

# Impact of Machine Learning on Safety Monitors

Francesco Terrosi

francesco.terrosi@unifi.it

University of Florence – Italy

Andrea Bondavalli

andrea.bondavalli@unifi.it

University of Florence – Italy

Lorenzo Strigini

l.strigini@city.ac.uk

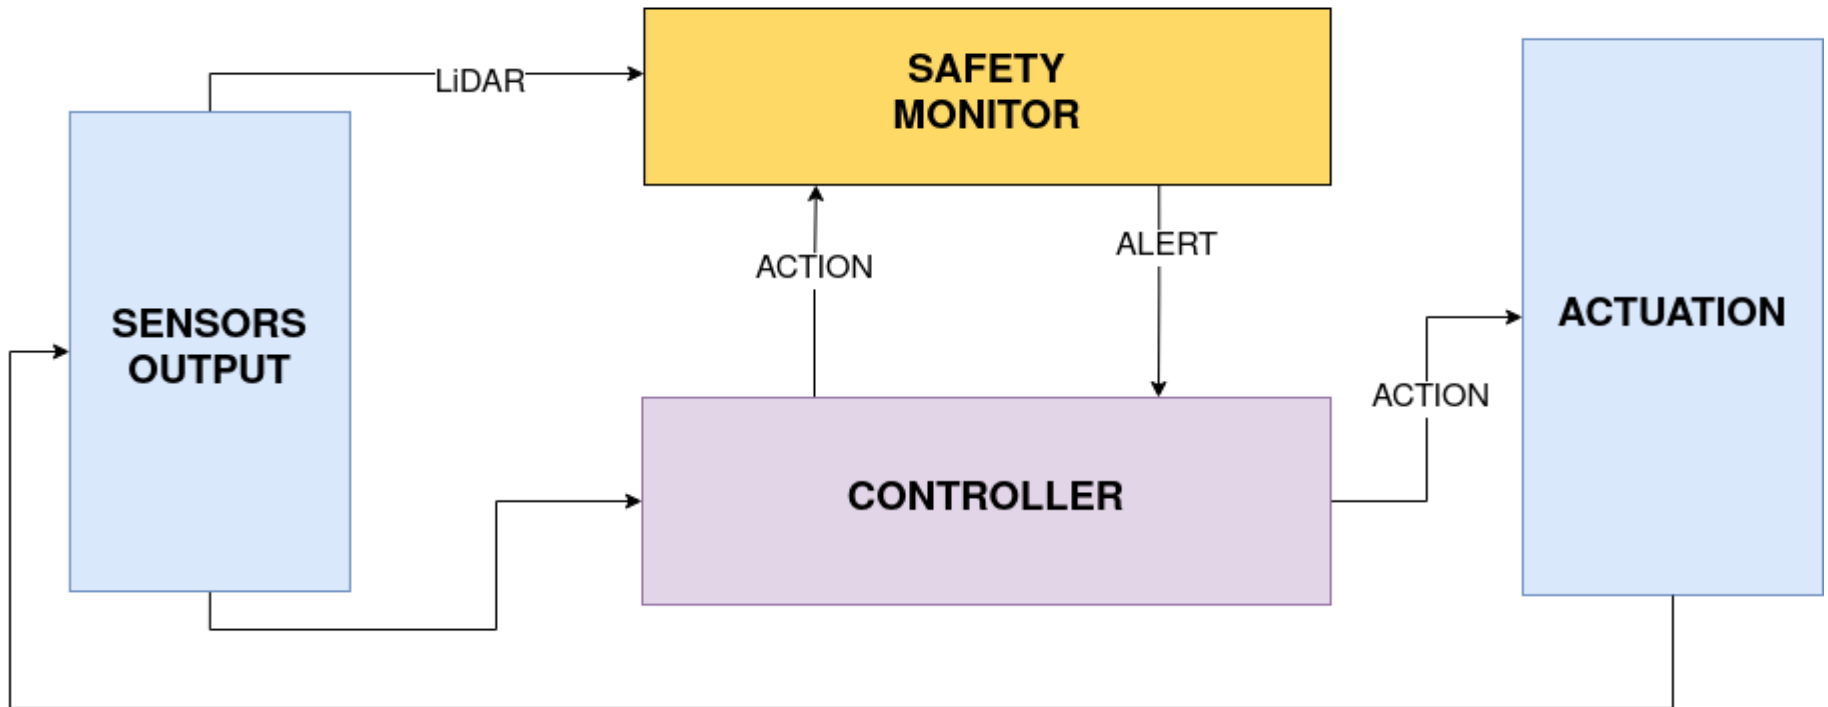
City, University of London, London



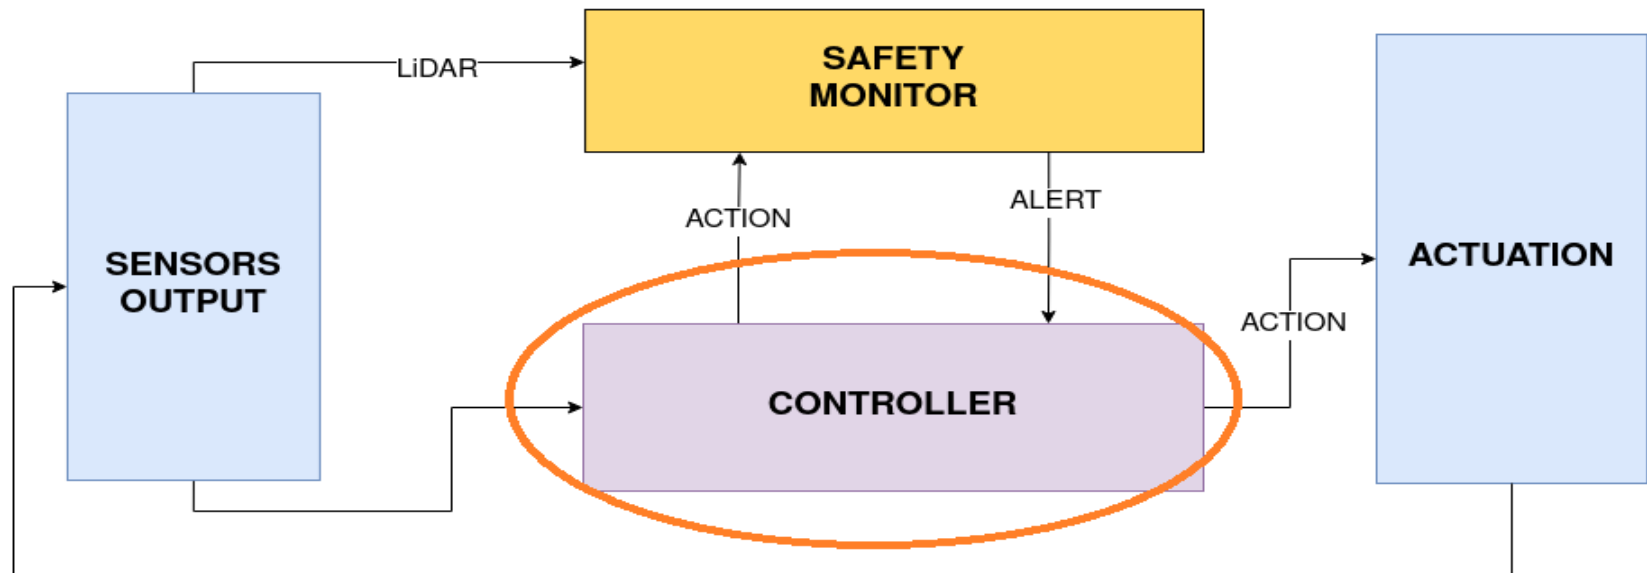
# Autonomous Vehicles - Overview

- AVs are critical systems with ultra-high dependability requirements
- Working at very close contact with humans
  - Need to minimize risk associated to “wrong actions” done by the vehicle
- Assessing their safety is mandatory, but far from easy

## Standard AV Architecture:

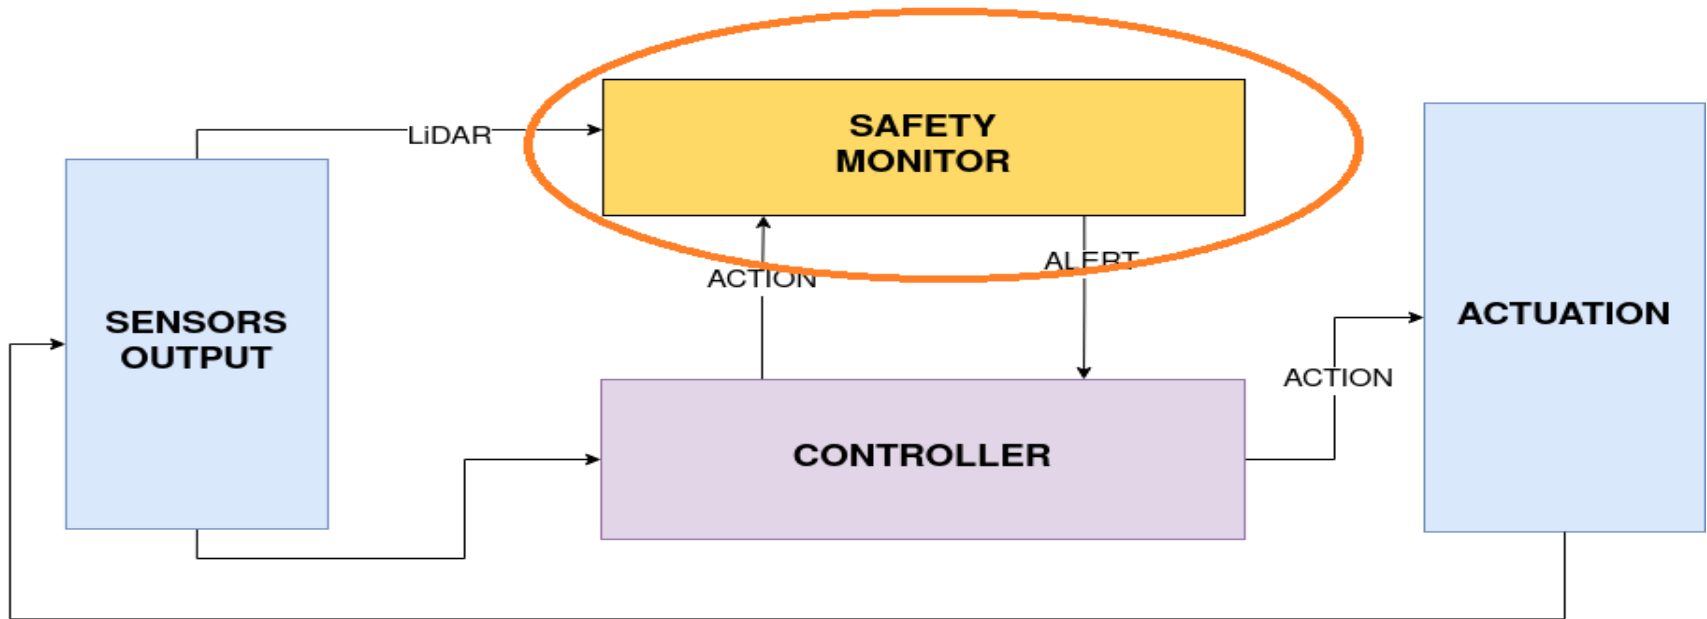


# Autonomous Vehicles - Overview



The Controller is a neural network (or an ensemble of) trained to drive. It is the primary component of the system i.e., its task is to make the car move safely obeying the laws of traffic

# Autonomous Vehicles - Overview



A Safety Monitor performs safety checks on the actions chosen by the Controller, on the basis of independent elaborations of sensors data.

In contrast with the Controller, it can't be trained. Its actions can change only if reimplemented



# Controller – Monitor Problem

- ▶ However, it's theoretically possible that the performances of the Safety Monitor changes as a result of the training of the Controller!
- ▶ In one case, the trained Controller learns to handle state that were not covered by the Safety Monitor, thus improving the global safety of the system.
- ▶ But it also may happen that during the training the Controller “concentrates” its failures in an area in which the Safety Monitor is useless, thus reducing the global safety of the system!



# Experimental Activity

- ▶ The Controller is a neural network trained with the Deep Deterministic Policy Gradient (DDPG) algorithm.
  - The implementation is provided by the “*Framework for Reinforcement Learning Coach*”, an open-source project developed by Intel’s AI Lab
  
- ▶ A Safety Monitor is in charge of detecting hazards
  - Processing LiDAR data, if the Controller breaks the minimum stopping distance rule the monitor will trigger a safety-brake

# Simulation Environment

- ▶ The vehicle was tested by simulation in CARLA, an open-source urban driving simulator







# Controller Testing

- ▶ The Controller is tested in “Scenarios” at different stage of the training.
  - All the commands generated by the neural network are recorded for *repeatability*.
  
- ▶ A “Scenario” is composed of a starting point, a set of target destinations and the seeds used in RNGs. Each scenario (150 in total) has 4 difficulties:
  1. Standard
  2. Increased amount of pedestrians
  3. Increased amount of vehicles
  4. Increased amount of pedestrians and vehicles



# Controller Testing

- ▶ If a *collision happens* or *all the target destinations are reached*, the test is stopped.
- ▶ Running simulations steps at a *fixed time-step* allows to compute measures such as Mean Time To Failure or Mean Distance Between Failures
- ▶ The Reward and Q function are good indicators in reinforcement learning to check whether the network is improving



# Safety Monitor Testing

- ▶ The runs previously recorded are repeated identically, but now the Safety Monitor is deployed on the system
- ▶ The use of a fixed time-step is mandatory to have a reference time: we can compute the *time t* necessary for the Monitor to prevent a collision, if happened
- ▶ The safety-brake is enabled only after *time t*, but the alerts raised are recorded during the run

# Safety Monitor Testing

► The Confusion Matrix for the Monitor's prediction is computed as follows:

- **True Positive (TP):** After the safety-brake is enabled, if the Safety Monitor prevents the crash is considered a true positive.
- **False Negative (FN):** Every collision not prevented is considered a false negative.
- **False Positive (FP):** Each iteration in which the Safety Monitor raised an alarm before time  $t$ .
- **True Negative (TN):** Each iteration before time  $t$  in which the Safety Monitor did not raise an alarm.



# Safety Monitor Testing

► From the confusion matrix we can compute many prediction metrics such as:

– Precision (P) =  $\frac{TP}{TP+FP}$

– Recall (R) =  $\frac{TP}{TP+FN}$


– Accuracy (A) =  $\frac{TP+TN}{P+N}$

– Matthew's Correlation Coefficient (MCC) =

$$\frac{TP*TN-FP*FN}{\sqrt{(TP+FP)*(TP+FN)*(TN+FP)*(TN+FN)}}$$


# Results - Controller

► The Controller was tested at 5 different stages  $C_1 \dots C_5$





# Results – Controller



# Results – Safety Monitor

- The MCC is decreasing: the more the Controller is trained the more the Monitor is close to random guessing


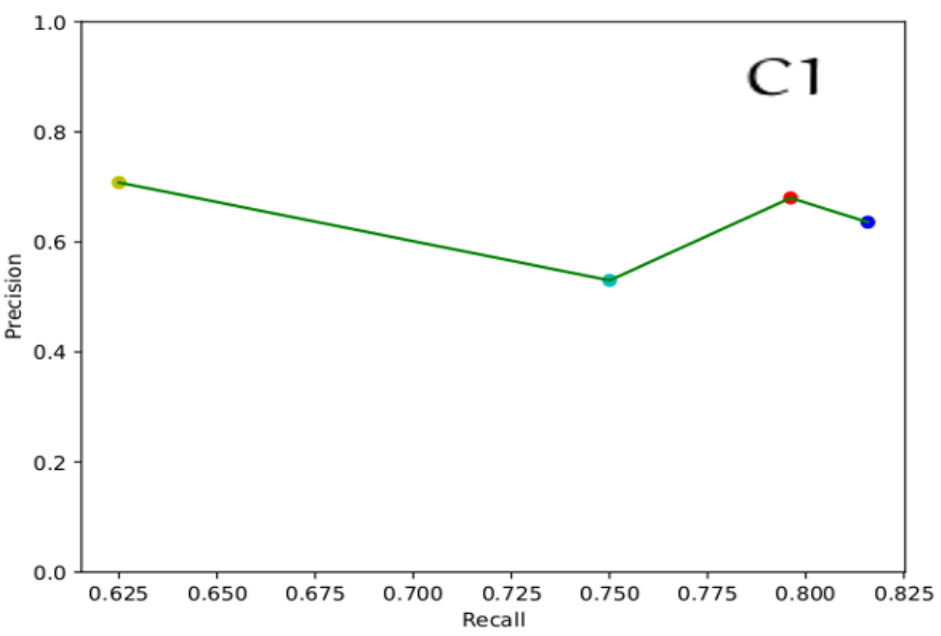
	$C_1$	$C_2$	$C_3$	$C_4$	$C_5$
Recall (TP rate)	0.75	0.82	0.75	0.77	0.74
FP rate	0.0029	0.0029	0.0062	0.0069	0.0058
A	0.995	0.992	0.993	0.992	0.994
P	0.63	0.14	0.17	0.15	0.10
MCC	0.68	0.34	0.35	0.34	0.28





# Results – Safety Monitor

## Precision-Recall Curves comparison for Controllers $C_1$ and $C_5$



- Standard
- Pedestrians
- Vehicles
- Pedestrians & Vehicles



# Consequences

- ▶ There are many example of crashes that makes safety monitors mandatory.
- ▶ However there are report of minor incidents related to false alarms.
- ▶ Arizona accident, 2018. The safety-brake was disabled to "reduce the potential for erratic vehicle behaviour".



# Conclusions

In this work we presented an experimental method putting emphasis on the *emergence* that can be observed by combining Machine Learning and “static” software

The common architectures that pair machine learning components with Safety Monitors need joint empirical validation of the whole system

This work also puts emphasis on the *need for design and analysis techniques* that address this issue effectively