# Safety Challenges for Autonomous Vehicles in the Absence of Connectivity

Akhil Shetty
Department of Mechanical Engineering
UC Berkeley

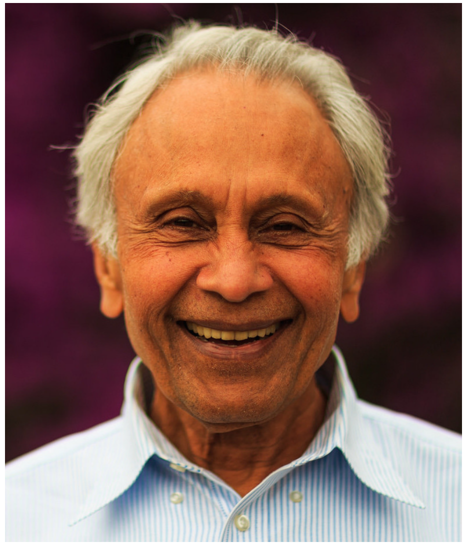



# Promise of Automated Vehicles (AVs)


- >1M lives lost each year due to traffic crashes
- Aim to eliminate crashes by replacing human drivers
- >$100B invested in the AV dream
- Will save 600k lives and $230B in economic costs by 2045
- AV future just around the corner…

Sources: NHTSA, Brookings, Strategy Analytics
Image: Freepik

# Reality Check

NATIONAL

## Nearly 400 car crashes in 11 months involved automated tech, companies tell regulators

June 15, 2022 · 1:26 PM ET

By The Associated Press

## It's 2020. Where are our self-driving cars?

In the age of AI advances, self-driving cars turned out to be harder than people expected.

TRANSPO / AUTONOMOUS CARS / TECH

## Cruise robotaxi collides with fire truck in San Francisco, leaving one injured

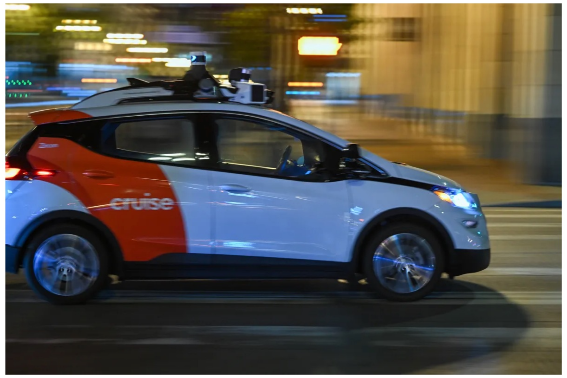


Photo by Tayfun Coskun / Anadolu Agency via Getty Images

/ One week after California approved 24/7 paid robotaxi services in San Francisco, a crash occurred between an autonomous Cruise taxi and a city fire department truck late at night.

By Umar Shakir, a news writer fond of the electric vehicle lifestyle and things that plug in via USB-C. He spent over 15 years in IT support before joining The Verge.

Aug 18, 2023, 9:49 AM PDT | 41 Comments / 41 New

## 17 fatalities, 736 crashes: The shocking toll of Tesla's Autopilot

Tesla's driver-assistance system, known as Autopilot, has been involved in far more crashes than previously reported

## San Francisco Experiencing Traffic Jam After Cruise, Waymo Get Green Light to Operate 24/7 Robotaxi Services

*Cruise's self-driving cars reportedly caused traffic issues in San Francisco's North Beach area.*



Images: Vox, Associated Press, Wall Street Journal, The Verge, Tech Times

# A Well-Known Statistic

94% of all crashes caused by human error.

# A Well-Known Statistic?

~~94%~~ of all crashes caused by ~~human error~~.

# What the NHTSA report actually says…

*"The critical reason was assigned to drivers in an estimated 2,046,000 crashes that comprise 94 percent of the NMVCCS crashes at the national level. However, in* ***none of these cases was the assignment intended to blame the driver*** *for causing the crash."*

# What the NHTSA report actually says...

*"The critical reason was assigned to drivers in an estimated 2,046,000 crashes that comprise 94 percent of the NMVCCS crashes at the national level. However, in* ***none of these cases was the assignment intended to blame the driver*** *for causing the crash."*

If not human error, what causes these crashes?

# Why do Crashes Occur?


- Vehicle Failure
  - Steering, braking, wheel defects
  - 1% of all crashes [1]
- Overt Driver Error
  - Drunk, Speeding, Traffic Violation
  - <60% of all crashes [1, 2]
- (Seemingly) Rational Driving
  - Occlusions, misjudging behavior
  - Uncertainty about state/behavior of neighboring vehicles

Information Gap!

Source: [1] NHTSA 2019
[2] Phil Koopman, Safe Autonomy Blog.

# A Taxonomy of Crashes

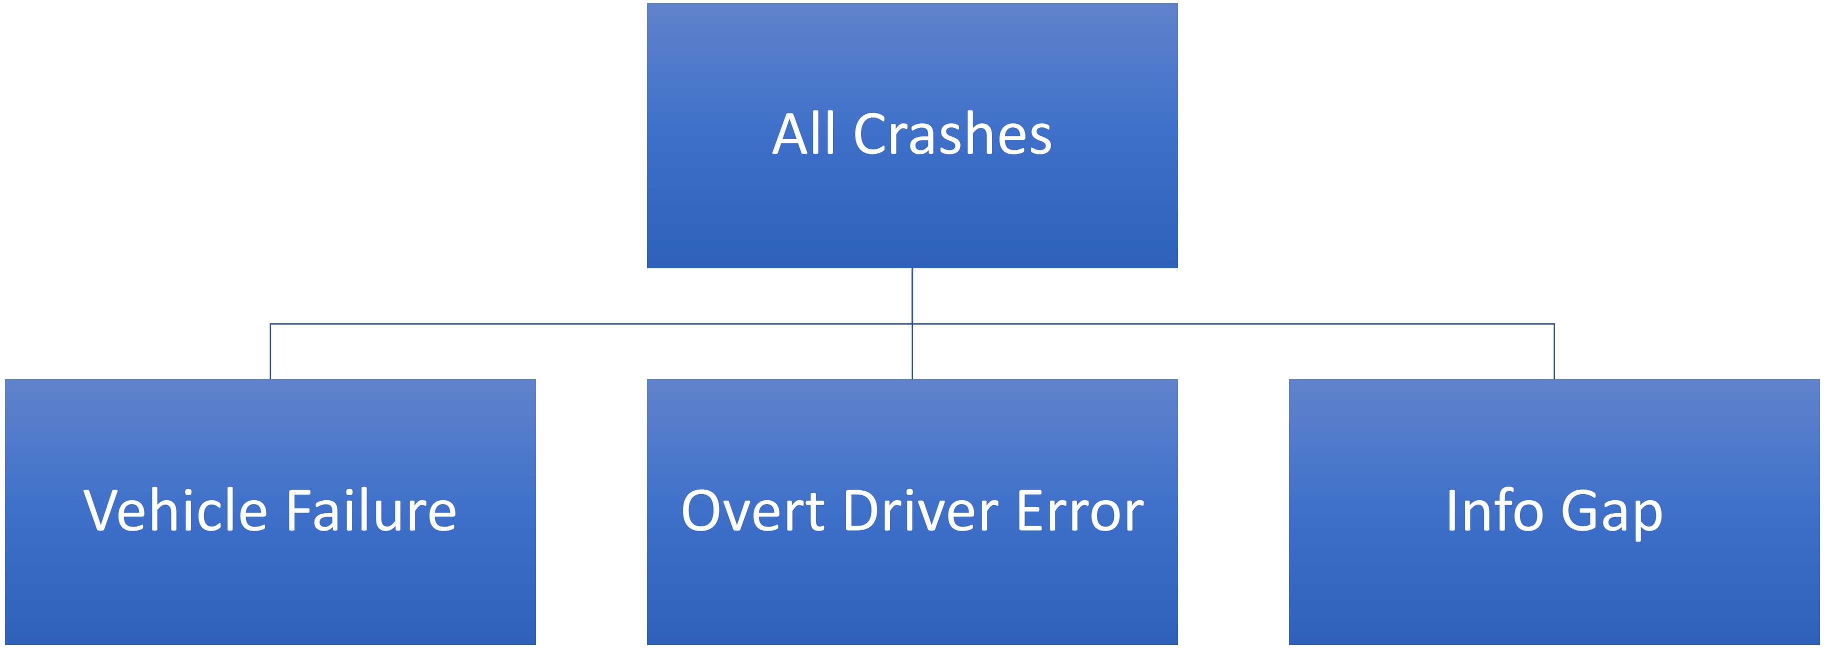


- Better Auto tech will eliminate

- Better AV tech will eliminate

?

# Information Gap

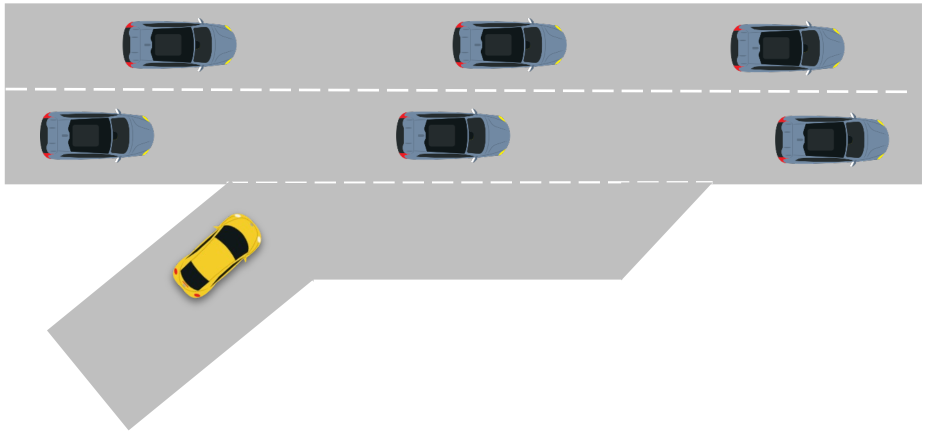

Merging – Behavior Uncertainty

- Our safety depends on positions and actions of neighboring vehicles
- If perfectly observed, AV can hope to be safe – a technology problem
- Partially observed in typical traffic scenarios
  - Occlusions - partial state
  - Behavior Prediction – only current action observed

# How can an AV be safe?

- Statistical techniques on partial observations to estimate state/action [1-3]
  - Can reduce crash probability but cannot guarantee safety

- Worst case safety
  - Mobileye's RSS Framework [4]
  - Robust to worst case across all possible states and reasonable behaviors – guaranteed safety

- Bridge information gap
  - I2V/V2V communication [5]
  - (Most) AVs do not use this – only rely on within vehicle sensing

Source: [1] Sadigh et al., Planning for autonomous cars that leverage effects on human actions; [2] Yu et al., Occlusion-aware risk assessment for autonomous driving in urban environments, [3] Schmerling et al., Multimodal probabilistic model-based planning for human-robot interaction; [4] Shashua et al., On a formal model for safe, scalable self driving cars; [5] Grembek et al., Making intersections safer with I2V communication

# The Holy Grail of Full Autonomy


Most AV companies do not rely on vehicle-to-vehicle (V2V) or infrastructure-to-vehicle (I2V) communication:

- Can deploy faster
- Cost effective
- Avoid security vulnerabilities due to malicious agents
- Humans don't need it, why should AVs?



Source: https://www.digi.com/blog/post/what-is-connected-vehicle-technology-and-use-cases

# Our Central Thesis

I2V/V2V communication is crucial for ensuring that AVs avoid crashes due to information gaps

- Worst case safety: too strict
  - AV cannot make left-turns or merge into traffic
- Estimation from partial observations
  - Cannot guarantee safety even allowing for small crash probability

# An Information Gap Crash


Source: YouTube

# Left Turn under Occlusion

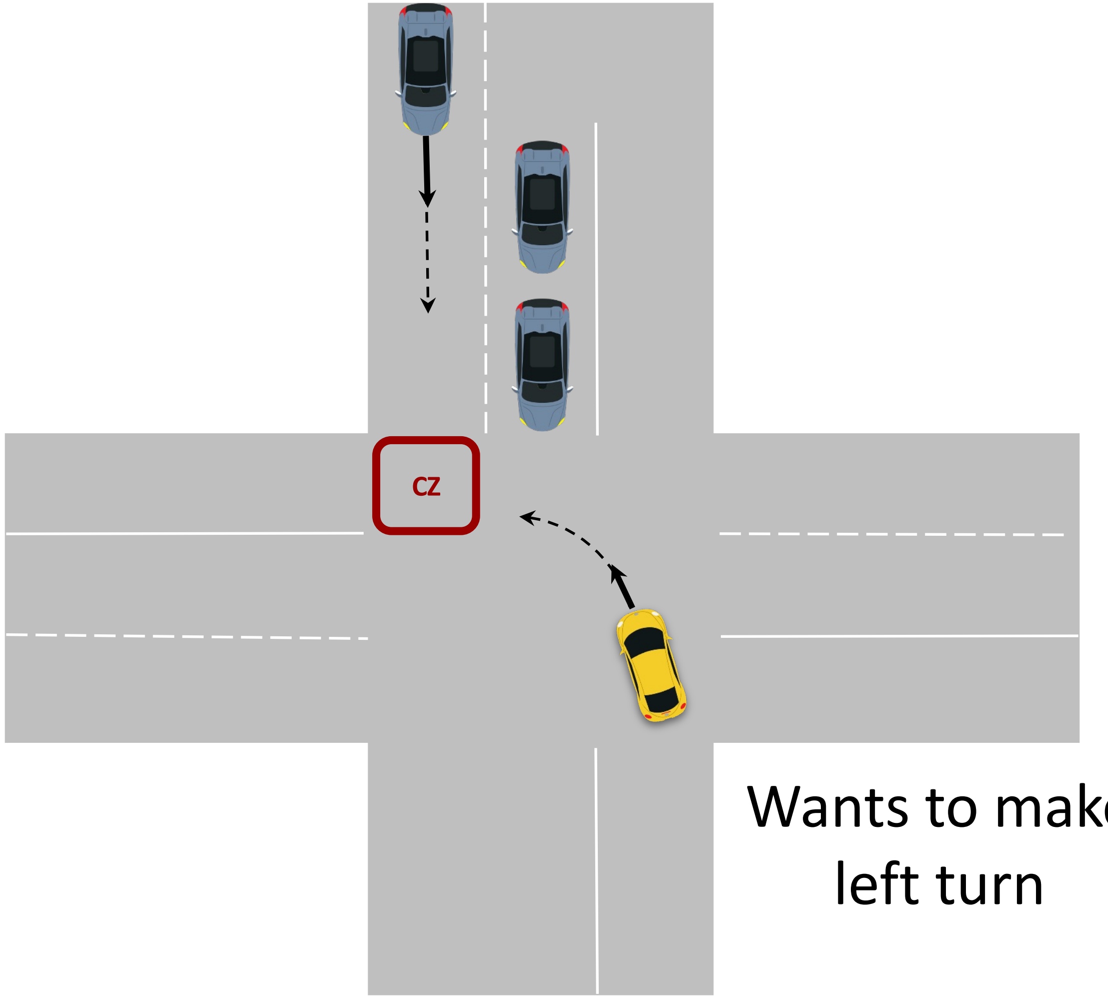


Wants to make
left turn

# Left Turn under Occlusion

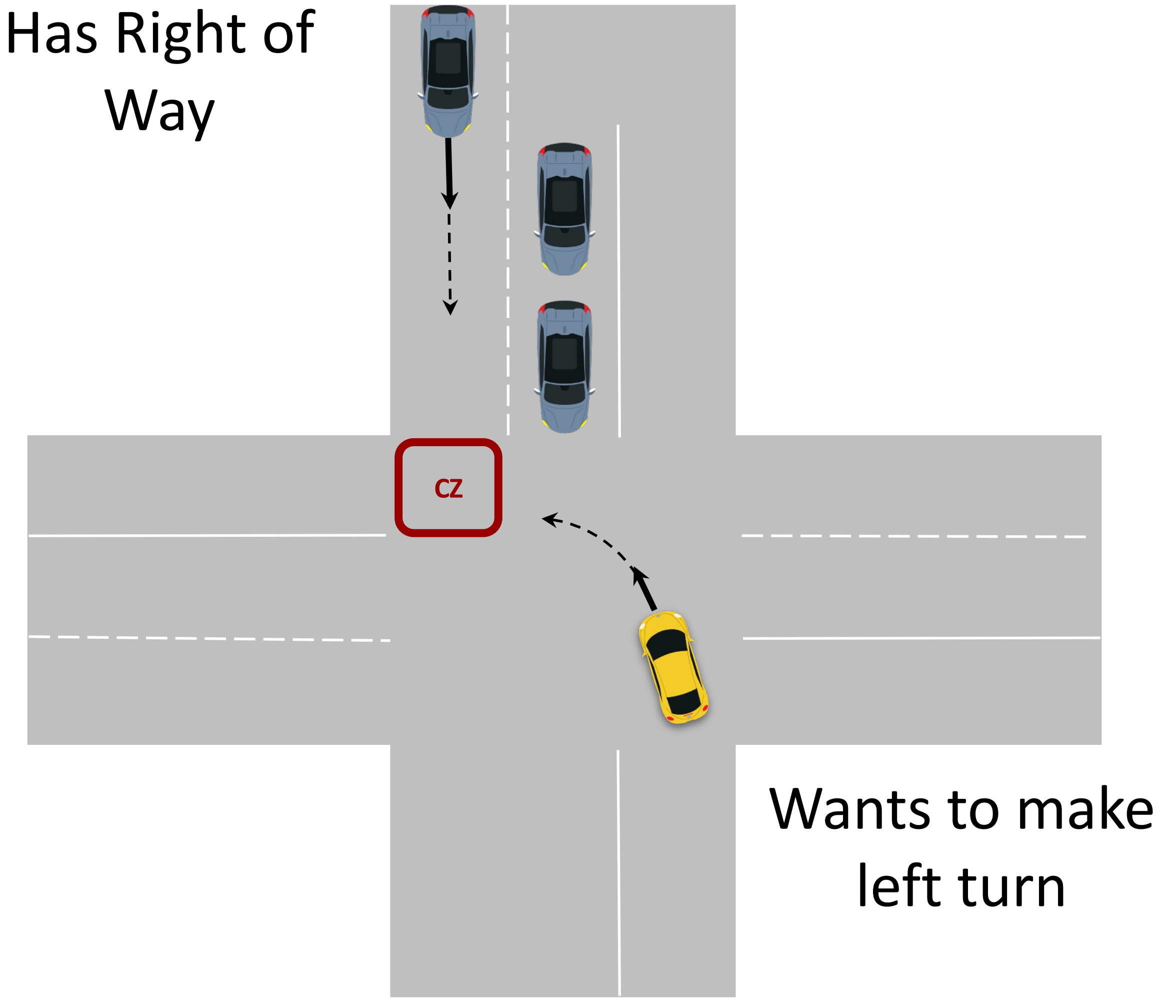

# Left Turn under Occlusion

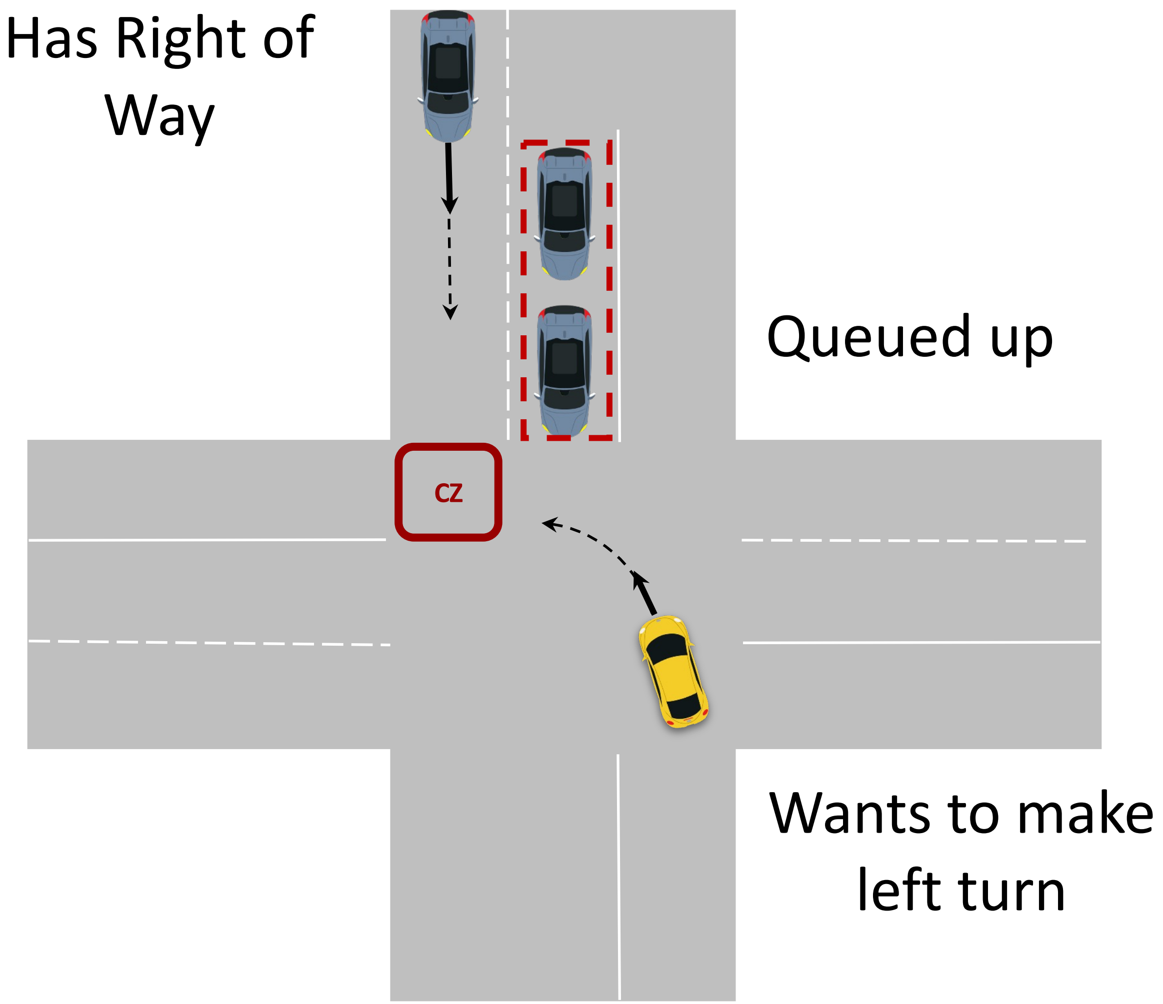

# Left Turn under Occlusion

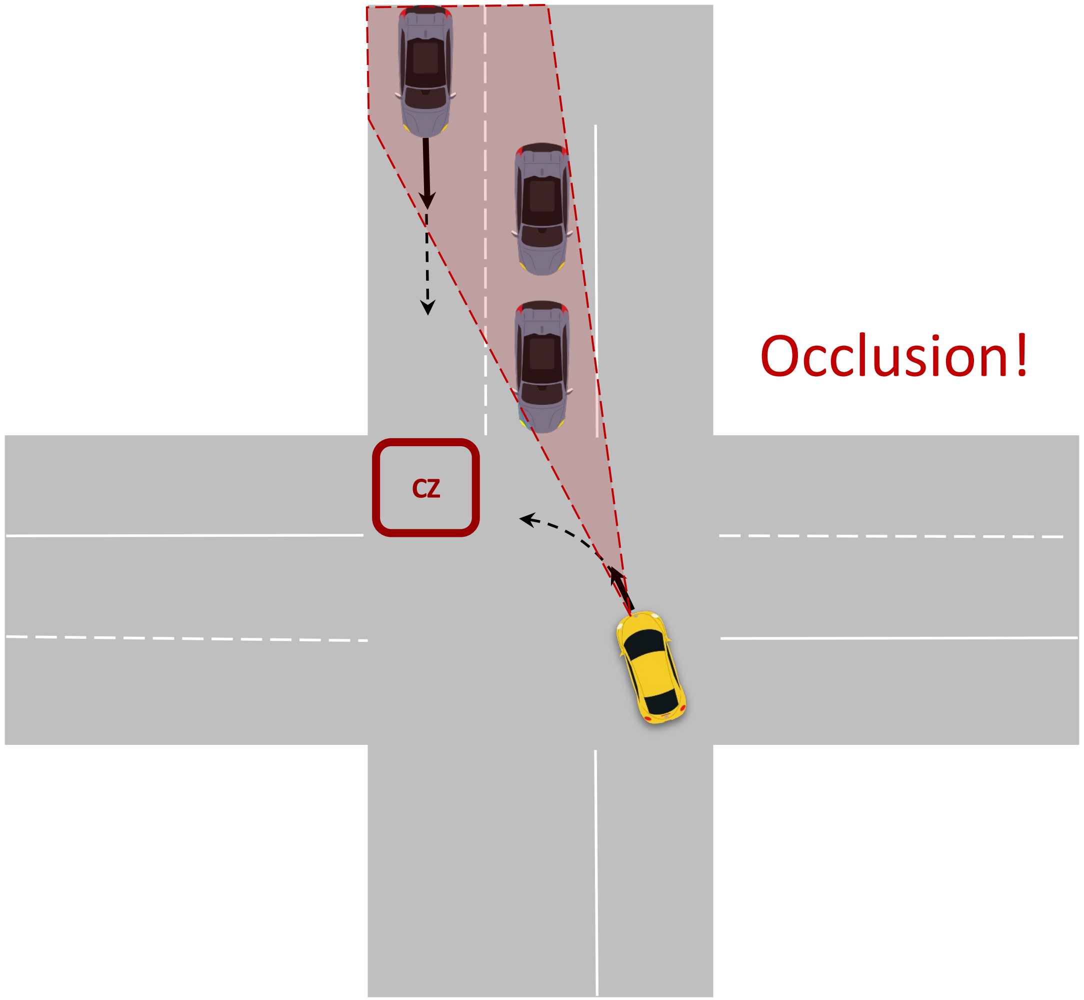

# Left Turn under Occlusion

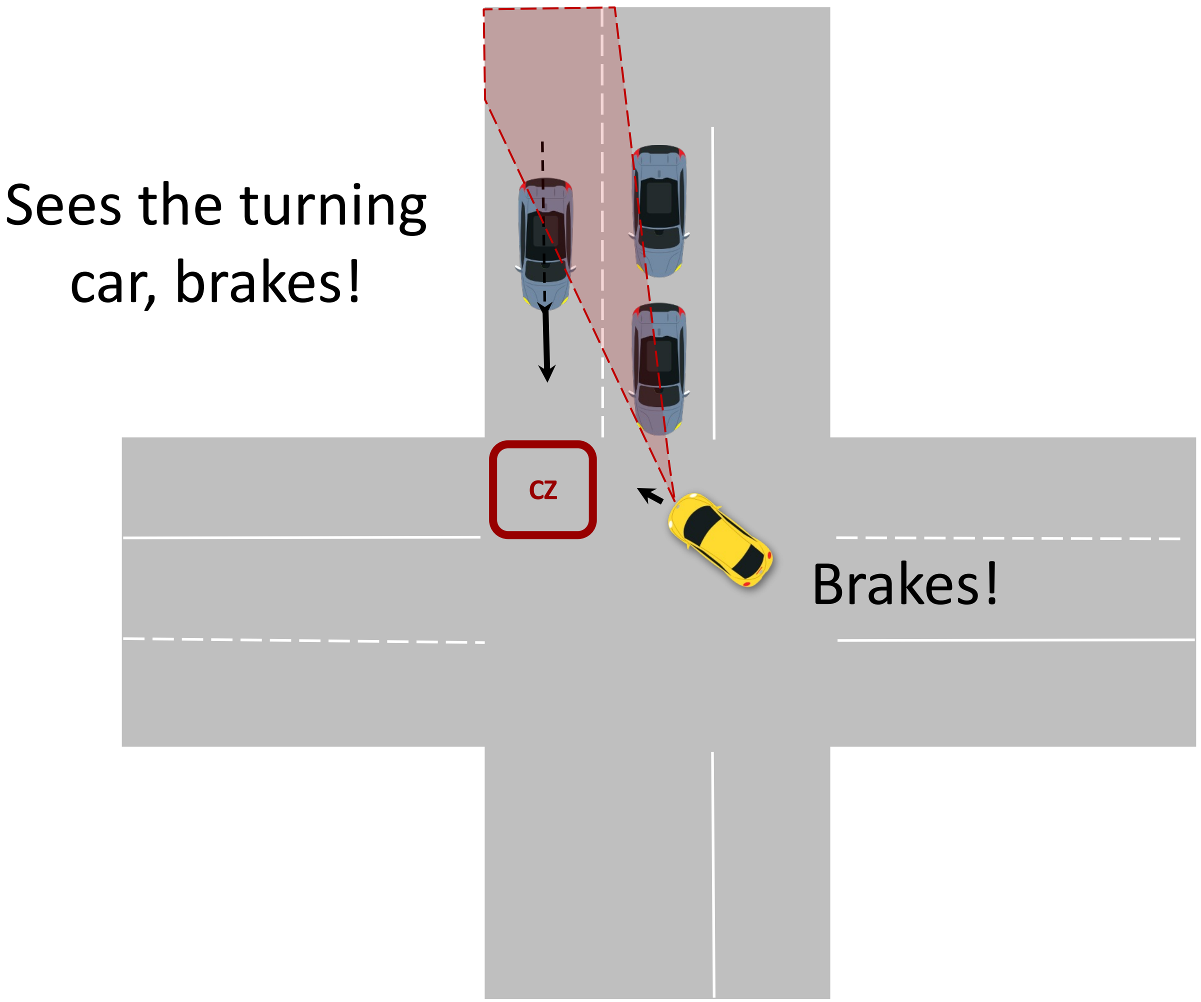

# Left Turn under Occlusion

# Can an AV avoid this fate?

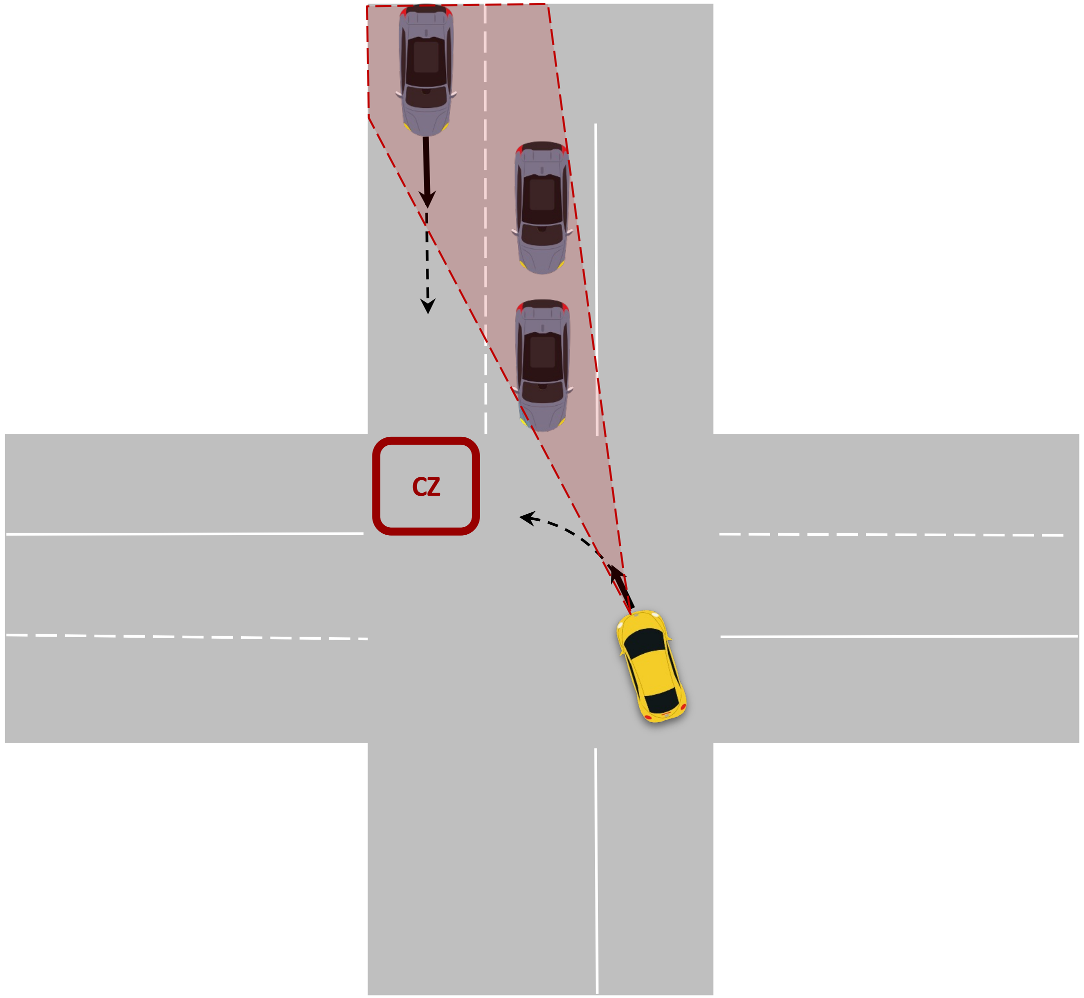

# Can an AV avoid this fate?

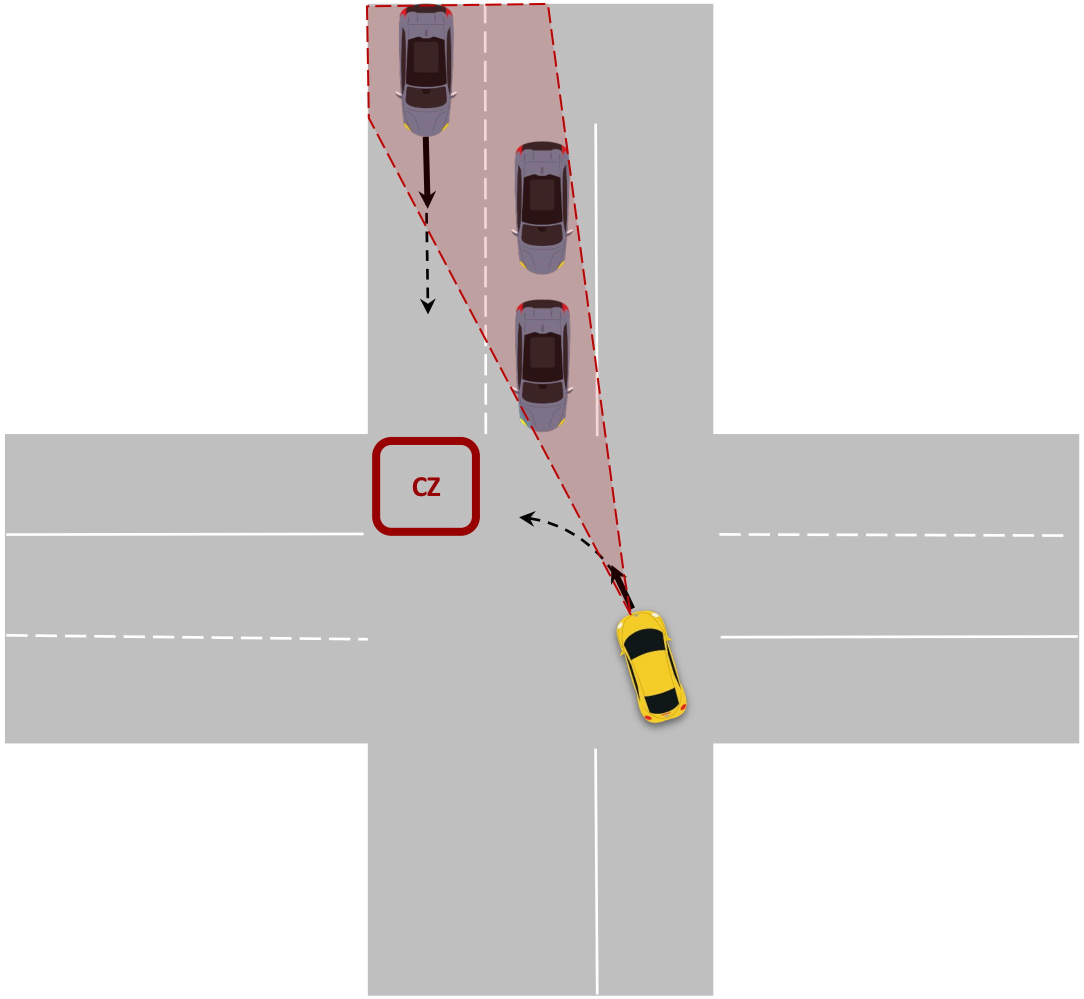


Uber AV Crash in Tempe, AZ
(March 2017)



Source: Atlantic https://www.theatlantic.com/technology/archive/2018/03/can-you-sue-a-robocar/556007/

# Can an AV avoid this fate?

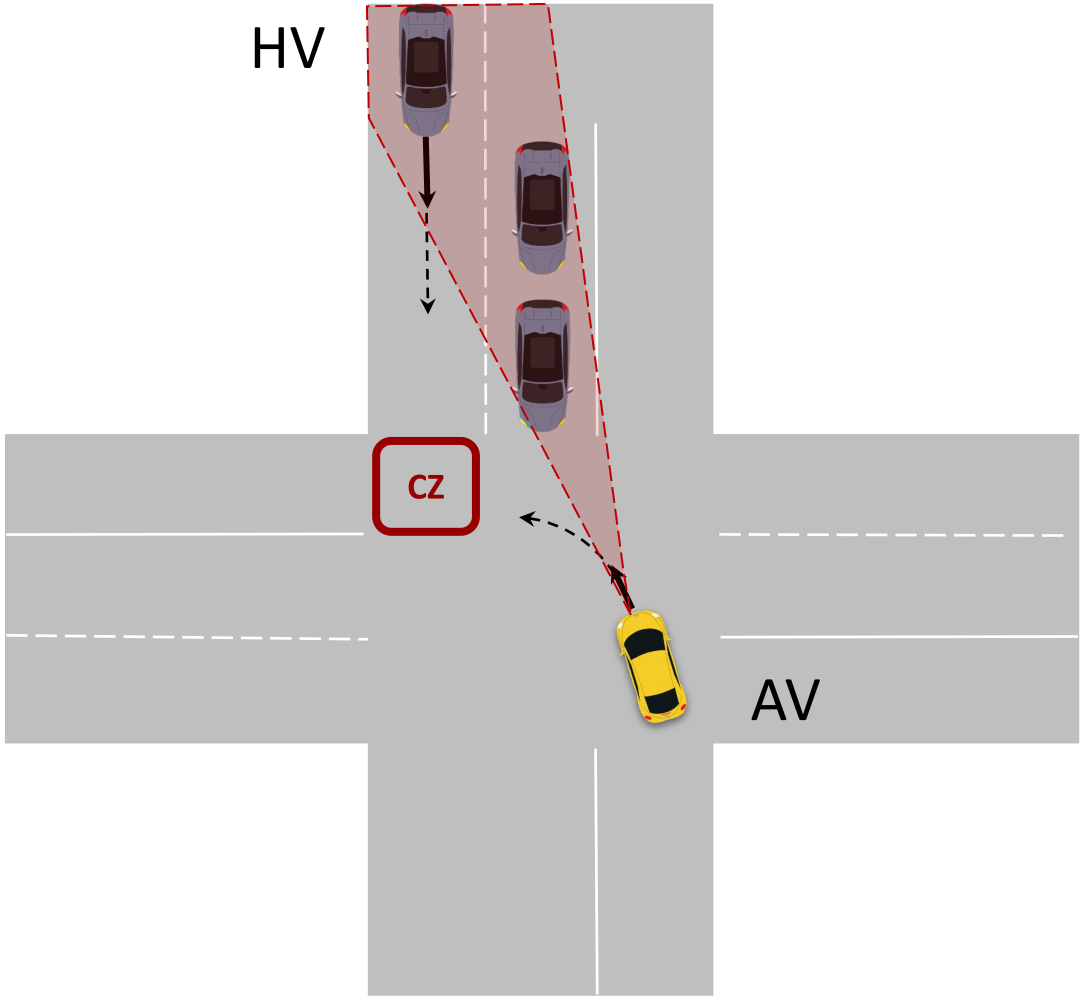


Assumptions

- AV (Yellow) has perfect sensing/perception capabilities – can "see" perfectly in its field of view
- Blue cars are human driven (HV)
  - follow traffic rules
  - no overt errors
  - best reaction time
- HV will go at its original speed until it sees the AV – attempts to brake to a stop

# Can an AV avoid this fate?

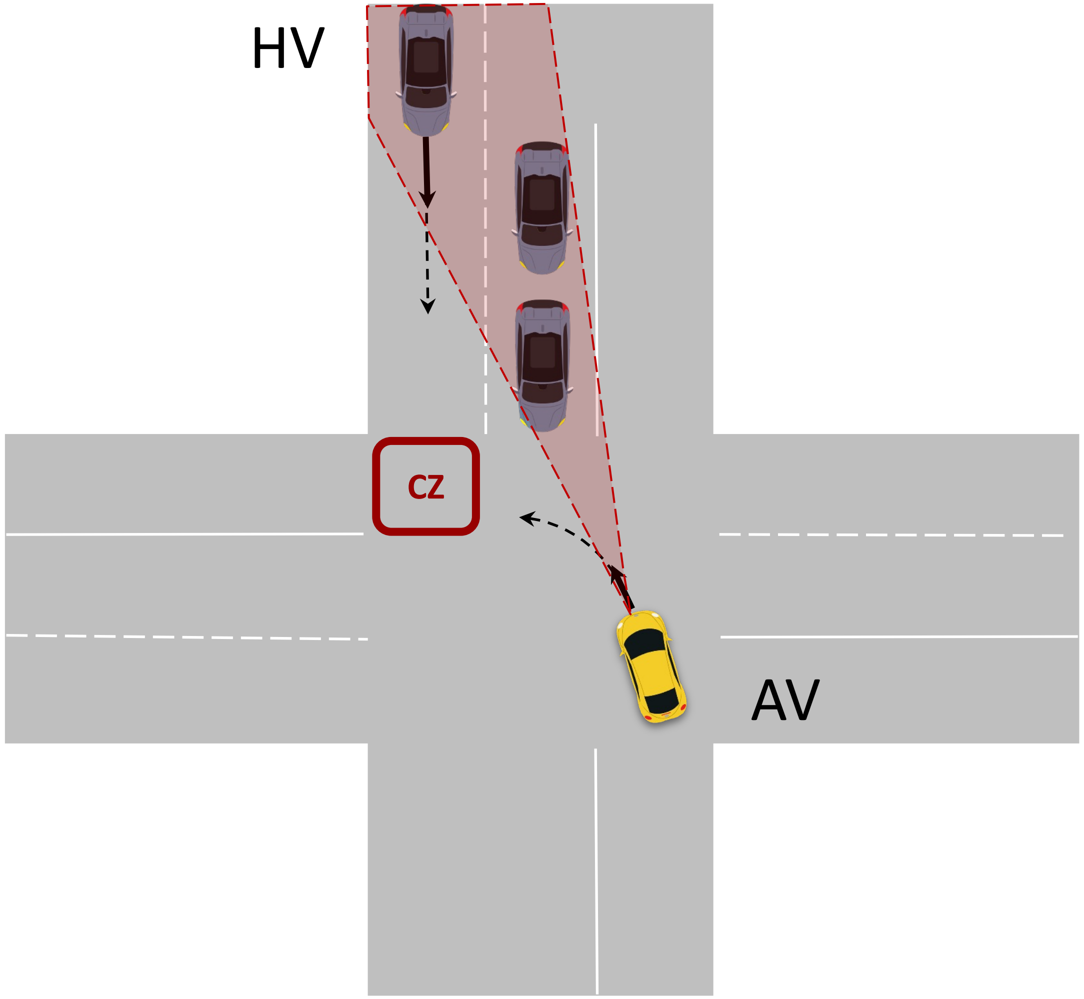


- What can the AV do to avoid crash?
  - Not much once it decides to turn: too close to CZ to stop in time
- AV can only be safe if it turns when HV is far enough from CZ or it is slow enough
- But, AV doesn't know HV's position or speed while deciding to turn!

Information Gap!



Source: Atlantic https://www.theatlantic.com/technology/archive/2018/03/can-you-sue-a-robocar/556007/

# Safety despite Information Gap

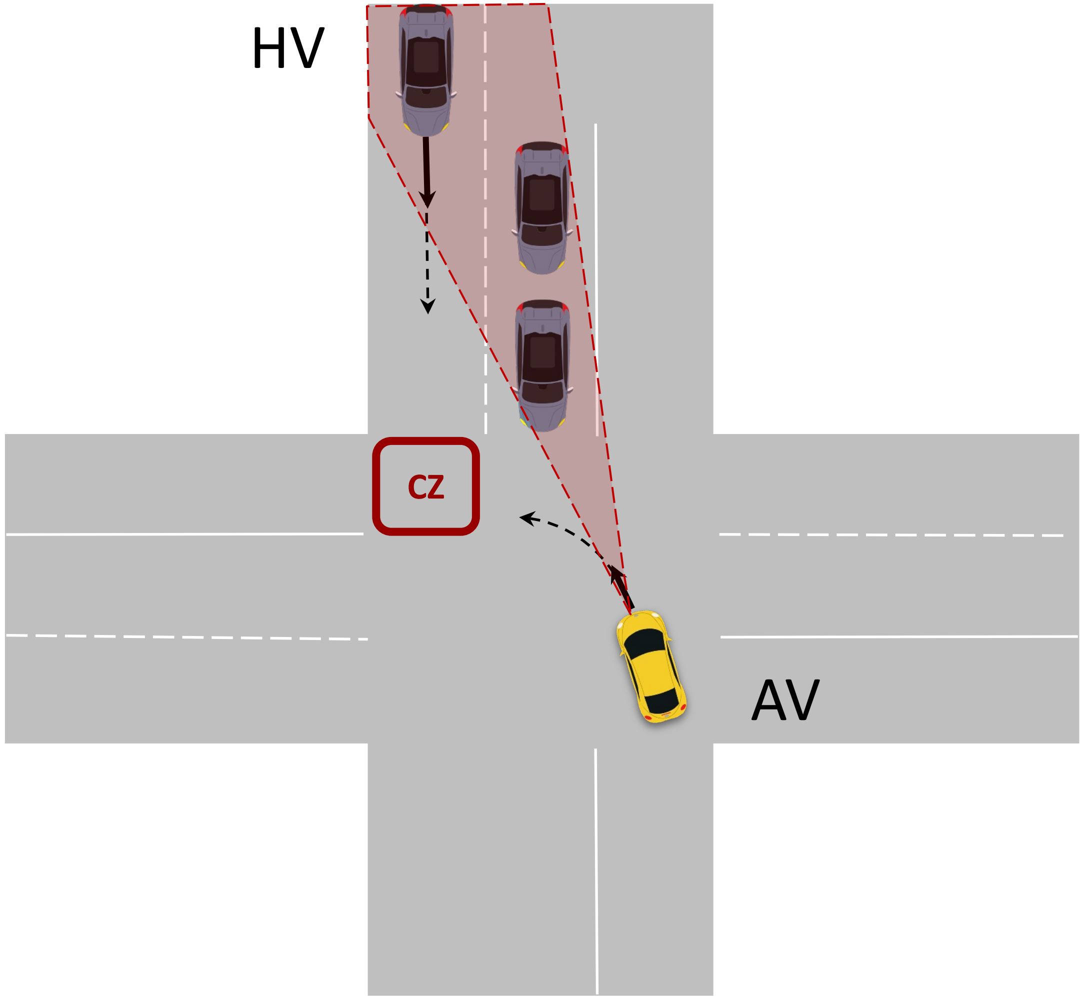


- Worst Case Safety
- Estimation from Partial Observations



Source: Atlantic https://www.theatlantic.com/technology/archive/2018/03/can-you-sue-a-robocar/556007/

# Worst Case Safety

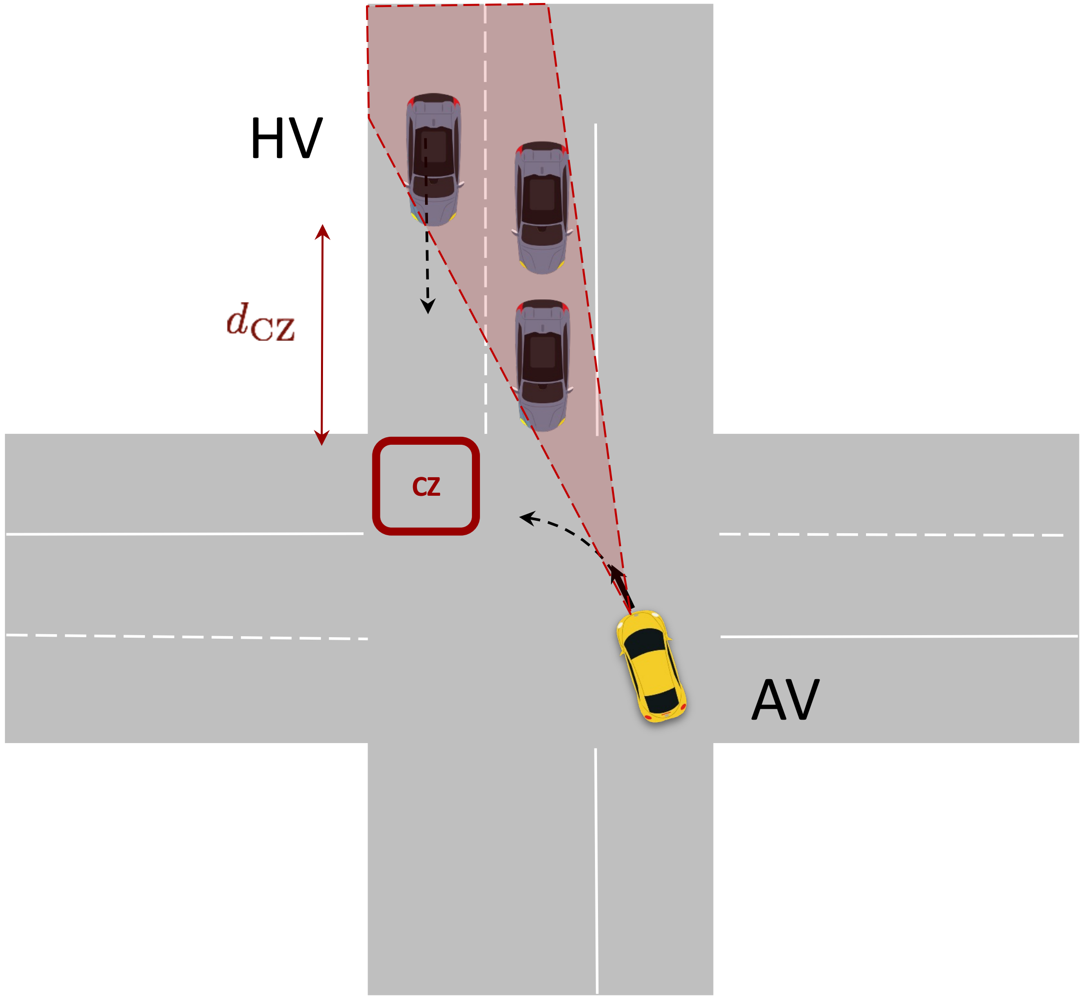


- Worst case: HV at the edge of AV's occluded field of view
- How slow should the HV be to stop in time?
  - HV's speed must be <17 mph
- Typical speed limit – 30 mph

AV cannot make left turn if it wants to be worst-case safe!

# Estimation from Partial Observations

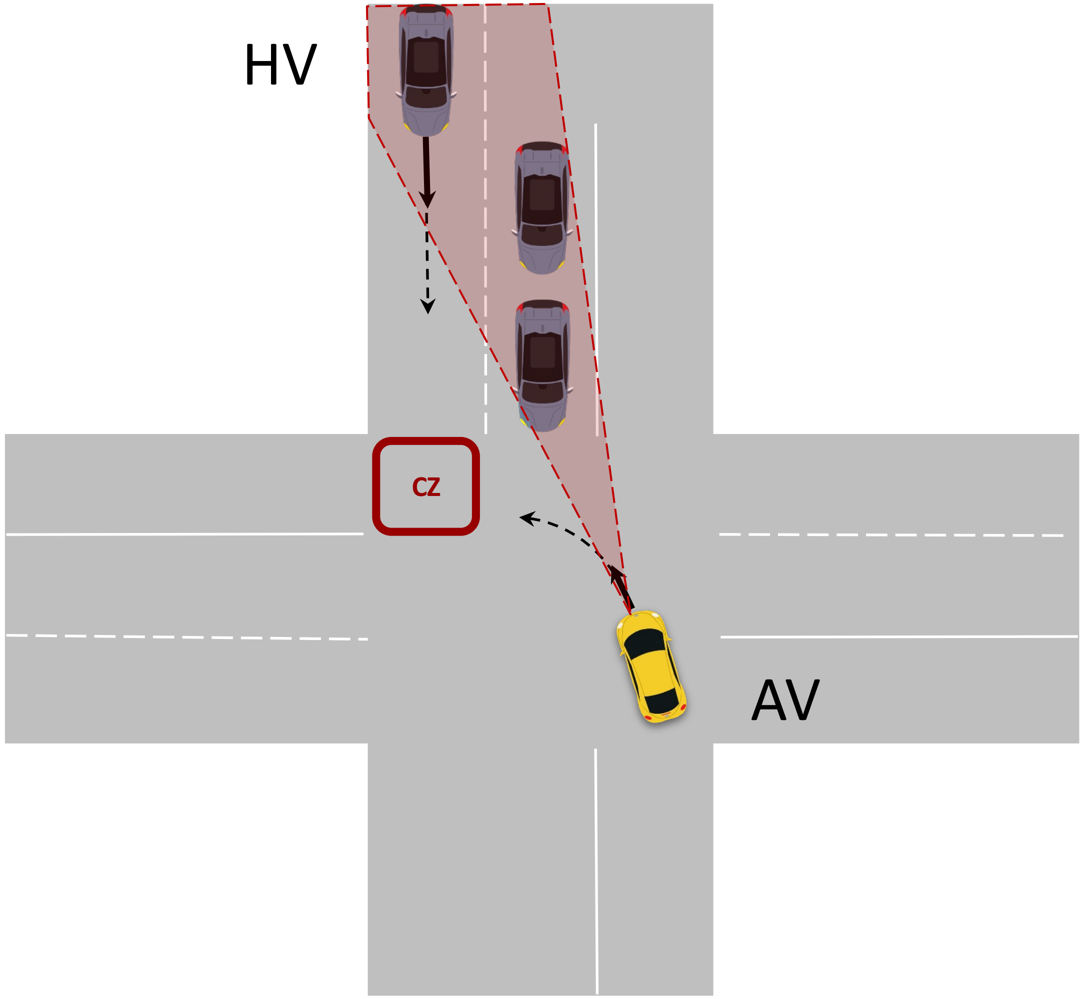


- Suppose AV is willing to accept a low crash probability
- Information Gap: HV's Position and Speed
- Partial Observation: AV observes HV flow while it waits to turn
- How can the AV use this information?

# Estimation from Partial Observations

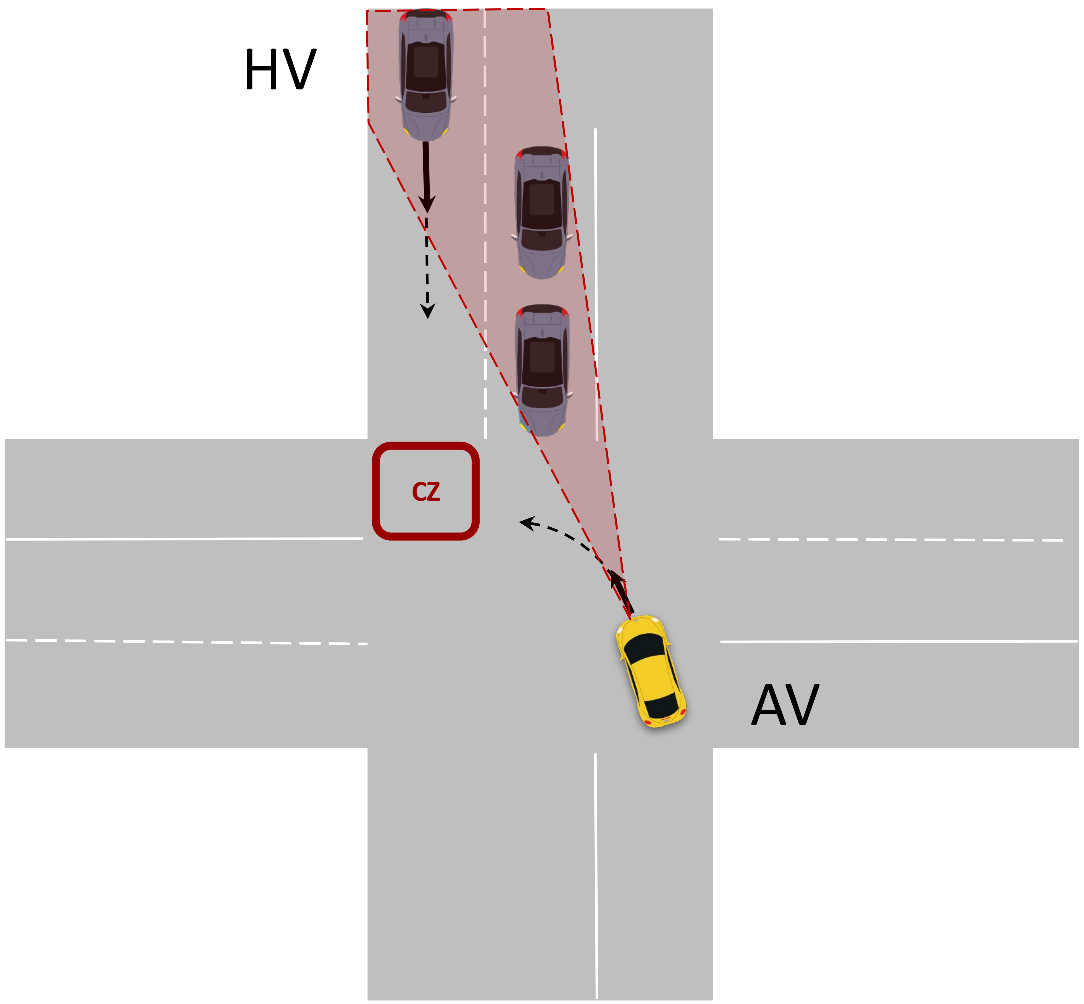


- Model for HVs – Poisson arrivals with rate λ at 25 mph
- AV does not know λ but can estimate by observing HV flow

AV's Decision:

| High λ | Low λ |
|---|---|
| Do not Turn | Turn |

# Estimation from Partial Observations

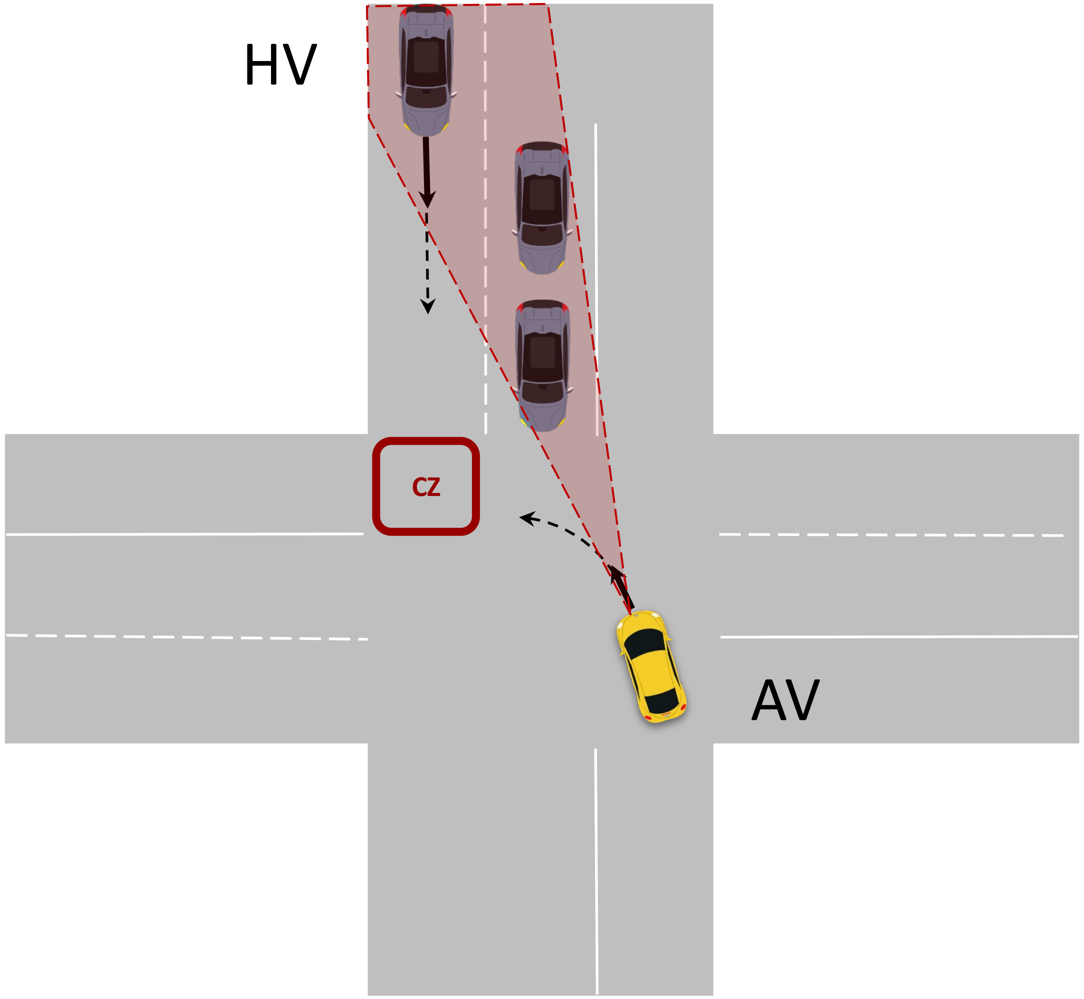


- Formulate this as a hypothesis testing problem for the AV

H0: High λ — Do not Turn

H1: Low λ — Turn

How long does the AV need to wait to be confident enough to turn?

# Estimation from Partial Observations

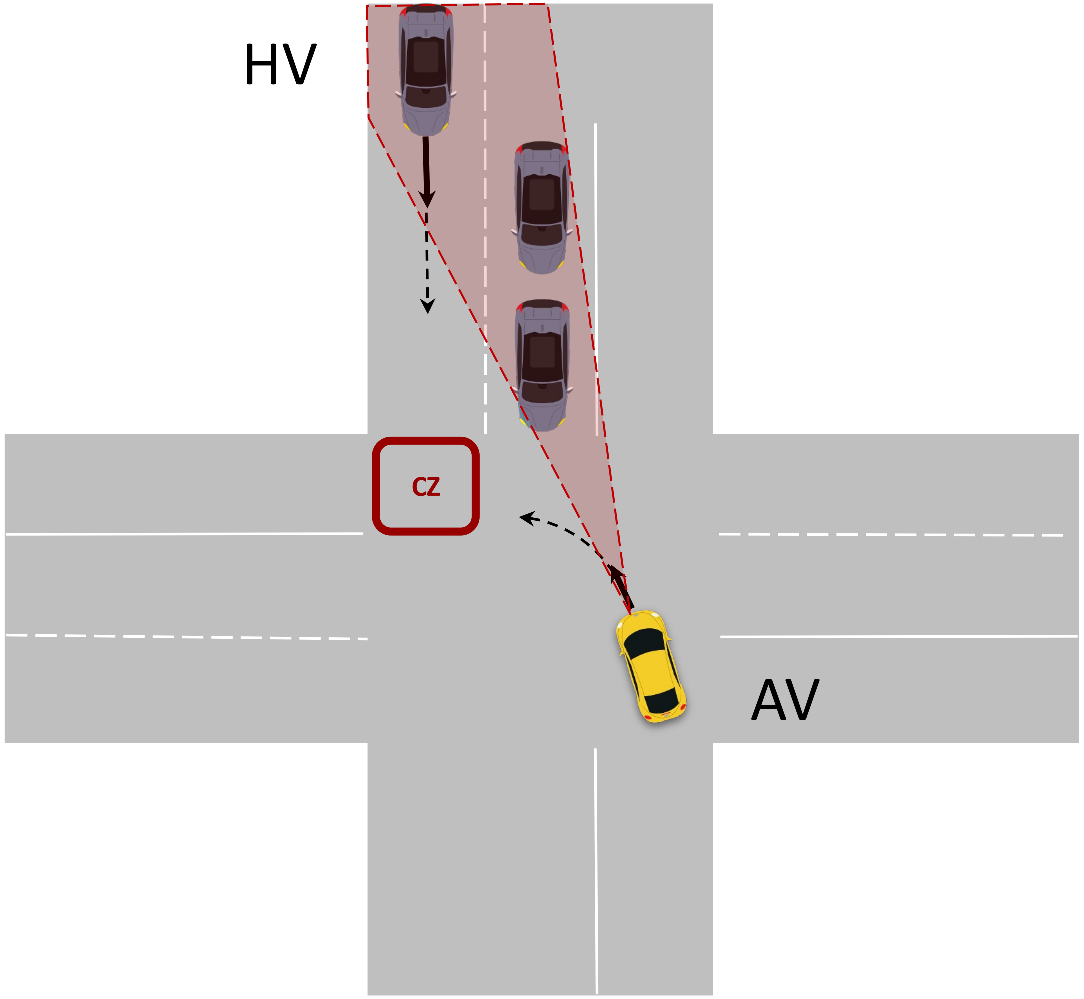


- Formulate this as a hypothesis testing problem for the AV

H0: High λ — Do not Turn

H1: Low λ — Turn

How long does the AV need to wait to be confident enough to turn?

AV would have to wait >7 min!

# Common Information Gap Scenarios

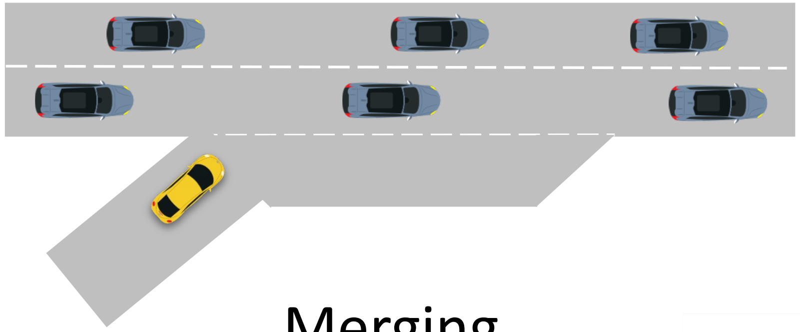

Merging

(Behavior)

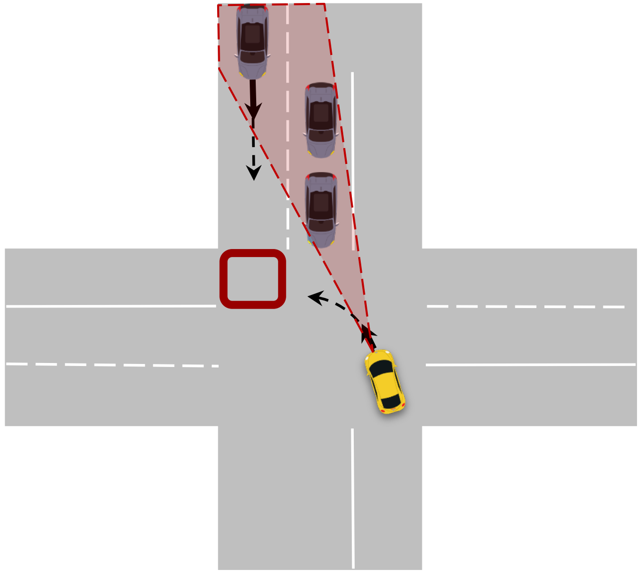

Occluded Turns (State)

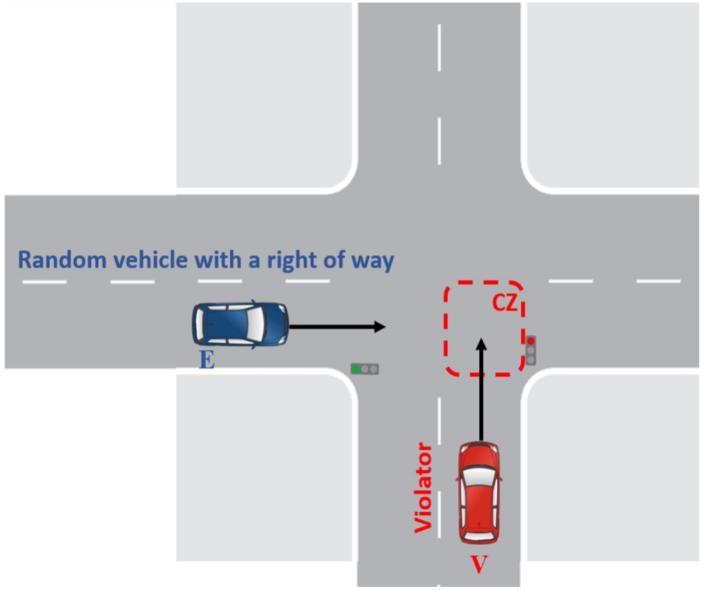


Traffic Violations (State + Behavior)

# Connectivity is Crucial

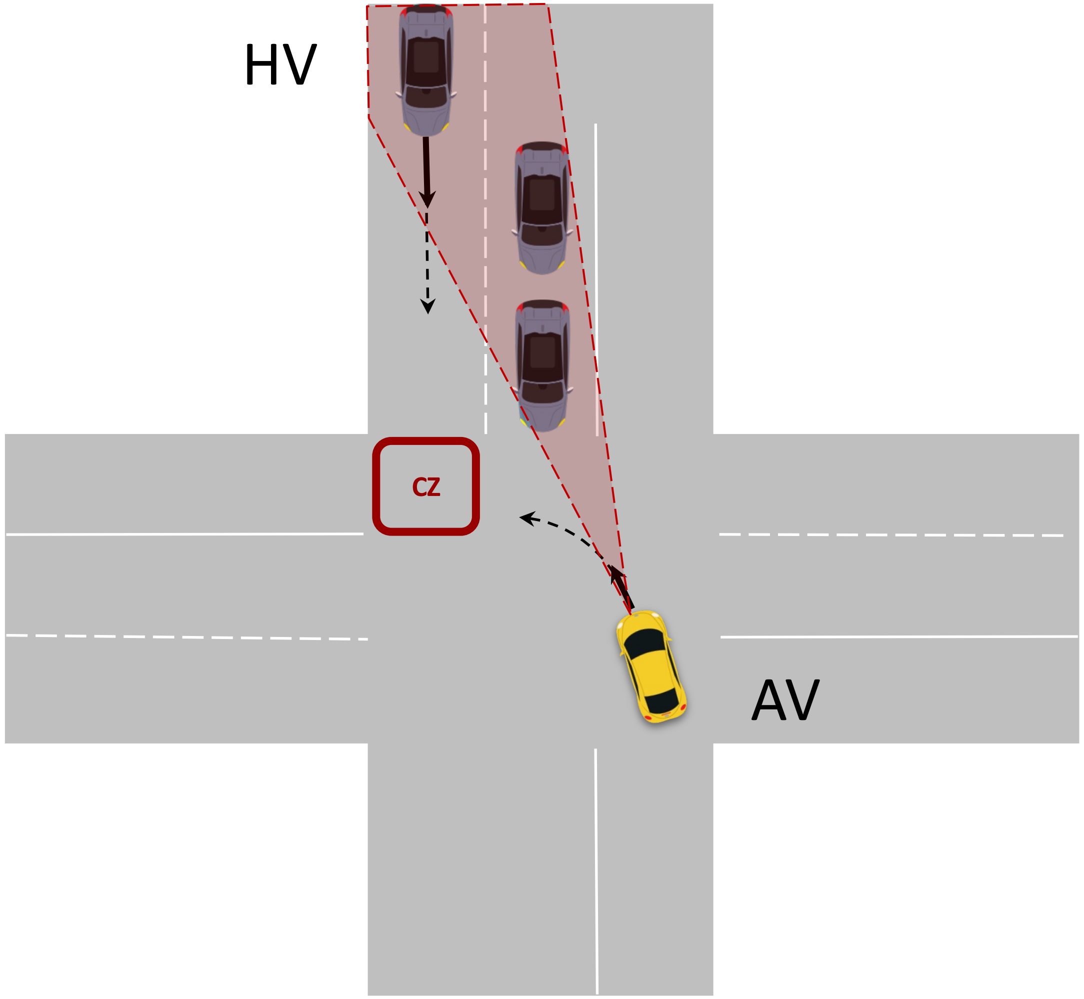


- No information gap if AV uses I2V connectivity
- A sensor placed 50 m away from intersection could alert the AV about a potential collision
  - This technology is already available!

Connectivity is a prerequisite for a safe AV future.

# Risk Assessment of Automated Vehicles

# AV Testing Approaches

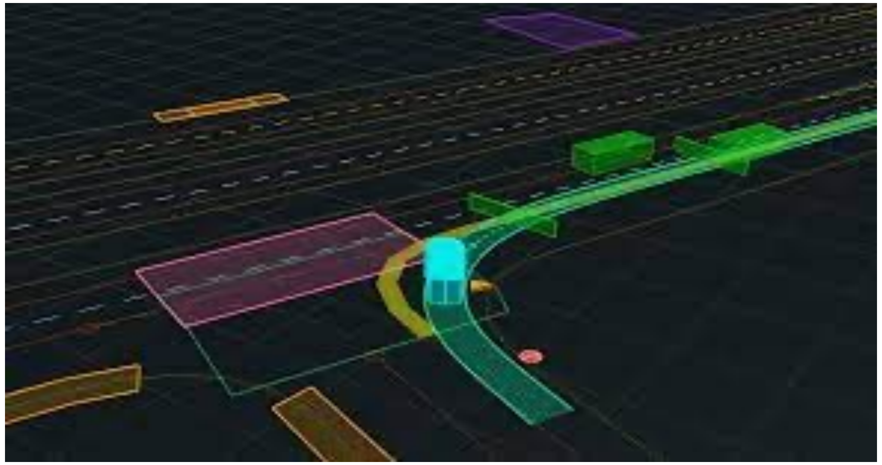

- Testing via Simulation
  - Simulator is only an approximation
- Closed Course Testing
  - Covers only a small subset of all real-world scenarios
- On-road Testing
  - Gold standard
  - Millions of miles covered by major AV companies
  - Compare human and AV crashes per mile?

Human crash rate: 1 every 500K miles



Images: Business Insider, Waymo

# Metrics can be Misleading

Human crash rate: 1 every 500K miles

Maintaining lane crash rate: 1 every 1.6M “maintaining lane” miles

Left turn crash rate: 1 every 91K “left turn” miles

- Crashes per mile doesn’t take into account maneuvers, locations, road conditions, vehicle/pedestrian behavior
- Varies significantly based on maneuver:
  - Left turn crashes 170 times as likely as crashes while maintaining lane (per mile)
- AVs have only been tested in limited domains – performance may not generalize to all contexts
- Can create a false sense of security

# Not all miles driven are equal

Risk assessment must take into account the context in which miles were driven



Images: Dreamstime, UCLA Health

# Not all miles driven are equal

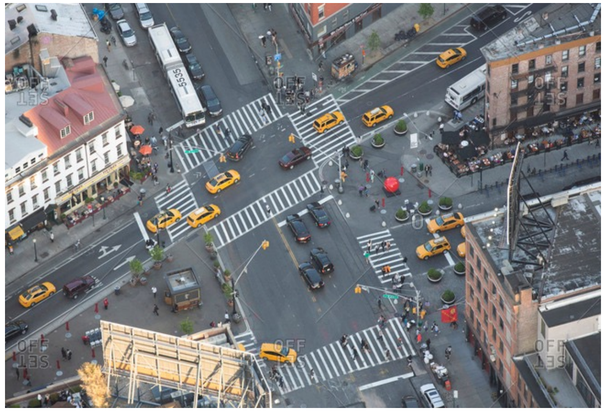


How should we account for diversity of driving contexts?



Images: AMA Insider, CarTrade, LA Magazine, Shutterstock, BBC

# Driving Context


Left turn on Hawthorne Blvd and Sepulveda Blvd

*Driving context*: description of driving scenario based on factors that influence crash risk

- Eg., maneuver, location, traffic density, time of day, road user behavior

Can be described at various levels:

- No context: *AV covers 100 ft*
- Maneuver level: *AV makes a left turn*
- (Maneuver, location, time of day) level: A northbound *AV on Hawthorne Blvd makes a left turn on to westbound Sepulveda Blvd during morning peak hour traffic*

# Driving Context


Left turn on Hawthorne Blvd and Sepulveda Blvd

- The more descriptive the driving context:
  - more specific risk assessment
  - lesser data available for assessment
- Suitable level of abstraction required to capture driving diversity in risk assessment

# Driving Context


Left turn on Hawthorne Blvd and Sepulveda Blvd

- The more descriptive the driving context:
  - more specific risk assessment
  - lesser data available for assessment
- Suitable level of abstraction required to capture driving diversity in risk assessment

No Context → Crashes per mile

What safety metrics follow from considering driving context?

# Maneuver Level Crash Analysis

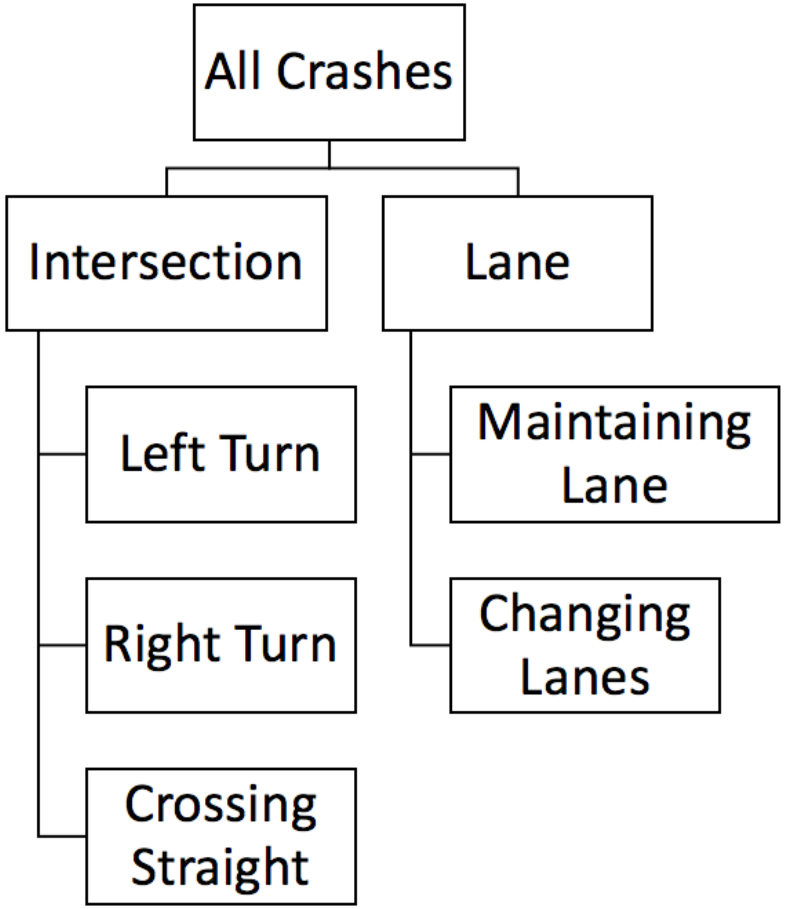

# A Simple Maneuver-level Crash Risk Model

$p_m$ : AV's crash probability during maneuver type m

$$\underbrace{C_m}_{\text{Total crashes involving maneuver}} \sim \mathrm{Bin}(\underbrace{N_m}_{\text{Frequency of maneuver}}, p_m)$$

MLE estimate of crash probability during maneuver

$$\hat{p}_m = \frac{C_m}{N_m}$$

# How significant is the effect of driving maneuver on crash risk?

# Leveraging Human Driving Data

MLE estimate of crash probability during maneuver

$$\hat{p}_m = \frac{C_m}{N_m}$$

Police Reports

Connected Vehicle Data

- Do not have AV testing data
- Humans drive billions of miles every year across diverse driving contexts
- Data available:
  - Police reported crashes
  - Vehicle maneuver frequency (eg., left/right turns, lane changes)
- Provides a useful risk prior for regions in which AVs haven’t been tested

# Torrance Case Study

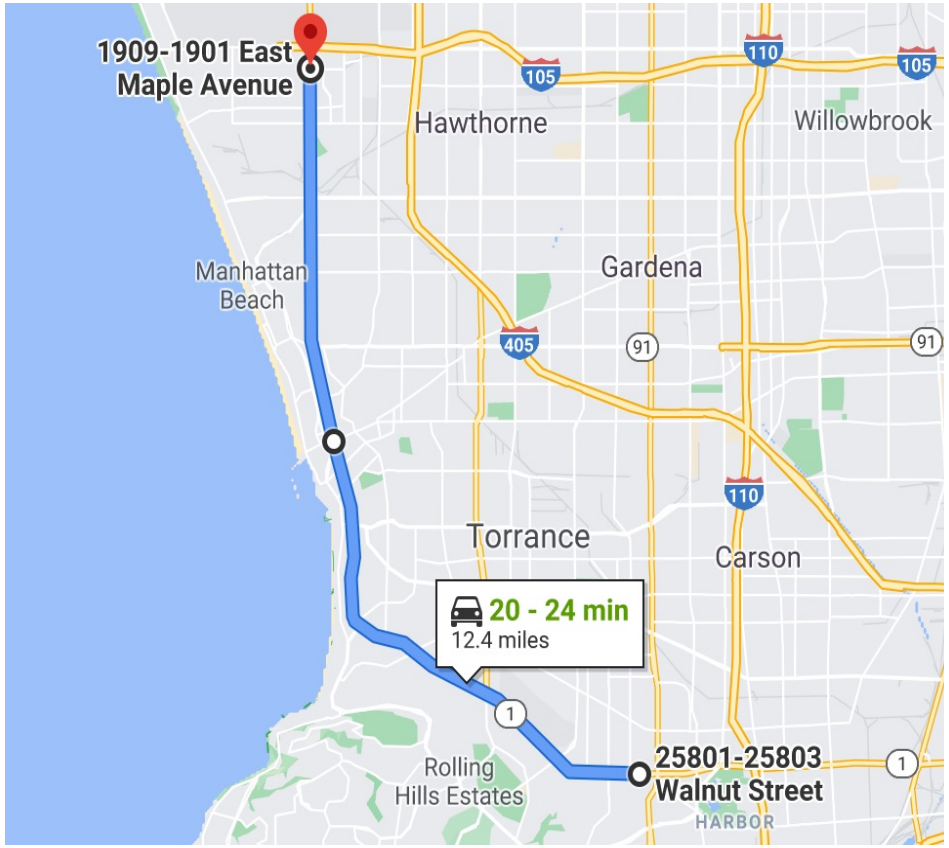


- 12.4 mile stretch of Pacific Coast Highway in Torrance, California
  - Data available for 29 signalized intersections along this stretch
- Maneuver level frequency data from Wejo Ltd. and Sensys Networks
- Crash data from police reports
- Can compute crash risk estimates for the average human driver

Images: Google Maps

# Maneuver Level Crash Risk

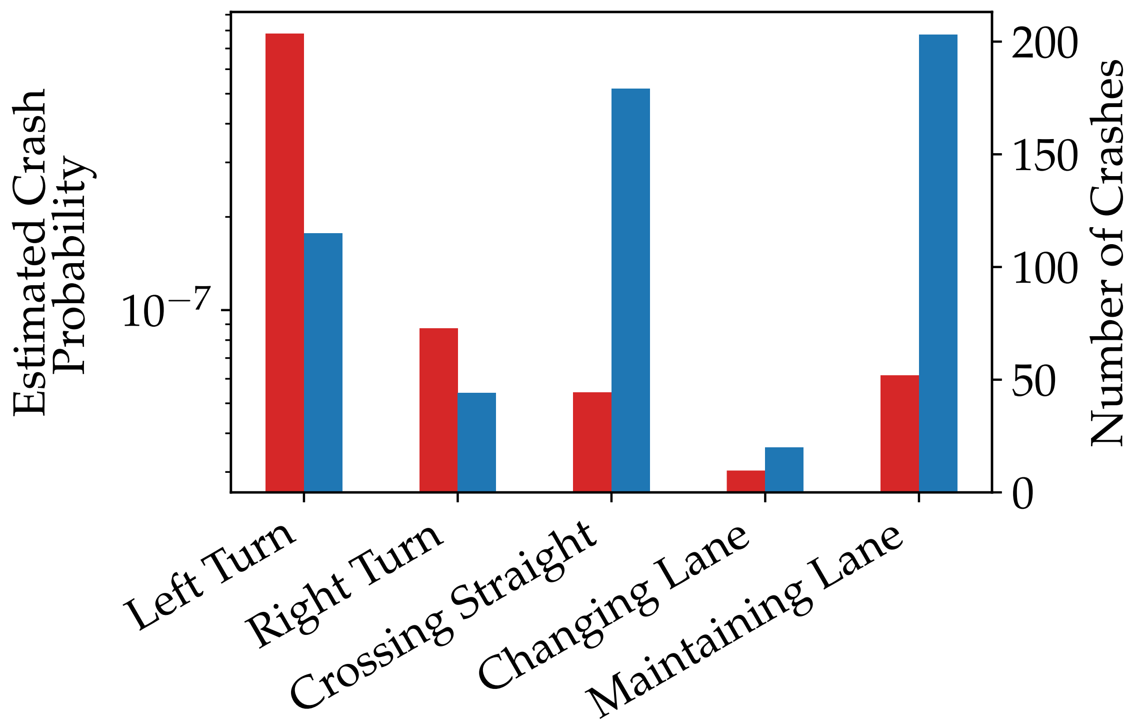


- Maneuvers involved in highest number of crashes (maintaining lane) do not have the highest crash probability (left turn)
- Significant variation in estimated crash probability across maneuvers

# Effect of Location

## Diamond Street

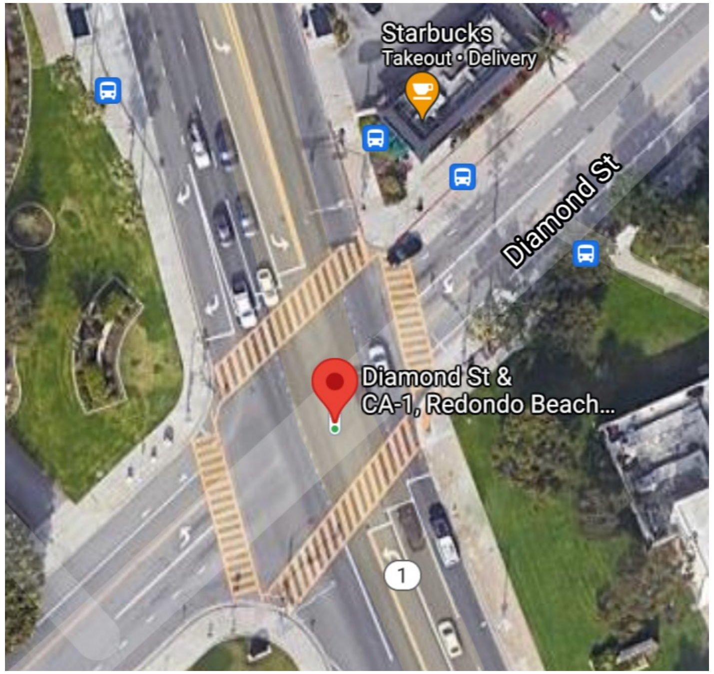


$$\hat{p}_{\mathrm{lt}}^{\mathrm{D}} = 6.7 \times 10^{-6}$$

## Palos Verdes Boulevard

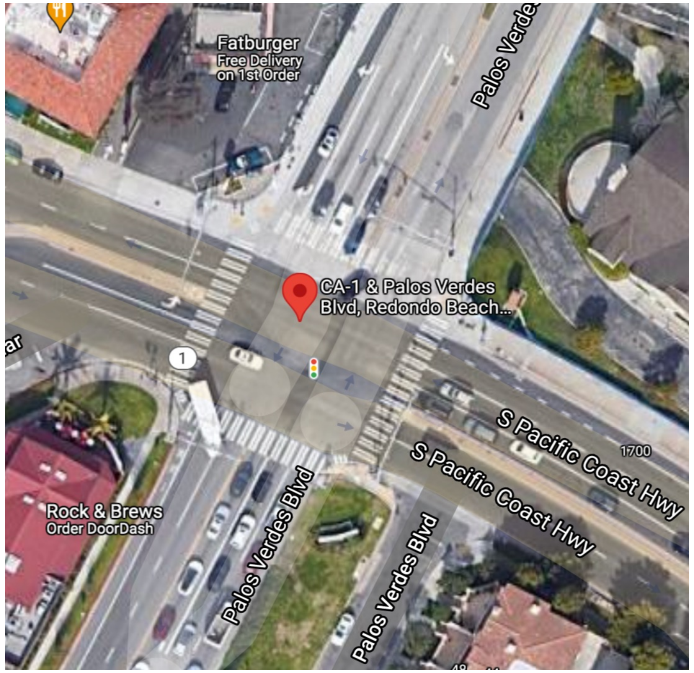


$$\hat{p}_{\mathrm{lt}}^{\mathrm{P}} = 9.3 \times 10^{-8}$$

70 times larger!

Crash risk varies considerably with location



Images: Google Maps

# San Francisco Crash Risk

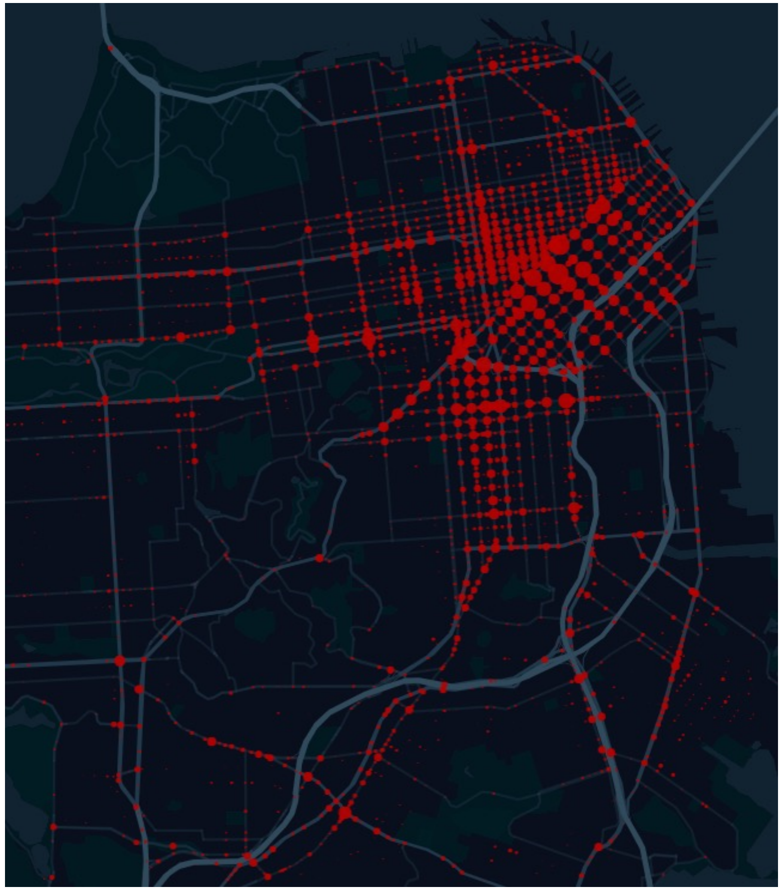

All Crashes

Left Turn Crash Risk

# San Francisco Crash Risk

Left Turn Crash Risk

Right Turn Crash Risk

# Route Risk

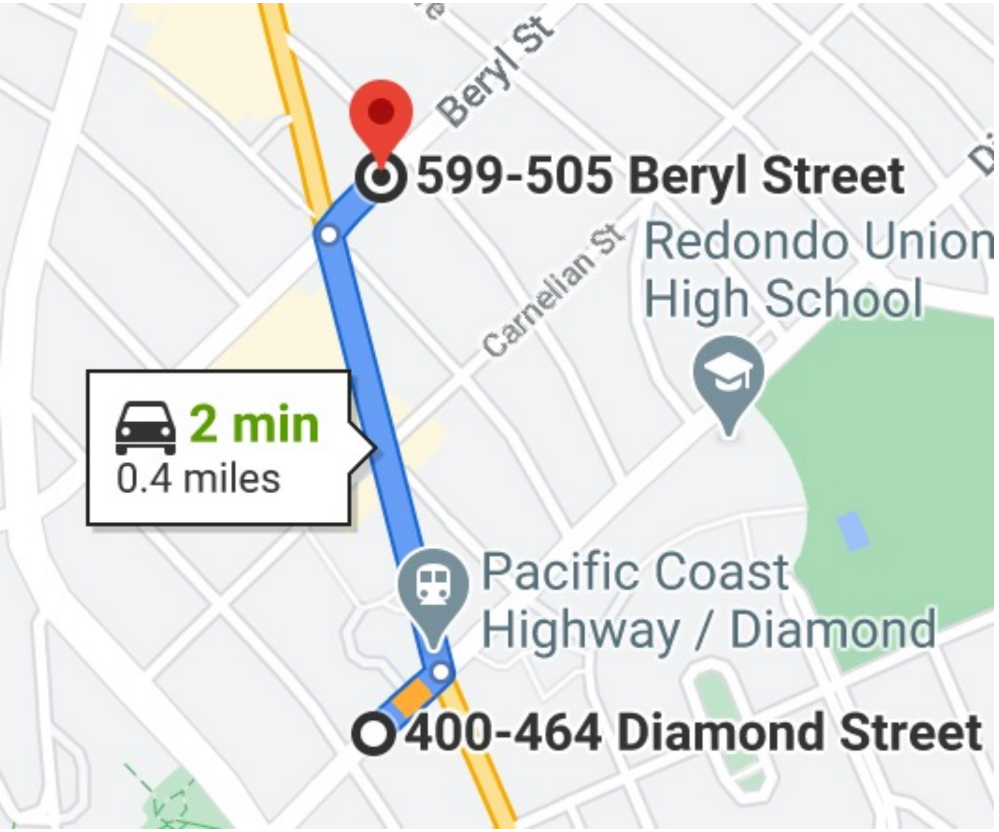




Images: Google Maps

# Route Risk

A route can be viewed as a sequence of (maneuver, location) pairs

Location in
$t^{th}$ step

$$R = \{r_t\}_{t=0}^{L} \qquad\qquad r_t = (m_t, i_t)$$

Maneuver in
$t^{th}$ step

$$P(\text{Crash along route } R) = 1 - \Pi_{t=0}^{L}(1 - P_t),$$

$$\approx \sum_t \hat{p}_{m_t}^{i_t}$$

Sum of crash probabilities
over all steps

# Route Risk

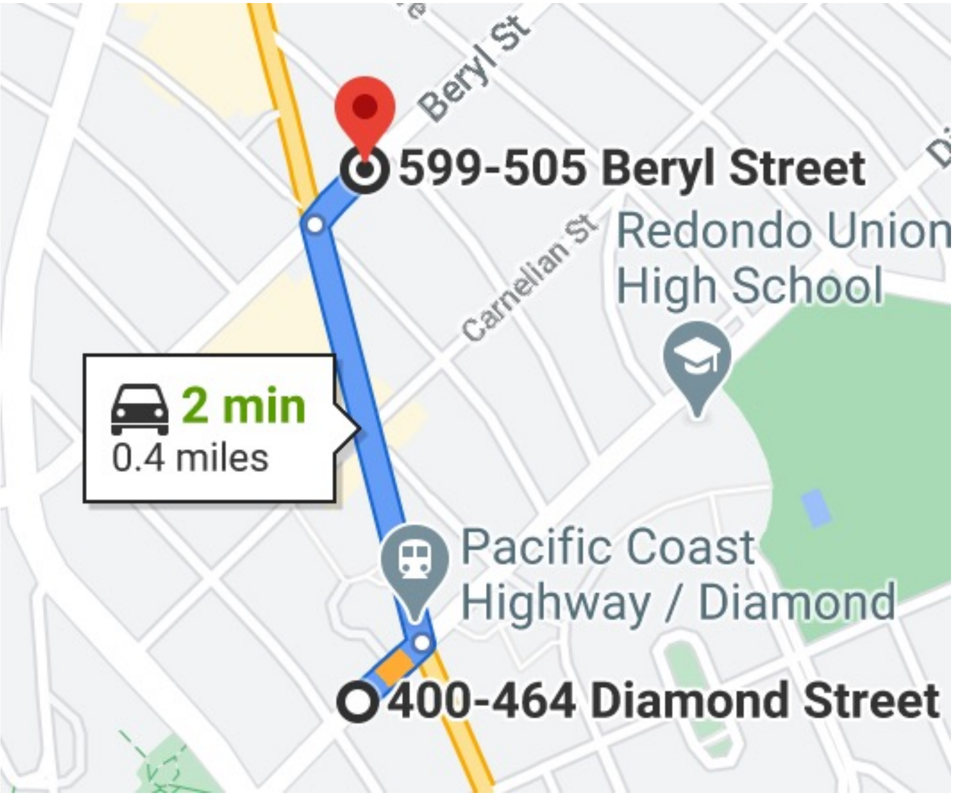


Diamond St to Beryl St on Pacific Coast Highway

Route:

1. Left turn from Diamond St. onto PCH
2. Staying in lane between Diamond St. and Carnelian St
3. Going straight through the Carnelian St. intersection
4. Lane change leading up to Beryl St. intersection
5. Right turn from PCH onto Beryl St.

Probability of crash during route:
$6.8 \times 10^{-6}$ (Left turn crash probability at Diamond St = $6.7 \times 10^{-6}$)

Images: Google Maps

# Economic Cost

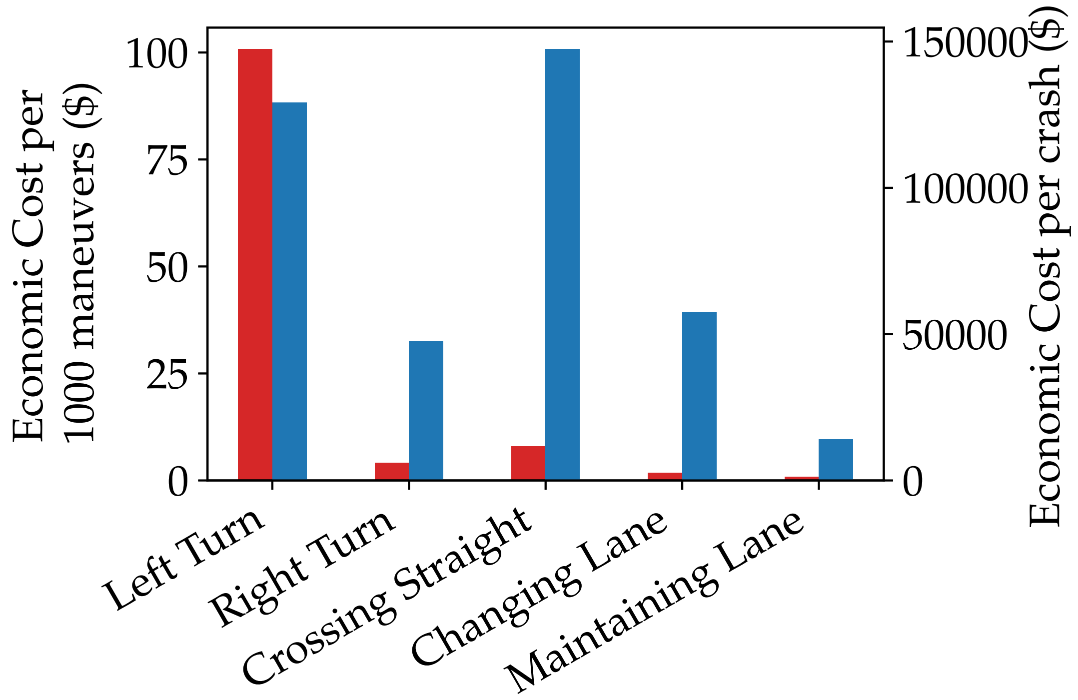


Left turn crashes have the highest economic cost per maneuver

- Economic cost of maneuver type m

Economic cost per crash

$$\hat{E}_m = \hat{p}_m e_m$$

Crash probability

- Compare left and right turns:
  - $100.8 per 1000 left turns vs $4.2 per 1000 right turns
  - Explains why UPS routes do not involve left turns

# Policy Implications

- Investment in Connected Infrastructure
  - Can be expensive -> Public-private partnerships
  - Cybersecurity research
- Safety metrics must account for diversity
  - Disengagements per mile: perverse incentives
- More data disclosure necessary
  - Types of maneuvers, environments
  - How does performance generalize?
- Driving test for AVs?
- Phased deployment from “easy” to “difficult” ODDs

# References


- A. S., M. Yu, A. Kurzhanskiy, O. Grembek, H. Tavafoghi, P. Varaiya, ``Safety Challenges for Autonomous Vehicles in the Absence of Connectivity'', Transportation Research: Part C, 2021
- A. S., H. Tavafoghi, A. Kurzhanskiy, K. Poolla, P. Varaiya, ``Risk Assessment of Autonomous Vehicles across Diverse Driving Contexts'', IEEE ITSC 2021
- A. S., H. Tavafoghi, A. Kurzhanskiy, K. Poolla, P. Varaiya, ``Automated Vehicle Safety and Deployment: Lessons from Human Crashes'', IEEE ICCVE 2022

# Thank you!

Email: shetty.akhil@berkeley.edu