EVALUATING ROBUSTNESS OF DEEP REINFORCEMENT LEARNING FOR AUTONOMOUS DRIVING

By Ege Bayram E.Bayram-1@student.tudelft.nl
Supervisors: Matthijs Spaan, Moritz Zanger
TU Delft

1. Research Question

How do domain randomizations influence training and the robustness of final policies under various testing conditions?

2. Background Information

Deep Q-Network (DQN)

- Q-learning
- ϵ-greedy policy
- Experience replay

Domain Randomization (DR)

- Train in a variety of environments to generalize to the target domain
- Visual and dynamic DR

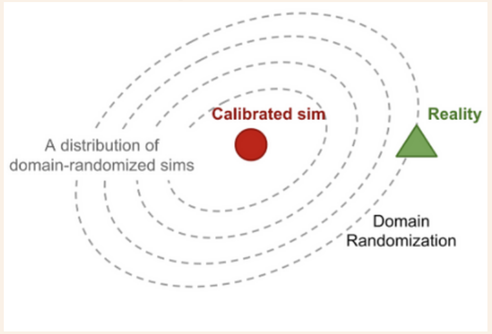

Figure 1: Conceptual illustration of domain randomization [5]

Figure 2: Examples of visual DR

4.1 Results for Visual DR

Training

Overall DR is underperforming

High variance during training

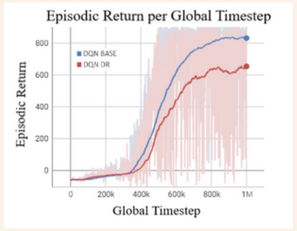

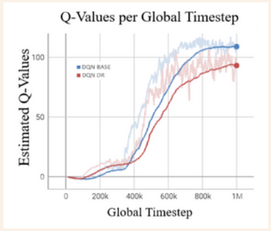

Testing

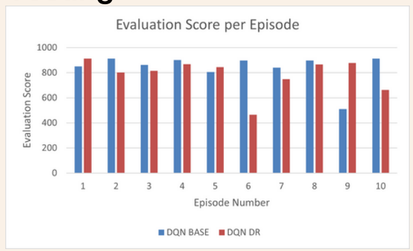

DR is underperforming during testing as well

Less variance difference compared to training data

	BASE	DR
MEAN	838.8	786.3
STDEV	120.6	133.9

4.2 Results for Dynamic DR

Training

Could not converge in 500k steps, low performance and high variance

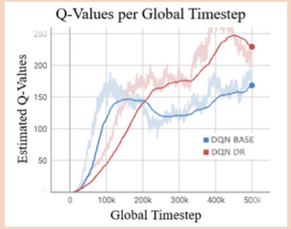

Overestimation of DR

Testing

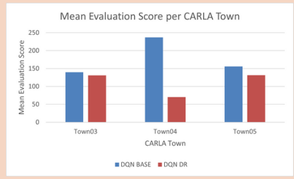

Lower mean in all cases, special case in Town04

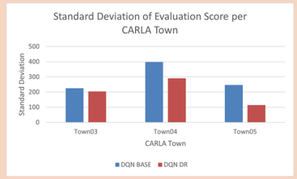

Lower standard deviation, higher stability

6. Conclusions

Visual DR in CarRacing-v2 shows a decrease in training performance and does not show a notable difference in robustness

Dynamic DR in CARLA shows worse training performance but an increase in robustness at the expense of rewards.

3. Methodology

DQN as base algorithm

Visual DR

- 1 million steps on OpenAI's CarRacing-v2
- Randomize background and track color at every episode
- 10 randomly generated tracks for robustness testing

Dynamic DR

- 500k steps on CARLA
- Randomize steering angle, acceleration value, and out-of-lane threshold every episode
- 10 episodes on 3 towns with random starting positions for robustness testing

5. Limitations

Hardware limitations delayed the setup for CARLA

Time limitations caused CARLA DR training to end before converging

Related Literature

[1] Bhairav Mehta, Manfred Diaz, Florian Golemo, Christopher J. Pal, and Liam Paull. Active domain randomization, 2019.
[2] Xue Bin Peng et al. “Sim-to-Real Transfer of Robotic Control with Dynamics Randomization”. In: 2018 IEEE International Conference on Robotics and Automation (ICRA). IEEE, May 2018. doi: 10.1109/icra.2018.8460528. url: https://doi.org/10.1109%2Ficra.2018.8460528.
[3] Fereshteh Sadeghi and Sergey Levine. Cad2rl: Real single-image flight without a single real image, 2017
[4] Josh Tobin et al. Domain Randomization for Transferring Deep Neural Networks from Simulation to the Real World. 2017. arXiv: 1703.06907 [cs.RO].
[5] Lilian Weng. “Domain Randomization for Sim2Real Transfer”. In: lilianweng.github.io (2019). url: https://lilianweng.github.io/posts/2019-05-05-domain-randomization/.