Controlling autonomous cars using ML

Maciej Dziubiński
27-11-2021

Center4ML

Agenda

- Introduce yourself
- The fundamental problems in autonomy
- Perception
- Localization
- Planning
- **Control**
- (Bird’s-eye view)

Introduce yourself

2016: Udacity Self-Driving Car Nanodegree

2017: 1st place in the F1/10th competition
(with Karol Majek and Łukasz Sztyber)

Projects:

- **2017-2018**: driving in the CARLA simulator
- **2019**: recreating the top-down view from side cameras
- **2019**: driving around a pond

Most recently:

Indy Autonomous Challenge

EuroRacing team's LI site

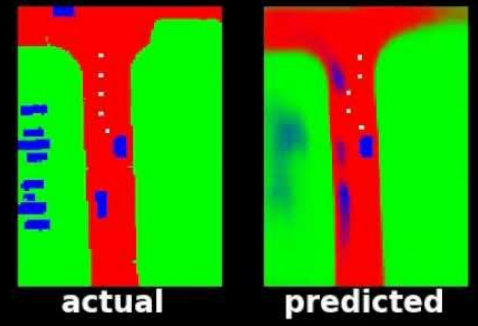

Fundamental sub-problems in autonomous cars

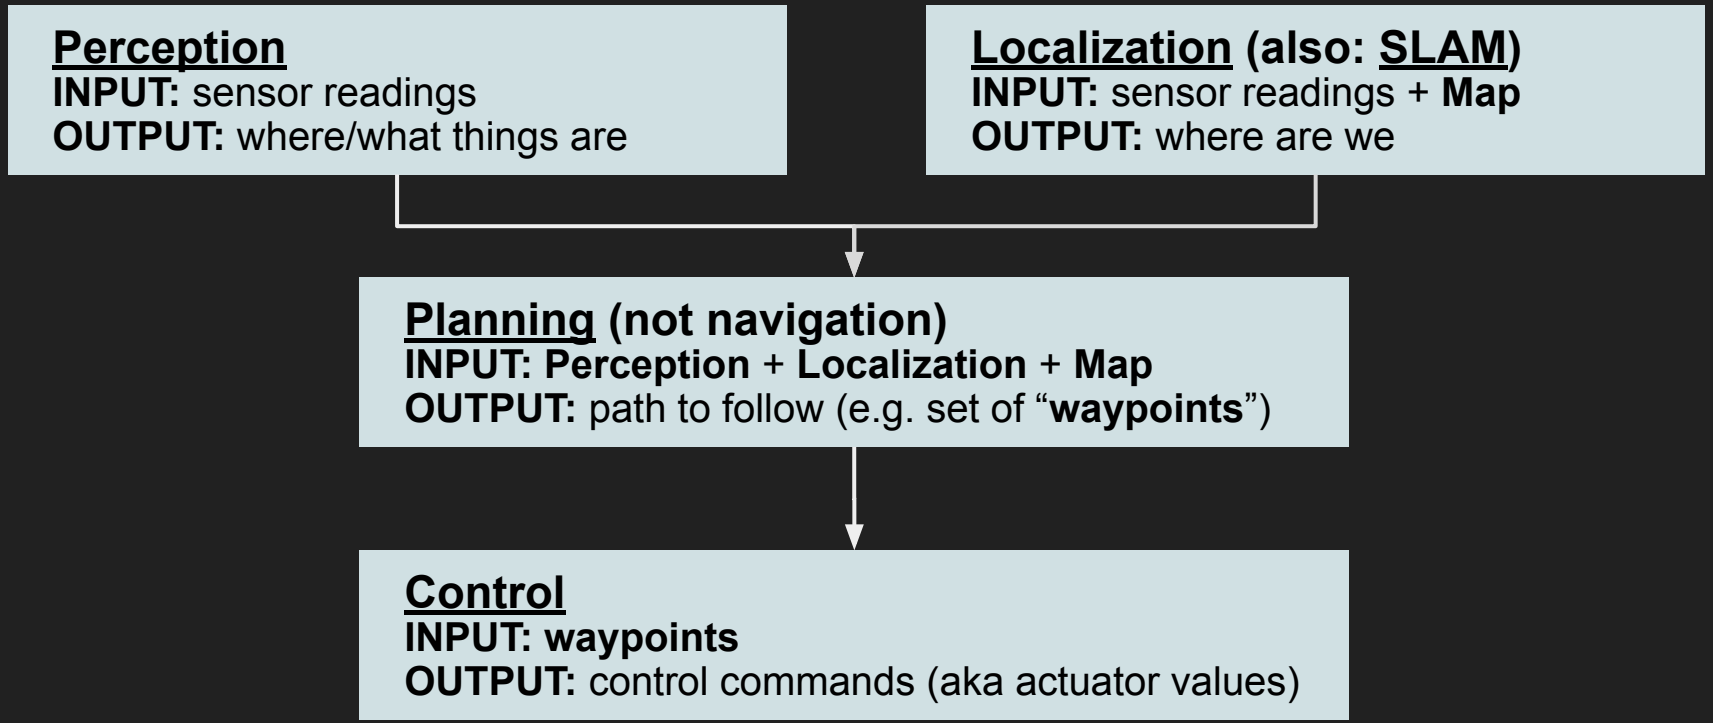

Perception & Machine Learning

- Probably the best sub-problem for ML
- Needs to be fast
- Sensors: cameras, LIDAR, radar, sonar, …
- On the right: Tesla + panoptic segmentation

Localization

- "No ML" sub-problem (or is it?)
- Needs to be fast (and precise)
- Sensors: cameras, LIDAR, radar, sonar, …
- On the right: Particle filter + SLAM

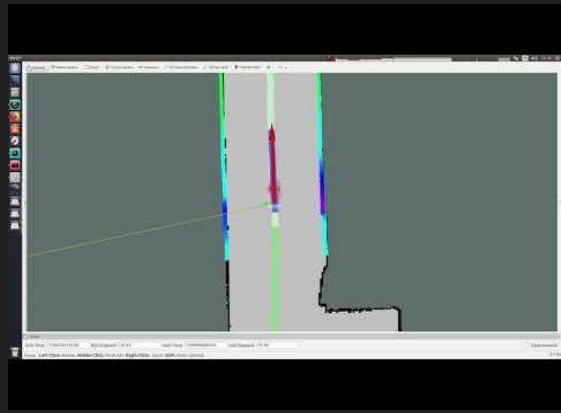

SLAM = Simultaneous Localization and Mapping

- In particular Visual SLAM (**VSLAM**)
- Needs to be fast (and precise)
- Sensors: cameras, LIDAR, radar, sonar, …
- On the right: Particle filter + SLAM

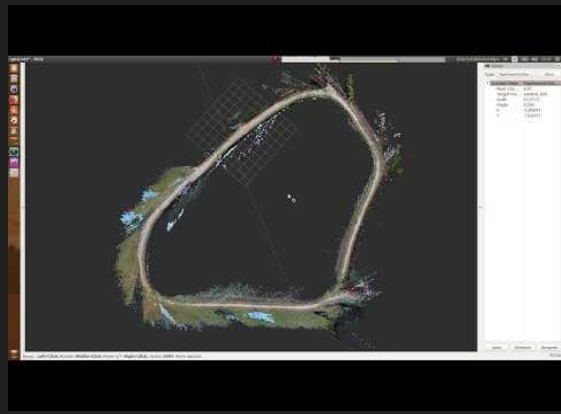

Planning

- Need to foresee what's going to happen up to 5sec into the future (at ~50kmph) -- requires motion prediction
- [Argoverse motion forecasting competition]
- [Kaggle competition (from Lyft)]

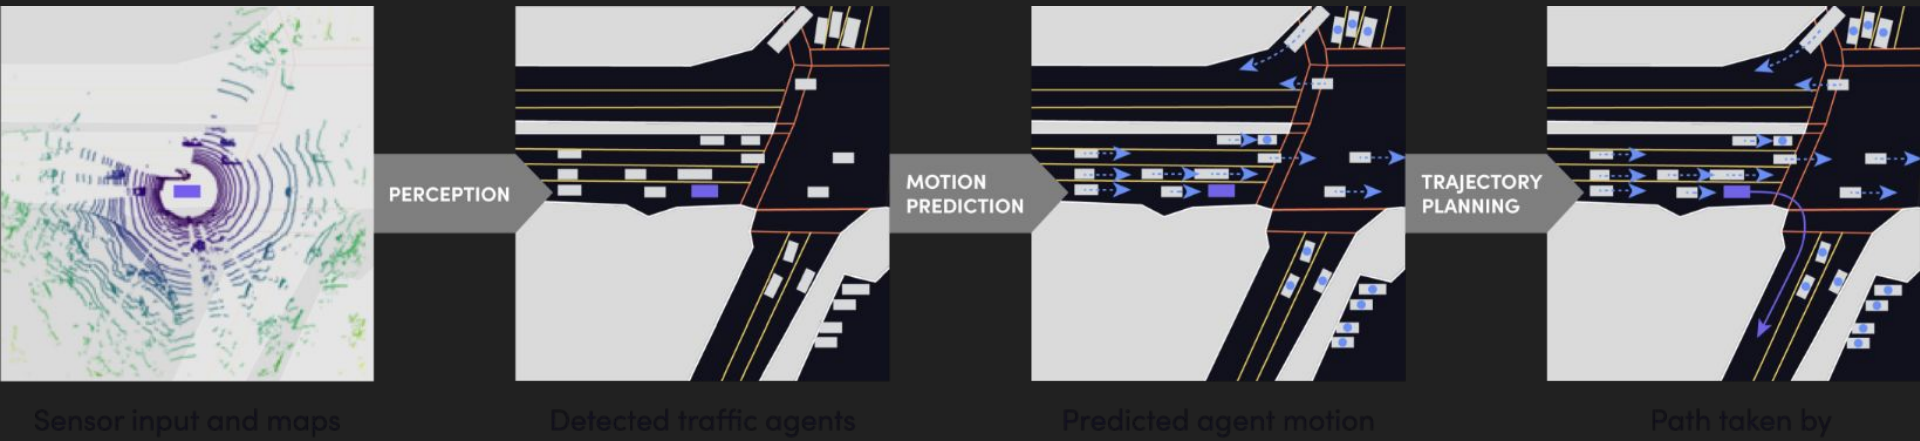

Control

- That is: control the steering angle and throttle such that you “check” a set of consecutive waypoints

- Typically framed as an Optimal Control Problem:
  | **Optimal control** theory is a branch of mathematical
  | optimization that deals with finding a control for a dynamical
  | system over a period of time such that an objective function
  | is optimized

- Reinforcement Learning is perfect for this kind of setup (on the right: Soft Actor Critic)

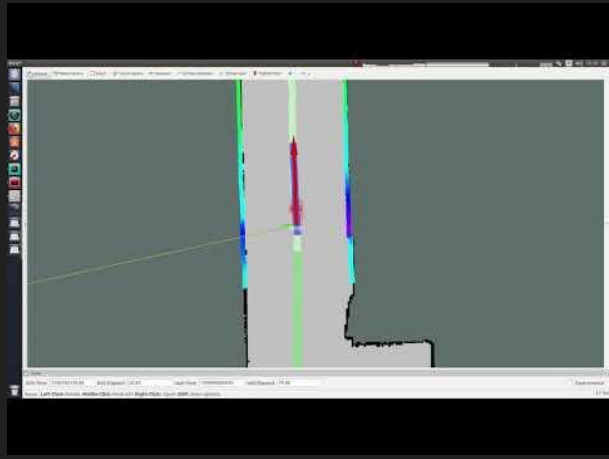

Bird’s Eye View

- Project input from side cameras into a bird’s-eye view
- My approach was to take semantic segmentation ground truth input from four cameras (in the CARLA simulator)

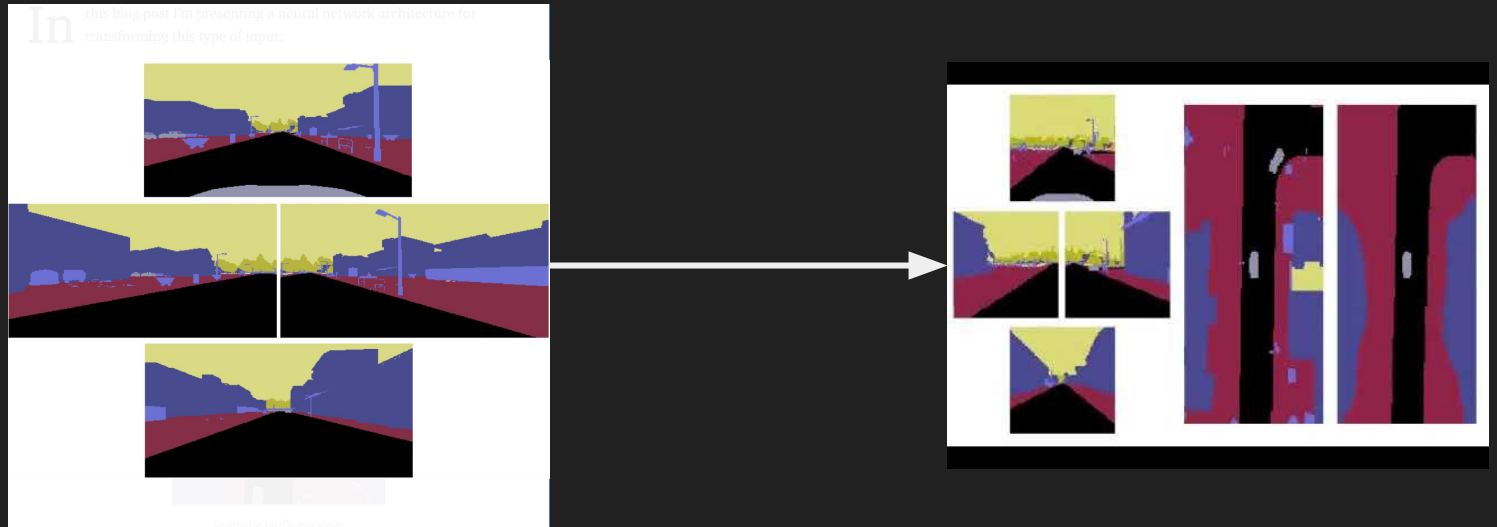

Future directions

- Tesla example (on the right)
- **Latest and greatest**

Thank you for your attention

Questions?