Obstacle Avoidance for Autonomous Racing using Reinforcement Learning

ESL
Electronic Systems Laboratory

UNIVERSITEIT • STELLENBOSCH • UNIVERSITY

Department of Electrical & Electronic Engineering

Benjamin Evans

Dr. W. Jordaan, Prof. H Engelbrecht

September 4th, 2020

UNIVERSITEIT
STELLENBOSCH
UNIVERSITY

Electronic Systems Laboratory

Table of Contents

UNIVERSITEIT STELLENBOSCH UNIVERSITY

UNIVERSITEIT • STELLENBOSCH • UNIVERSITY

Introduction & Background

• • •

What is F1tenth?

Electronic Systems Laboratory

UNIVERSITEIT
STELLENBOSCH
UNIVERSITY

F1tenth Racing

You have a car and the goal is to race as fast as possible around the track

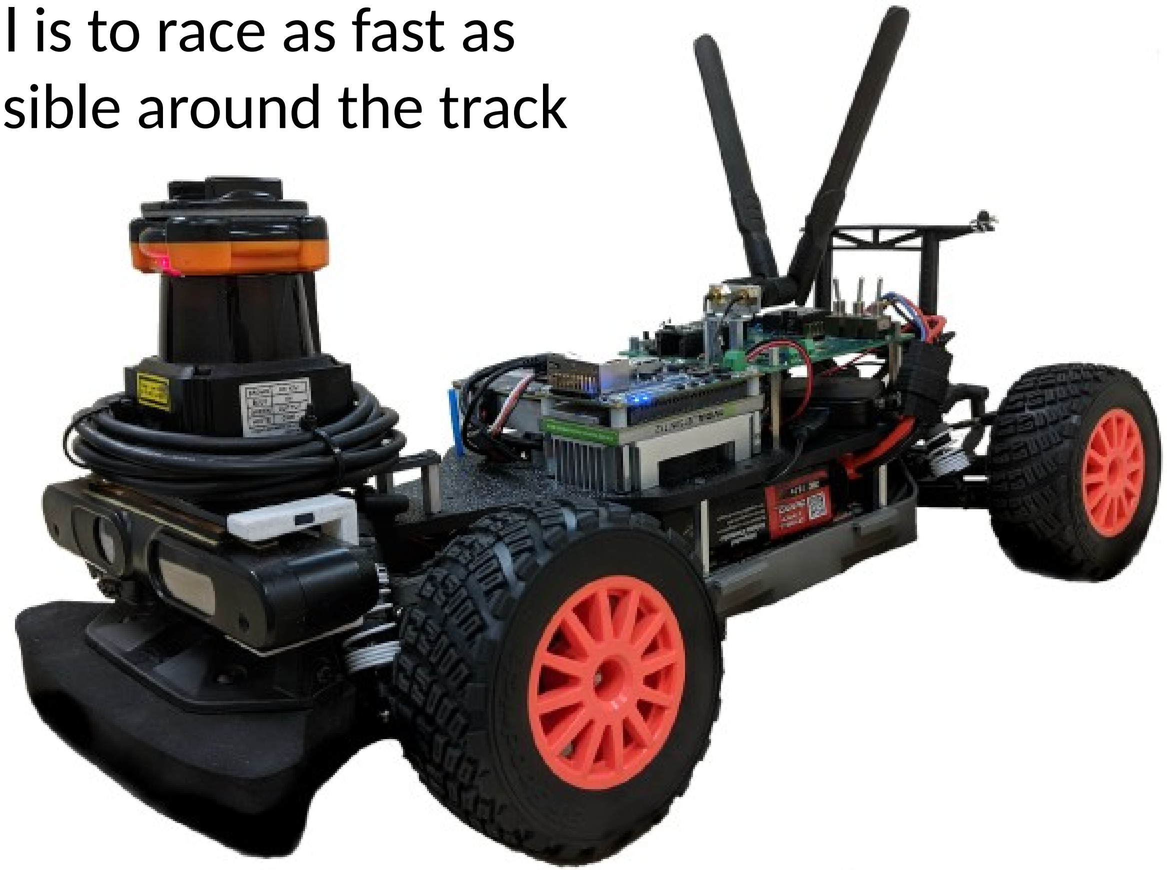

ESL
Electronic Systems Laboratory

F1 Tenth Demo

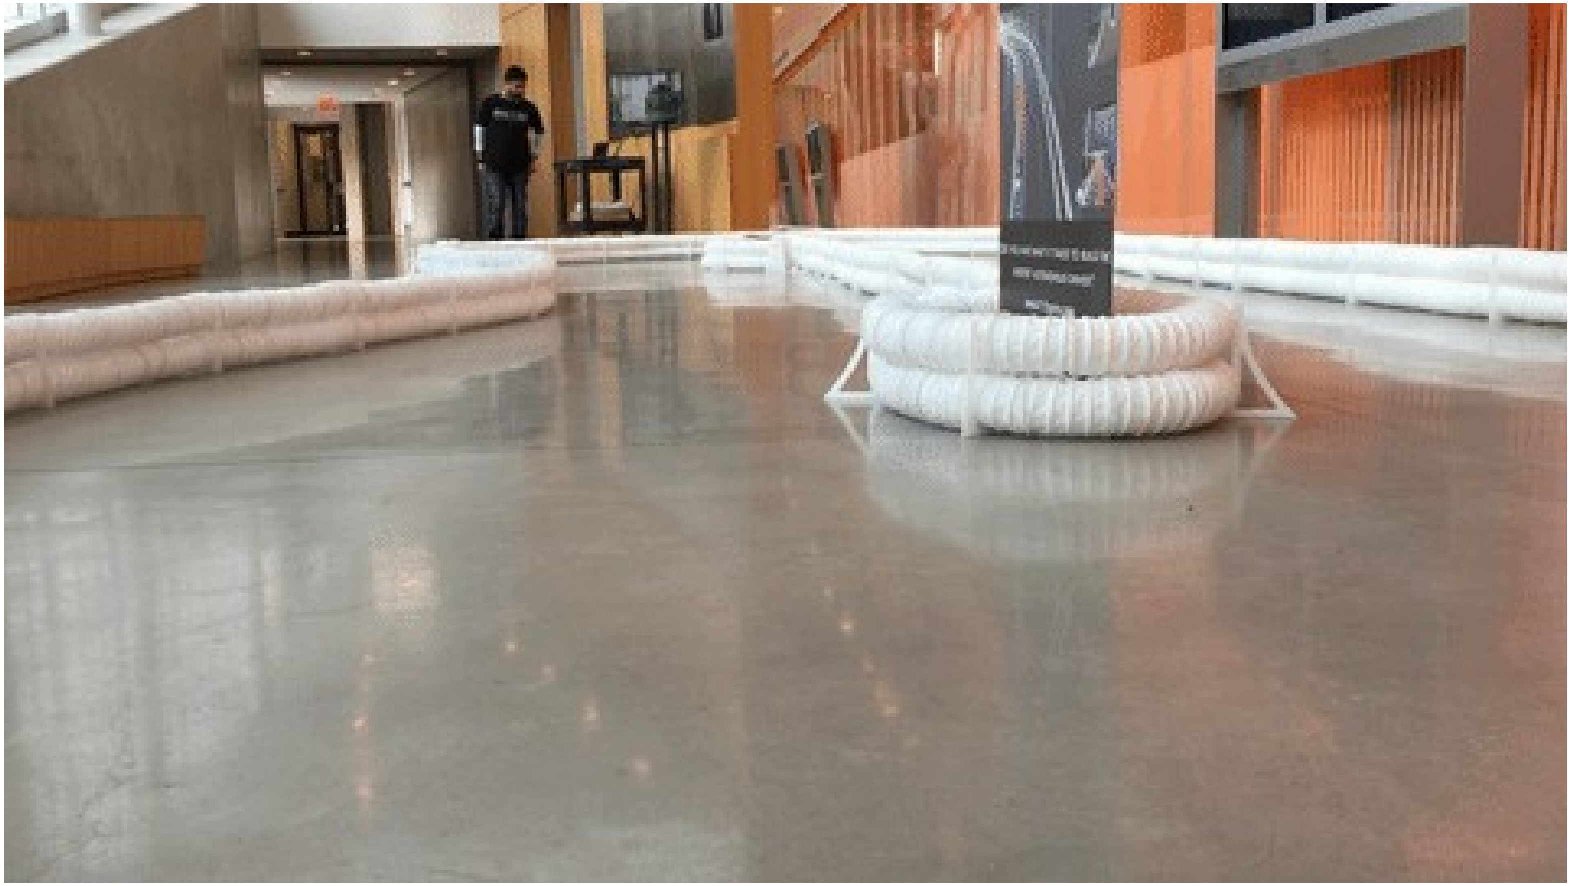

UNIVERSITEIT
STELLENBOSCH
UNIVERSITY

F1tenth Components

- Usually split into three parts
  1. Perception: what is happening around you
  2. Planning: deciding what to do
  3. Control: executing your plan

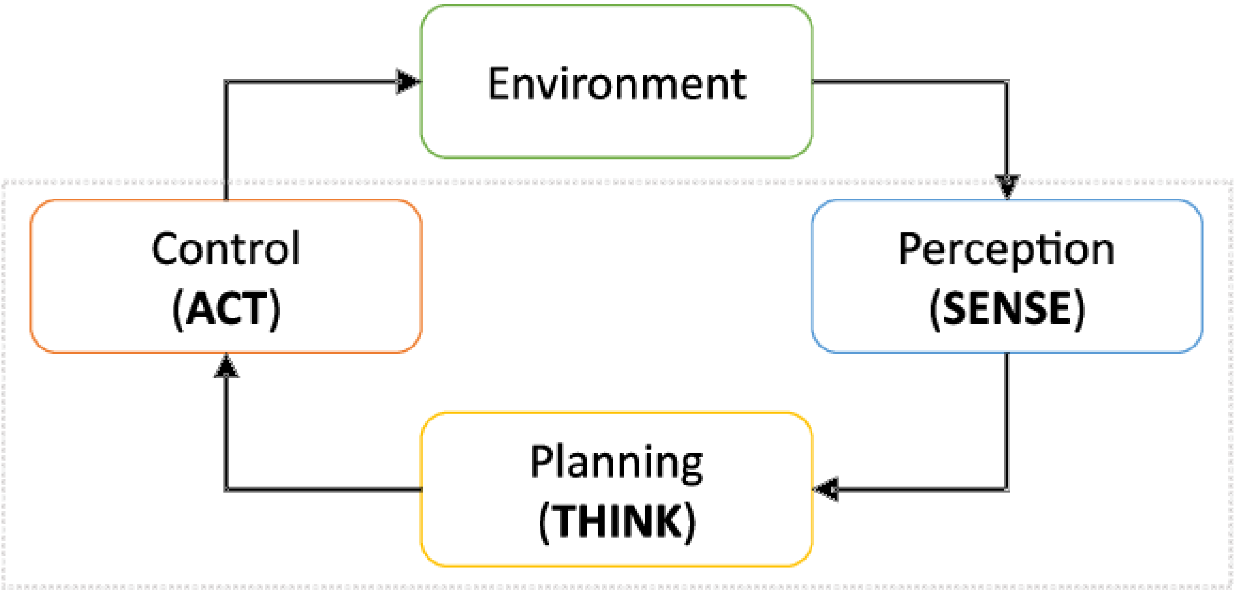

Electronic Systems Laboratory

Planning to Win

- Planning is done in two stages
- Global planning
  - Done before the race
  - Takes in a map of the track
  - Produces low resolution waypoints through the course
  - Usually uses a path finder like A* with an optimization step
- Local planning
  - Done on a continuous basis through the race
  - Relies on an <u>update obstacle map</u> built with a SLAM algorithm
  - Uses the updated map and current state of vehicle to determine control references for the control system.
  - Output: velocity magnitude, heading direction.

UNIVERSITEIT STELLENBOSCH UNIVERSITY

SLAM

- As the vehicles race, they use SLAM algorithms to "simultaneously, localize and map" the environment.
  - Localize: determine their location.
  - Map: build a map of the environment around them.

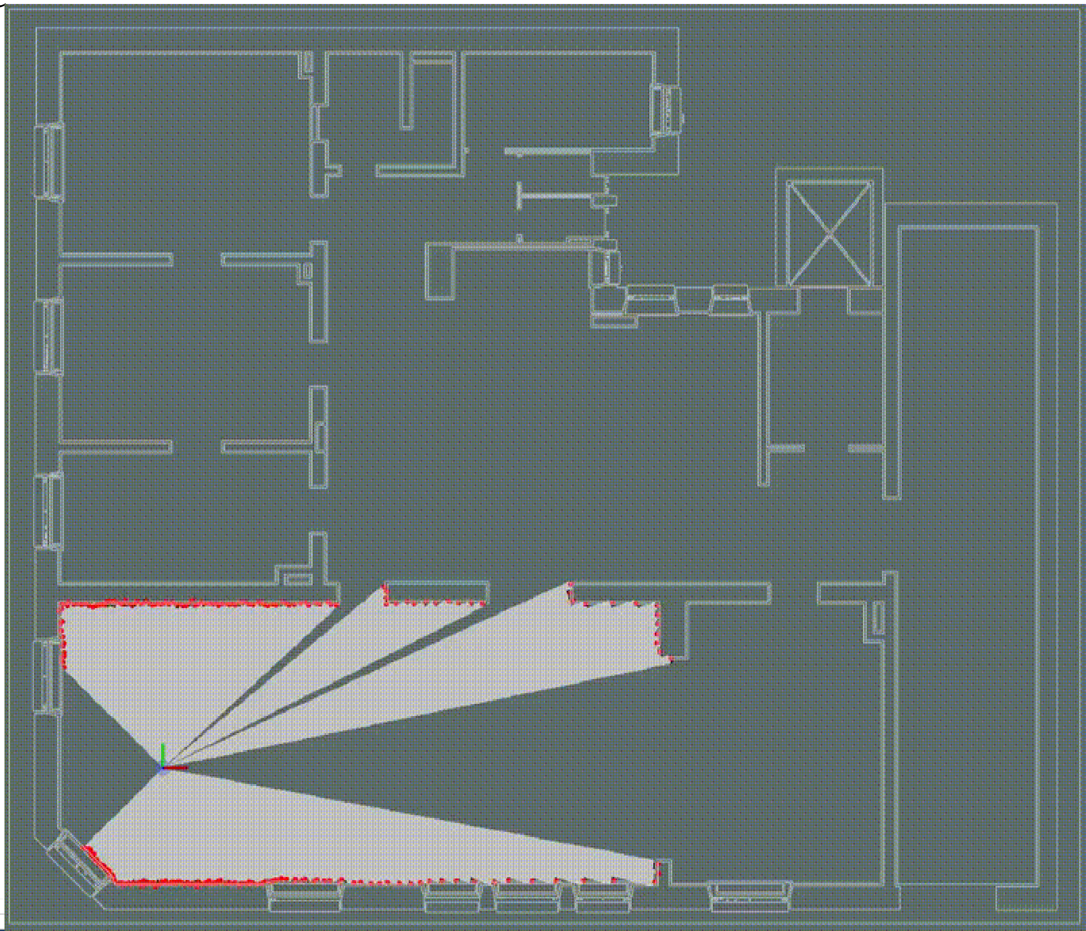

Problems with SLAM

- Takes a long time
  - This is a problem in racing where you need your system to run fast
- Requires a lot of processing power
  - Sometimes have limited processing power
- Complex to implement
  - Requires advanced algorithms
- Struggles to map dynamic targets
  - It doesn't know which parts of the environment are moving and which are stationary.

Electronic Systems Laboratory

UNIVERSITEIT STELLENBOSCH UNIVERSITY

Problem Definition

- Create a method of racing that can plan without explicitly maintaining an updated obstacle map using SLAM.
- Specifically, the system must be able to avoid unknown obstacles in the environment.
- NB: must be practically possible.
- Two step problem
  - Avoid static obstacles
  - Avoid dynamic obstacles

UNIVERSITEIT • STELLENBOSCH • UNIVERSITY

Solutions

• • •

Electronic Systems Laboratory

UNIVERSITEIT STELLENBOSCH UNIVERSITY

Solution Framework

- Reinforcement learning is theoretically able to learn from experience to map sensor readings to actions.
- Generally RL uses neural networks to approximate the state space and determine the relevant actions.

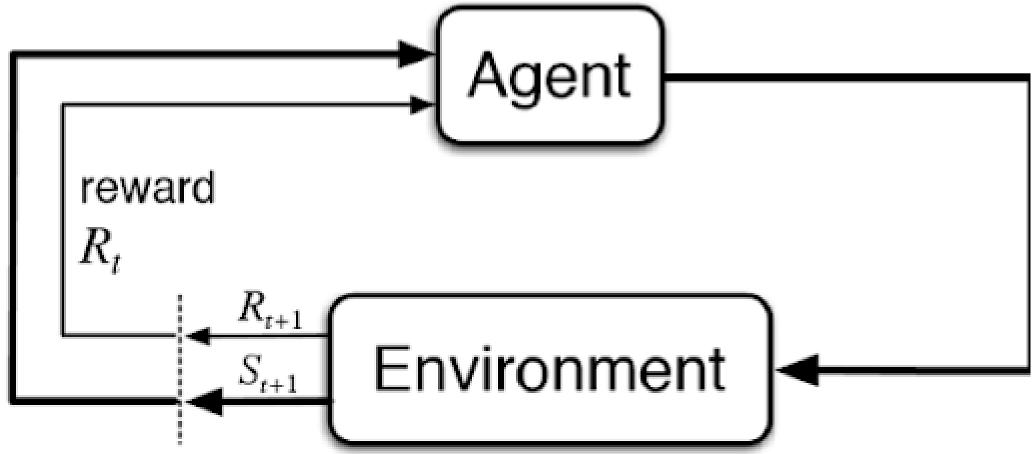

Electronic Systems Laboratory

Solution 1: Parallel agent

Solution 1: RL Agent in Parallel

- We use the local planner with the known track map that assumes that there are no obstacles
- We then send all that as the state to the agent which can modify the local planner references.
  - Either the magnitude or heading

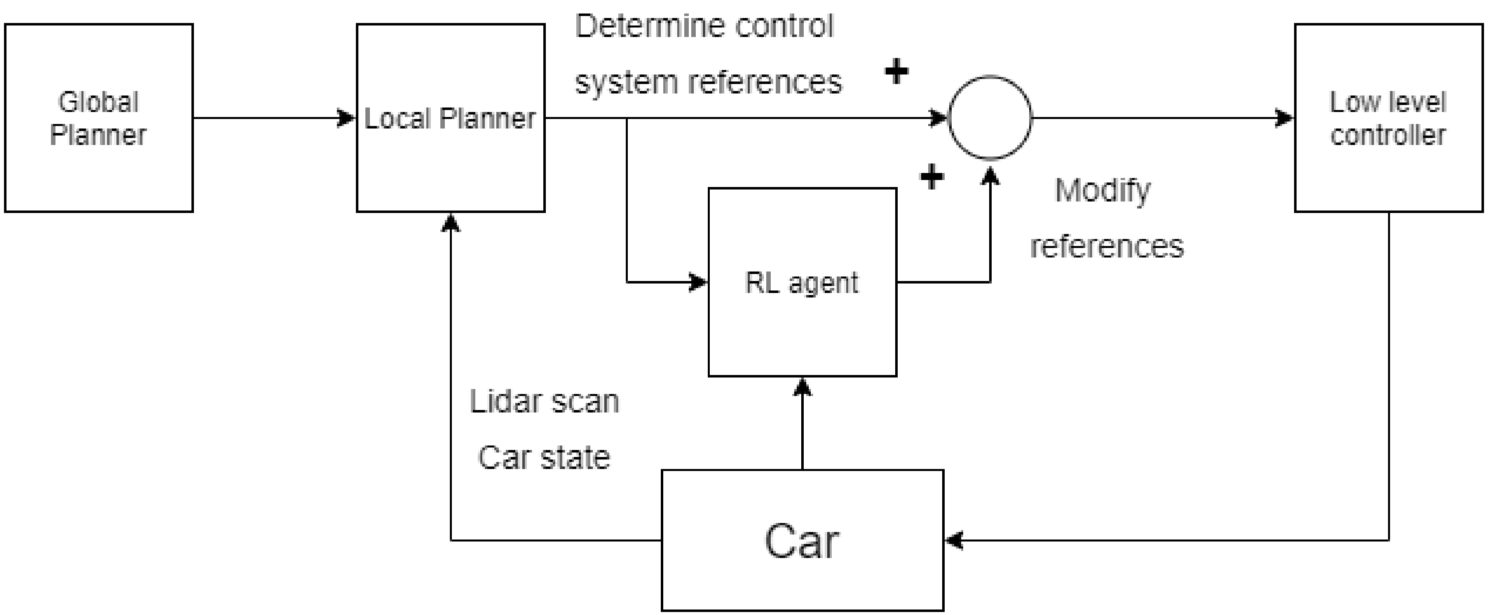

UNIVERSITEIT STELLENBOSCH UNIVERSITY

Solution 1:
Parallel agent

Training of Parallel Agent

- Aim is to encourage it to swerve just enough from the local planner to not crash.
- Training environment
  - Custom built env of a room with obstacles
  - Train on an actual racetrack
- Reward (scaled to [-1, 1]):
  - -1 for crashing
  - -beta * deviation from local planner
- Train in simulation because we can't afford to crash real cars.

Solution 2: Replacement Agent

- We replace the local planner with a RL agent
- The RL agent gets the lidar readings in and determines magnitude and velocity references.

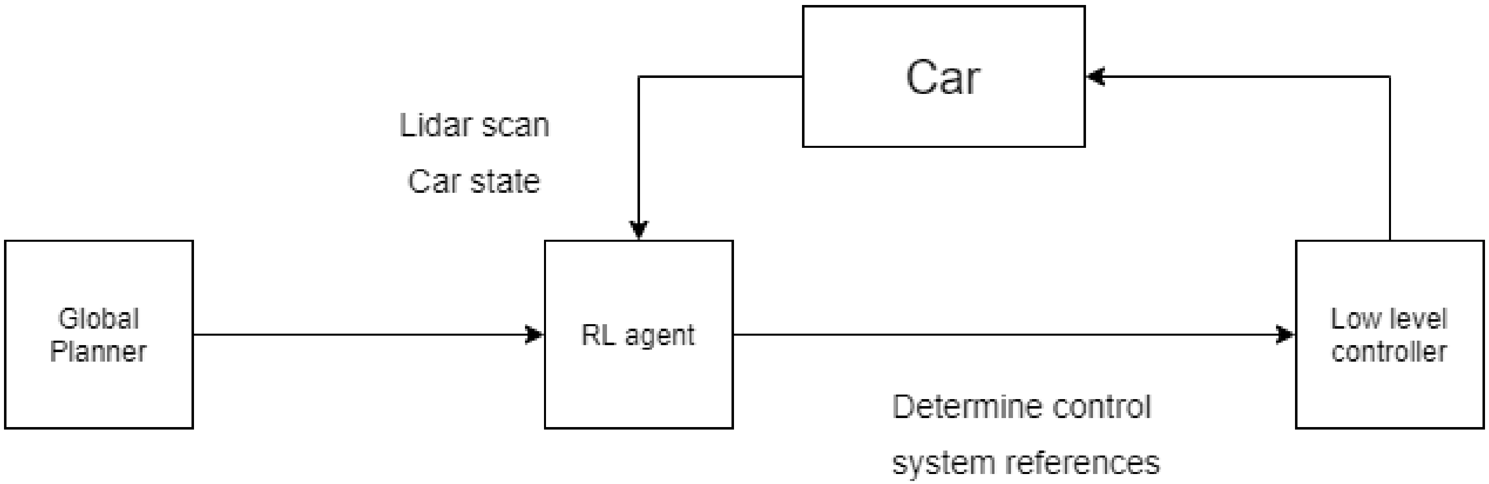

Electronic Systems Laboratory

Solution 2: Replacement agent

Training of Replacement Agent

- Generate trajectories using map with known obstacles and optimal control.
- Use supervised (imitation) learning to learn to imitate the controller.
- State:
  - Relative global plan target
  - Current velocity
  - Lidar Scan
- Action space:
  - Continuous speed and heading references
- Train in simulation or in practice because we are following an optimal control strategy in training
  - Use additional external sensors that are able to supply a perfect map of the environment

UNIVERSITEIT STELLENBOSCH UNIVERSITY

Electronic Systems Laboratory

Current progress

Current Progress

- Plan
  - Build solution architectures in a simple simulator
  - Upgrade simulation to use ROS and Gazebo
  - Run physical tests
- Current result
  - Replacement Architecture
  - Learning to avoid obstacles
- Next step
  - Improve simulation quality

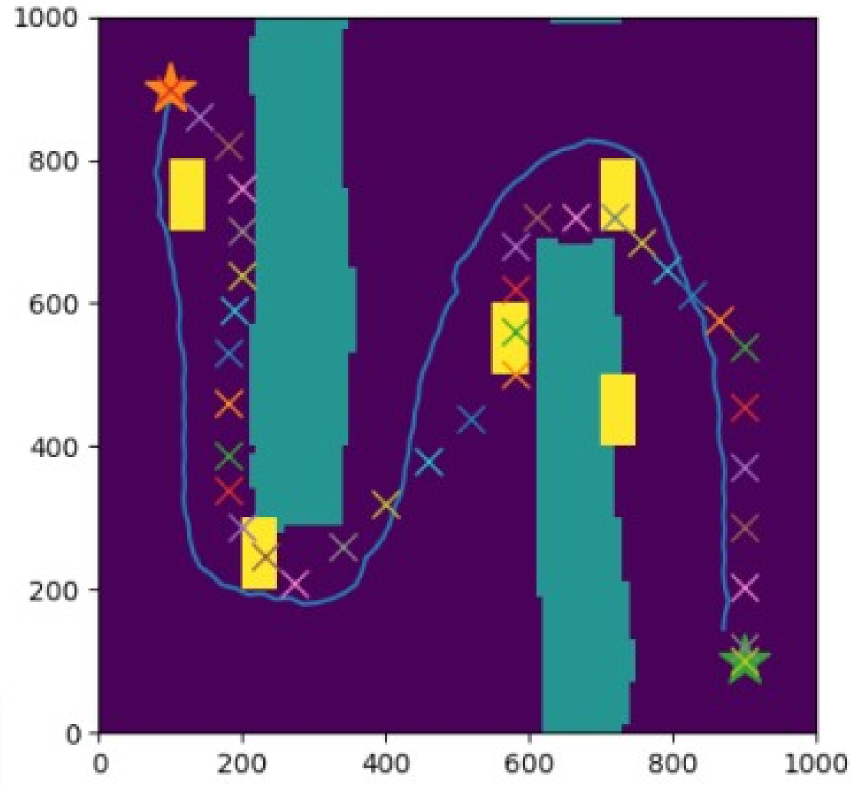

UNIVERSITEIT STELLENBOSCH UNIVERSITY

Electronic Systems Laboratory

The End

Questions?

UNIVERSITEIT • STELLENBOSCH • UNIVERSITY

Department of Electrical & Electronic Engineering

UNIVERSITEIT
STELLENBOSCH
UNIVERSITY