Pushing the Limits: Running at 3.2m/s on Cassie

Jeremy Dao, Helei Duan, Kevin Green, Jonah Siekmann, Yesh Godse, Jonathan Hurst, Alan Fern
Collaborative Robotics and Intelligent Systems Institute, Oregon State University
{daoje, duanh, greenkev, siekmanj, godsey, jonathan.hurst, alan.fern}@oregonstate.edu

I. Background

Recent bipedal locomotion learning research has found success in describing bipedal gaits as alternating swing and stance phases through a periodic cost function [1]. By periodically penalizing foot forces (swing) and foot velocities (stance) according to a probabilistic function defined by a set of "gait parameters" like swing ratio and a phase offset, it is possible to learn a policy that can create all bipedal gaits.

Our further work finds that for the simpler task of just walking we are able to use a much simpler formulation that results in a similar reward structure. Furthermore, this simplification allows the policy to further specialize for a single task and achieve much higher running speeds of 3.2 m/s, the highest speed of a unsupported bipedal robot achieved so far.

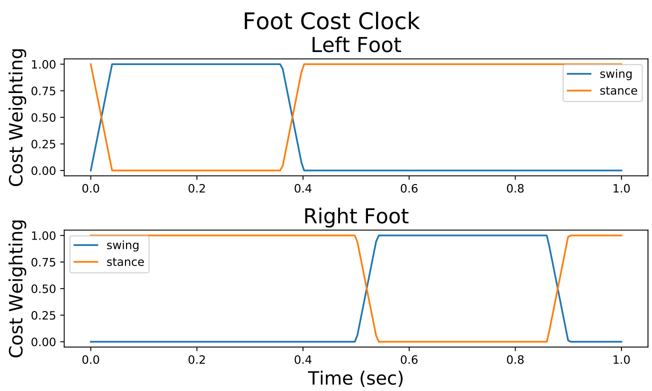

Fig. 1: An example of the piece-wise linear clock used to describe the swing and stance cost weightings for each foot. This clock is constructed with a swing ratio of 40% and a cycle time of 1 second.

II. Methods

Our methods build off of our previous work presented in Siekmann et. al [1] with a few simplifications to the reward and learning structure.

The primary modification we make is to replace the probabilistic function with just a piecewise-linear cyclic function. This function, which is still dependent on a given swing ratio, then defines the weighting of the foot force and foot velocity cost components. A plot of an example clock with a swing ratio of 40% and a cycle time of 1 second is shown in Fig. 1

We use a feedforward neural network ([1] uses an LSTM recurrent neural network) trained using the PPO algorithm in a MuJoCo simulation of Cassie. During training we arbitrarily assign (using visual heuristics) a single given swing ratio and stepping frequency for each possible commanded speed. We train on speeds between 0 and 4 m/s. For the corresponding swing ratio and stepping frequencies commanded at each speed refer to Table I. For speeds 1m/s to 3m/s we linearly increase swing ratio from 0.4 to 0.8 and stepping frequency from 1Hz to 1.5Hz in relation to the speed.

Fig. 2: A picture of Cassie running at 3.2m/s.

Speed Range [m/s]	Swing Ratio [%]	Step Frequency [Hz]
0 - 1	40	1
1 - 3	40 - 80	1 - 1.5
3 - 4	80	1.5

TABLE I: Heuristics of gait parameters used in our running policy.

III. Results

Using our method, we are able to train policies that are capable of walking at speeds between 0 and 4 m/s in simulation and can achieve a maximum speed of 3.2m/s on hardware. We are able to transfer these simulation trained policies directly to hardware. We believe this to be state of the art performance on hardware for bipedal locomotion, achieving the fastest running speed on a such a robot so far. A video of our hardware results can be viewed at: https://www.youtube.com/watch?v=PcGhlbkneGU

Acknowledgment

This work is supported by NSF Grant No. IIS-1849343, DGE-1314109 and DARPA Contract W911NF-16-1-0002. Thanks to Intel's vLab team for computing resources.

References

[1] Jonah Siekmann, Yesh Godse, Alan Fern, and Jonathan Hurst. Sim-to-Real Learning of All Common Bipedal Gaits via Periodic Reward Composition, 2021.