Engineering @ Cruise

cruise

Agenda

- Introductions: Why Self-Driving Cars?
- AV Vehicles
- Ground Truth
- Data Infrastructure
- Machine Learning Platform
- Product Engineering

Our Mission

We’re building the world’s most advanced self-driving vehicles to safely connect people with places, things, and experiences they care about.

Cruise Behaviors

Stay safe

Stay focused

Own it

Seek truth

Work together

Be humble

Why Self-Driving Cars?

- 1.35 million people die each year
- Road crashes are leading cause of death for ages 5-29
- Passenger vehicles emit 4.6 tons CO2/yr
- Americans spend about 293 hrs driving per year → 7 40-hr weeks!
- 42 hrs sitting in traffic :(

1. https://www.who.int/news-room/fact-sheets/detail/road-traffic-injuries
2. https://www.epa.gov/greenvehicles/greenhouse-gas-emissions-typical-passenger-vehicle
3. https://www.automotive-fleet.com/136735/americans-spend-an-average-of-17-600-minutes-driving-annually
4. https://bestlifeonline.com/traffic-commute-times/

Why Self-Driving Cars?

Less human error: greater safety

Lower carbon emissions

Save driving time!

Fewer cars for more people

Less need for parking

Cruise Vehicle History

About Me

Austin Harris

- Hardware reliability engineer
- I interned at Cruise!
- Been at Cruise ~7 months
- Graduated with Mech E degree from University of California, Davis in 2018
- Love cars, love bikes, love going fast

Audi S4

- Retrofitted 2x Audi S4's with RP-1

Nissan Leaf

- Retrofitted 5x Nissan Leafs working towards fully autonomous system

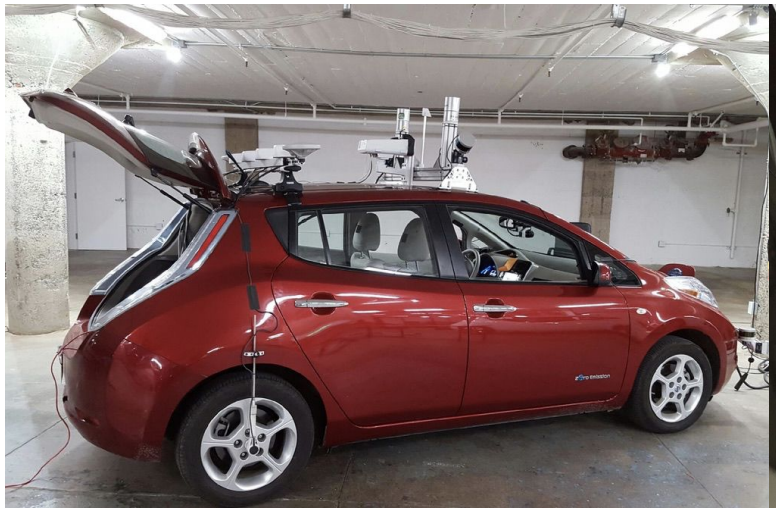

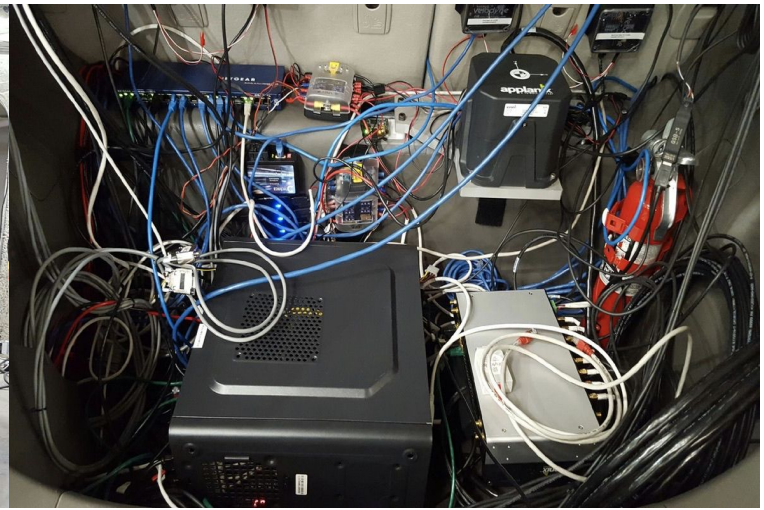

Track 1

- Retrofitted Chevy Bolts with fully autonomous system

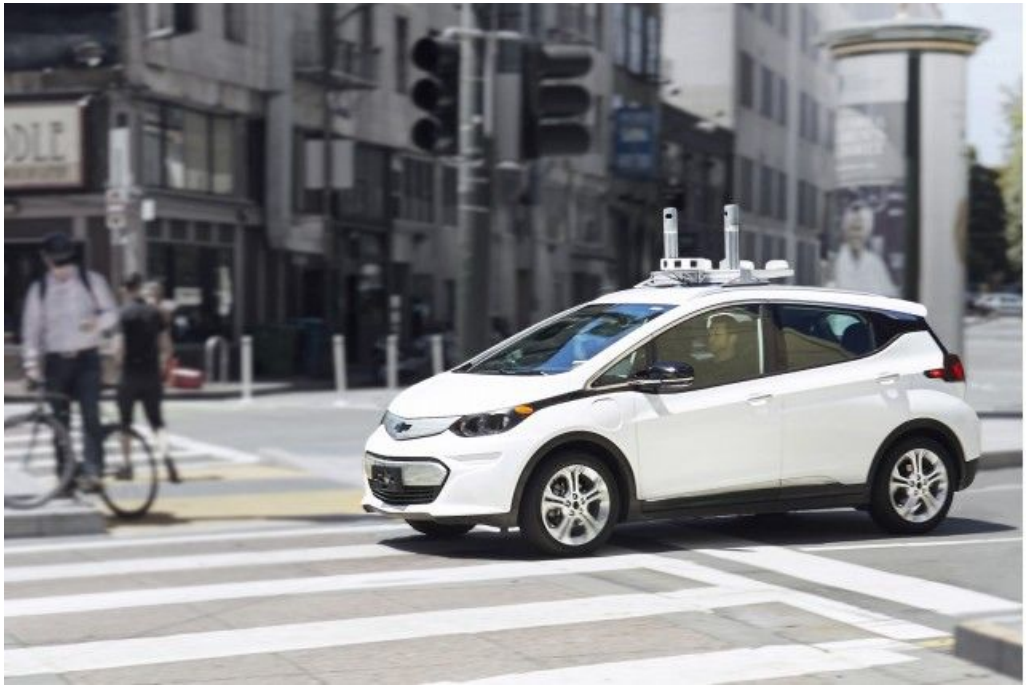

Track 2

- Retrofitted Bolts (with some changes) with fully autonomous system
  - Vehicles assembled by GM

Track 3

- **Custom designed Bolt with deeply integrated autonomous system and redundant systems**

T3 Systems Overview

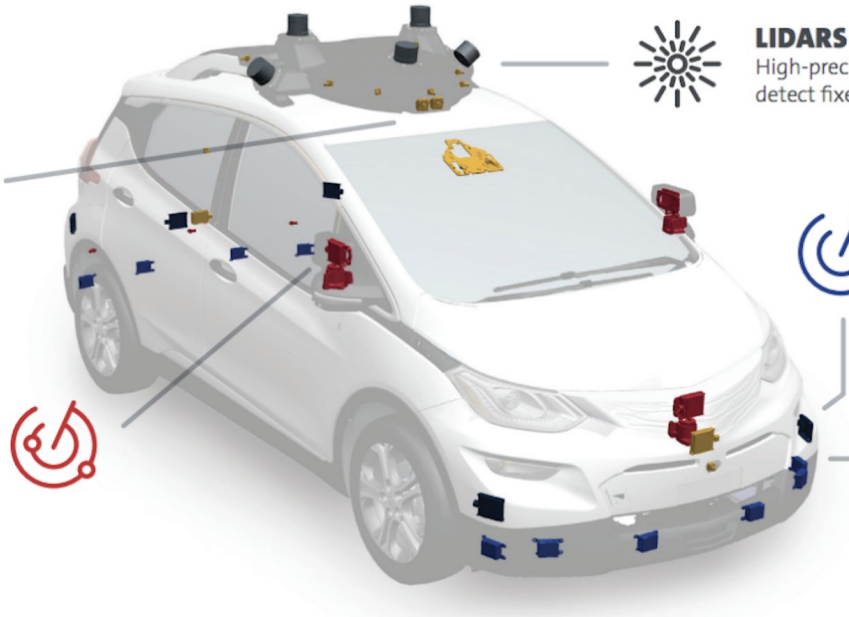

LIDARS
High-precision laser sensors that detect fixed and moving objects

CAMERAS
Detect and track pedestrians / cyclists, traffic lights, free space and other features

LONG-RANGE RADARS
Detect vehicles and measure velocity

ARTICULATING RADARS
Detect moving vehicles at long range over a wide field of view

SHORT-RANGE RADARS
Detect objects around the vehicle

What does the Vehicles team do?

Hardware Reliability Team

- Fleet availability
- Fleet reliability
- Drive product improvements

Examples

- Develop technician documentation
- Root cause hardware issues
- Develop emergency response processes
- Take findings from root cause analysis and drive product requirements/improvements to GM

Calibration & Development Team

- Develop & test sensor calibration scripts
- Develop & test hardware required for calibrations
- Vehicle bring-up development

Examples

- Improve reliability of lidar calibration script
- Test & improve calibration scenes for object cameras
- Define bring-up processes to hand off to GM

Fun Fact

How much power does our autonomous driving computer consume?

Ground Truth

Groups at Cruise (Engineering)

AV Platform	Building the nerves and muscles of the AV
Core Infrastructure	Deploy highly scalable systems
Data Infrastructure	Enable insights from our extensive data collection
Data Science	Use our diverse data to improve business and engineering decisions
Engineering Productivity	Build a world-class development environment at Cruise
Ground Truth	Enable our AVs to understand the world
Mapping	Enable our AVs to navigate the world
ML Platform	Unlock ML development efficiency to accelerate our development
Product Engineering	Launch the world's best AV ridesharing service
Remote Assistance	Ensure we are always able to get the AV unstuck

Groups at Cruise (Engineering)

AV Platform	Building the nerves and muscles of the AV
Core Infrastructure	Deploy highly scalable systems
Data Infrastructure	Enable insights from our extensive data collection
Data Science	Use our diverse data to improve business and engineering decisions
Engineering Productivity	Build a world-class development environment at Cruise
Ground Truth	**Enable our AVs to understand the world**
Mapping	Enable our AVs to navigate the world
ML Platform	Unlock ML development efficiency to accelerate our development
Product Engineering	Launch the world's best AV ridesharing service
Remote Assistance	Ensure we are always able to get the AV unstuck

About Me

Victoria Lo

- Full-Stack Software Engineer
- Ground Truth Engineering team
- Been at Cruise for < 1.5 years
- Graduated UC Berkeley in 2015
- San Francisco Bay Area native, but familiar with the Toronto area/GTA

- How do cars know what they are seeing?
- How do machine learning models know what the “right” answer is?
- How do tests/QA know how to evaluate the car code?
- How does data science know what decisions to make?
- How do we know how well we are driving?

Ground Truth provides labeled data.

What's Ground Truth?

- Literally, the truth of what's on the ground.
- Let's peek at what the cars see...
- Camera:

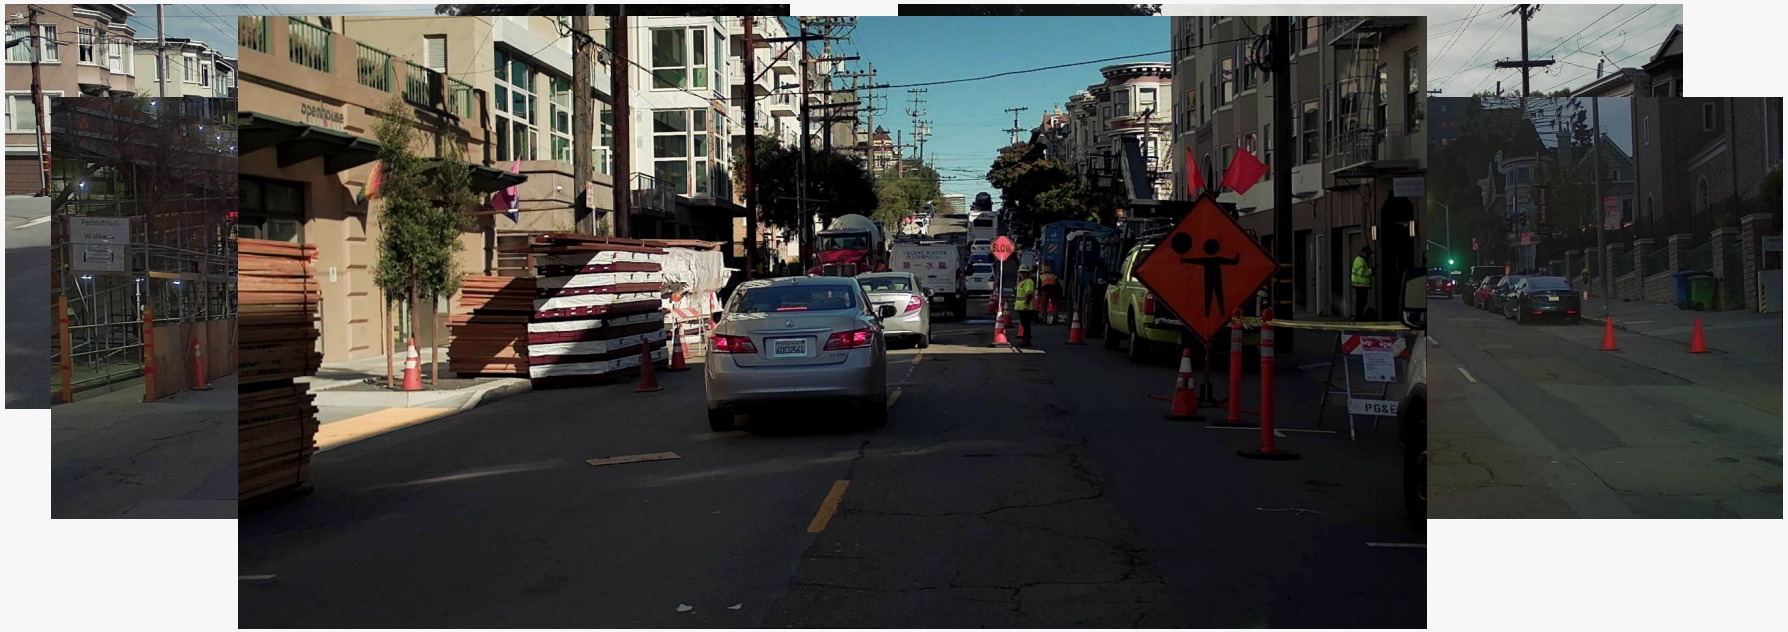

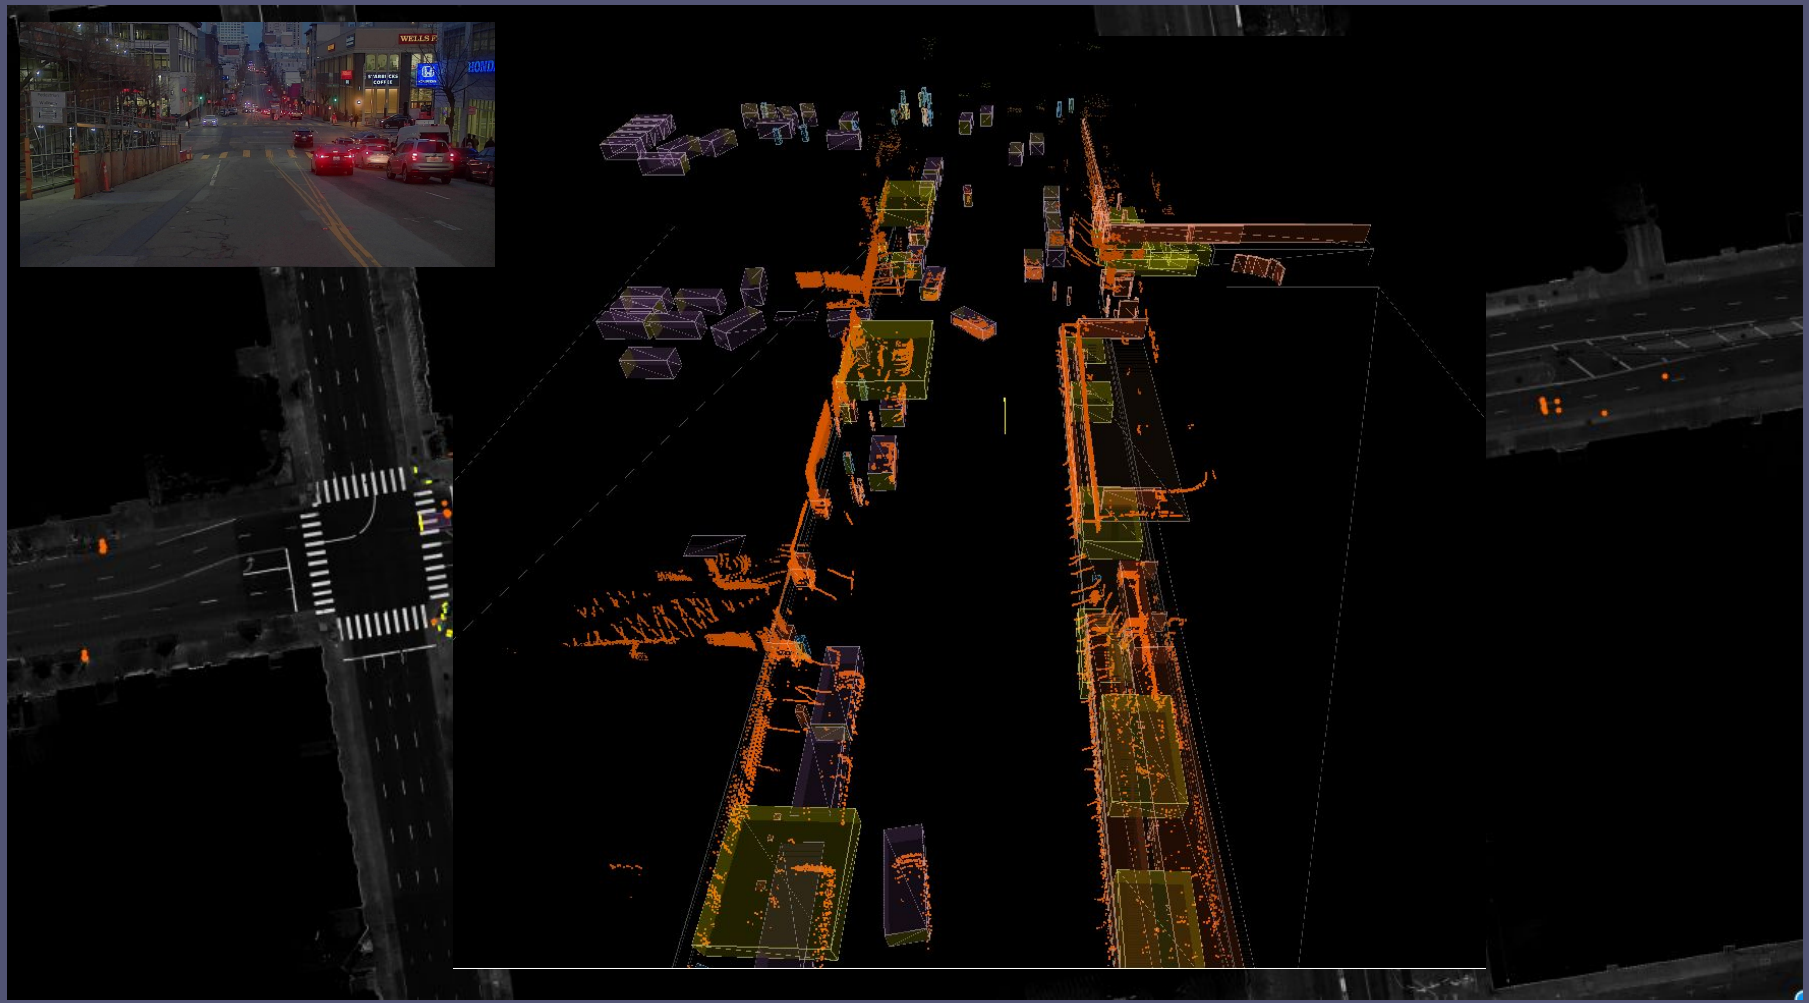
WELLS F
STARBUCKS
COFFEE
HONDA

Tools

Types of Labeling

- Clustering
  - Which points in this
    shape of this thing?
- Segmentation
  - Label all of this thing
  - What's the shape an
- Classification
  - Does this thing exist in the image? Or
    what is this thing?

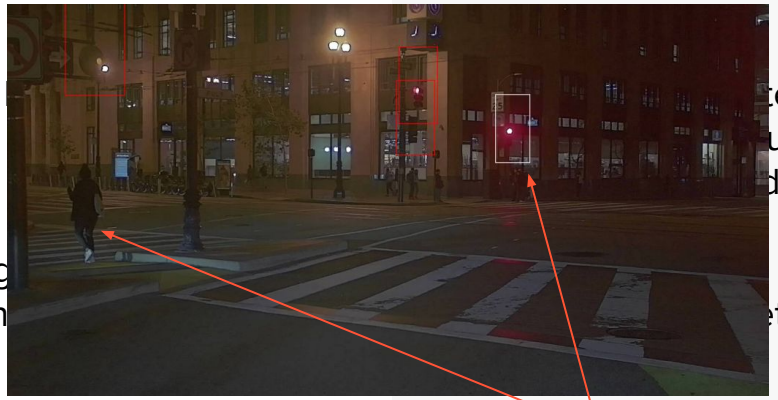

Sensors

- ection and Ranging
- uds!
- dy is a blob.

- etection and Ranging

- Camera
  - Photos!
  - Traffic lights, pedestrians, etc.
  - Everybody is a blur.

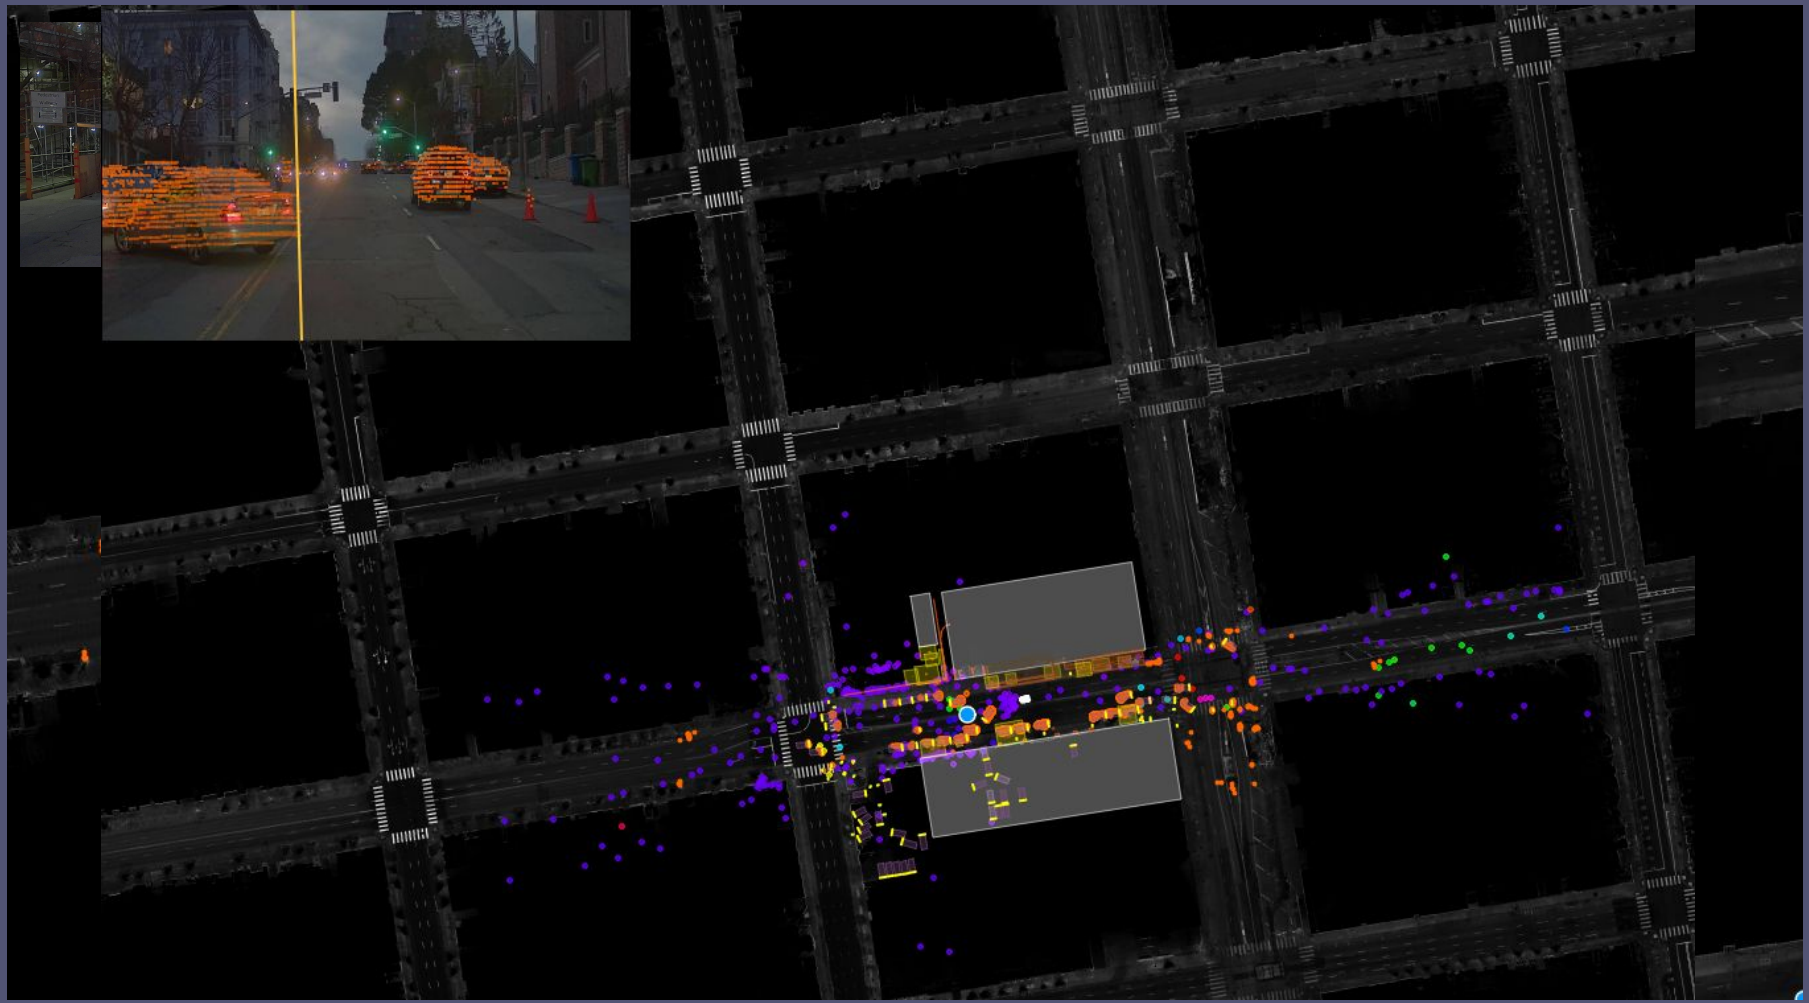

Teams/Pods in Ground Truth

Engineering

- Point Cloud (“PoClo”)
- Hollywood
- Data Foundation
- Platform
- REC (Ride Event Classification)
- Automation

Other

- Technical Program Manager (TPM)
- Product Manager - Technical (PMT)
- Operations
  - Point Cloud
  - Hollywood
  - Growth

Stack (Technologies)

Front-End

- React
- Redux
- Immutable.js
- Flow
- Three.js
- Openlayers/WebGL

Back-End

- Go
- Python
- Postgres & S3
- Docker
- AWS
- Google Cloud Platform - coming soon!

Why Cruise? Why Do I Love Cruise?

- Self-driving cars.
  - It's ***THE FUTURE***.
  - Ride free in a self-driving car and make your friends jealous!

- Delicious lunch and huge snack selection!



- Cool projects
  - YOUR work is going to contribute to self-driving cars!
  - We've had a Waterloo intern for the past few months on Data Foundation

- Cool, smart coworkers and awesome culture!
  - Memes and bubble tea ("boba") and Super Smash Brothers

- San Francisco is amazing
  - Pretty good weather: 45 - 85 °F (7 - 29 °C) in San Francisco year-round and no snow
  - Tons of food, hiking, Silicon Valley, Lake Tahoe, Napa

Questions?

Ask me anything afterwards!

Data @ Cruise

Groups at Cruise (Engineering)

AV Platform	Building the nerves and muscles of the AV
Core Infrastructure	Deploy highly scalable systems
Data Infrastructure	Enable insights from our extensive data collection
Data Science	Use our diverse data to improve business and engineering decisions
Engineering Productivity	Build a world-class development environment at Cruise
Ground Truth	Enable our AVs to understand the world
Mapping	Enable our AVs to navigate the world
ML Platform	Unlock ML development efficiency to accelerate our development
Product Engineering	**Launch the world's best AV ridesharing service**
Remote Assistance	Ensure we are always able to get the AV unstuck

About Me

Guozhen La

- Software Engineer
- Been at Cruise for < 1 year
- Data Infrastructure team
- MLP Team
- Graduated Waterloo ECE in 2018

AV Data

The Cruise AV generates thousands of different messages across nodes

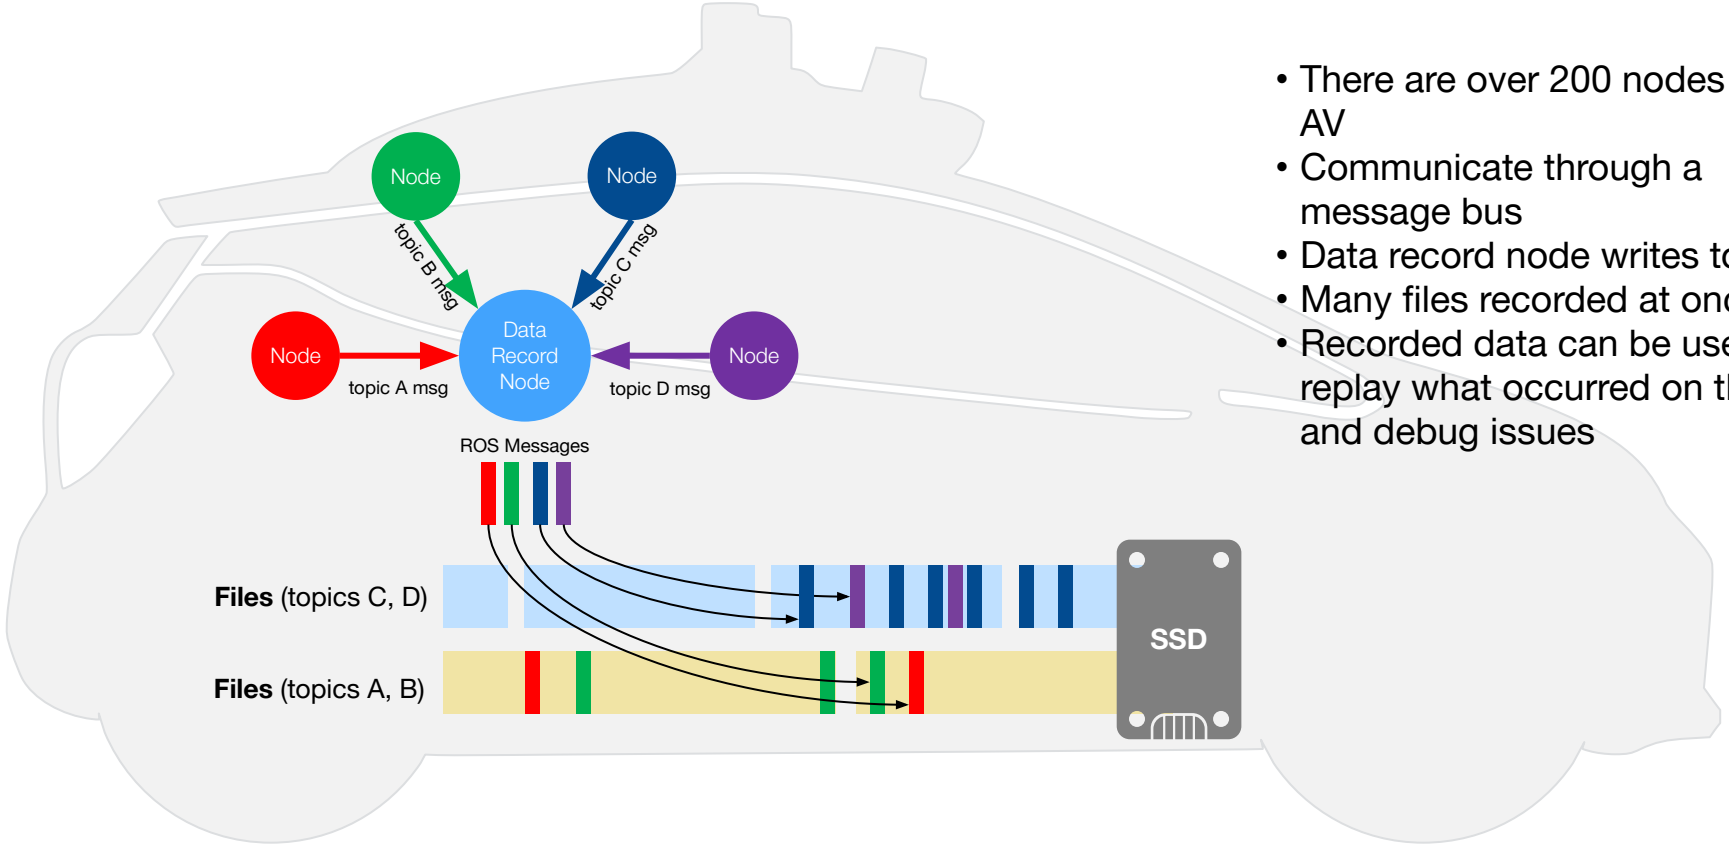

- There are over 200 nodes in the AV
- Communicate through a message bus
- Data record node writes to disk
- Many files recorded at once
- Recorded data can be used to replay what occurred on the AV and debug issues

AV Data Record Volumes

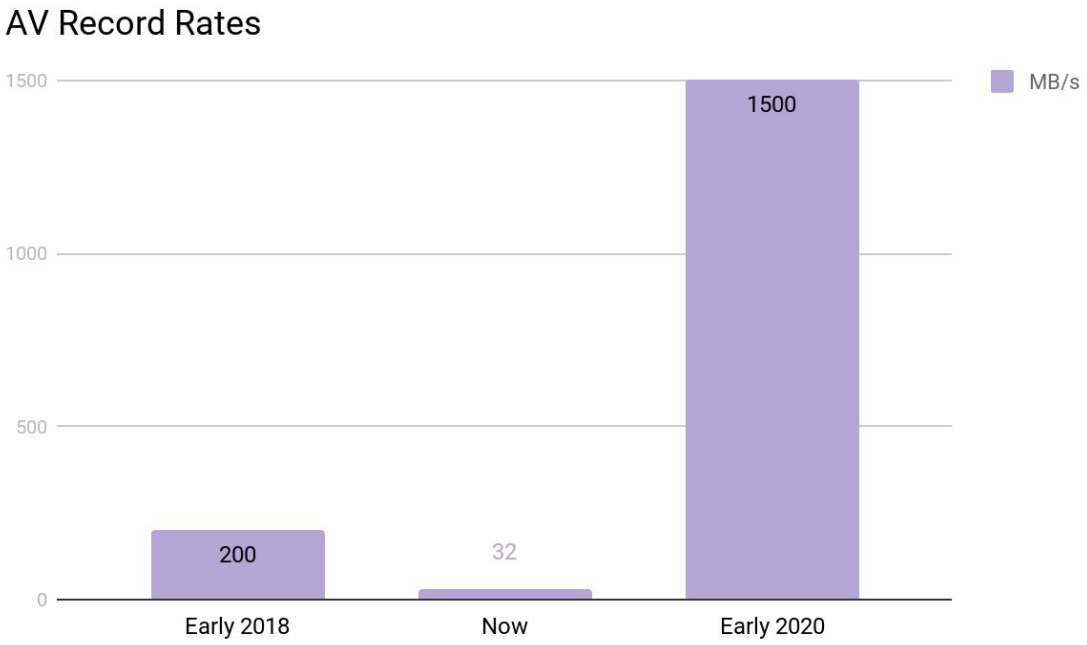

AV Data Pipeline

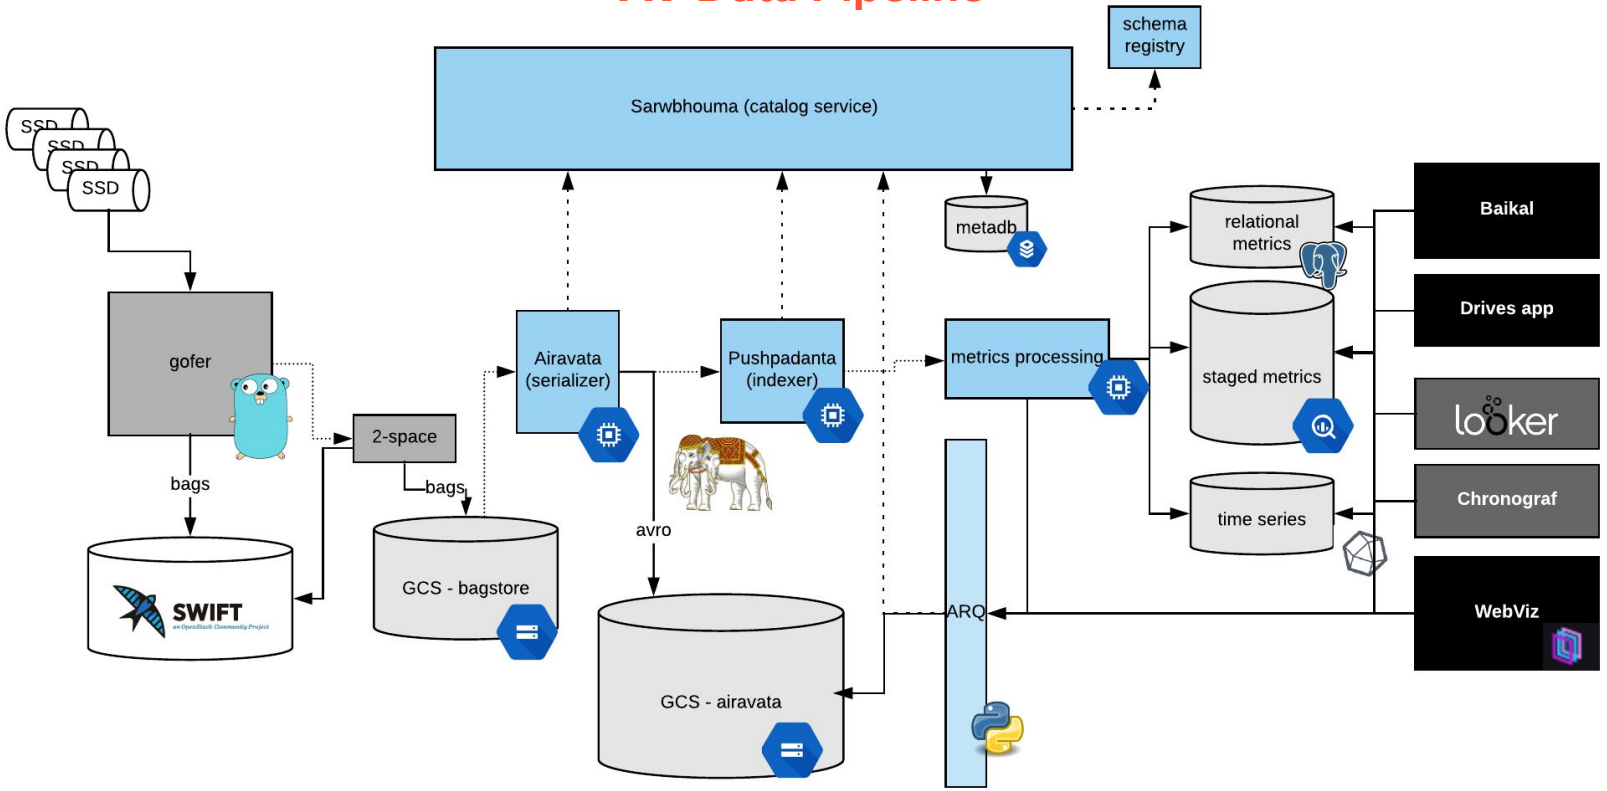

cruise

Customers

Other data @ Cruise

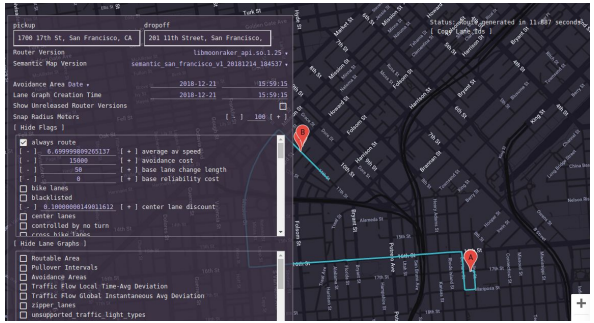

Maps & Routing

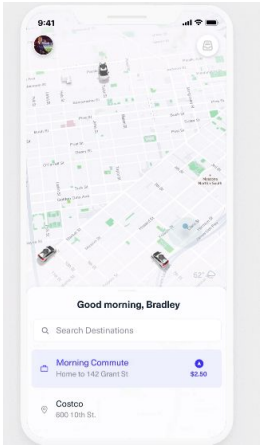

Ridesharing

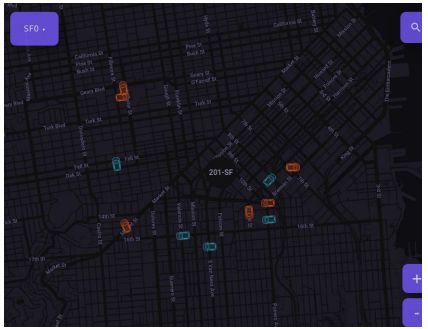

Fleet management

Simulated/Test

Labeled datasets

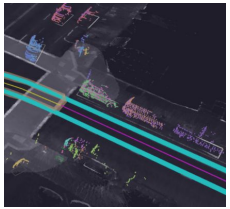

cruise

Machine Learning Platform

Groups at Cruise (Engineering)

AV Platform	Building the nerves and muscles of the AV
Core Infrastructure	Deploy highly scalable systems
Data Infrastructure	Enable insights from our extensive data collection
Data Science	Use our diverse data to improve business and engineering decisions
Engineering Productivity	Build a world-class development environment at Cruise
Ground Truth	Enable our AVs to understand the world
Mapping	Enable our AVs to navigate the world
ML Platform	**Unlock ML development efficiency to accelerate our development**
Product Engineering	Launch the world's best AV ridesharing service
Remote Assistance	Ensure we are always able to get the AV unstuck

cruise

Machine-Learning and Autonomous Vehicles

Think about a robot car trying to make a right-turn in a quiet neighborhood

cruise

Machine-Learning and Autonomous Vehicles

Think about a robot car trying to make a right-turn in a quiet neighborhood

cruise

Machine-Learning and Autonomous Vehicles

Let's look at some other examples

Can you come up with heuristics to solve, say....

cruise

Machine-Learning and Autonomous Vehicles

^ This ^

cruise

Machine-Learning and Autonomous Vehicles

^ Or this ^

cruise

Cruise leverages ML across multiple scales

- Machine-Learning is a critical component of Cruise's AV software
- Examples include...

Cruise leverages ML across multiple scales

In our individual AVs

- Object Detection
- Trajectory Prediction

For our AV Fleets

- Routing/ETA
- Fleet Dispatch/Management

To support other ML pipelines

- ML-driven auto labeling

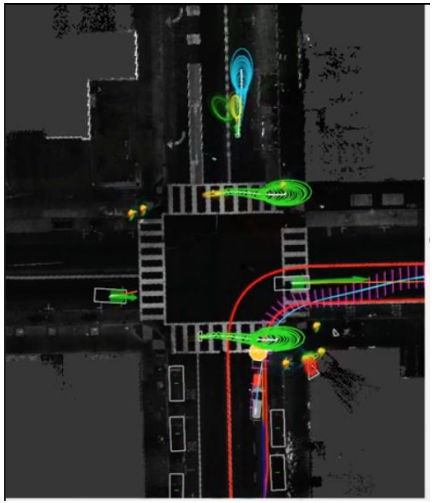

Cruise leverages ML across multiple scales

In our individual AVs

- Object Detection
- Trajectory Prediction

For our AV Fleets

- Routing/ETA
- Fleet Dispatch/Management

To support other ML pipelines

- ML-driven auto labeling

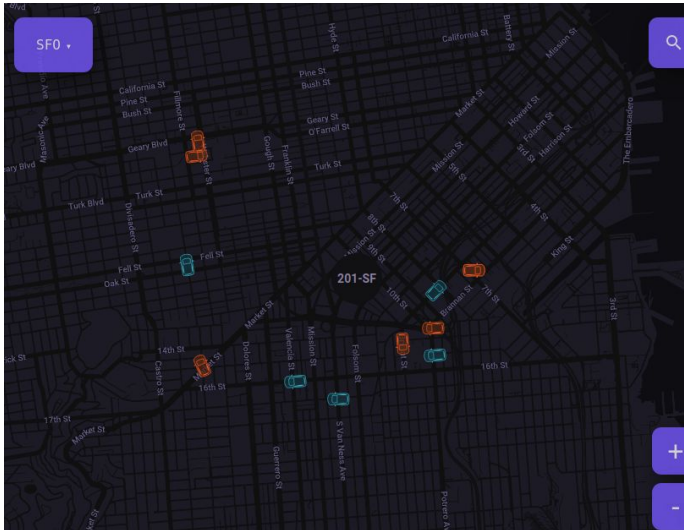

Cruise leverages ML across multiple scales

In our individual AVs

- Object Detection
- Trajectory Prediction

For our AV Fleets

- Routing/ETA
- Fleet Dispatch/Management

To support other ML pipelines

- ML-driven auto labeling

cruise

Cruise leverages ML across multiple scales

In our individual AVs

- Object Detection
- Trajectory Prediction

For our AV Fleets

- Routing/ETA
- Fleet Dispatch/Management

To support other ML pipelines

- ML-driven auto labeling

And many more!

ML-DevOps differs from “traditional” SW-DevOps

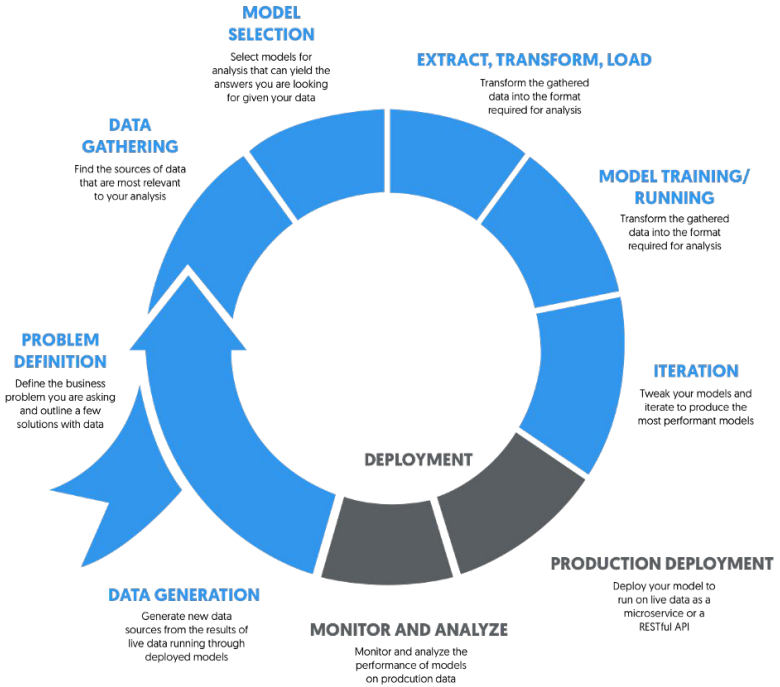

ML-DevOps: “Data is your lifeblood”

cruise

ML-DevOps differs from “traditional” SW-DevOps

Machine Learning:
The High-Interest Credit Card of Technical Debt

D. Sculley, Gary Holt, Daniel Golovin, Eugene Davydov,
Todd Phillips, Dietmar Ebner, Vinay Chaudhary, Michael Young
{dsculley,gholt,dgg,edavydov}@google.com
{toddphillips,ebner,vchaudhary,mwyoung}@google.com
Google, Inc

ML systems have an increased risk of technical debt

- Entanglement
- Data Dependencies
- Visibility Debt (Undeclared Consumers)
- Lack of common frameworks => Lots of Glue Code

ML Platforms in Industry

Dedicated ML-Platforms are a nascent but critical component of modern day SW companies

Examples:

- Michelangelo (Uber)
- FBFlow (Facebook)
- TFX (Google)
- Bighead (AirBnB)

ML Platform @ Cruise

Empower machine learning teams at Cruise to efficiently experiment, frequently deploy and safely accelerate our development.

cruise

Machine Learning Platform @ Cruise

Three Core Product Areas

Data Preparation

Discover data and prepare datasets

Training

Experiment and learn quickly

Deployment

Deploy to the cloud or car

Machine Learning Platform @ Cruise

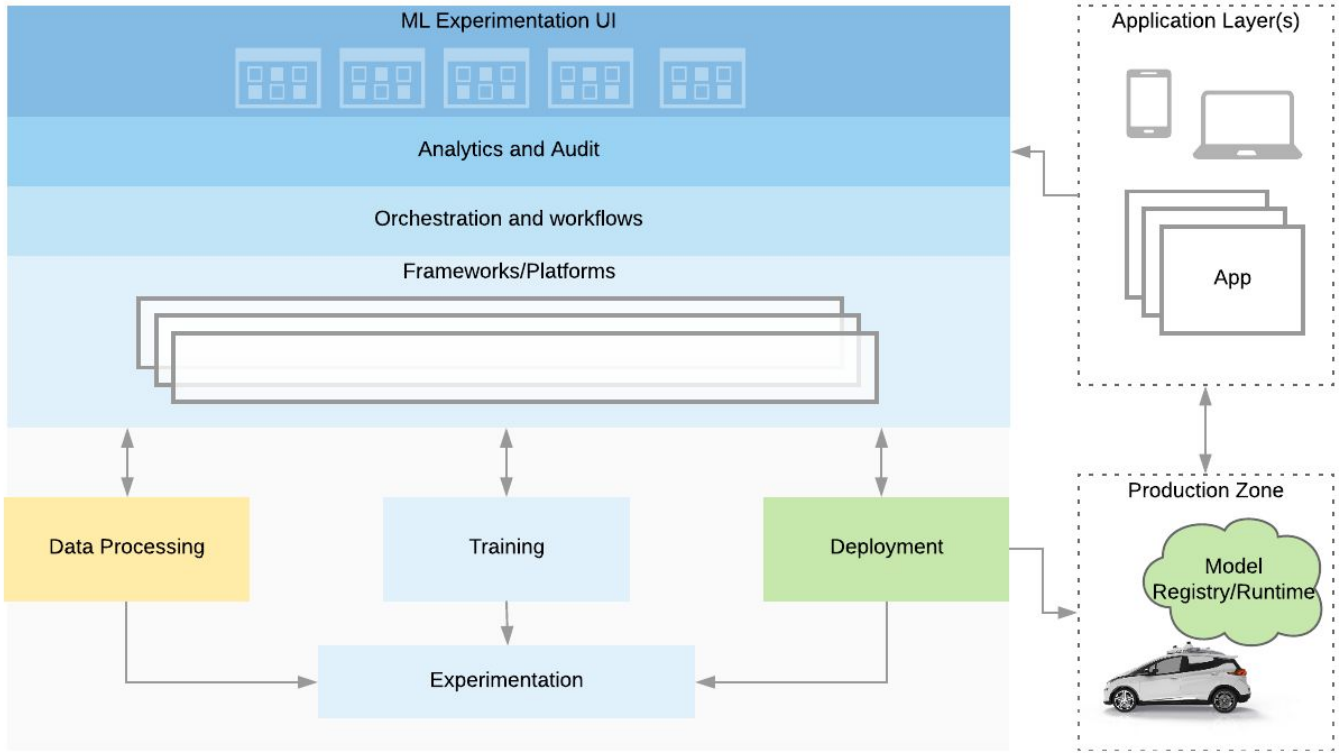

cruise

Product Engineering

About Me

Teresa Lin

- Joined Cruise Sept 2018
- Full-stack software engineer
- Product Engineering - Rideshare Backend
- MIT BS 2015, MS 2018
- New to the SF Bay Area

Groups at Cruise (Engineering)

AV Platform	Building the nerves and muscles of the AV
Core Infrastructure	Deploy highly scalable systems
Data Infrastructure	Enable insights from our extensive data collection
Data Science	Use our diverse data to improve business and engineering decisions
Engineering Productivity	Build a world-class development environment at Cruise
Ground Truth	Enable our AVs to understand the world
Mapping	Enable our AVs to navigate the world
ML Platform	Unlock ML development efficiency to accelerate our development
Product Engineering	**Launch the world's best AV ridesharing service**
Remote Assistance	Ensure we are always able to get the AV unstuck

cruise

Teams in Product Engineering

- Mobile (iOS, Android)
- In-Car Experience (Tablet)
- Dispatch
- Simulation Test Creation
- Fleet Maintenance
- Mapping

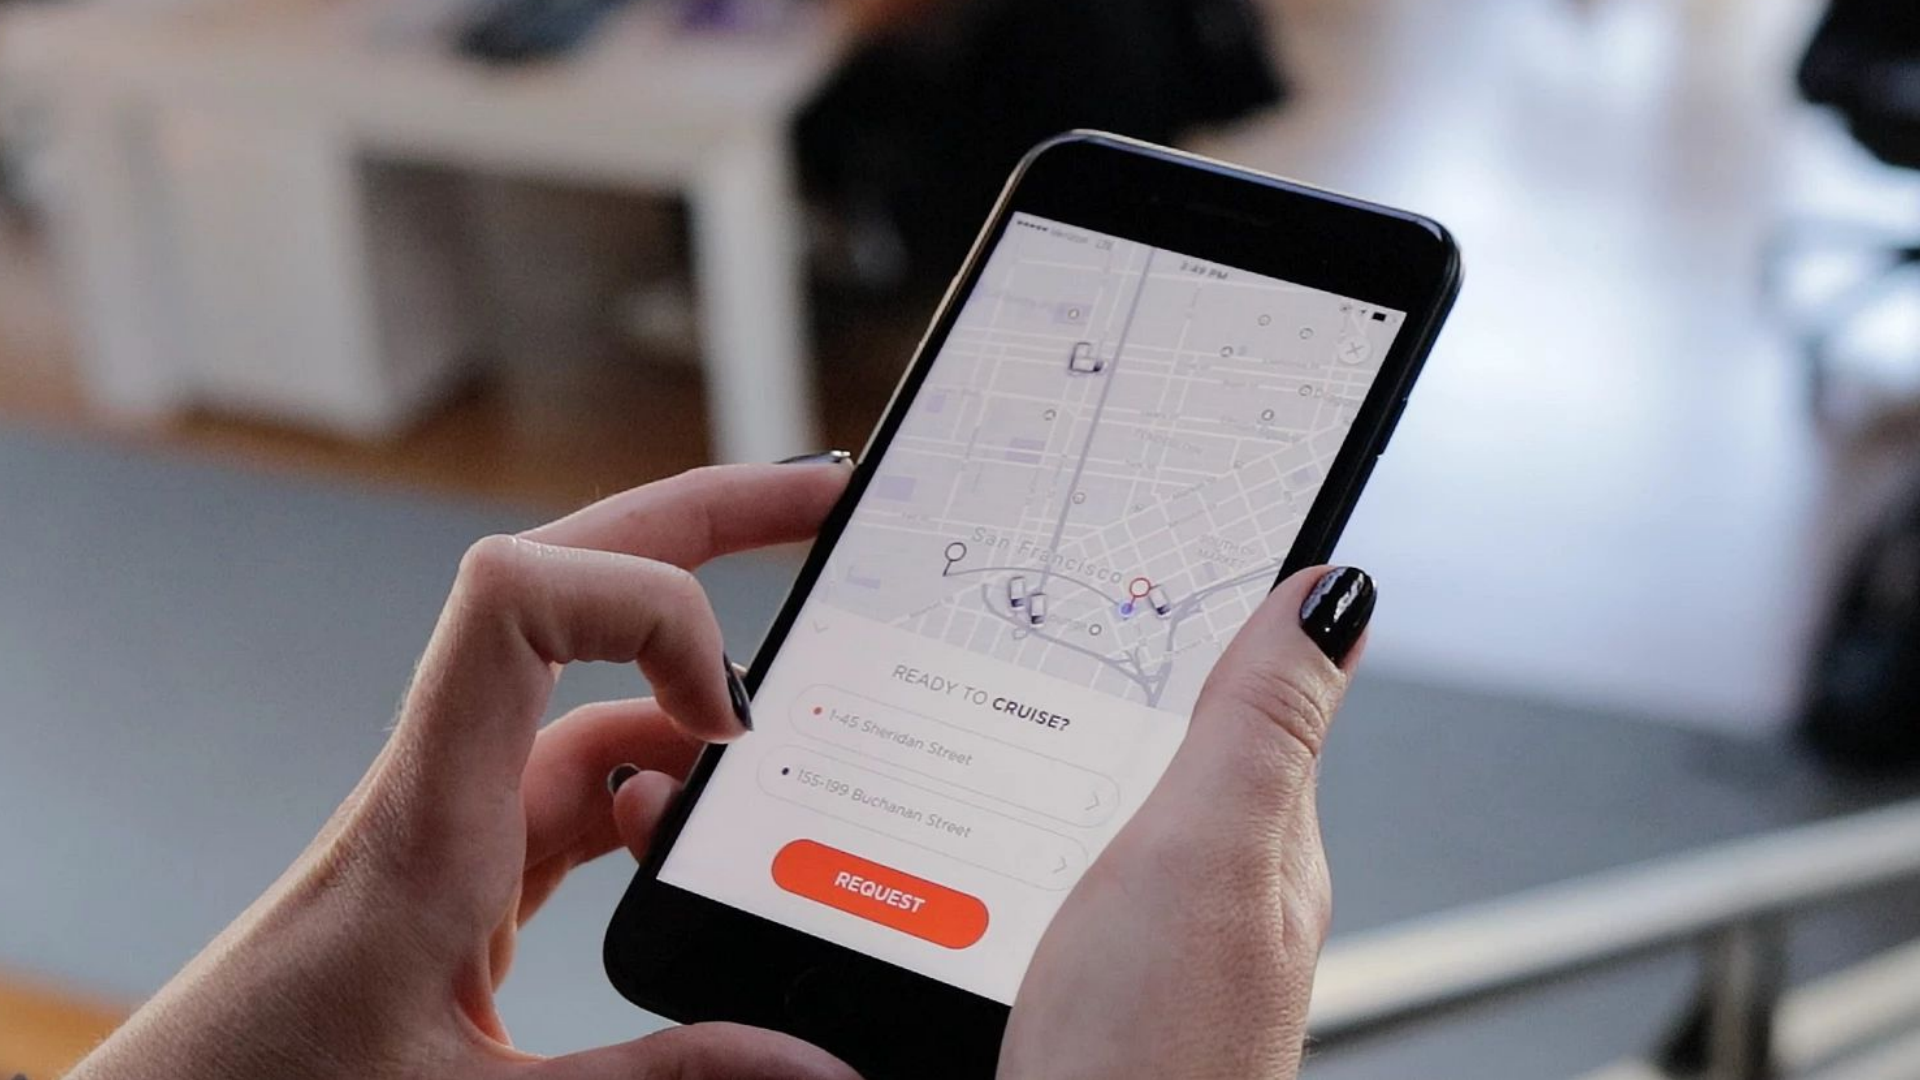
San Francisco
READY TO CRUISE?
1-45 Sheridan Street
155-199 Buchanan Street
REQUEST

Rideshare

Mobile (iOS, Android), In-Car Experience, Dispatch

Connect customers with Cruise, by delivering an *efficient, easy to use,* and *empowering*
autonomous ride-hailing experience

cruise

Rideshare

Mobile (iOS, Android), In-Car Experience, Dispatch

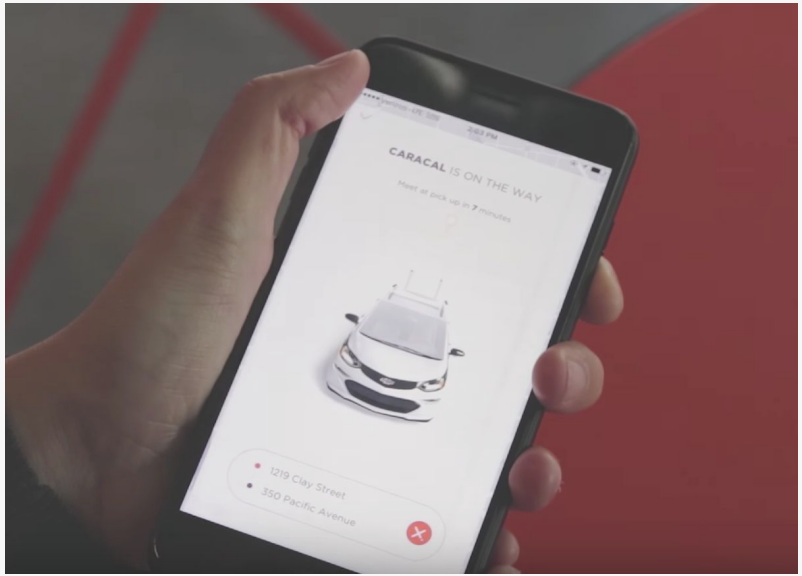

Continually updates route to most efficient one known

Familiar ridesharing interface

Detailed user experience testing*

Designing for accessibility*

Rideshare

Mobile (iOS, Android), In-Car Experience, Dispatch

Safety is #1!

- User unlocks doors
- Pullovers anytime
- In-car tablet
- Customer support
- Remote assistance*

Rideshare

Mobile (iOS, Android), In-Car Experience, Dispatch

Safety, Comfort

- Conveys what car sees
- Allows trip modifications and ride adjustments
- Integrated customer support

Rideshare

Mobile (iOS, Android), In-Car Experience, Dispatch

How do you route cars all over SF?

- Use data science models to position vehicles ideally for next trip
- Router on car is loaded with map of drivable area, lane information, etc.
- Router determines points en-route to destination

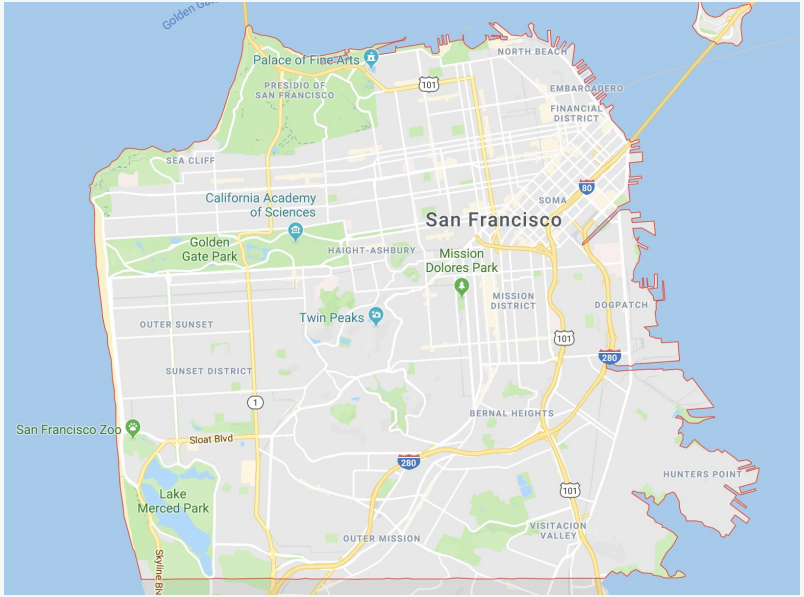

Rideshare

Mobile (iOS, Android), In-Car Experience, Dispatch

How do you track cars all over SF?

- Vehicles send information about current location, trip, etc
- Router in the cloud takes in map versions and vehicle information & determines where the vehicle will go next, ETA, etc.
- Cache map versions for faster lookup

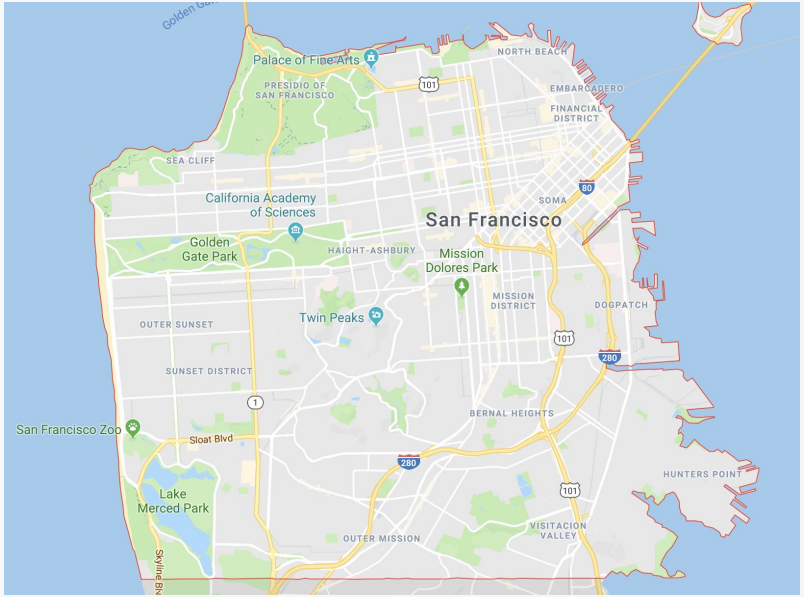

Regression Testing

How do you test changes to AV software?

Do you take a ride.. ?

How do you know if your change worked?

What if your change broke 100000 other features??

Expediting Simulation Test Creation

Simulation Test Creation

Sim tests are run on files specifying a “scenario”

Build and maintain a GUI that speeds up scenario generation for QA

Exports the scenario as necessary file for 2D sim tests

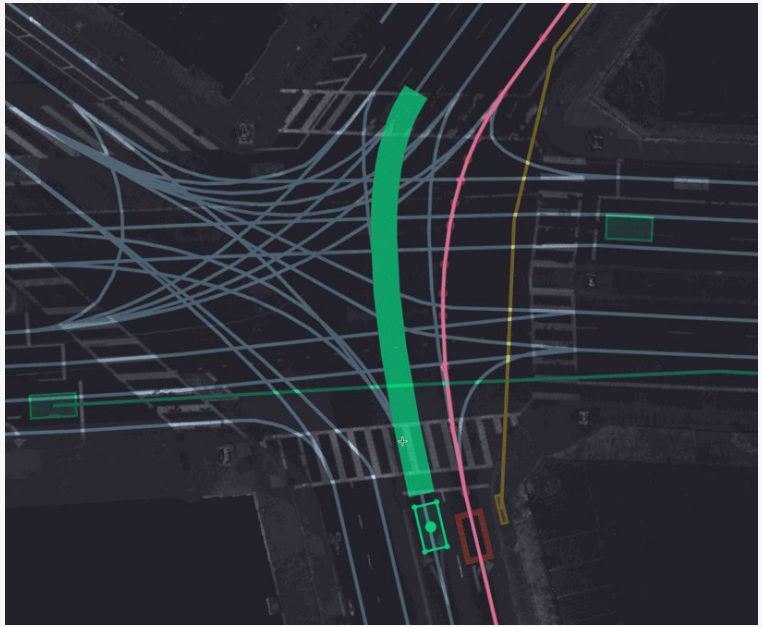


cruise
JOIN THE DRIVERLESS REVOLUTION

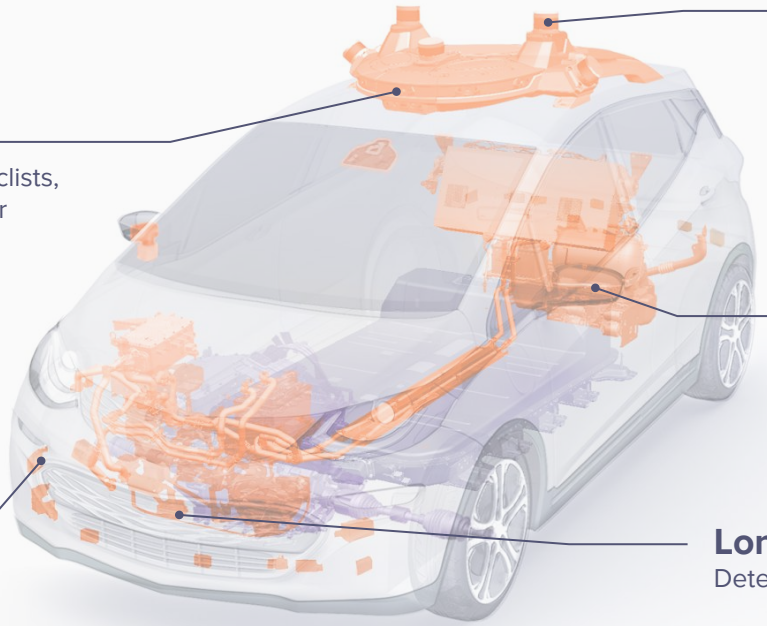

Lidar
High-precision laser sensors that detect fixed and moving objects

Cameras
Detect and track pedestrians/ cyclists, traffic lights, free space and other features

Articulating radars
Detect moving vehicles at long range over a wide field of view

Short-range radars
Detect objects around the vehicle

Long-range radars
Detect vehicles and measure velocity

Maintenance and Fleet Operations

Fleet Maintenance

100+ vehicles of different..

- Model
- Tracking numbers
- Hardware profile
- Software profile
- Use case
- Status
- Operational region
- ...and more!

Maintenance and Fleet Operations

Fleet Maintenance

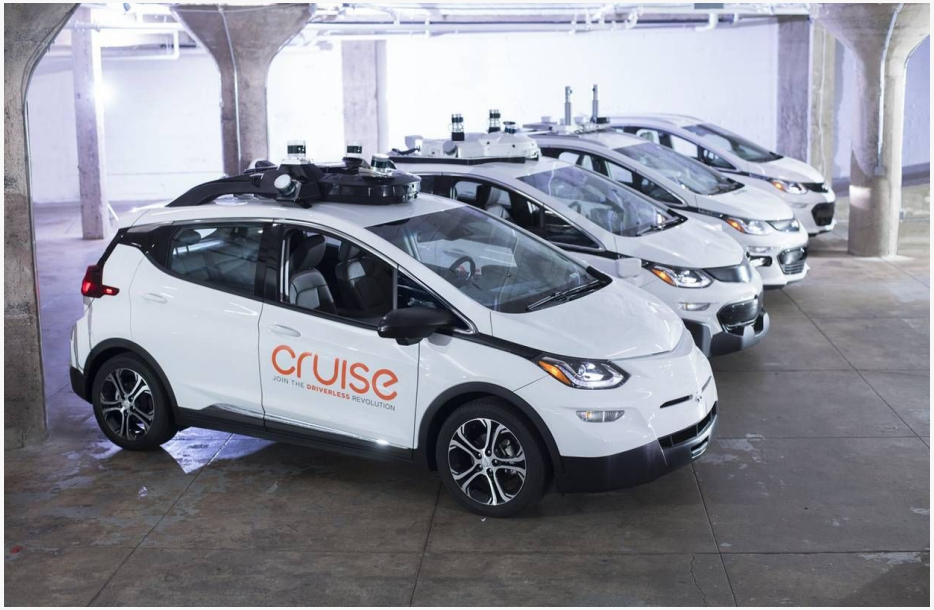

Engineers need to run tests on-car

Operators need to build software stack onto cars quickly

As number of cars scales, so must processes to maintain cars and engineering

Characterizing SF

Mapping

Driving in SF is hard!

- One-way streets
- Tram and bus only lanes
- Do not park lanes
- Placement of traffic lights
- Double parked vehicles
- Pedestrians
- Cyclists
- etc, etc, etc!

Characterizing SF

Mapping

Vehicles specifically assigned to map SF

Important to track map releases, areas to re-map, etc.

Downstream customers rely on accurate and existent maps

Why I Work at Cruise

Impact, impact, impact!

- Challenging problems
- Fast turnarounds
- Lots of ownership

Work with amazing people!

- Values ground us
- Communities within Cruise

Cool perks!

- Take rides in Cruise cars
- Company wide demos
- Snacks!