
LocoMuJoCo: A Comprehensive Imitation Learning Benchmark for Locomotion

Firas Al-Hafez¹ Guoping Zhao² Jan Peters^{1,3,4,5} Davide Tateo¹

¹Intelligent Autonomous Systems Group, TU Darmstadt

²Locomotion Laboratory, TU Darmstadt

³German Research Center for AI (DFKI)

⁴Centre for Cognitive Science

⁵Hessian.AI

{name.surname}@tu-darmstadt.de

Abstract

Imitation Learning (IL) holds great promise for enabling agile locomotion in embodied agents. However, many existing locomotion benchmarks primarily focus on simplified toy tasks, often failing to capture the complexity of real-world scenarios and steering research toward unrealistic domains. To advance research in IL for locomotion, we present a novel benchmark designed to facilitate rigorous evaluation and comparison of IL algorithms. This benchmark encompasses a diverse set of environments, including quadrupeds, bipeds, and musculoskeletal human models, each accompanied by comprehensive datasets, such as real noisy motion capture data, ground truth expert data, and ground truth sub-optimal data, enabling evaluation across a spectrum of difficulty levels. To increase the robustness of learned agents, we provide an easy interface for dynamics randomization and offer a wide range of partially observable tasks to train agents across different embodiments. Finally, we provide handcrafted metrics for each task and ship our benchmark with state-of-the-art baseline algorithms to ease evaluation and enable fast benchmarking. The code and videos can be found here:

<https://github.com/robfirmas/loco-mujoco>.

1 Introduction

Imitation Learning (IL) plays a pivotal role in machine learning, especially in robotics, enabling rapid skill acquisition without requiring manual skill programming or reward function tuning. Recent advancements, particularly in Adversarial IL and Inverse Reinforcement Learning (IRL), have expanded the applicability of IL algorithms to complex, high-dimensional tasks. In locomotion, IL is essential due to the inherent complexity of defining reward functions and limitations of classical approaches in challenging environments. However, the lack of standardized benchmarks and datasets is a significant issue in the field. Many IL works use locomotion tasks for evaluation but define their own unique environments and expert datasets. This variability hampers evaluation consistency, conceals method strengths and weaknesses, and impedes reproducibility. While many IL benchmarks are available, particularly for manipulation, we still lack a proper benchmark for the locomotion tasks. Existing benchmarks for this setting either tackle toy tasks, are very far from realistic settings and real robotic platforms or are not designed for IL, therefore lacking proper datasets.

To address this challenge, we present LocoMuJoCo, a Python-based benchmark tailored for locomotion within the context of imitation learning (IL). Our benchmark is designed to enhance method usability by offering compatibility with Gymnasium [22], a widely adopted interface for Reinforcement Learning (RL) algorithms. Additionally, we provide native support for the Mushroom-RL [6]

library. Furthermore, we supply a collection of baseline approaches capable of yielding reasonable solutions for the benchmark tasks. LocoMuJoCo encompasses various locomotion tasks, including humanoid locomotion, quadruped locomotion, and musculoskeletal human models. For each of these environments, we provide a wide variety of tasks and comprehensive datasets, including a) realistic noisy motion capture data mapped to the respective embodiment, b) ground truth expert data with actions for many environments, and c) ground truth sub-optimal expert data with actions for regularization in offline IRL or preference-based IL. This wide variety of tasks and datasets allows the benchmark to cover a broad spectrum of complexities, ranging from relatively simple tasks to extremely challenging problems yet to be fully solved. To propel application to real-world robotic systems, LocoMuJoCo allows the users to easily include domain randomization, reducing the sim-to-real gap. Finally, our benchmark provides a reward function for every task, allowing the user to measure the agent’s performance. Alternatively, the user can define his own reward function to use it either as a performance metric or to use the whole benchmark for pure RL.

Related Work. RL Bench [13] is a benchmarking platform for testing and evaluating RL algorithms in the context of robotic manipulation. It offers standardized tasks and environments and provides a planner to solve each task, allowing the use of this benchmark for IL. It comes with a realistic Franka Panda robot arm model, allowing zero-shot transfer of the policies to the real world [1]. Meta-World [24] a benchmark for multi-task and meta RL. However, it does not provide expert demonstration, rendering it unsuitable for imitation learning. Robosuite [26] is a manipulation benchmark focusing on modularity, which provides utilities for collecting human demonstrations for IL. Simitate [14] is a benchmark explicitly designed for IL, where a dataset of 1938 sequences of human tasks are provided together with a simulator. Franka Kitchen [9] is another multi-task manipulation benchmark based on the Franka Panda robot, in which multiple tasks have to be solved in parallel. While an unlabelled dataset is available, it is not suitable for IL. Adroit [18] is a small manipulation benchmark using a floating human hand, which also comes with a set of demonstrations. Gym-Mujoco [3] and DeepMind Control Suite [23] are two popular reinforcement benchmarks encompassing many locomotion tasks, ranging from 2D-walker to simple humanoid tasks. Both of the benchmarks do not explicitly focus on the IL setting and only consider toy tasks with dynamics fine-tuned for RL. D4RL [7] is an offline RL benchmark, including datasets for, inter alia, Franka Kitchen, Adroit, Gym-Mujoco, and DeepMind Control Suite. For each task, multiple datasets represent agents of different performances, which can be used for IL. Finally, the Myosuite [4] is a RL benchmark for general muscle-actuated systems, including a humanoid. However, the benchmark does not provide expert data for imitation learning. None of the aforementioned benchmarks focuses on IL for realistic locomotion, emphasizing the need for such a benchmark.

2 Benchmark Design

In this section, we introduce the framework of our benchmark, which has been meticulously crafted for versatility. This benchmark accommodates various scenarios of IL, such as cases involving IL without access to expert actions or dealing with correspondence mismatches of different embodiments. It encompasses tasks spanning a range of difficulty levels and offers user-friendliness, effectively addressing the diverse requirements of IL in realistic locomotion scenarios.

Environments are the core building blocks of this benchmark and are presented in Figure 1. At the time of writing, the benchmark includes 12 environments. The environments consist of 6 base environments: the Talos robot, the Atlas robot, the Unitree H1 humanoid, the Unitree A1 quadruped, a torque-actuated skeleton model, and a musculoskeletal human model. Additionally, both human models can be scaled to 4 different human sizes to represent an infant (~ 2-year-old), a child (~ 5-year-old), a teenager (~ 12-year-old), and an adult. To generate the different models, the links (i.e., the bones) are scaled linearly, the masses are scaled cubically, the inertias are scaled quintically, and the actuator torques or muscle forces are scaled quadratically w.r.t. the height of the human model. The skeleton models are converted models from [10], the Talos model is converted from the official library¹, the Atlas model is converted from the one provided by the Open Source Robotics Foundation², and the Unitree H1 and A1 models are taken from the MuJoCo menagerie [5].

¹https://github.com/pal-robotics/talos_robot

²<https://github.com/osrf/drcsim>




Figure 1: Overview of environments. Each task is defined by a certain dataset in an environment, e.g., the Talos carry or the muscle humanoid running task. Currently, LocoMuJoCo encompasses 12 environments with a total of 27 tasks.

Tasks are defined by a dataset representing the expert’s behavior. For example, the musculoskeletal human environment encompasses two tasks, walking and running, represented by two datasets. Each environment comes with at least two tasks, resulting in 27 tasks in total, with more to come in the future. Each dataset can be quickly loaded and replayed to allow easy inspection. Furthermore, we construct Partially Observable Markov Decision Process (POMDP) environments by randomizing the carry weight in the Atlas, Talos, and the Unitree H1 environment or randomizing the height of the skeleton model. In these cases, state masks are provided, which can be used to hide the important information regarding these tasks – weight of mass or type of humanoid. In doing so, the user can choose which part of his algorithm can get access to privileged information and which not. A typical use case is the teacher-student scenario, in which the policy only receives an observation – the masked state – and the critic gets the complete state information. Finally, we provide the possibility of initializing the simulation from either a randomly sampled or a specifically chosen expert demonstration, which has proven crucial in imitation learning settings with environment interactions [15, 17].

Datasets not only express the tasks that are supposed to be solved but are also a tool to accommodate various scenarios of IL defining a broad spectrum of difficulties. Within this benchmark, we cover the following IL paradigms: a) learning with/without access to expert actions; b) learning under mismatches between the expert’s and agent’s embodiments; c) learning with sub-optimal expert states and actions. The first paradigm is of particular importance since it is often easy to access an expert’s observation – e.g., the kinematic trajectory of a human walking – but it is often impossible to access its actions – the underlying muscle actuation. The second paradigm is essential since it is easy to collect data for some embodiment – human running – but often difficult or even impossible under some other – Atlas running. Hence, we provide realistic motion capture datasets for all environments, which naturally come with embodiment mismatches, while also providing ground truth datasets for most environments. The third paradigm is relevant in the preference-based IL setting, where preference-ranked expert datasets are used to deal with settings in which sub-optimal demonstrations are available [20]. Together with the environments and tasks, these three paradigms can be combined to define the complexity of the IL setting, allowing the user to decide on the difficulty. This makes our benchmark versatile and suitable for different IL algorithms ranging from the simplest methods – such as behavioral cloning – to state-of-the-art methods. We collected the data via the Qualisys motion capture system, positioning a comprehensive set of markers to enable the capture of whole-body movements. Subsequently, we used the marker-based datasets to calculate human joint kinematics. We performed this computation using the OpenSim software platform, using a comprehensive full-body model, adapted from the previous work [25]. In contrast, we generated the datasets for the Unitree A1 using Model Predictive Control (MPC).

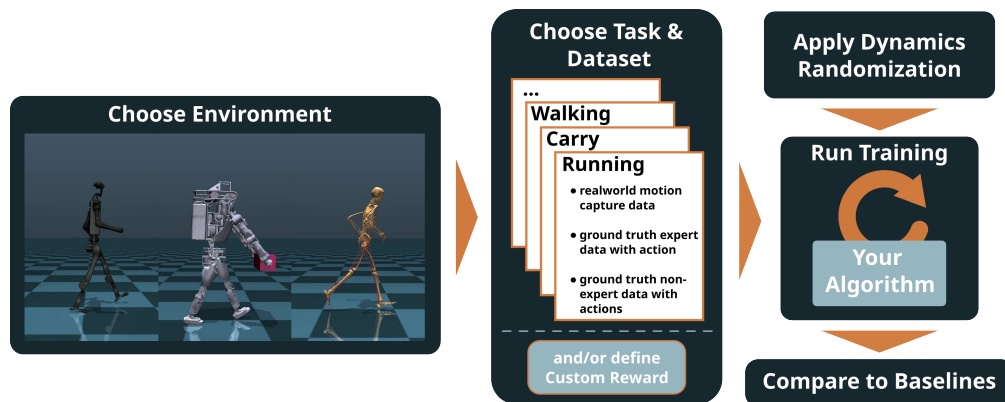


Figure 2: Training pipeline of LocoMuJoCo. First, an environment is chosen. Then, a task and dataset are chosen, and the training is started optionally with dynamics randomization. Finally, the performance of the algorithm can be compared to the expert performance or the performance of one of the provided baseline algorithms.

Training. After choosing the task and dataset, the user has the option to define a dynamics randomization configuration. LocoMuJoCo supports randomization of different parameters, such as the properties of the joints, the inertia, or the friction between bodies. The user can choose whether to sample these parameters from a Gaussian with a provided standard deviation or from a uniform distribution with provided ranges. Furthermore, we provide simple handcrafted reward functions, which are based on the underlying goal of an expert’s task. For example, we define the reward to be the difference in the center of mass velocity between the agent and the expert in the Talos walking task. Despite this reward being simple and not useful for RL to train realistic gaits, it has proven very effective for evaluation in the IL setting. Alternatively, the user can define a custom reward function, making this benchmark also suitable for pure RL. Finally, the user has the choice of implementing his algorithm in the framework of his choice by using the Gymnasium interface or by using Mushroom-RL.

Evaluation. To show the feasibility of the benchmarks, we provide a set of baselines of classical IRL and adversarial IL approaches. We base all the implementations of these approaches on the Mushroom-RL library. For the Adversarial IL setting we provide Generative Adversarial Imitation Learning (GAIL) [11], Variational Adversarial Imitation Learning (VAIL) [16] and Generative Adversarial Imitation from Observation (GAIfO) [21]. In the IRL setting we consider Inverse soft Q-Learning (IQ-Learn) [8] and the Least Squares Inverse Q-Learning (LS-IQ) [2] approaches. Furthermore, we provide the Soft Q-Imitation Learning (SQIL) [19] algorithm, a simple IL approach that uses fixed rewards for expert and demonstration trajectories. The whole pipeline of our method is presented in Figure 2.

3 Conclusion

In this paper, we introduced LocoMuJoco, a novel benchmark for imitation learning in locomotion tasks. Addressing a notable gap in the field, our benchmark enables users to evaluate IL algorithms across various locomotion tasks, spanning a wide range of difficulty levels – from relatively straightforward to highly challenging tasks. Our benchmark covers most of the open problems in IL, including learning from real-world demonstrations or without expert actions. We provide fine-tuned baselines for most tasks using state-of-the-art IL algorithms. LocoMuJoco is easily extensible and provides straightforward interfaces to common RL libraries, such as Gymnasium and Mushroom-RL. In future works, we aim to address a prevalent challenge in the majority of imitation learning (IL) benchmarks – namely, devising effective methods for quantifying the quality of the cloned behavior. The current implementation provides a simple reward function measuring the overall gait quality. Our experiments show that a higher reward corresponds to a better imitation of the learned gait. Yet, deriving metrics grounded in biomechanical principles or in theoretical principles that account for the divergence between probability distributions is essential to propel research in this field.

References

- [1] Firas Al-Hafez and Jochen Steil. Redundancy resolution as action bias in policy search for robotic manipulation. In *Proceeding of the Conference on Robot Learning*, London, United Kingdom, November 2021.
- [2] Firas Al-Hafez, Davide Tateo, Oleg Arenz, Guoping Zhao, and Jan Peters. LS-IQ: Implicit reward regularization for inverse reinforcement learning. In *ICLR 2023* [12].
- [3] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. OpenAI Gym, 2016.
- [4] Vittorio Caggiano, Huawei Wang, Guillaume Durandau, Massimo Sartori, and Vikash Kumar. Myosuite: A contact-rich simulation suite for musculoskeletal motor control. In *Proceedings of The 4th Annual Learning for Dynamics and Control Conference*, Palo Alto, California, United States, June 2022.
- [5] MuJoCo Menagerie Contributors. MuJoCo Menagerie: A collection of high-quality simulation models for MuJoCo, 2022.
- [6] Carlo D’Eramo, Davide Tateo, Andrea Bonarini, Marcello Restelli, and Jan Peters. Mushroomrl: Simplifying reinforcement learning research. *Journal of Machine Learning Research*, 22:1–5, 2021.
- [7] Justin Fu, Aviral Kumar, Ofir Nachum, George Tucker, and Sergey Levine. D4rl: Datasets for deep data-driven reinforcement learning, 2020.
- [8] Divyansh Garg, Shuvam Chakraborty, Chris Cundy, Jiaming Song, and Stefano Ermon. Iq-learn: Inverse soft-q learning for imitation. In *Proceeding of the Thirty-fifth Conference on Neural Information Processing Systems*, Sydney, Australia, December 2021.
- [9] Abhishek Gupta, Vikash Kumar, Corey Lynch, Sergey Levine, and karol Hausman. Relay policy learning: Solving long-horizon tasks via imitation and reinforcement learning. In *Proceeding of the Conference on Robot Learning*, Osaka, Japan, November 2019.
- [10] Samuel R. Hamner, Ajay Seth, and Scott L. Delp. Muscle contributions to propulsion and support during running. *Journal of Biomechanics*, 43(14):2709–2716, 2010.
- [11] Jonathan Ho and Stefano Ermon. Generative adversarial imitation learning. In *Proceeding of the Thirtieth Conference on Neural Information Processing Systems*, Barcelona, Spain, December 2016.
- [12] *Proceeding of the International Conference on Learning Representations*, Kigali, Rwanda, May 2023.
- [13] Stephen James, Zicong Ma, David Rovick Arrojo, and Andrew J. Davison. Rlbench: The robot learning benchmark & learning environment. *IEEE Robotics and Automation Letters*, 5:3019–3026, 2022.
- [14] Raphael Memmesheimer, Ivanna Kramer, Viktor Seib, and Dietrich Paulus. Simitate: A hybrid imitation learning benchmark. In *Proceedings of the 2021 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Macau, November 2019.
- [15] Xue Bin Peng, Pieter Abbeel, Sergey Levine, and Michiel van de Panne. Deepmimic: Example-guided deep reinforcement learning of physics-based character skills. *ACM Transactions on Graphics*, 37(4):143:1 – 143:14, 2018.
- [16] Xue Bin Peng, Angjoo Kanazawa, Sam Toyer, Pieter Abbeel, and Sergey Levine. Variational discriminator bottleneck: Improving imitation learning, inverse rl, and gans by constraining information flow. In *Proceeding of the International Conference on Learning Representations*, New Orleans, United States, May 2019.
- [17] Xue Bin Peng, Ze Ma, Pieter Abbeel, Sergey Levine, and Angjoo Kanazawa. Amp: Adversarial motion priors for stylized physics-based character control. *ACM Transactions on Graphics*, 40:1 – 20, 2021.

- [18] Aravind Rajeswaran, Vikash Kumar, Abhishek Gupta, Giulia Vezzani, John Schulman, Emanuel Todorov, and Sergey Levine. Learning complex dexterous manipulation with deep reinforcement learning and demonstrations. In *Robotics: Science and Systems*, Pittsburgh, Pennsylvania, United States, June 2018.
- [19] Siddharth Reddy, Anca D. Dragan, and Sergey Levine. Sqil: Imitation learning via reinforcement learning with sparse rewards. In *Proceeding of the International Conference on Learning Representations*, Virtual, May 2020.
- [20] Aleksandar Taranovic, Andras Kupcsik, Niklas Freymuth, and Gerhard Neumann. Adversarial imitation learning with preferences. In ICLR 2023 [12].
- [21] Faraz Torabi, Garrett Warnell, and Peter Stone. Generative adversarial imitation from observation. In *Proceeding of the International Conference on Machine Learning*, Long Beach, California, July 2019.
- [22] Mark Towers, Jordan K. Terry, Ariel Kwiatkowski, John U. Balis, Gianluca de Cola, Tristan Deleu, Manuel Goulão, Andreas Kallinteris, Arjun KG, Markus Krimmel, Rodrigo Perez-Vicente, Andrea Pierré, Sander Schulhoff, Jun Jet Tai, Andrew Tan Jin Shen, and Omar G. Younis. Gymnasium, March 2023.
- [23] Saran Tunyasuvunakool, Alistair Muldal, Yotam Doron, Siqi Liu, Steven Bohez, Josh Merel, Tom Erez, Timothy Lillicrap, Nicolas Heess, and Yuval Tassa. dm_control: Software and tasks for continuous control. *Software Impacts*, 6:100022, 2020.
- [24] Tianhe Yu, Deirdre Quillen, Zhanpeng He, Ryan Julian, Karol Hausman, Chelsea Finn, and Sergey Levine. Meta-world: A benchmark and evaluation for multi-task and meta reinforcement learning. In *Conference on Robot Learning (CoRL)*, 2019.
- [25] Guoping Zhao, Martin Grimmer, and Andre Seyfarth. The mechanisms and mechanical energy of human gait initiation from the lower-limb joint level perspective. *Scientific Reports*, 11(1):22473, 2021.
- [26] Yuke Zhu, Josiah Wong, Ajay Mandlekar, Roberto Martín-Martín, Abhishek Joshi, Soroush Nasiriany, and Yifeng Zhu. robosuite: A modular simulation framework and benchmark for robot learning. In *arXiv preprint arXiv:2009.12293*, 2020.