

Deep Reinforcement Learning for autonomous driving

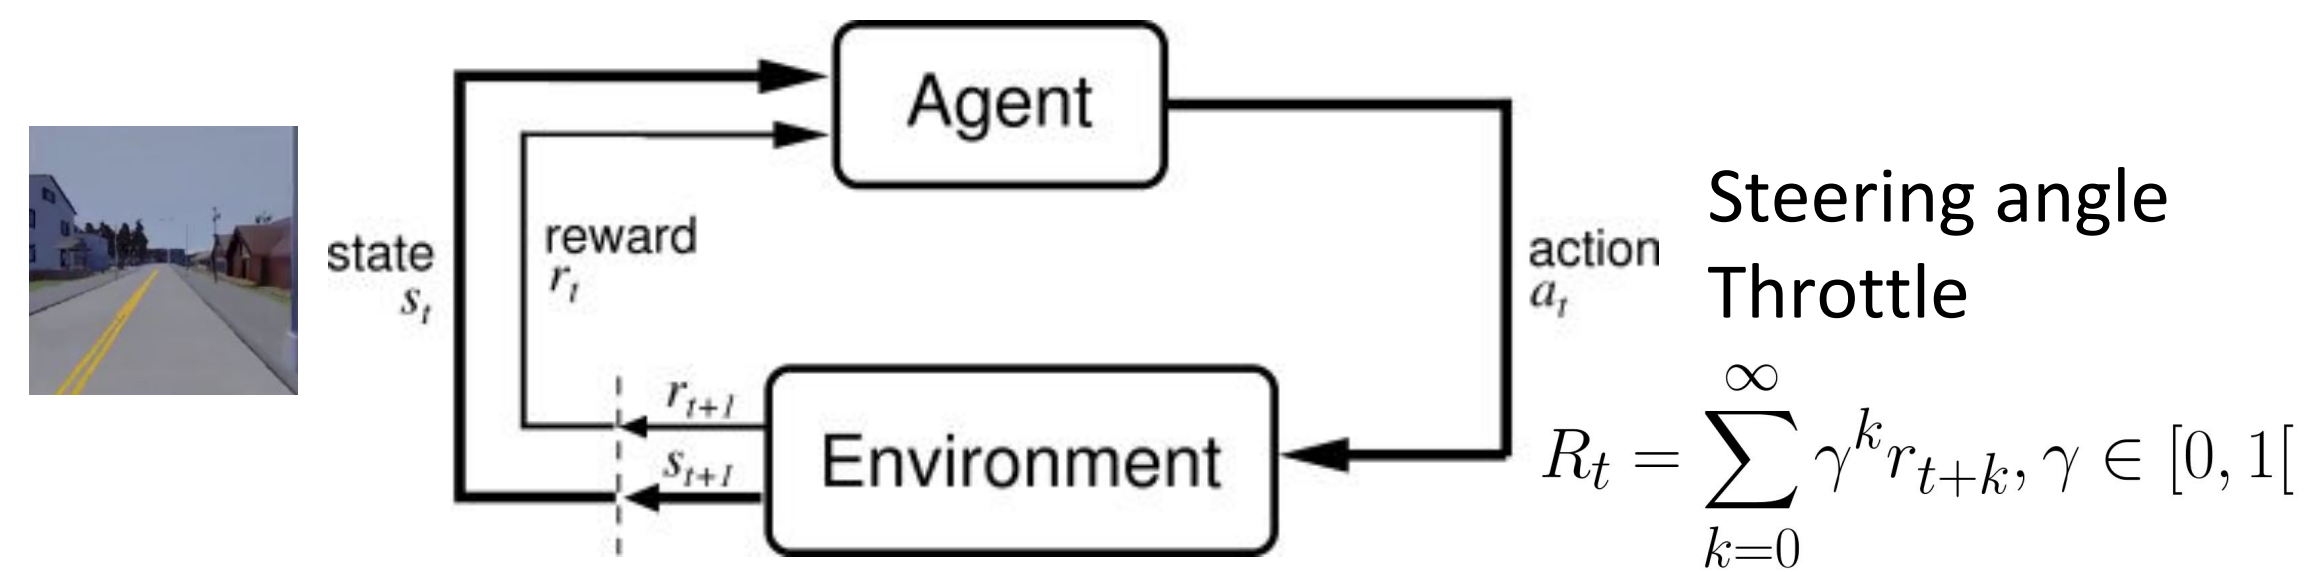
Marin Toromanoff

Phd Student at Mines ParisTech (Pr F. Moutarde) and Valeo (Dr E. Wirbel)
marin.toromanoff@mines-paristech.fr




Reinforcement Learning Setup:

- Agent tries to maximize sum of accumulated rewards



New algorithm developed, Rainbow-IQN:

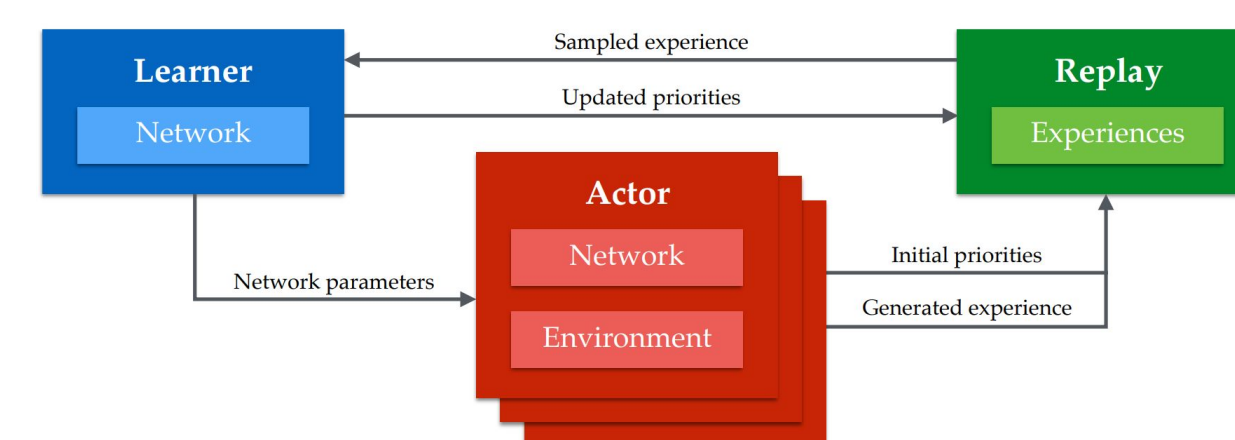
- Rainbow [1]:**
 - Current state-of-the-art on Atari
 - Combination of 6 improvements over DQN [2]
 - Include Distributional RL, C51 algorithm [3]



- Implicit Quantiles Network (IQN) [4]:**
 - Improvement over C51
 - Predict quantiles

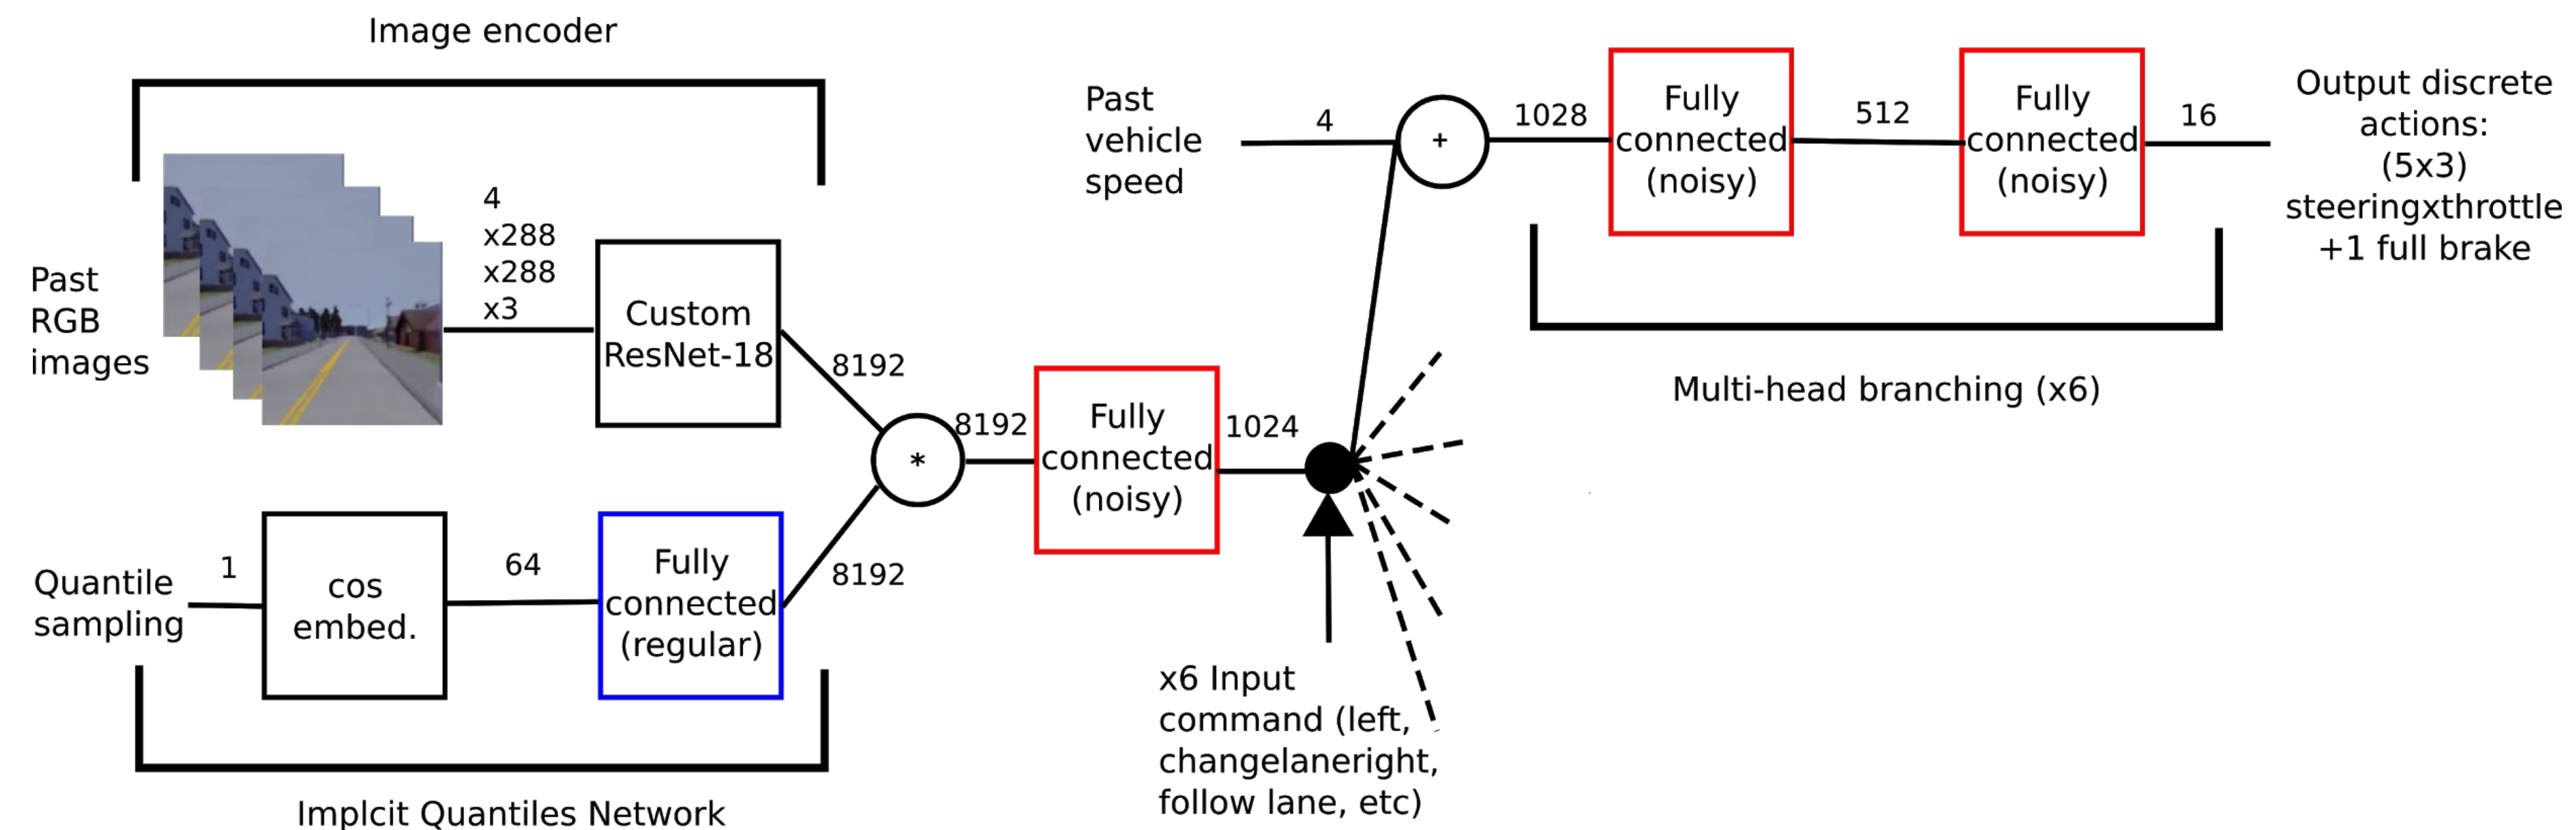
Making it distributed, Rainbow-IQN Ape-X:

- Distributed Prioritized Experience Replay (Ape-X) [5]:**
 - Multi-agent training
 - Allow multi-town training




Applying Deep Reinforcement Learning (DRL) to autonomous driving: Network architecture

- Major issue: Traffic light**
 - DQN-like network take 84x84 grayscale image as input
 - Necessity of using way bigger network
 - Use a resnet-18 (10 times more weights)
- Handle orders with multi-head branching [6]



Applying DRL to autonomous driving: Reward shaping

- Reward speed:**
 - Maximum when reaching desired speed
 - Desired speed adapts to situation
- Reward lateral:**
 - Maximum when center of the lane
 - Agent death when too far from lane center
- Episode termination:**
 - Collision on pedestrians, vehicles, others
 - Running red light
 - Too far from lane center
 - Stuck (if no reason to stop)



Challenges solved but not detailed

- Network way bigger (for DRL):**
 - How to make RL converge fast enough?
- Replay memory size:**
 - Set to 1M for DQN-like algorithm
 - Images used are 30 times bigger than previously used
- Value-based RL:**
 - discrete actions (Q-Learning algorithm)
 - stabilization of prediction need
 - lot of oscillation still remains

[1] M. Hessel et al., *Rainbow: Combining improvements in deep reinforcement learning*, 2018
 [2] V. Mnih et al., *Playing Atari with Deep Reinforcement Learning*, 2013
 [3] M. G. Bellemare et al., *A distributional perspective on reinforcement learning*, 2017
 [4] W. Dabney et al., *Implicit quantile networks for distributional reinforcement learning*, 2018
 [5] Horgan, J. Quan et al., *Distributed prioritized experience replay*, 2018
 [6] Codevilla et al., *End-to-end driving via Conditional Imitation Learning*, 2017