Leveraging Spring Mass Locomotion to Guide Learned Walking Controllers

Kevin Green, Yesh Godse, Jeremy Dao, Ross L. Hatton, Alan Fern, and Jonathan Hurst
Collaborative Robotics and Intelligent Systems Institute, Oregon State University
{*greenkev, godsey, daoje, ross.hatton, alan.fern, jhurst*}@*oregonstate.edu*

***Abstract*—In this work, we describe an approach to achieving dynamic legged locomotion on physical robots which uses existing reduced order models of locomotion to guide reinforcement learning. Ultimately, our goal is that this work leads to a control hierarchy where the highest-level behaviors are planned through reduced-order models and lower level controllers utilize a learned policy that can bridge the gap between the idealistic model and the physical robot. Here we present the learned dynamic walking controller by showing that a range of walking motions from reduced-order models can be used as the command and primary training signal. Instead of a reactive planner, we use a library of optimized periodic walking motions from an actuated SLIP model. The resulting policies do not attempt to naively track the motion as a traditional trajectory tracking controller would, but instead balance immediate motion tracking with long term stability. The resulting controller is demonstrated on a physical Agility Robotics Cassie series bipedal robot at speeds up to 1.2 m/s. This work builds the foundation of a generic, dynamic learned walking controller that can be applied to many different tasks.**

I. Motivation

Our ultimate goal is an agile, dynamic legged robot utilizing a control hierarchy consisting of a terrain aware motion planner that issues motion commands to a low level reactive learned locomotion controller. Learned controllers have incredible potential to create dynamic locomotion, but to be integrated into this type of control hierarchy we need an effective control interface [2]. A compelling interface is the motion of reduced-order models of locomotion because they capture the core dynamics of locomotion including foot placement and underactuation [1]. Here we present a method of using reduced-order models of walking to direct high quality, transferable walking controllers (Fig. 1) and demonstrate its effectiveness on a Cassie series robot from Agility Robotics.

II. Methods

As a stand in for a more sophisticated motion planner we use a library of reduced order model motions in our control hierarchy. These motions are optimized periodic trajectories of a 3D bipedal actuated SLIP model between 0 and 2 m/s.

We then train a policy using PPO in a MuJoCo simulation of the robot and state observer. The policy we use takes the form of a two hidden layer neural-network. The reward function we use is mostly taken from the reference motion: 30% matching the body velocity of the model trajectory, 30% matching the foot placement relative to the body and 10% minimizing overall body position drift. The remaining 30% of the reward function consist of minimizing foot rotation and a smooth action reward.

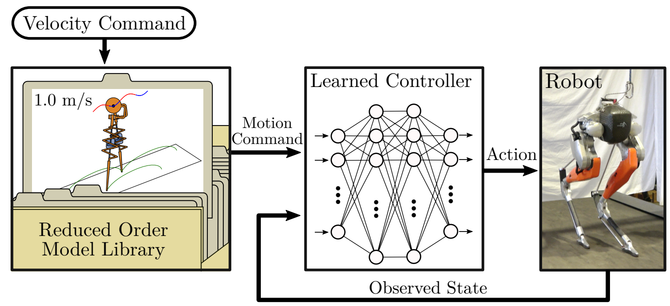

Fig. 1: Our proposed control hierarchy where a learned controller is commanded using reduced-order model motions.

III. Results

We directly transfer the policies trained in simulation to hardware, demonstrating that this approach can achieve a strong sim-to-real transfer. We observe that the produced motion resembles the underlying spring-mass motion with slight oscillations in the pelvis velocity and changes in leg length directly corresponding to the same variations in the motion of the reduced-order model. In simulation the overall error in foot positions is within 10 cm at high speeds, and as low as 3 cm for lower speeds which shows that the reference trajectories are acting as guides for the induced behavior. The success of this work will enable many different opportunities for higher level planning in future work. Now that we have a policy capable of following a desired reduced-order model motion, we can work to extend this to generate policies that follow any arbitrary reduced-order model trajectory. Having reactive planning decisions like navigation and obstacle avoidance happen at the reduced-order model level would make it significantly easier to achieve fully autonomous agile legged robots.

Acknowledgments

This work was supported by DARPA contract W911NF-16-1-0002, NSF Grant No. CMMI-1653220, IIS-1849343 and DGE-1314109. Thanks to Intel Labs for computing resources and support

References

[1] Igor Mordatch, Martin de Lasa, and Aaron Hertzmann. Robust physics-based locomotion using low-dimensional planning. *ACM Transactions on Graphics*, 2010.

[2] Zhaoming Xie, Patrick Clary, Jeremy Dao, Pedro Morais, Jonathan Hurst, and Michiel van de Panne. Learning Locomotion Skills for Cassie: Iterative Design and Sim-to-Real. In *3rd Conference on Robotic Learning*, 2019.