Università della Svizzera italiana

Software Institute

Towards Anomaly Detectors that Learn Continuously

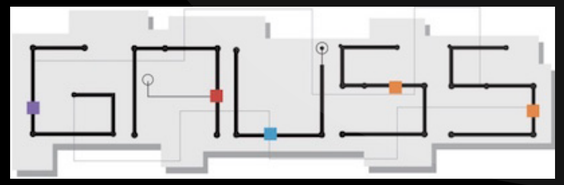

Andrea Stocco

@tsigalko18

Paolo Tonella

@paolo_tonella

- Lane keeping & detection
- Object recognition
  - vehicles
  - pedestrians
- Traffic Sign Recognition

Cruise uses Continual Learning to predict intent during on-road driving

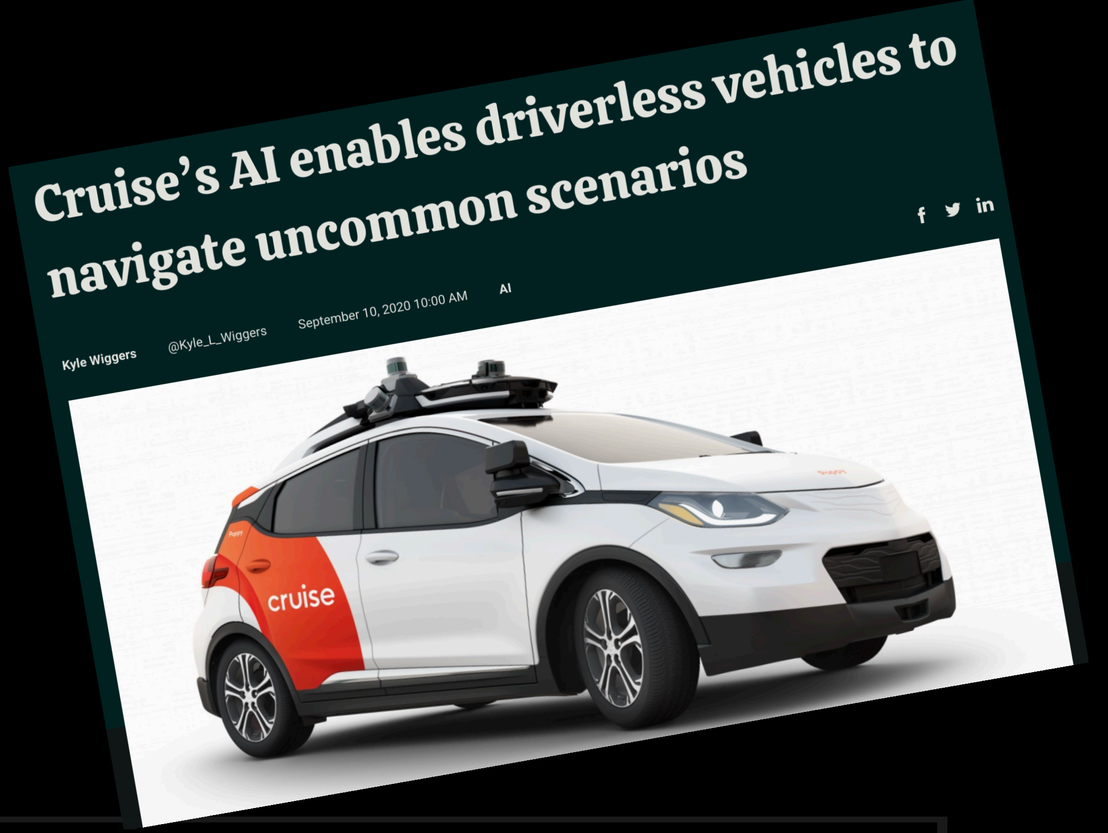

Tesla continually records in-field data to run shadow testing of their autopilot

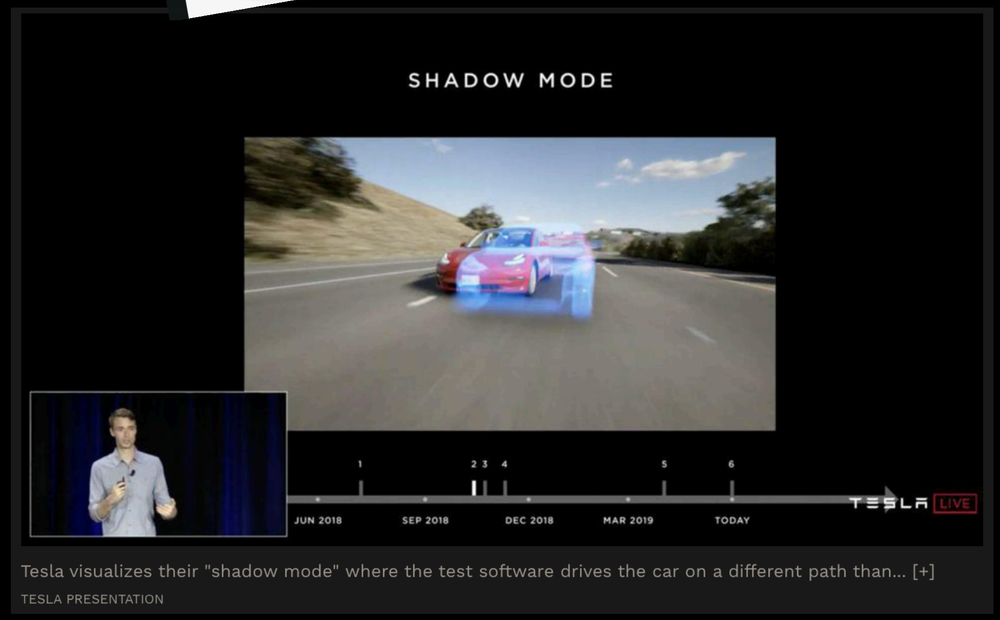

Sunny (nominal conditions)

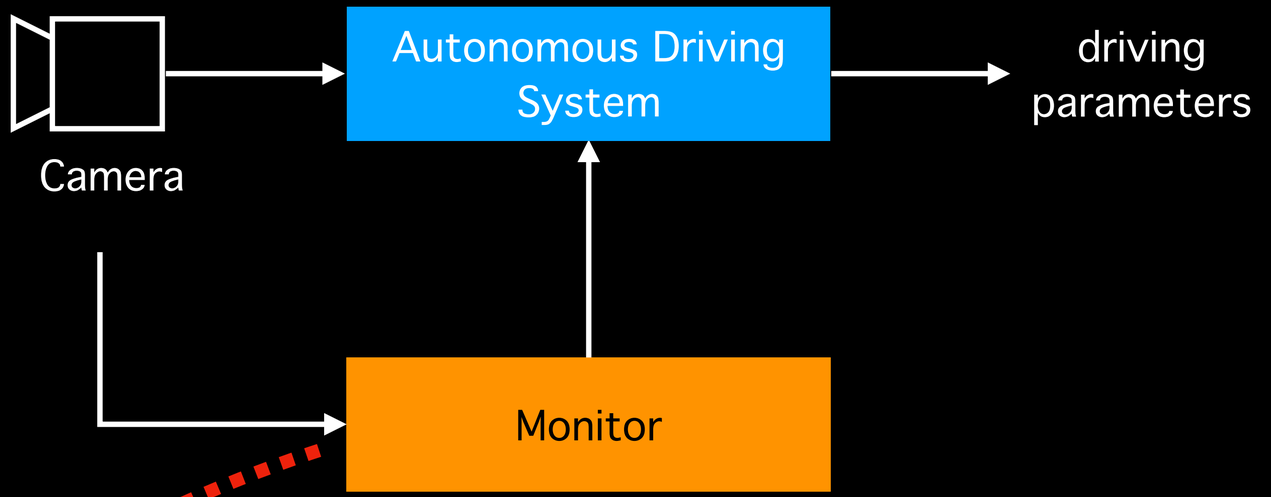

Stocco et al. Misbehaviour Prediction for Autonomous Vehicles. ICSE '20.

Night + Snow (unseen conditions)

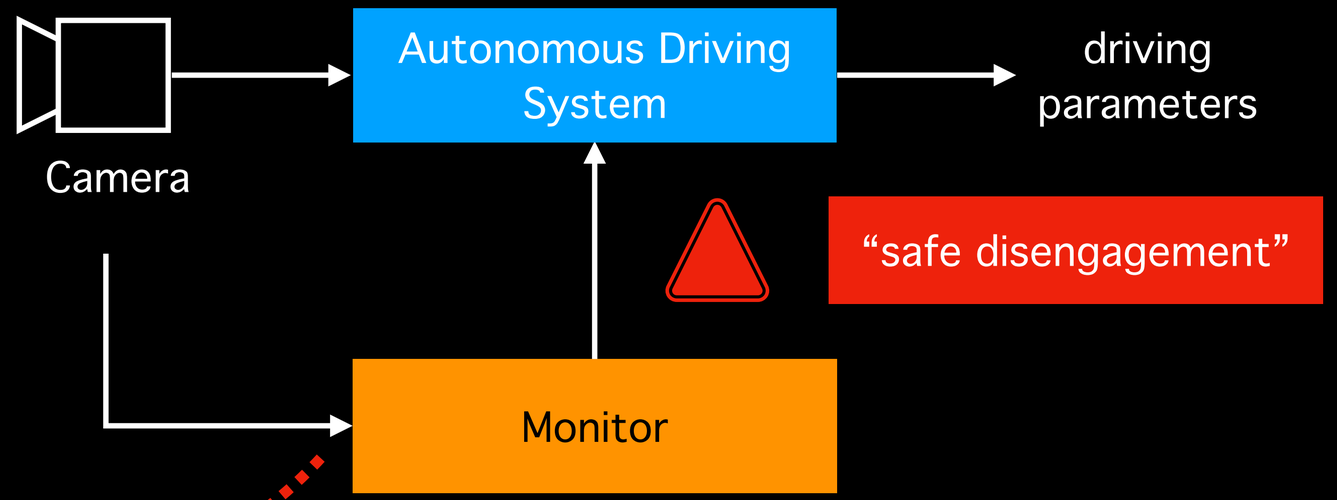

Stocco et al. Misbehaviour Prediction for Autonomous Vehicles. ICSE ’20.

underrepresented inputs (bridge)

novel classes of data (light rain)

FP = context is nominal &&
anomaly detector raises an alarm

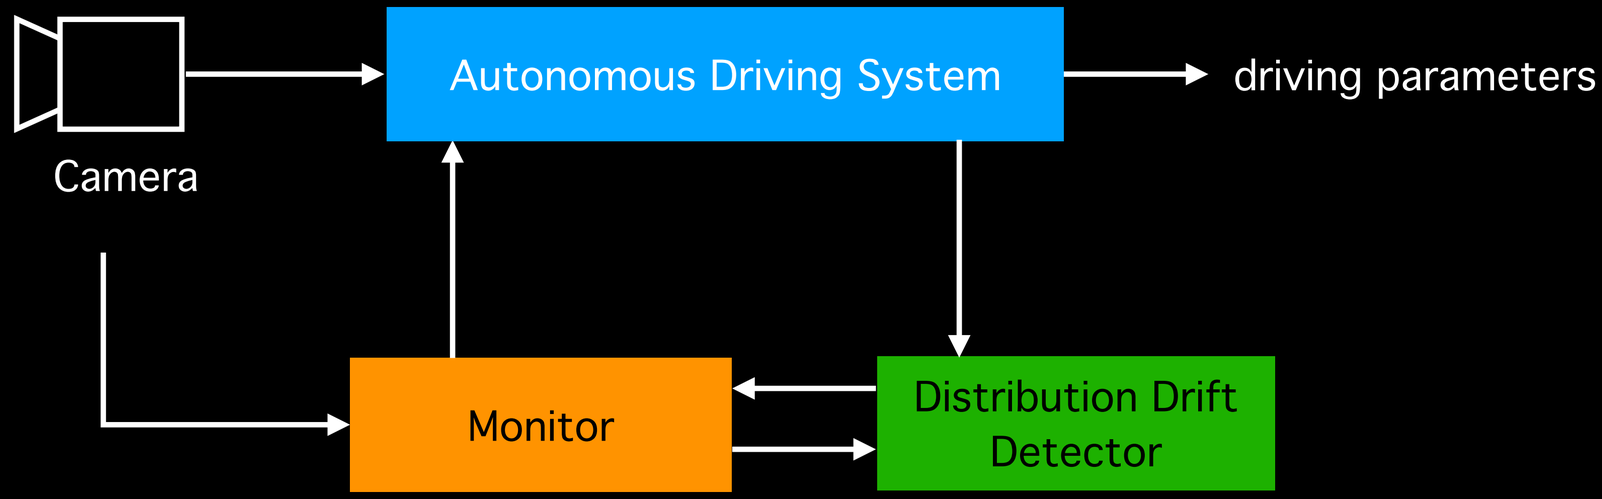

Can we use in-field data to train a better monitor ?

Udacity sim

Dave-2 NVIDIA

Monitor (SelfOracle)

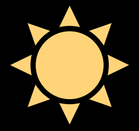

Distribution Drift Detector

Predictive Uncertainty (Unc)

Lateral Position (CTE)

LFP = Unc/CTE is good &&
anomaly detector raises an alarm

Two Experiments

Novelty Detection

Class Imbalance

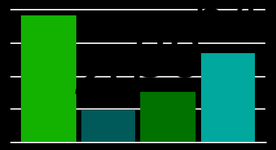

Metric

% detected LFP

before and after adaptation

Udacity sim

Dave-2

Monitor (SelfOracle)

Distribution Drift Detector

Predictive Uncertainty (Unc)

Lateral Position (CTE)

LFP = Unc/CTE is good &&
anomaly detector raises an alarm

Two Experiments

Novelty Detection

Class Imbalance

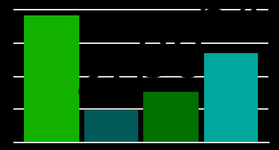

Metric

% detected LFP

before and after adaptation

Udacity sim

Dave-2

Monitor (SelfOracle)

Distribution Drift Detector

Predictive Uncertainty (Unc)

Lateral Position (CTE)

LFP = Unc/CTE is good &&
anomaly detector raises an alarm

Two Experiments

Novelty Detection

Class Imbalance

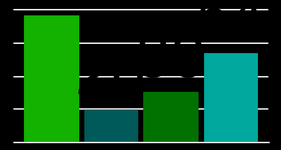

Metric

% detected LFP

before and after adaptation

Udacity sim

Dave-2

Monitor (SelfOracle)

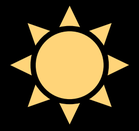

Distribution Drift Detector

Predictive Uncertainty (Unc)

Lateral Position (CTE)

LFP = Unc/CTE is good &&
anomaly detector raises an alarm

Two Experiments

Novelty Detection

Class Imbalance

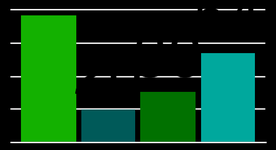

Metric

% detected LFP

before and after adaptation

Can we use in-field data to train a better monitor ?

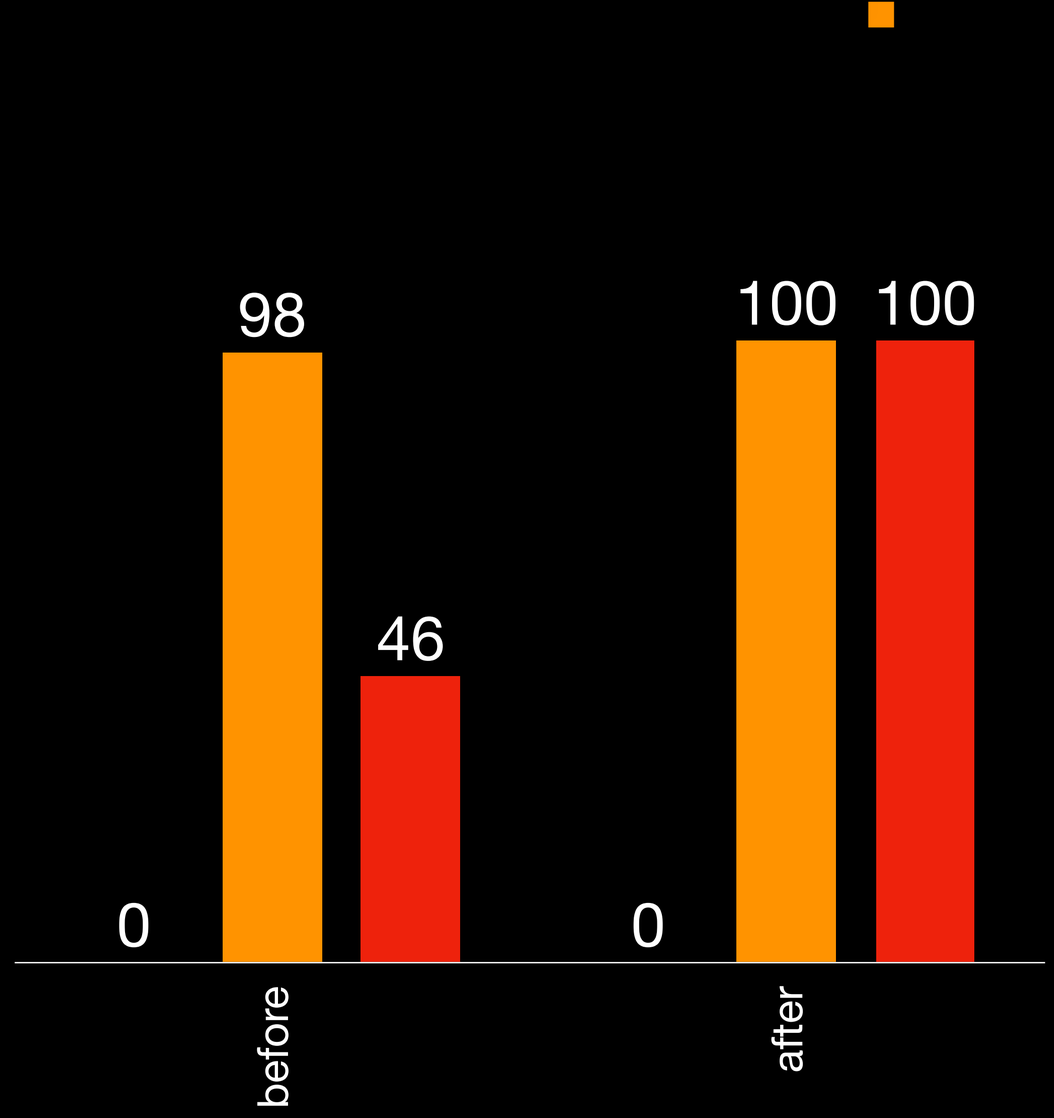

Class Imbalance

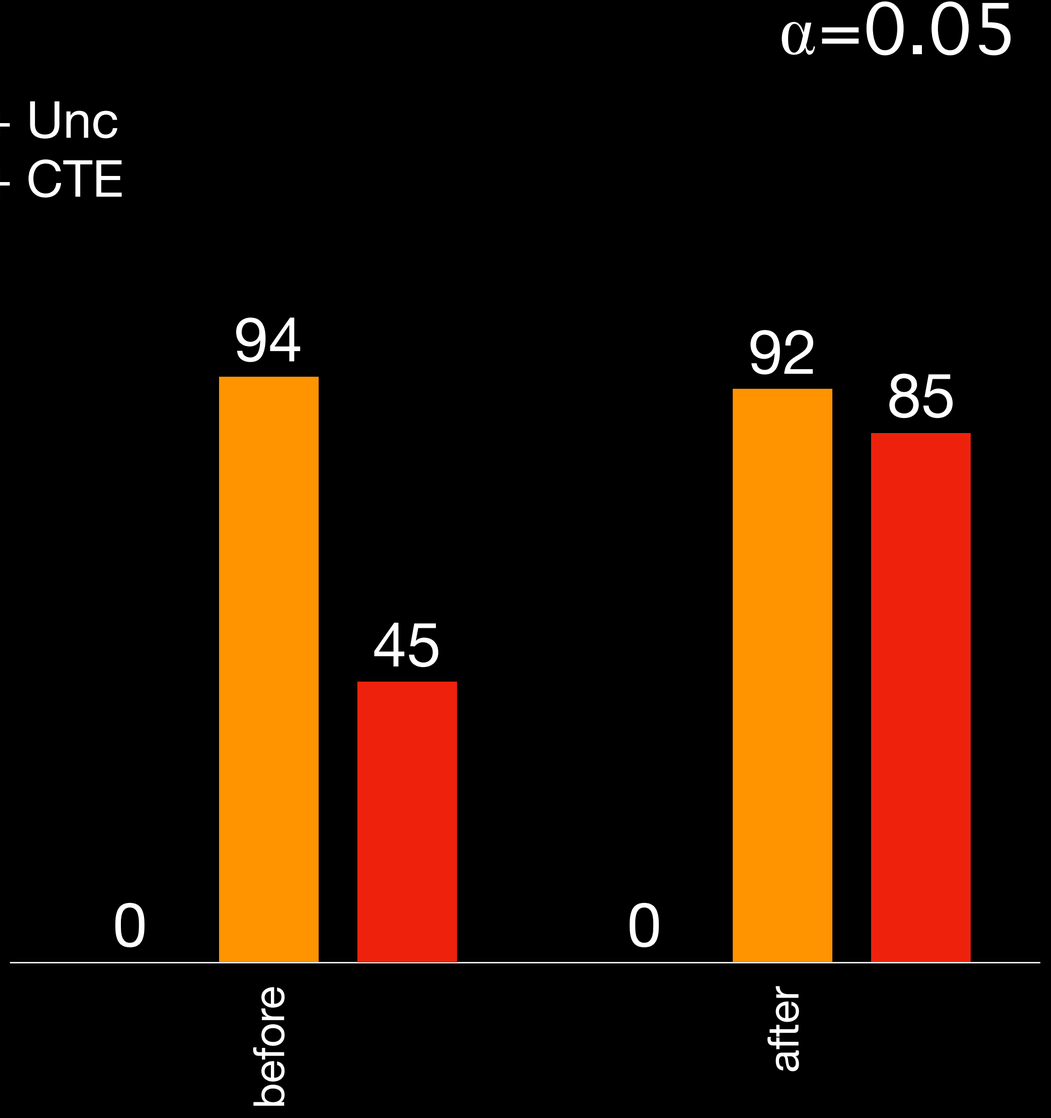

Novelty Detection

Università della Svizzera italiana

Software Institute

Towards Anomaly Detectors that Learn Continuously

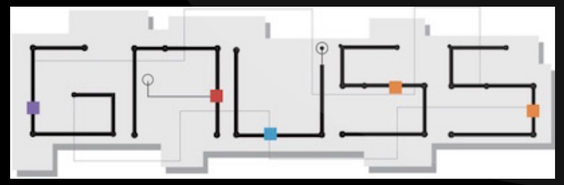

Andrea Stocco

@tsigalko18

Paolo Tonella

@paolo_tonella