



Prof. Philip Koopman

**Carnegie
Mellon
University**

Autonomous Ground Vehicle Safety: Lessons To Learn

December 2023

www.Koopman.us

PHILIP KOOPMAN

**HOW SAFE IS
SAFE ENOUGH?**

Measuring and Predicting
Autonomous Vehicle Safety



- 2023 was a pivotal year for robotaxis
 - Summer of hype
 - Winter of discontent
 - Will we see a glorious summer 2024?
- Top lessons for AV success
 - More robust criteria to decide whether safe enough/good enough to deploy
 - Create more sustainable trust



“AV” = “Autonomous Vehicle” (no human driver needed)

Why Is AV Safety Complicated?

■ Public expectations


- Expect super-human machine performance
- Trust too easily given, backlash when broken

■ Technical challenges

- Machine Learning safety is work in progress
- Statistical approach vs. high severity rare events

■ Still working on resolving a culture clash

- Autonomy researchers: it's all about the cool small-scale demo
- Silicon Valley: move fast + break things
- Automotive: blame driver for not mitigating equipment failures
- Regulators: test-centric; weak digital safety expertise



Robotaxis: "Safety Is Our #1 Priority"



Because
Safety is
Urgent™

Autonomous Driving
Technology Can Save
Lives and Improve
Mobility

<https://waymo.com/safety/>

cruise

Safety first,
always

<https://getcruise.com/safety/>



Safety Drives Us

Motional is developing safe
autonomous vehicles.

<https://motional.com/safety-philosophy>

ZOOX

A new bar for safety

Safety isn't just part of what we do. It's why we're here.

<https://zoox.com/safety/>

Getting Past the AV Safety Rhetoric


- Nobody knows when/if Autonomous Vehicles (AVs) will be safer than human drivers
 - Current studies promising, but limited
 - Reduced fatality rates are aspirational
- Some humans drive drunk
 - On average they are still good and adaptable
- But computers lack common sense
 - Machine Learning is brittle to novelty
- Computer drivers can be imperfect
 - Net balance for computer drivers uncertain
 - **Safety must be engineered, not assumed**



Early 2023: Driverless 1 Million Miles

Waymo And Cruise have Both Hit 1M Miles With No Driver, But Waymo Publishes Detailed Safety Data

Brad Templeton · Feb 28, 2023, 12:00pm EST



<https://bit.ly/46G07Gg>

Forbes

In January 2023, Waymo reached **1 million rider-only miles**



No reported injuries



Only 2 collisions that met the criteria for inclusion in NHTSA's CISS



18 minor contact events



55% of all events were the result of a human driver hitting a stationary Waymo vehicle



Human drivers violated road rules and/or behaved dangerously in every vehicle-to-vehicle event



10% of all events happened at night



No intersection-related events



No events involving vulnerable road users


Waymo: Feb. 2023.
<https://bit.ly/3N5F6xF>

Waymo passenger injury August 2, 2023:
<https://bit.ly/47Z9pyb>

Both emphasize "at fault" crashes

Updated Human Ridehail Benchmark vs Cruise AVs in 1M

Collision Counts in San Francisco



Cruise: Sept. 2023.
<https://bit.ly/47W1DVR>

Are Robotaxis Safer?

- Nobody knows when Autonomous Vehicles (AVs) will be safer than human drivers
 - Press releases overstate company study results
 - Study limits: at-fault crashes; ride-hail driver baseline
- Fatality safety rates take 300+ Million miles
 - Currently about 5 million miles/company in S.F.
 - Current studies are predictions and extrapolations
 - Reduced fatality rates are still aspirational
- Declaring safety “victory” at this point is like claiming a medal ... after the first mile in a marathon



Summer of Robotaxi Hype: July 2023

Humans are terrible drivers

42,795 Americans were killed in car crashes last year

You might be a good driver, but many of us aren't. People cause millions of accidents every year in the US. Cruise driverless cars are designed to save lives. Our cars were involved in 92% fewer collisions as the primary contributor* They also never drive distracted, drowsy or drunk.

<https://twitter.com/kvogt/status/1679517290847694848>



Kyle Vogt
@kvogt

Follow


We ran this full-page ad in @nytimes and several local papers today.

Human drivers aren't good enough. America can do better, and it is time we fully embrace AVs.



Last edited 11:45 AM · Jul 13, 2023 · 956K Views


“Better than Human” Narrative Erodes



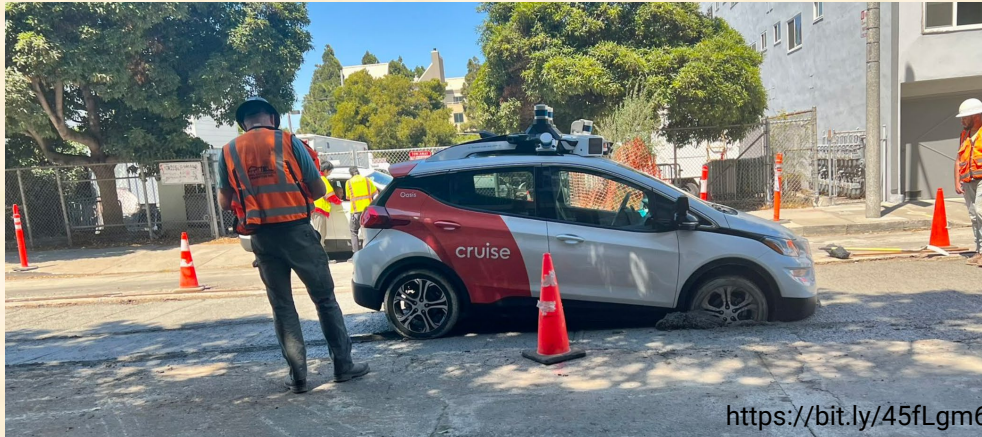
August 2023

March 2023

Are software bugs the new drunk driving?



August 2023




City of San Francisco Concerns

AV driving that interferes with emergency response



Emergency Response (SFFD) Impact Incidents by Type (Jan 1 – Sept 27, 2023)



Summer Into Fall

- July “Coning”:
 - Putting a traffic cone on the hood of a robotaxi as a form of protest
- August 10:
 - California PUC approves expanded robotaxi services


CityLab
Perspective

Bloomberg

CityLab
Perspective

<https://bit.ly/46G606r>

Why Robotaxis Can Make Cities Safer



An urban policy analyst at the Manhattan Institute argues in favor of expanding autonomous ride-hailing in San Francisco and beyond.

By Jordan McGillis September 20, 2023

Why Robotaxis Have No Place in the City



Two San Francisco activists explain why they're trying to halt the city's autonomous ride-hailing vehicles in their tracks.

September 20, 2023

By Aditya Bhumbra and Mingwei Samuel

2023 Winter of Discontent

Los Angeles Times

General Motors recalls all Cruise robotaxis after one dragged a pedestrian



General Motors is updating the software of its Cruise robotaxi vehicles after one struck and dragged a pedestrian in San Francisco last month, according to documents posted by safety regulators Wednesday. (Paul Sancya / Associated Press)

- October 2, 2023 crash
 - Human-driven vehicle hits pedestrian
 - Cruise robotaxi dragged pedestrian *after* initial emergency stop was completed
 - Allegations of Cruise attempting cover-up
 - Oct. 24, CA DMV suspends Cruise permits

What Happens Next?

- Cruise will resolve its situation one way or another
 - Spending cuts, single city operations, slower scale-up
 - Other outcomes are possible
 - Success will require hard pivot into safety
- Other companies will continue to develop
 - Heavy trucks, local deliveries
 - Low speed shuttles
 - Level 2+++ (almost fully self-driving)
- The following lessons apply to all the companies



Lessons To Learn

#1: Net “better than human driver” is just a start

- Difficult to prove up front
- Mishaps undermine credibility of claim



■ Also need to also address:

#2: Negligent driving

#3: Risk transfer

#4: Recall issues

#5: Safety standards

#6: Ethical & equity concerns

#7: Sustainable trust

Aug. 2023:
Injury crash with fire truck.
Reduced pace of Cruise testing due to CA DMV request to cut fleet operations in half.

#2: Avoid Negligent Driving

- **“Negligent” robotaxi driving involves:**
 - Establishing a duty of care to other road users
 - Was a loss event caused by breach of duty of care?
 - Would a human driver have been negligent?
 - Statistical safety arguments are irrelevant here
 - “Safe” drivers don’t get free passes to run red lights
- **October 2 Cruise pedestrian mishap:**
 - Robotaxi arguably should have increased caution
 - Pedestrian in crosswalk (even if in adjacent lane)
 - Other car hitting pedestrian was readily predictable
 - Robotaxi should not have moved with pedestrian under vehicle



#3: Avoid Risk Transfer

- What if children at greater risk?
 - Or disabled pedestrians?
 - Or bicyclists? Etc.
- Caution – this particular article is controversial
 - Regardless, this is an important safety constraint
- Avoid increasing any group's risk
 - Extra effort decreasing risk to vulnerable groups



#4: Fine-Grain Risk & Regulators

- Want to avoid regulatory recalls
 - “Undue Risk” in the small – specific issues
 - Informed by test-centric standards
- Recalls are specific, not net risk
 - Rolling through stop signs
 - Phantom braking
 - Malfunctioning display console
 - Software safety & net risk are mostly beyond regulatory scope
- Regulators struggling to predict safety outcomes in advance
 - 2020 Proposal to require industry safety standards is inactive



Part 573 Safety Recall Report

#5: Standards Set Expectation of Safety

SYSTEM SAFETY	ANSI/UL 4600		Safety Beyond Dynamic Driving	HIGHLY AUTOMATED VEHICLE SAFETY CASE ANSI/UL 4600
DYNAMIC DRIVING FUNCTION	ISO 21448	SaFAD/ISO TR 4804	Environment & Edge Cases	
FUNCTIONAL SAFETY	ISO 26262		Equipment Faults	
CYBER-SECURITY	SAE J3061	SAE 21434	Computer Security	
VEHICLE SAFETY	FMVSS	NCAP	Basic Vehicle Functions	

REQUIRED

ROAD TESTING SAFETY
SAE J3018

#6: Ethical & Equity Concerns

- Ride Hail made promises ... with disappointing results
 - Why will this turn out any differently?
- Equity concerns:
 - Labor issues (e.g., displaced ride-hail/taxi drivers)
 - Will disabled community access really happen?
 - Cheap taxis undermine public transit
 - Municipal preemption / no local control of issues
- Ethical & related concerns
 - Testing risk imposed upon vulnerable people
 - Long-term aspirational safety incurs real short-term risk
 - No required independent safety technical oversight



[Dall-e]

#7: Sustainable Trust

- Trust-degrading rhetoric:
 - “Robotaxis won’t make stupid driving mistakes”
 - Relentless blame of human drivers
- Trust-degrading actions:
 - Lobbying for municipal preemption
 - Redacting & withholding information
- Toward increasing trust:
 - Talking with (not “at”) stakeholders
 - More transparency on incidents & corrective actions
 - Accepting proportional responsibility for loss events
 - Stating release criteria in advance & tracking metrics




Ford VSSA 2021 <https://bit.ly/3njionT>

Lessons for AV Industry Success

1. Net “better than human driver” only a starting point
2. Avoid negligent driving behavior
3. Avoid risk transfer to vulnerable populations
4. Fine-grain regulatory risk management
5. Conform to industry safety standards
6. Address ethical & equity concerns
7. Build sustainable trust

More talks here:

<https://users.ece.cmu.edu/~koopman/>



[General Motors]