Augmented Vehicular Reality (**AVR**)

Hang Qiu, Fawad Ahmad, Fan Bai, Marco Gruteser, Ramesh Govindan

USC University of Southern California

RUTGERS THE STATE UNIVERSITY OF NEW JERSEY

GM

Reliable Autonomous Driving

Human drivers can achieve high reliability

- ~ 100,000,000 miles between fatalities

Can autonomous cars reach the same level of reliability?

Reliability with Advanced Sensing

LiDAR

Stereo Camera

Radar

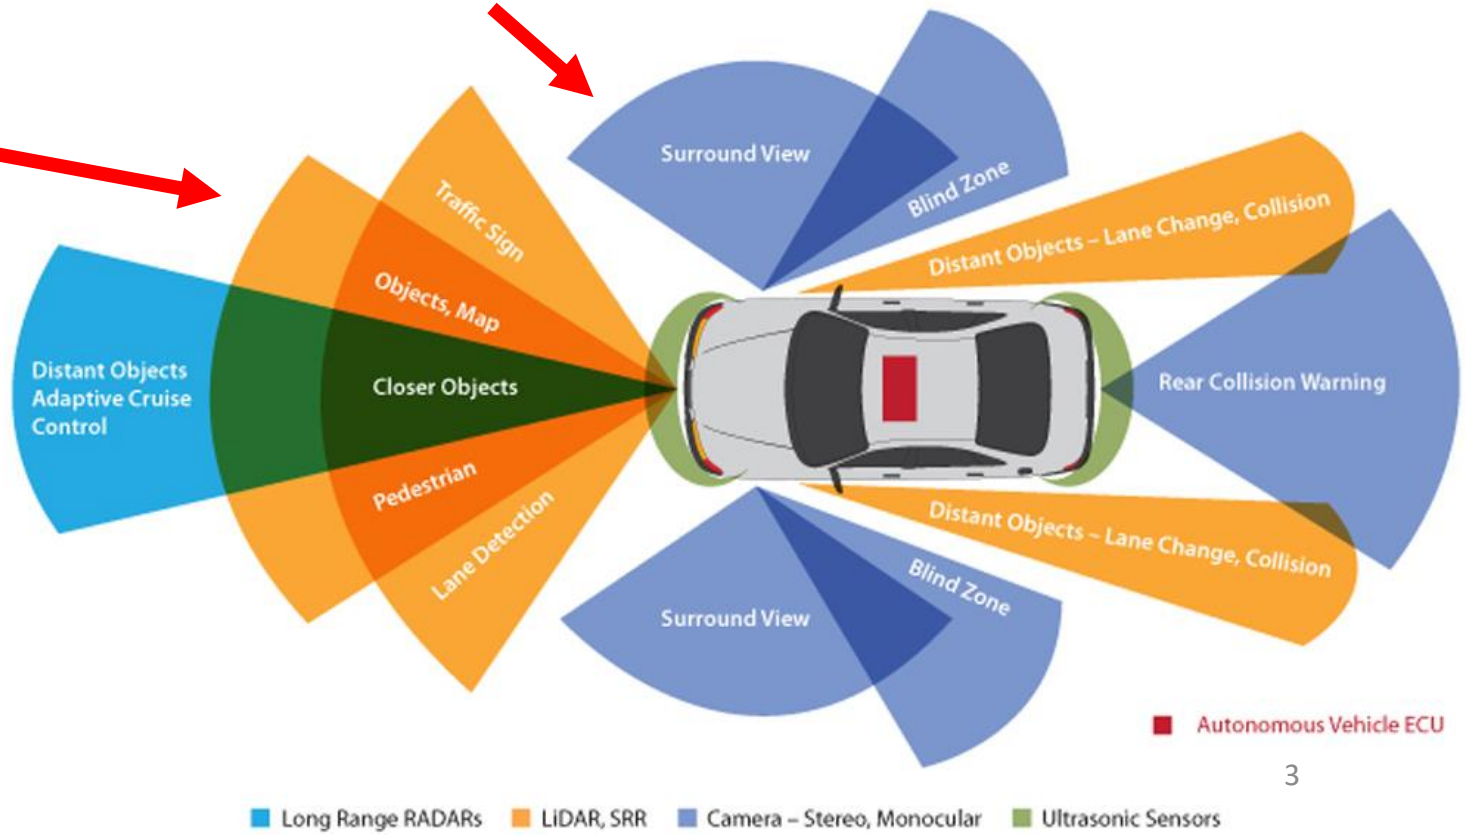

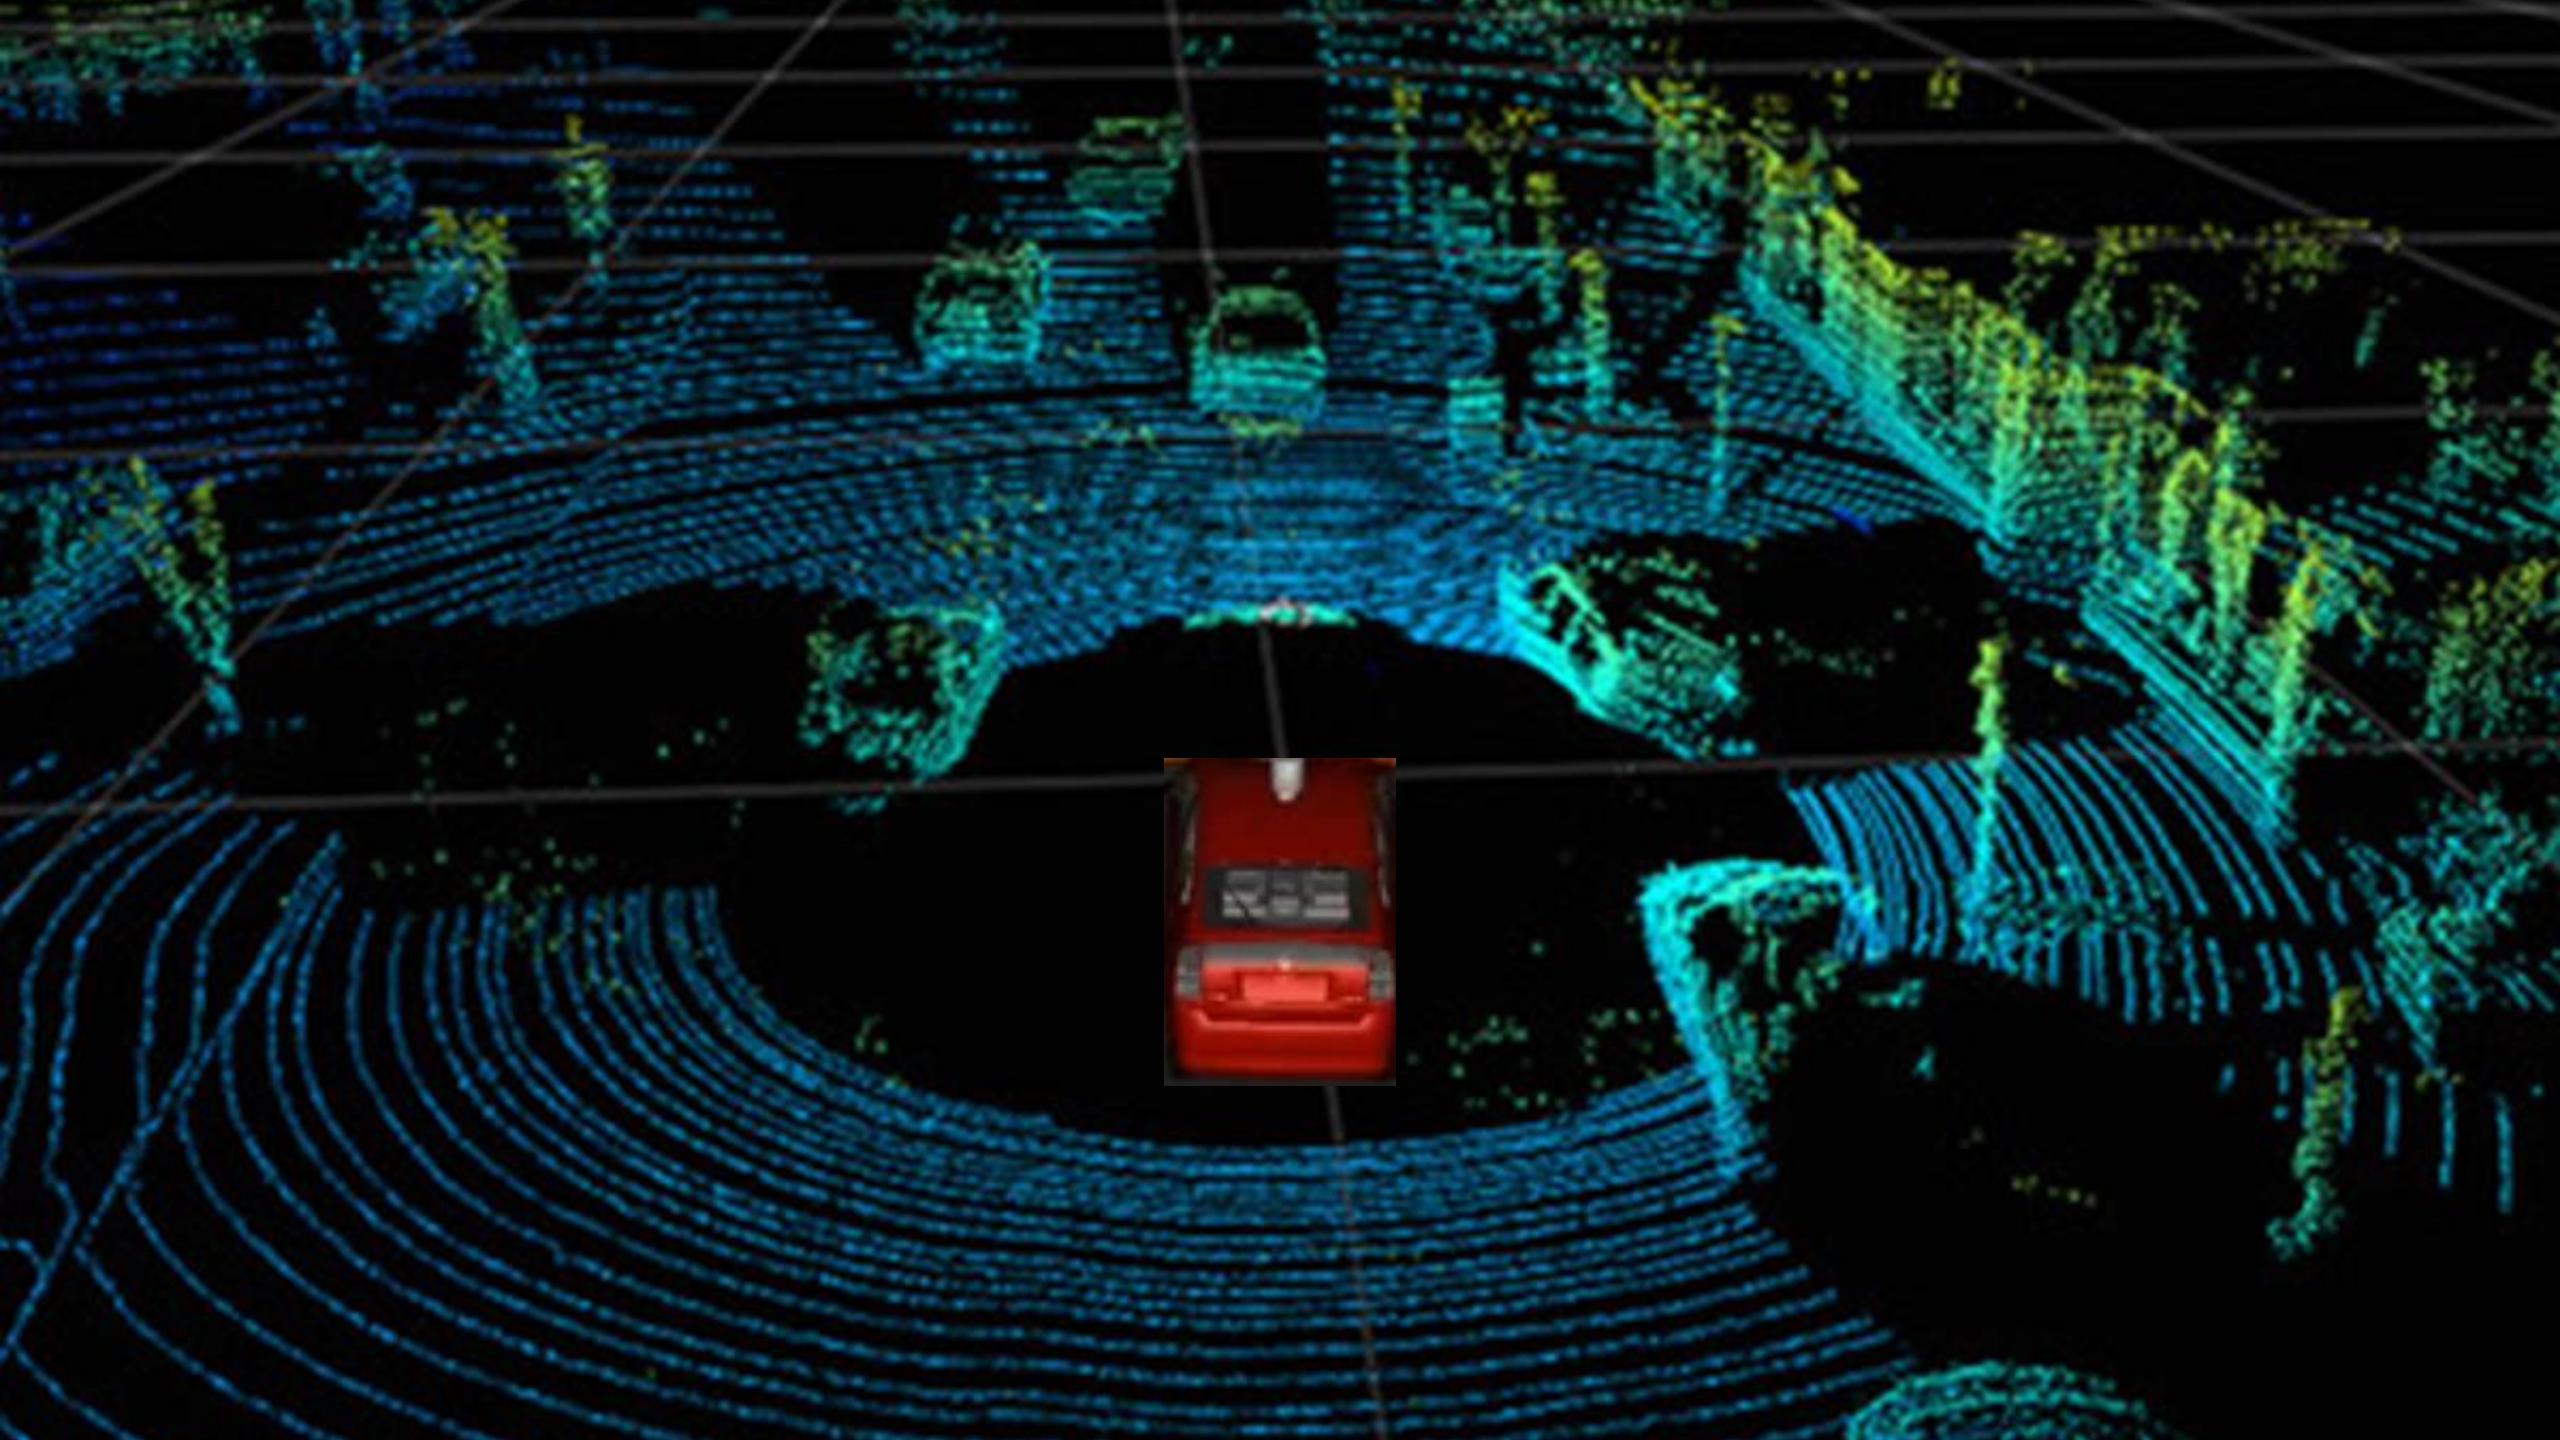

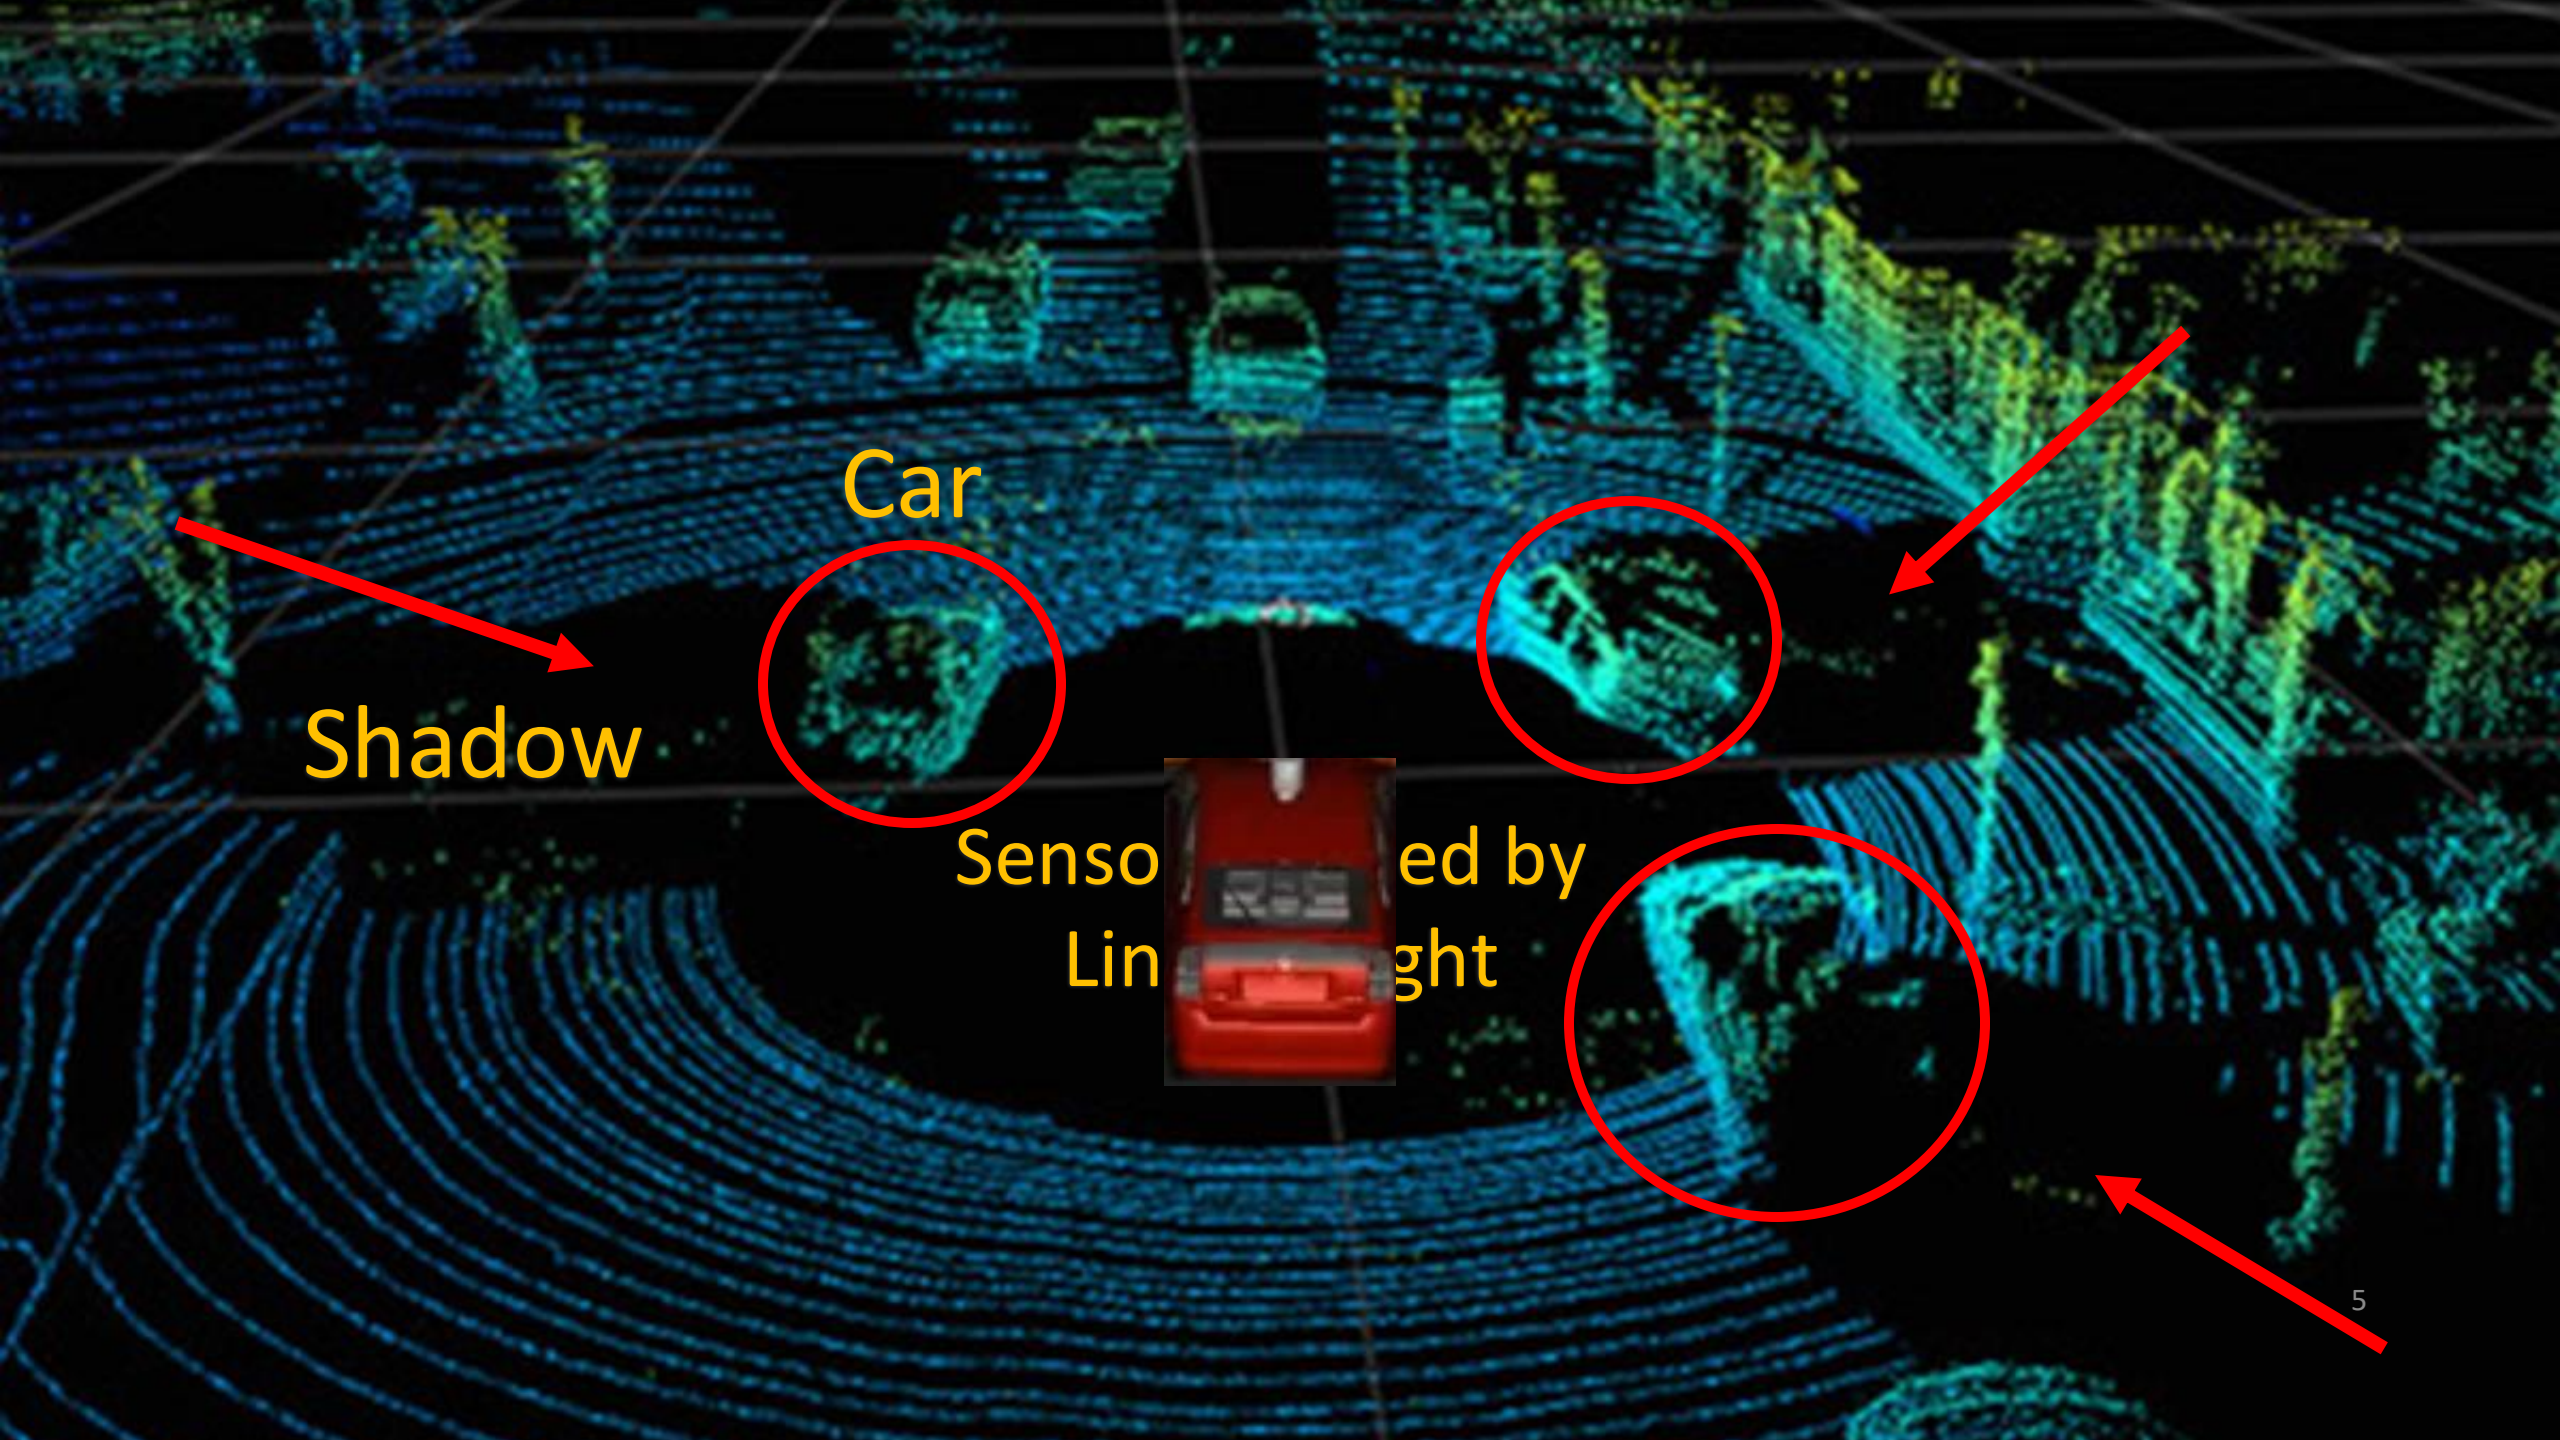
Car
Shadow
Senso ed by
Lin ght

3D perception is limited by occlusion

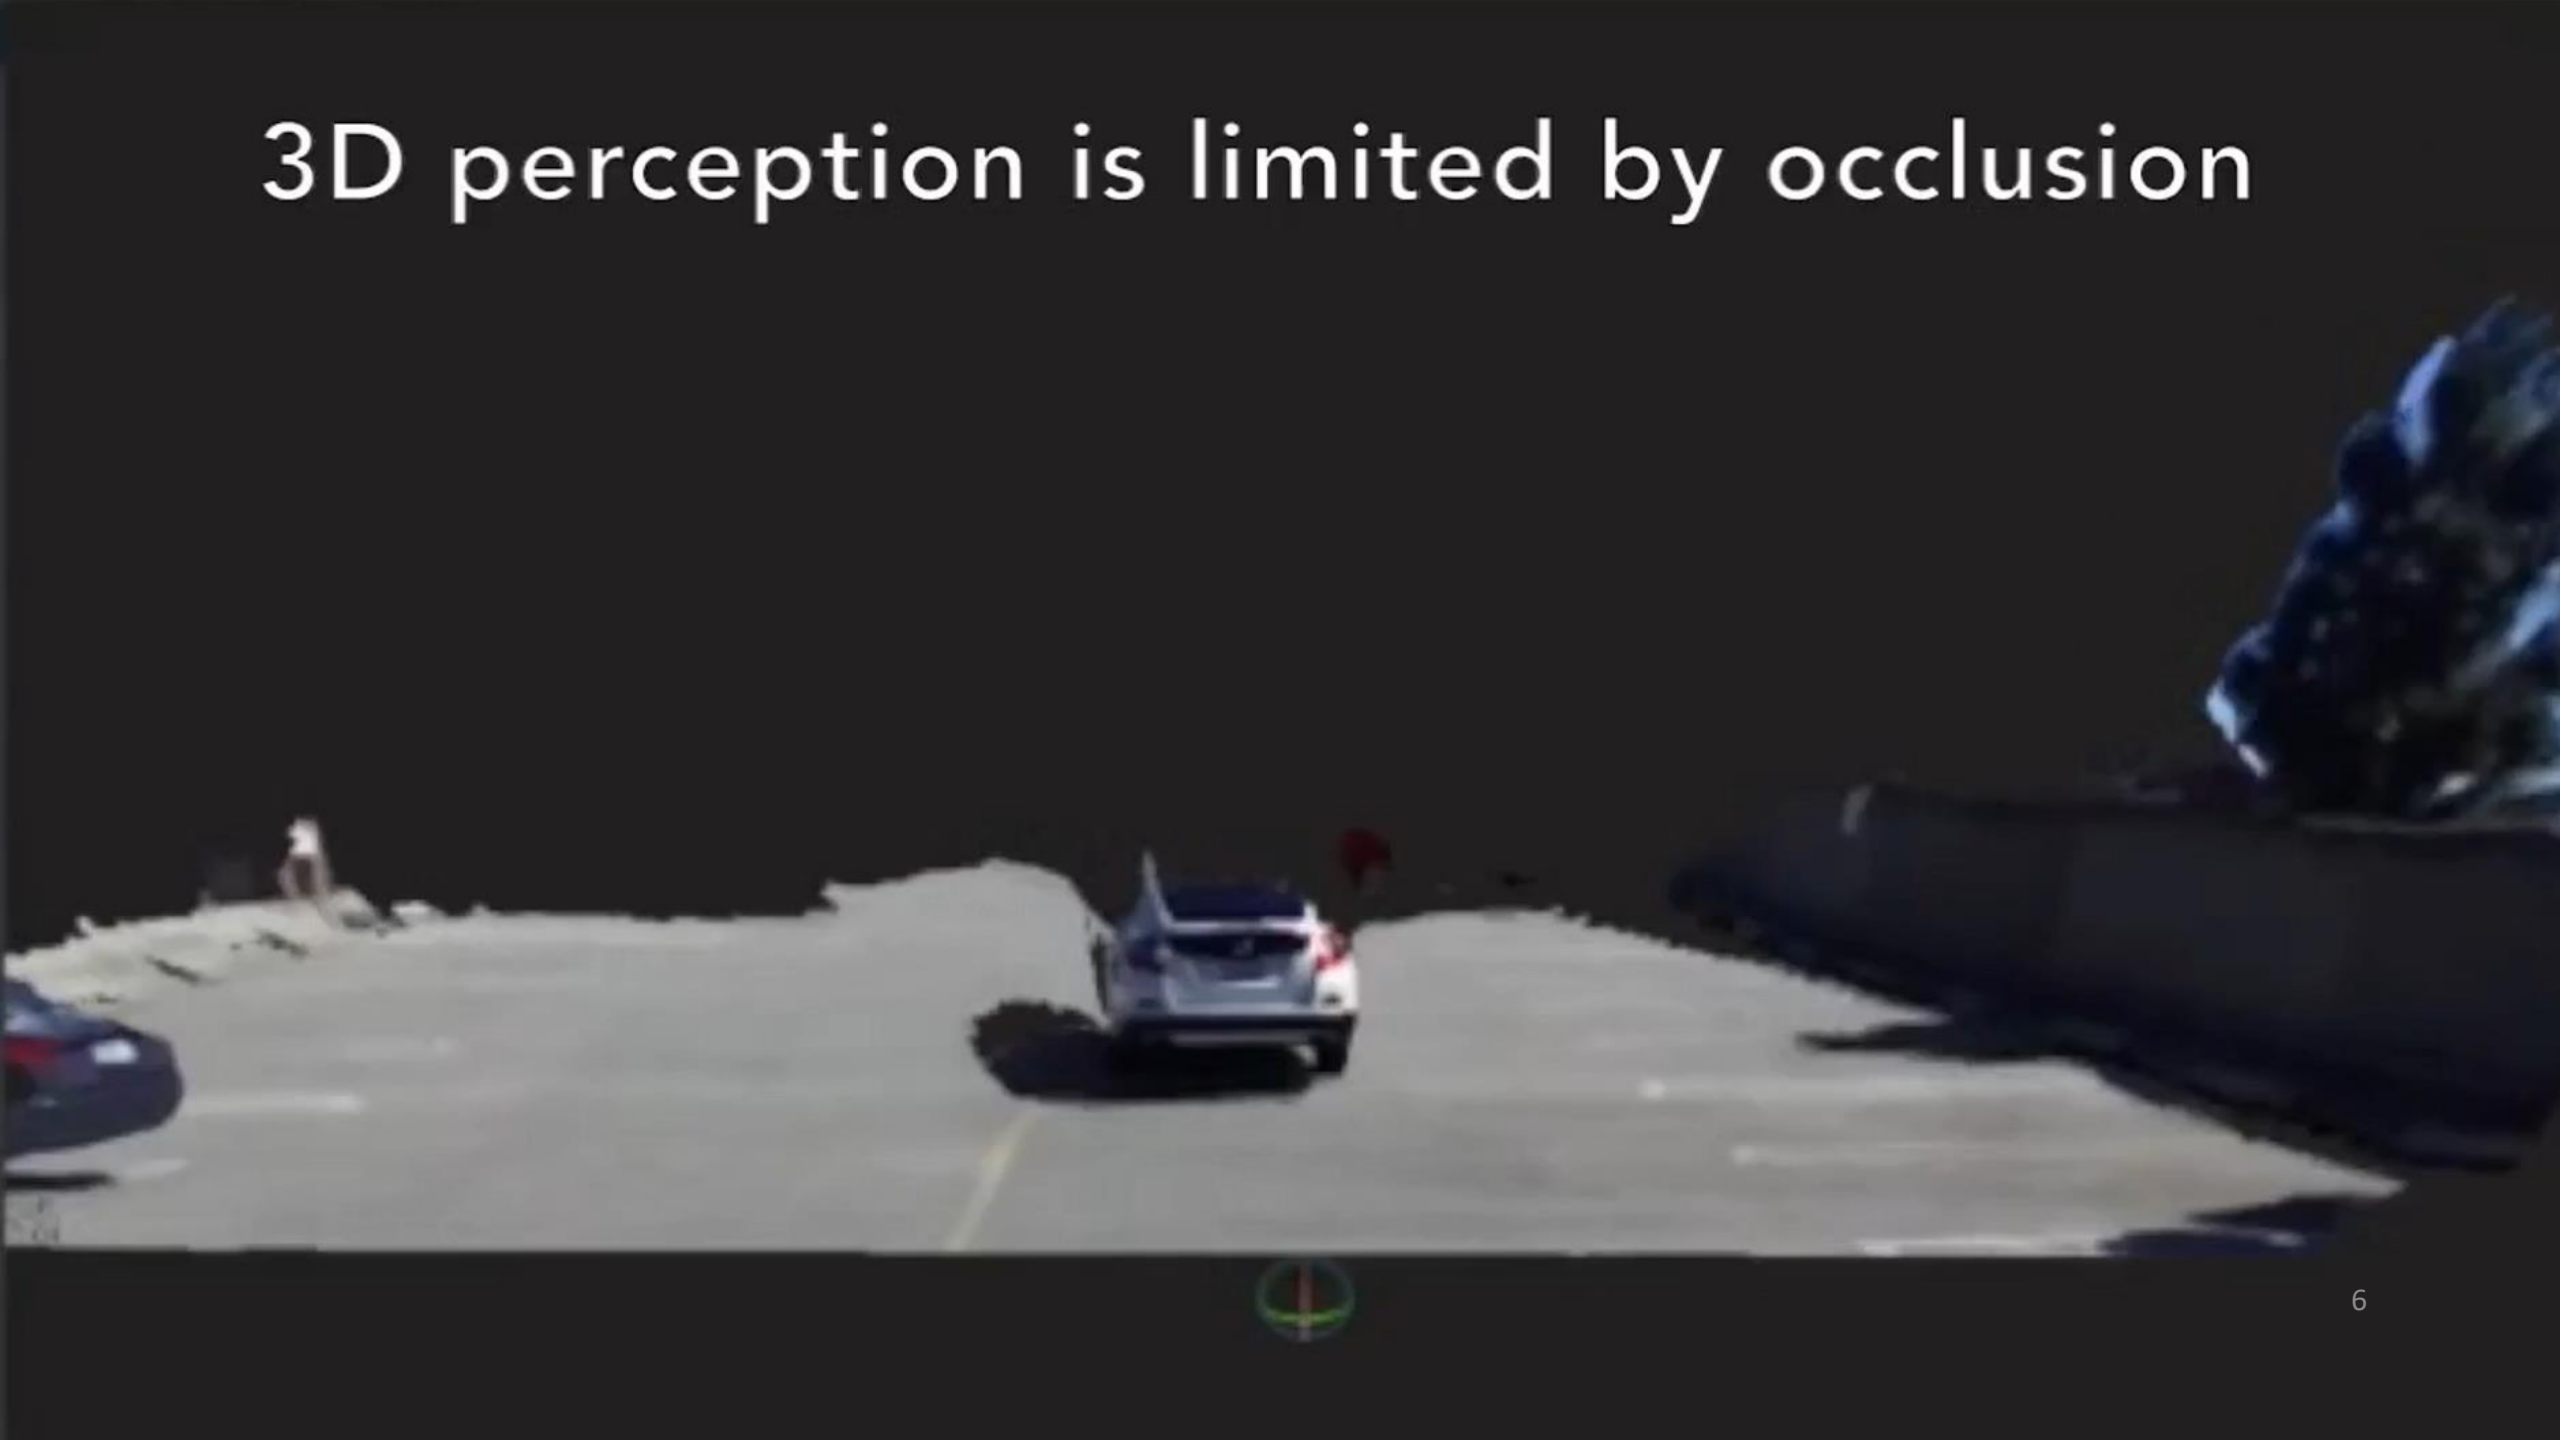

3D perception is limited by occlusion

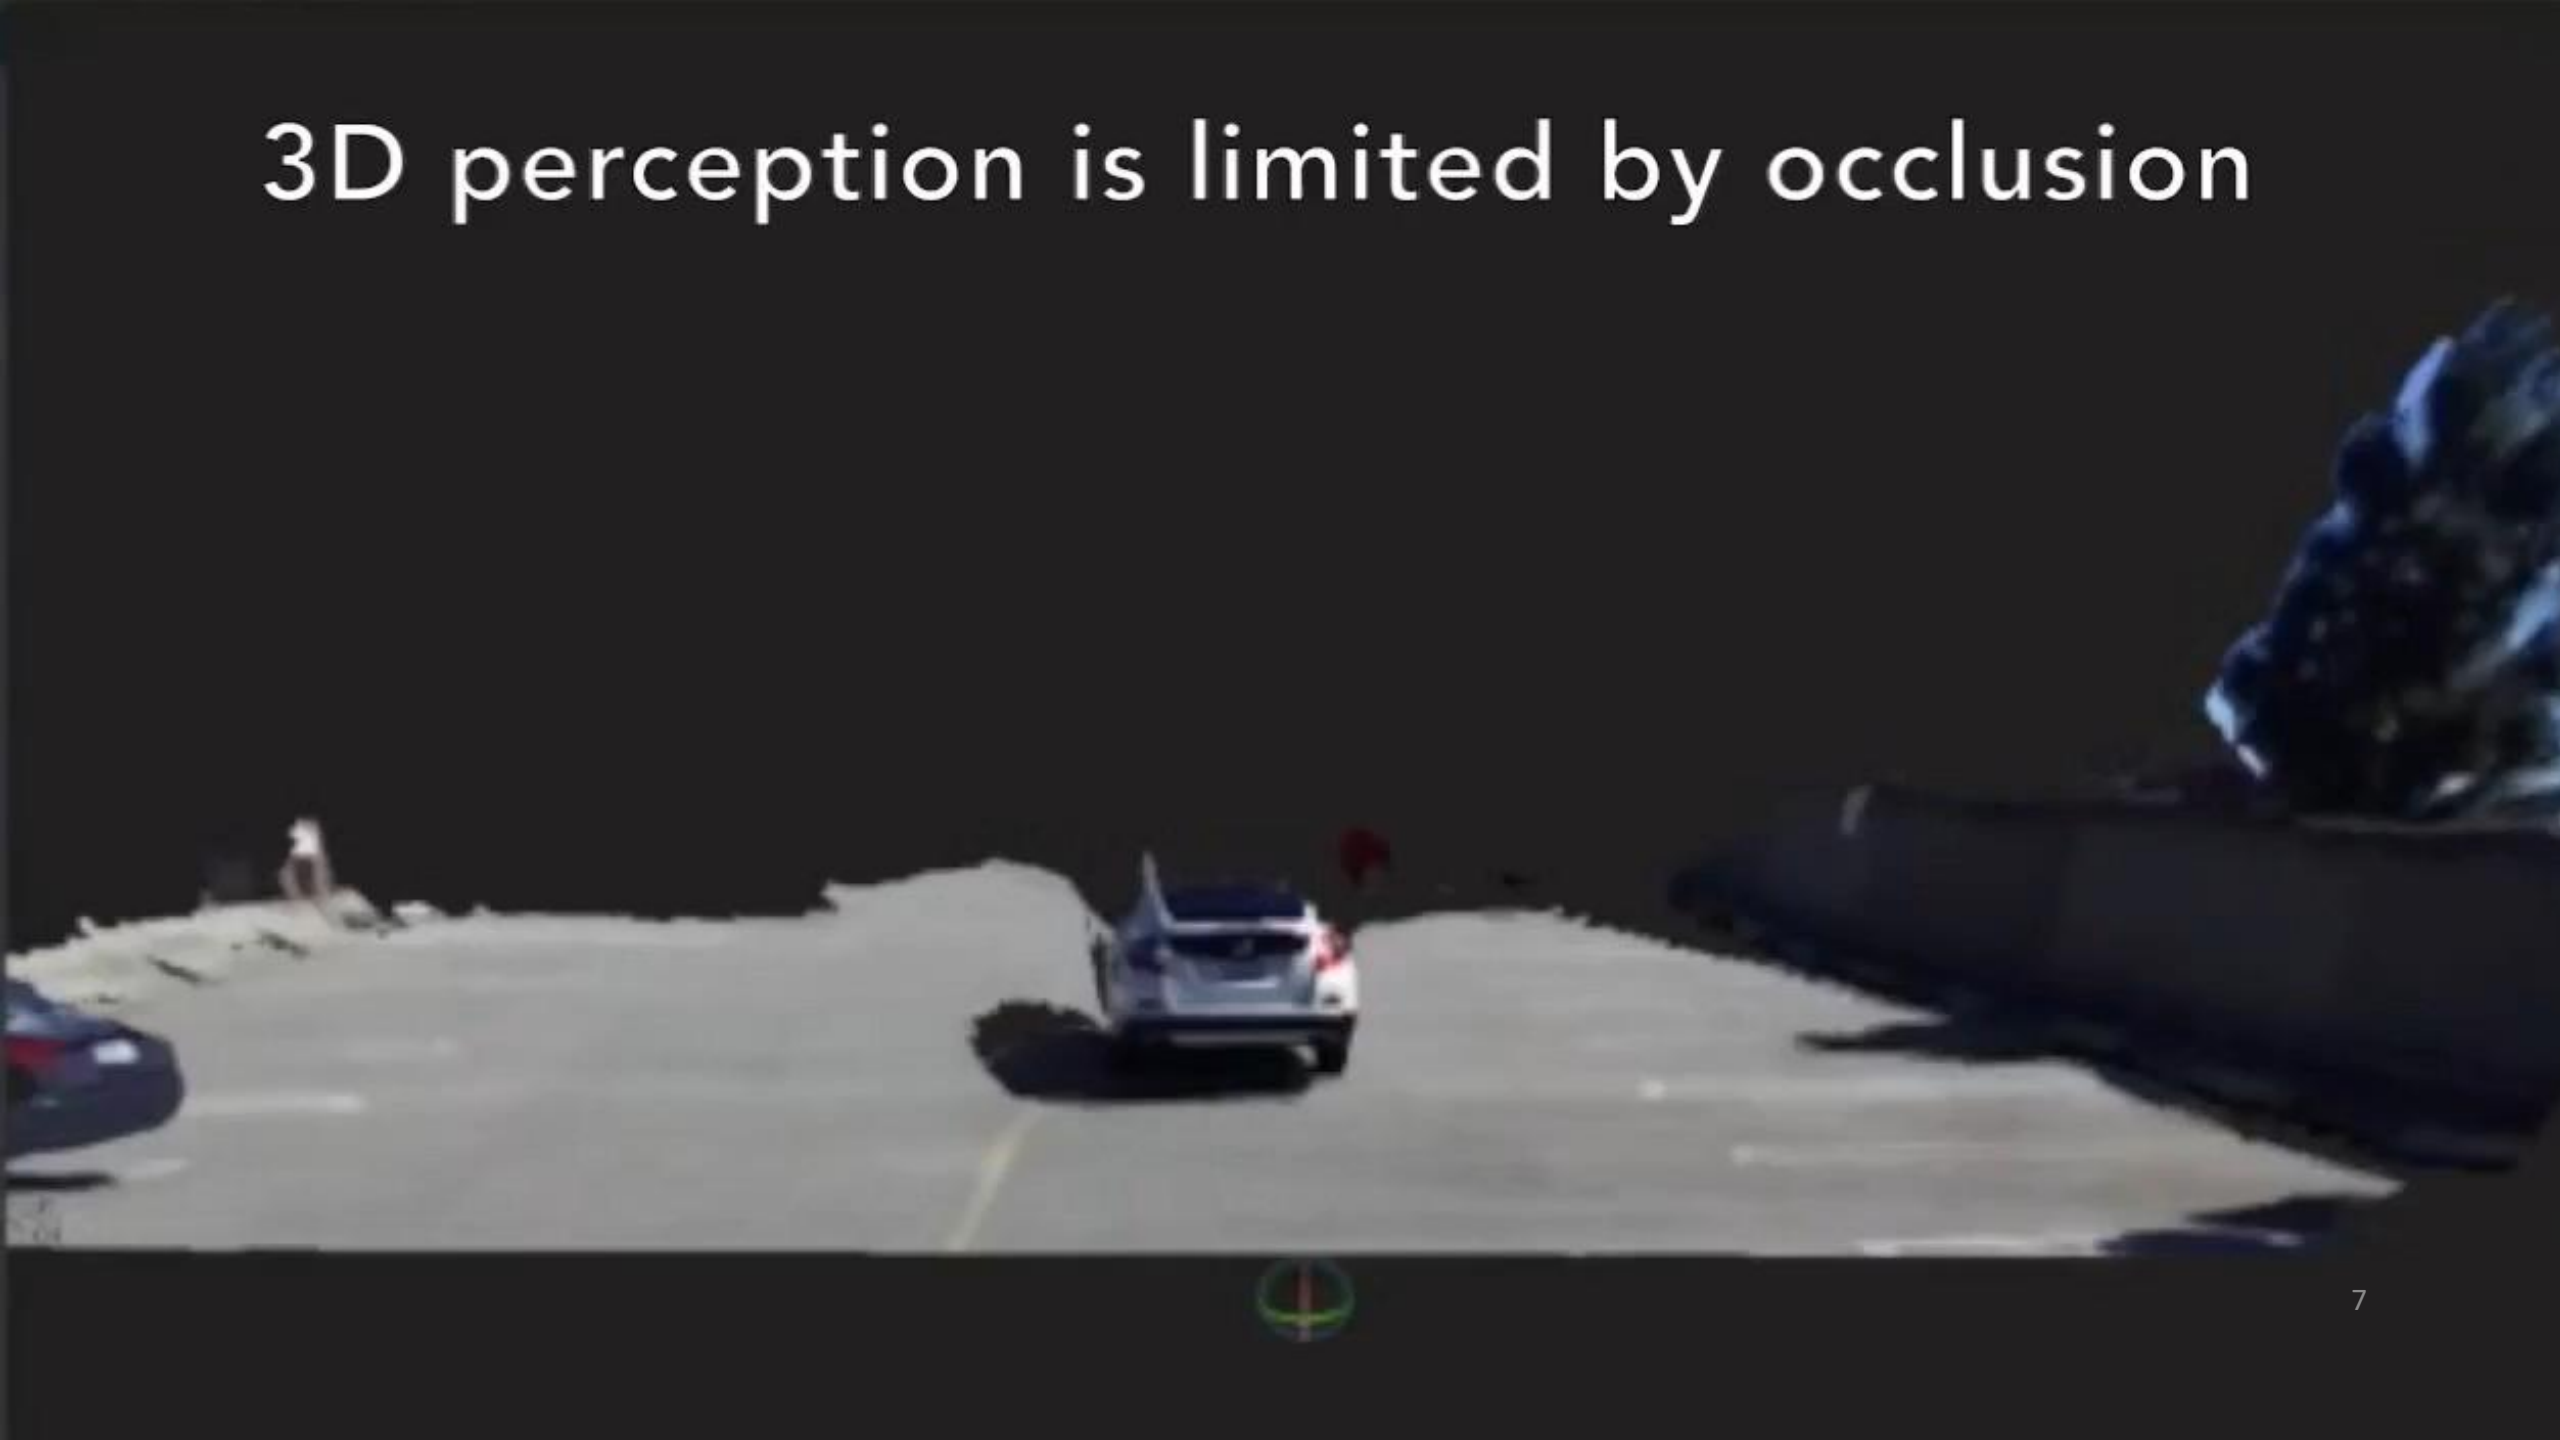

Extended Vision

Extended Vision Improves Safety

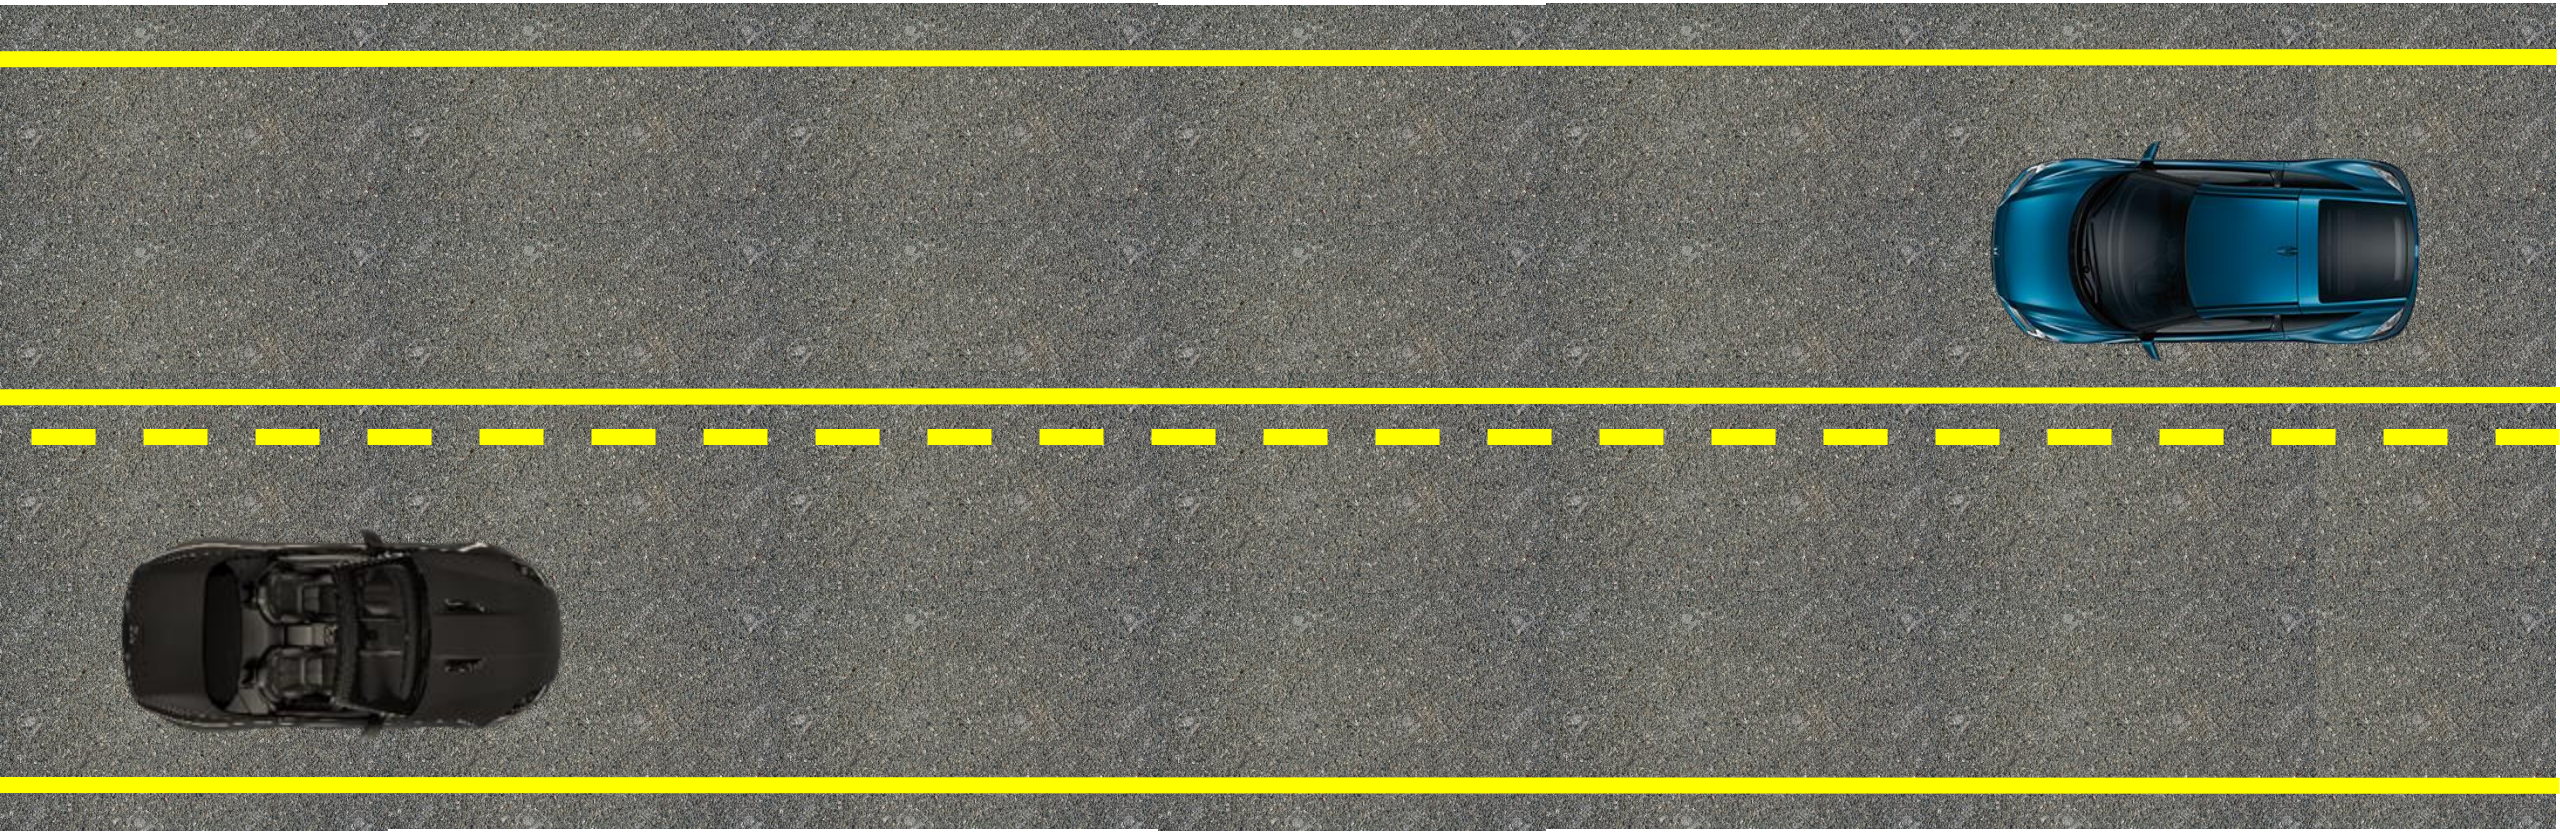

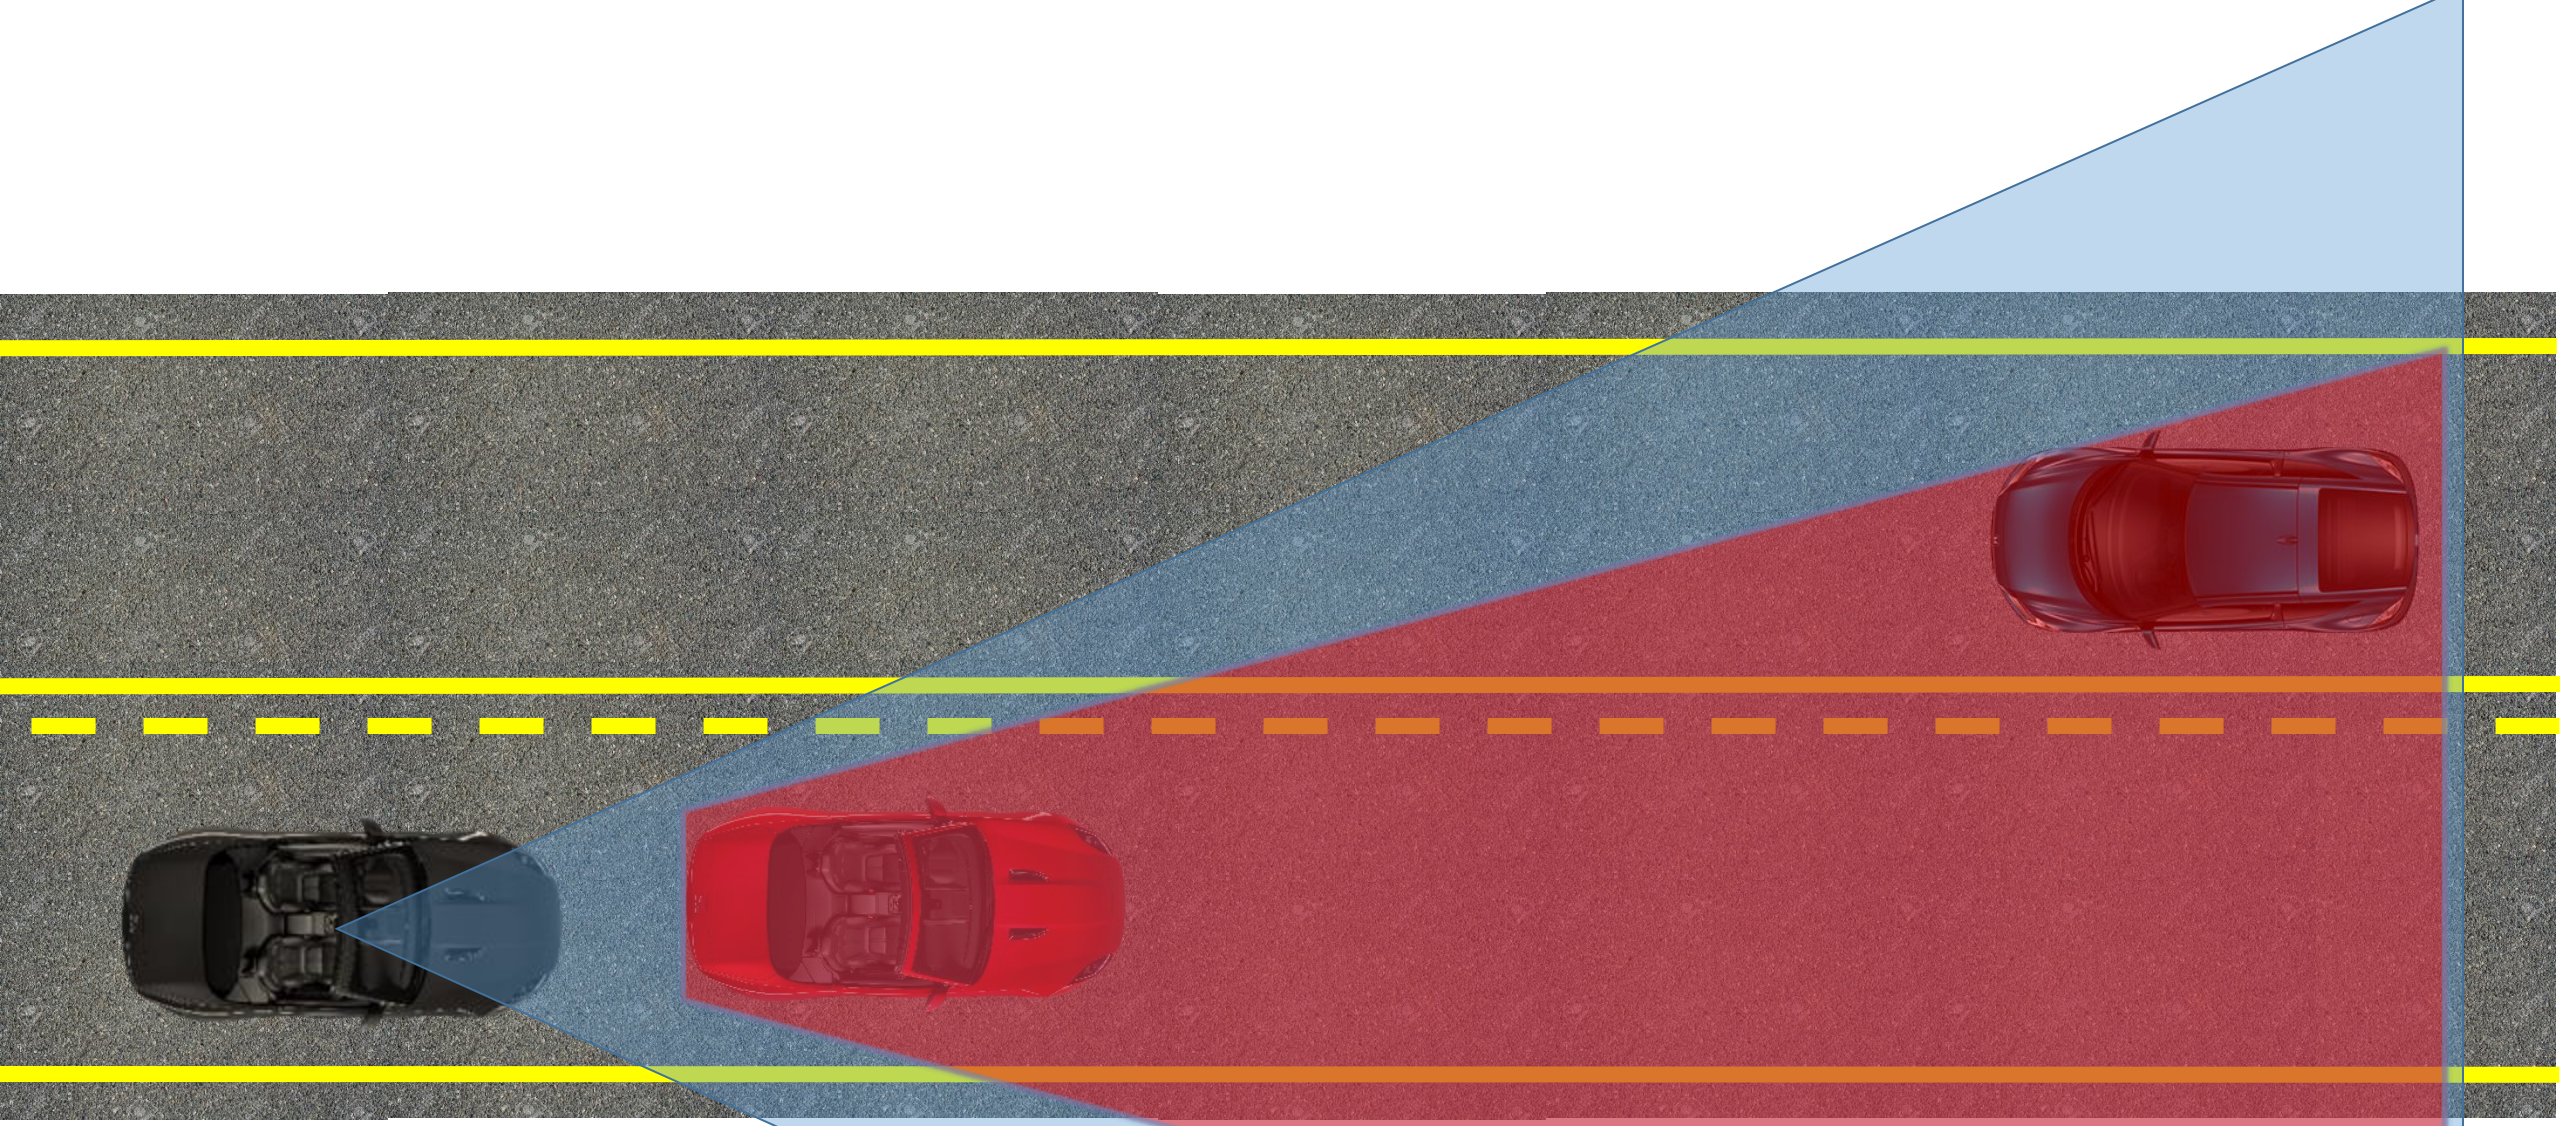

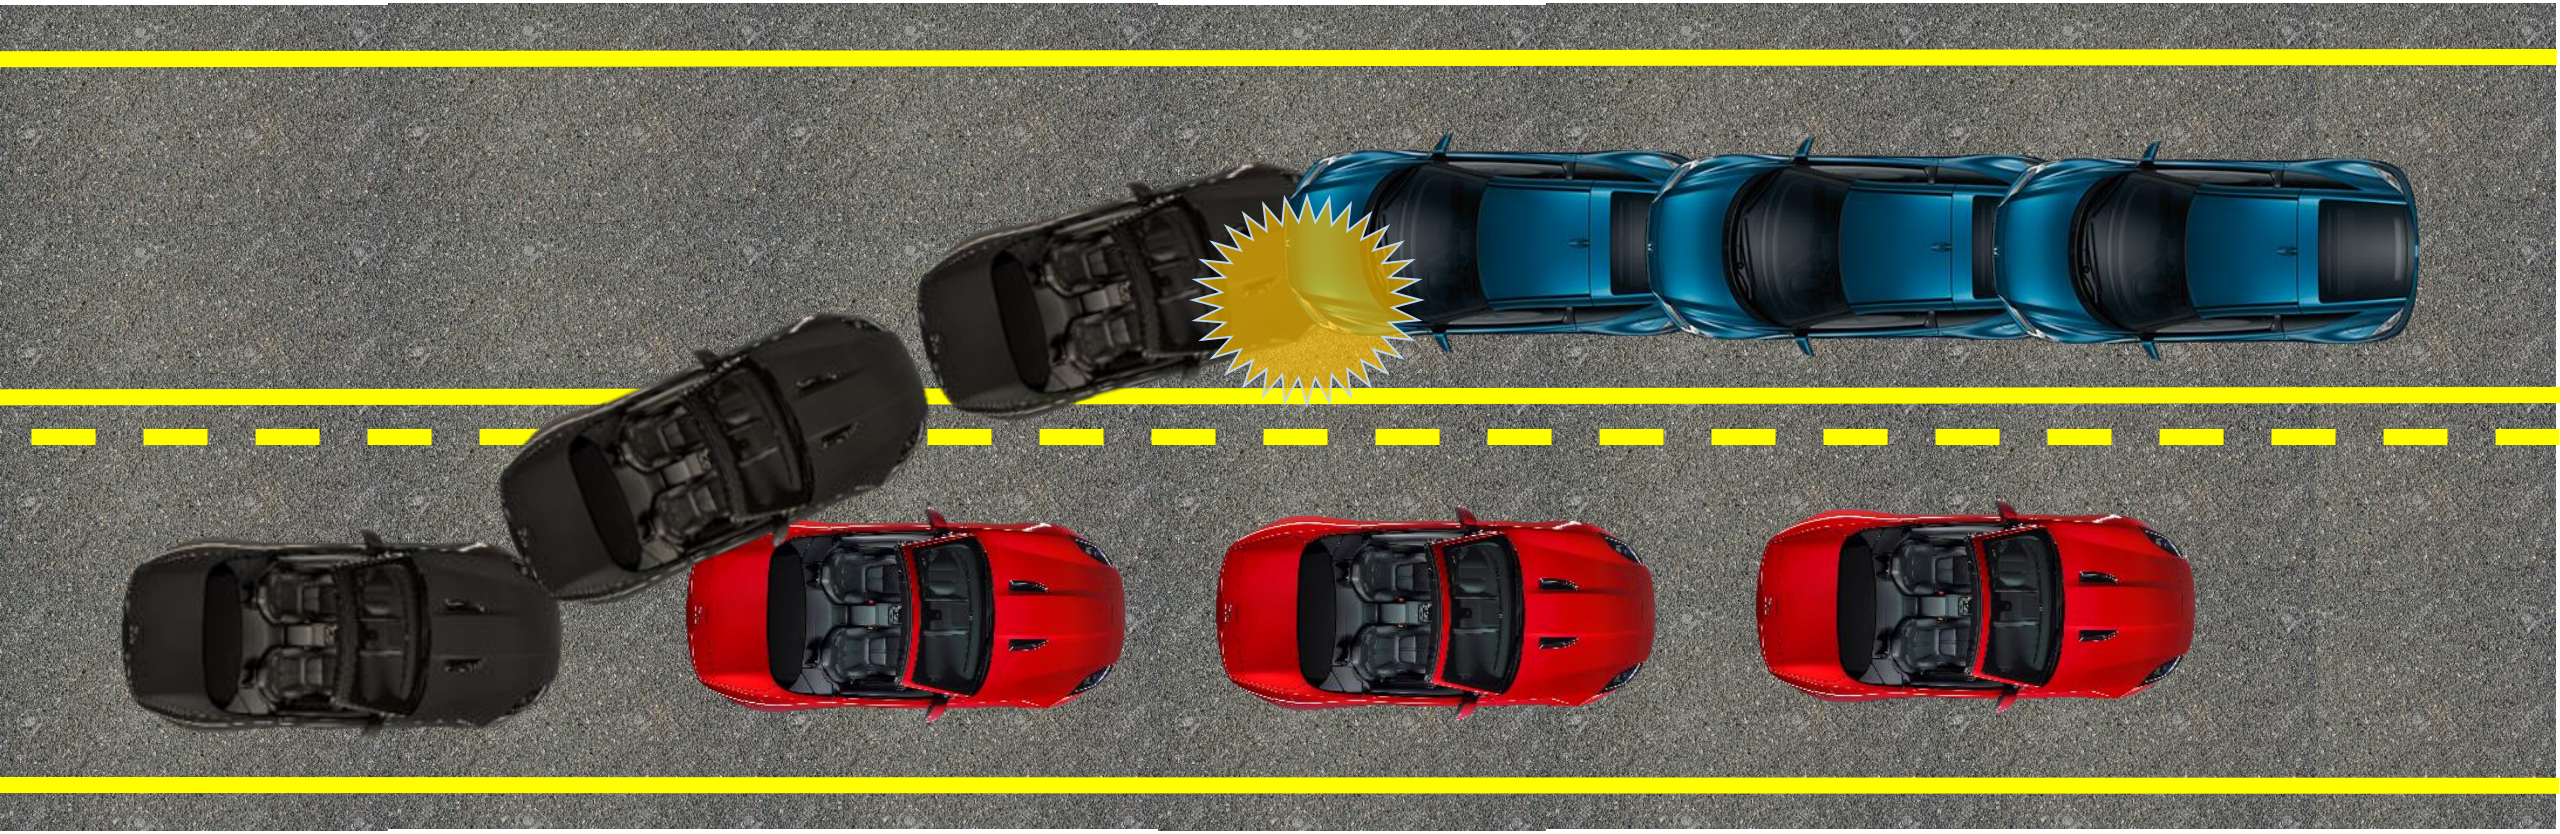

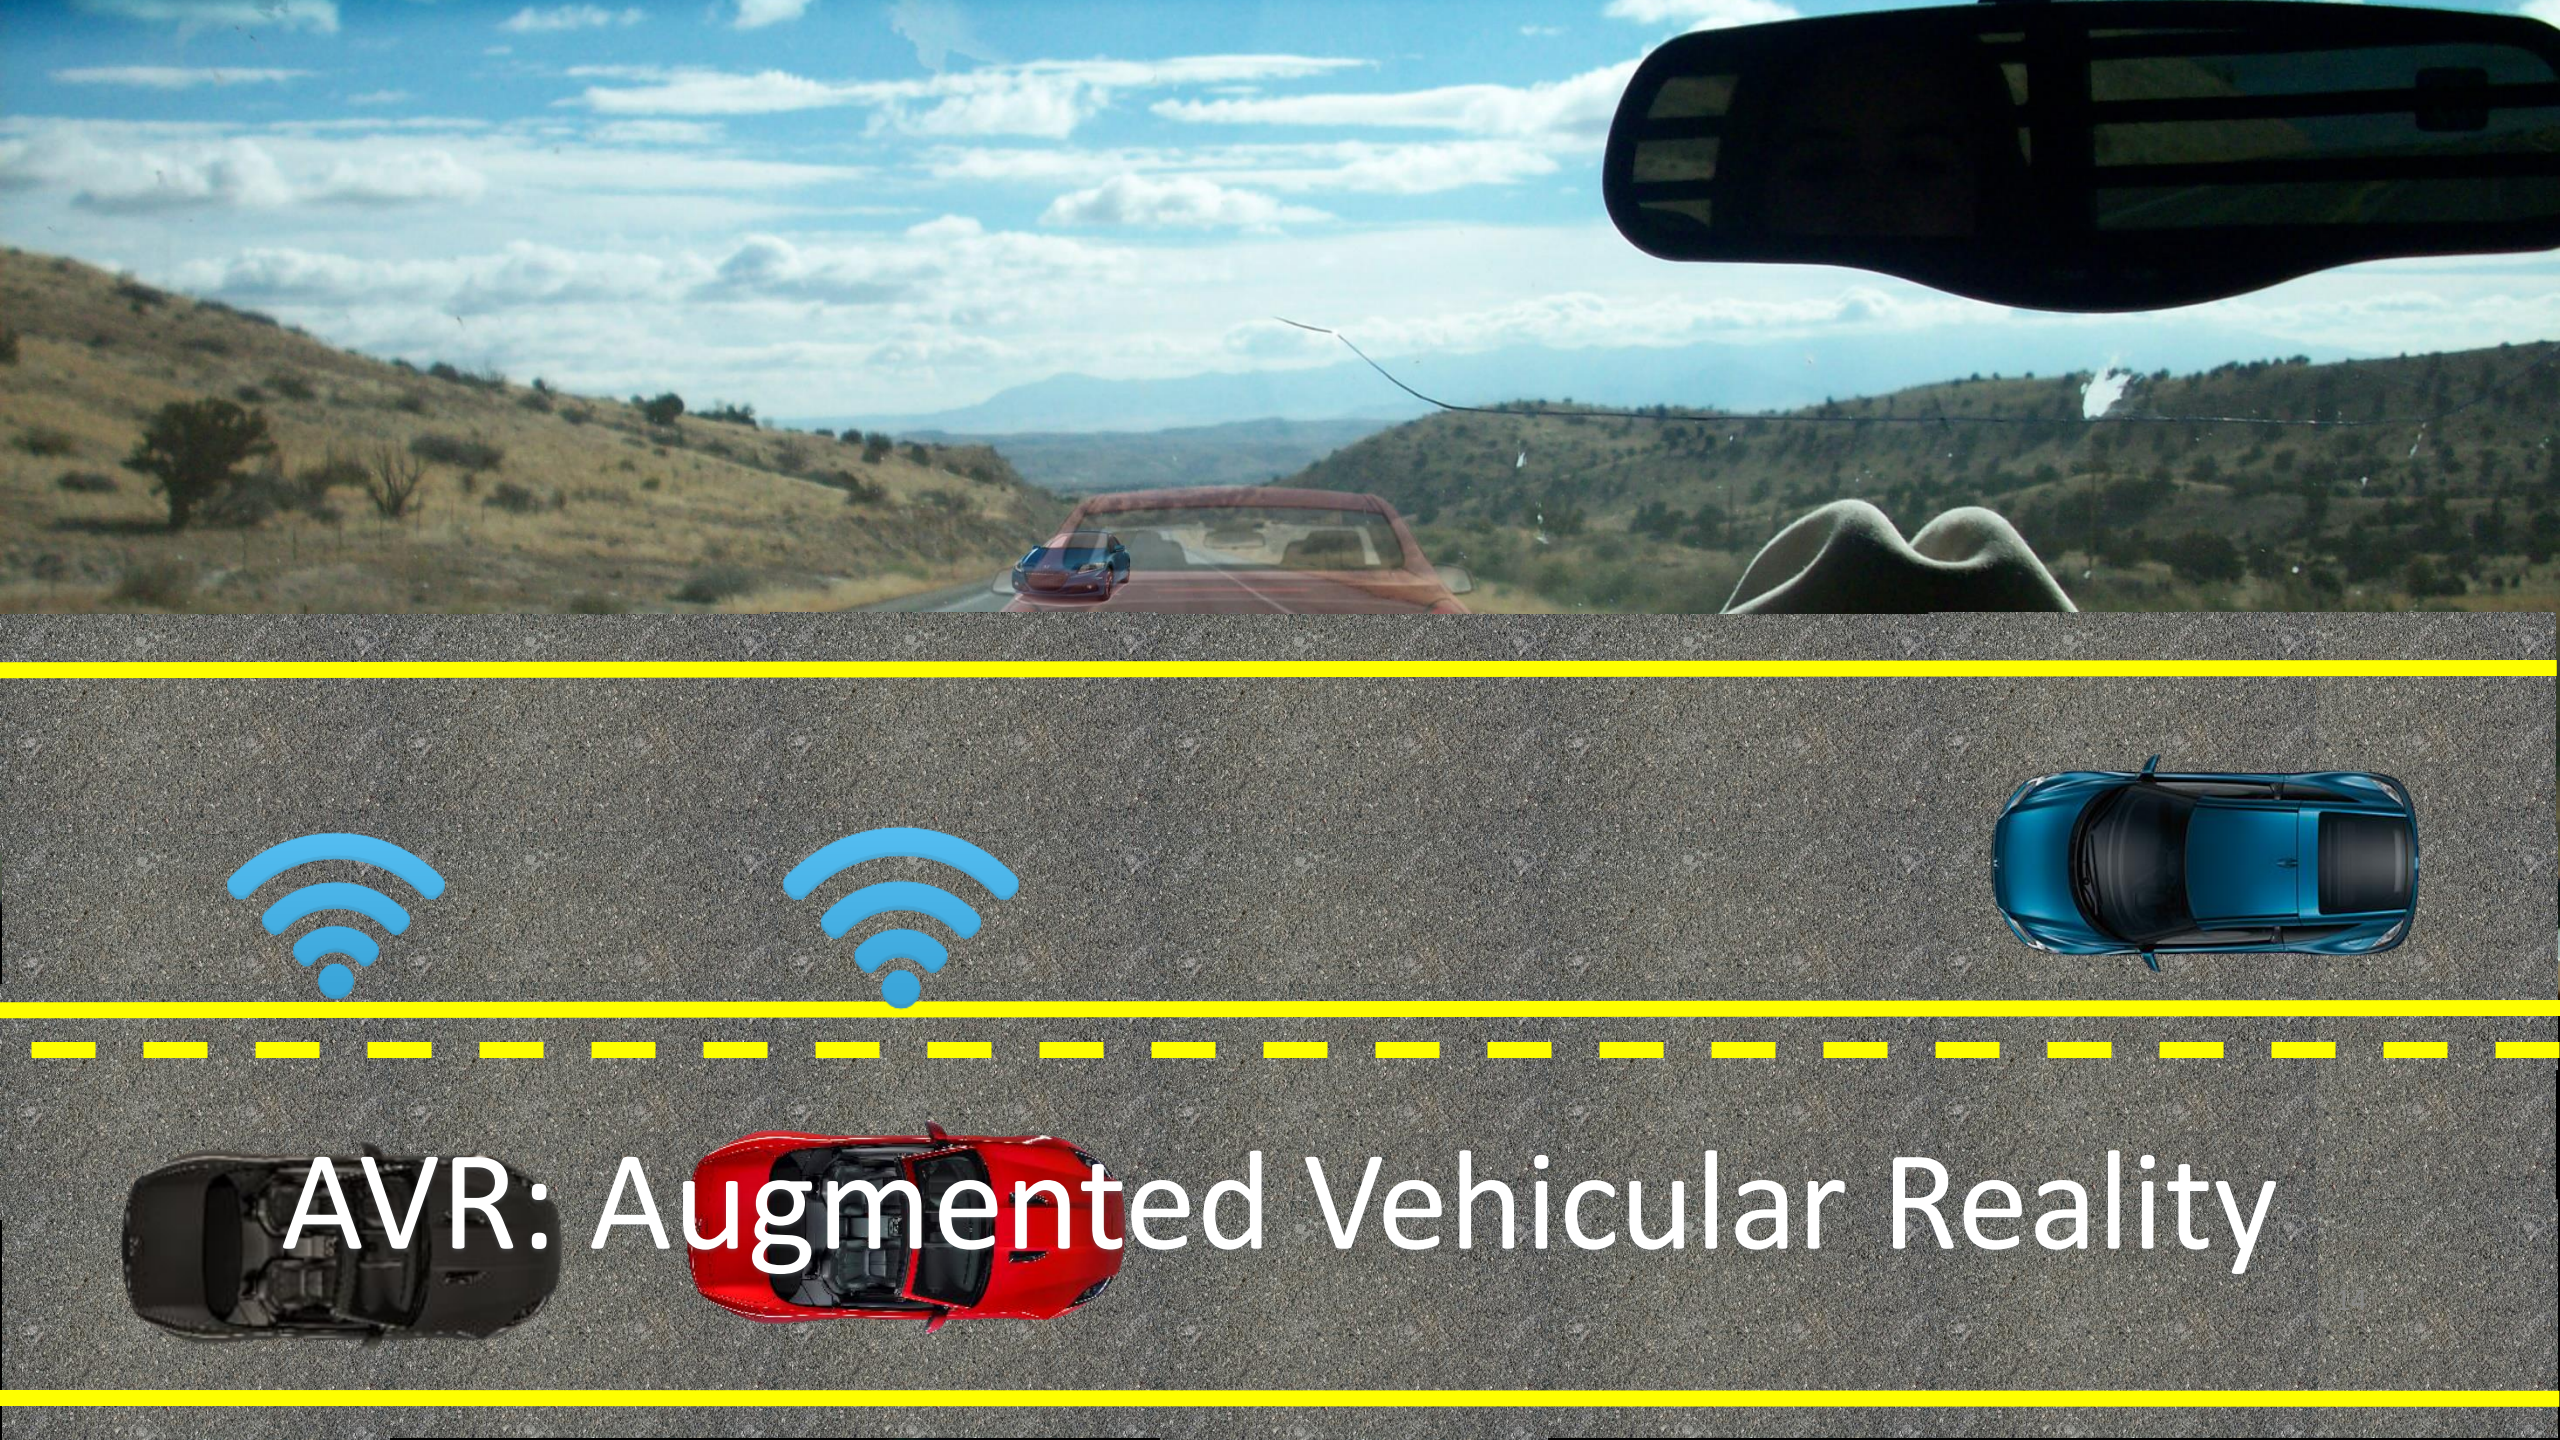
AVR: Augmented Vehicular Reality

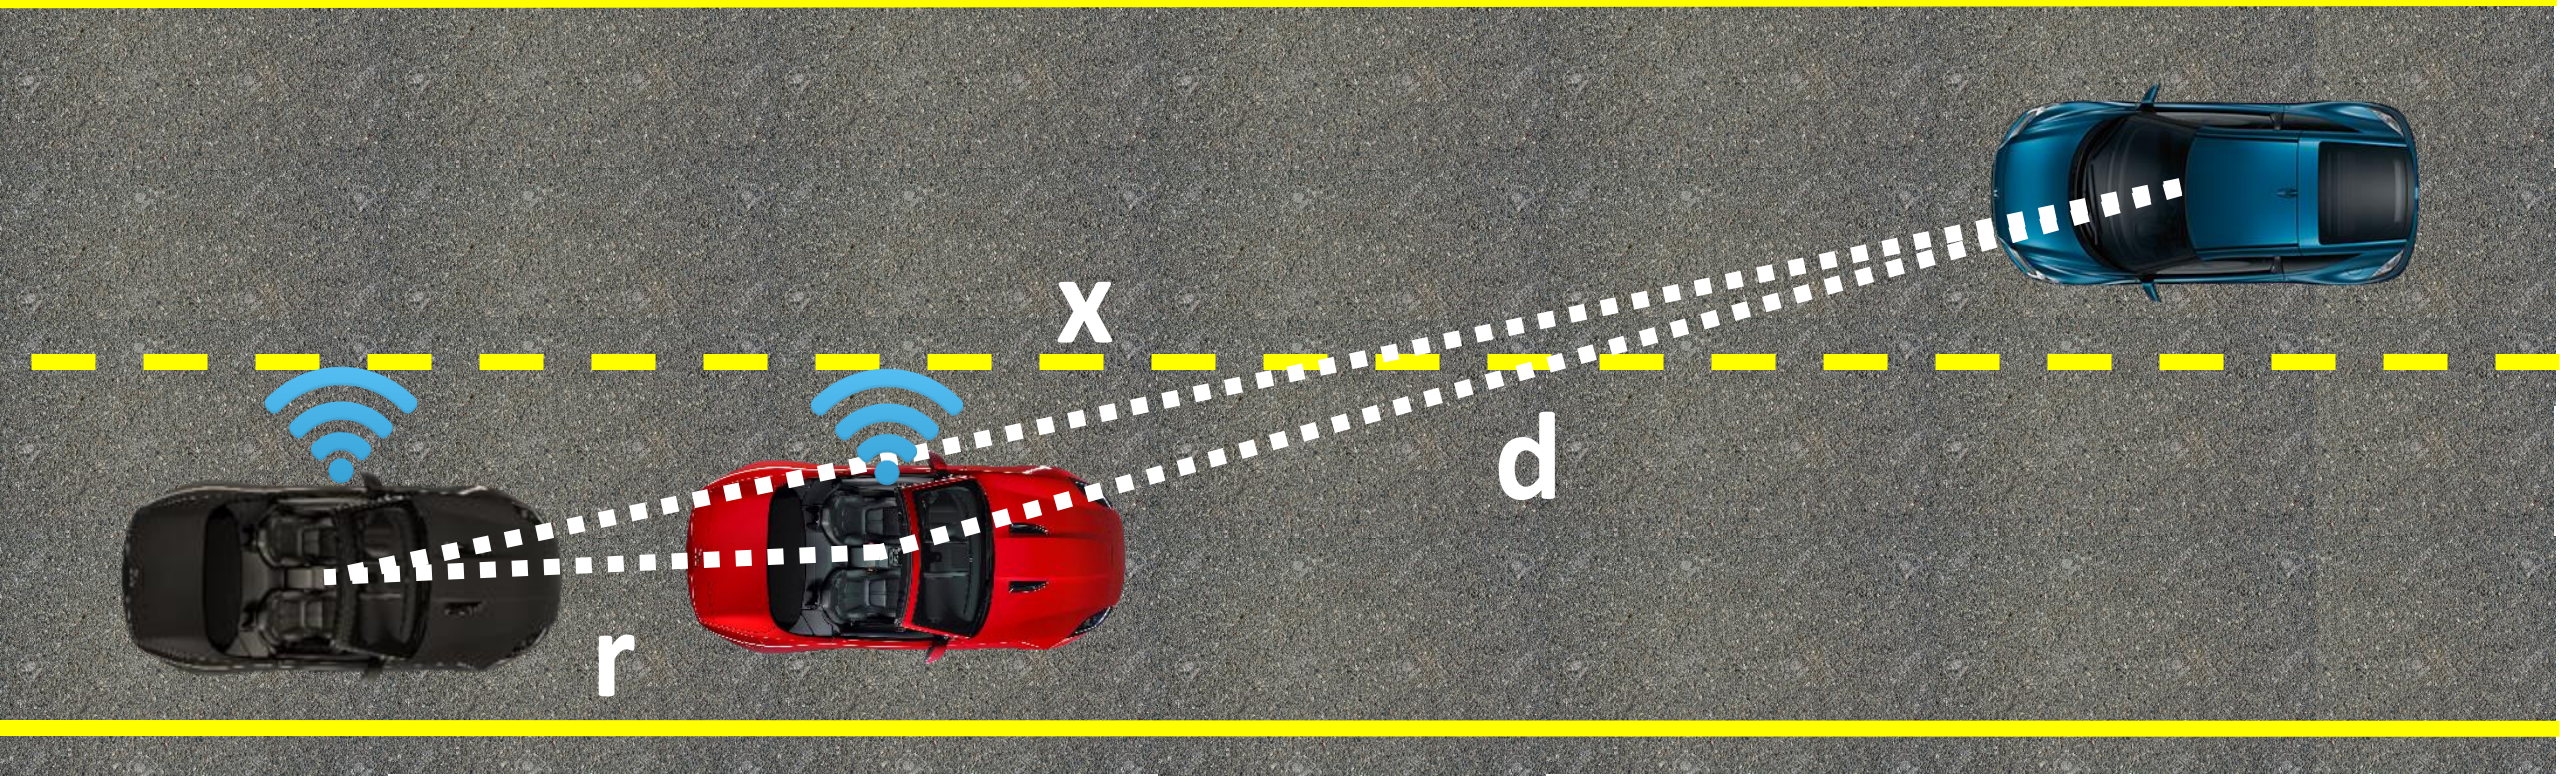

AVR: Challenges

- Accurate Positioning
- Extremely Limited V2V Bandwidth
- Low Latency

AVR Contributions

- Accurate Positioning
  - Relative Positioning using 3D Feature Map
  - Perspective Transformation
- Extremely Limited V2V bandwidth
  - Dynamic Object Extraction
  - Adaptive Motion Vector Transmission
- Low Latency
  - Pipeline Optimization
  - Motion Prediction to Hide Latency

AVR Contributions

- Accurate Positioning
  - Relative Positioning using 3D Feature Map
  - Perspective Transformation
- Extremely Limited V2V bandwidth
  - Dynamic Object Extraction
  - Adaptive Motion Vector Transmission
- Low Latency
  - Pipeline Optimization
  - Motion Prediction to Hide Latency

3D Perception: The Point Cloud

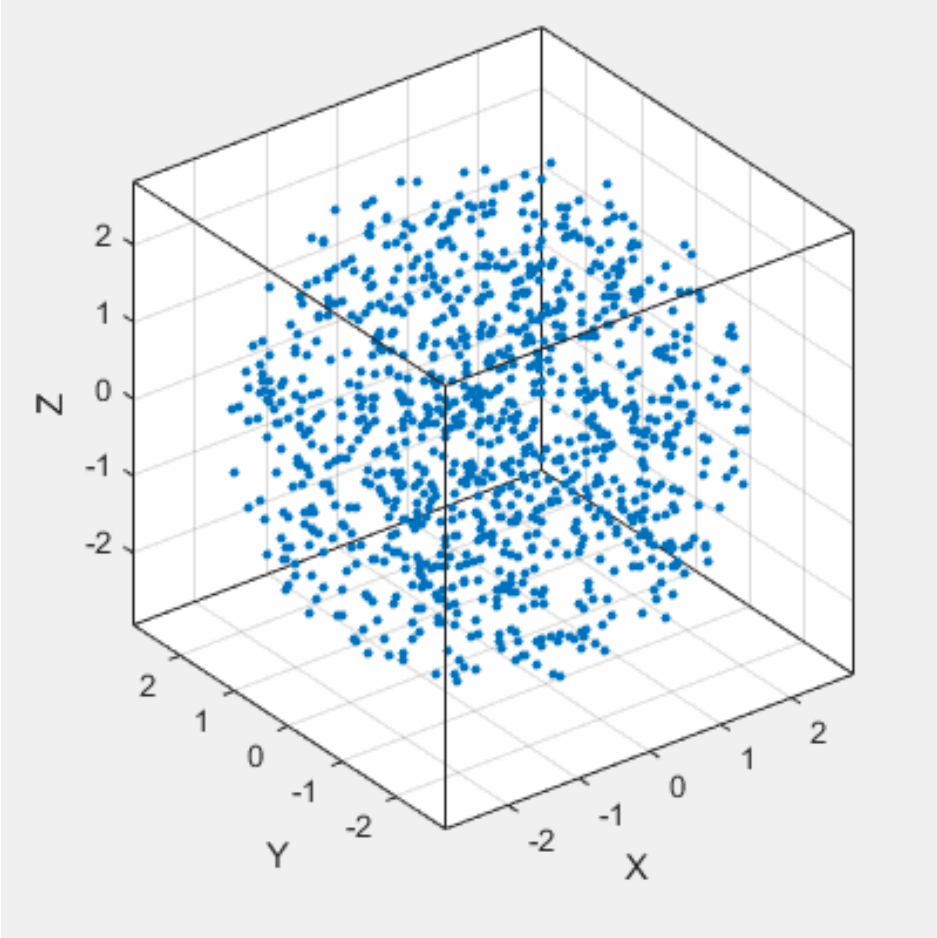

- Voxel
  - Position (x, y, z)
  - Color (R, G, B)
  - Reflection Intensity

Relative Positioning

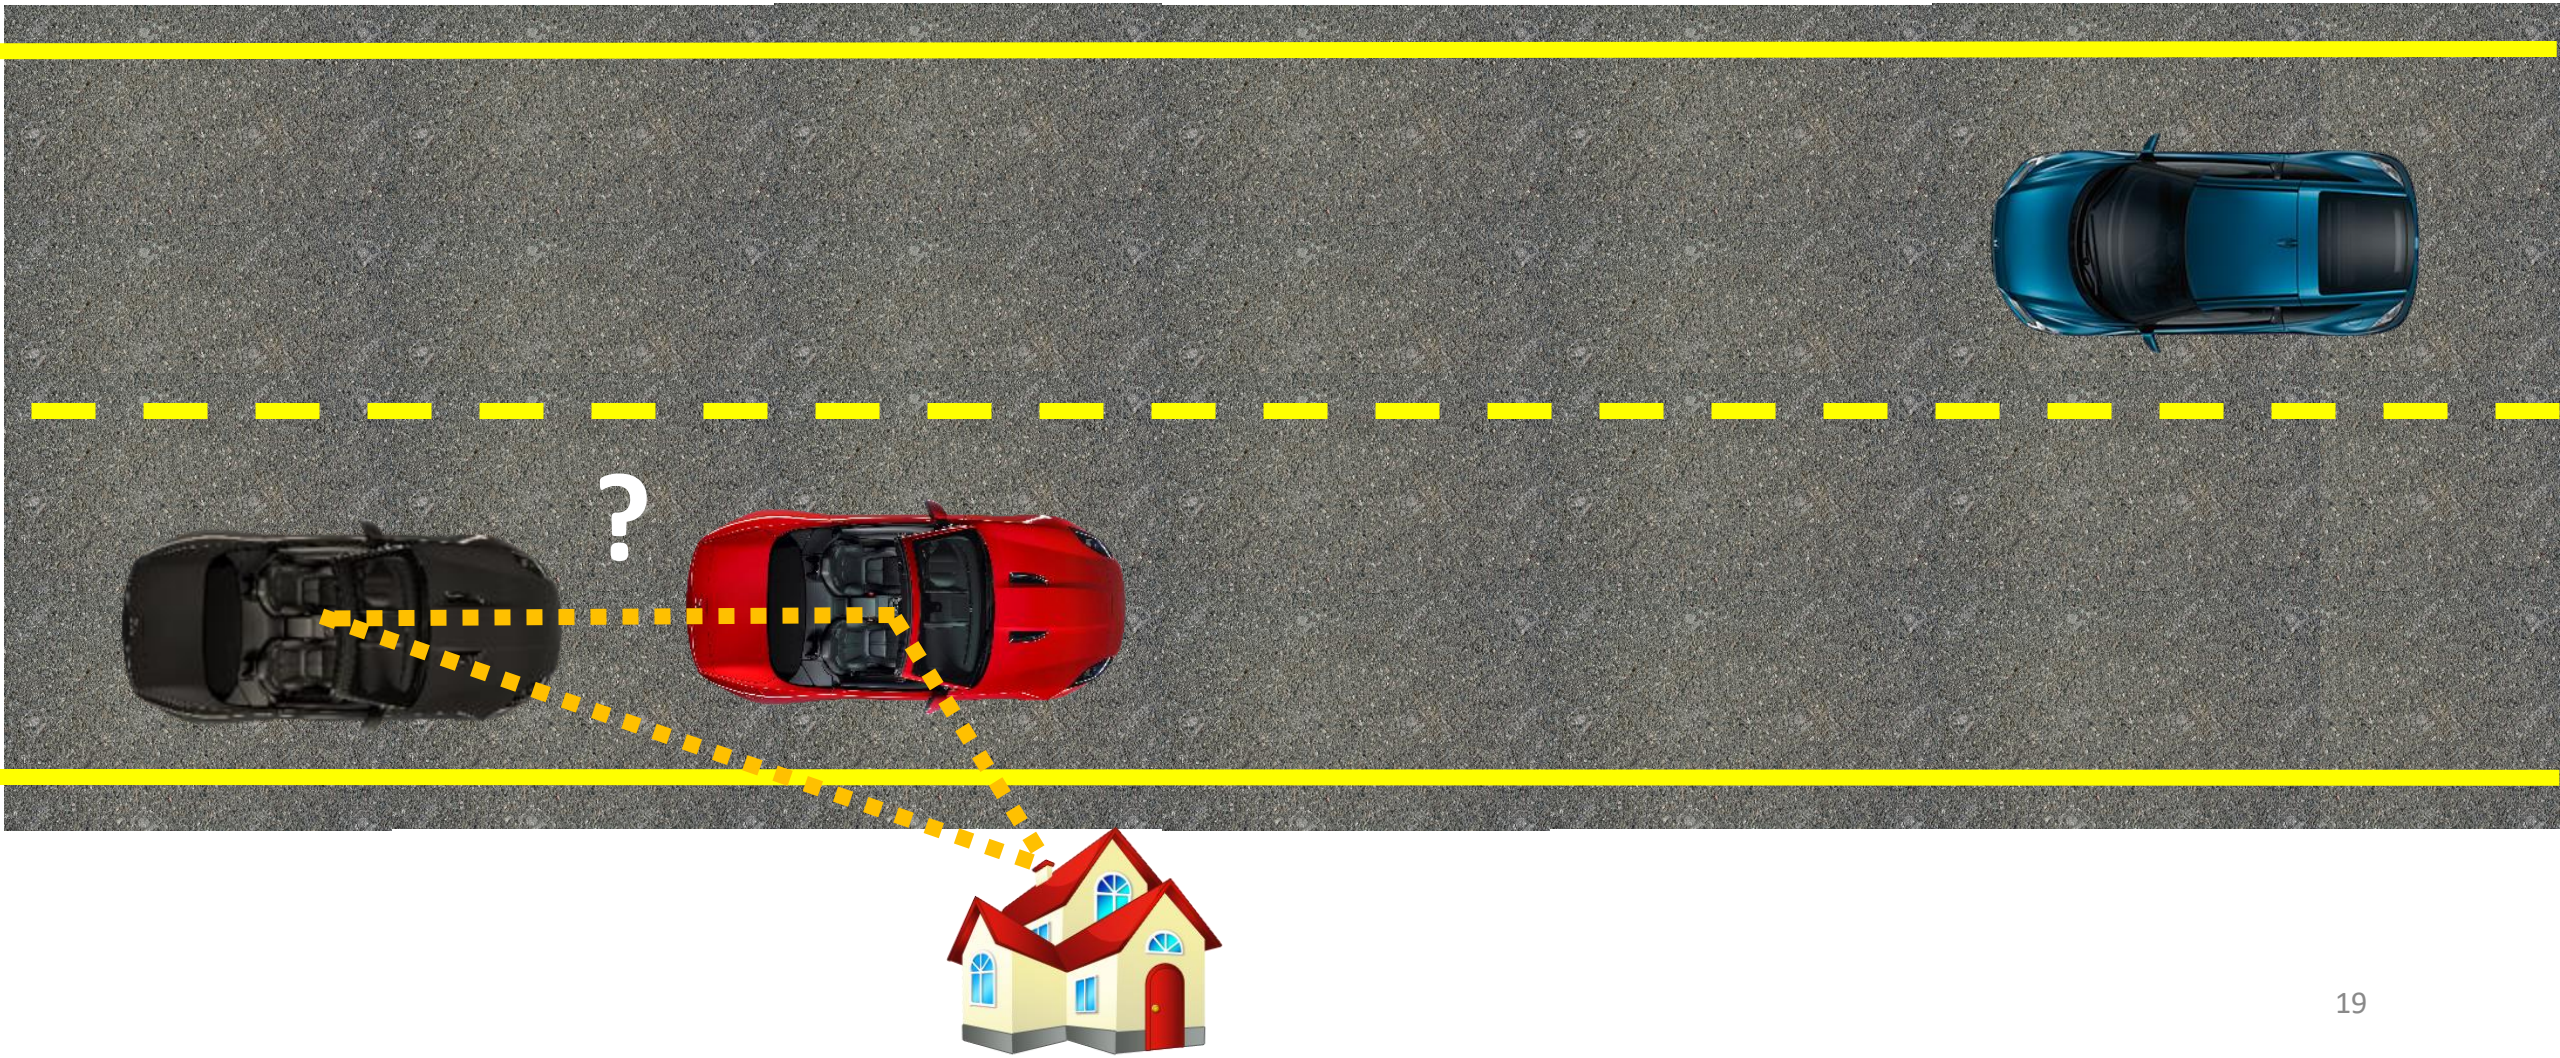

Sparse Features for Localization

Collecting the 3D feature map

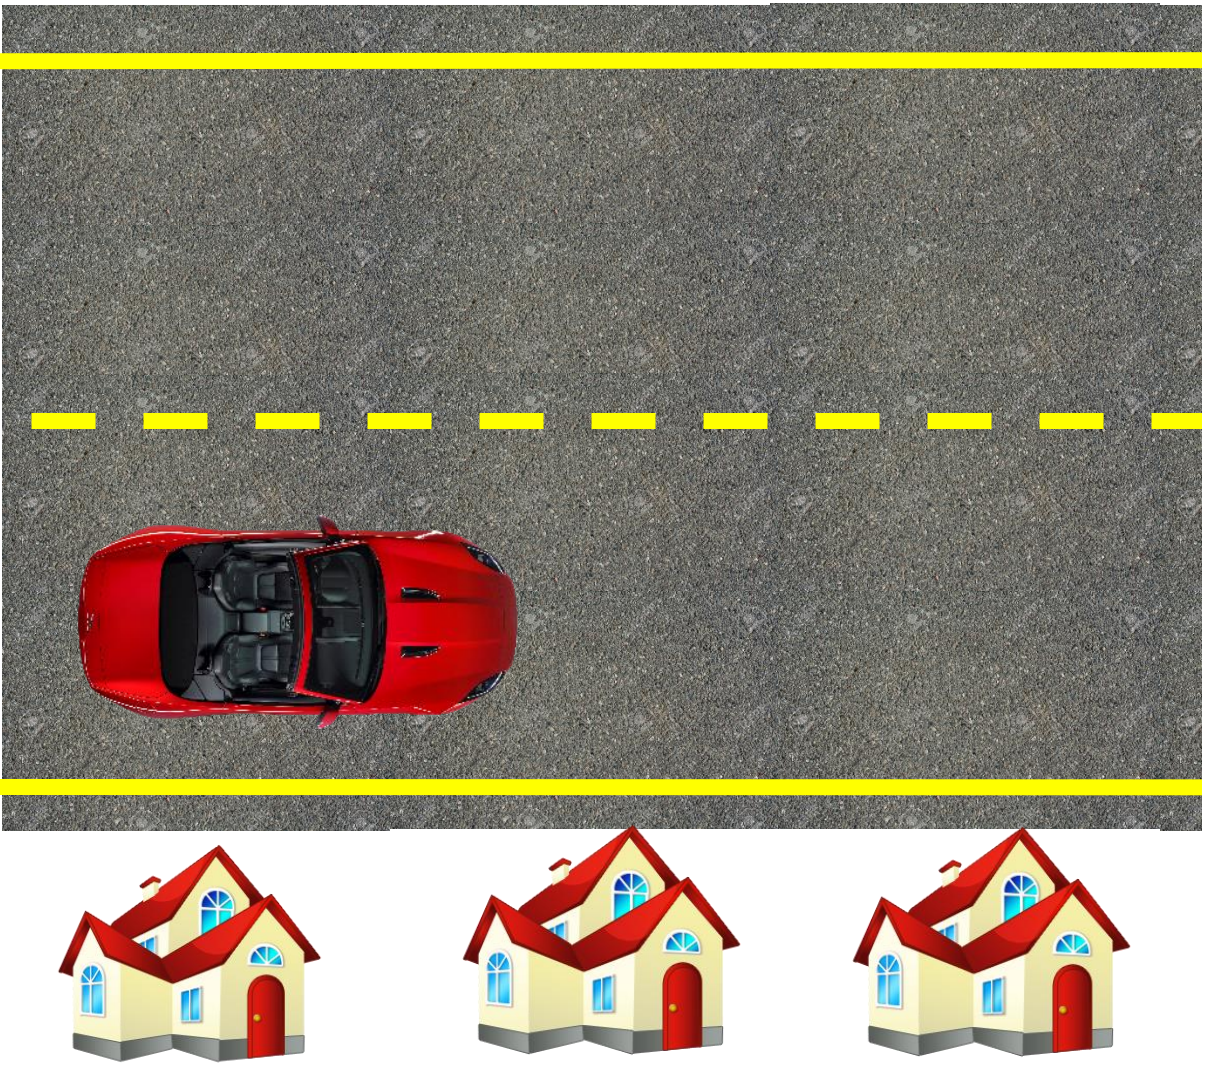

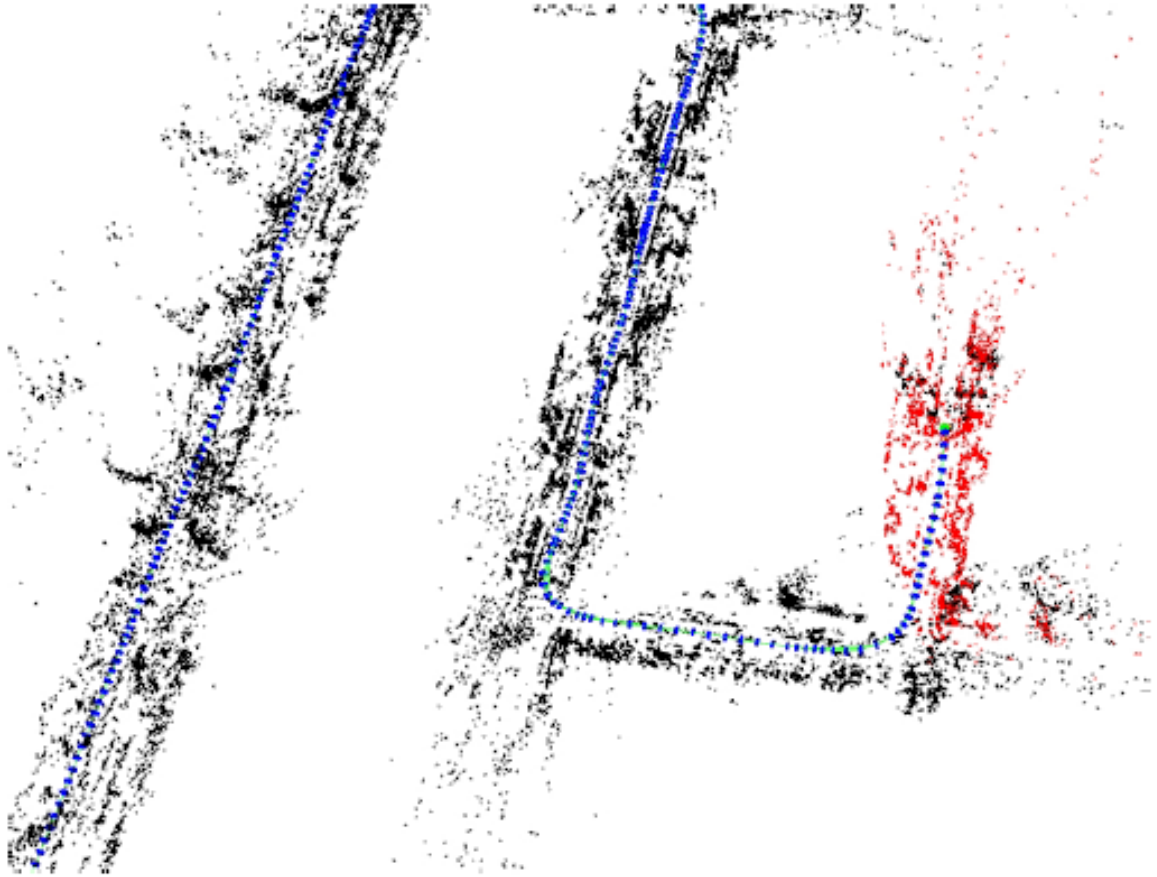

AVR uses SLAM to build a 3D feature map

Localization using sparse 3D map

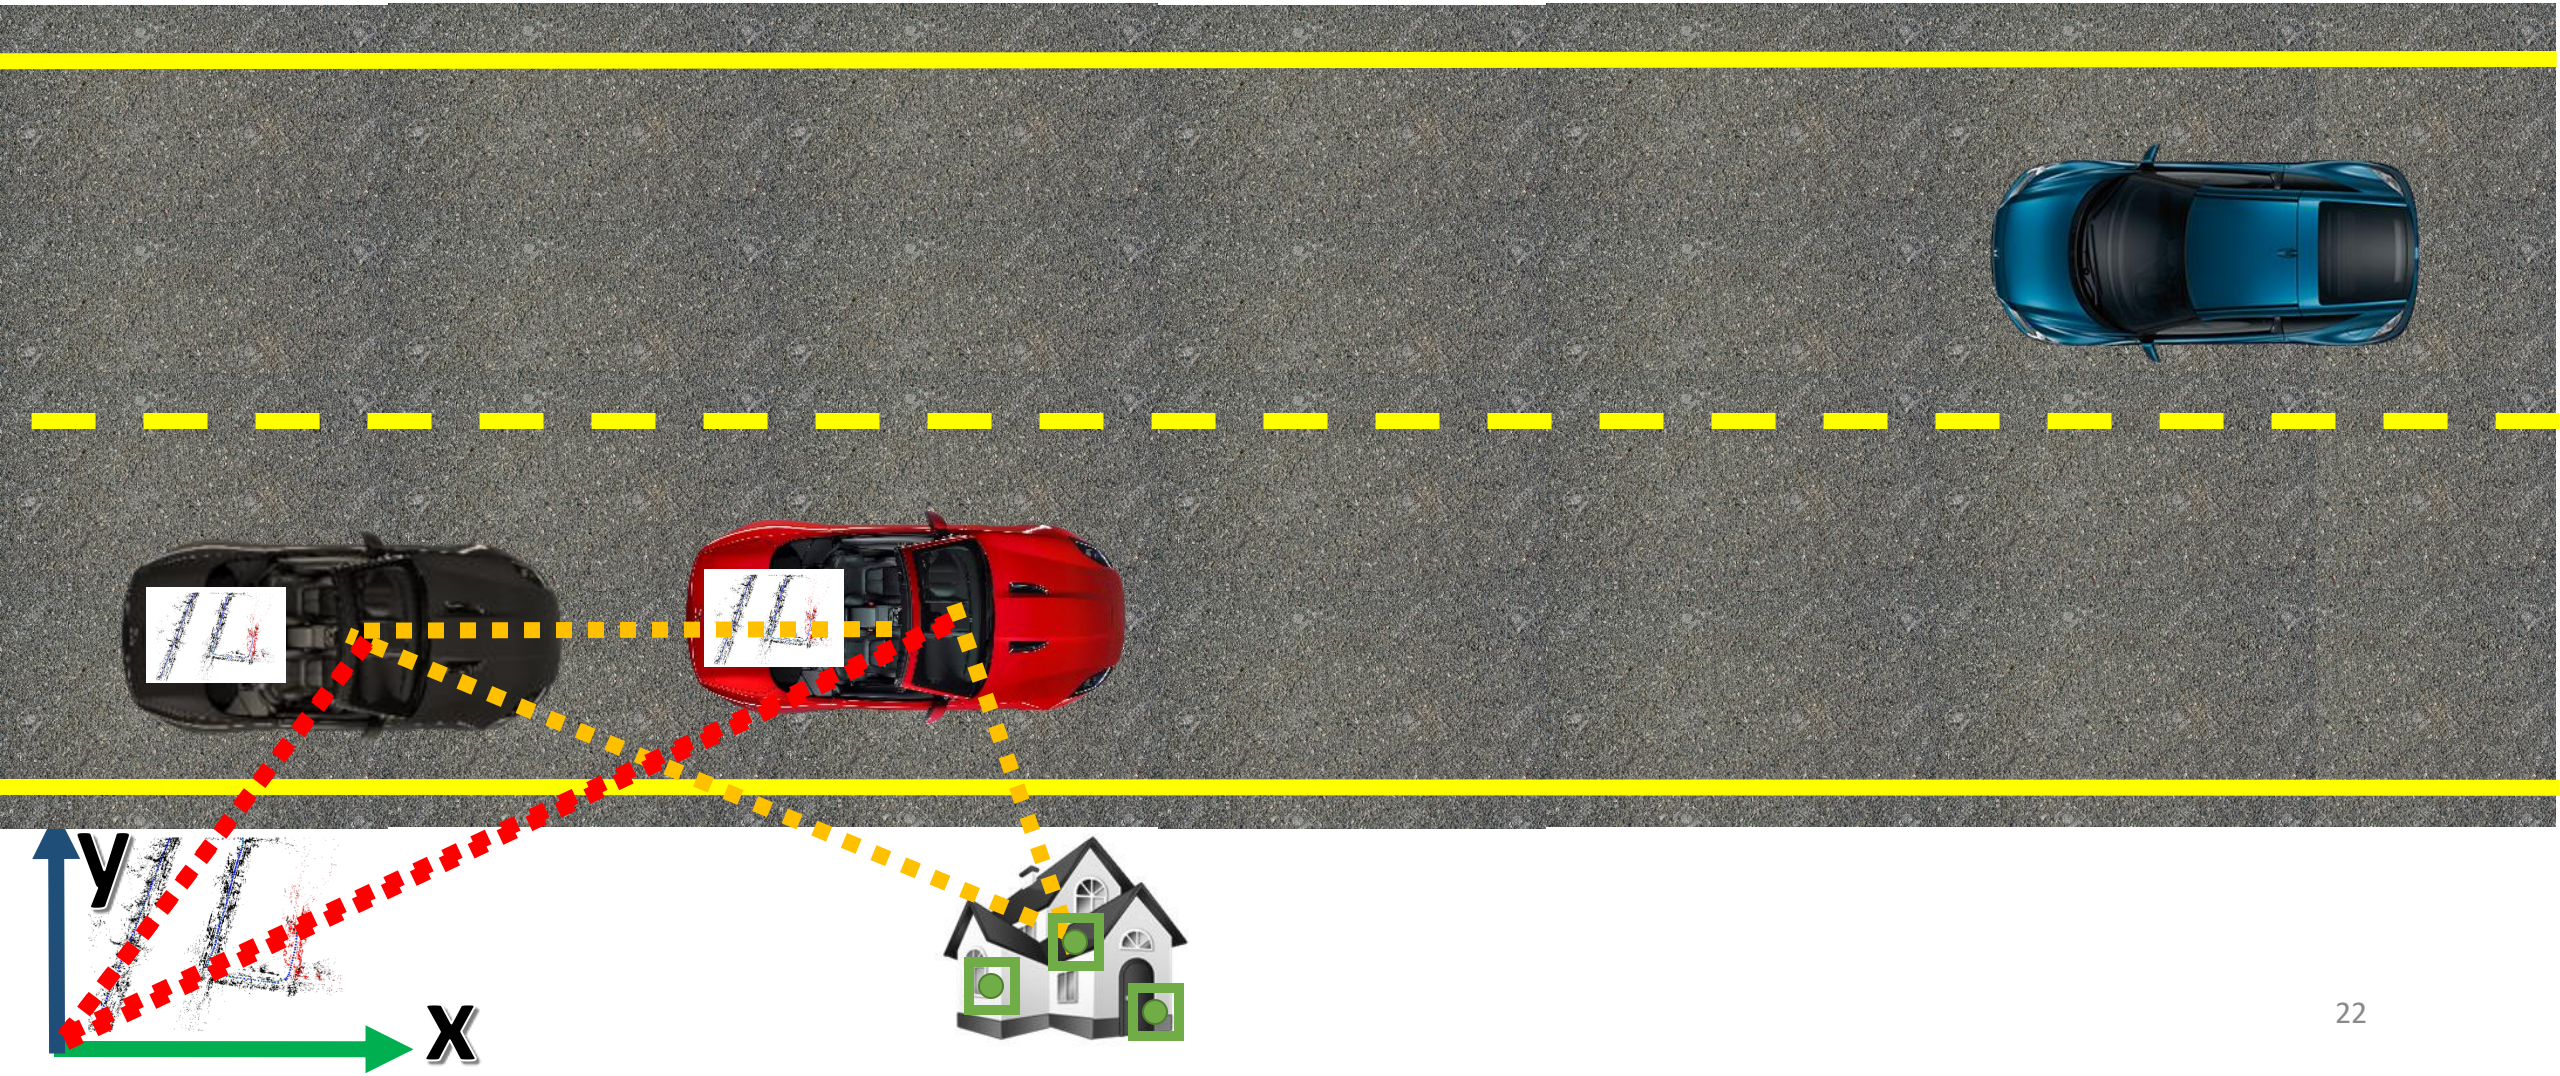

Perspective Transformation

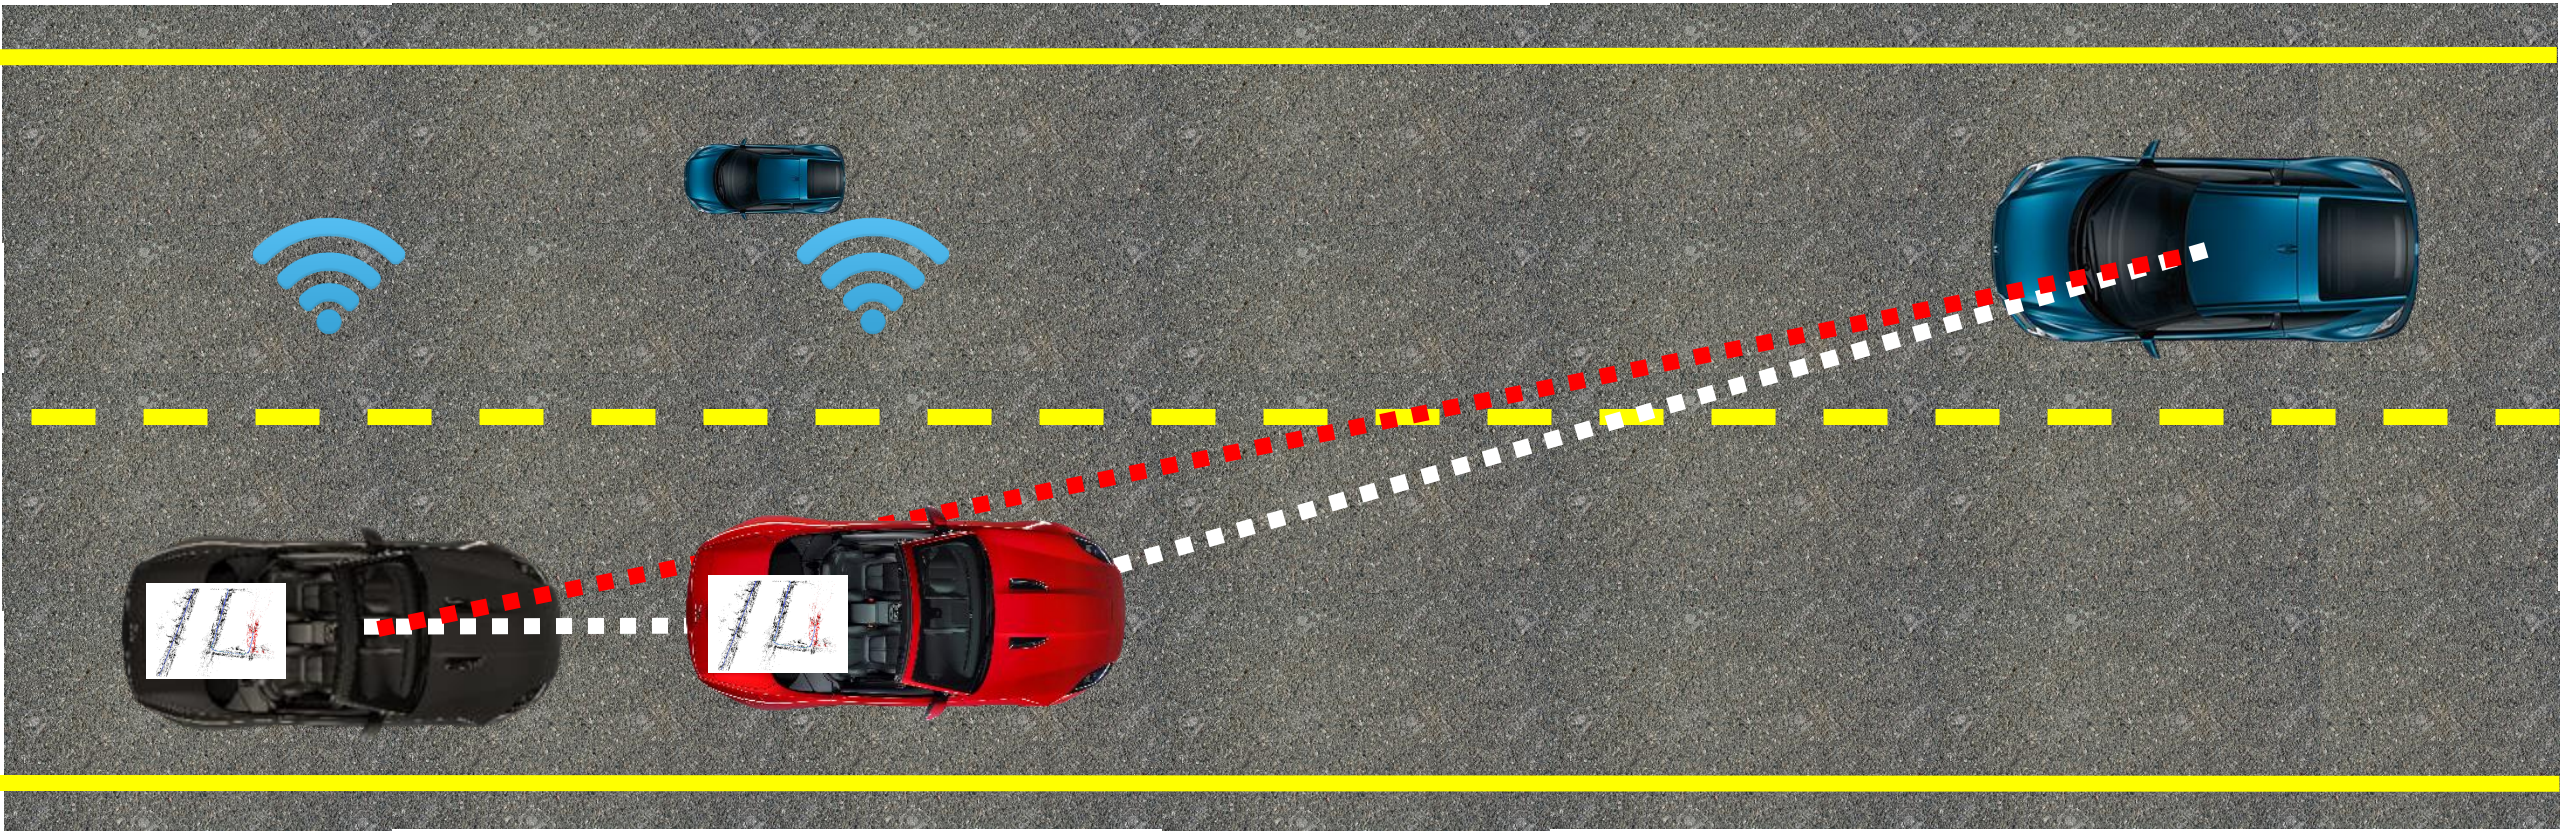

AVR Contributions

- Accurate Positioning
  - Relative Positioning using 3D Feature Map
  - Perspective Transformation
- Extremely Limited V2V bandwidth
  - Dynamic Object Extraction
  - Adaptive Motion Vector Transmission
- Low Latency
  - Pipeline Optimization
  - Motion Prediction to Hide Latency

The Bandwidth

	Data Size	Bandwidth at 30fps
Full Point Cloud	4.91 MB	1178.4 Mbps
A Single Car	0.33 MB	79.2 Mbps

DSRC	802.11g
6 ~ 27 Mbps	6 ~ 54 Mbps

Dynamic Object Detection

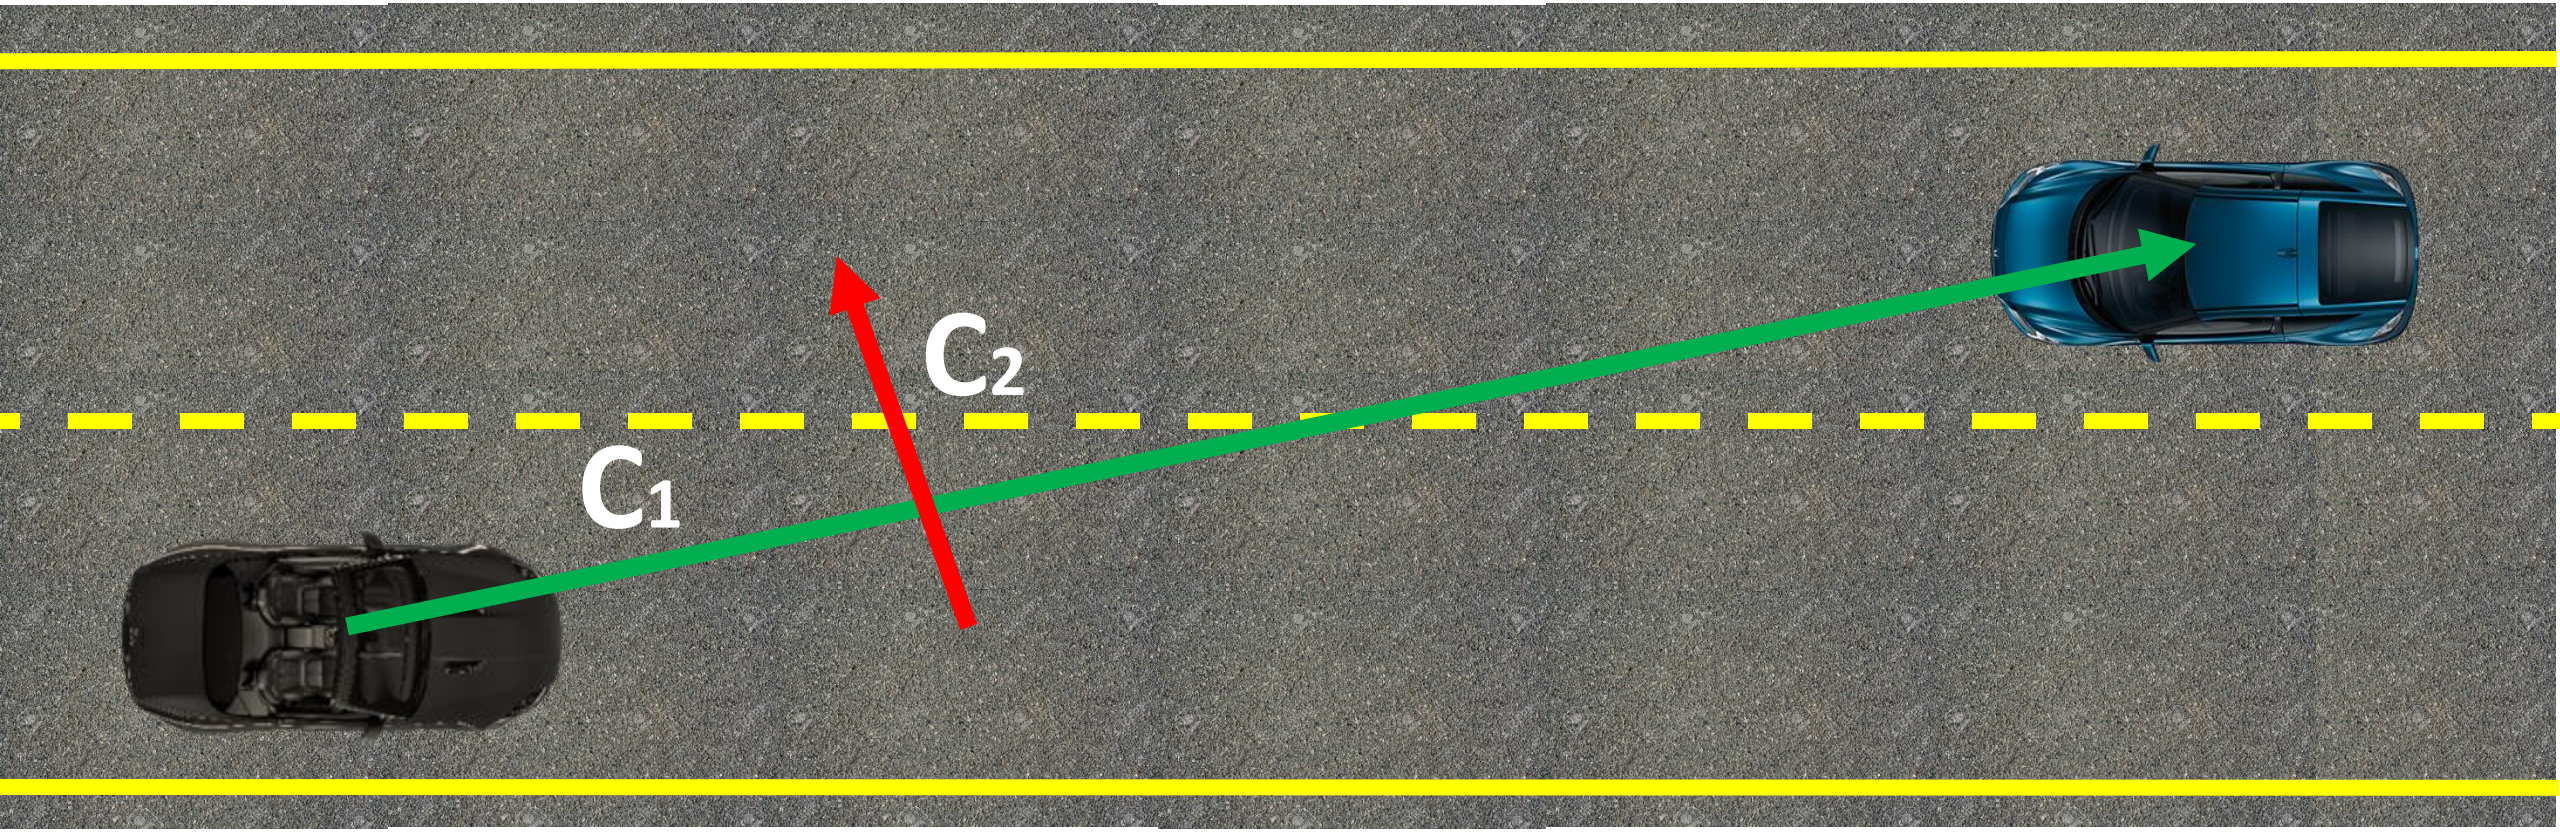

Dynamic Object Detection

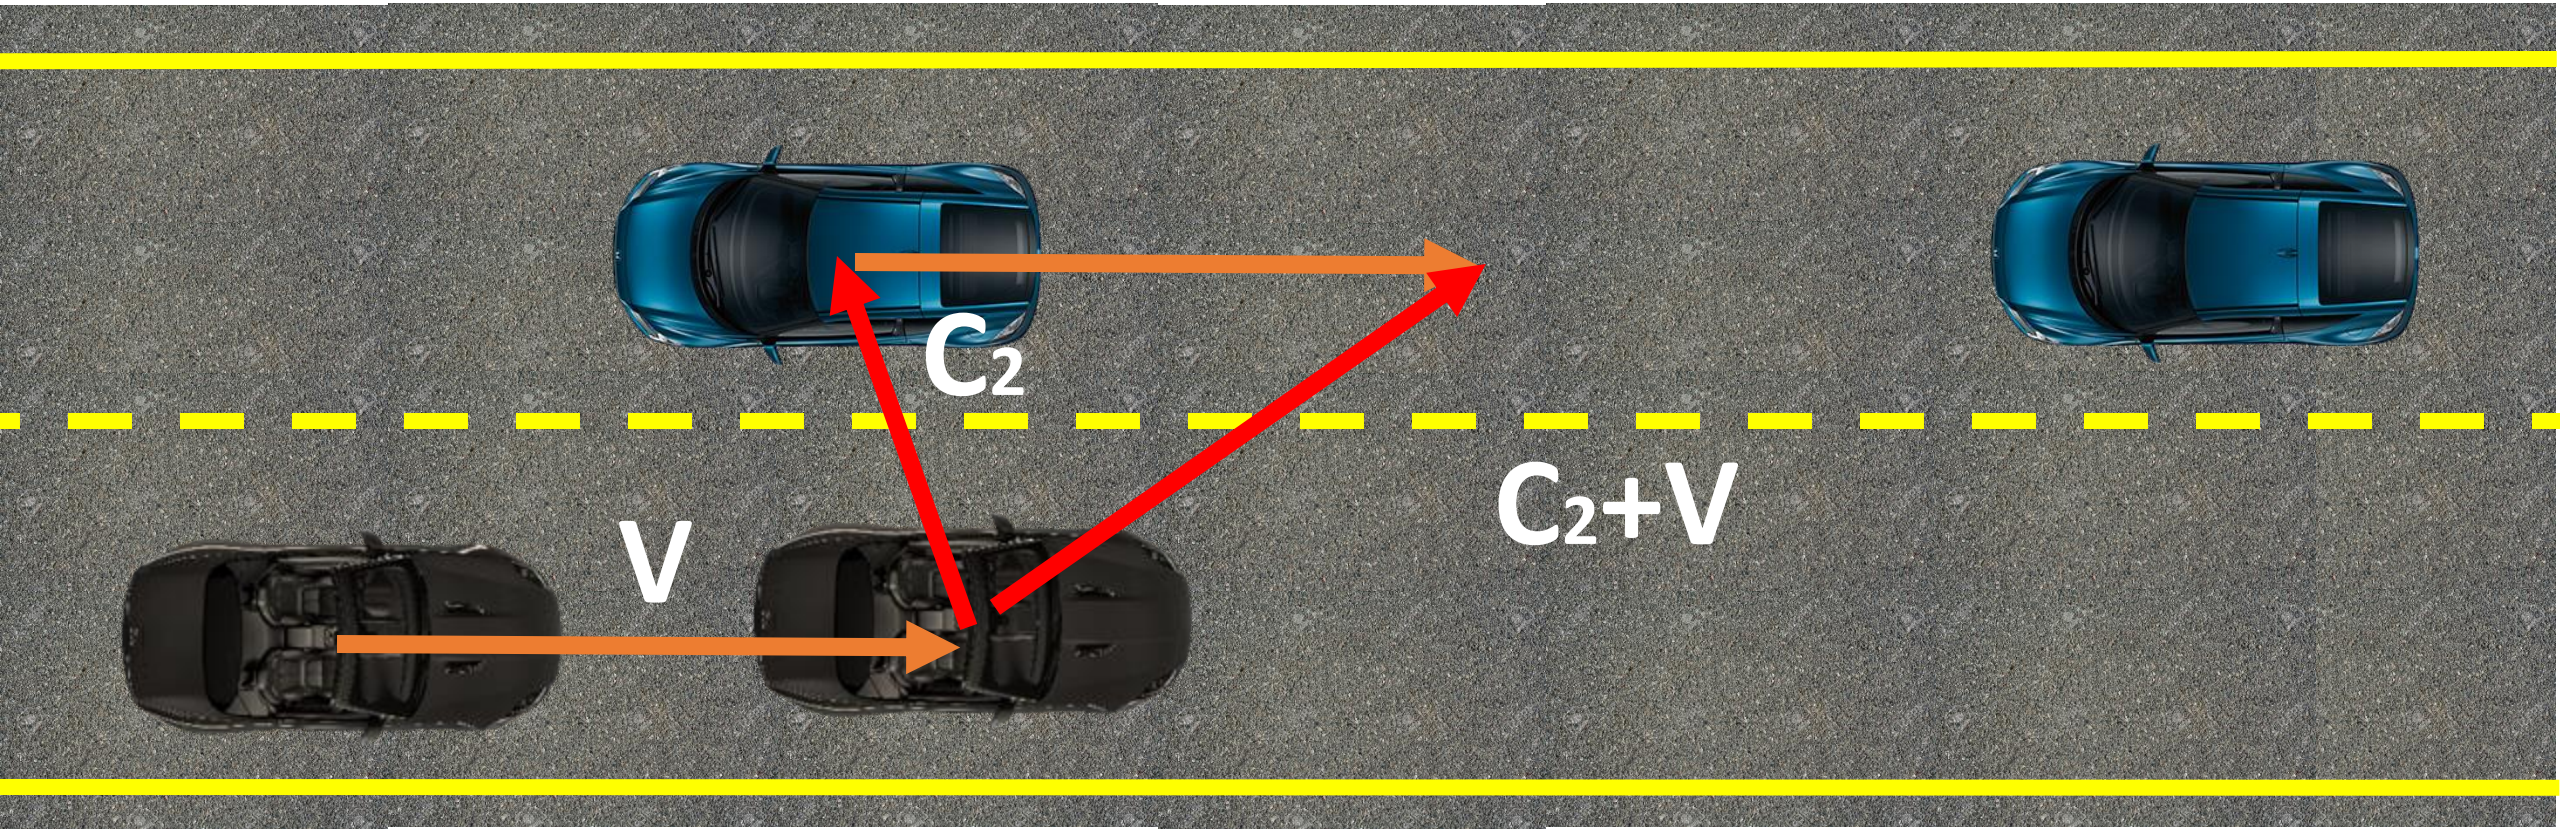

Estimating Absolute Motion

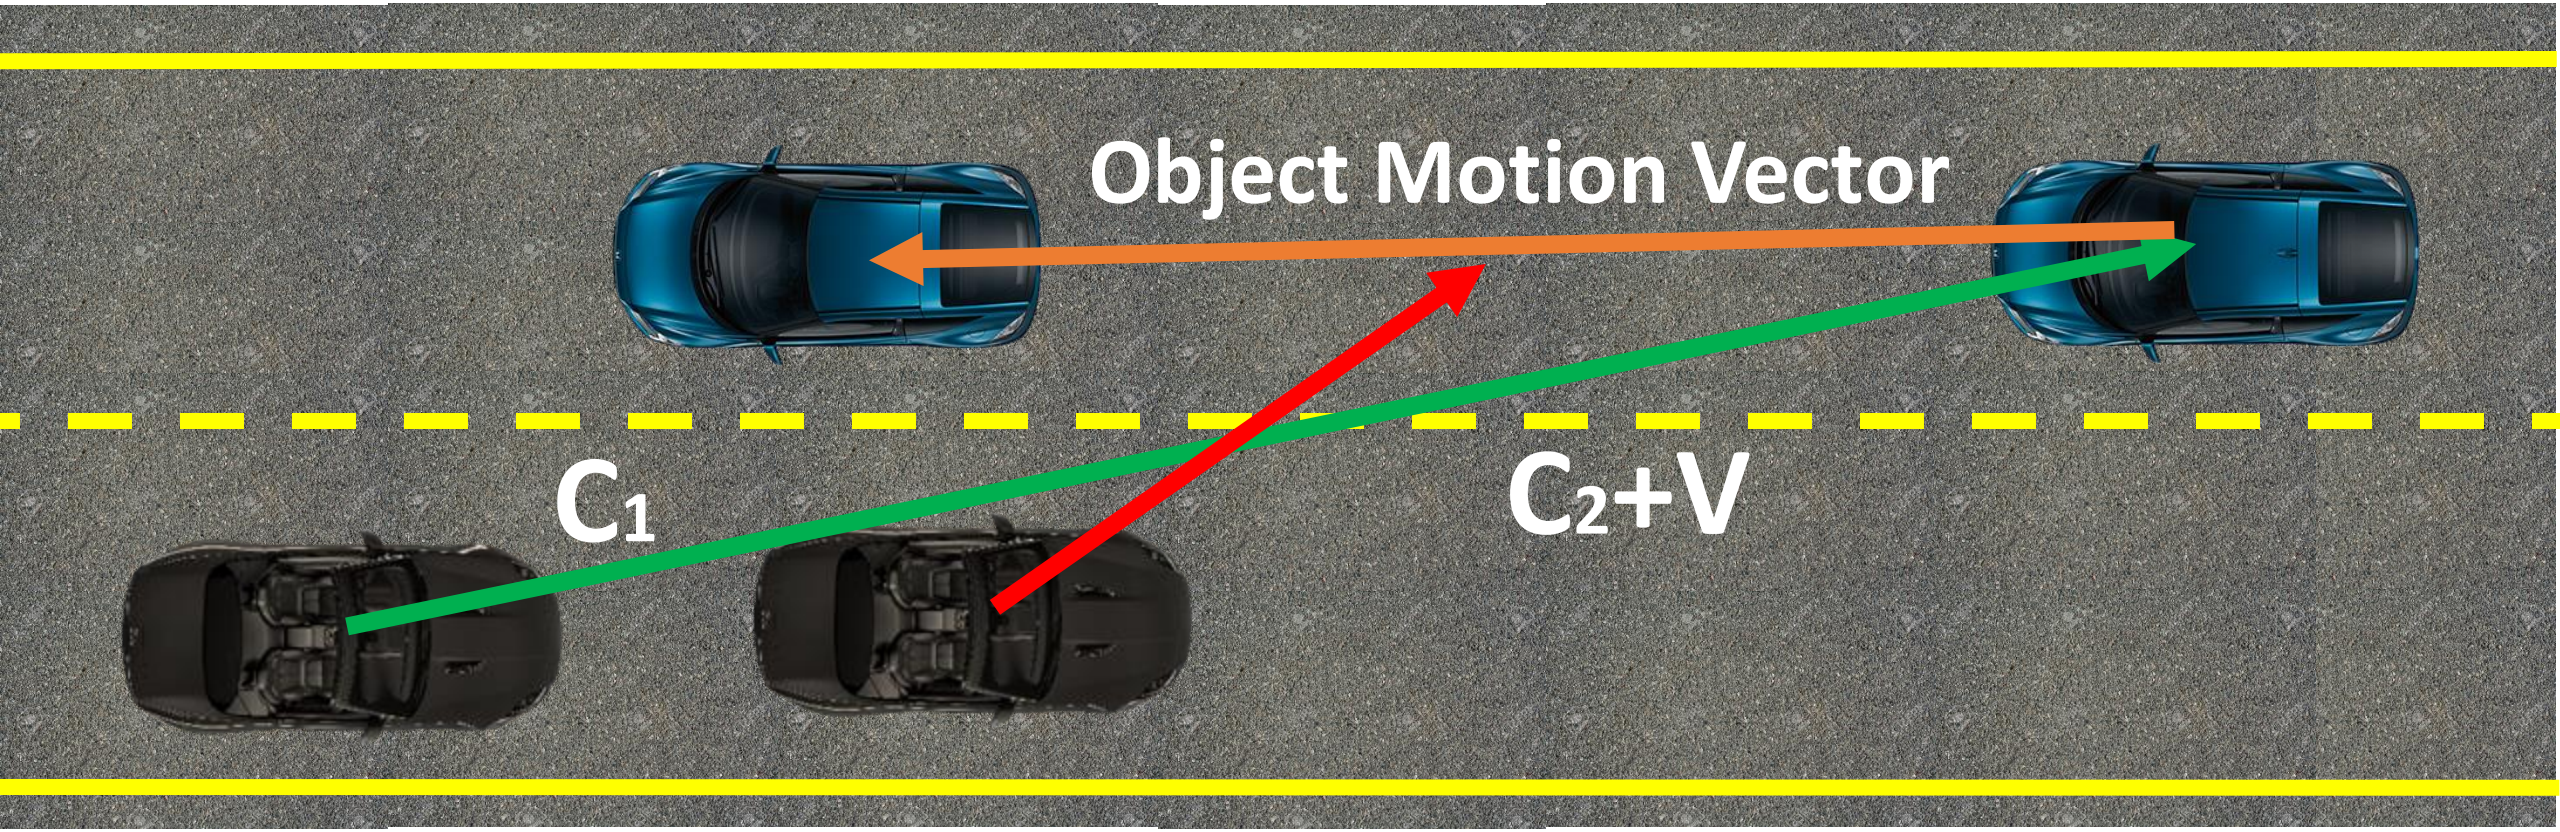

Isolating Dynamics from Statics

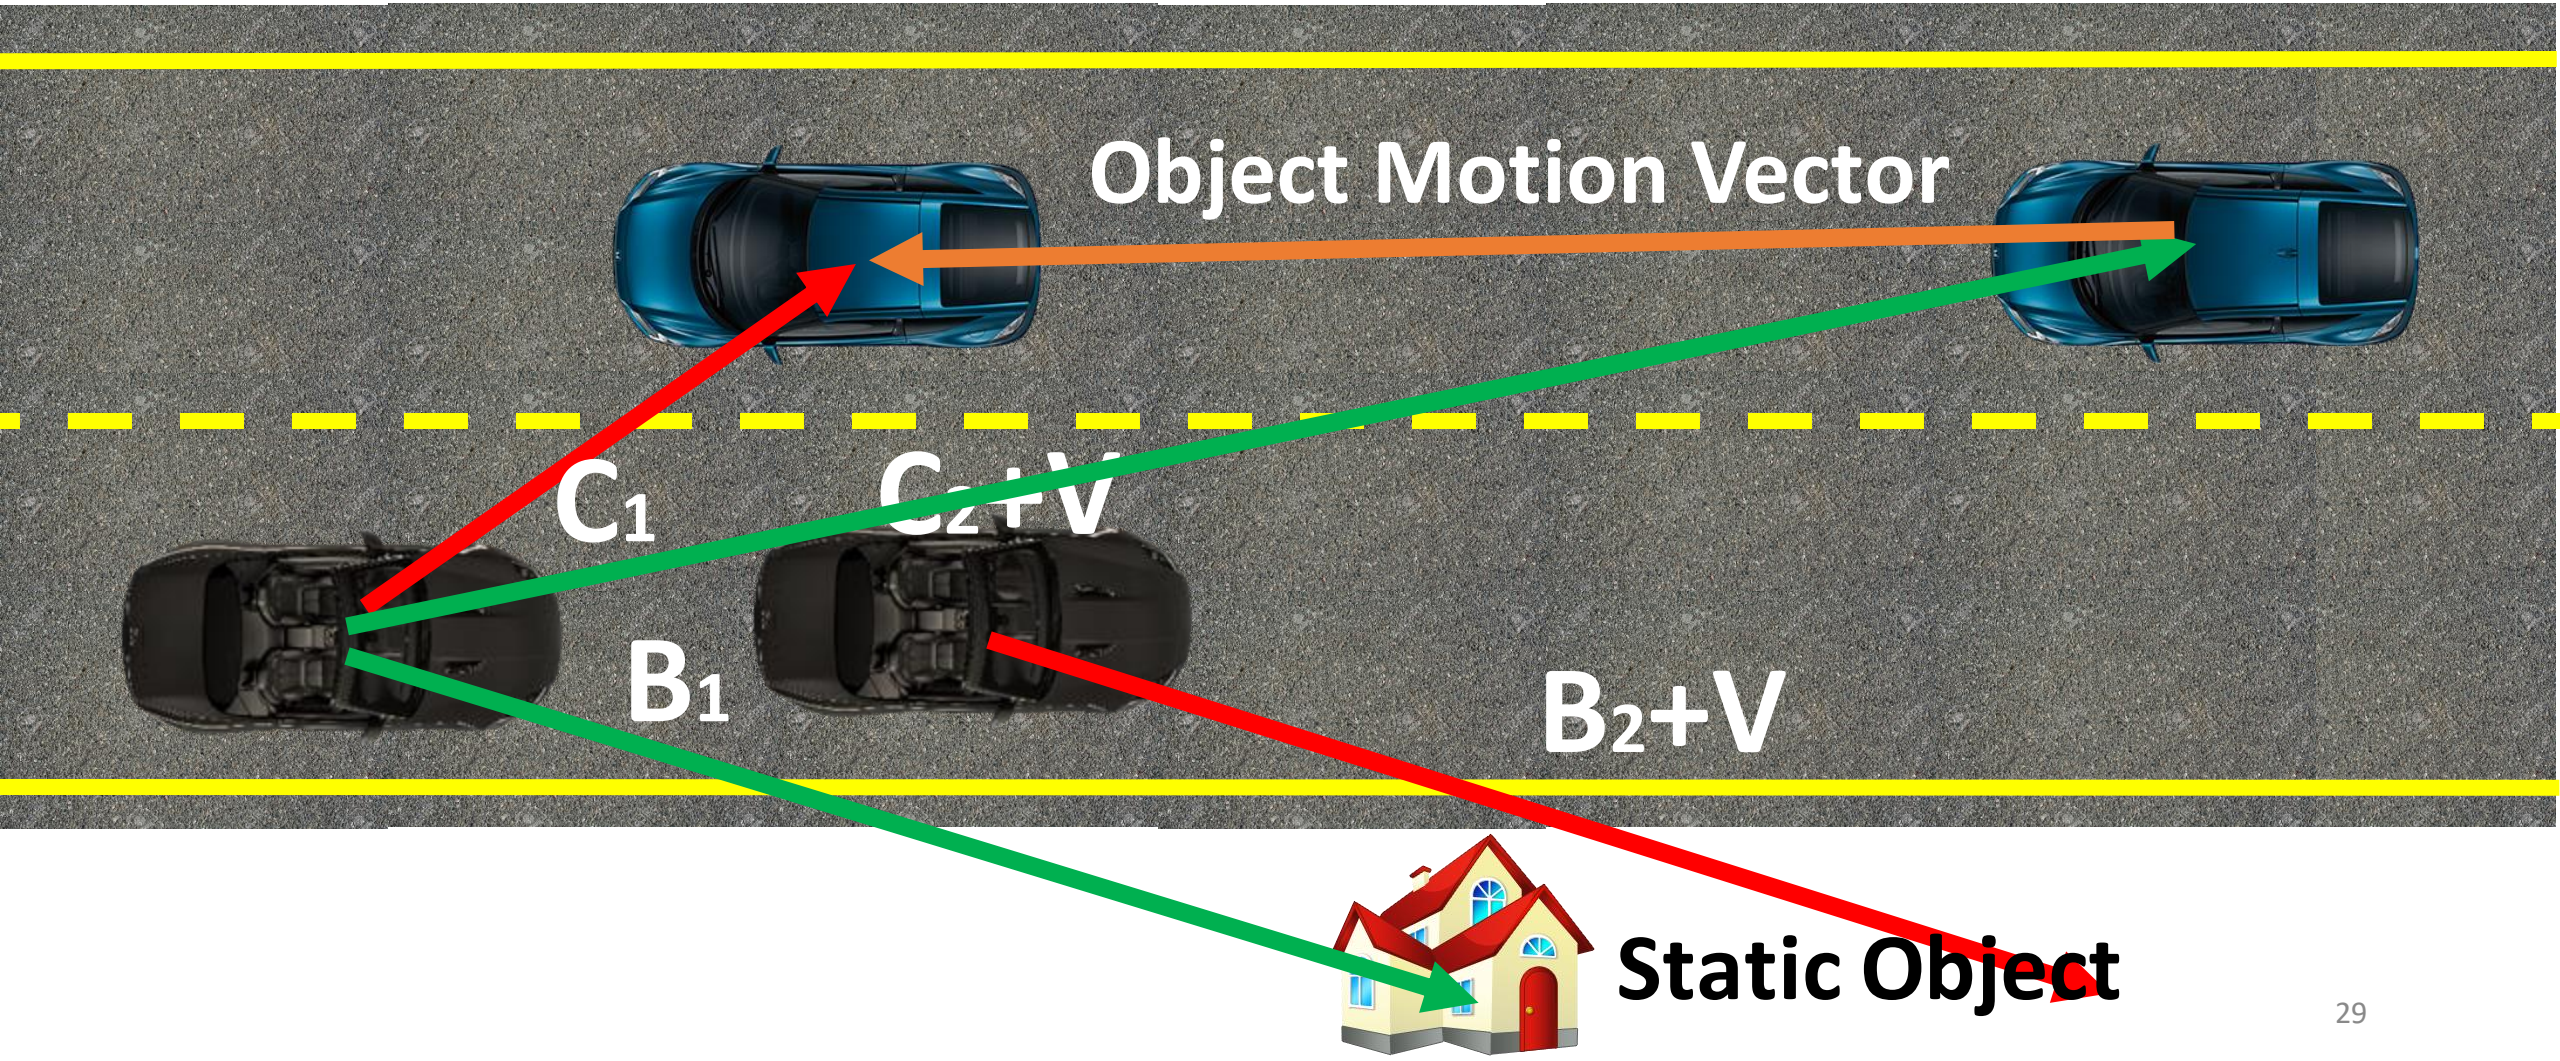

Adaptive Transmission

Transmitting Motion Vectors Adaptively based on Bandwidth Availability.

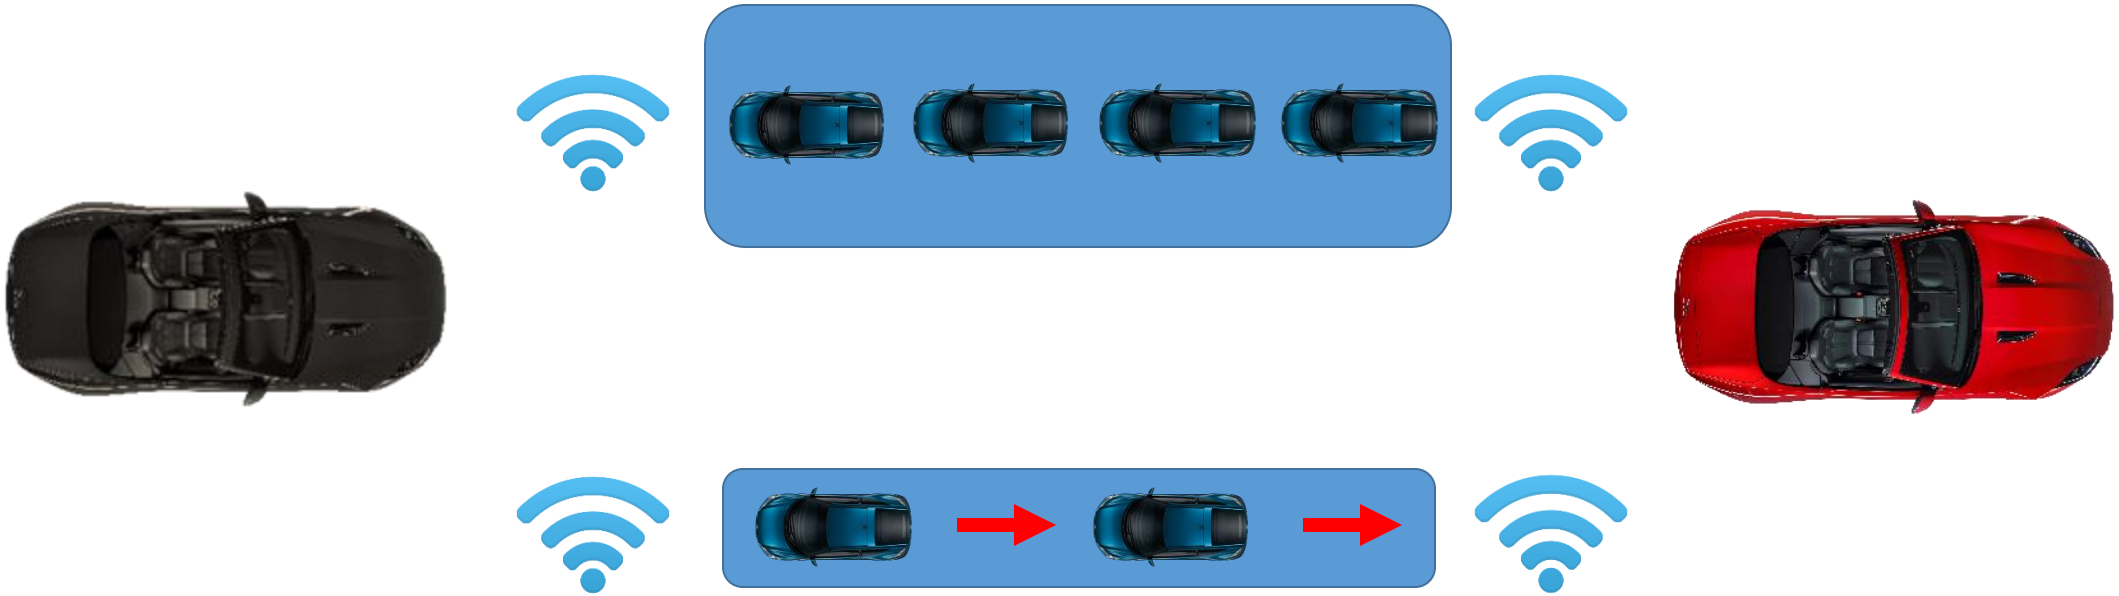

AVR Contributions

- Accurate Positioning
  - Relative Positioning using 3D Feature Map
  - Perspective Transformation
- Extremely Limited V2V bandwidth
  - Dynamic Object Extraction
  - Adaptive Motion Vector Transmission
- Low Latency
  - Pipeline Optimization
  - Motion Prediction to Hide Latency

Pipeline Optimization

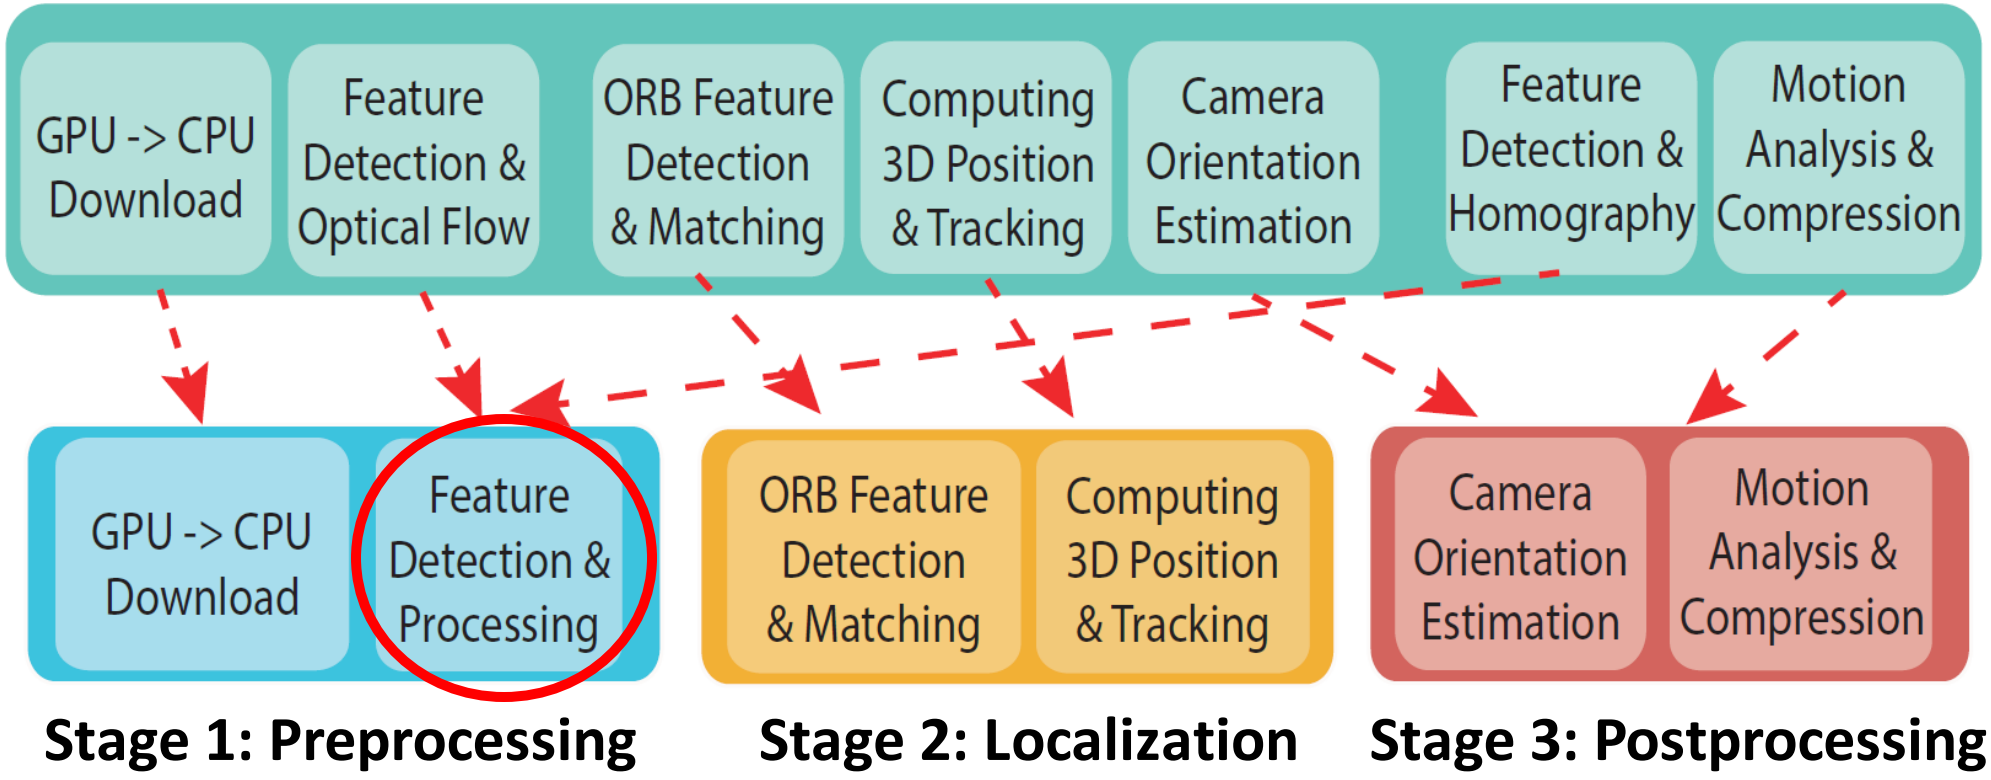

Decouple computation and reuse intermediate results to cut latency.
Distribute computation evenly to 3 stages for high throughput.

AVR Evaluation

- The Benefits of AVR in ADAS and Autonomous Driving
- End-to-end Performance
- Reconstruction Accuracy
- Bandwidth / Accuracy Tradeoff
- The Impact of Delay
- Motion Estimation Accuracy
- Throughput And Latency

AVR Evaluation

- The Benefits of AVR in ADAS and Autonomous Driving
- End-to-end Performance
- Reconstruction Accuracy
- Bandwidth / Accuracy Tradeoff
- The Impact of Delay
- Motion Estimation
- Throughput And Latency

AVR Setup

Alienware Laptop
- I7 CPU, 4.4GHz
- 16GB DDR4 RAM
- Nvidia 1080p GPU, 2560 CUDA cores.

ZED Stereo Camera

TP-Link Talon 7200
- 802.11 g/n/ac/ad
- Wireless Distributed System (WDS)

AVR Setup

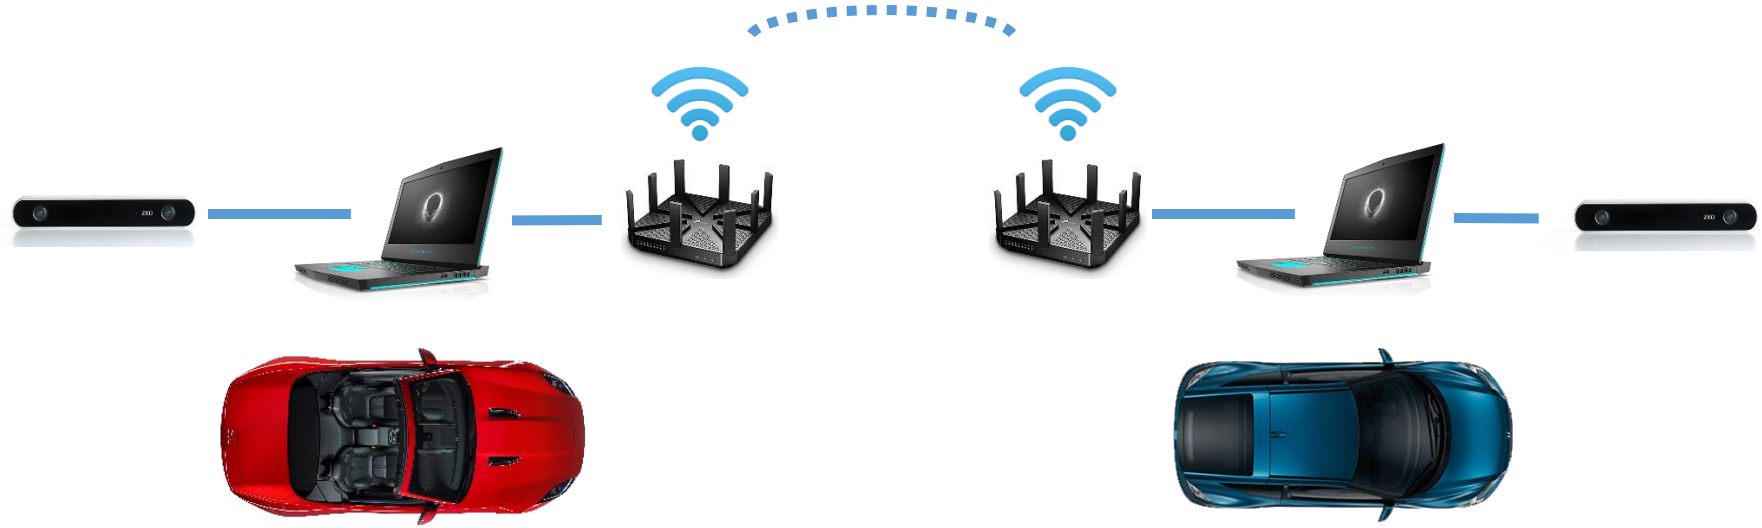

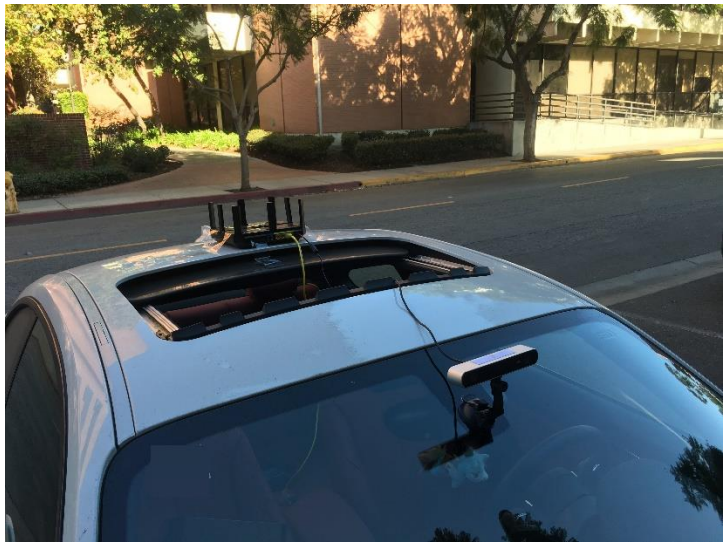

AVR Setup

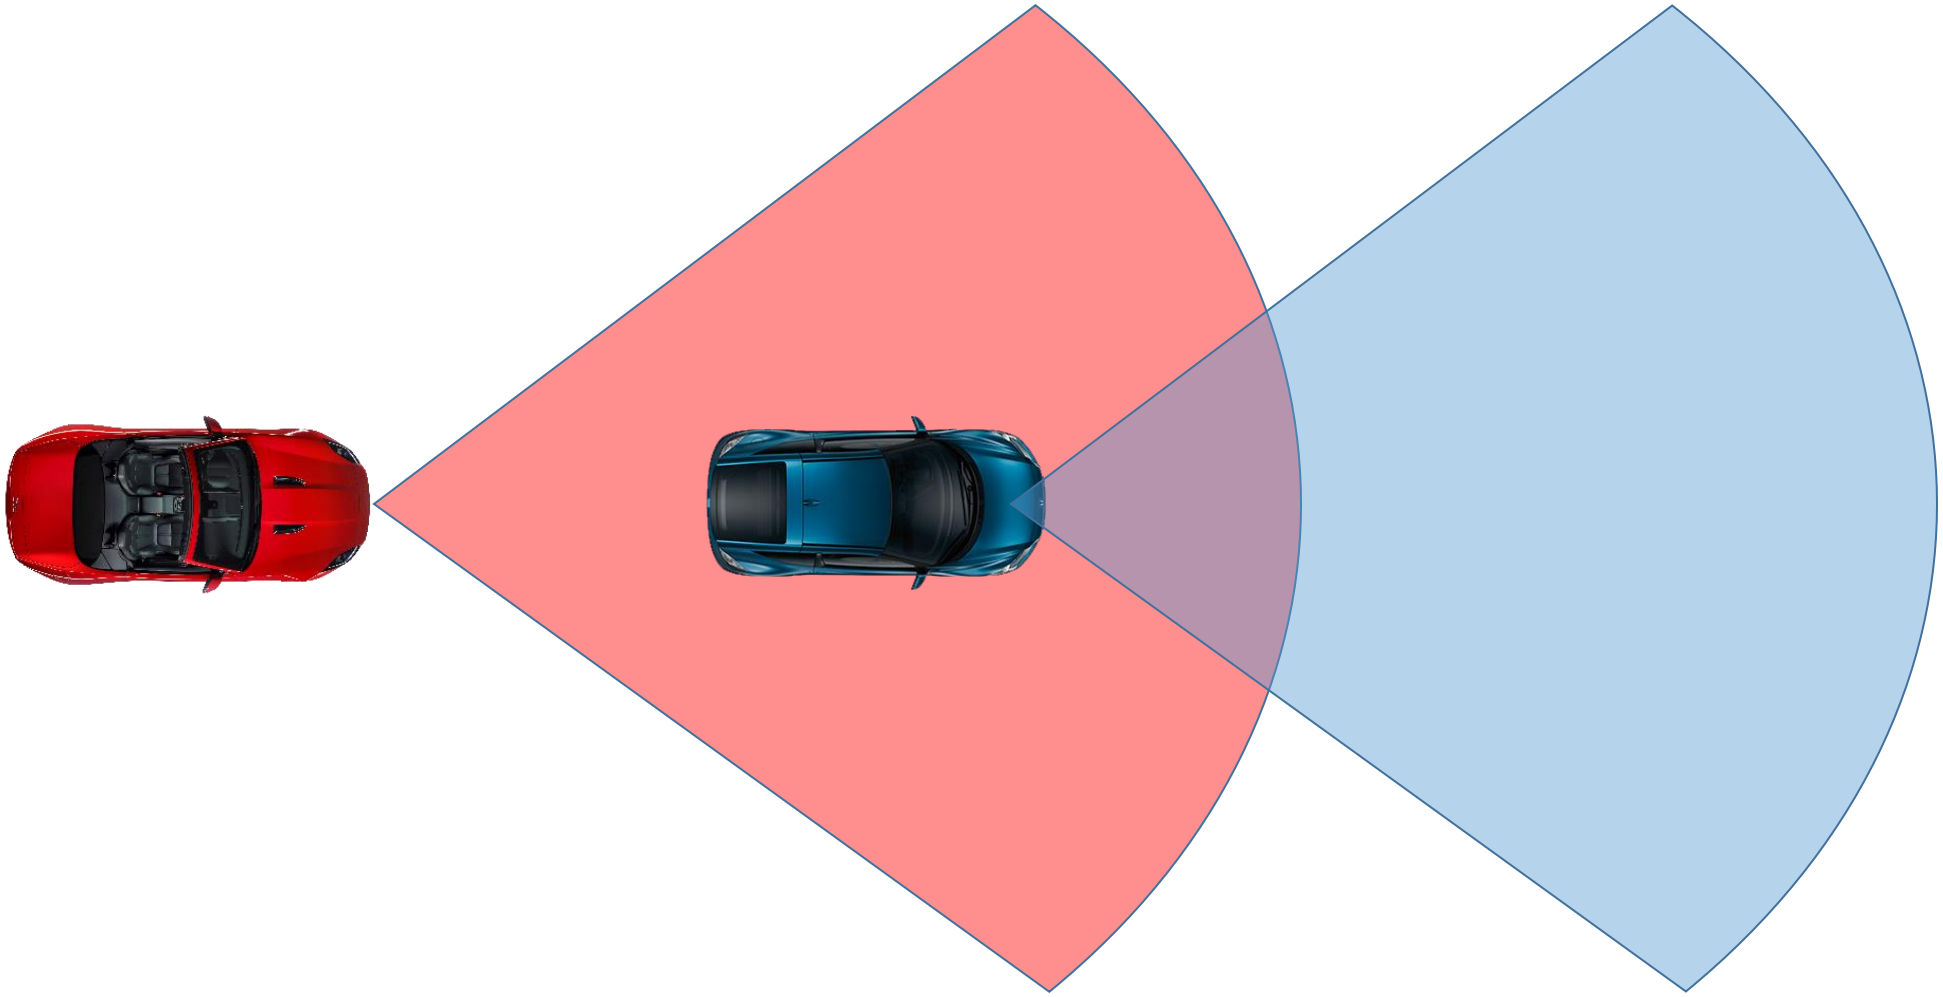

AVR Setup

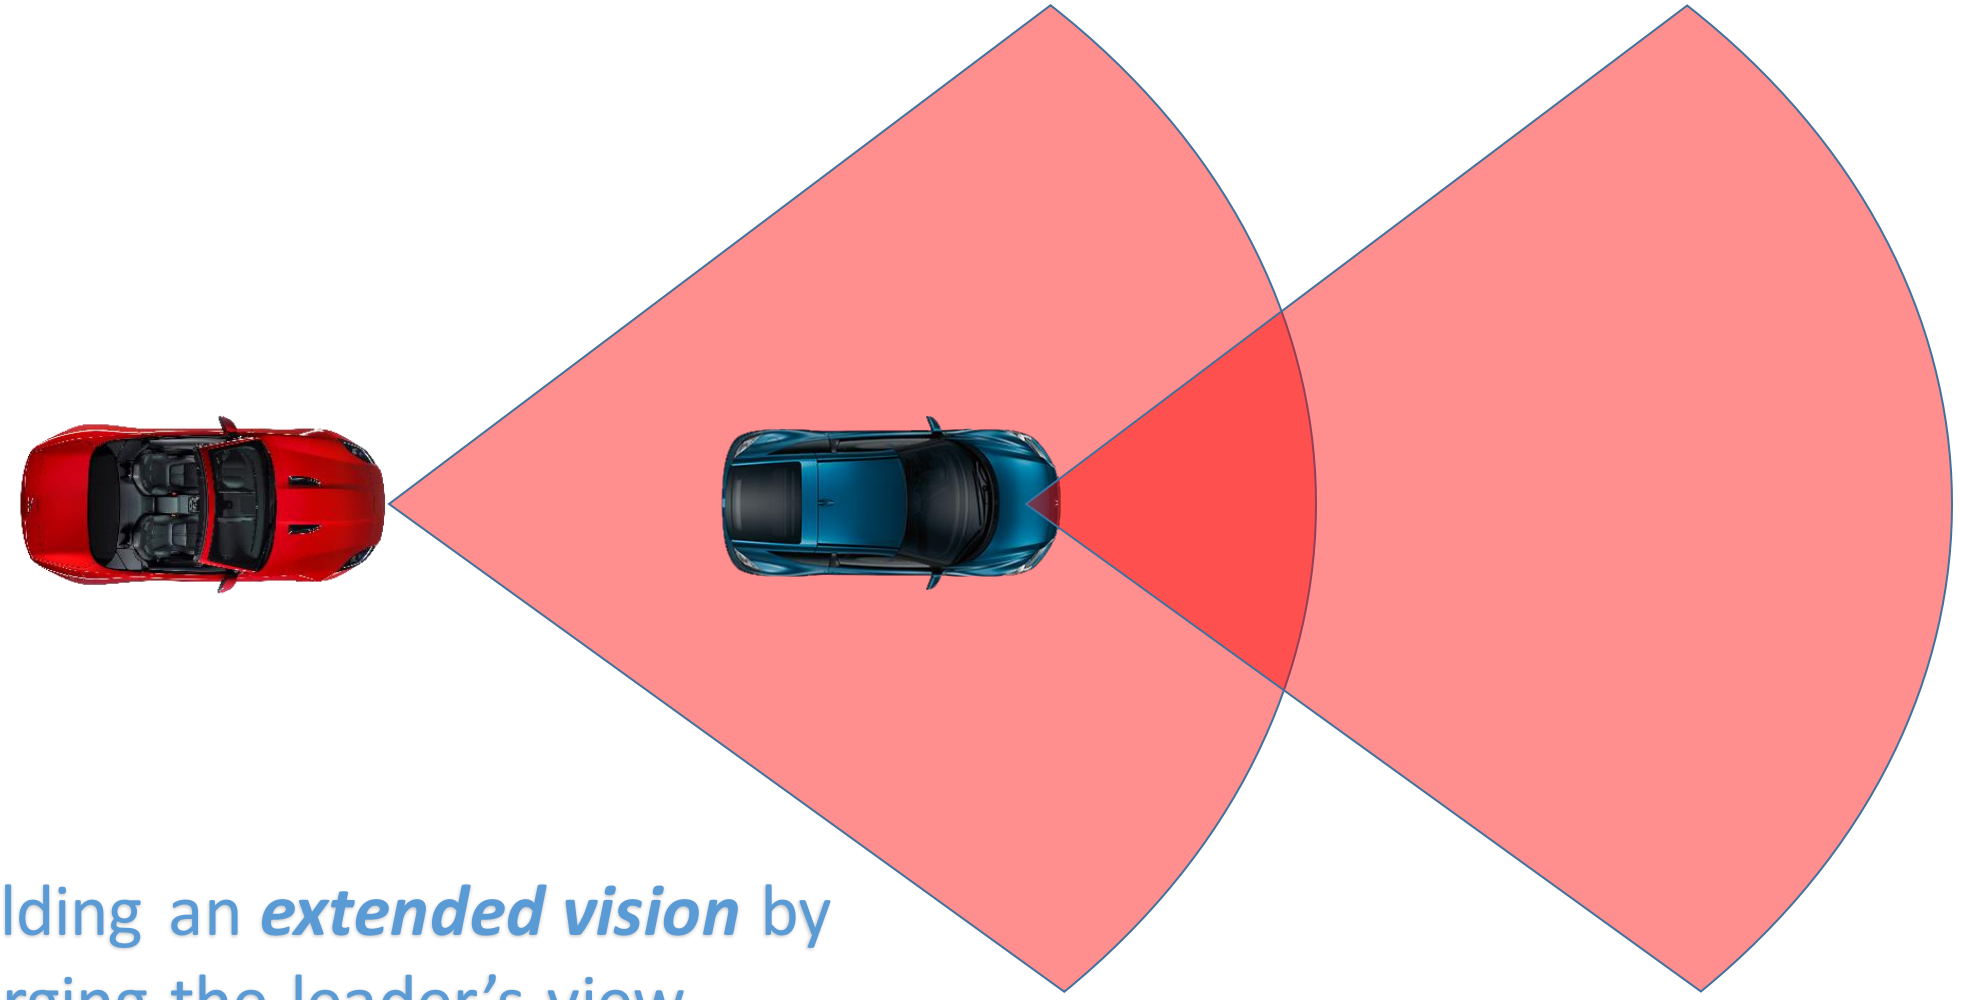

Building an ***extended vision*** by merging the leader's view.

Benefits in Autonomous Driving

Drivable Space (in blue)

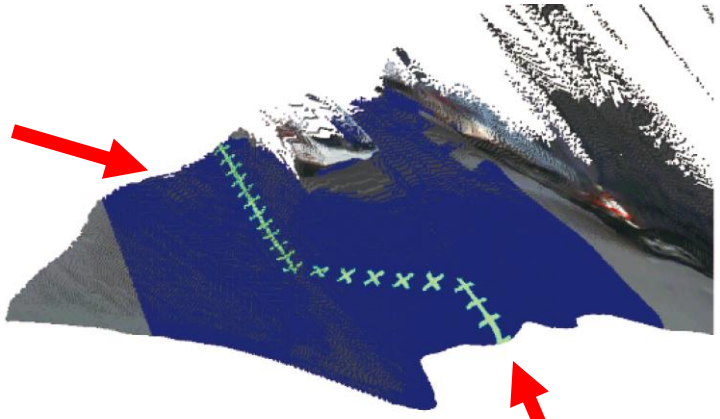

a) Without AVR

Planned Path (in green)

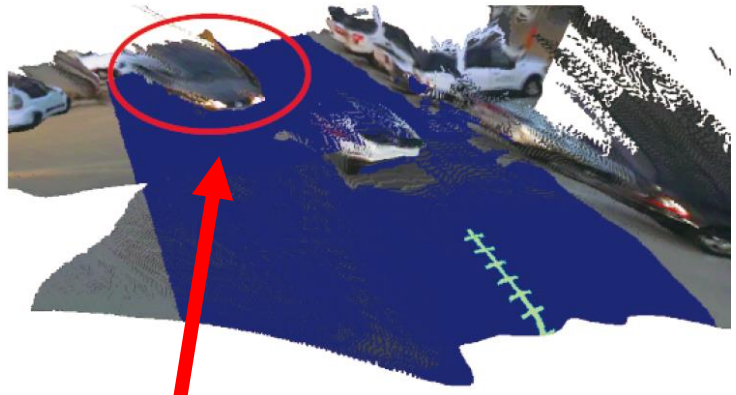

b) With AVR

Follower can see oncoming vehicle (in red circle)

Benefits in Autonomous Driving

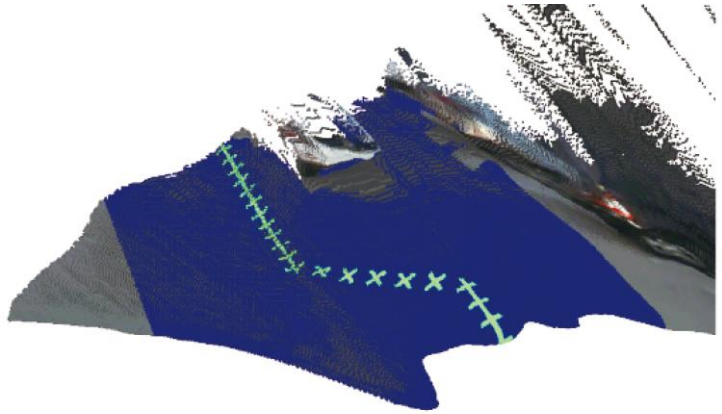

a) Without AVR

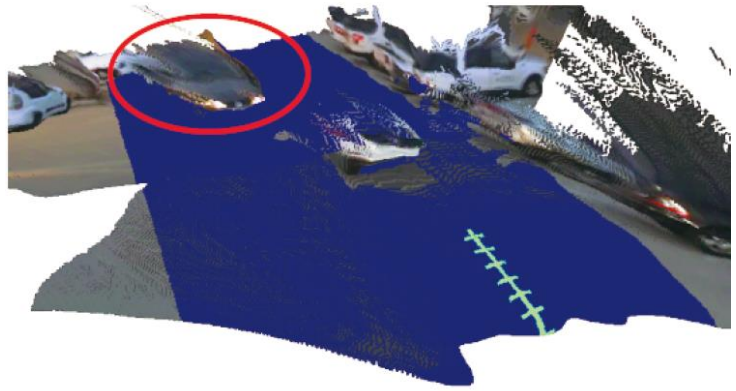

b) With AVR

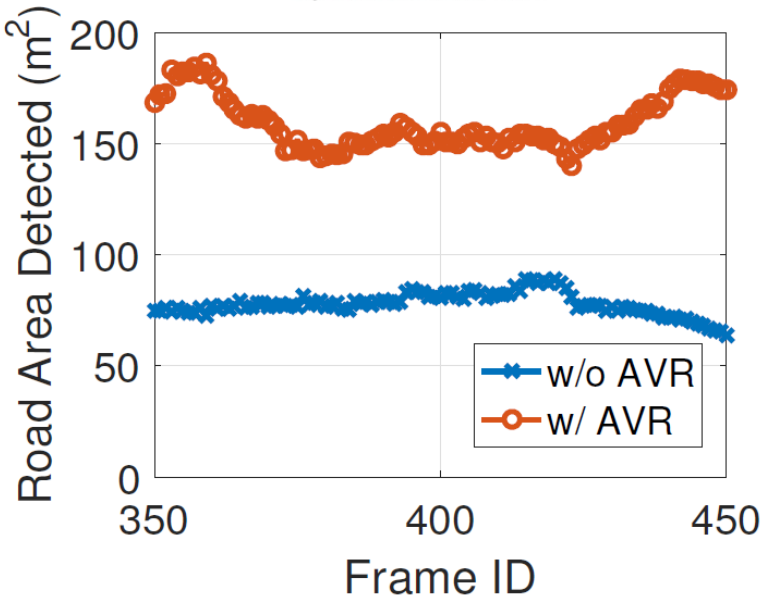

AVR can double the drivable space detected.

Benefits in Autonomous Driving

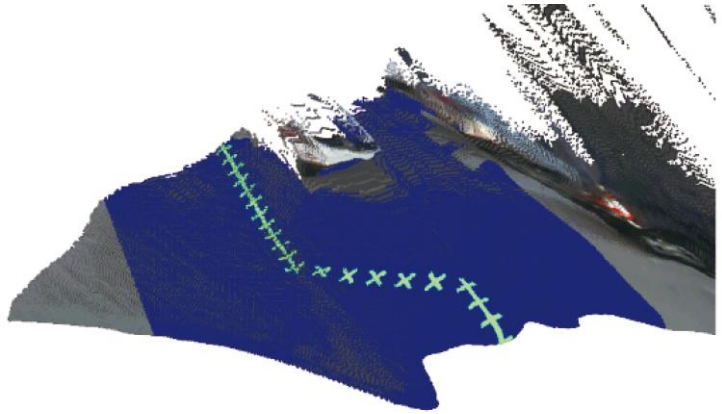

a) Without AVR

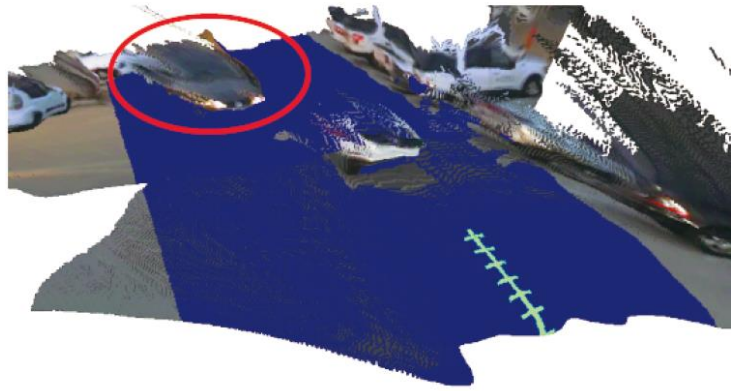

b) With AVR

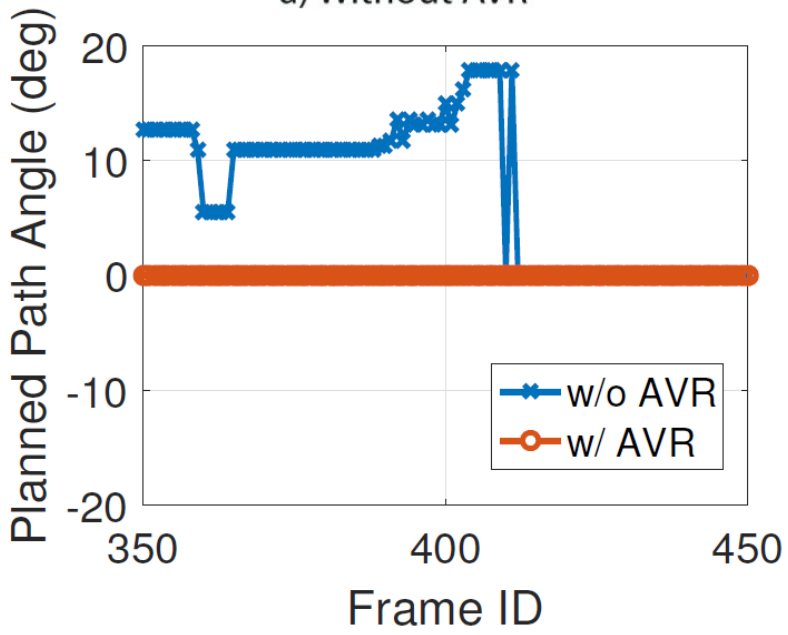

AVR can help path planner to avoid dangerous / wasteful lane change maneuver.

End-to-end Result

Mean	Throughput (Mbps)	Motion Vector Streak	Motion Vector / Frame
Full	367.02	2.80	1.17
Dynamic	34.95	1	0.003

- Full: sharing the point cloud of the whole scene.
- Dynamic: sharing only the dynamic part of the scene.

AVR reduces 10x of bandwidth requirement in dynamic mode.
AVR adapts smoothly to available bandwidth.

End-to-end Result

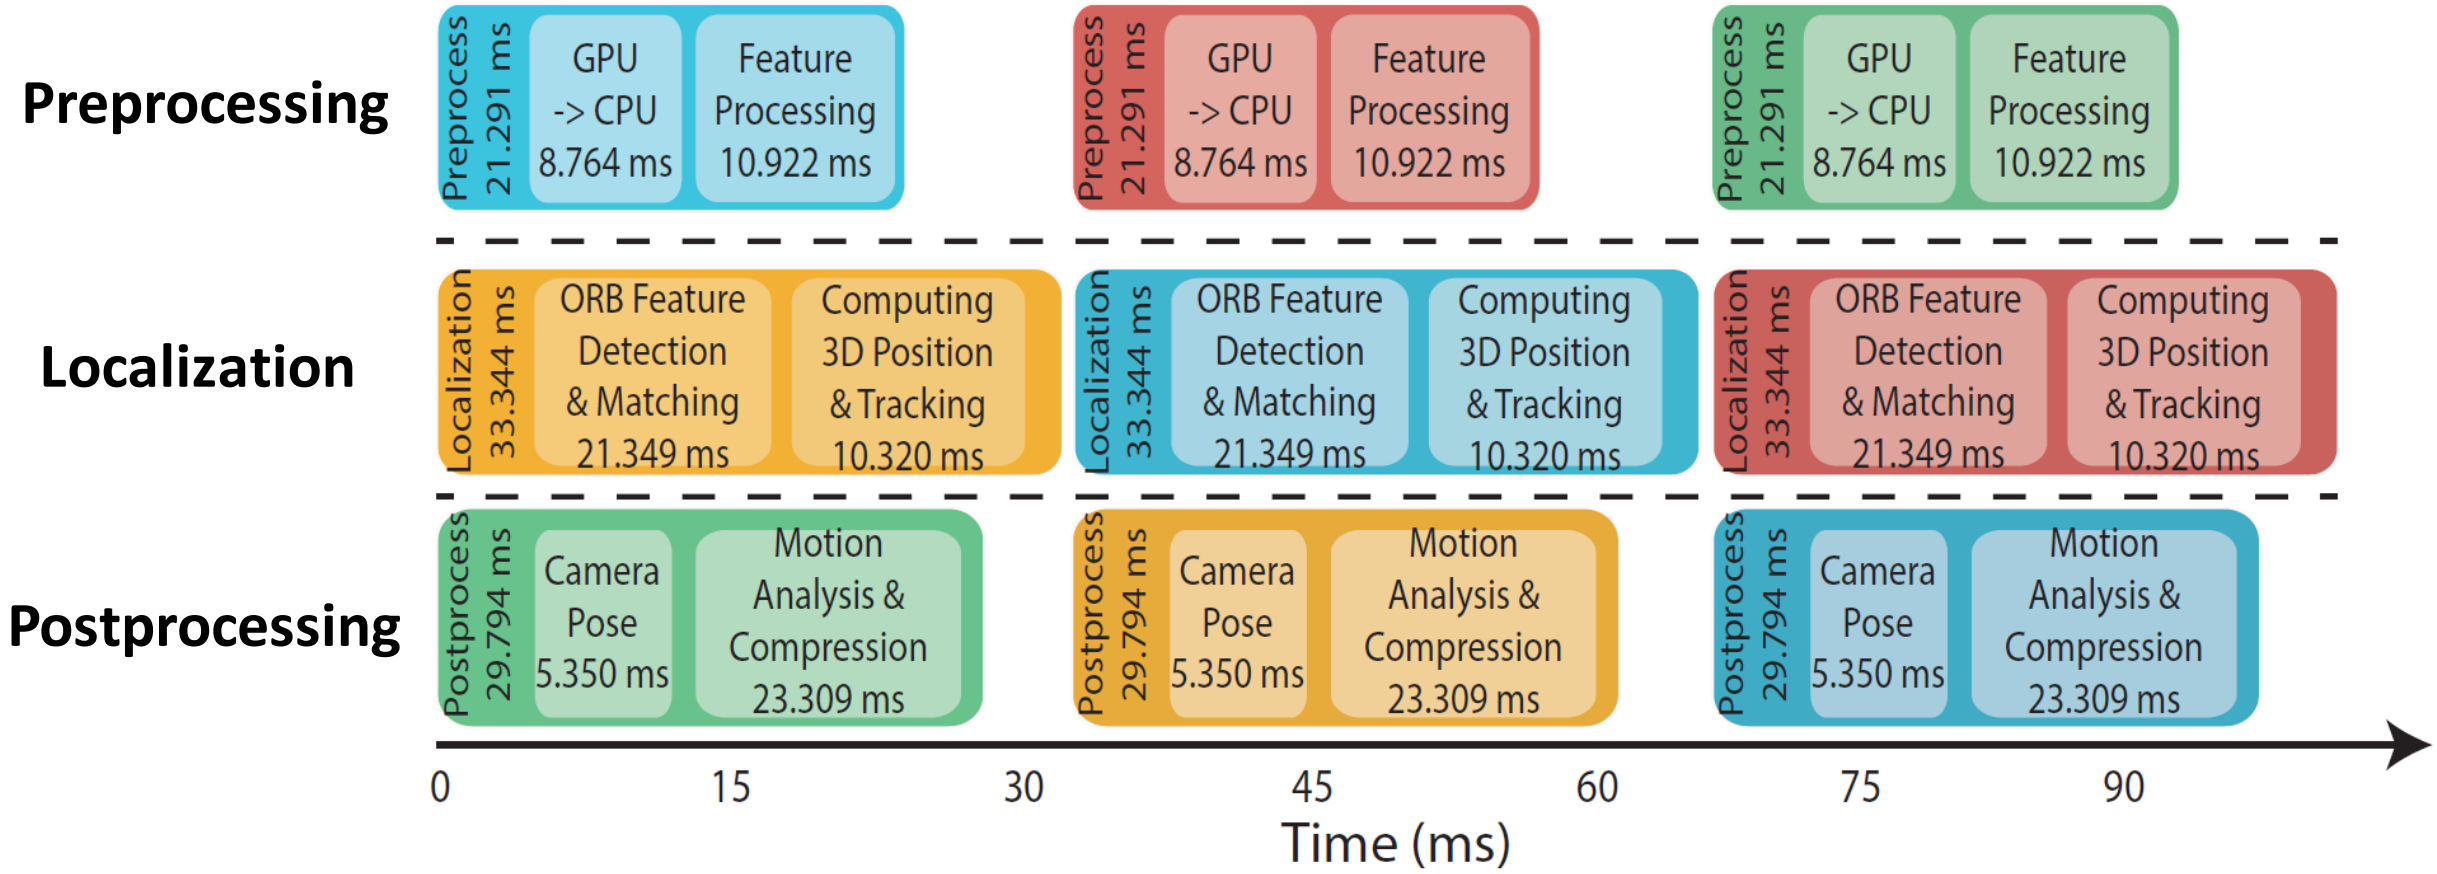

AVR achieves 30 fps using a 3-stage pipeline

AVR induces 96 ms processing delay

Reconstruction Accuracy

Error @ 20mph	Static Objects	Moving Objects
Median	0.027 m	0.070 m
90th Percentile	0.045 m	0.193 m

Caveat: perfect calibration, more details in the paper.

AVR achieves high reconstruction fidelity.
Major source of error: motion estimation, camera calibration.

Limitation & Future Work

- Stereo Camera vs. LiDAR
  - Accuracy
  - Range
  - Robustness to poor visibility
- Cooperative AVR

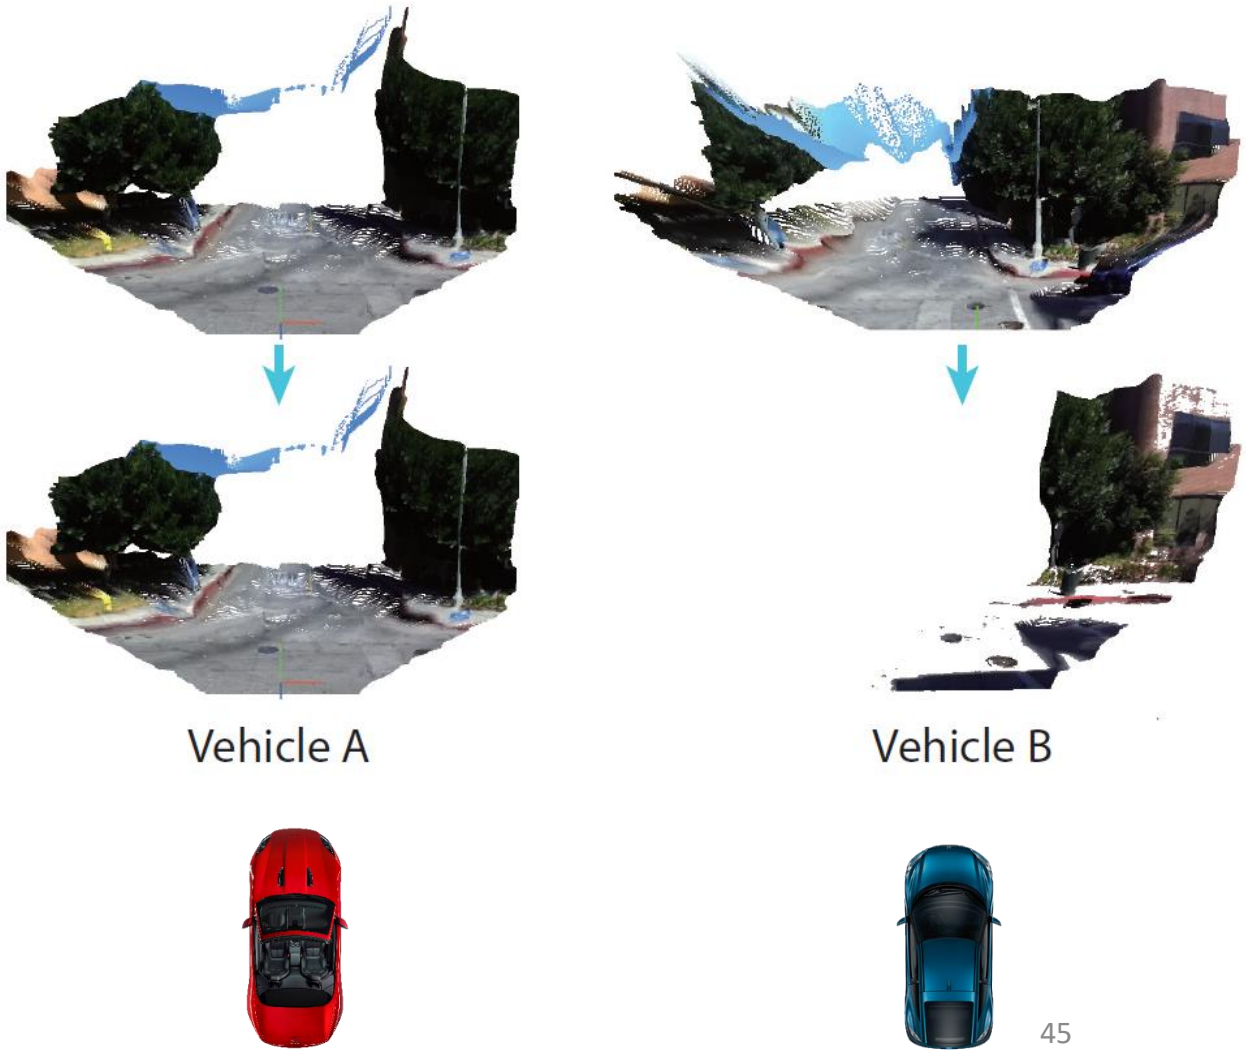

AVR Demo

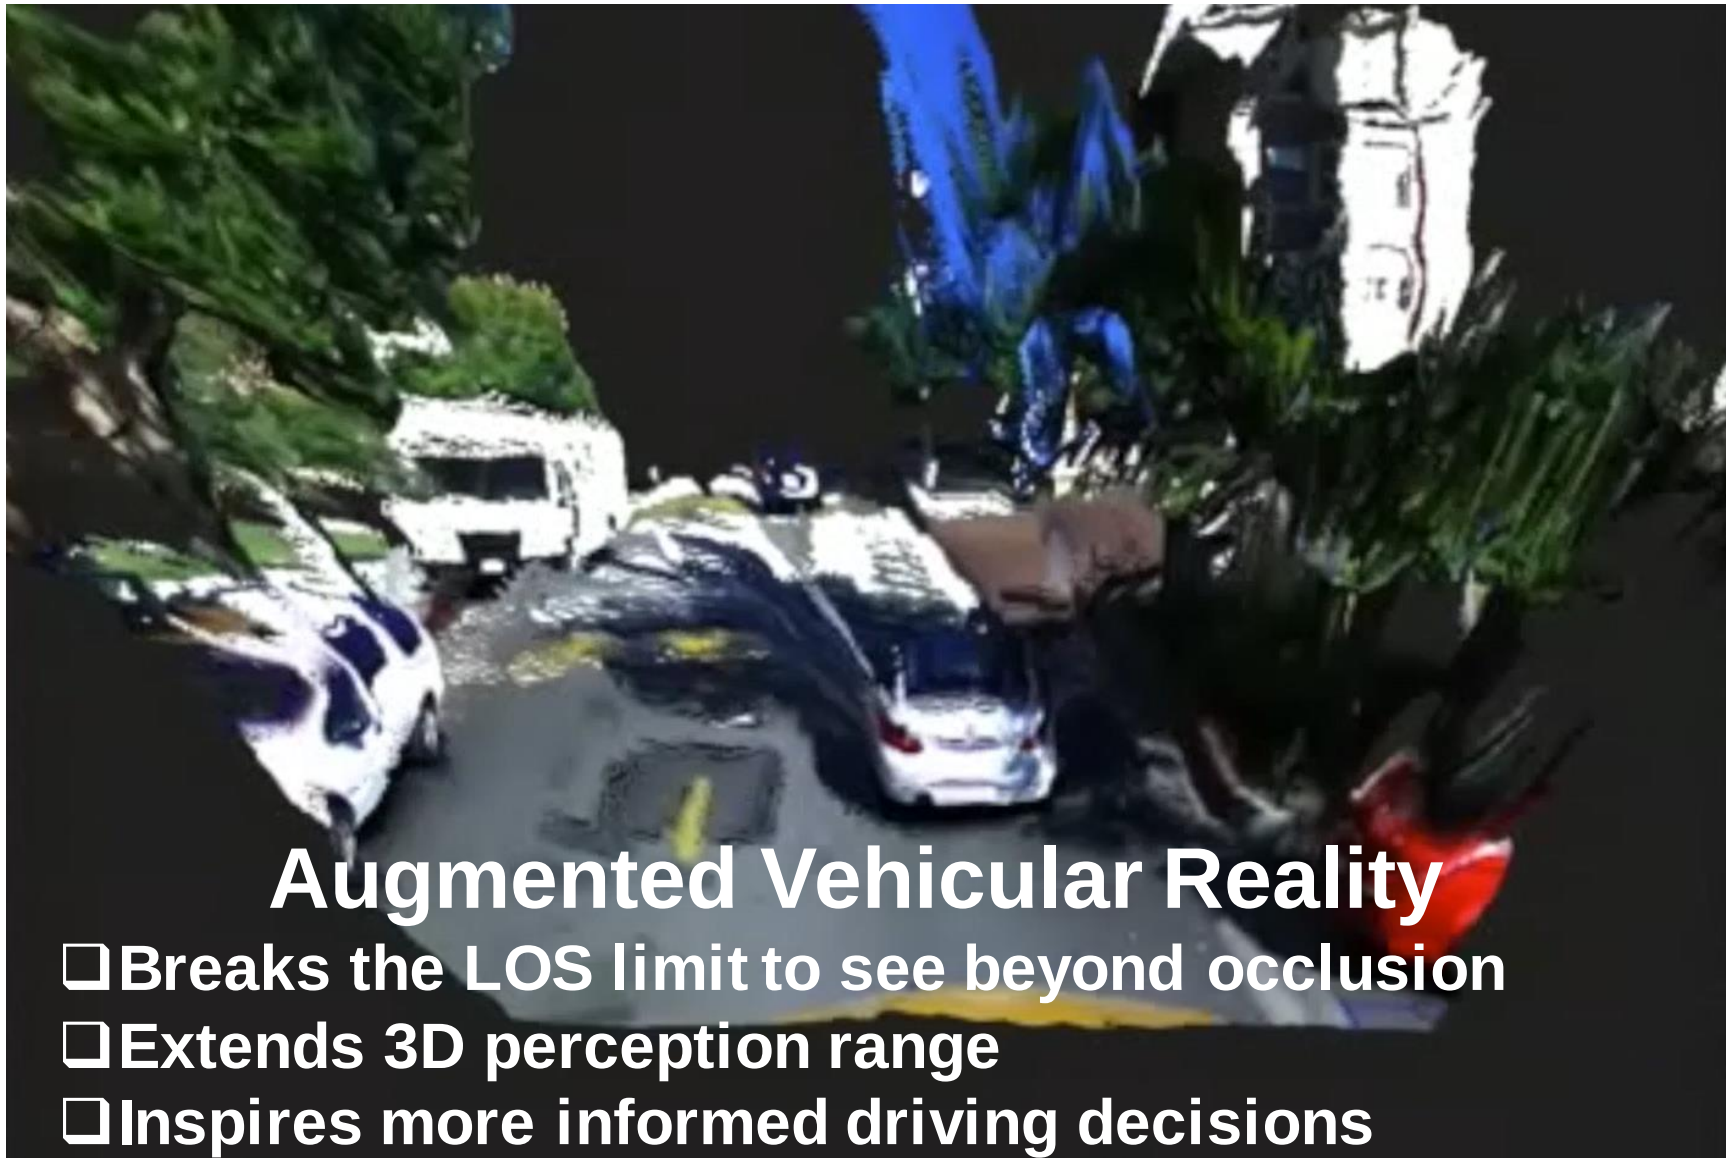
Augmented Vehicular Reality
❑Breaks the LOS limit to see beyond occlusion
❑Extends 3D perception range
❑Inspires more informed driving decisions


Q&A
Thank You!
Q&A
Code & Dataset
https://github.com/hanqiu/AVR16
Q&A