Motional

Challenges in Assured Autonomy for Self-driving Cars

Eric M. Wolff
Principal Research Scientist
Motional Autonomy Research

Why is programming a self-driving car hard?

- Sense
  - “See” with cameras, lidar, and radar
  - High-dimensional input (pixels + pointclouds), occlusions, so many types of objects
- Plan
- Act

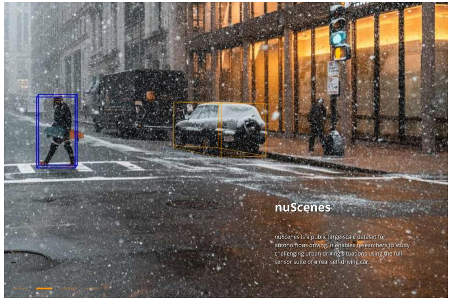

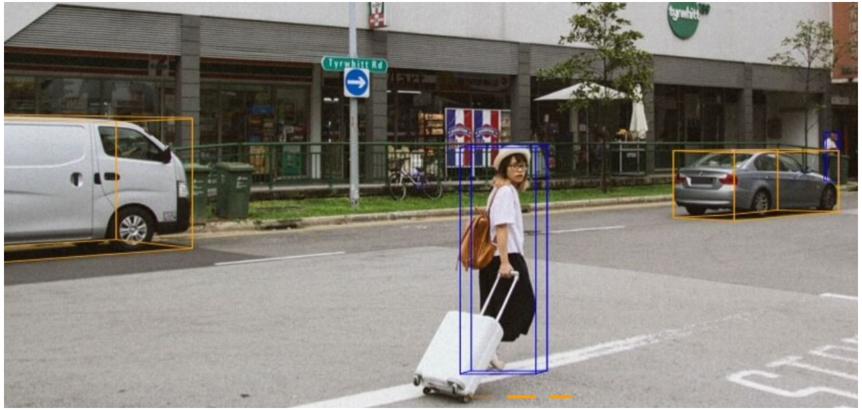

Why is programming a self-driving car hard?

- Sense
    - “See” with cameras, lidar, and radar
    - High-dimensional input (pixels + pointclouds), occlusions, so many types of objects
- Plan
    - Behavior depends on the scene context
        - Other cars, pedestrians, bicyclists, etc
        - Road geometry and markings (e.g., bus stops, stop signs)
    - Need to anticipate what others will do
    - Rules-of-the-road are not complete or followed
- Act
    - Need to drive smoothly and comfortably

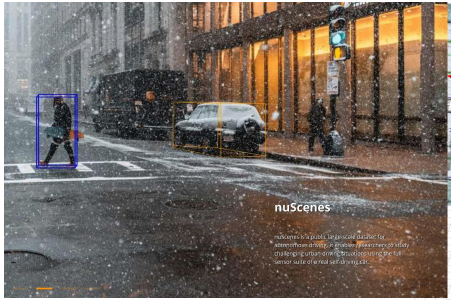

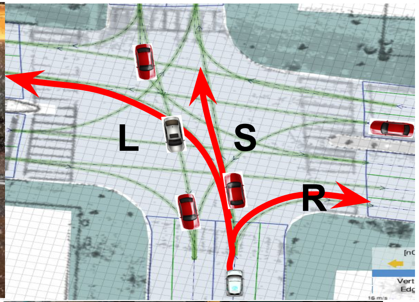

Machine Learning is a key component of self-driving cars

- When to use ML?
  - Lots of data
  - Stable data distribution
  - No good analytical model

Machine Learning is a key component of self-driving cars

- When to use ML?
  - Lots of data
  - Stable data distribution
  - No good analytical model

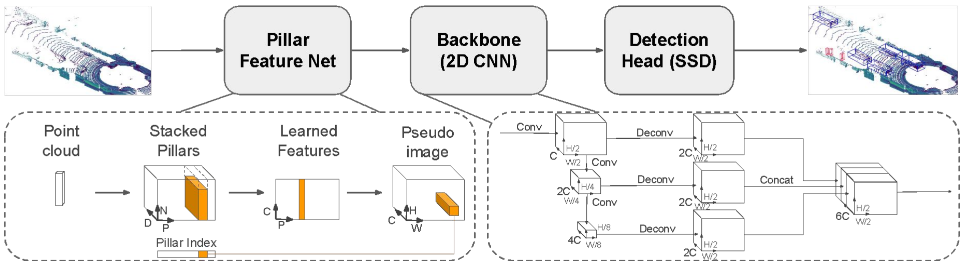

A. Lang et al. "PointPillars: Fast Encoders for Object Detection from Point Clouds." CVPR 2019 https://arxiv.org/abs/1812.05784

- Where is ML typically used in self-driving cars?
  - Perception (heavily)
  - Motion prediction (moderately)
  - Planning (partially)

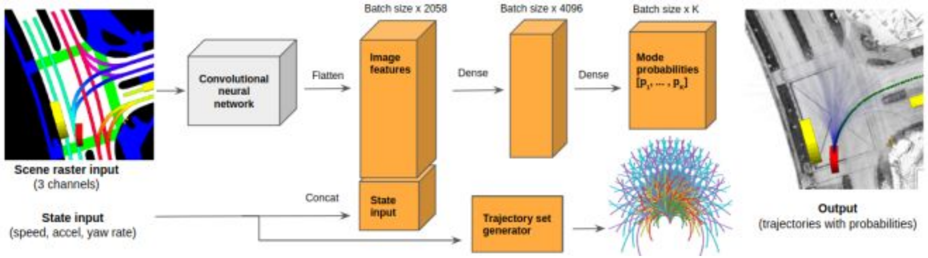

T. Phan et al. "CoverNet: Multimodal Behavior Prediction using Trajectory Sets." CVPR 2019. https://arxiv.org/abs/1911.10298

How to assure safety of ML systems?

- There is no silver bullet
- A key challenge is rare events (the long tail)

How to assure safety of ML systems?

- There is no silver bullet
- A key challenge is rare events (the long tail)
- Safety certification of ML systems
    - Compute (easy)
    - Data pipeline (med)
    - Model (hard)
        - Meets performance targets across a comprehensive test set
        - Non-ML safety checks and limits
        - Within operating domain?
        - Uncertainty estimates

How to assure safety of ML systems?

- There is no silver bullet
- A key challenge is rare events (the long tail)
- Safety certification of ML systems
  - Compute (easy)
  - Data pipeline (med)
  - Model (hard)
    - Meets performance targets across a comprehensive test set
    - Non-ML safety checks and limits
    - Within operating domain?
    - Uncertainty estimates
- Shared data and evaluation across the industry?
  - Eval and metrics must be carefully determined
  - Research: www.nuscenes.org, www.argoverse.org, www.waymo.com/open
  - Safety
    - www.pegasusprojekt.de/en/
    - hwww.safetypool.ai/

How to convince the public?

- Multiple stakeholders
  - Users
  - Other road users
  - Local, state, and federal government
- Safety is not apparent from a test drive
- Can we adapt current auto regulations? Or look to aerospace?

Conclusions

- Machine Learning is a core part of self-driving cars
  - Perception
  - Prediction
  - Planning
- How to assure safety of an ML system?
  - A grand challenge
  - Compute, data pipeline, and model
  - Comprehensive evaluation and data sharing

You can reach me at **eric.wolff@motional.com**