End-to-End Driving in a Realistic Racing Game with Deep Reinforcement Learning

[1] Valeo [2] Inria

Etienne Perot[1] Maximilian Jaritz[1,2] Marin Toromanoff[1] Raoul de Charette[2]

OVERVIEW

Learning direct control from image

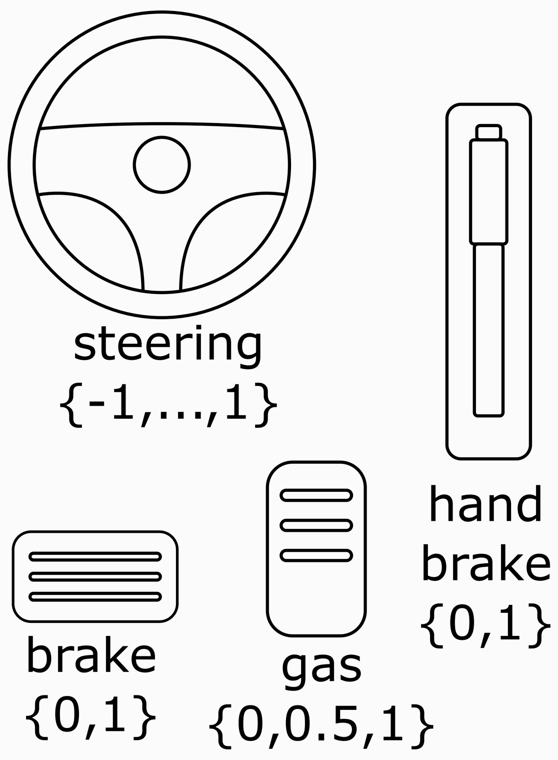

Challenges

- Full control
- Realistic graphics, complex dynamics
- Racing setup: fast driving, drifting

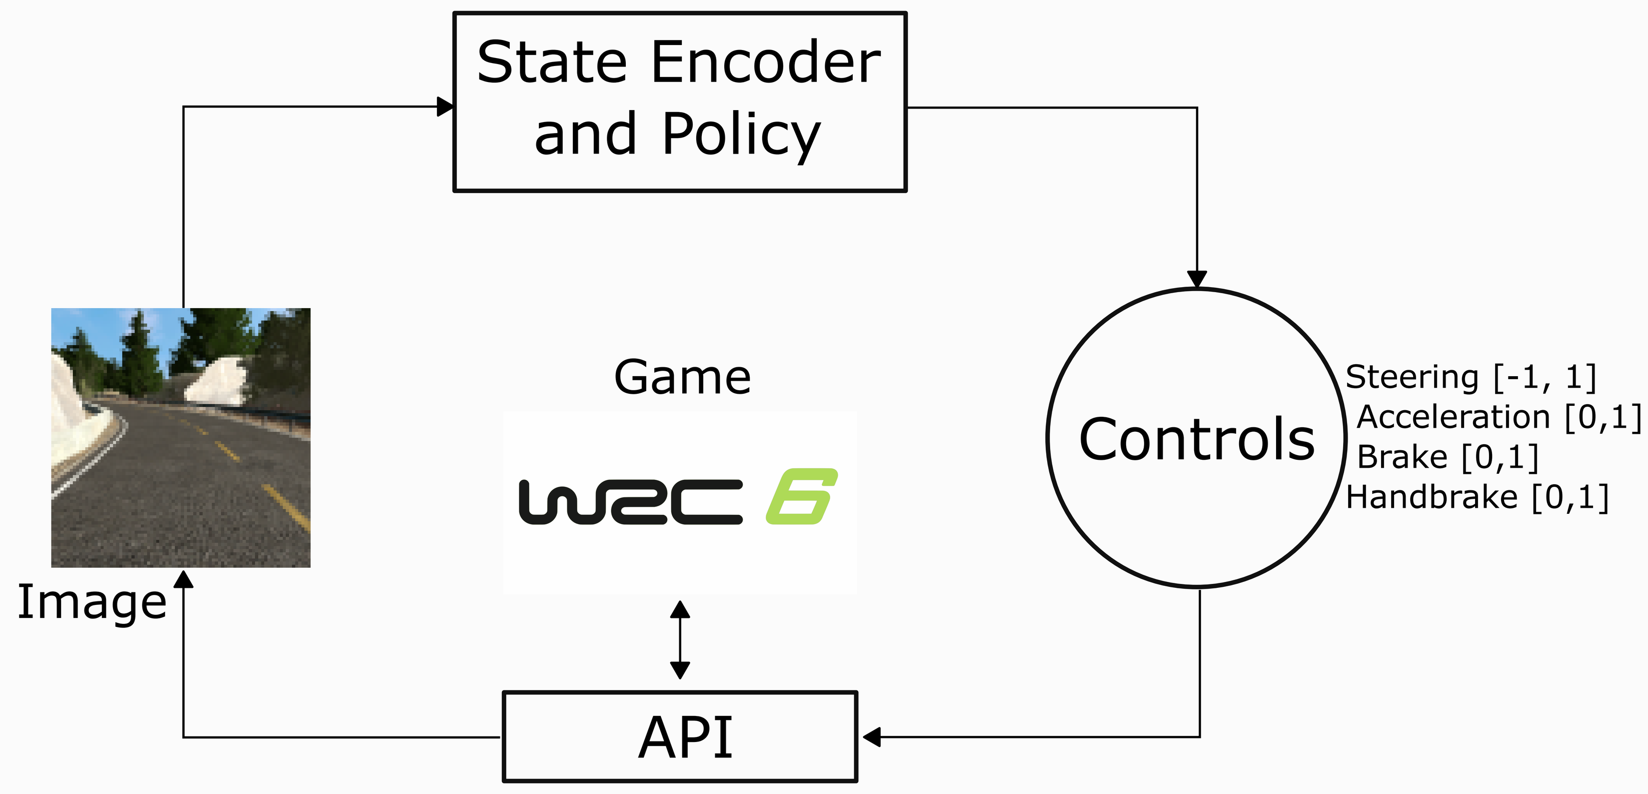

Dedicated API

- Receive image, speed and angle, send controls

METHODOLOGY

Related works

Deep reinforcement learning [1]

TORCS

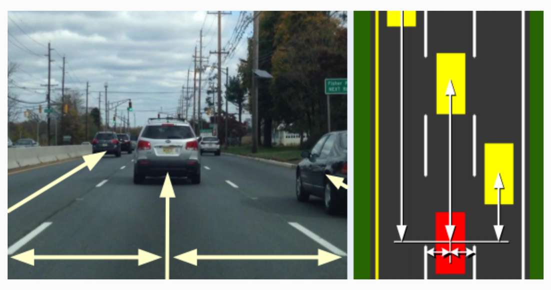

Direct perception [2]

TORCS

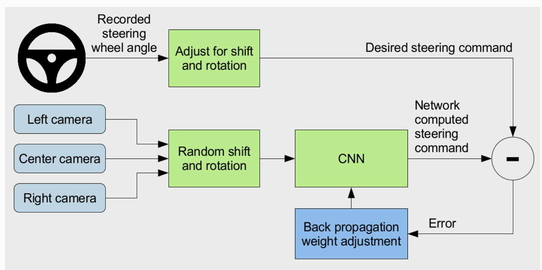

Behavioral cloning for lateral control [3]

Transformed real images

State Encoder

- LSTM takes additional inputs (previous action, current speed)
- Outputs action probabilities (softmax)

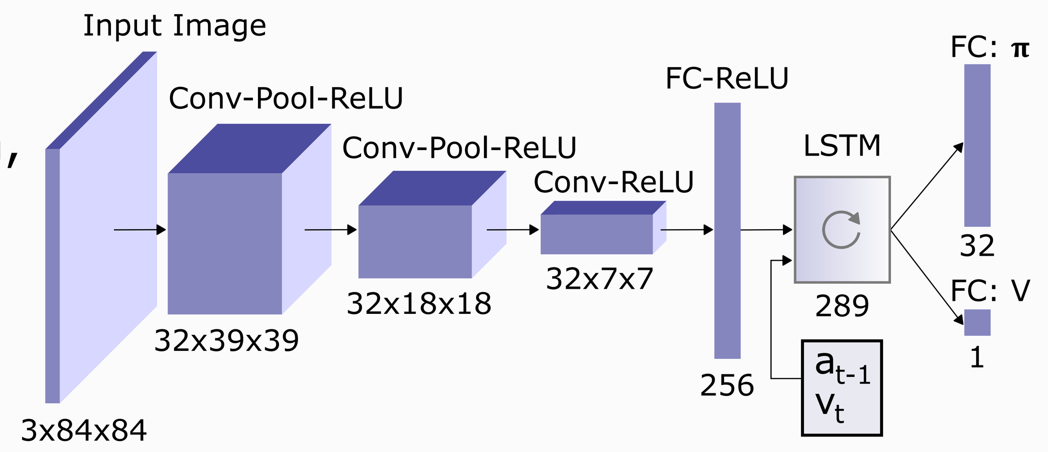

A3C Algorithm [1]

- Asynchronous: No experience replay needed
- Multi-agent: Parallel computing

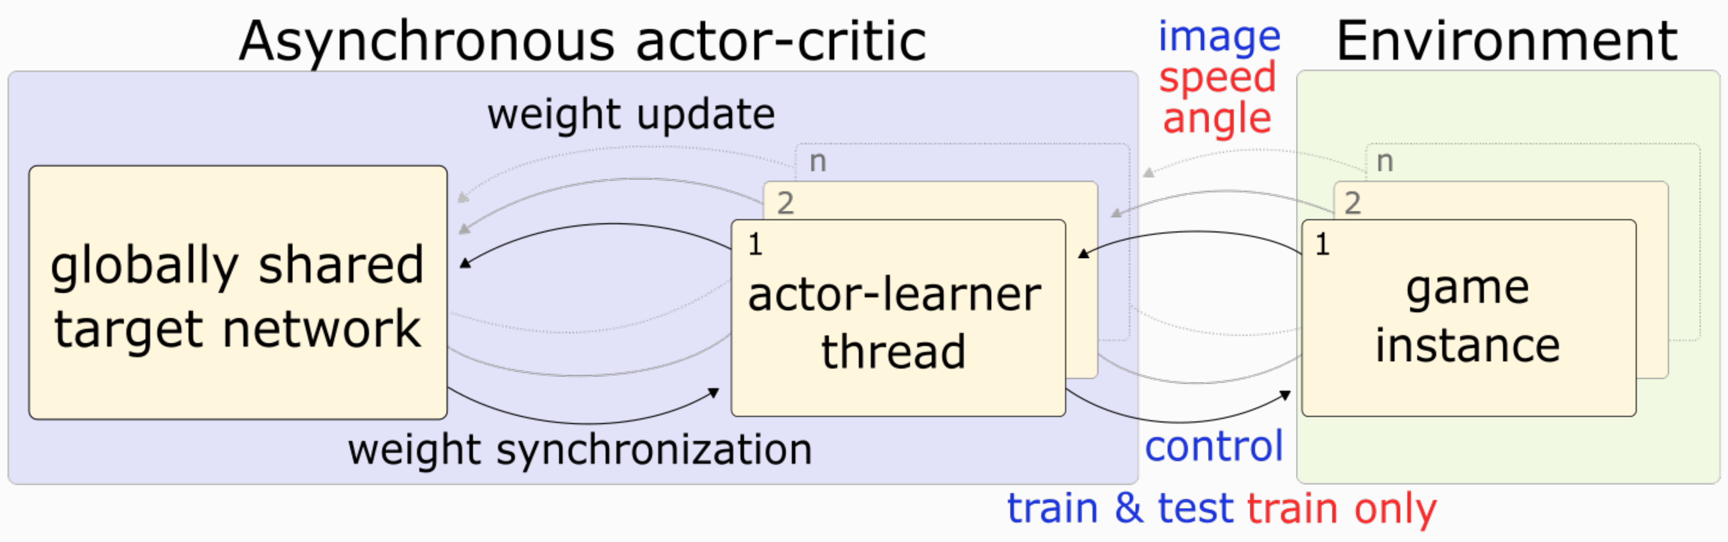

Reward

- Frame-wise, function of speed and angle
- New: add distance from the middle of the track as penalty

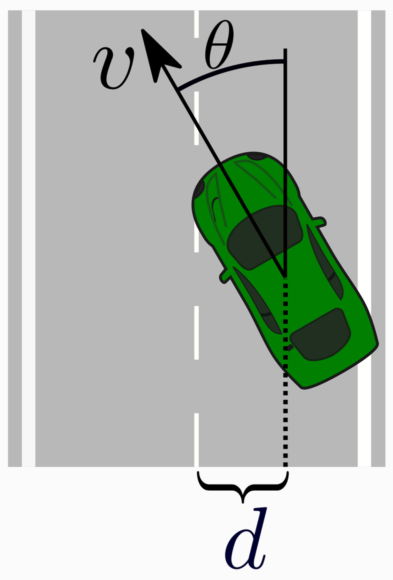

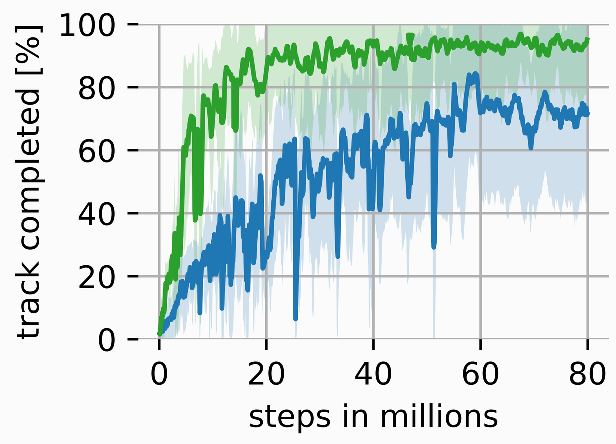

Reward used
$R = v(\cos\theta - d)$

VS

Mnih et al. [1]
$R = v\cos\theta$

EXPERIMENTS

Joint training on 3 tracks with different graphics and physics

Task specific evaluation metrics

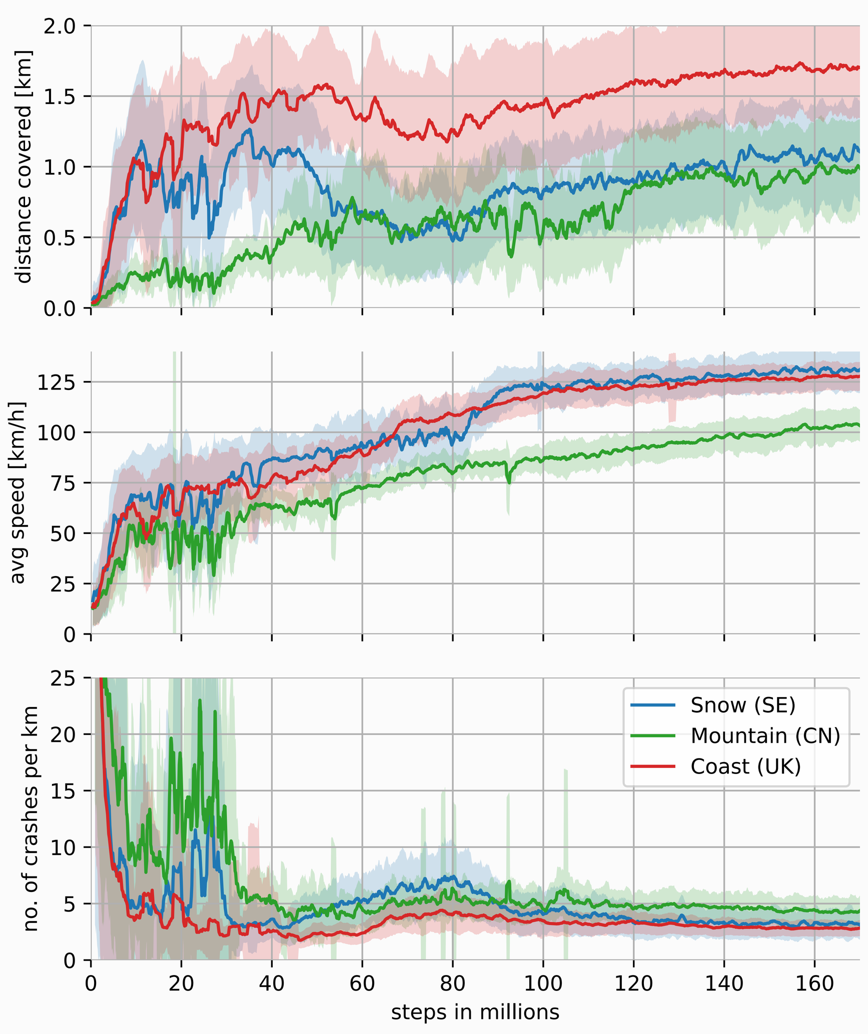

Performance after 170 million steps

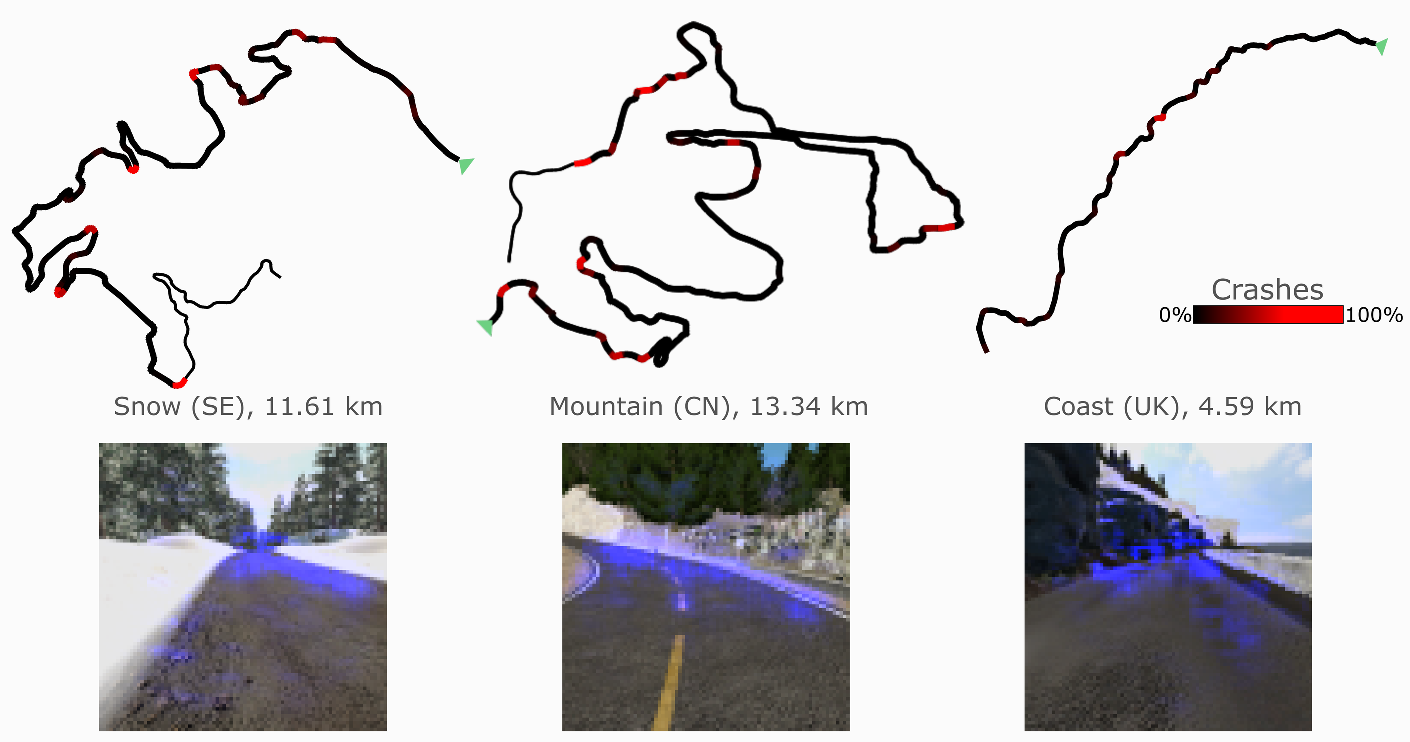

- Trained with 15 agents (5 per track)
- Average run: 121 km/h, 1.29 km, 2.5 crashes/km

Test on unseen track

Generalization

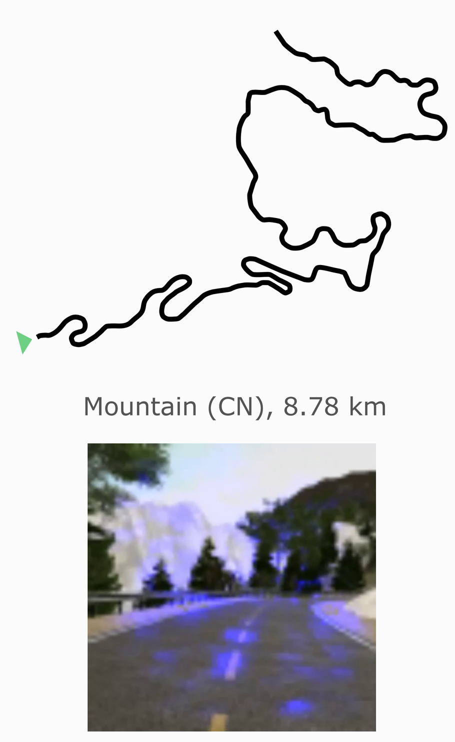

- Qualitative performance in video

VIDEO

https://youtu.be/e9jk-IBWFlw

[1] Mnih et al., Asynchronous methods for deep reinforcement learning, ICML 2016
[2] Chen et al., DeepDriving: Learning affordance for direct perception in autonomous driving, ICCV 2015
[3] Bojarski et al., End to end learning for self-driving cars, arxiv 2016