

## SAFETY & SECURITY

### PHYSICAL CAR & DRIVER SAFETY +ACCIDENT DETECTION
drivebetter, splitsecnd, neXar, GUARDIAN Optical Technologies, phrame, LyfeLens, SMART AUTOLABS, tourmalinelabs, iproc, eLSys, FORESIGHT

### EMOTION, FATIGUE & ALCOHOL DETECTION + DISTRACTION AVOIDANCE
CORTEX LABS, eyeris, ALMA care, CARMA, saverone, SMART wheel, Radiomize, EASYC, carvi

### CYBERSECURITY
CYMOTIVE, cognomotiv, VisualThreat, ARGUS, Olympus Sky, Karamba Security, TrIAlium, Arilou

### INTRUSION, TRACKING & RECOVERY
trackinstop, nonda, gmt Connect, ROADEYE, caruma

## IN-CAR INTELLIGENCE + ASSISTANCE

### VEHICLE DIAGNOSTICS & PREDICTIVE MAINTENANCE + SENSOR-BASED VEHICLE SAFETY
TrueMotion, BRAIQ, IQ, metromile, LIGHT METRICS

### PASSENGER-FOCUSED SENSORS (INCLUDING USAGE-BASED INSURANCE)
Cartics, Zubie, nauto, SONICLUE, PREDICTSYSTEMS, NEBULA, engie

### INFOTAINMENT + DISPLAY
link motion, HUDIFY, CARROBOT, FormCar, swedspot, OverDrive, ROAV, FICONIC SOLUTIONS, RIGHTWARE, PELAGICORE, CLOUDCAR

### PERSONAL / VOICE ASSISTANCE
navideck, AUTOMATIC, dash, GRIN, cariq

### NAVIGATION ASSISTANCE + PEDESTRIAN ANALYSIS & COMMUNICATIONS
WAYRAY, PERCEPTIVE AUTOMATA, navdy, hudly

## AUTONOMY

### AUTOMATION SYSTEM
AutoX, drive.ai, POLYSYNC, OXBOTICA, oryx, ARGO AI, AIMOTIVE, PILOT, cruise, nuTonomy, vector.ai, MOMENTA, cognivue, AIDRIVING

### MAPPING, SIMULATION, & IMAGE RECOGNITION / ANNOTATION
ivi, minds.ai, NOMOKO, Mapper.ai, CARMERA, TOMTOM, IMPROBABLE, UDACITY, DEEPMAP, understand.ai, Civil Maps

### AUTONOMOUS VEHICLE MAKER + TOOLS
nuro, WAYMO, ZOOX, FARADAY FUTURE

## INFRASTRUCTURE + CONNECTED CAR

### SENSOR NETWORKING INFRASTRUCTURE (V2V, V2X) - LPWA, CELLULAR, WIFI
DIGIMONDO, SAVARI, ZiFiSense, EVVOS, Actility

### CONNECTED CAR - DATA, PLATFORM, SOFTWARE
ACTIVESCALER, INRIX, SPARC, mobiliuz, carvoyant, vinli, otonomo, bright box, VEHICLE DATA SCIENCE, mojio

### FLEET + TRAFFIC MANAGEMENT
BESTMILE, Cloud Your Car, synovia, Fleetmatics, urban engines, Motionlogic, Travel AI, Automile, Zendrive, GLYDEL, locoNav, Immense Simulations, MOTORWAY BUDDY, mentis, FR, Brisk Synergies

### OTA CAR SOFTWARE UPDATION + SMART PHONE ENABLEDTELEMATICS
FIXD, DRIVEBOT, VOYOMOTIVE, ZENE, aryaq

## INTELLIGENT MANUFACTURING

### NEW/ADVANCED MATERIALS
AQUARIUS, NAWA TECHNOLOGIES, SIRRUS, NANOSTEEL

### RAPID PROTOTYPING - 3D PRINTING, MODULARIZATION, OPEN SOURCE
Carbon, nanocore, DIVERGENT, OSVehicle, stratasys, VADER

### ADVANCED / AUTOMATED ASSEMBLY LINE
Grabit, SYMBIO, CLEARPATH, ROBOMOTIVE, rethink robotics.

### MATERIAL CHARACTERIZATION & TESTING
SynTouch, SASTRA, VIDI, OPTOFIDELITY

## ONBOARD SENSORS

### LOCATION - GIS, PRECISION POSITIONING, PATH PLANNING
swift NAVIGATION, radiosense, EXO TECHNOLOGIES

### VISION / CAMERA
DEEP VISION, Carnegie Robotics, machines with vision, ROADSENSE, DEEPSCALE, Chronocam, VAYAVISION

### LIDAR
faceI, INNOVIZ, PHANTOM INTELLIGENCE, VOYAGE, QUANERGY, AEYE, TRILUMINA, LeddarTech, Velodyne LiDAR, STROBE

### RADAR
OCULii, TOPOSENS, RFISEE, 6TH SENSE, SDS

## SPECIALTY VEHICLES

### 2-WHEELERS
mobike, smarthalo, superpedestrian, Zagster

### PUBLIC TRANSPORT
navya, Varden Labs, AURO, EASY MILE

### TRUCKS / FREIGHT
JAYBRIDGE ROBOTICS, OTTO MOTORS, PELOTON, TESLA, FLEXPORT, TRUELITE TRACE, STARSKY ROBOTICS

### FLIGHT
WRIGHT ELECTRIC, ZUNUM Aero, ZEE, JOBY, Kitty Hawk, LILIUM AVIATION, AEROMOBIL

### OTHER: HYPERLOOP, PERSONAL MOBILITY
hyperloop one, ninebot, WHILL

Search

Bloomberg Businessweek

August 17, 2021, 3:00 AM PDT  Corrected August 19, 2021, 2:25 PM PDT

# Waymo Is 99% of the Way to Self-Driving Cars. The Last 1% Is the Hardest

Vox recode

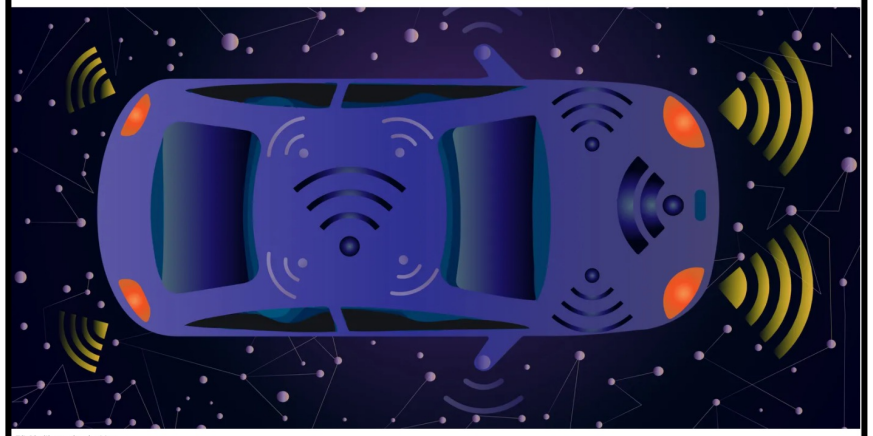

Efi Chalikopoulou for Vox

## Why we're still years away from having self-driving cars

Self-driving cars were expected to roll out by 2021. Here's what we need to solve and build first.

By Eric Adams | Sep 25, 2020, 3:40pm EDT

The New York Times

# *The Costly Pursuit of Self-Driving Cars Continues On. And On. And On.*

Many in Silicon Valley promised that self-driving cars would be a common sight by 2021. Now the industry is resetting expectations and settling in for years of more work.

# Personal Robotics

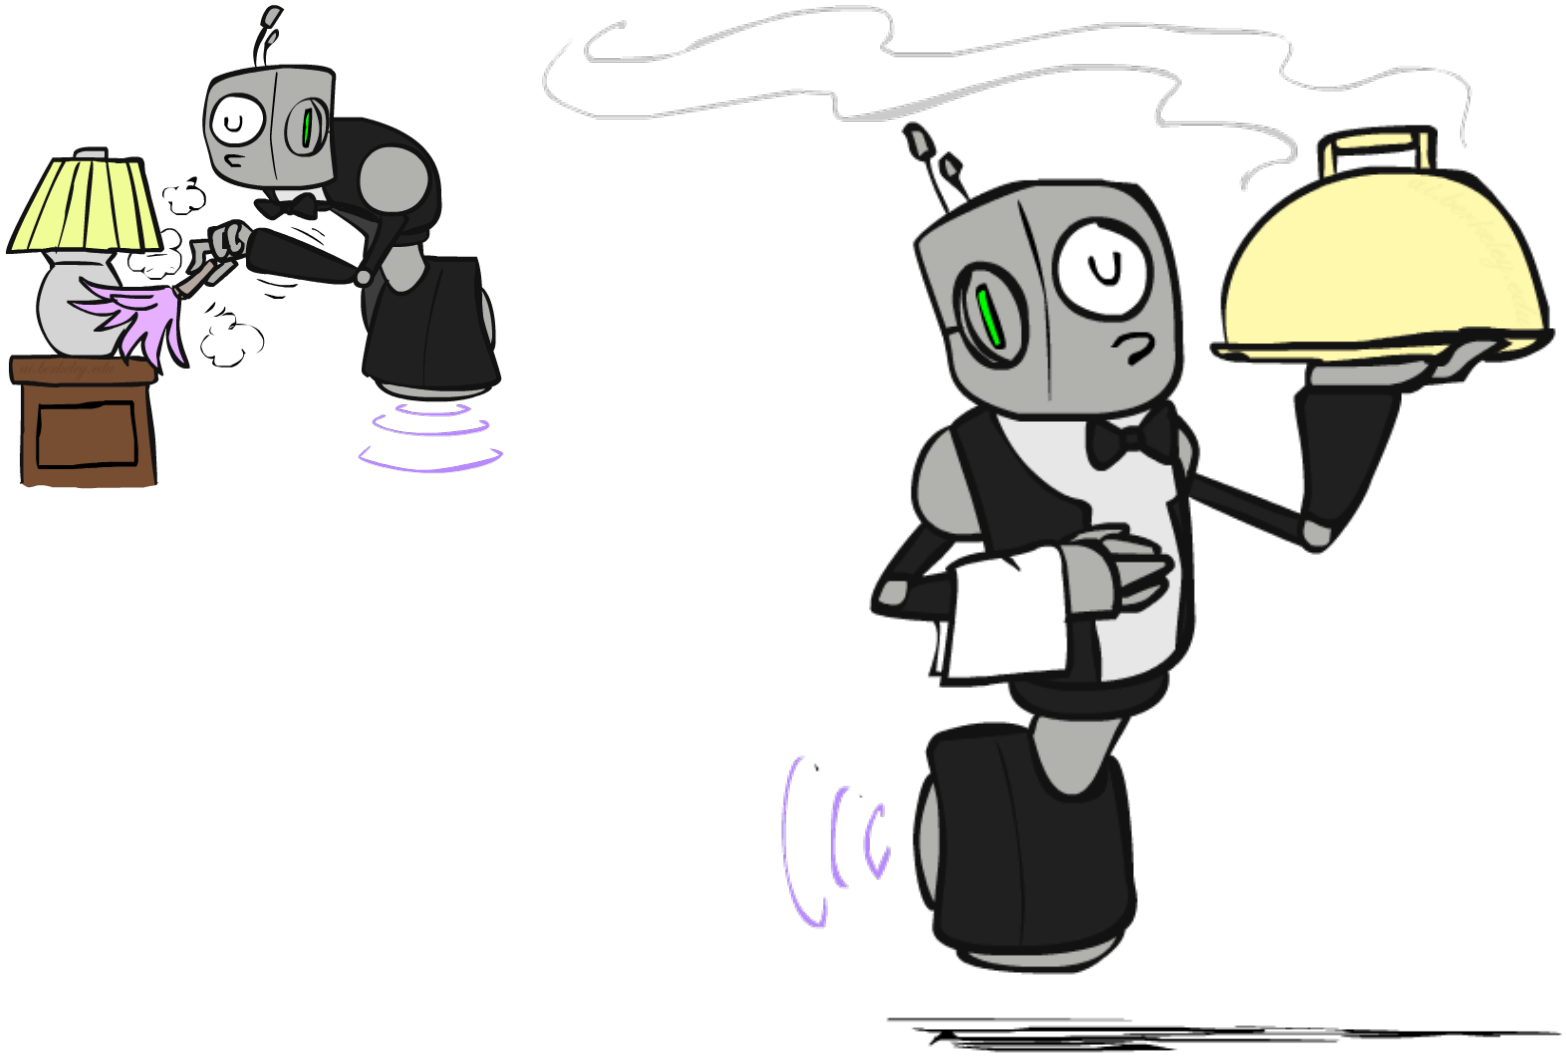

# PR-1

[Wyrobek, Berger, van der Loos, Salisbury, ICRA 2008]

# Challenge Task: Robotic Laundry

# How about a range of skills?

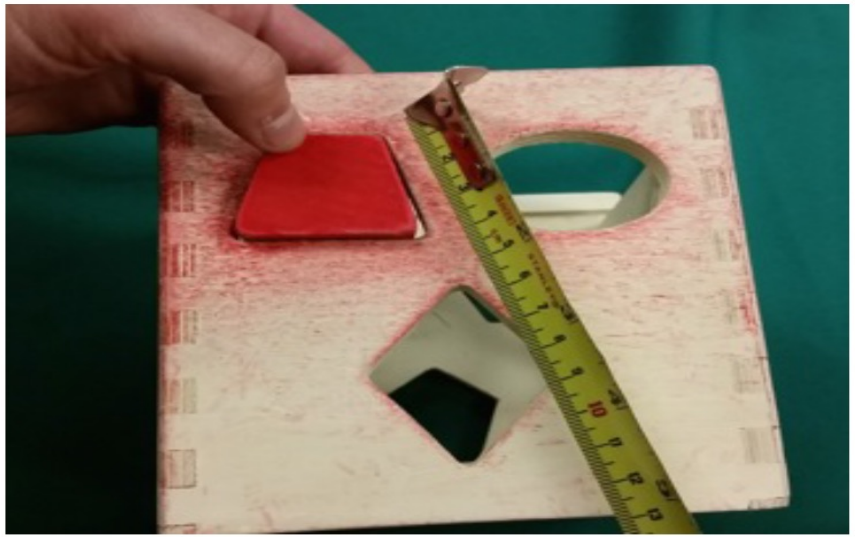

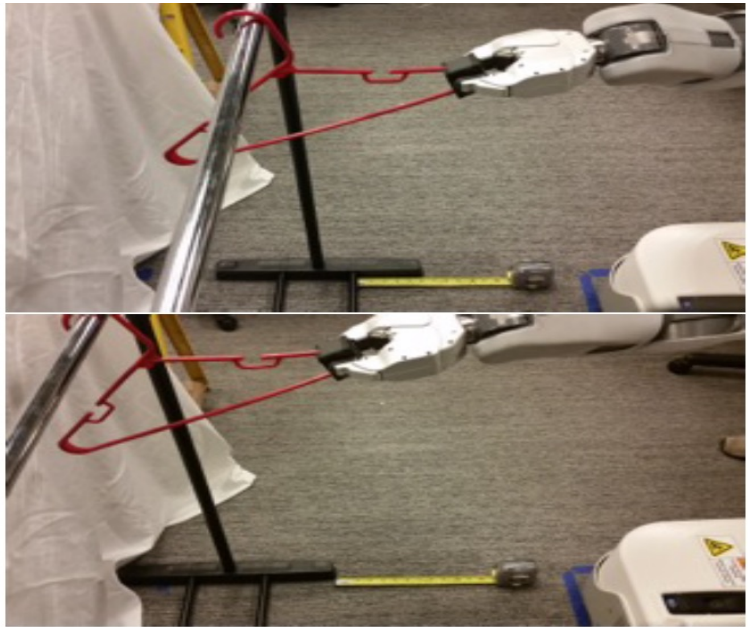

[Levine*, Finn*, Darrell, Abbeel, JMLR 2016]

# Reinforcement Learning

[Levine*, Finn*, Darrell, Abbeel, JMLR 2016]

# Learned Skills

[Levine*, Finn*, Darrell, Abbeel, JMLR 2016]


[Levine et al, 2016]

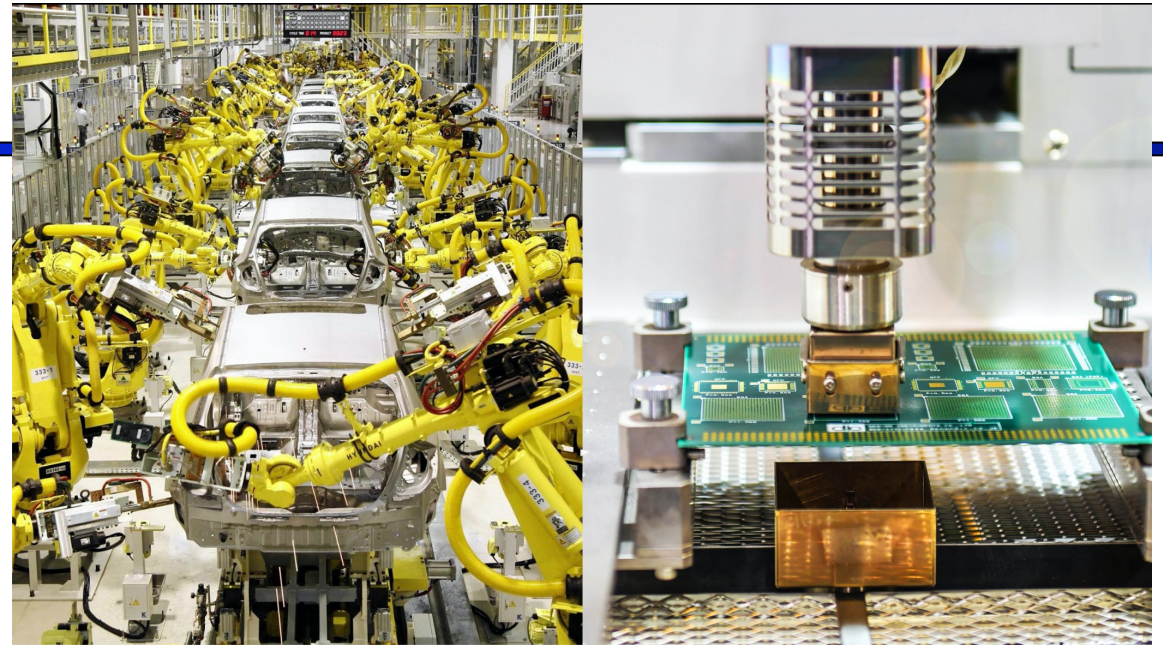

# Traditional Robotic Automation

Highly structured environments

Pre-programmed, deterministic

# Next Generation: AI Robotic Automation

Less structured environments

Intelligent, reactive behaviors

# Many start-ups are going after the opportunities

## Recycling

AMP ROBOTICS
ZENROBOTICS
SMARTER SORTING
WASTE ROBOTICS
greyparrot

## Farming

AGROBOT ROBOTIC HARVESTERS
abundant ROBOTICS
GRoboMac Green Robot Machinery Private Limited
FF Robotics
METOMOTION
TRAPTIC

## Warehousing

covariant
BERKSHIRE GREY
DEXTERITY
Ambidextrous
nomagic
IAM ROBOTICS
nimble.ai
KINDRED
RIOS INTELLIGENT MACHINES
SOFT ROBOTICS
vicarious
PICKNIK
RIGHTHAND ROBOTICS
READY ROBOTICS
PLUS ONE ROBOTICS
OSARO
Robotics
BostonDynamics


BAND-AID
TYLENOL
RSVP
MOUTH RINSE
ESTÉE LAUDER

# Concrete Major Opportunities

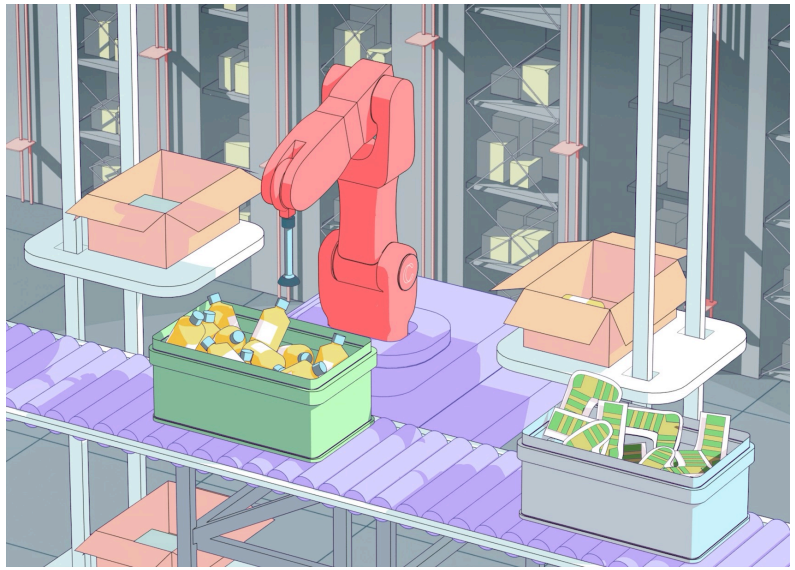

order picking
decanting

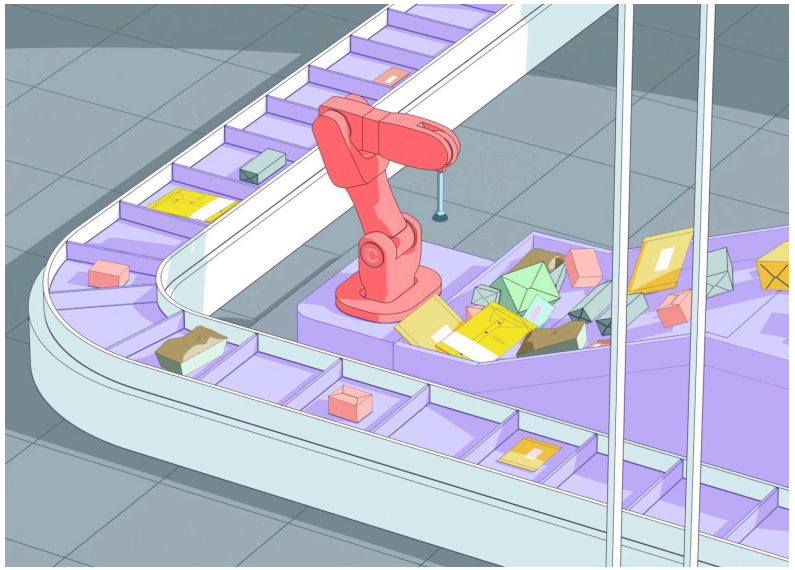

parcel induction
apparel induction

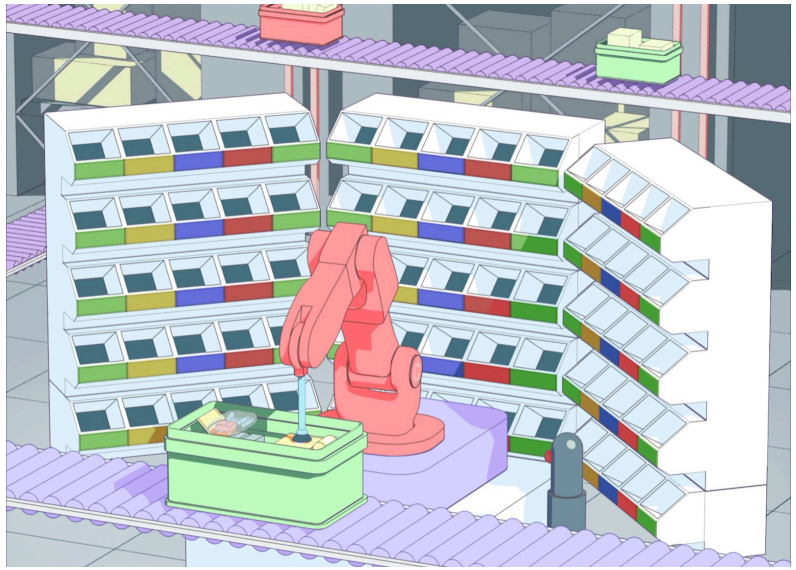

putwall
kitting


Covariant
KNAPP
KNAPP

# Next Time:

- AI Research Frontiers
- Course wrap-up
- Where to go after 188