Efficient Focus for Autonomous Agents

15 MAR 2024

Bram Grooten, PhD Candidate

Supervisors: Decebal Mocanu (University of Luxembourg), Mykola Pechenizkiy (TU/e)
Email: b.j.grooten@tue.nl   Twitter/X: @BramGrooten

NWO
TU/e EINDHOVEN UNIVERSITY OF TECHNOLOGY

Bram Grooten

PhD candidate

Started: Nov 2021

Internship at Sony AI
Research visit in Alberta

Master Applied Mathematics
Master Science Education

Bachelor Applied Mathematics (Boston, US)

Finding Focus in RL

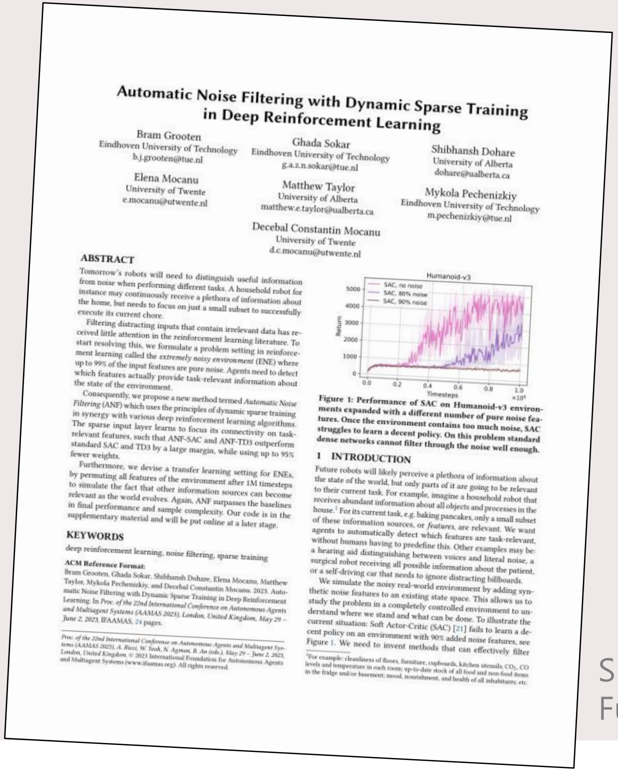

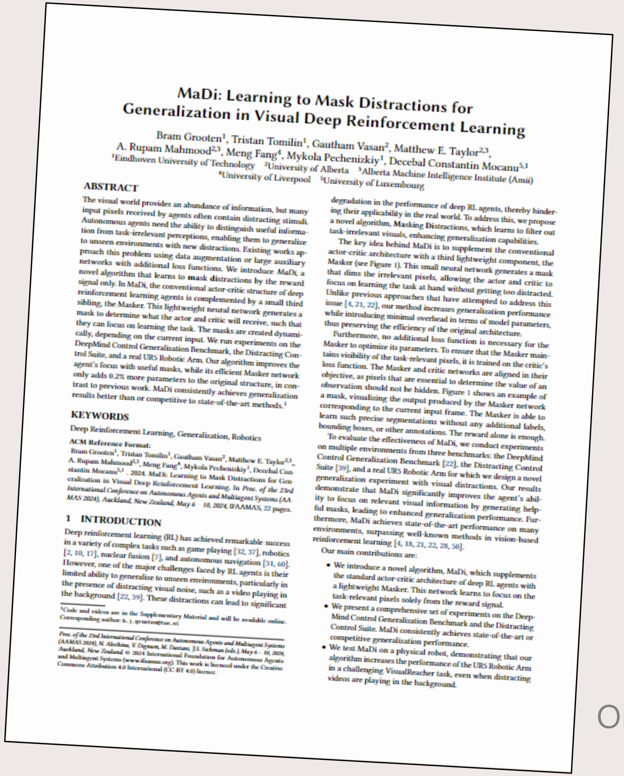

Oral at AAMAS'24

Spotlight at SNN workshop of ICLR'23
Full-paper at AAMAS'23

Why this research?

https://www.istockphoto.com/nl/foto/robotmeisje-dat-een-dienblad-houdt-en-voedsel-en-drank-in-moderne-binnenlandse-keuken-gm1309871457-399451605

https://www.pexels.com/photo/people-walking-on-times-square-during-night-time-12729170/

 Bram Grooten b.j.grooten@tue.nl

Types of noise

1. Uncertainty in perception

   measurement errors

2. Task-irrelevant percepts

   distracting information

https://eu.cjonline.com/story/news/local/2010/09/25/sun-poses-driving-challenges/16488302007/

https://www.pexels.com/photo/people-walking-on-times-square-during-night-time-12729170/

Bram Grooten b.j.grooten@tue.nl

Contents

Automatic Noise Filtering

Current situation:

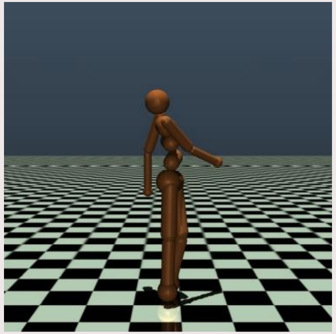
Humanoid-v3

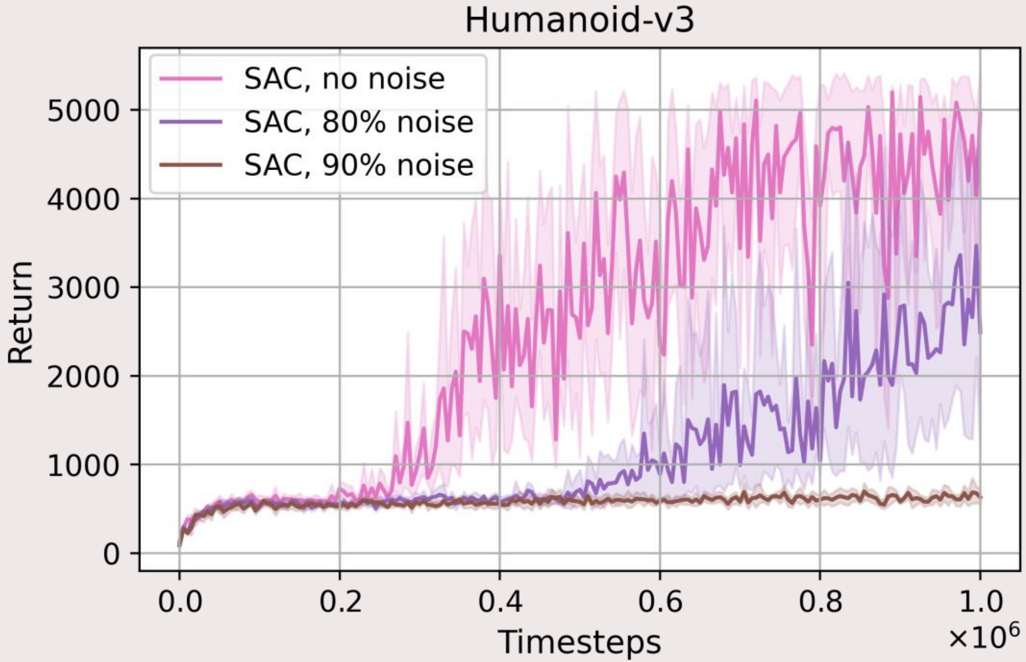

 Bram Grooten b.j.grooten@tue.nl

TU/e

How does ANF work?

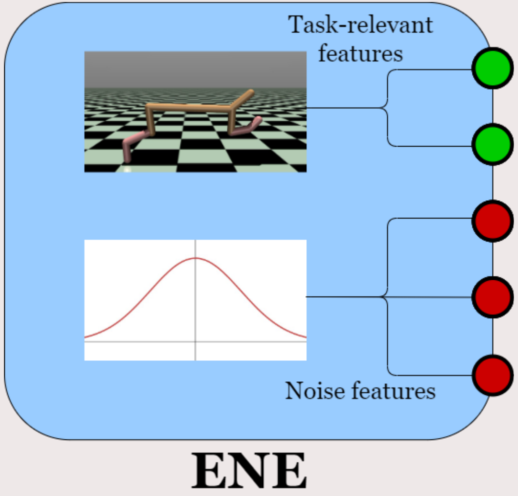

Extremely noisy environment:

simulate irrelevant information by adding many *noise features*

How does ANF work?

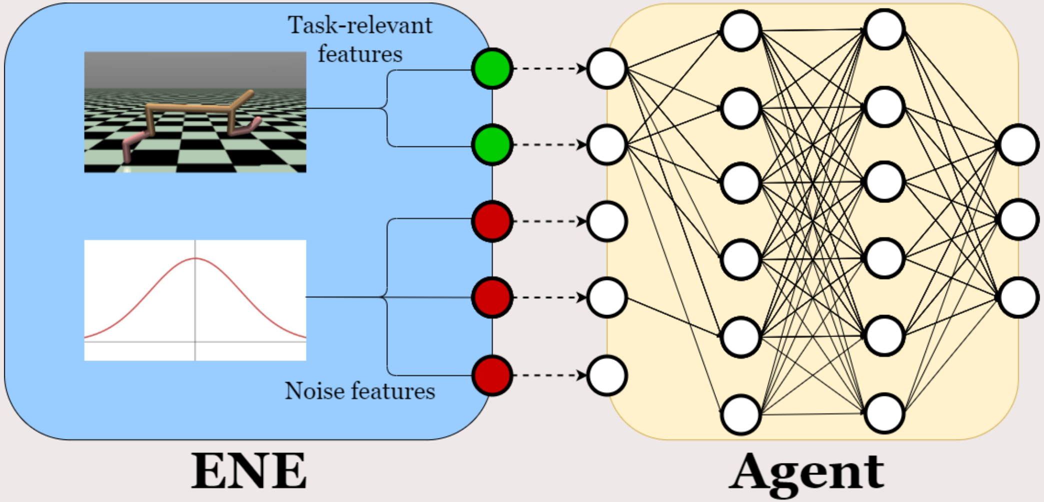

Automatic Noise Filtering:

adapt the input layer's connectivity to focus on relevant features

How does ANF work?

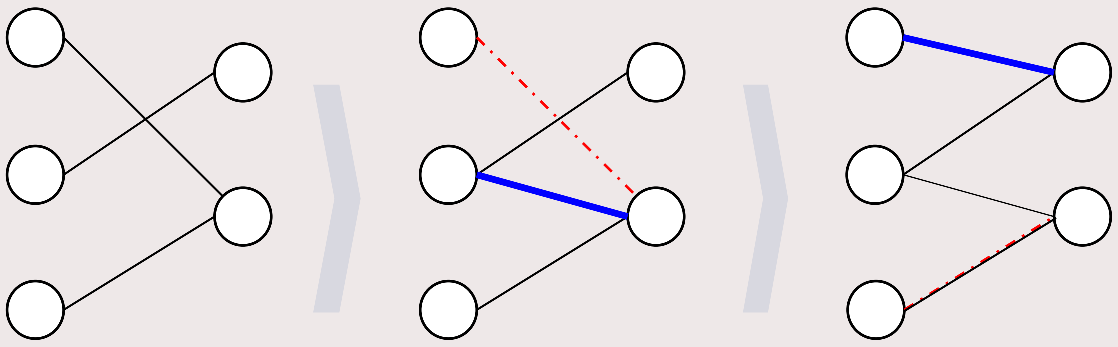

Dynamic Sparse Training (Drop - · - · - , Grow ——)

Sokar, G., et al. *Dynamic Sparse Training for Deep Reinforcement Learning*. IJCAI (2022).
Mocanu, D.C., et al. *Scalable training of artificial NNs with adaptive sparse connectivity inspired by network science*. Nature Comms. (2018).

Experiments

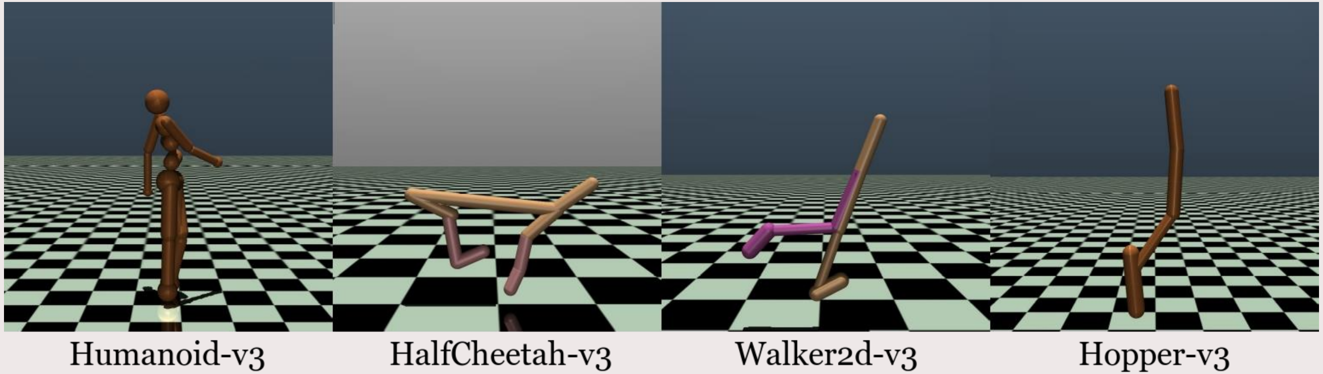

Table 1: State and action space dimensions.

Environment	State dim. *Original*	Action dim. *Original*	State dim. *ENE* ($n_f = .8$)	State dim. *ENE* ($n_f = .99$)
Humanoid-v3	376	17	1880	37600
HalfCheetah-v3	17	6	85	1700
Walker2d-v3	17	6	85	1700
Hopper-v3	11	3	55	1100

Bram Grooten b.j.grooten@tue.nl

Results

90% noise features:

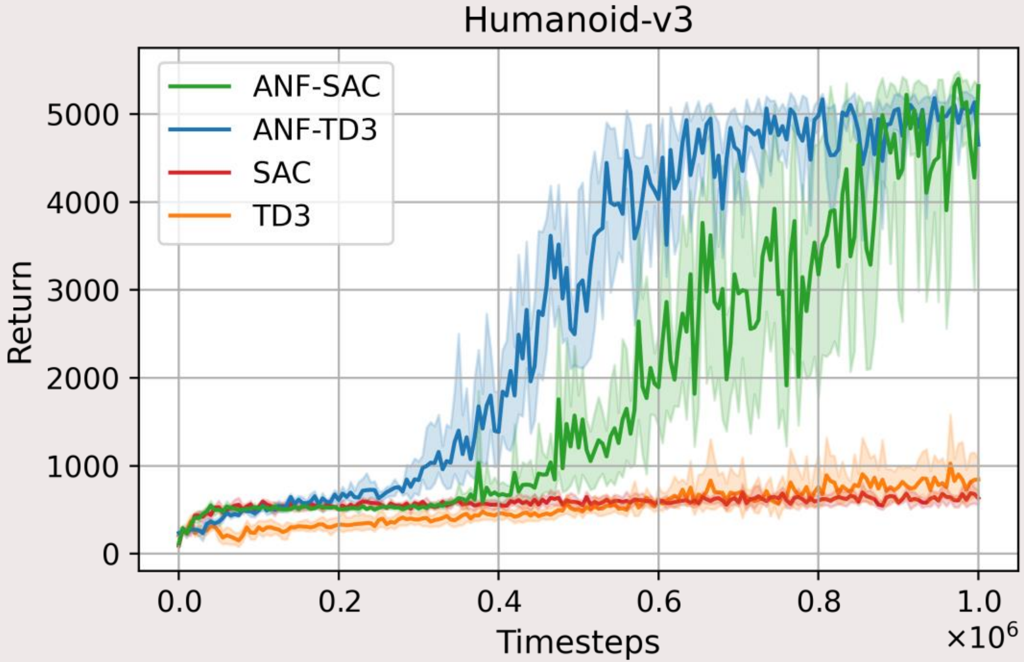

Bram Grooten b.j.grooten@tue.nl

Results

98% noise features:

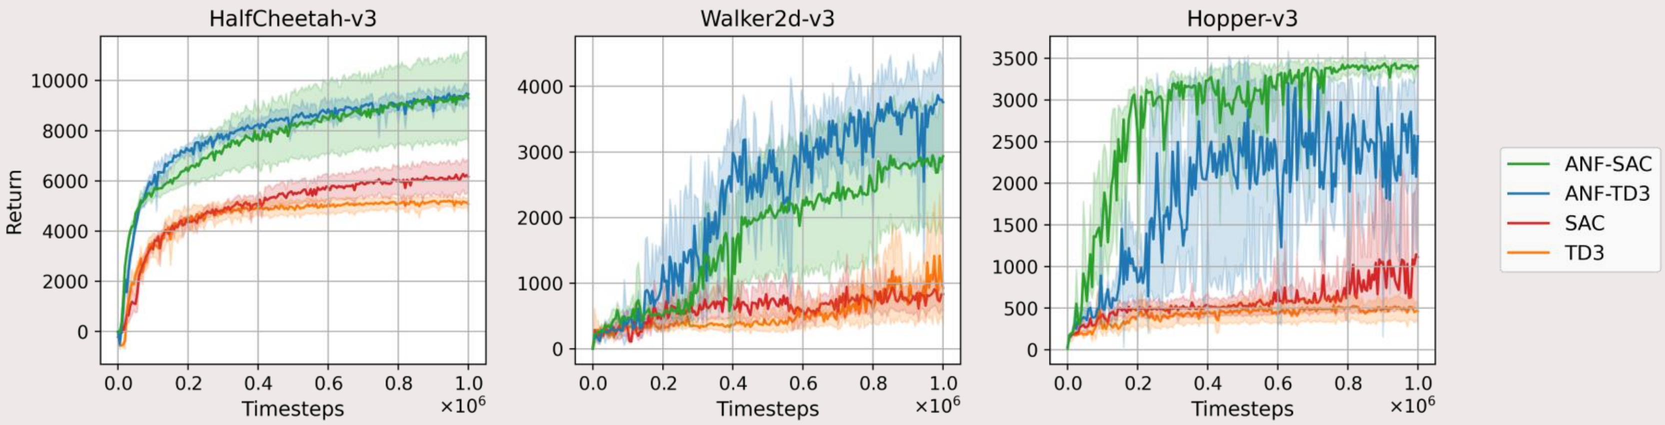

 Bram Grooten b.j.grooten@tue.nl

Results

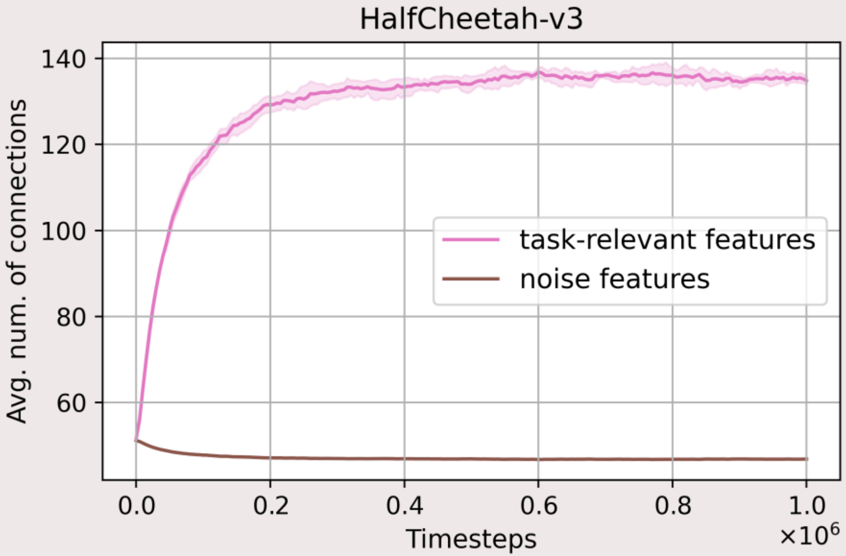

ANF adapts the connectivity to focus on task-relevant features

Louder noise

Previously:

Noise sampled from $N(0, 1)$

Now:

Sample from: $N(0, \sigma^2)$

Varying the *noise amplitude* σ

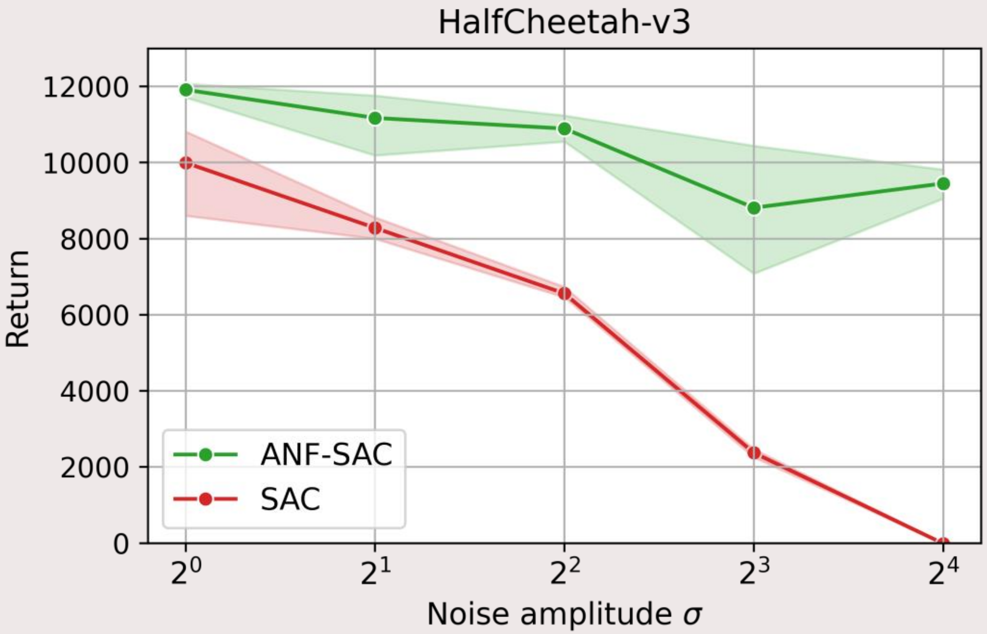

Louder noise

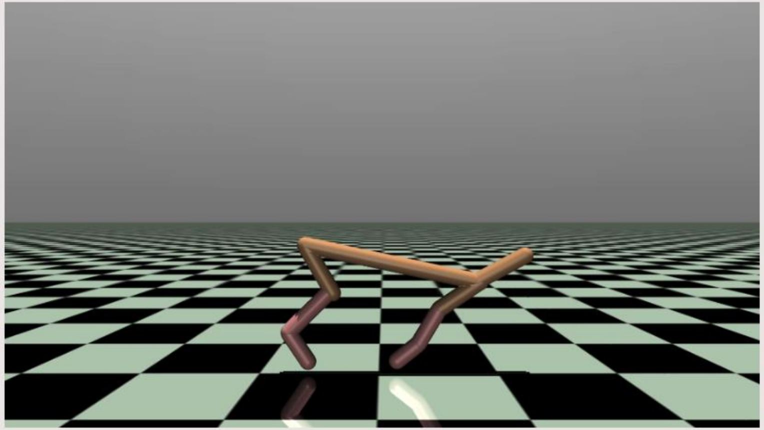

SAC

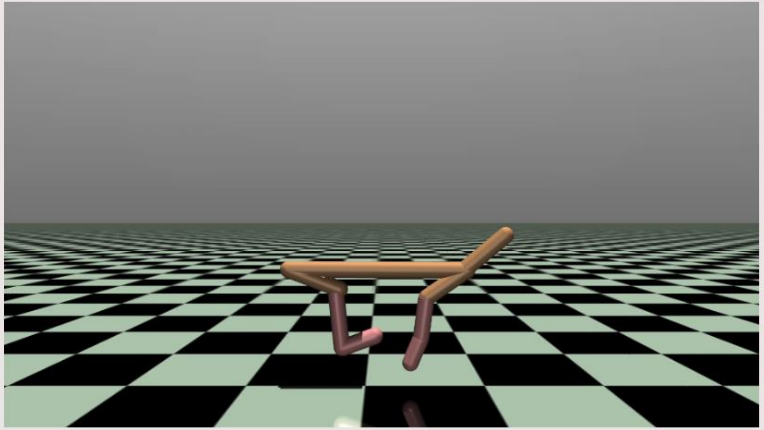

ANF-SAC

 Bram Grooten b.j.grooten@tue.nl

Louder noise

Source: twitter.com/wonderofscience/status/1542146051658104832

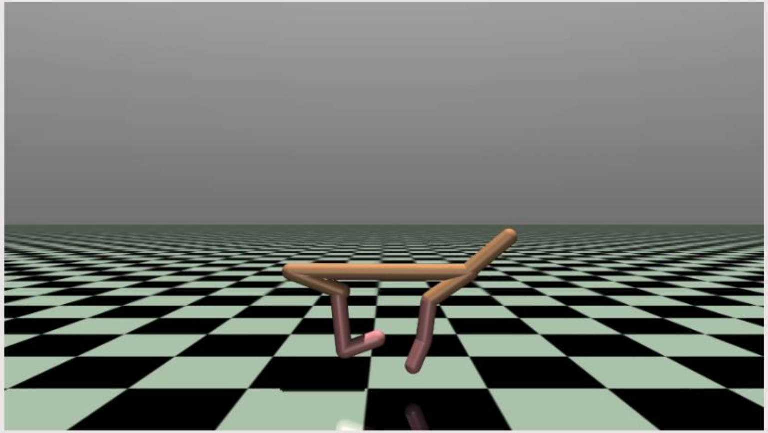

ANF-SAC

Transfer learning

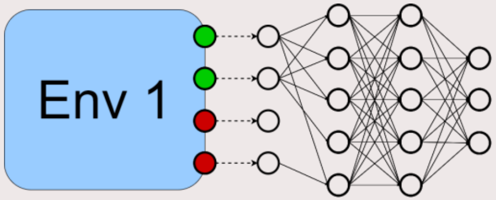

Transfer learning

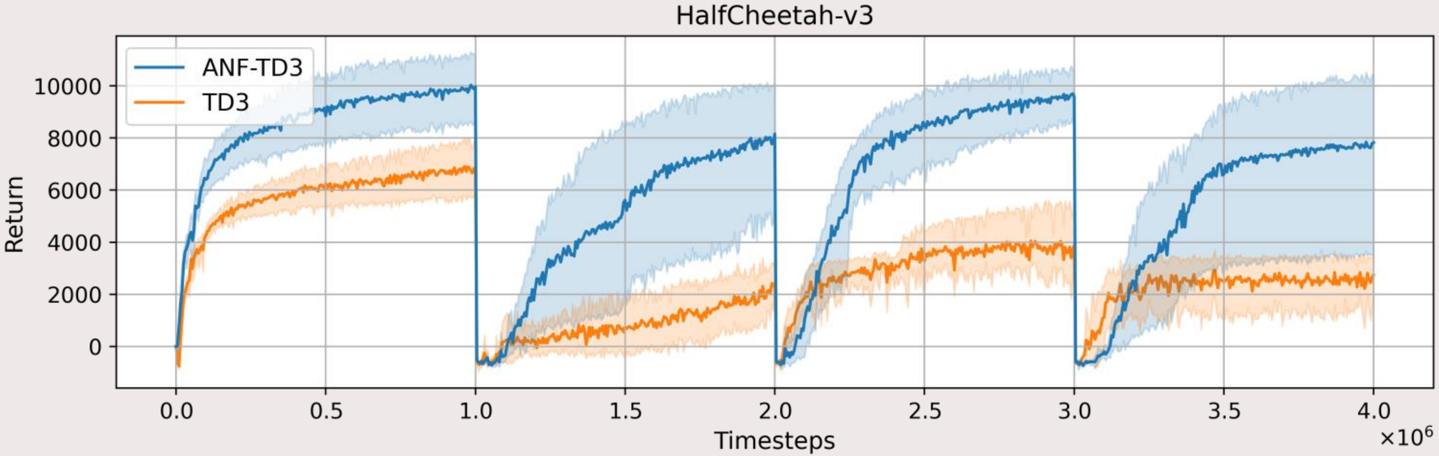

Every 1M timesteps the environment changes

Imitating real features

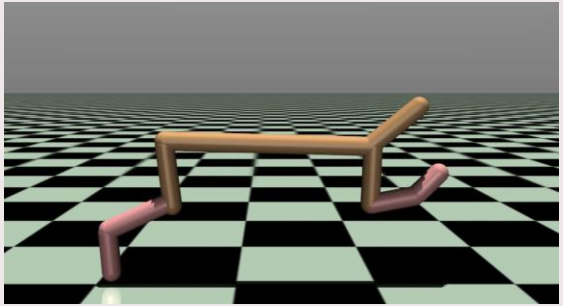

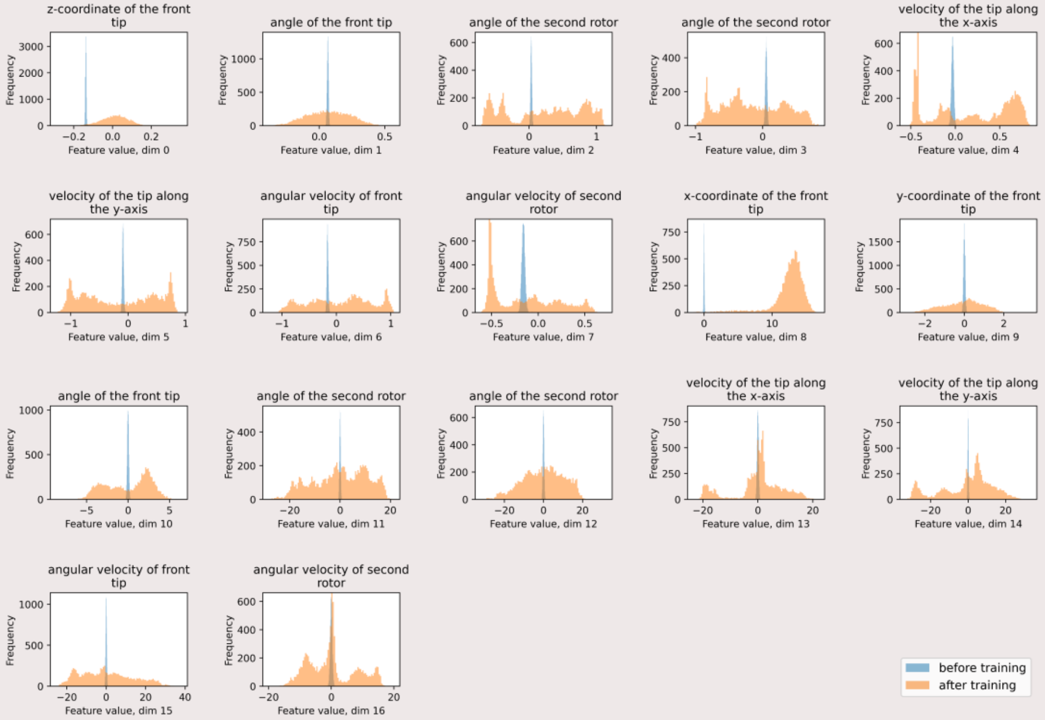

Imitating real features

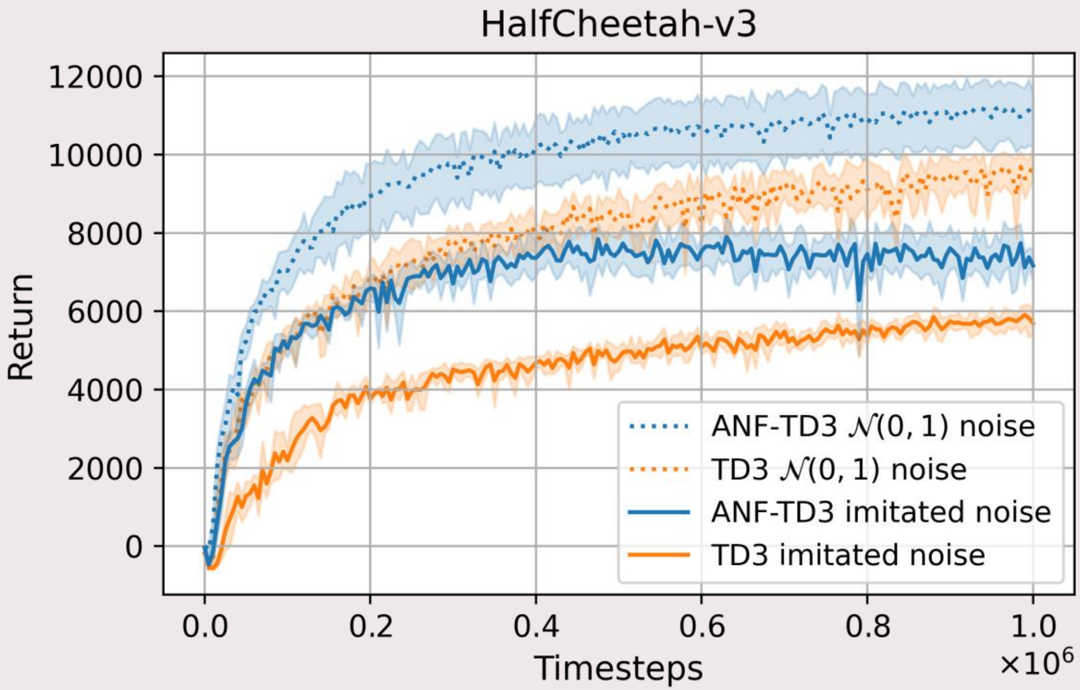

Motivation & Introduction

Automatic Noise Filtering (ANF)

Masking Distractions (MaDi)

Potential future work

Bram Grooten b.j.grooten@tue.nl

Masking Distractions (MaDi)

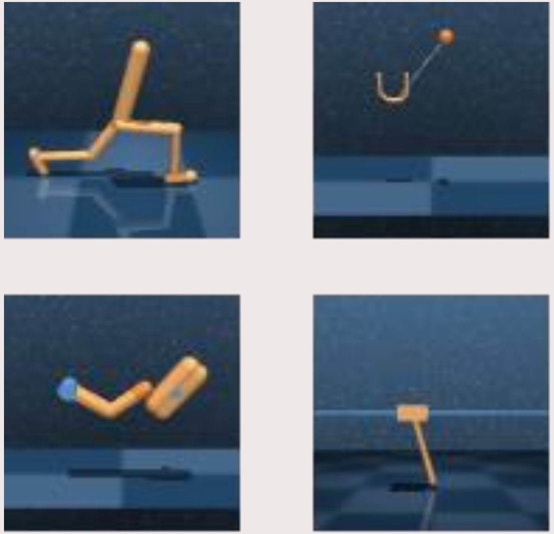

Training environments

Testing

Stone et al. *The Distracting Control Suite – A Challenging Benchmark for RL from Pixels.* (2021)

Bram Grooten b.j.grooten@tue.nl

How does MaDi work?

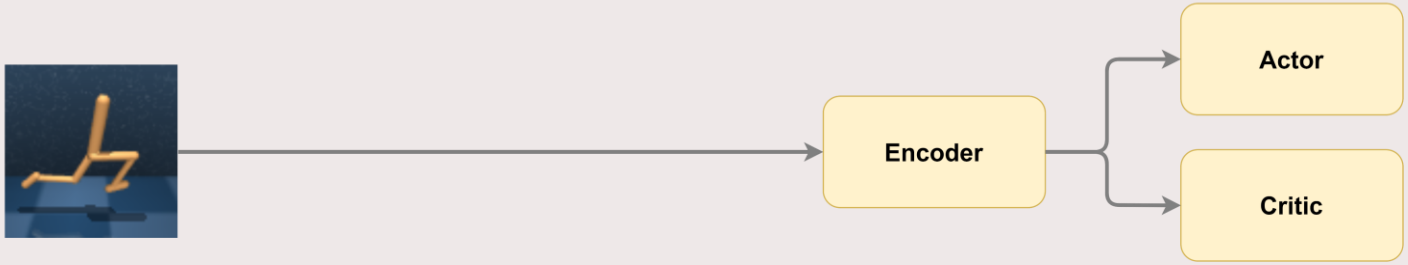

Hansen et al. *Stabilizing Deep Q-Learning with ConvNets and Vision Transformers under Data Augmentation.* (2021)

How does MaDi work?

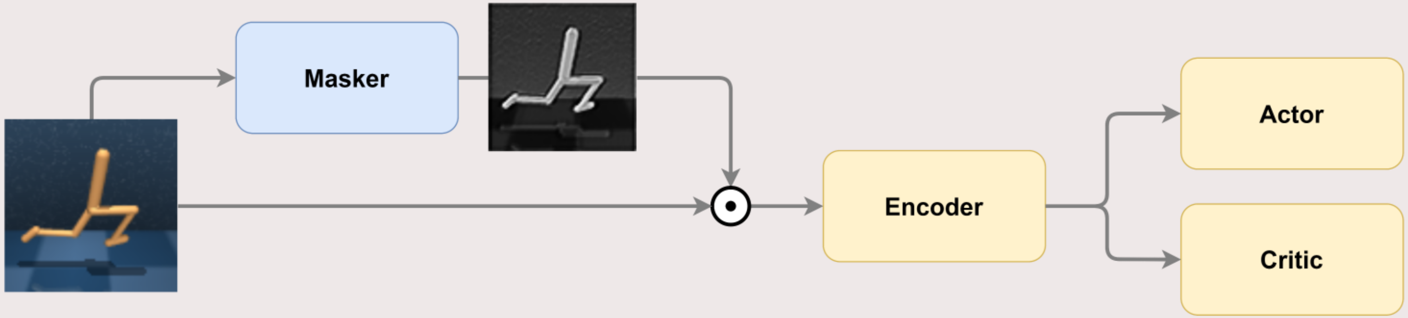

How does MaDi work?

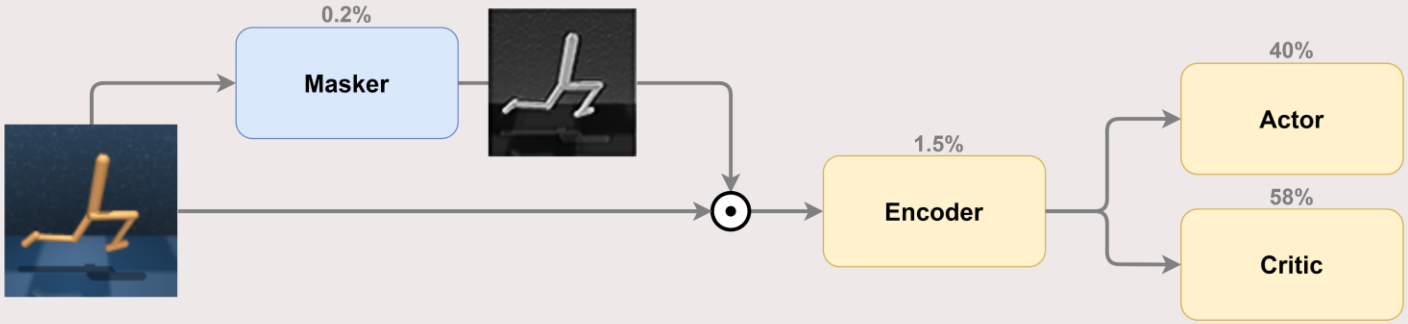

TU/e

How does MaDi work?

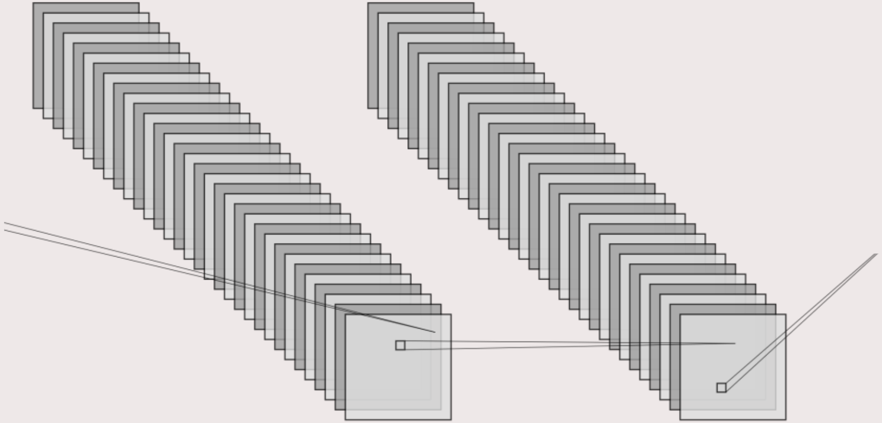

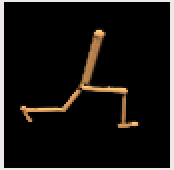

How does MaDi work?

Initialization:

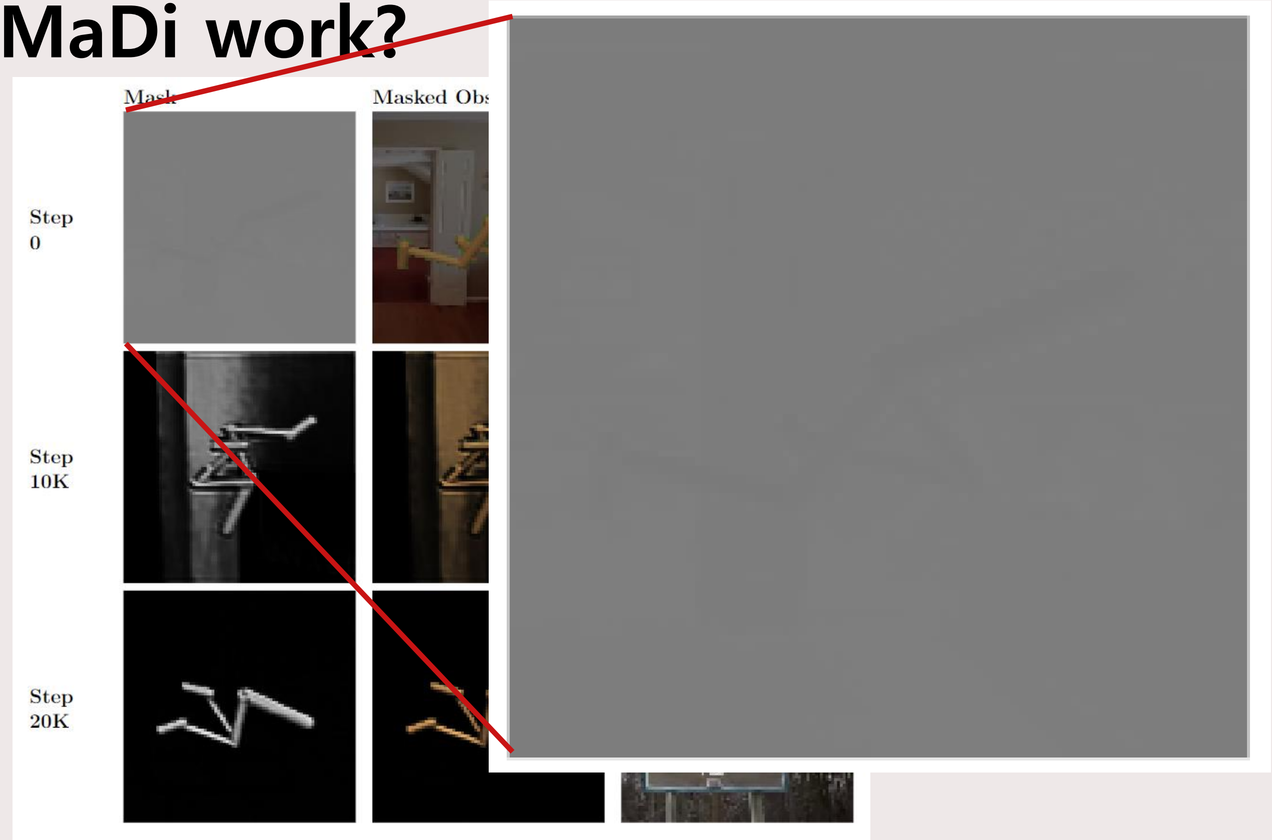

Results

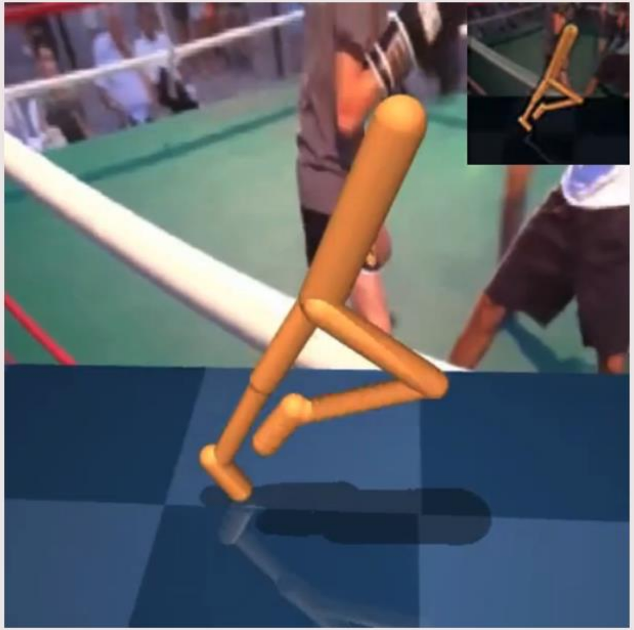

 Bram Grooten b.j.grooten@tue.nl

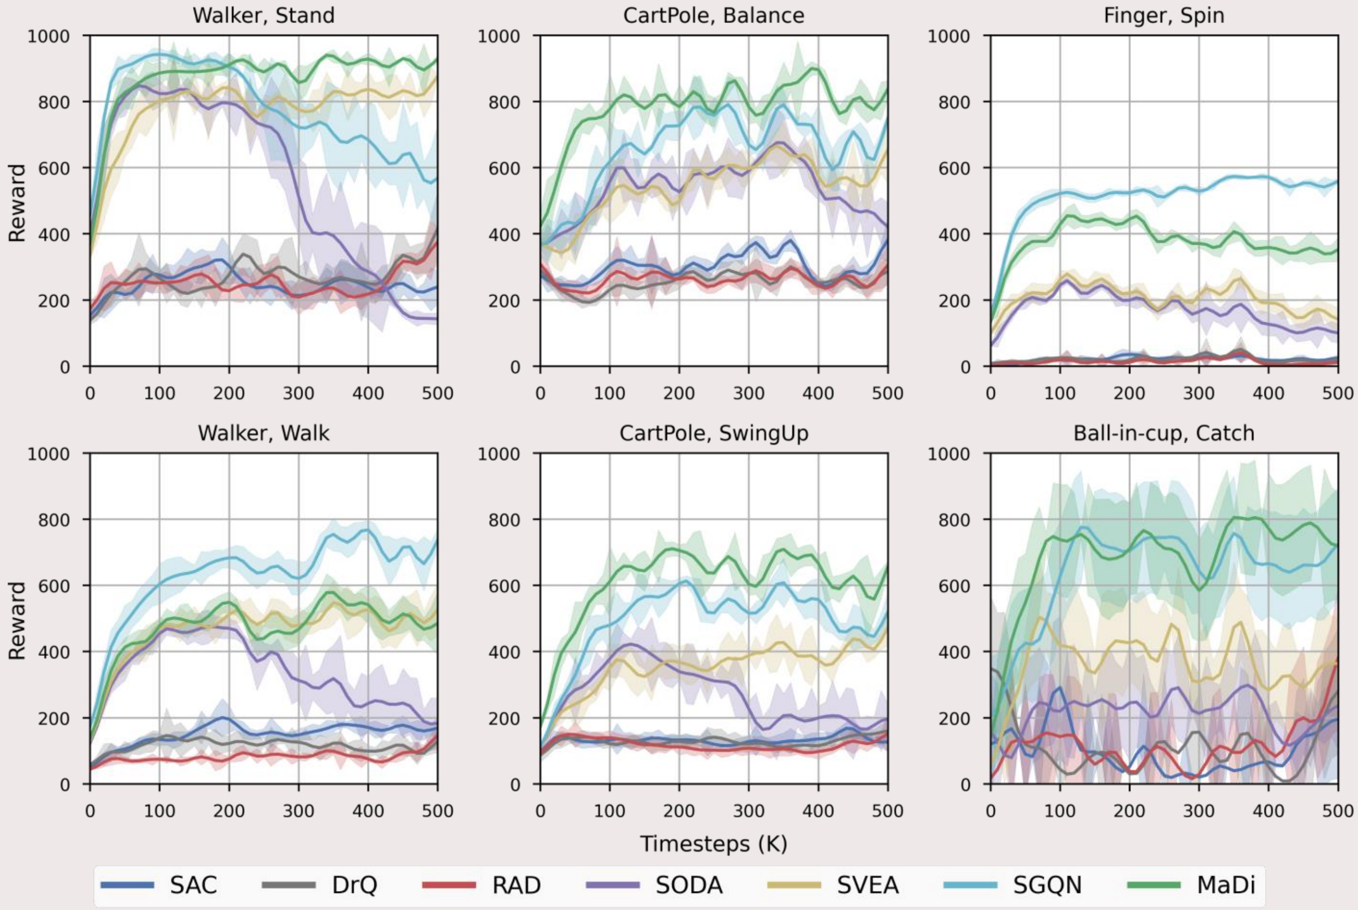
Walker, Stand
CartPole, Balance
Finger, Spin
Walker, Walk
CartPole, SwingUp
Ball-in-cup, Catch
Reward
Timesteps (K)
SAC
DrQ
RAD
SODA
SVEA
SGQN
MaDi

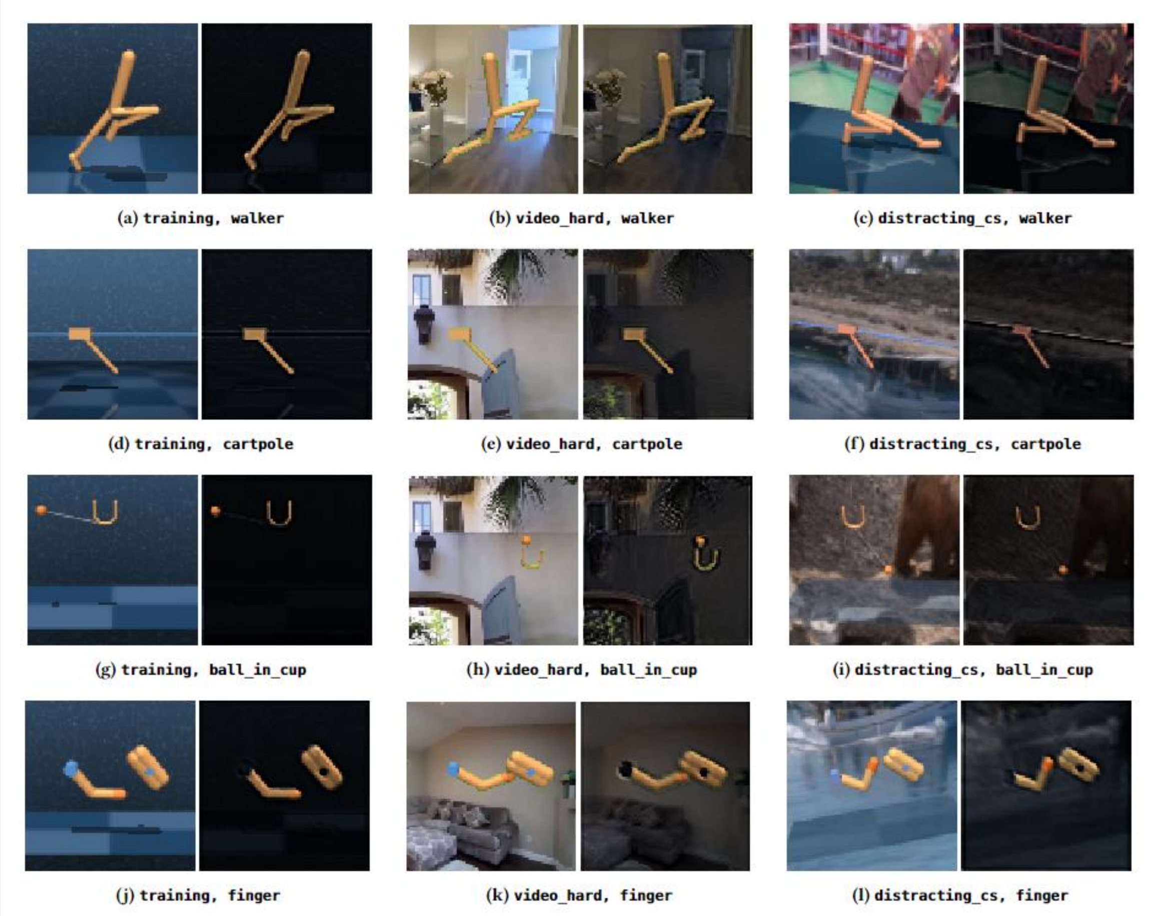
(a) training, walker
(b) video_hard, walker
(c) distracting_cs, walker
(d) training, cartpole
(e) video_hard, cartpole
(f) distracting_cs, cartpole
(g) training, ball_in_cup
(h) video_hard, ball_in_cup
(i) distracting_cs, ball_in_cup
(j) training, finger
(k) video_hard, finger
(l) distracting_cs, finger

Robotic experiments

UR5 Robotic Arm

VisualReacher task

Training asynchronously

Inputs:

- webcam view
- proprioception

Wang et al. Real-Time Reinforcement Learning for Vision-Based Robotics Utilizing Local and Remote Computers. (2023)

TU/e

Robotic experiments

Training

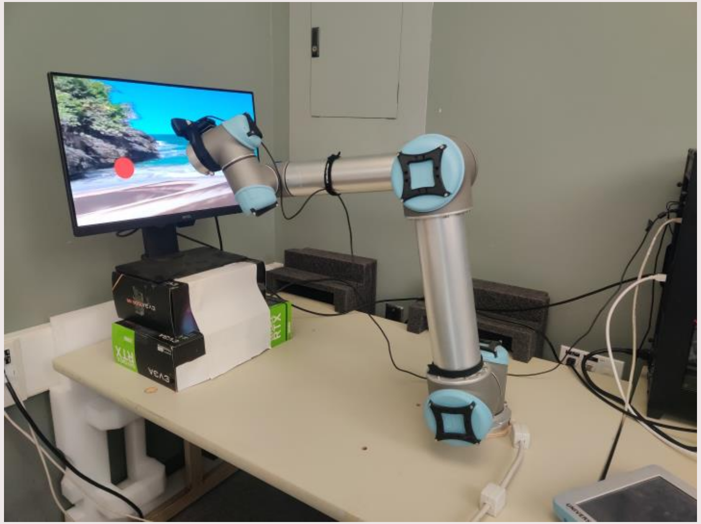

Testing

 Bram Grooten b.j.grooten@tue.nl

TU/e

Robotic experiments

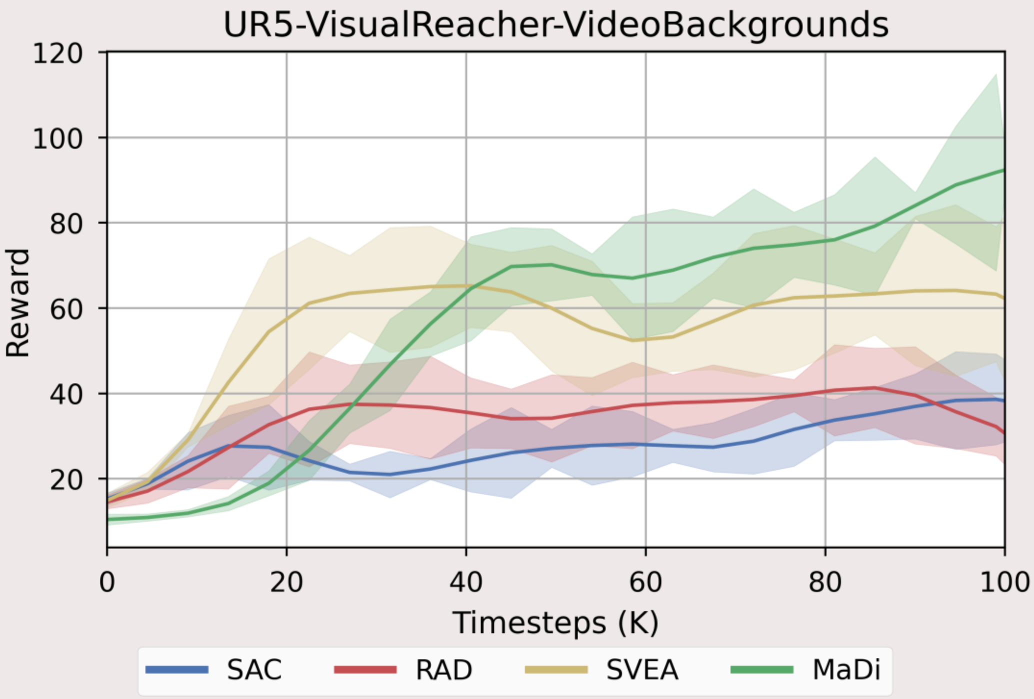

UR5-VisualReacher-VideoBackgrounds

MaDi

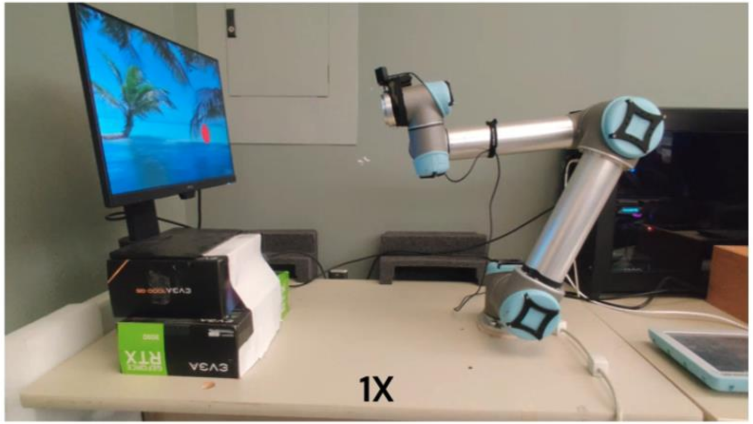

SVEA

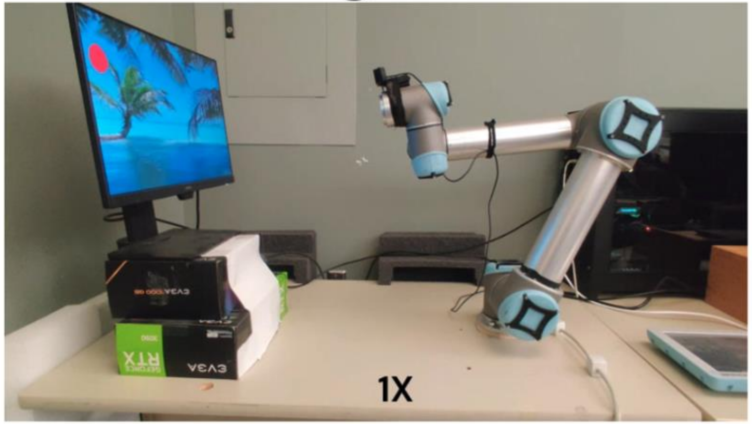

RAD

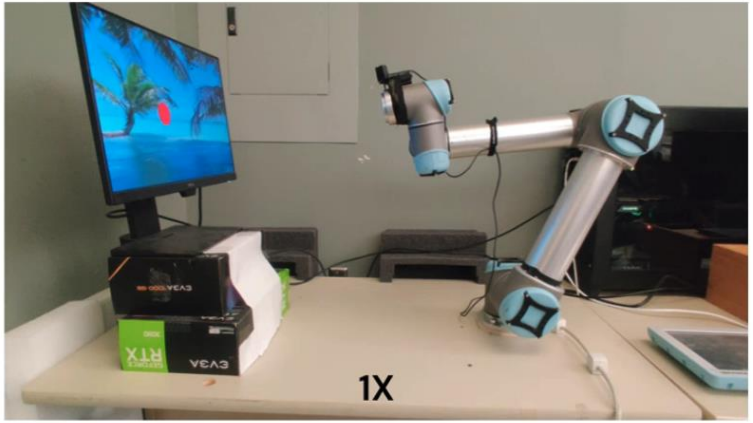

SAC

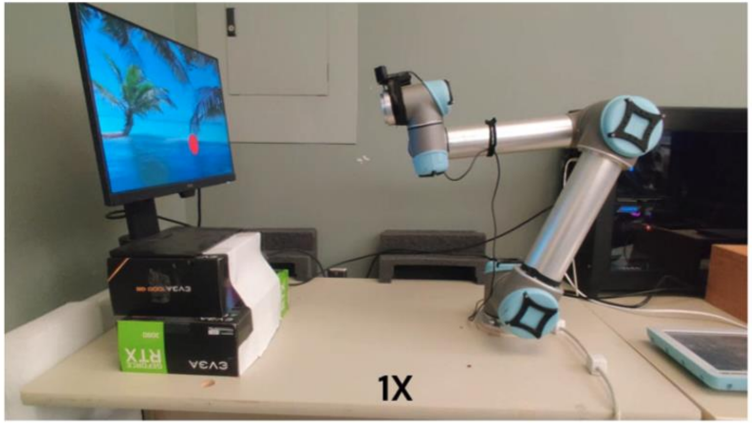

Masks

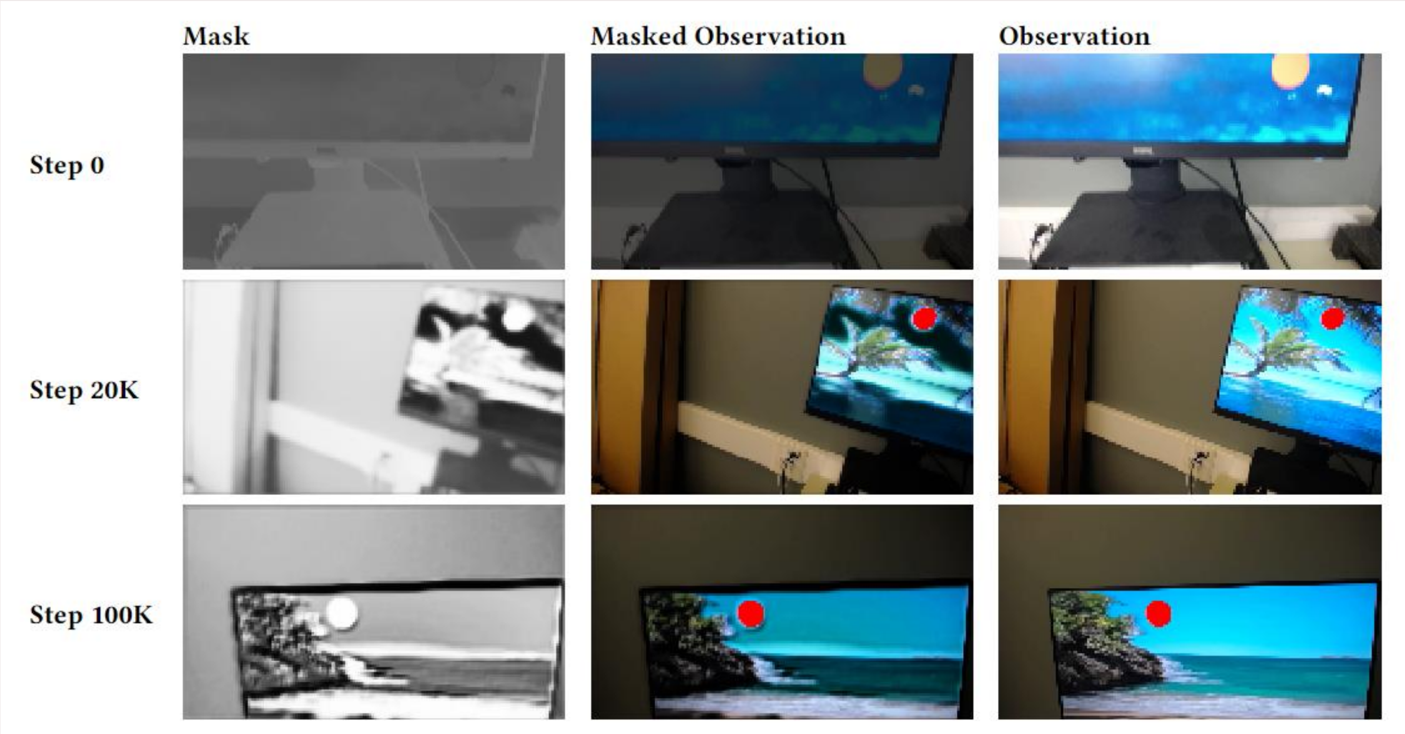

Mask
Masked Observation
Observation
Step 0
Step 20K
Step 100K

Motivation & Introduction

Automatic Noise Filtering (ANF)

Masking Distractions (MaDi)

Potential future work

Bram Grooten b.j.grooten@tue.nl

Potential future work

- Other benchmarks or robotics applications
- Sparsify later in the network
- Transfer Learning with MaDi
- Supervised Learning

Bram Grooten b.j.grooten@tue.nl

MaDi on other benchmarks

CARLA (Car Learning to Act)

Dosovitskiy et al. *CARLA: An Open Urban Driving Simulator.* CoRL (2017).

 Bram Grooten b.j.grooten@tue.nl

Sparsify later in the network

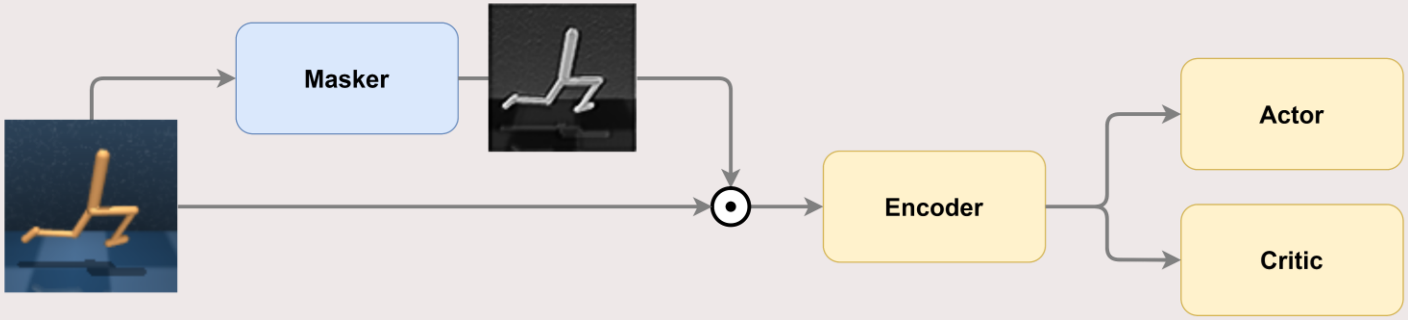

Sparsify later in the network

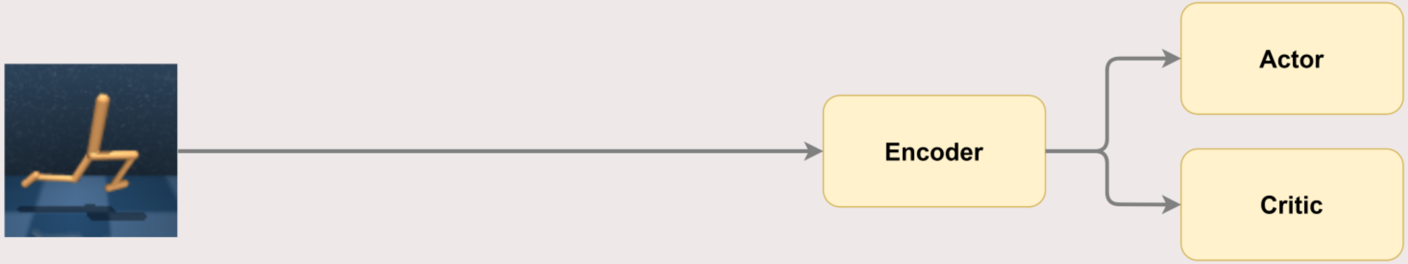

Sparsify later in the network

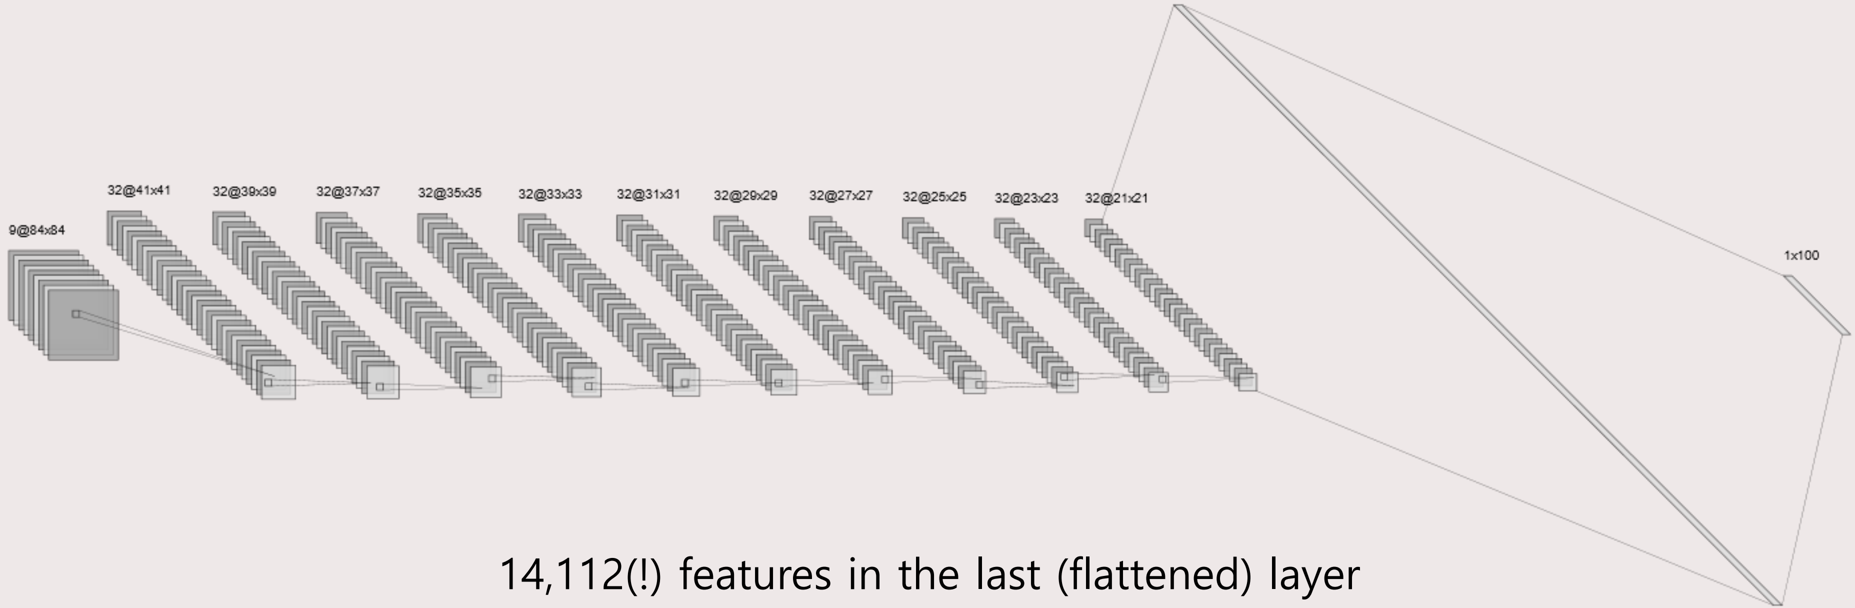

14,112(!) features in the last (flattened) layer

Made with: alexlenail.me/NN-SVG

Sparsify later in the network

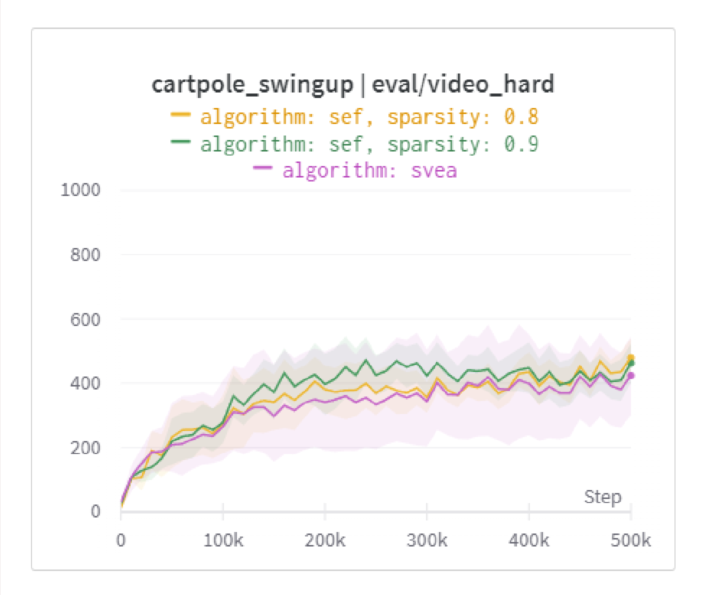

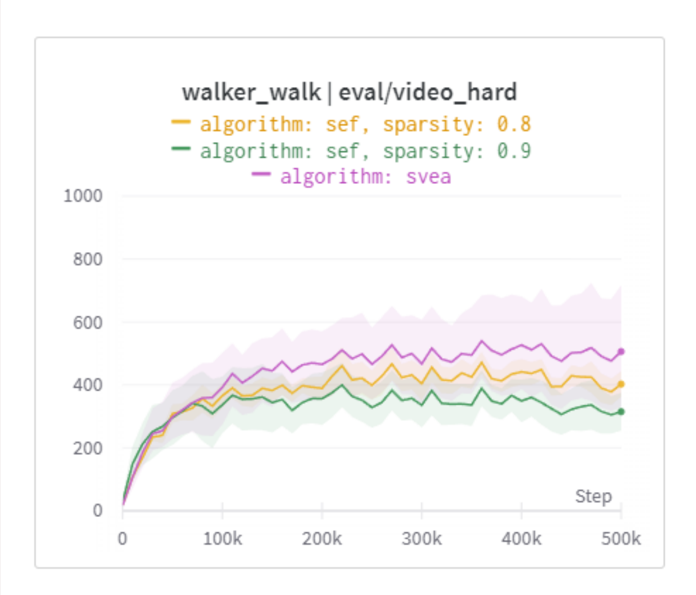

Transfer Learning with MaDi

walker_stand ➡ walker_walk

cartpole_balance ➡ cartpole_swingup

walker_stand ➡ cartpole_swingup (?)

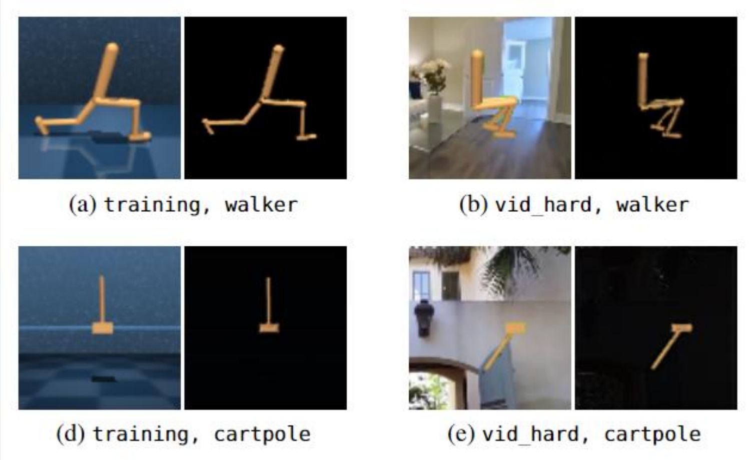

Bram Grooten b.j.grooten@tue.nl

Main take-away

ANF and MaDi outperform deep RL methods by learning to focus on task-relevant features.

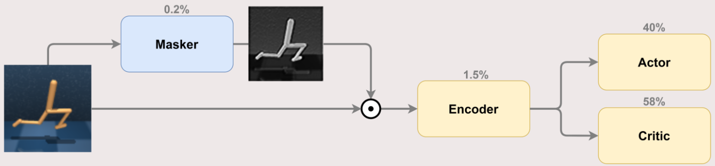

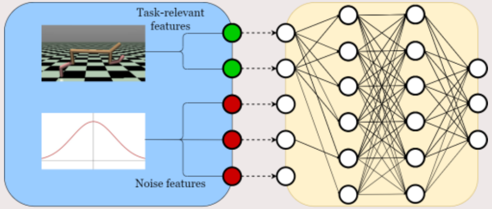

arxiv.org/abs/2302.06548

arxiv.org/abs/2312.15339

github.com/bramgrooten/automatic-noise-filtering

github.com/bramgrooten/mask-distractions

@BramGrooten

Bram Grooten b.j.grooten@tue.nl
TU/e

Tutorial

github.com/bramgrooten/automatic-noise-filtering/blob/main/tutorial.md

Sparse Training in Deep RL - Tutorial

This guide is designed for anyone interested in using sparse neural networks in deep reinforcement learning. It was part of the IJCAI 2023 tutorial T27: *Sparse Training for Supervised, Unsupervised, Continual, and Deep Reinforcement Learning with Deep Neural Networks*. In the following we will play around with some of the settings in this repository. I hope this will give you a feeling for the types of problems we are researching in this field.

Author: Bram Grooten.

Install

First `git clone` and install this repository, by following the instructions in the README.

Visualizing a trained agent

We have included a trained agent in the repository already, so let's view what this actually looks like! Go to this repository's main folder in your terminal (the directory with the `view_mujoco.py` file in it), make sure your venv is activated, and run:

```
python view_mujoco.py
```

A window should pop up showing you one episode of a running HalfCheetah! The camera doesn't track the agent by default, press `TAB` to do that.

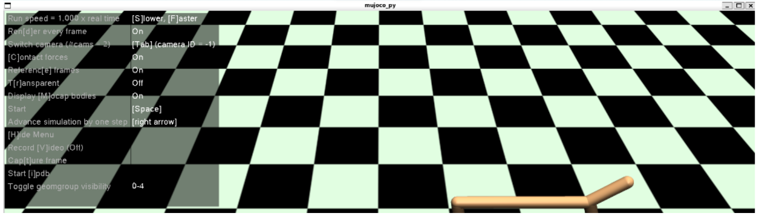

 Bram Grooten b.j.grooten@tue.nl

References

Dosovitskiy, A., Ros, G., Codevilla, F., Lopez, A., and Koltun, V. *CARLA: An Open Urban Driving Simulator.* Conference on Robot Learning (2017).

Grooten, B., Sokar, G., Dohare, S., Mocanu, E., Taylor, M., Pechenizkiy, M., and Mocanu, D.C. *Automatic Noise Filtering with Dynamic Sparse Training in Deep Reinforcement Learning*. AAMAS (2023).

Hansen, N., Su, H., Wang, X. *Stabilizing Deep Q-Learning with ConvNets and Vision Transformers under Data Augmentation*. NeurIPS (2021).

Mocanu, D.C., Mocanu, E., Stone, P., Nguyen, P., Gibescu, M., and Liotta, A. *Scalable training of artificial neural networks with adaptive sparse connectivity inspired by network science*. Nature communications (2018).

Sokar, G., Mocanu, E., Mocanu, D.C., Pechenizkiy, M., and Stone, P. *Dynamic Sparse Training for Deep Reinforcement Learning.* IJCAI (2022).

Stone, A., Ramirez, O., Konolige, K., and Jonschkowski, R. *The Distracting Control Suite – A Challenging Benchmark for Reinforcement Learning from Pixels*. arXiv preprint (2021).

Wang, Y., Vasan, G., and Mahmood, A.R. *Real-Time Reinforcement Learning for Vision-Based Robotics Utilizing Local and Remote Computers.* ICRA (2023).