When AI Meets AV

Challenges and Opportunities

Ching-Yao Chan
Co-Director
Berkeley DeepDrive, UC Berkeley

C2V2.AT Summer School , AI-Enabled Mobility
Sept. 14, 2020

Berkeley DeepDrive

Presentation Outline

- AI for AV – How and Where?
- AV/AI Research at UC Berkeley
  - Berkeley DeepDrive (BDD)
  - Selective Research Highlights
- AV+AI – Challenges and Opportunities
- Concluding Remarks

When AI meets AV

A Powerful Enabler

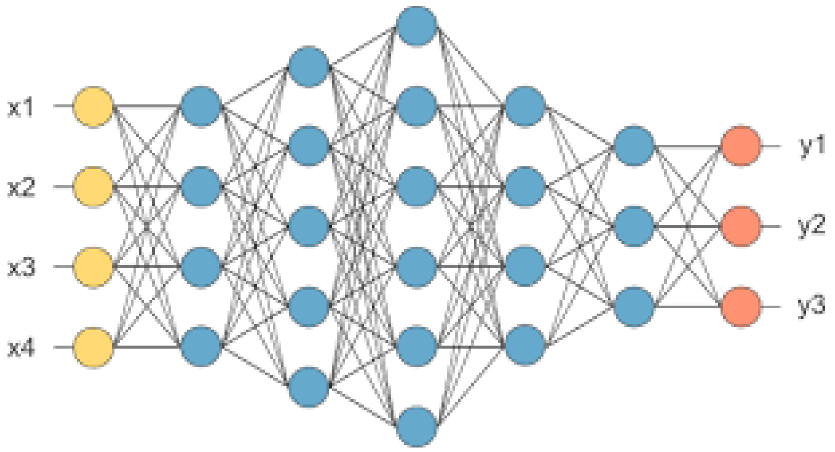

A Fitting Challenge

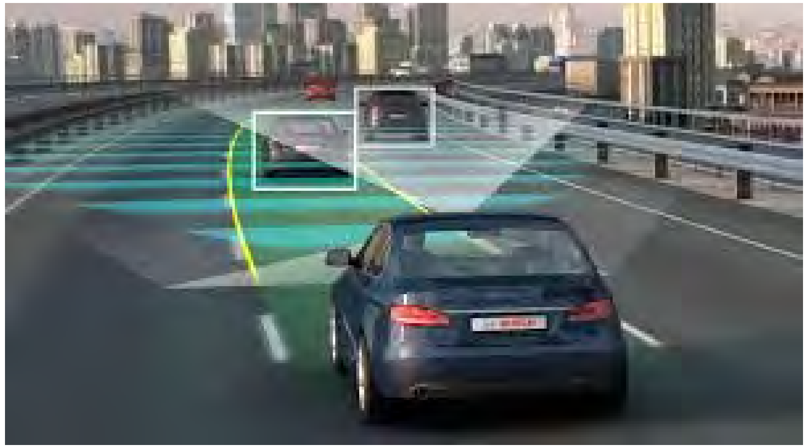

AI for AV, Where and How?

Berkeley DeepDrive

Automated Driving Systems (ADS)
- Functional Block Diagram

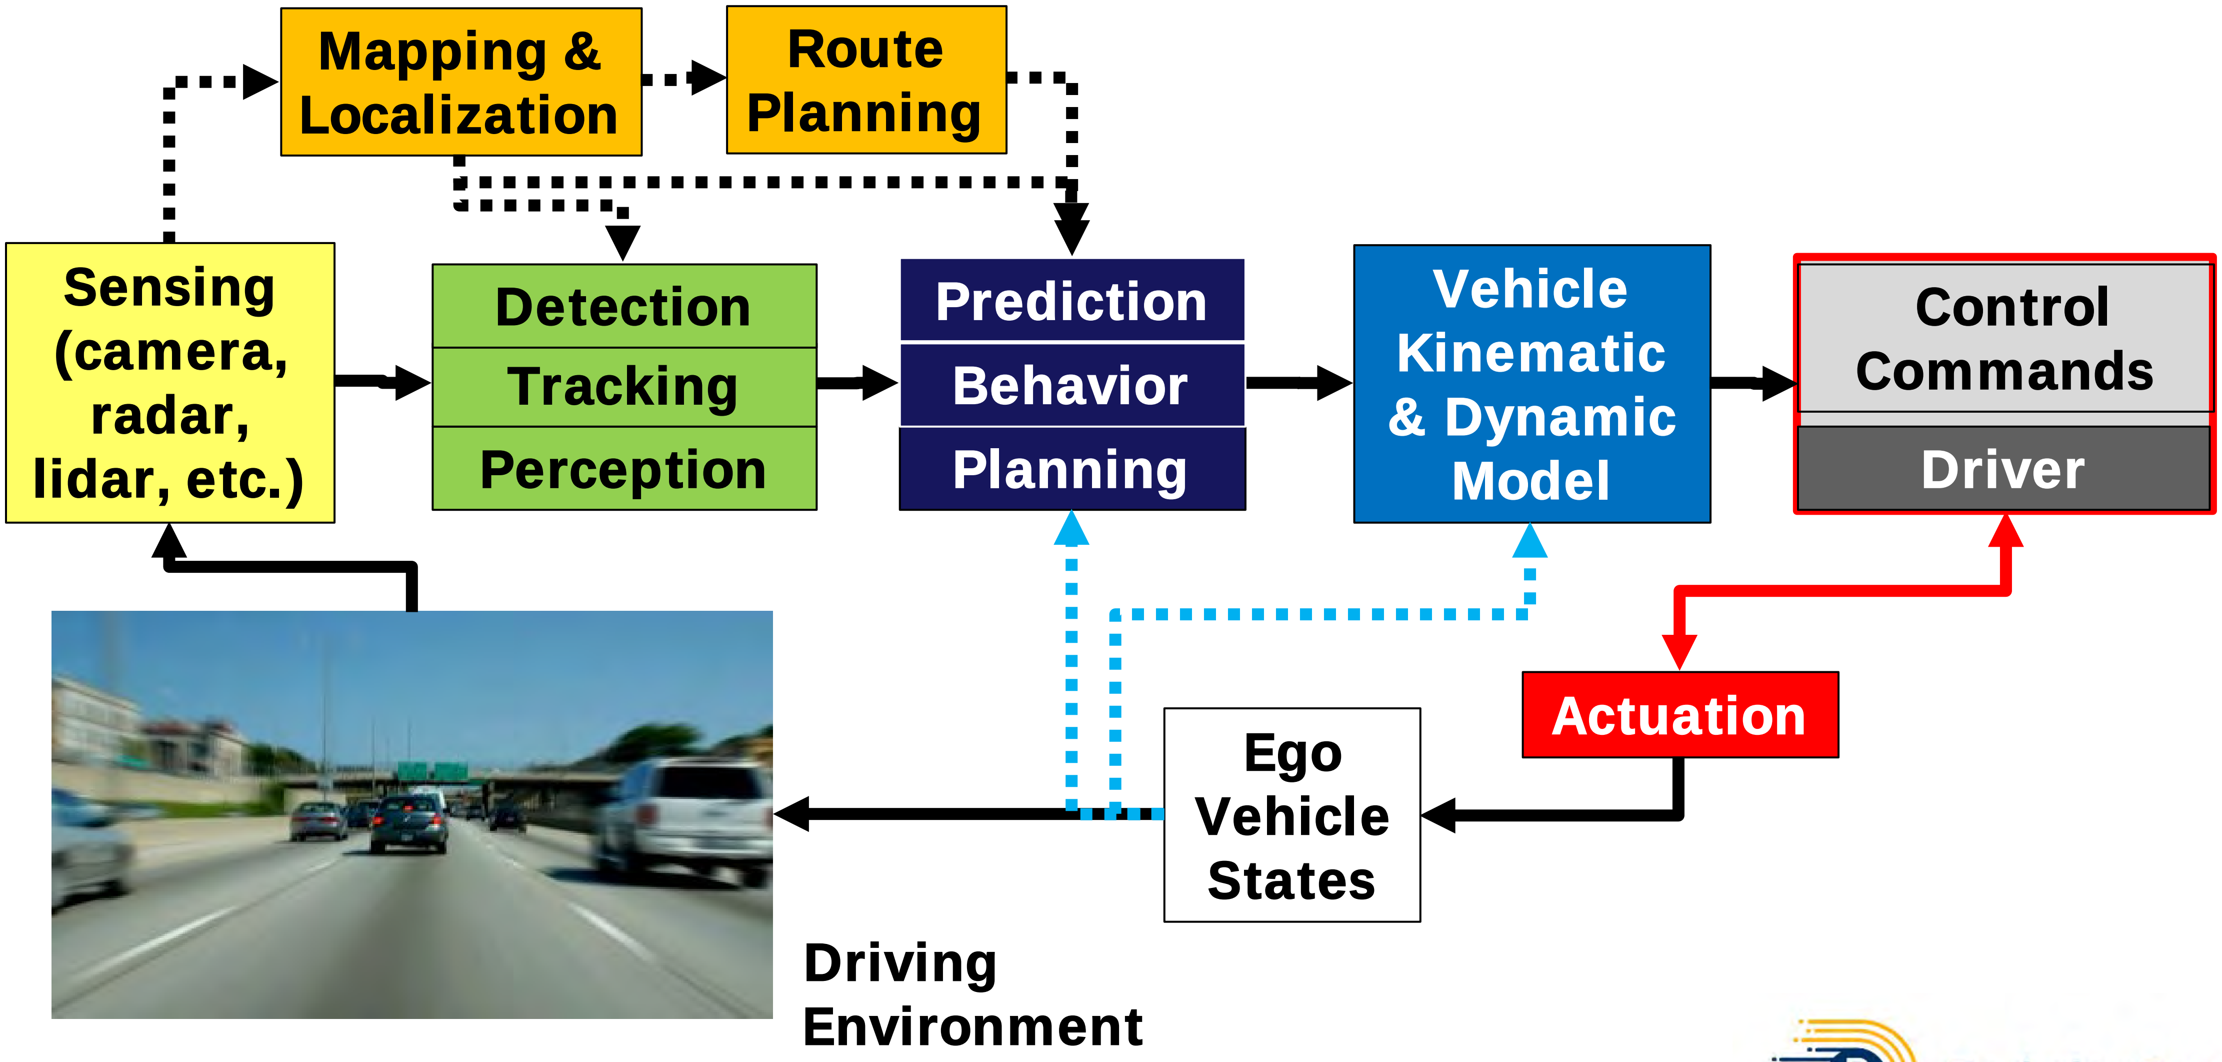

Berkeley DeepDrive

Automated Driving Systems (ADS)
System Integration and Functional Blocks

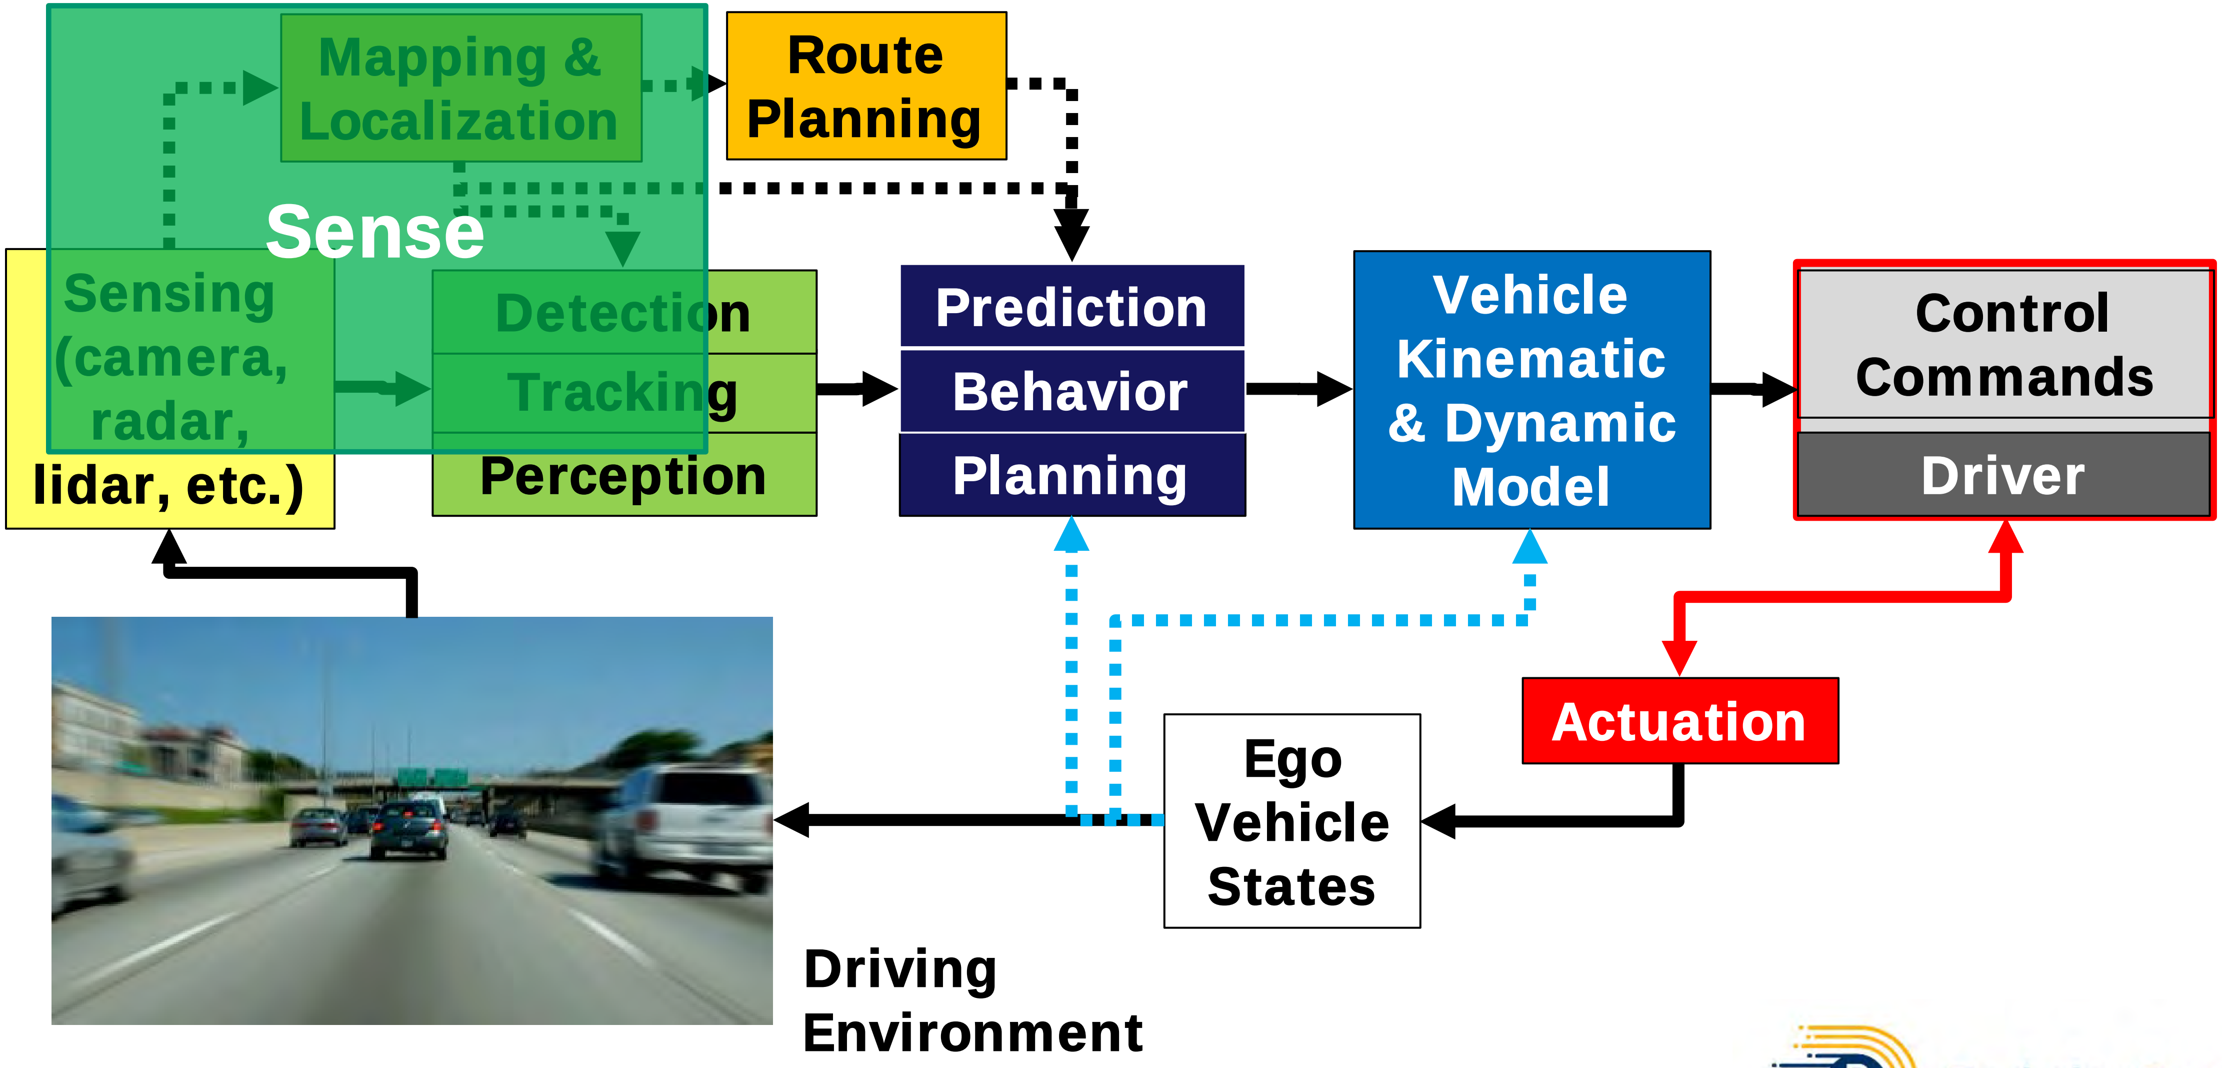

Berkeley DeepDrive

Perception by Deep Learning

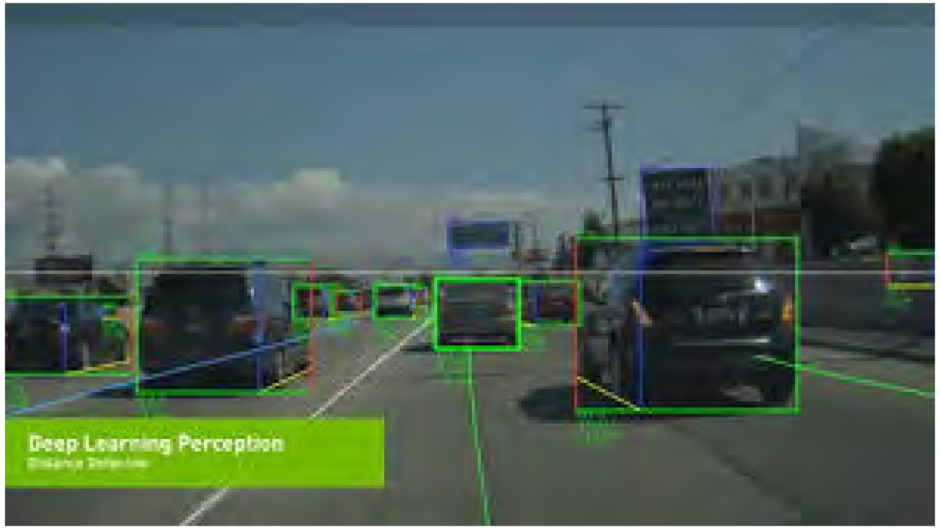

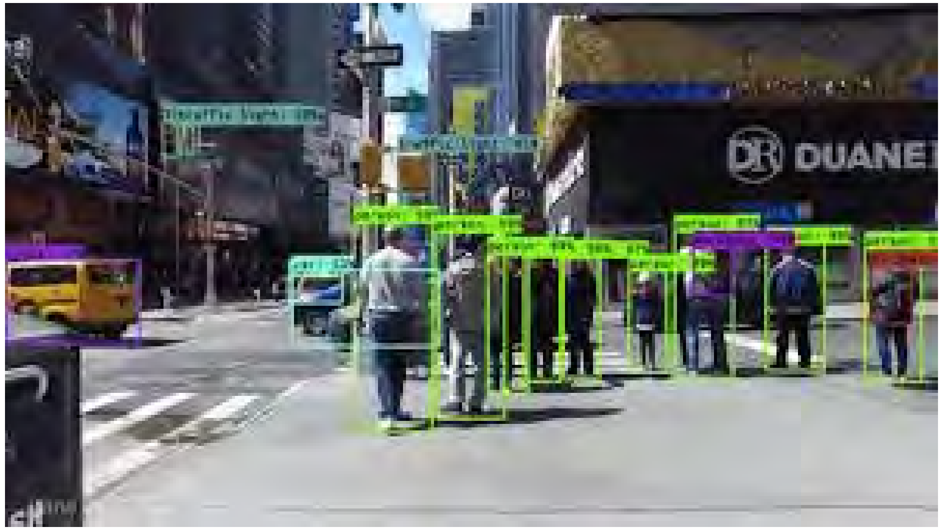

Berkeley DeepDrive

Automated Driving Systems (ADS)
System Integration and Functional Blocks

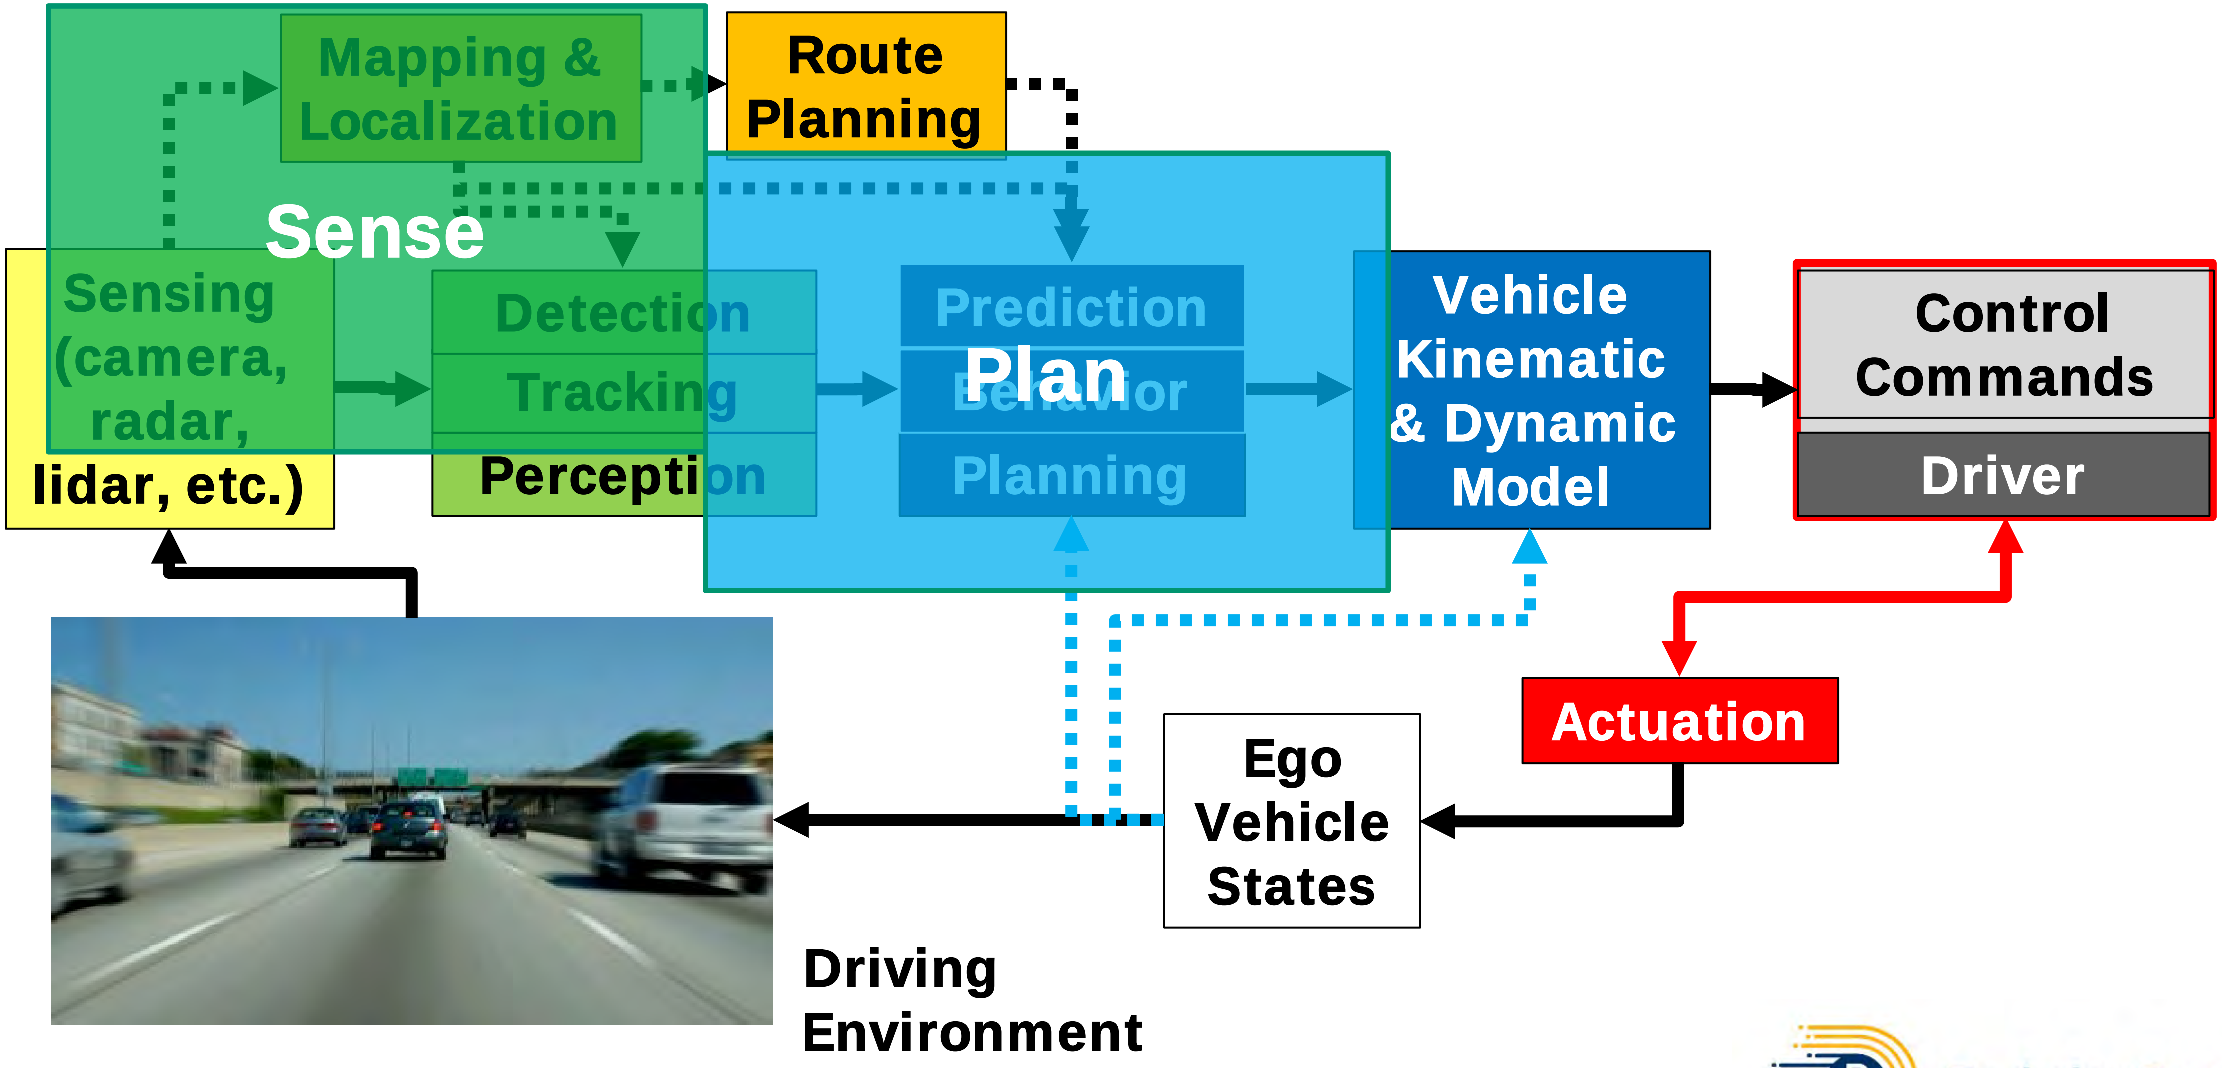

Berkeley DeepDrive

Decision Making and Planning

Machine learning approaches can enhance rule-based methods in

- **Prediction**
- **Behavior**
- **Trajectory planning**

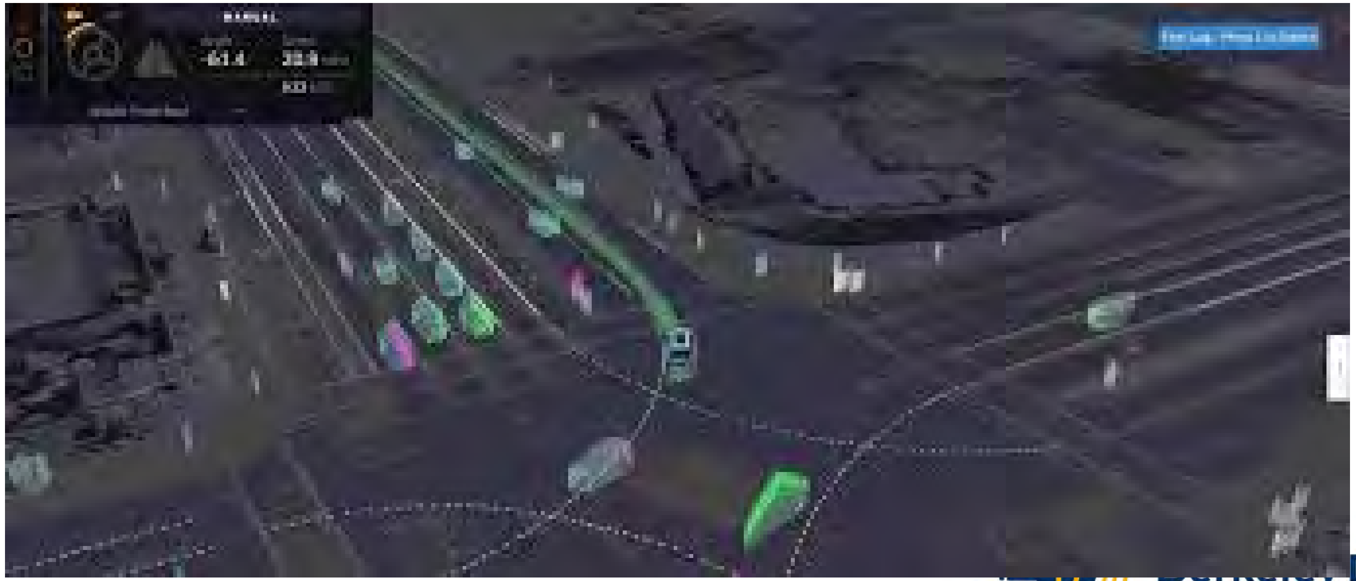

Berkeley DeepDrive

Automated Driving Systems (ADS)
System Integration and Functional Blocks

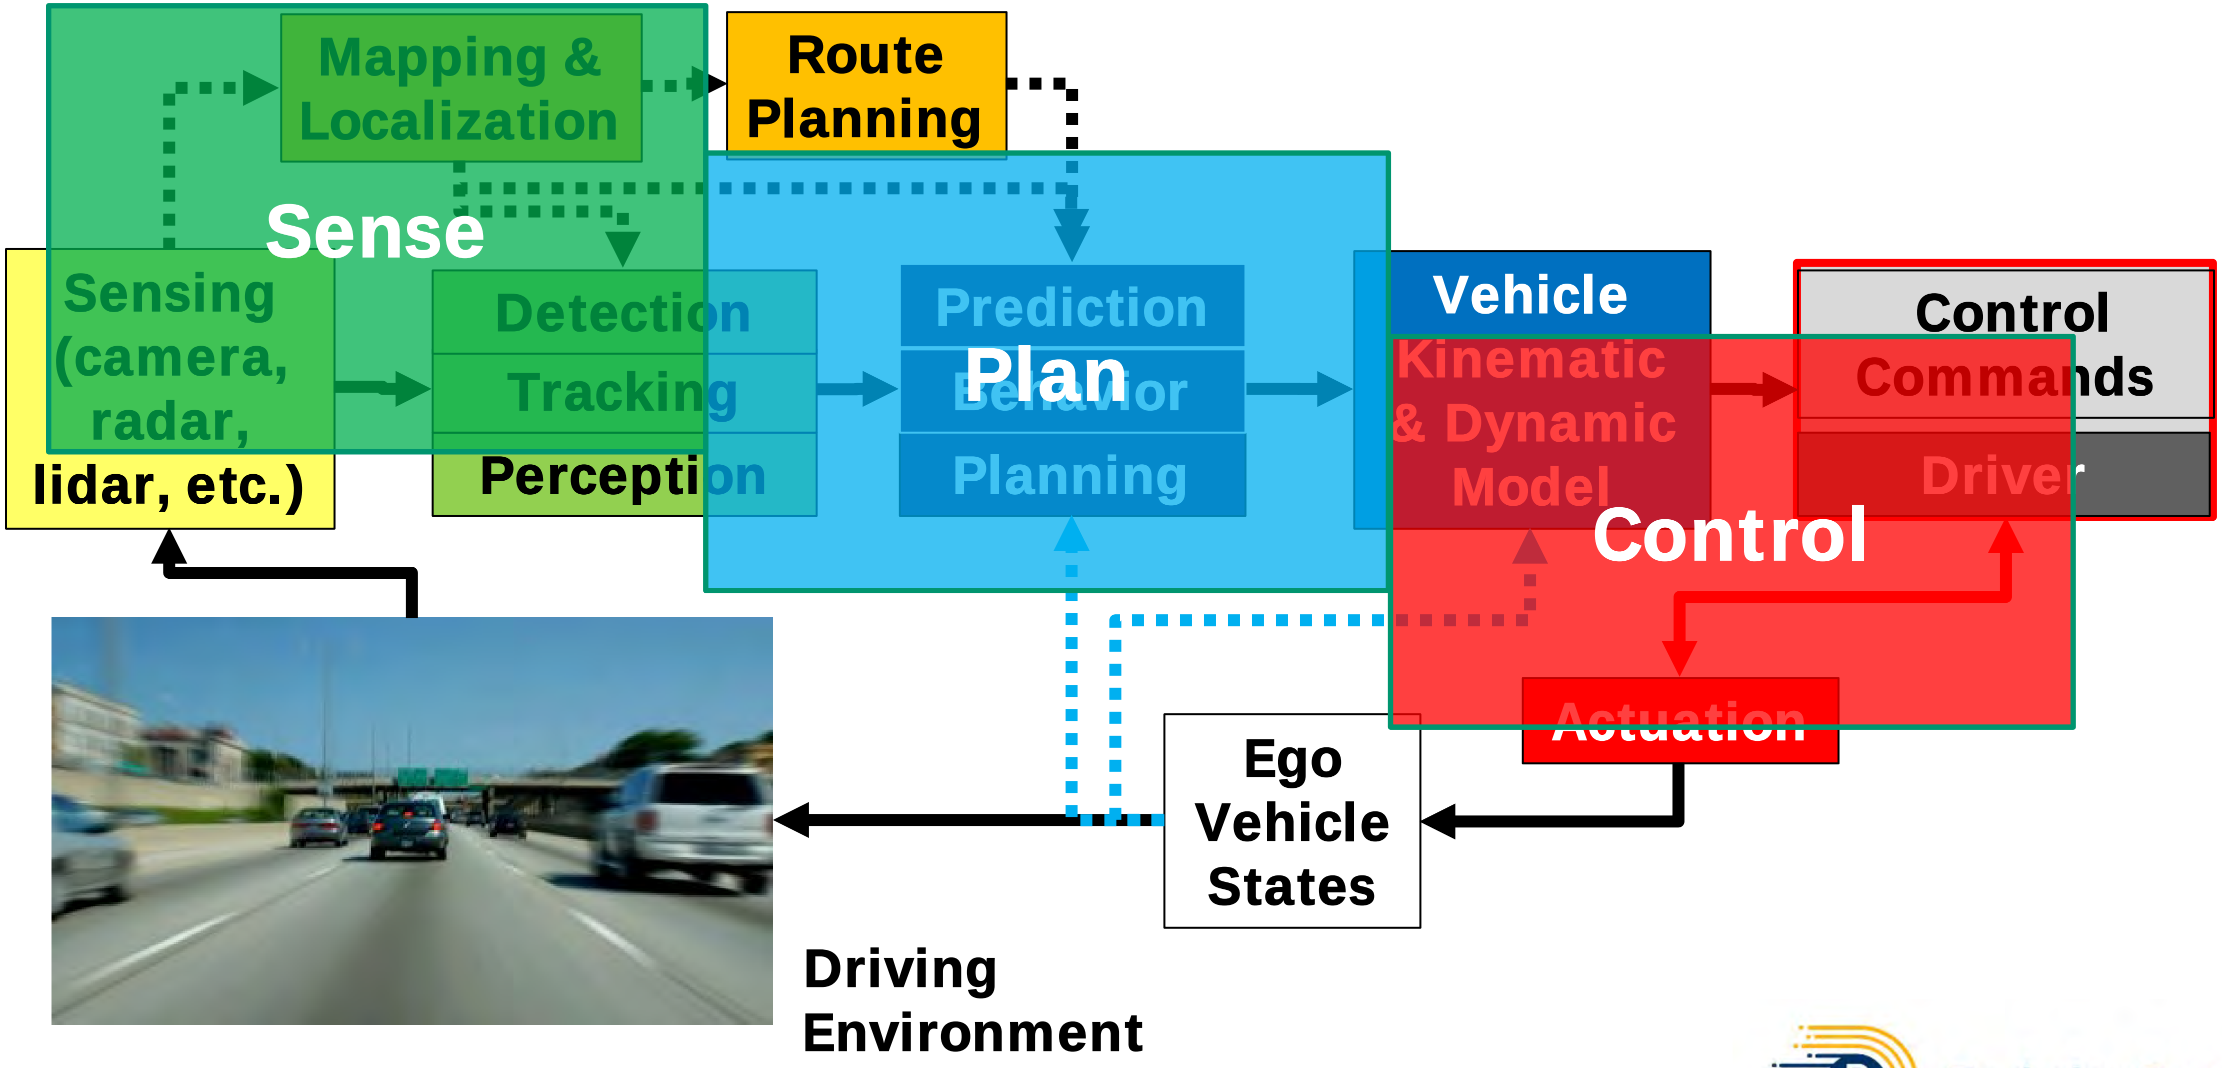

Berkeley DeepDrive

Control and Actuation

Deeping Learning for Intelligent Powertrain

I-Ming Chen, BDD, UC Berkeley

Behavior of the DQN controller for HEV powertrain:

- Reasonable gear shifting maneuver
- Hybrid mode at low vehicle speed
- ICE mode at high vehicle speed
- Braking regeneration

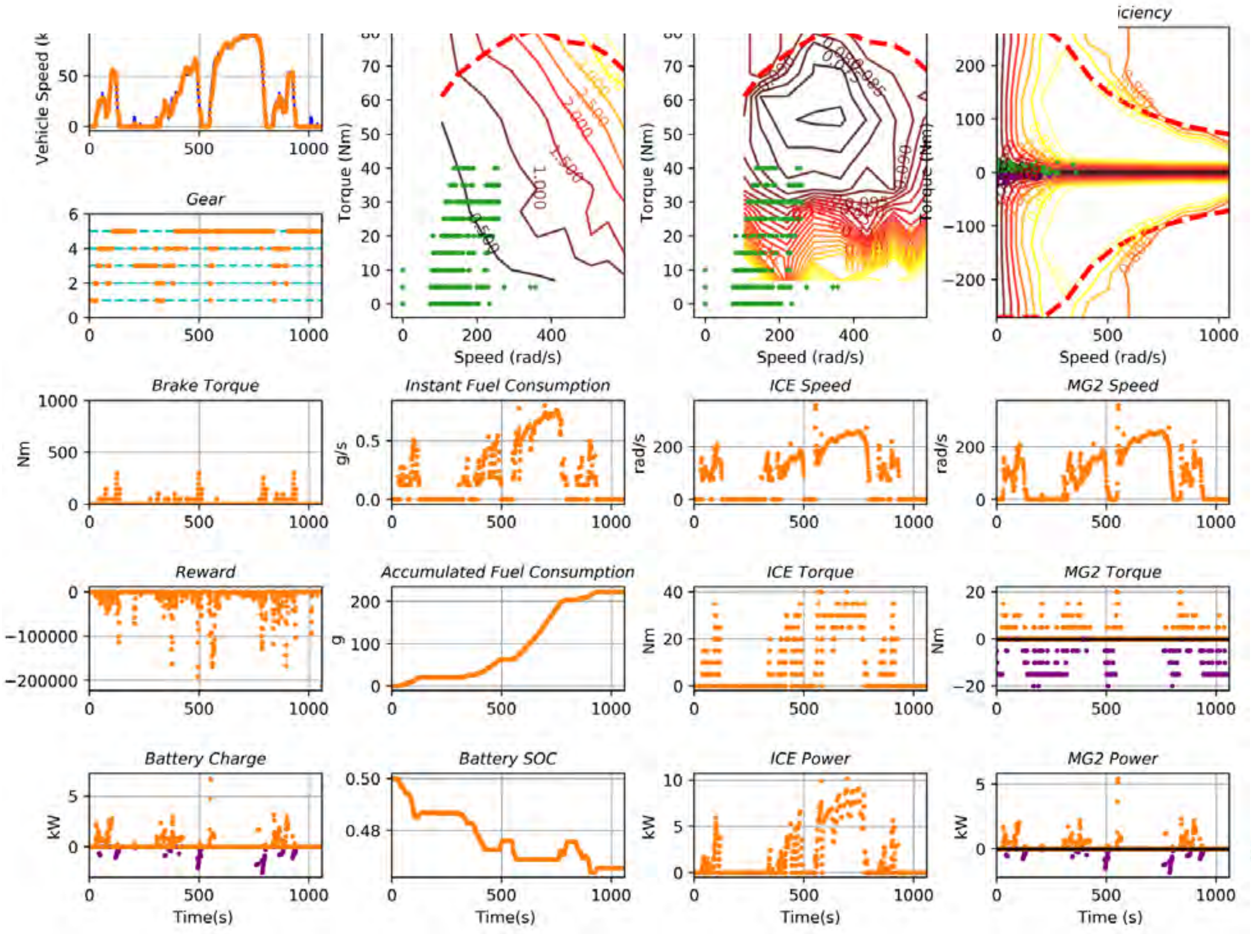

Drive Toward Intelligent Autonomy
Automated Driving Systems

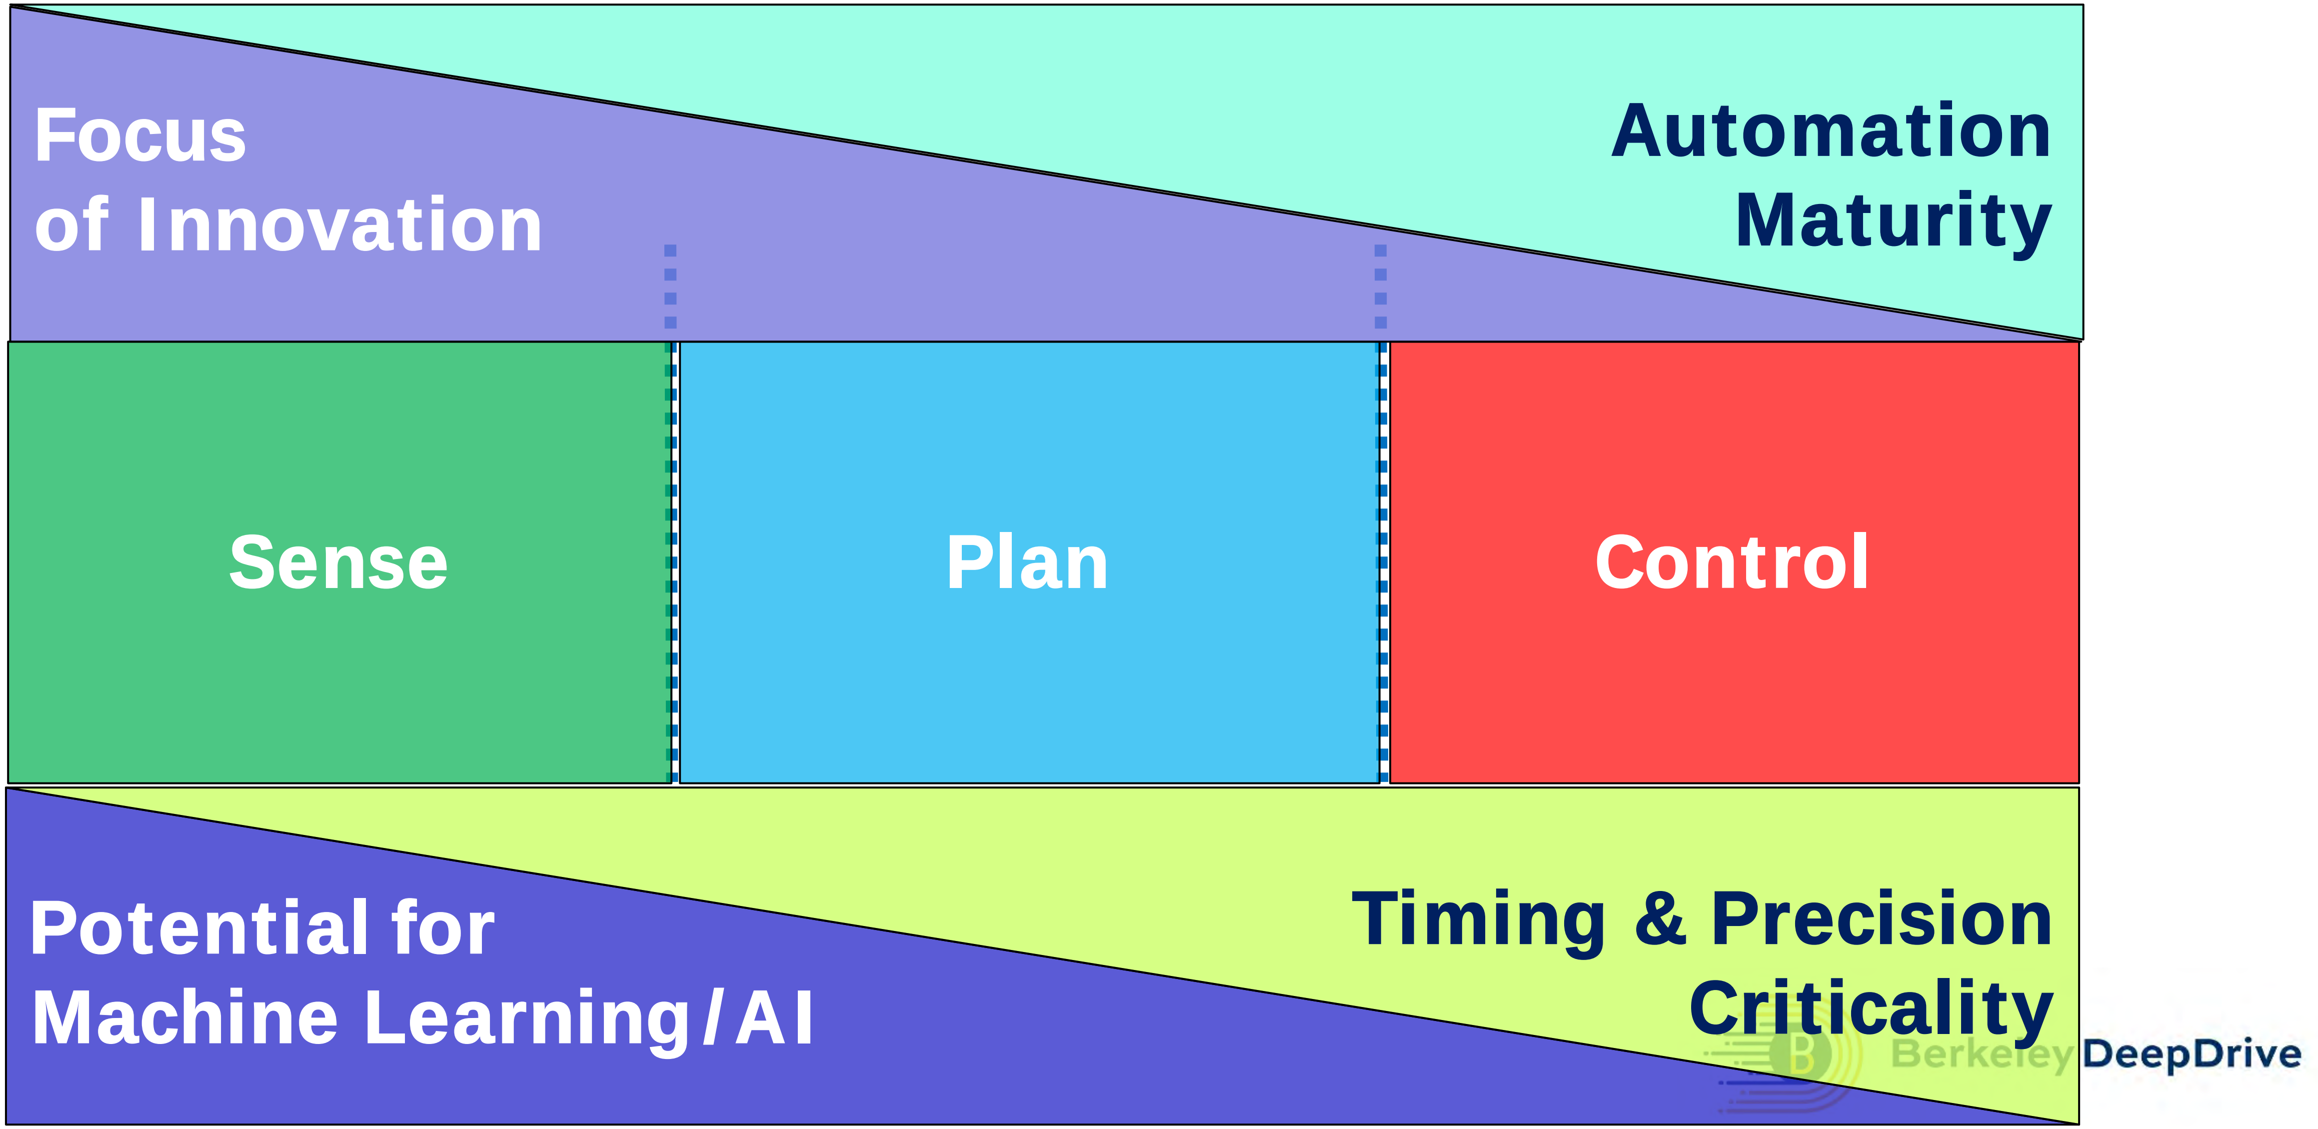

AI / Autonomous Driving Research at Berkeley

Berkeley DeepDrive

BAIR-BDD Community at Berkeley

- **Three Parallel Programs**
  - *Berkeley Artificial Intelligence Research (BAIR)*
  - *Berkeley DeepDrive (BDD)*
  - *BAIR Open Research Commons (BAIR ORC)*

Microsoft

SAMSUNG

Google

amazon

Alibaba.com

Qisda

DiDi

Panasonic

SONY

GM

美团

Meituan-Dianping

TOYOTA

Ford

Berkeley DeepDrive

Berkeley DeepDrive

Our Mission

Merge Machine Learning and Computer Vision to

Enable Intelligent Autonomy, with Applications for

Autonomous Driving and Robotics

BDD Research Domains

See deepdrive.berkeley.edu for Research Activities

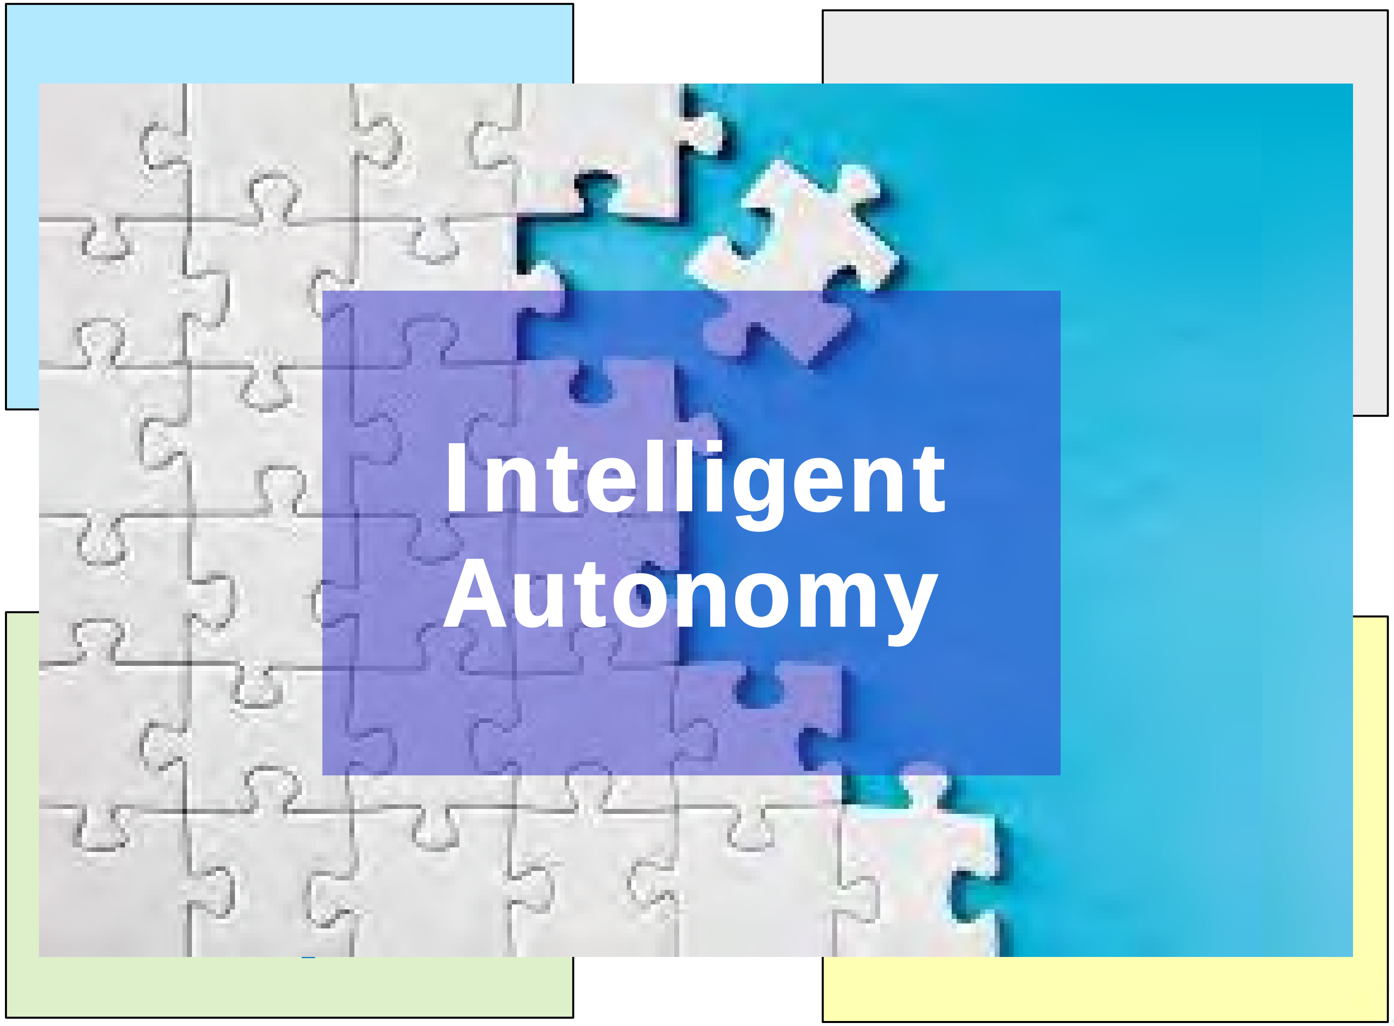

Berkeley DeepDrive

Selective Research Highlights

Berkeley DeepDrive

Selective Research Highlights

- *Data*
  - *From image to video*
  - *From object recognition to tracking*
  - *What if you have too much data?*
  - *...*

Sensing and Perception

Deep learning enables a renaissance in computer vision (images: BDD-100k)

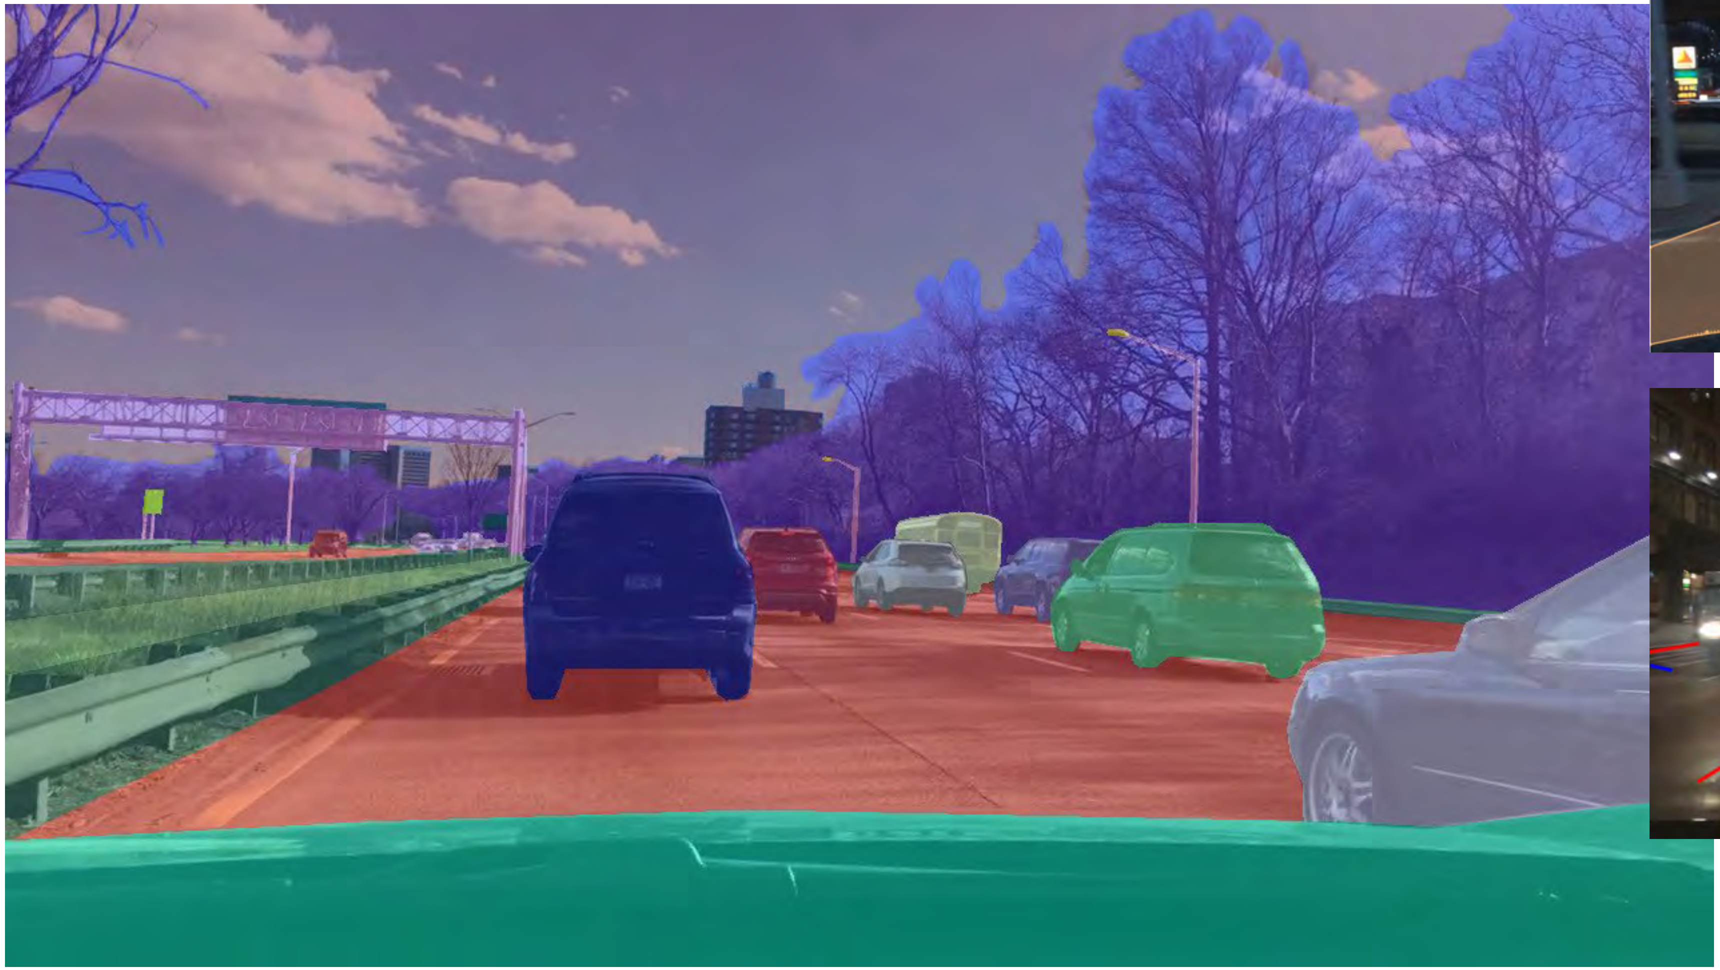

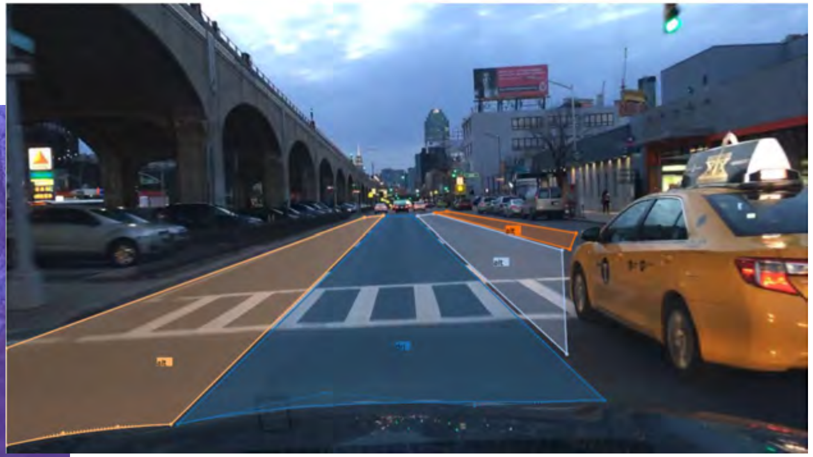

Berkeley DeepDrive

Has the machine seen this data before?

Scene Novelty Prediction from Unsupervised Discriminative Feature Learning,

Arian Ranjbar, Daniel Ye, Sascha Hornauer, Ching-Yao Chan, Stella Yu

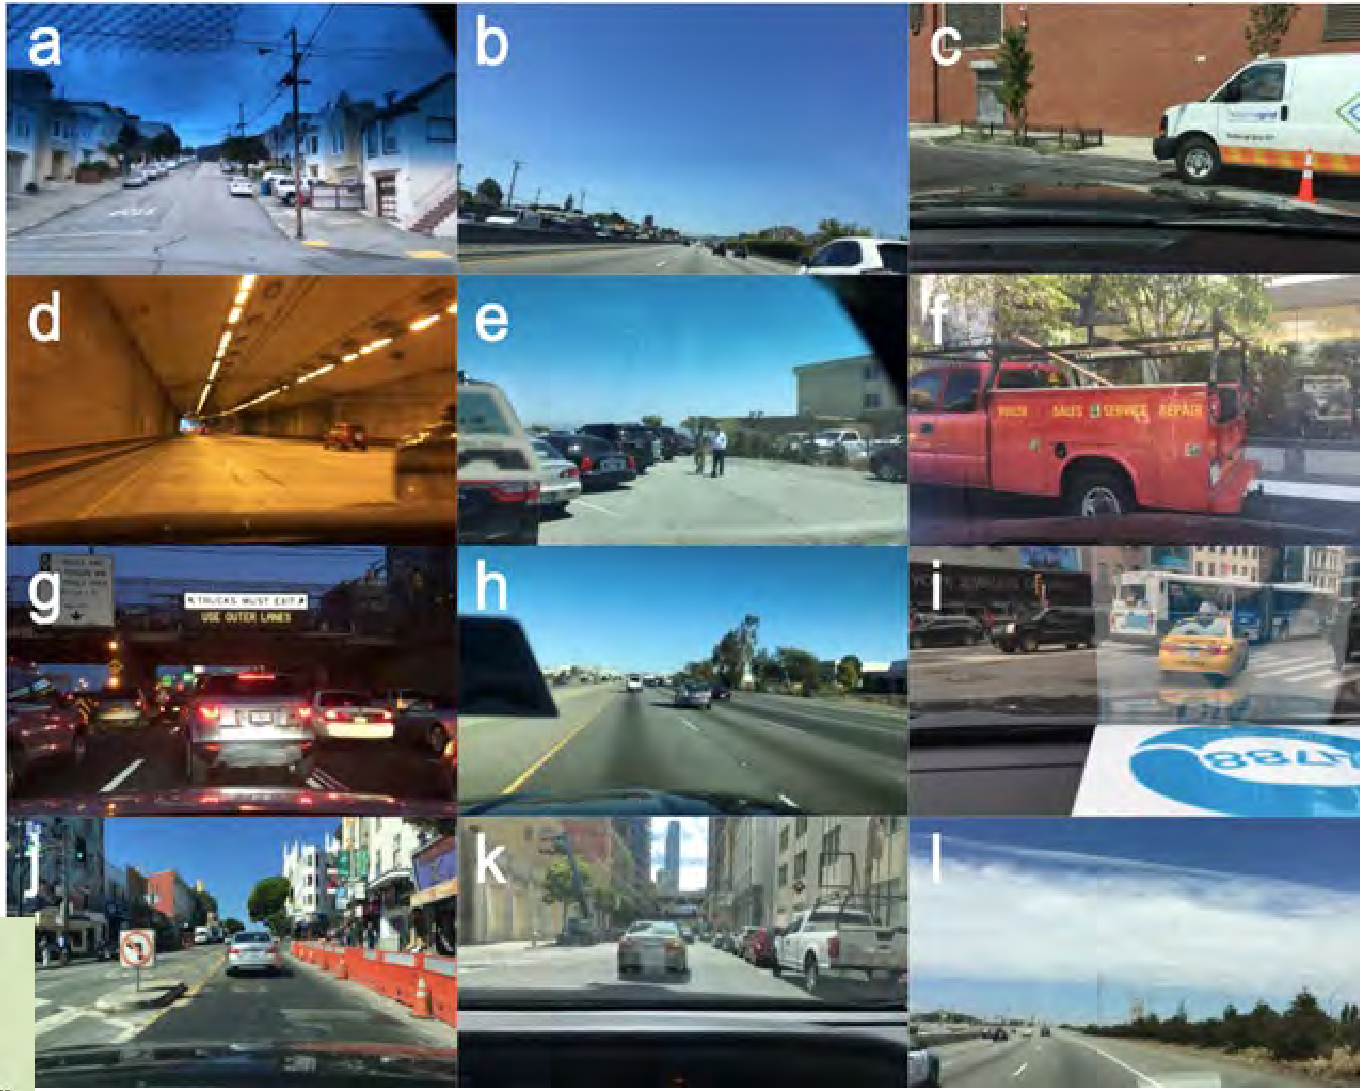

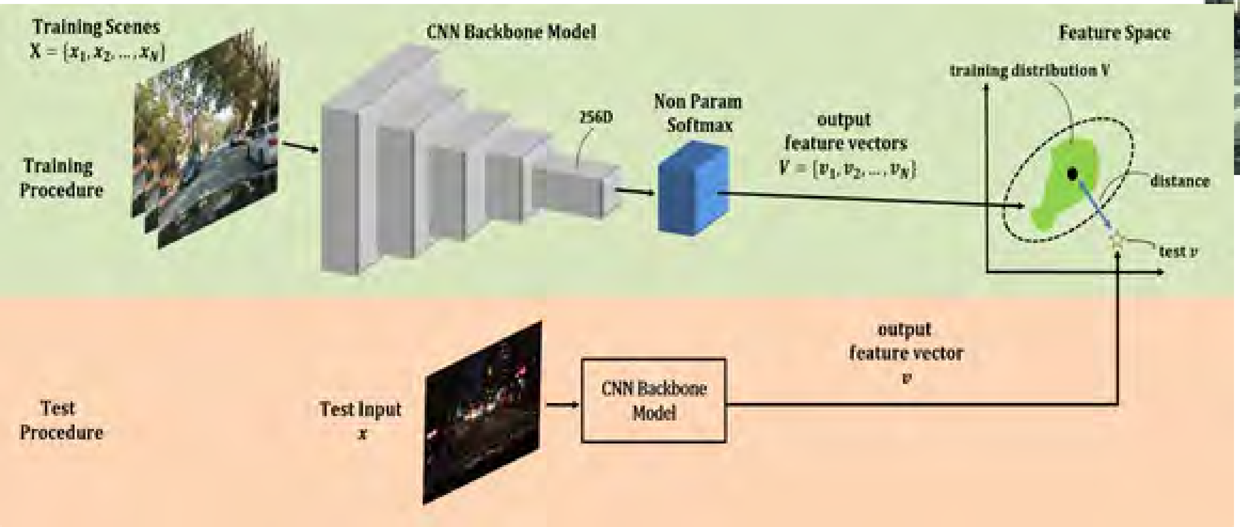

Source: Sarcha Hornauer
BDD Workshop: 09/2019

Berkeley DeepDrive

Retrievals of Similar Scenes

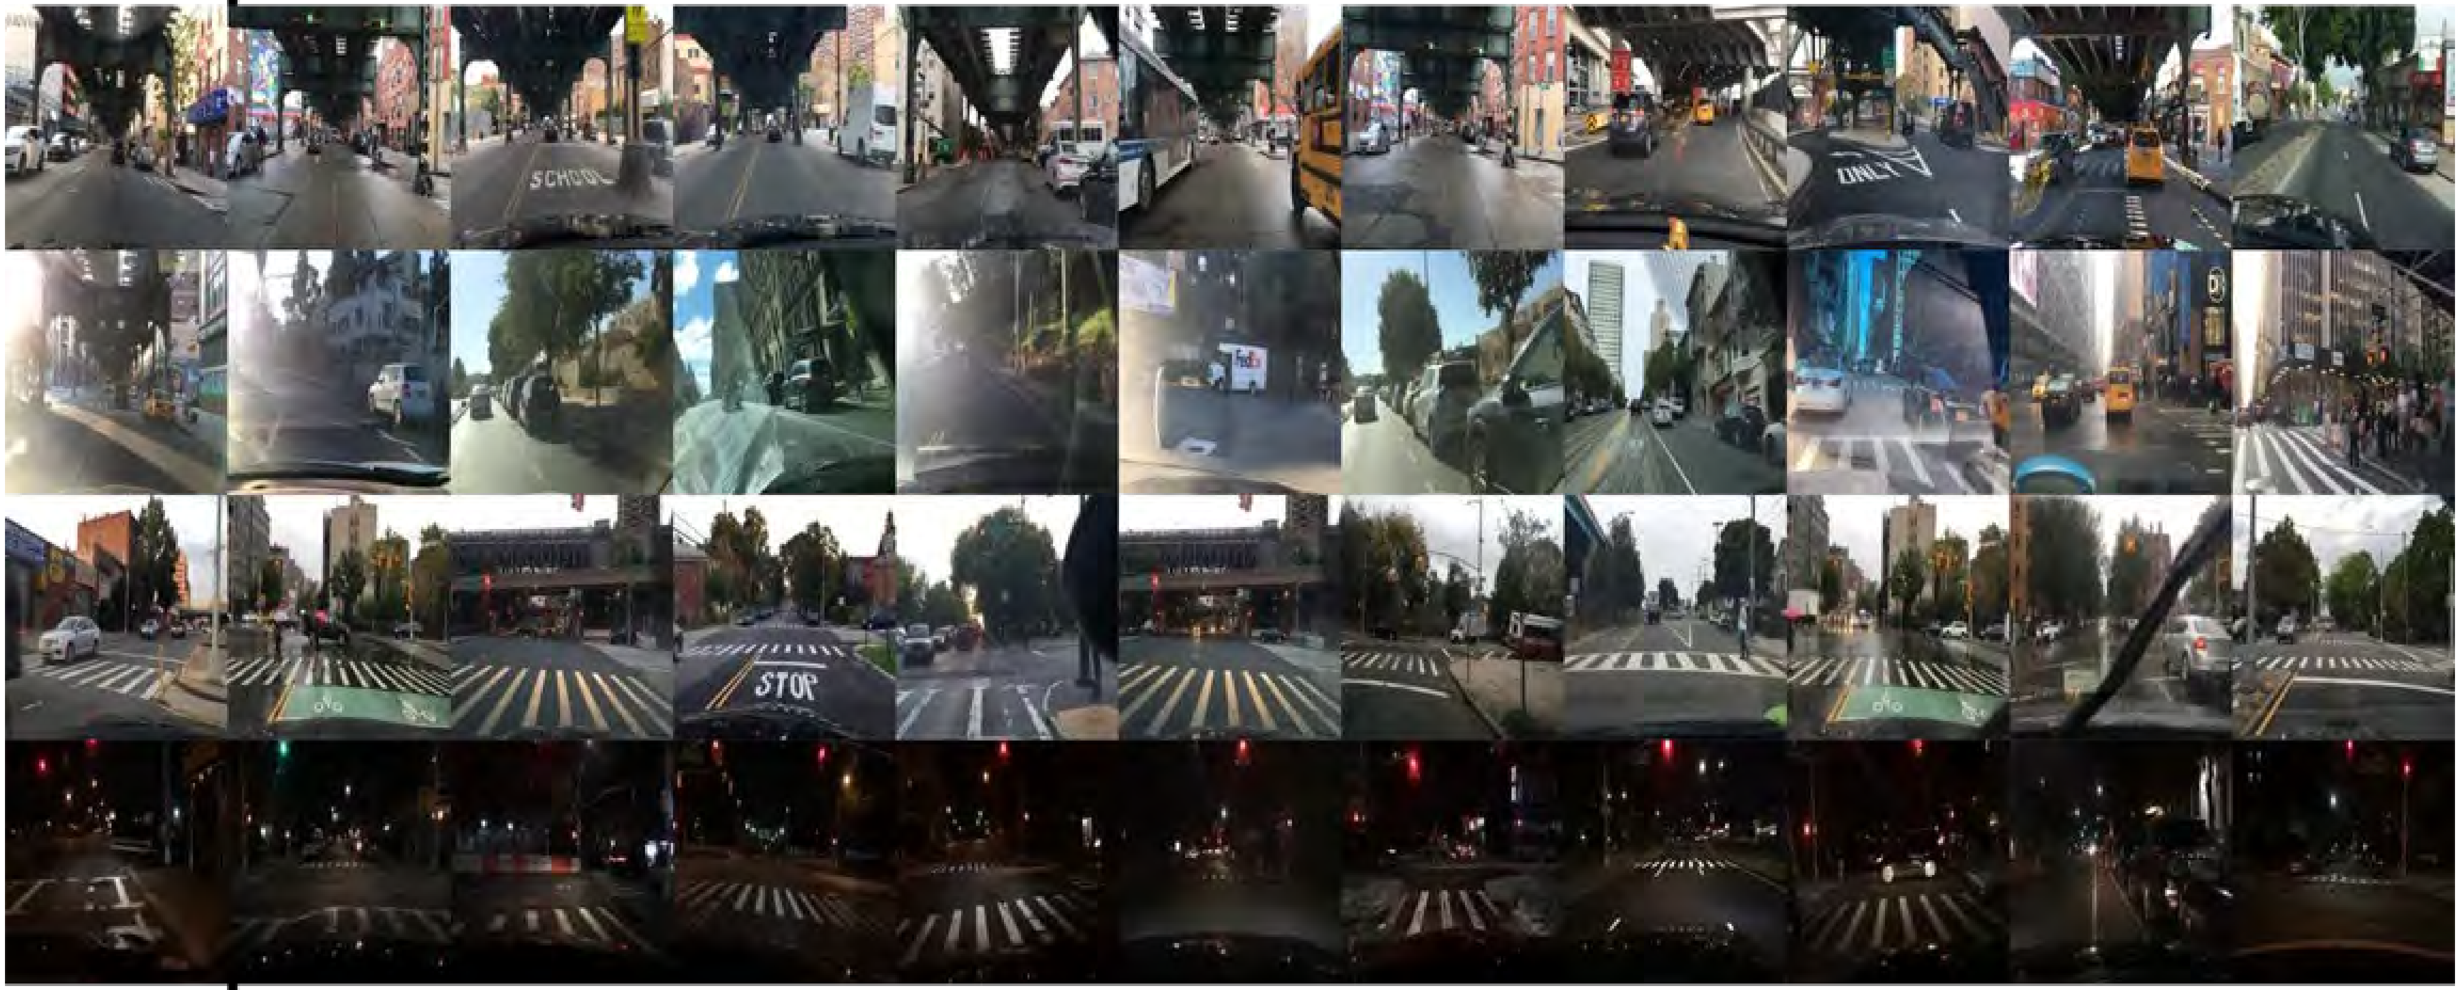

Source: Sarcha Hornauer
BDD Workshop: 09/2019

Berkeley DeepDrive

Selective Research Highlights

- *Data*
  - *From image to video*
  - *From object recognition to tracking*
  - *Confidence on data based on scene similarity*
  - *…*
- *Learning to drive*
  - *Learning reward and policy*
  - *Meta learning – What if you don't have enough data?*
  - *How to apply Deep Reinforcement Learning for Tracking Control?*

Decision-making Strategy based on Adversarial-IRL

Meta-Adversarial Inverse Reinforcement Learning:

- Learn decision on when and where to change lanes (AIRL)
- Learn multiple behaviors (Info-AIRL)
- Learn an *adaptable policy and a reward function* from demonstrations for quick and continuous learning *(Meta-AIRL)*

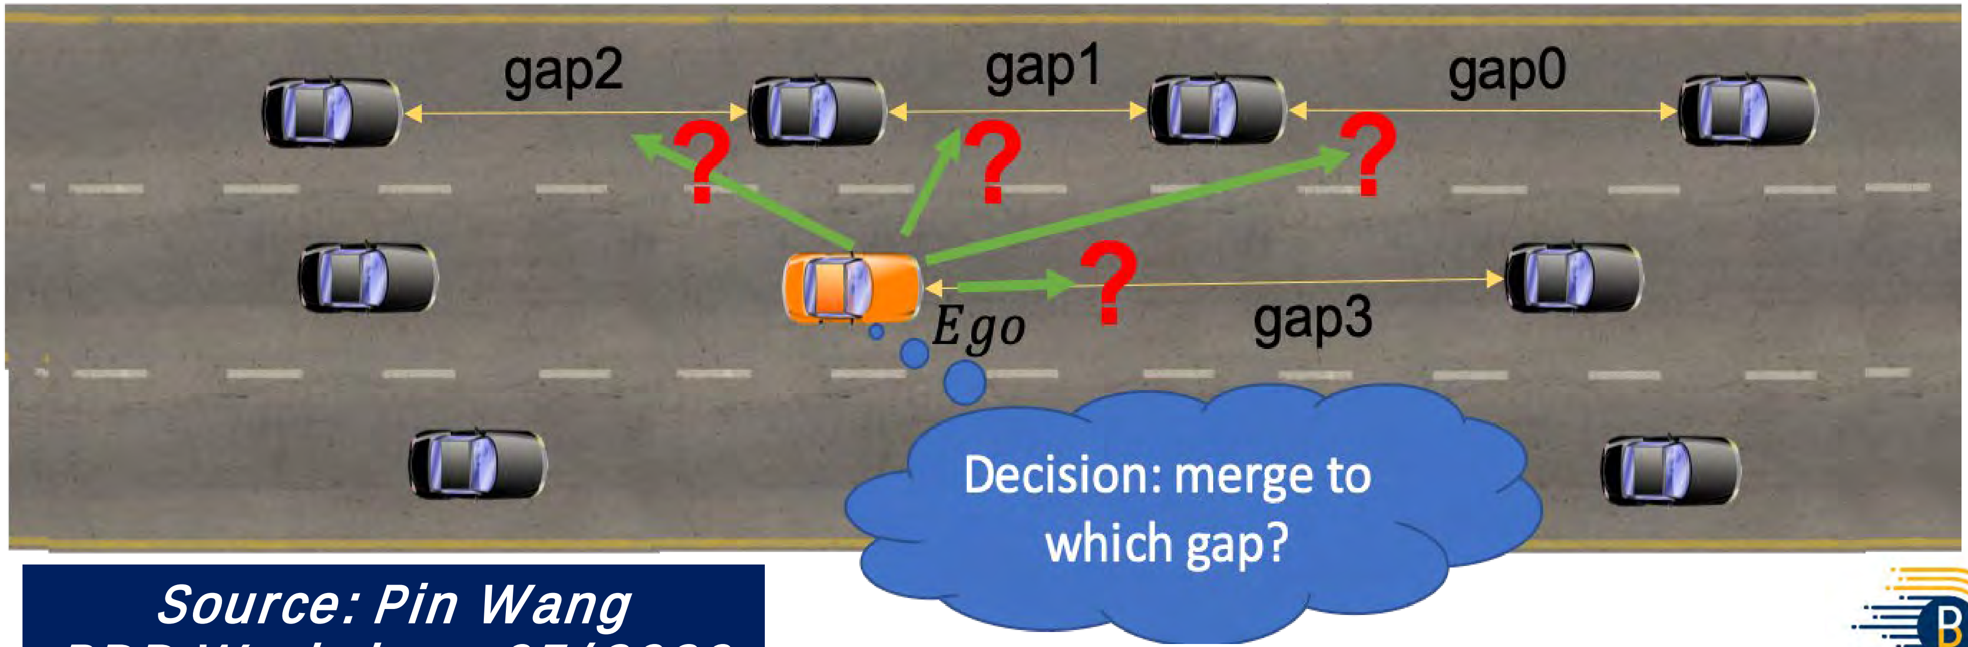

Source: Pin Wang
BDD Workshop: 05/2020

Berkeley DeepDrive

Learning to Drive with Different Styles

...

aggressive **conservative** **neutral**

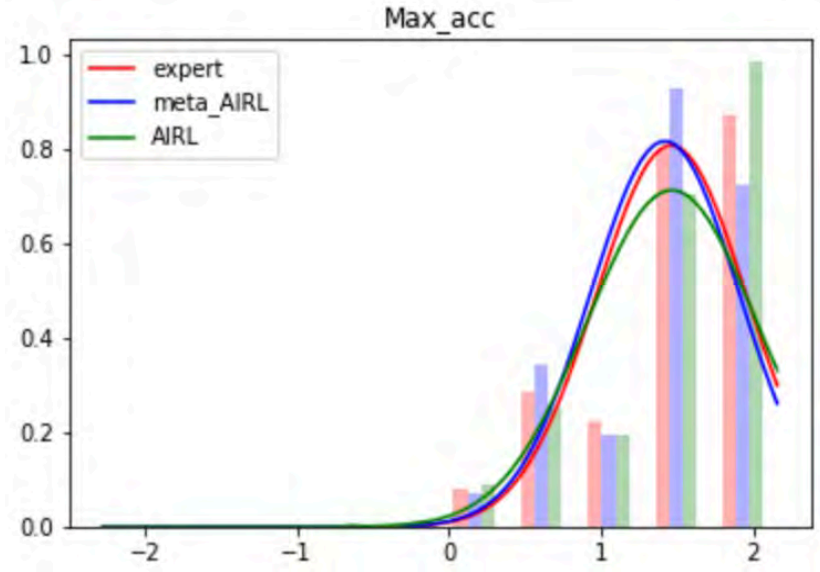

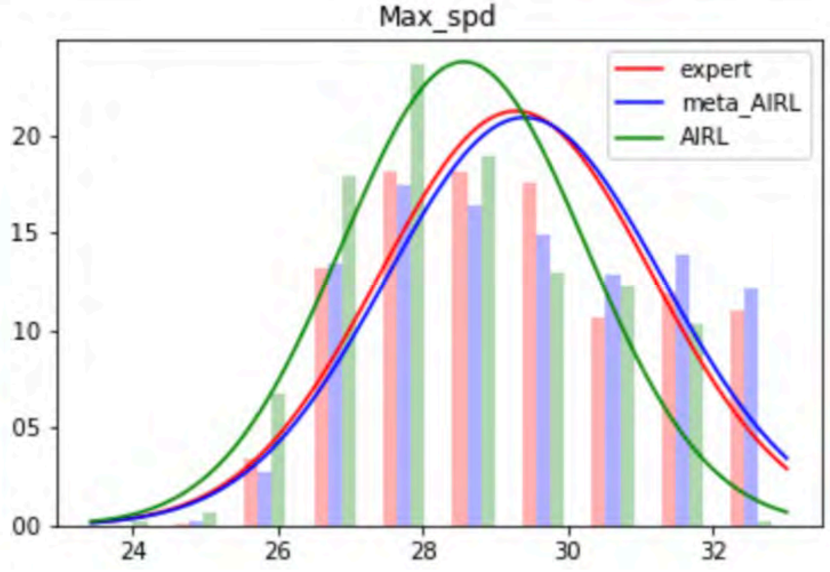

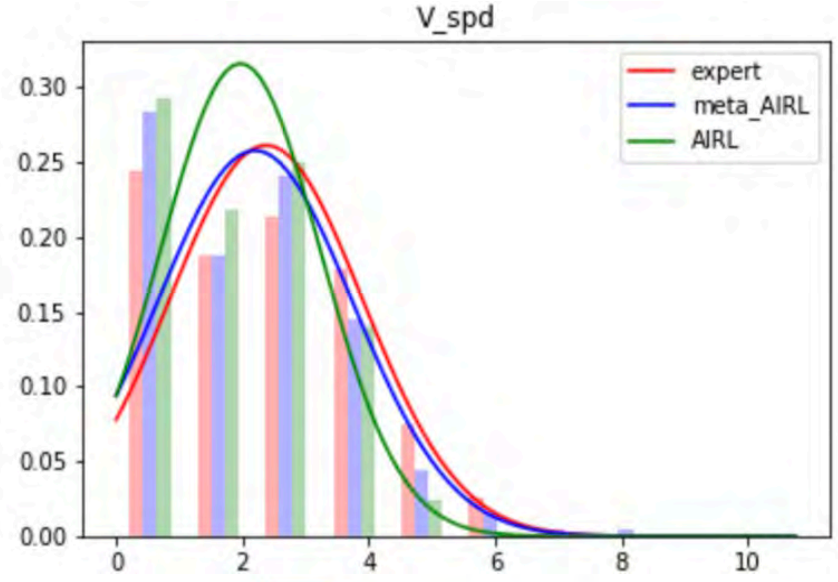

Source: Pin Wang
BDD Workshop: 05/2020

Berkeley DeepDrive

Meta Learning for Automated Lane Change Maneuvers

- **Meta-Learner (Outer-level)**
  - **Meta-learner enables efficient adaption of the lane change to new situations**
- **Task-Learner (Inner-level)**
  - **Task learner learns a decision-making strategy for automated lane change**

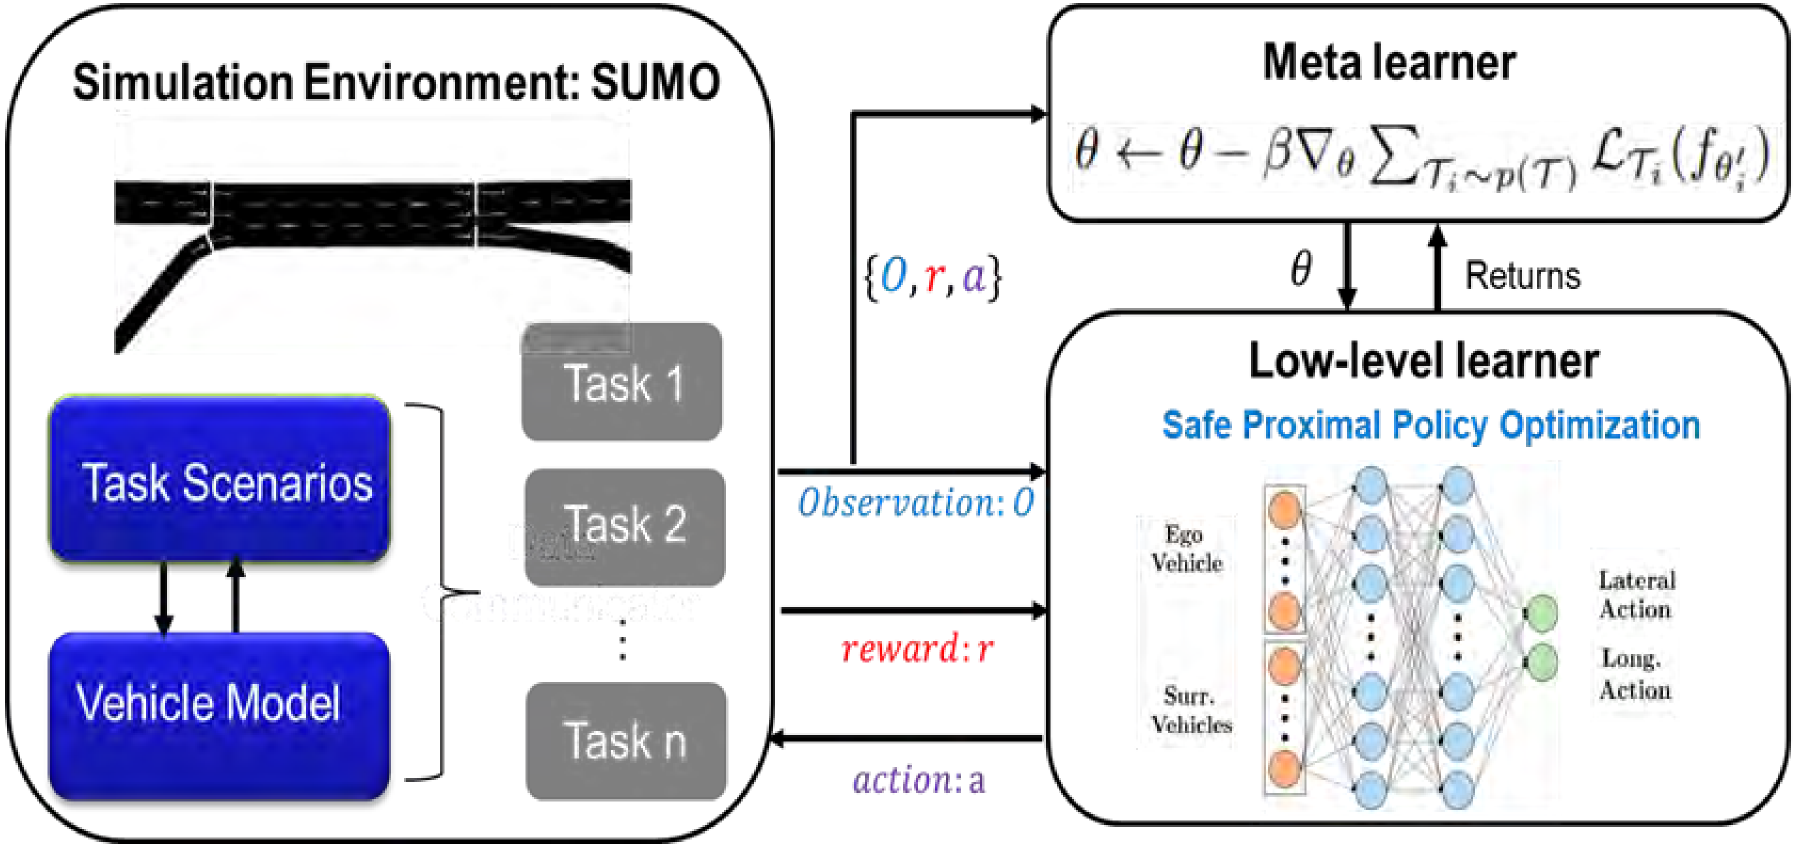

Source: Fei Ye
Project Update: 08/2020

Berkeley DeepDrive

Simulation at Different Traffic Conditions

- **Highway exit behavior under light, moderate, dense traffic conditions**
  - **Ego vehicle: red color**
  - **Other vehicles: blue color**

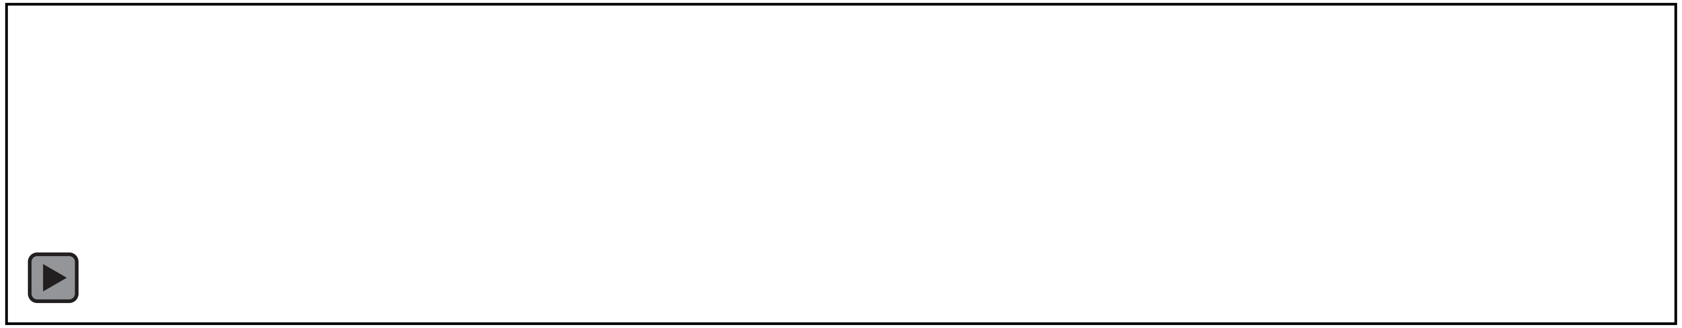

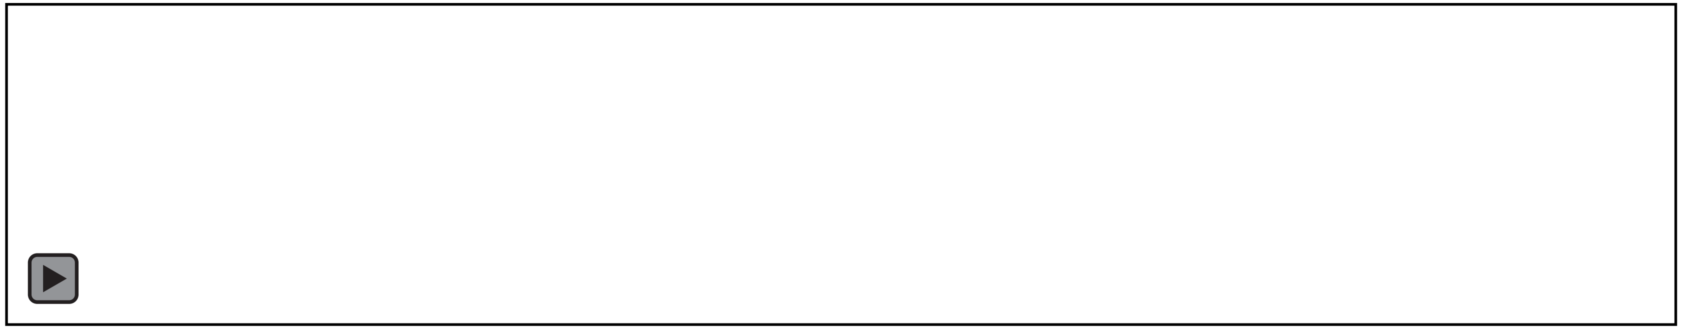

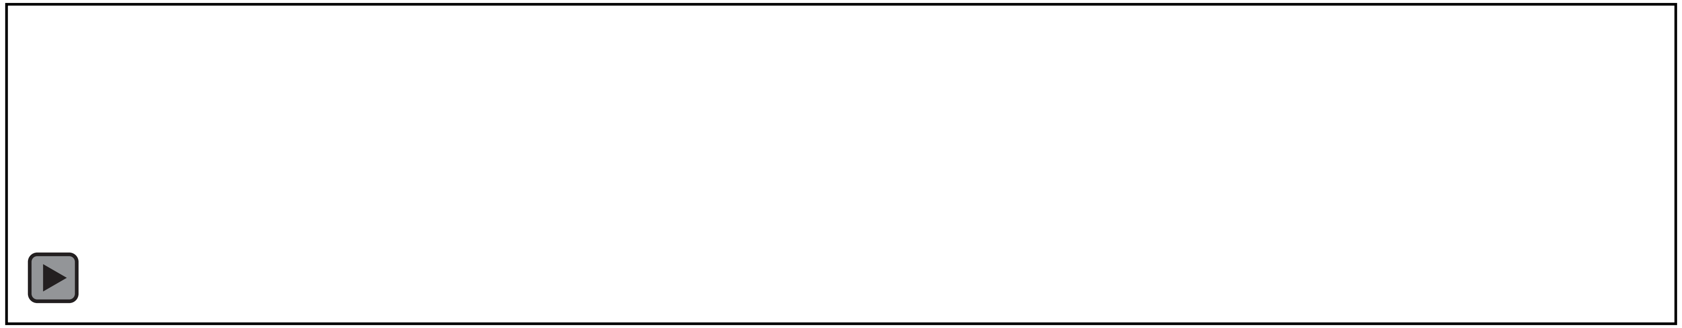

Source: Fei Ye
Project Update: 08/2020

Berkeley DeepDrive

Control Architecture for Dual Use of A Traditional Controller and A RL Controller in Parallel

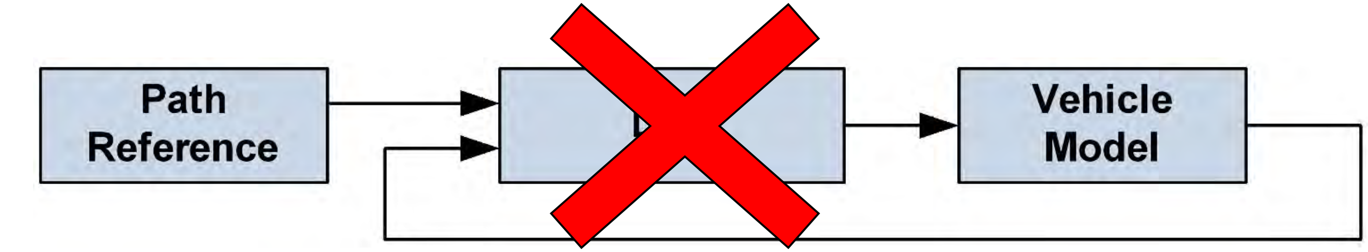

Problems:
- During training, one failing step can cause erratic tracking actions.
- When the reference path changes, DRL controller becomes unstable.

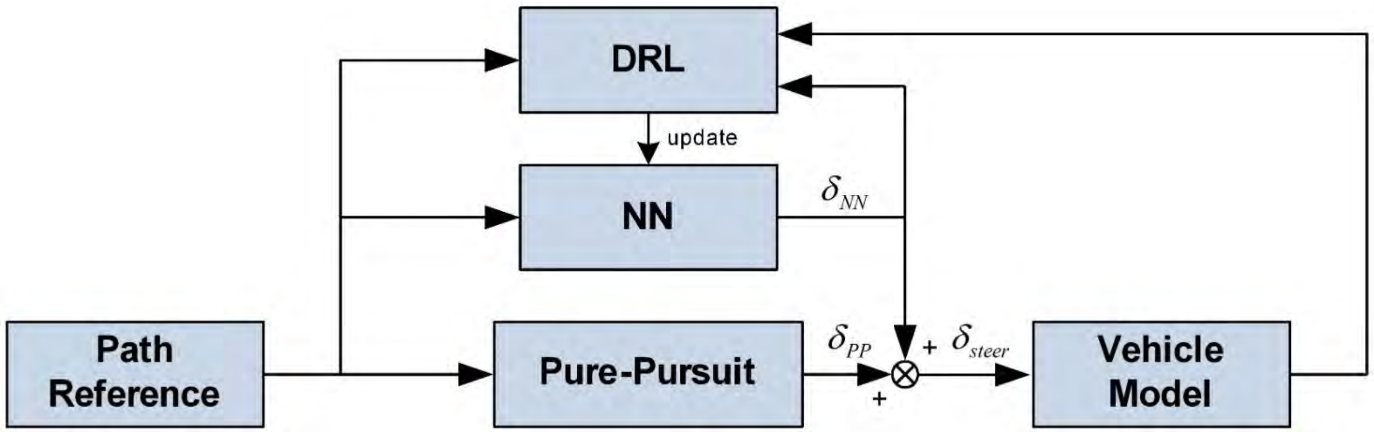

Revised Architecture:
- PP provides a basic path following capability guiding DRL.
- DRL provides adaptability and optimality.

- **RL output has the potential to supplement and/or enhance the traditional controller**

Source: I-Ming Chen Project Update: 04/2020

PP+DRL Controller versus MPC

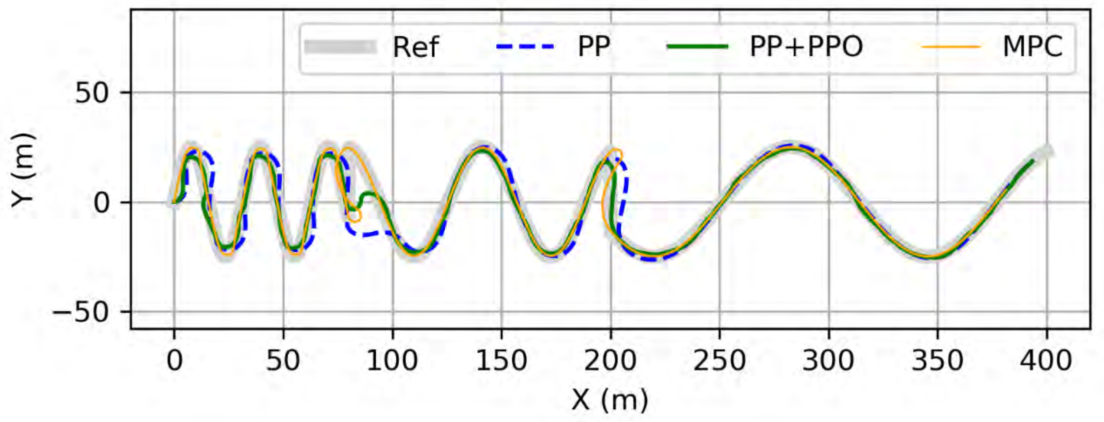

- The actual execution time of the simulation is 5.6s for PPO+PP and 138.4s for MPC with Intel Core i7-8565U CPU @ 1.80GHz.
- Considering the tracking accuracy and computing cost, the proposed PP+PPO is effective and suitable for real-time control.

Berkeley DeepDrive

Challenges in Realizing AV

- **What must be overcome?**
  - ➢ Technical
  - ➢ Safety concerns
  - ➢ User Expectation
  - ➢ Market Acceptance
  - ➢ Legal
  - ➢ ...

Challenges in Realizing AV

- Misuse & Malfunctioning

Tesla Fatality Incident (May 2016, Florida)

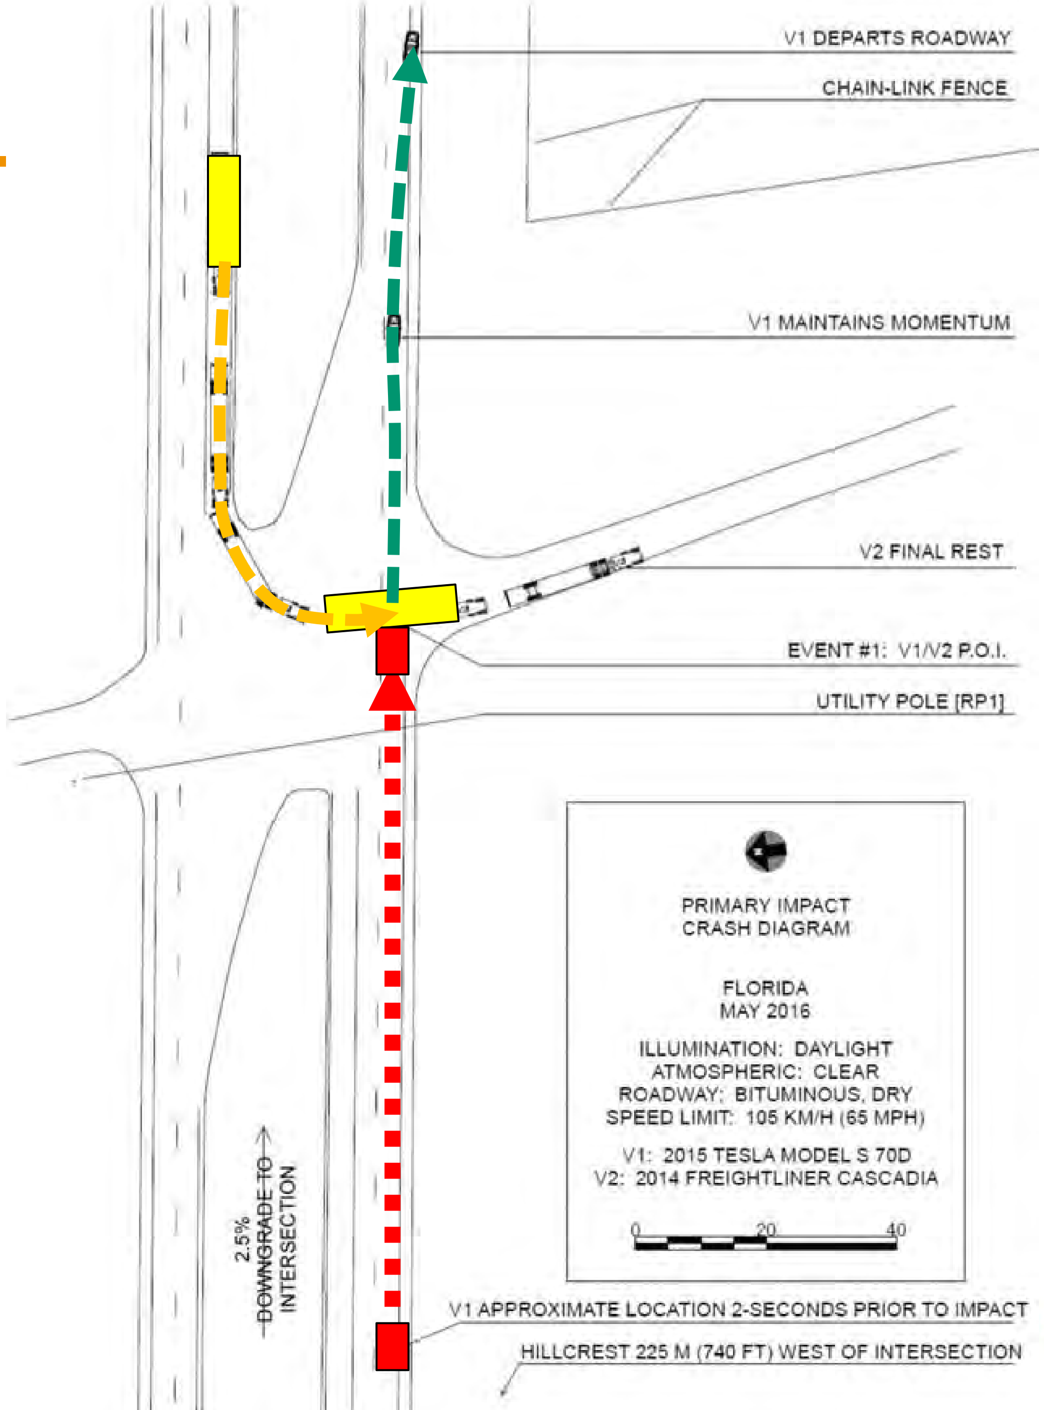

What Happened?

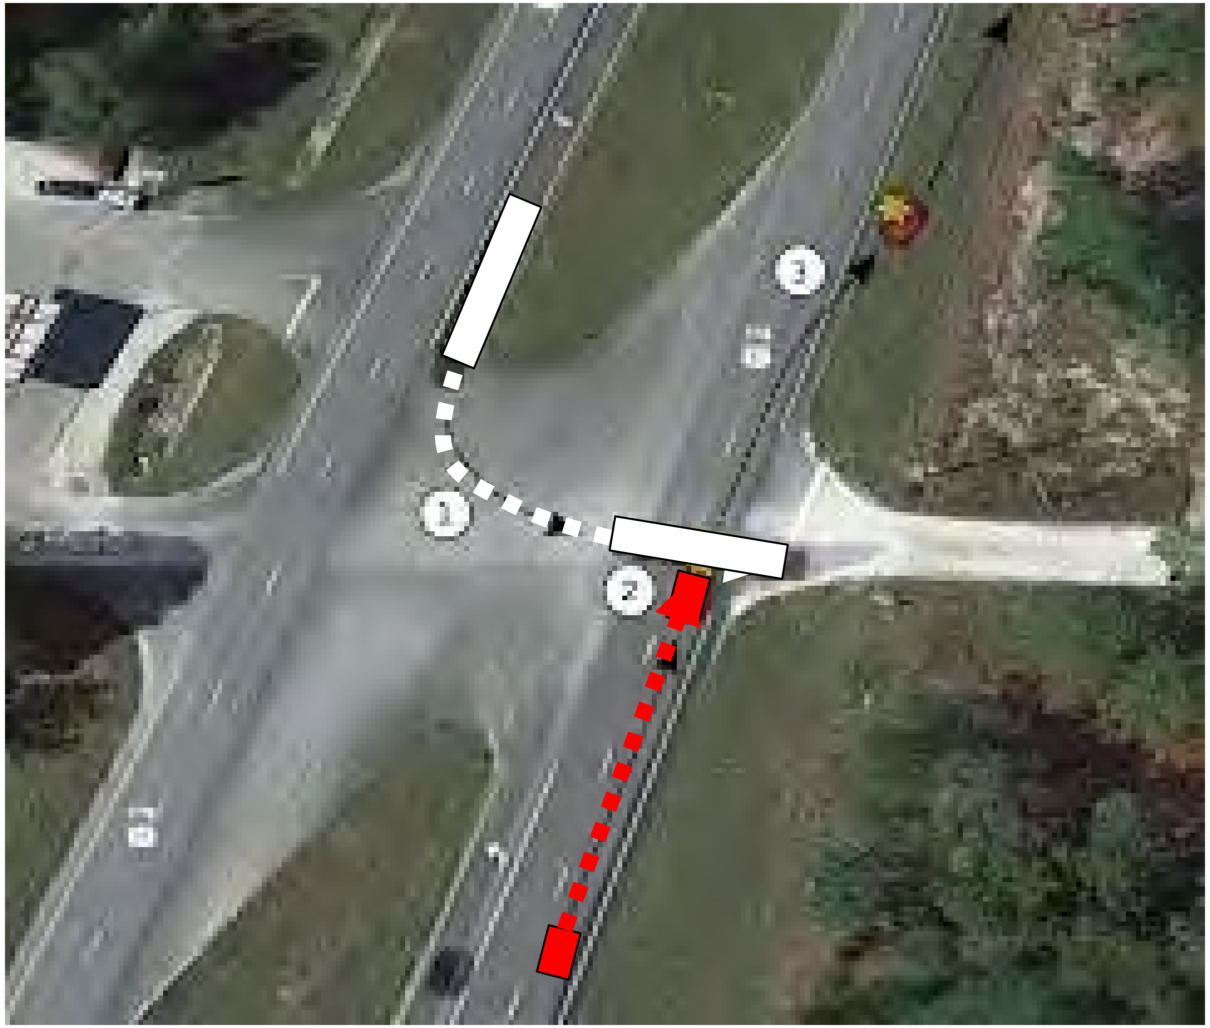

Berkeley DeepDrive

Tesla Fatality Incident (May 2016, Florida)

- Tesla cruise control set at 74 mph.
- Neither human nor the machine hits the brake.
- Driver's hands on wheel for only 25 seconds during 37-minute period
- Driver ignored 7 visual warnings and 6 audible warnings during the trip

Berkeley DeepDrive

déjà vu??

Tesla Fatality Incident (March 2019, Florida)

- Autopilot activated 10 seconds before
- Tesla cruise control set at 68 mph.

- Neither human nor the machine hit the brake.
- Almost a complete duplicate of 2016 accident,
  - Despite hardware and software updates

Tesla Fatality Crash (Mountain View, CA, March 2018)

Tesla 03/18 Fatality Crash

Berkeley DeepDrive

Approach to Divider

Errors in:

Lane recognition?

Control decision?

https://weibo.com/tv/v/63299a963968b9cbdb7ffc1d44c83b37?fid=1034:63299a963968b9cbdb7ffc1d44c83b37

Berkeley DeepDrive

Automated Driving Systems per SAE 3016

❖ Automation Level	❖ Human / Machine Roles
• I	• Feet or Hands off
• II	• Feet and Hands both off
• III	• Eyes (Perception & Response) off
• IV	• Mind off
• V	• Driver off

- We can say "machine is automated up to this level", but
- Do we have much say in "what drivers will do"?

Take Over, Please?

Berkeley DeepDrive

Challenges in Realizing AV

- ➢ Misuse and Associated Accidents
- ➢ **Safety Performance**

Berkeley DeepDrive

How Bad Are Human Drivers?

- **>90%** crashes caused by driver errors.
- In US the numbers are roughly **one fatality** per **80+M vehicle-miles travelled.** (per National Safety Council 2015-2016)
  - **80 million vehicle-miles per fatality ~= 20,000 miles/per year X 50 years X 80 life-times**
- How many ADS teams are confident that their systems can do better than the above?

Berkeley DeepDrive

How Safe is Safe Enough?

- **In comparison to other travel options**
  - **Are we creating a new travel mode, compatible with an existing and much better performing system?** **or**
  - **Are we providing a crew of drivers that can perform better than current (human) drivers on average?**

Here ADS

or

Here? ADS

How Safe is Safe Enough?

- If AVs are better than average drivers statistically (total numbers and severity of incidents) but sometimes make mistakes, can we accept and deploy them?
- What are additional safety risks?
- How small is acceptable?

Challenges in Realizing AV

- Misuse and Associated Accidents
- Safety
- **Expectation and Reality**

Berkeley DeepDrive

How Soon will it become widely Available?
(Source: Silicon Valley Mobility, 07/2017)

Autonomous driving market forecasts align with characteristics from other innovation and provide a perspective for the next 30 years

Realistically, this will take 2-3 decades, if not longer.

But, investors who have put in 10B’s want to have their return soon (now).

Errh, Overestimation?

Waymo is worth $~~175~~ 105 billion, Morgan Stanley says.

1. **Things are not happening fast.**
2. **Waymo will still be a leader by 2030.**

Source:
Various business news, 08/2018
Update, 09/2019

Berkeley DeepDrive

Will it (robot taxi) be here?

GM's Cruise Plans to Operate Autonomous Taxi ~~in 2019~~ Later

- ~~*The Street, 06/2018*~~
- *Various media reports, Postponed, 07/2019*

Waymo CEO Says Self-Driving Cars Won't Be Ubiquitous for Decades

John Krafcik, Waymo CEO

- Bloomberg, 11/2018

No Self-Driving Cars in my life time (probably)

Steve Wozniak, Apple Co-Founder

- Various Media, 10/2019

There'll Be Hundreds Of Self-Driving Cars On The Road In Five Years

Chris Urmson, Formerly Waymo, now CEO of Aurora

- Forbes, 10/2019

We will have self-flying cars before self-driving cars

Sebastian Thrun, Google self-driving team founder

- Various Media, 10/2019

AI is Great for AV, but ...

Challenges When AI Meets AV

- **Culture Clash**

DARPA Challenges and Google AV

- DARPA challenges in 2000s’ stirred up momentum in AV technologies
- Google established AV Google-X team
- Automakers and technology industry powered ahead the push into ADS
- Wall Street craze for mobility service enticed investments into AV

Berkeley DeepDrive

Automotive Industry versus Technology Industry

Automotive Sector	Technology Sector
• Conservative	• Aggressive
• Slow moving	• Fast paced turnover
• HW manufacturer	• Software innovator
• 4-5 year model cycle + 10-20 years usage	• 6 months – 2 years design and replacement
• Let’s make sure.	• We can do it!

- Two seemingly incompatible sectors, now competing in the same sector

Challenges when AV meets AI

- ➢ Culture Clash
- ➢ **Achieving common-sense AI for AV**

Berkeley DeepDrive

Real World is More Complex than Training

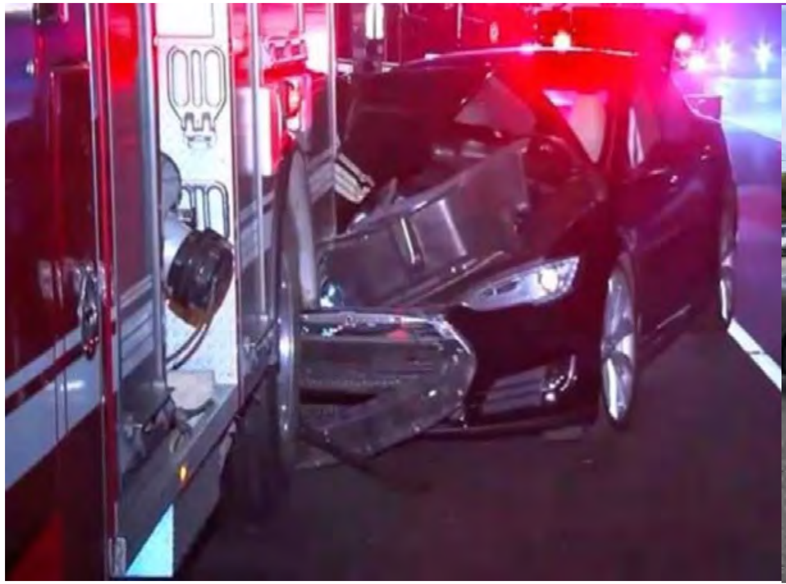

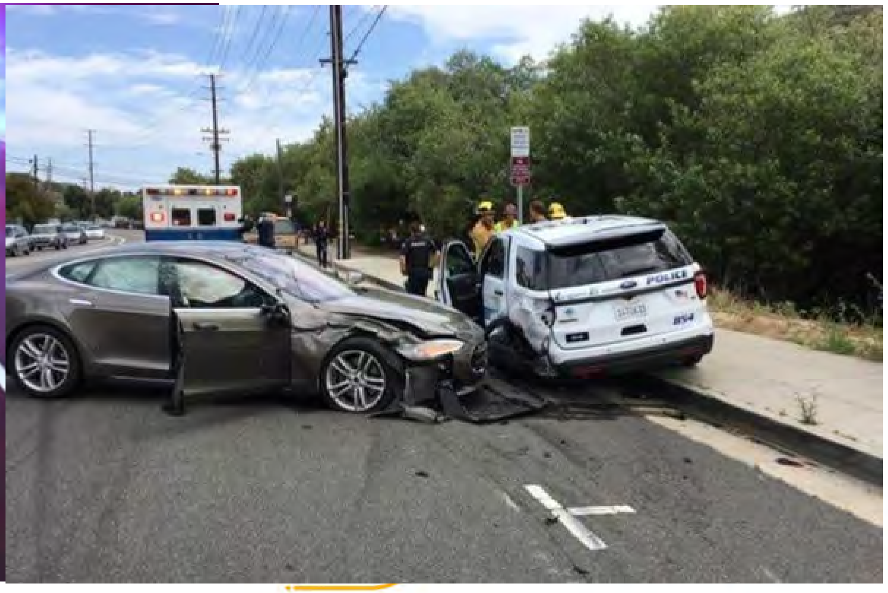

Unusual Suspects

ey DeepDrive

Challenges

Opportunities

Opportunities

- **Misuse & Reliability**

  → **New and Supplementary Technologies**

  - **lidar, driver monitoring, HD map, machine learning, ….**

Emerging and Improving Technologies

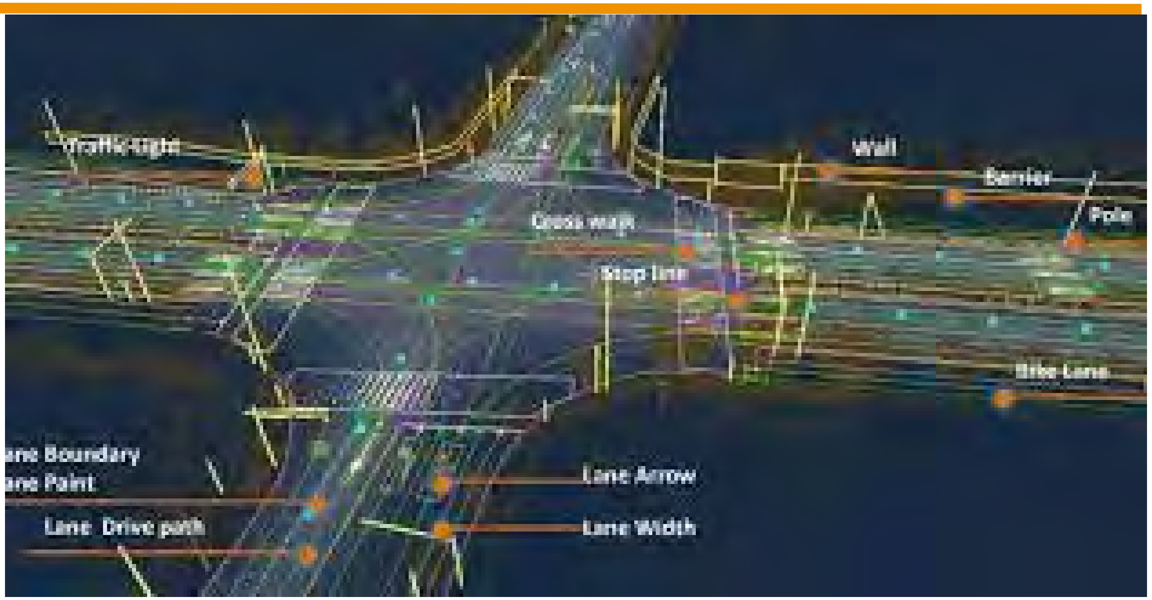

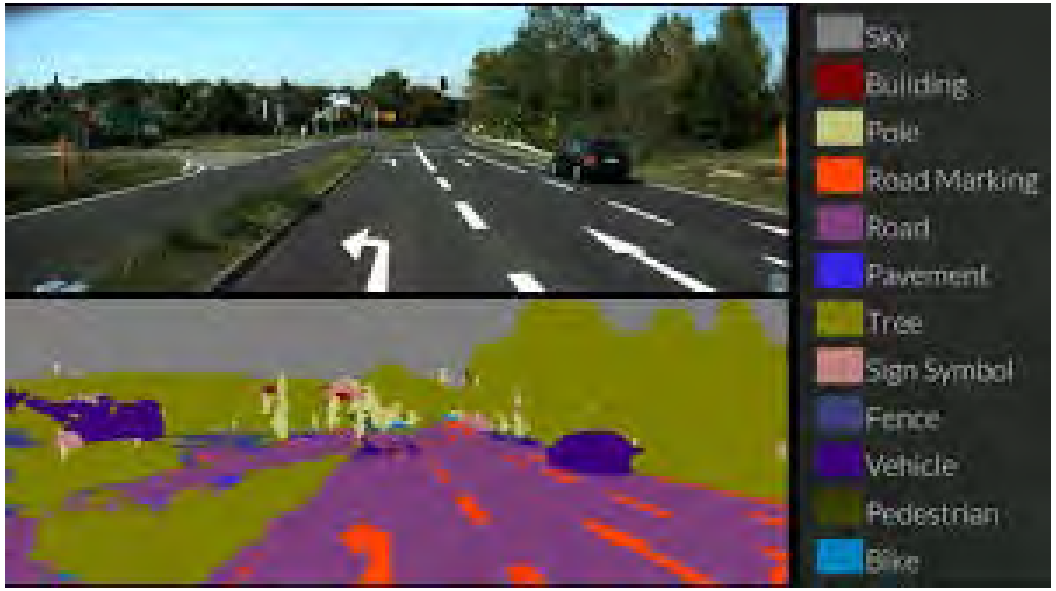

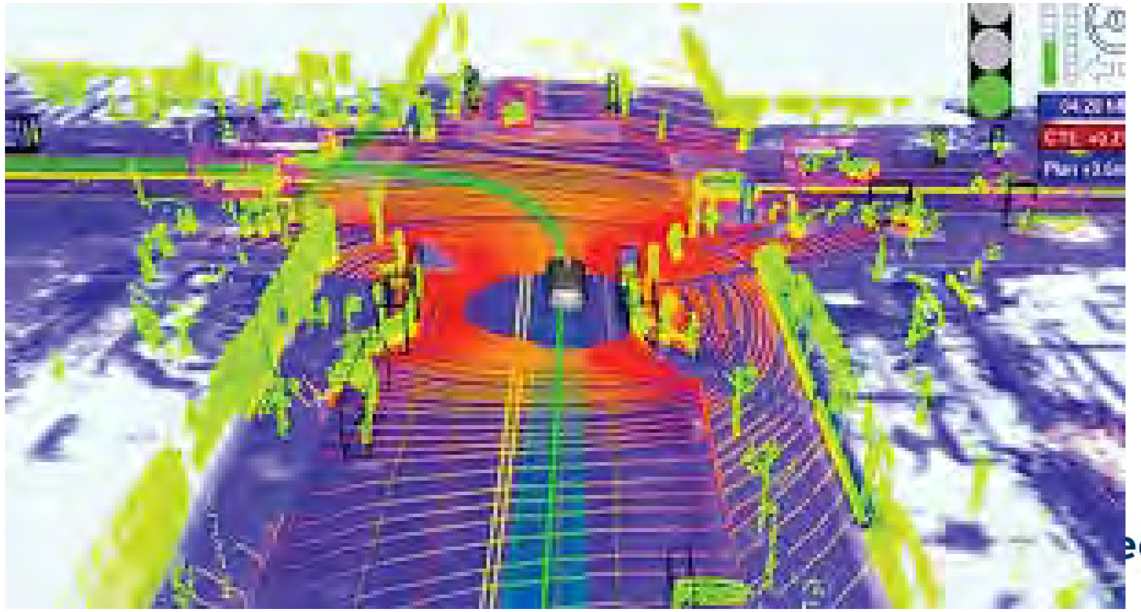

Opportunities

- **Safety**

    → **Risk Management (Limited ODD)**

Safety Management

- Safety Risk =
  Exposure X Severity X Controllability

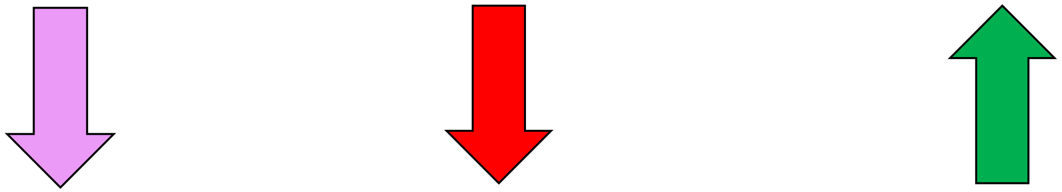

- Limited Operational Design Domain (ODD)
  - Robot Taxi in Geofenced Area
  - Low speed driverless shuttle
  - …

Opportunities

- **Expectation and Delayed Deployment**

  → Special use cases in mining, airport, factory, corporate campus, retirement villages, etc.

Near Term Deployment

Opportunities

- **Toward Human Intelligence**

  → Continuing Progress in Machine Learning

Berkeley DeepDrive

Machine Learning

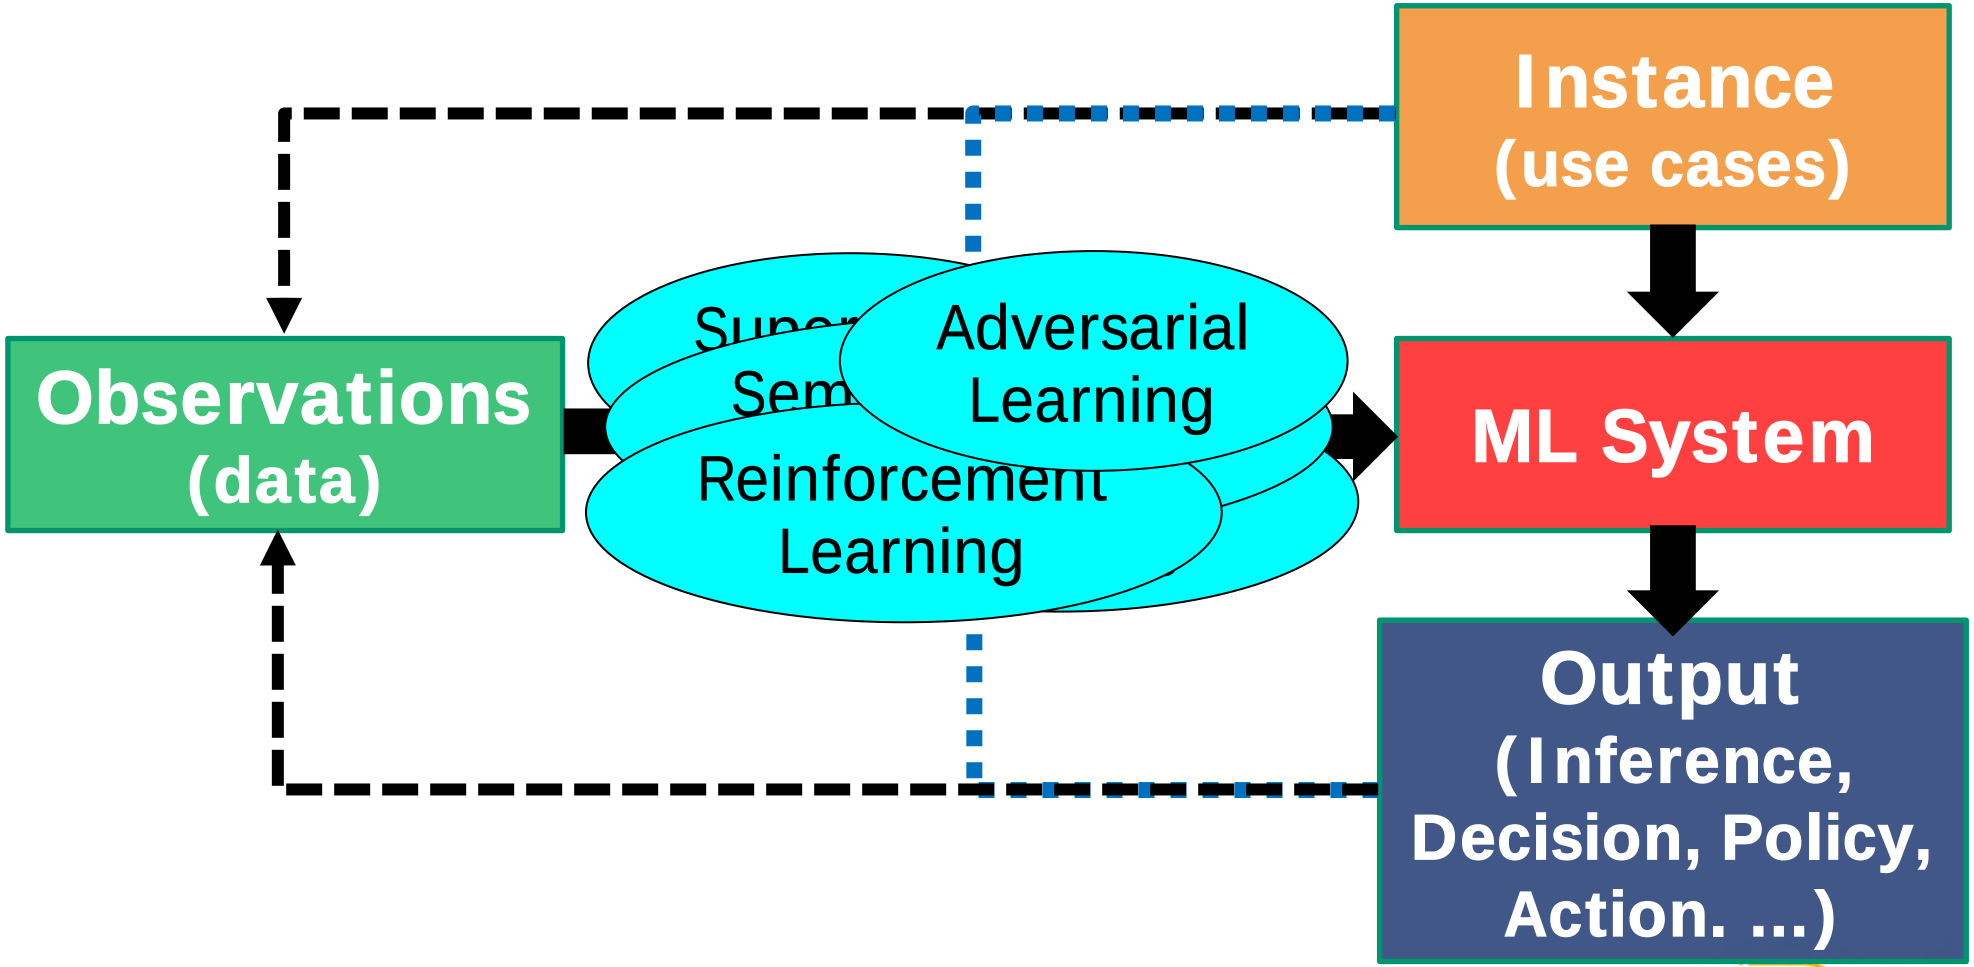

Berkeley DeepDrive

Lifelong Learning

Motivation:

Life-long learning that allows the machine to generalize and to adapt to new environment/ tasks

Approach:

While general AI is not here yet,

Combinational use of various machine learning techniques offers an alternative.

Berkeley DeepDrive

Concluding Remarks

Demo '97: Platoon Scenario

- Development by California PATH, General Motors, Delphi, Delco Electronics, Hughes and Caltrans
- Showcase a fully automated highway system
  - Lateral control based on magnetic guidance
  - Longitudinal control at 6 meters spacing
  - Automated lane changing
  - Coordination layer control

Then and Now

	Then (NAHSC, 1990s)	Now (ADS+AV, 2010s)


SOHJOA.FI
EZ10


prime

This is an exciting era for AI and AV.

A.I. – The revolution hasn't happened yet.

Michael Jordan, UC Berkeley

Berkeley DeepDrive

Towards Human-Level AI

- AI systems will eventually make better decisions than humans.
- Still missing:
  - Real understanding of language
  - Integration of learning with knowledge
  - Long-range thinking at multiple levels of abstraction
  - Cumulative discovery of concepts and theories
- When?
  - Date unpredictable

*Provably Beneficial AI, Stuart Russell, UC Berkeley, BDD Workshop, 11/2018

Can We Accept Imperfect Automation or AI?

Intelligence ≠ Perfection

Artificial or Human

Berkeley DeepDrive

What to Expect, as AI meets AV?

- *AV is a revolutionary transformation for the automotive and transportation industries.*
- *The revolution in AI will propel AV further forward in a significant manner.*

Berkeley DeepDrive

Thank you.

Ching-Yao Chan
cychan@berkeley.edu

https://deepdrive.berkeley.edu/

Berkeley DeepDrive