

# Learning control of underactuated double pendulum with Model-Based Reinforcement Learning

Dalla Libera, Alberto; Turcato, Niccolò; Giacomuzzo, Giulio; Carli, Ruggero; Romeres, Diego

TR2024-142    October 17, 2024

## Abstract

This report describes our proposed solution for the second AI Olympics competition held at IROS 2024. Our solution is based on a recent Model-Based Reinforcement Learning algorithm named MC-PILCO. Besides briefly reviewing the algorithm, we discuss the most critical aspects of the MC- PILCO implementation in the tasks at hand.

*Competition: AI Olympics With RealAIGym 2024*

© 2024 MERL. This work may not be copied or reproduced in whole or in part for any commercial purpose. Permission to copy in whole or in part without payment of fee is granted for nonprofit educational and research purposes provided that all such whole or partial copies include the following: a notice that such copying is by permission of Mitsubishi Electric Research Laboratories, Inc.; an acknowledgment of the authors and individual contributions to the work; and all applicable portions of the copyright notice. Copying, reproduction, or republishing for any other purpose shall require a license with payment of fee to Mitsubishi Electric Research Laboratories, Inc. All rights reserved.



# Learning control of underactuated double pendulum with Model-Based Reinforcement Learning

Niccolò Turcato<sup>1</sup>, Alberto Dalla Libera<sup>1</sup>, Giulio Giacomuzzo<sup>1</sup>, Ruggero Carli<sup>1</sup> and Diego Romeres<sup>2</sup>

**Abstract**— This report describes our proposed solution for the second AI Olympics competition held at IROS 2024. Our solution is based on a recent Model-Based Reinforcement Learning algorithm named MC-PILCO. Besides briefly reviewing the algorithm, we discuss the most critical aspects of the MC-PILCO implementation in the tasks at hand.

## I. INTRODUCTION

In this short report, we present the Reinforcement Learning (RL) [1] approach our team implemented to tackle the simulation stage of the second AI Olympics competition held at IROS 2024<sup>1</sup>. The algorithm we employed, Monte-Carlo Probabilistic Inference for Learning CONTROL (MC-PILCO) [2], is a Model-Based (MB) RL algorithm that proved remarkably data-efficient in several low-dimensional benchmarks, such as a cart-pole, a ball & plate, and a Furuta pendulum, both in simulation and real setups. MC-PILCO is also the algorithm that won the first edition of this competition [3]. MC-PILCO is part of the class of MB policy gradient algorithms. It exploits data collected by interacting with the system to derive a system dynamics model and optimizes the policy by simulating the system, rather than optimizing the policy directly on the system’s data. When applied to physical systems, this approach can be highly performing and more data-efficient than Model-Free (MF) solutions.

This paper is organized as follows: Section II introduces the goal and the settings of the competition. Section III presents the MC-PILCO algorithm. Section IV reports the experiments that have been performed, finally Section V concludes the paper.

## II. GOAL OF THE COMPETITION

The challenge considers a 2 degrees of freedom (dof) underactuated pendulum [4] with two possible configurations. In the first configuration, also called Pendubot, the first joint, namely, the one attached to the base link is active, and the second is passive. Instead, in the second configuration, also named Acrobot, the first joint is passive and the second is actuated. For each configuration, the competition’s goal is to derive a controller that performs the swing-up and stabilization in the unstable equilibrium point of the systems. Both robots are underactuated, which makes the task particularly challenging from the control point of

view. The systems are simulated at 500 Hz with a Runge-Kutta 4 integrator for a horizon of  $T = 10$  s. The competition is composed of 2 stages. In the first stage, namely the simulation stage, controllers are assessed based on performance and robustness scores in the simulated system. In the second stage, namely the real hardware stage, the participating teams test their controllers on the real system, with the possibility of retraining their learning-based controllers. The winners of the competition are chosen based on the performance and the reliability of the submitted controllers.

## III. MC-PILCO FOR UNDERACTUATED ROBOTICS

In this section, firstly we review MC-PILCO, secondly, we discuss its application to the considered problem.

### A. MC-PILCO review

MC-PILCO is a MB policy gradient algorithm, in which GPs are used to estimate system dynamics and long-term state distributions are approximated with a particle-based method.

Consider a system with evolution described by the discrete-time unknown transition function  $f : \mathbb{R}^{d_x} \times \mathbb{R}^{d_u} \rightarrow \mathbb{R}^{d_x}$ :

$$\mathbf{x}_{t+1} = f(\mathbf{x}_t, \mathbf{u}_t) + \mathbf{w}_t, \quad (1)$$

where  $\mathbf{x}_t \in \mathbb{R}^{d_x}$  and  $\mathbf{u}_t \in \mathbb{R}^{d_u}$  are respectively the state and input of the system at step  $t$ , while  $\mathbf{w}_t$  is an independent white noise describing uncertainty influencing the system evolution. As usual in RL, a cost function  $c(\mathbf{x}_t)$  encodes the task to be accomplished. A policy  $\pi_\theta : \mathbf{x} \rightarrow \mathbf{u}$  that depends on the parameters  $\theta$  selects the inputs applied to the system. The objective is to find policy parameters  $\theta^*$  that minimize the cumulative expected cost, defined as follows,

$$J(\theta) = \sum_{t=0}^T \mathbb{E}[c(\mathbf{x}_t)], \quad (2)$$

where the initial state  $x_0$  is sampled according to a given probability  $p(\mathbf{x}_0)$ .

MC-PILCO consists of a series of attempts, known as trials, to solve the desired task. Each trial consists of three main phases: (i) model learning, (ii) policy update, and (iii) policy execution. In the first trial, the GP model is derived from data collected with an exploration policy, for instance, a random exploration policy.

In the model learning step, previous experience is used to build or update a model of the system dynamics. The policy update step formulates an optimization problem whose objective is to minimize the cost in eq. (2) w.r.t. the parameters

<sup>1</sup>Department of Information Engineering, University of Padova, Italy.

<sup>2</sup>Mitsubishi Electric Research Laboratories, Cambridge, MA, USA  
Correspondence to niccolo.turcato@phd.unipd.it

<sup>1</sup><https://ai-olympics.dfki-bremen.de/>

of the policy  $\theta$ . Finally, in the last step, the current optimized policy is applied to the system and the collected samples are stored to update the model in the next trials.

In the rest of this section, we give a brief overview of the main components of the algorithm and highlight their most relevant features.

1) *Model Learning*: MC-PILCO relies on GP Regression (GPR) to learn the system dynamics [5]. In our previous work, [2], we presented a framework specifically designed for mechanical systems, named speed-integration model. Given a mechanical system with  $d$  degrees of freedom, the state is defined as  $\mathbf{x}_t = [\mathbf{q}_t^T, \dot{\mathbf{q}}_t^T]^T$  where  $\mathbf{q}_t \in \mathbb{R}^d$  and  $\dot{\mathbf{q}}_t \in \mathbb{R}^d$  are, respectively, the generalized positions and velocities of the system at time  $t$ . Let  $T_s$  be the sampling time and assume that accelerations between successive time steps are constant. The following equations describe the one-step-ahead evolution of the  $i$ -th degree of freedom,

$$\dot{q}_{t+1}^{(i)} = \dot{q}_t^{(i)} + \Delta_t^{(i)} \quad (3a)$$

$$q_{t+1}^{(i)} = q_t^{(i)} + T_s \dot{q}_t^{(i)} + \frac{T_s}{2} \Delta_t^{(i)} \quad (3b)$$

where  $\Delta_t^{(i)}$  is the change in velocity. MC-PILCO estimates the unknown function  $\Delta_t^{(i)}$  from collected data by GPR. Each  $\Delta_t^{(i)}$  is modeled as an independent GP, denoted  $f^i$ , with input vector  $\tilde{\mathbf{x}}_t = [\mathbf{x}_t^T, \mathbf{u}_t^T]^T$ , hereafter referred as GP input. Given an input-output training dataset  $\mathcal{D}^{(i)} = \{\tilde{\mathbf{X}}, \mathbf{y}^{(i)}\}$ , where the inputs are  $\tilde{\mathbf{X}} = [\tilde{\mathbf{x}}_1^T, \dots, \tilde{\mathbf{x}}_n^T]^T$ , and the outputs  $\mathbf{y}^{(i)} = [y_1^{(i)}, \dots, y_n^{(i)}]^T$  are measurements of  $\Delta_t^{(i)}$  at time instants  $t = 0, \dots, T_{tr}$ , GPR assumes the following probabilistic model,

$$\mathbf{y}^{(i)} = f^i(\tilde{\mathbf{X}}) + \mathbf{e}, \quad (4)$$

where vector  $\mathbf{e}$  accounts for noise, defined a priori as zero mean independent Gaussian noise with variance  $\sigma_i^2$ . The unknown function  $f^i$  is defined a priori as a GP with mean  $m_{\Delta}^{(i)}$  and covariance defined by a kernel function  $k(\tilde{\mathbf{x}}_{t_i}, \tilde{\mathbf{x}}_{t_j})$ , namely,  $f^i(\tilde{\mathbf{X}}) \sim N(m_{\Delta}^{(i)}, K_{\tilde{\mathbf{X}}\tilde{\mathbf{X}}})$ , where the element of  $K_{\tilde{\mathbf{X}}\tilde{\mathbf{X}}}$  at row  $r$  and column  $j$  is  $E[\Delta_{t_r}^{(i)}, \Delta_{t_j}^{(i)}] = k(\tilde{\mathbf{x}}_{t_r}, \tilde{\mathbf{x}}_{t_j})$ . The mean function  $m_{\Delta}^{(i)}$  can be derived from prior knowledge of the system, or can be set as the null function if no information is available. Instead, as regards the kernel function, one typical choice to model continuous functions is the squared-exponential kernel:

$$k(\tilde{\mathbf{x}}_{t_i}, \tilde{\mathbf{x}}_{t_j}) := \lambda^2 e^{-\|\tilde{\mathbf{x}}_{t_i} - \tilde{\mathbf{x}}_{t_j}\|_{\Lambda}^{-1}} \quad (5)$$

where  $\lambda$  and  $\Lambda$  are trainable hyperparameters tunable by maximizing the marginal likelihood (ML) of the training samples [5].

As explained in [5], the posterior distributions of each  $\Delta_t^{(i)}$  given  $\mathcal{D}^{(i)}$  are Gaussian distributed, with mean and variance expressed as follows:

$$\begin{aligned} \mathbb{E}[\hat{\Delta}_t^{(i)}] &= m_{\Delta}^{(i)}(\tilde{\mathbf{x}}_t) + K_{\tilde{\mathbf{x}}_t \tilde{\mathbf{X}}} \Gamma_i^{-1} (\mathbf{y}^{(i)} - m_{\Delta}^{(i)}(\tilde{\mathbf{X}})) \\ \text{var}[\hat{\Delta}_t^{(i)}] &= k_i(\tilde{\mathbf{x}}_t, \tilde{\mathbf{x}}_t) - K_{\tilde{\mathbf{x}}_t \tilde{\mathbf{X}}} \Gamma_i^{-1} K_{\tilde{\mathbf{X}} \tilde{\mathbf{x}}_t} \\ \Gamma_i &= K_{\tilde{\mathbf{X}} \tilde{\mathbf{X}}} + \sigma_i^2 I \end{aligned} \quad (6)$$

Then, also the posterior distribution of the one-step ahead transition model in (3) is Gaussian, namely,

$$p(\mathbf{x}_{t+1} | \mathbf{x}_t, \mathbf{u}_t, \mathcal{D}) \sim \mathcal{N}(\boldsymbol{\mu}_{t+1}, \boldsymbol{\Sigma}_{t+1}) \quad (7)$$

with mean  $\boldsymbol{\mu}_{t+1}$  and covariance  $\boldsymbol{\Sigma}_{t+1}$  derived combining (3) and (6).

2) *Policy Update*: In the policy update phase, the policy is trained to minimize the expected cumulative cost in eq. (2) with the expectation computed w.r.t. the one-step ahead probabilistic model in eq. (7). This requires the computation of long-term distributions starting from the initial distribution  $p(\mathbf{x}_0)$  and eq. (7), which is not possible in closed form. MC-PILCO resorts to Monte Carlo sampling [6] to approximate the expectation in eq. (2). The Monte Carlo procedure starts by sampling from  $p(\mathbf{x}_0)$  a batch of  $N$  particles and simulates their evolution based on the one-step-ahead evolution in eq. (7) and the current policy. Then, the expectations in eq. (2) are approximated by the mean of the simulated particles costs, namely,

$$\hat{J}(\theta) = \sum_{t=0}^T \left( \frac{1}{N} \sum_{n=1}^N c(\mathbf{x}_t^{(n)}) \right) \quad (8)$$

where  $\mathbf{x}_t^{(n)}$  is the state of the  $n$ -th particle at time  $t$ .

The optimization problem is interpreted as a stochastic gradient descend problem (SGD) [7], applying the reparameterization trick to differentiate stochastic operations [8].

The authors of [2] proposed the use of dropout [9] of the policy parameters  $\theta$  to improve exploration and increase the ability to escape from local minima during policy optimization of MC-PILCO.

## B. MC-PILCO for underactuated robotics

The task in object presents several practical issues when applying the algorithm. The first one is that the control frequency requested by the challenge is quite high for a MBRL approach. Indeed, high control frequencies require a high number of model evaluations which increases the computational cost of the algorithm. Generally, this class of systems can be controlled at relatively low frequencies, for instance, [2] and [10] derived a MBRL controller for a Furuta Pendulum at 33 Hz. Indeed, in the real hardware stage of the first edition of the competition, the MC-PILCO controller was trained to work at 33 Hz. However, the physical properties of the simulated systems (no friction) make the system particularly sensitive to the system input. For these reasons, we selected a control frequency of 50 Hz.

The second issue is that controllers are evaluated by a performance and robustness score. In the robustness test, the characteristics of the system and data acquisition vary. This is an issue for data-driven solutions like MC-PILCO since retraining of the controller is not allowed. For this reason, we decided to focus only on solving the swingup task on the nominal system, even if in our previous work we showed that MC-PILCO can be robust to noise and filtering by including these effects in the simulation.

Since the nominal model of the system is available to develop the controller, we use the forward dynamics function of the plant as the prior mean function of the change in velocity for each joint. The available model is

$$B\mathbf{u}_t = M(\mathbf{q}_t)\ddot{\mathbf{q}}_t + n(\mathbf{q}_t, \dot{\mathbf{q}}_t), \quad (9)$$

where  $M(\mathbf{q}_t)$  is the mass matrix,  $n(\mathbf{q}_t, \dot{\mathbf{q}}_t)$  contains the Coriolis, gravitational and damping terms, and  $B$  is the actuation matrix, which is  $B = \text{diag}([1, 0])$  for the Pendubot and  $B = \text{diag}([0, 1])$  for the Acrobot. We define then

$$m_\Delta(\tilde{\mathbf{x}}_t) = \begin{bmatrix} m_\Delta^{(1)} \\ m_\Delta^{(2)} \end{bmatrix} := T_s \cdot M^{-1}(\mathbf{q}_t)(B\mathbf{u}_t - n(\mathbf{q}_t, \dot{\mathbf{q}}_t)) \quad (10)$$

as the mean function in eq. (6). It is important to point out that eq. (10) is nearly perfect to approximate the system when  $T_s$  is sufficiently small, but it becomes unreliable as  $T_s$  grows. In particular, with  $T_s = 0.02$  s the predictions of eq. (10) are not good enough to describe the behavior at the unstable equilibrium. The inaccuracies of the prior mean are compensated by the GP models. To cope with the large computational burden due to the high number of collected samples, we implemented the GP approximation Subset of Regressors, see [11] for a detailed description.

An important aspect of policy optimization is the particles initialization, in this case, it is guaranteed that the system will always start at  $\mathbf{x}_0 = \bar{0}$ , therefore the initial distribution can be set to  $p(\mathbf{x}_0) \sim \mathcal{N}(\bar{0}, \epsilon I)$  with  $\epsilon$  in the order of  $10^{-4}$ .

The cost function must evaluate the policy performance w.r.t. the task requirements, in this case, we want the system to reach the position  $\mathbf{q}_G = [\pi, 0]^T$  and stay there indefinitely. A cost generally used in this kind of system is the saturated distance from the target state:

$$c_{st}(\mathbf{x}_t) = 1 - e^{-\|\mathbf{q}_t - \mathbf{q}_G\|_{\Sigma_c}^2} \quad \Sigma_c = \text{diag}\left(\frac{1}{\ell_c}, \frac{1}{\ell_c}\right), \quad (11)$$

with  $\ell_c = 3$ . Notice that this cost does not depend on the velocity of the system, just on the distance from the goal state, but it does encourage the policy to reach the goal state with zero velocity.

The policy function that is used to learn a control strategy is the general purpose policy from [2]:

$$\pi_\theta(\mathbf{x}_t) = u_M \tanh\left(\sum_{i=1}^{N_b} \frac{w_i}{u_M} e^{-\|\mathbf{a}_i - \phi(\mathbf{x}_t)\|_{\Sigma_\pi}^2}\right) \quad (12)$$

$$\phi(\mathbf{x}_t) = [\dot{\mathbf{q}}_t^T, \cos(\mathbf{q}_t^T), \sin(\mathbf{q}_t^T)]^T$$

with hyperparameters  $\theta = \{\mathbf{w}, A, \Sigma_\pi\}$ , where  $\mathbf{w} = [w_1, \dots, w_{N_b}]^T$  and  $A = \{\mathbf{a}_1, \dots, \mathbf{a}_{N_b}\}$  are, respectively, weights and centers of the  $N_b$  Gaussians basis functions, whose shapes are determined by  $\Sigma_\pi$ . For both robots, the dimensions of the elements of the policy are:  $\Sigma_\pi \in \mathbb{R}^{6 \times 6}$ ,  $\mathbf{a}_i \in \mathbb{R}^6$ ,  $w_i \in \mathbb{R}$  for  $i = 1, \dots, N_b$ , since the policy outputs a single scalar. In the experiments, the parameters are initialized as follows. The basis weights are sampled uniformly in  $[-u_M, u_M]$ , the centers are sampled uniformly in the image of  $\phi$  with  $\dot{\mathbf{q}}_t \in [-2\pi, 2\pi]$  rad/s. The matrix  $\Sigma_\pi$  is initialized to the identity. Given the ideal conditions considered in this




Fig. 1: Simulation of the Pendubot system (500 Hz), under control of the policy trained with MC-PILCO.

simulation, for the purpose of the challenge, the control switches to an LQR controller after the swing-up. Under ideal circumstances, the LQR controller has the capability to stabilize the system at an unstable equilibrium by exerting zero final torque. The switching condition is obtained by checking if the system's state is within the LQR's region of attraction.

#### IV. EXPERIMENTS

In this section, we briefly discuss how the algorithm was applied to both systems and show the main results. We also report the optimization parameters used for both systems, all the parameters not specified are set to the values reported in [2]. All the code was implemented in Python with the PyTorch [12] library.

For both robots, we use the model described in Section III-A.1, with mean function from eq. (10) and kernel function from eq. (5). The max torque  $u_M$  was set to conservative values, to improve the performance score of the controller. The policy optimization horizon was set much lower than the horizon required for the competition, this allows to reduce the computational burden of the algorithm, moreover, it pushes the optimization to find policies that can execute a fast swing-up. We exploit dropout in the policy optimization as a regularization strategy, to yield better policies.

##### A. Pendubot

The policy for the Pendubot swing-up was optimized for a horizon of  $T = 3.0$  s, with  $u_M$  set to 25% of the torque limit of the actuator. The Controller's strategy is depicted in fig. 1, in fig. 3 (left) we report the robustness bar charts. This controller has a performance score of 0.48 and a robustness score of 0.61. In table I we compare our controller's score with other tested control strategies.

Controller	Perf. score	Rob. score	Avg. score
TVLQR	0.526	0.767	0.647
MC-PILCO	0.48	0.61	0.545
iLQR MPC stab.	0.353	0.674	0.514
iLQR Riccati	0.536	0.255	0.396

TABLE I: Penubot scores comparison.




Fig. 2: Simulation of the Acrobot system (500 Hz), under control of the policy trained with MC-PILCO.

### B. Acrobot

The policy for the Acrobot swing-up was optimized for a horizon of  $T = 2.0$  s, with  $u_M$  set to 25% of the torque limit of the actuator. The Controller’s strategy is depicted in fig. 2, in fig. 3 (right) we report the robustness bar charts. This controller has a performance score of 0.316 and a robustness score of 0.25. In table II we compare our controller’s score with other tested control strategies.




Fig. 3: Pendubot (left) and Acrobot (right) robustness bar charts.

## V. CONCLUSIONS

In both systems, our MBRL approach is able to solve the task with very good swing-up time, comparable with the result of the first edition. The performance scores of the controllers are not very competitive, w.r.t. the baselines in the leaderboard<sup>2</sup>, since the score penalizes energy consumption, velocity and torque smoothness, which are not penalized in the cost function in eq. (11). As already seen in the first edition, MC-PILCO controllers are sensitive to parameter changes and delays, since they are not present in the data seen in training. However, the possibility of retraining on real hardware with few trials is a great strength of this approach. Lastly, the controllers proved a certain level of robustness when subject to actuation perturbations.

Controller	Perf. score	Rob. score	Avg. score
TVLQR	0.504	0.607	0.556
iLQR MPC stab.	0.345	0.343	0.344
MC-PILCO	0.316	0.25	0.283
iLQR Riccati	0.396	0.138	0.267

TABLE II: Acrobot scores comparison.

## ACKNOWLEDGEMENTS

Alberto Dalla Libera was supported by PNRR research activities of the consortium iNEST (Interconnected North-East Innovation Ecosystem) funded by the European Union Next GenerationEU (Piano Nazionale di Ripresa e Resilienza (PNRR) – Missione 4 Componente 2, Investimento 1.5 – D.D. 1058 23/06/2022, ECS.00000043). This manuscript reflects only the Authors’ views and opinions, neither the European Union nor the European Commission can be considered responsible for them.

## REFERENCES

- [1] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction*. MIT press, 2018.
- [2] F. Amadio, A. Dalla Libera, R. Antonello, D. Nikovski, R. Carli, and D. Romeres, “Model-based policy search using monte carlo gradient estimation with real systems application,” *IEEE Transactions on Robotics*, vol. 38, no. 6, pp. 3879–3898, 2022.
- [3] F. Wiebe, N. Turcato, A. Dalla Libera, C. Zhang, T. Vincent, S. Vyas, G. Giacomuzzo, R. Carli, D. Romeres, A. Sathuluri, M. Zimmermann, B. Belousov, J. Peters, F. Kirchner, and S. Kumar, “Reinforcement learning for athletic intelligence: Lessons from the 1st “ai olympics with realaigym” competition,” in *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, IJCAI-24* (K. Larson, ed.), pp. 8833–8837, International Joint Conferences on Artificial Intelligence Organization, 8 2024. Demo Track.
- [4] F. Wiebe, S. Kumar, L. J. Shala, S. Vyas, M. Javadi, and F. Kirchner, “Open source dual-purpose acrobot and pendubot platform: Benchmarking control algorithms for underactuated robotics,” *IEEE Robotics Automation Magazine*, vol. 31, no. 2, pp. 113–124, 2024.
- [5] C. E. Rasmussen, “Gaussian processes in machine learning,” in *Summer school on machine learning*, pp. 63–71, Springer, 2003.
- [6] R. E. Caflisch, “Monte carlo and quasi-monte carlo methods,” *Acta numerica*, vol. 7, pp. 1–49, 1998.
- [7] L. Bottou, “Large-scale machine learning with stochastic gradient descent,” in *Proc of COMPSTAT’2010*, pp. 177–186, Springer, 2010.
- [8] D. P. Kingma and M. Welling, “Auto-encoding variational bayes,” *arXiv preprint arXiv:1312.6114*, 2013.
- [9] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, “Dropout: a simple way to prevent neural networks from overfitting,” *JMLR*, vol. 15, no. 1, pp. 1929–1958, 2014.
- [10] F. Amadio, A. D. Libera, D. Nikovski, R. Carli, and D. Romeres, “Learning control from raw position measurements,” 2023.
- [11] J. Quiñero-Candela and C. E. Rasmussen, “A unifying view of sparse approximate gaussian process regression,” *Journal of Machine Learning Research*, vol. 6, no. 65, pp. 1939–1959, 2005.
- [12] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer, “Automatic differentiation in pytorch,” 2017.

<sup>2</sup>[https://dfki-ric-underactuated-lab.github.io/real\\_ai\\_gym\\_leaderboard/](https://dfki-ric-underactuated-lab.github.io/real_ai_gym_leaderboard/)