Real2Sim Transfer using Differentiable Physics

Eric Heiden*, David Millard*, and Gaurav S. Sukhatme
Department of Computer Science
University of Southern California
{heiden,dmillard,gaurav}@usc.edu

***Abstract*—Accurate simulations allow modern machine learning techniques to be applied to robotics problems, with sample-collection runtimes orders of magnitudes faster than the real world. Current reinforcement learning approaches require laborious manual calibration of carefully designed models, or, in a model-free context, vast amounts of training data to acquire such accurate models from real-world trials. In this work, we introduce a new layer in the deep learning toolbox that imposes a strong inductive bias to generate physically accurate predictions of rigid-body dynamics and allows for the automatic inference of system parameters given an ad-hoc model description.**

I. Introduction

Reinforcement learning (RL) enables robots to learn robust policies from experience. Since most state-of-the-art RL algorithms suffer from a prohibitively high sample complexity to train in the real world, policies are typically trained in a simulator first before being transferred onto the actual system. Discrepancies between simulation and reality impair the performance of such policies. There are many approaches to solving this transfer learning problem, including improving simulation fidelity and domain randomization techniques [1].

Instead of open-loop Sim2Real transfer, in this work, we investigate how simulators can be used as a model of the real world that can be updated from experience. By maintaining a probabilistic representation of physical parameters, simulators can play a role in designing exploration policies that excite unseen areas of the physical parameter space. In this way, simulators and trained policies can jointly improve in a Real2Sim and Sim2Real feedback process. The Real2Sim aspect, which this paper focuses on, is an approach to system identification [2] where the parameters of a model are inferred from real-world observations.

This work is based on our recent article on a differentiable physics engine [3] that is currently under review. This paper differs significantly by focusing on the important robot learning problem of Sim2Real transfer. We present new experiments where we predict the motion of a real-world double pendulum and show how probabilistic estimation can be leveraged to attain a consistent model of the real world that opens avenues to a principled strategy of domain randomization, a commonly used technique in transfer learning.

II. Rigid Body Dynamics

Throughout this paper, we estimate quantities of a kinematic chain of rigid bodies that are connected by joints. Following [4], we implement the Articulated Body Algorithm which computes the forward dynamics, i.e., the joint accelerations at the next time step given the current joint positions and velocities. Subsequently, we advance the system dynamics using semi-implicit Euler integration. Having implemented the entire physics engine in the C++ automatic differentiation framework Stan Math [5] allows us to compute accurate gradients of any quantity in our simulator, achieving fast convergence with gradient-based optimizers that leverage our engine to estimate model parameters, optimize trajectories, or derive control policies. Throughout this paper, we denote a single step of our physics simulator by $f_\theta(\cdot)$ – a function conditioned on model parameters θ which returns the next world state given the current.

III. Real2Sim Transfer

To demonstrate the capability of our approach to estimate physical parameters of chaotic systems in the real world, we estimate the kinematic parameters of a compound pendulum. The dynamics of a pendulum are fully determined by the length of each link. Using a VICON motion capture system, we obtain sub-millimeter accurate positional trajectories.

Estimating the length of each link is trivial given the Cartesian coordinates of the pendulum masses. Instead, we exercise our approach by estimating the lengths from a trajectory $\tau = \{\mathbf{q}_0, \cdots, \mathbf{q}_T\}$, where $\mathbf{q}_t$ are the generalized joint positions at time step t. We model the double pendulum using spherical joints whose 3D rotations are defined by quaternions. Given $\mathbf{q}_0 = \mathbf{q}_0^*$, we minimize

$$\mathcal{L} = \sum_t ||z(f_\theta(\mathbf{q}_{t-1})) - z(\mathbf{q}_t^*)||_2^2, \tag{1}$$

where z computes unit heading vectors for given quaternions, and $\mathbf{q}_t^*$ are the joint coordinates of the real pendulum.

Leveraging the differentiability of f_θ, we employ the gradient-based Adam optimizer to estimate the link lengths. We converge after ca. 80 epochs for trajectories 10 steps long, sampled at 25 Hz (cf. Fig. 1). Our performance is hampered

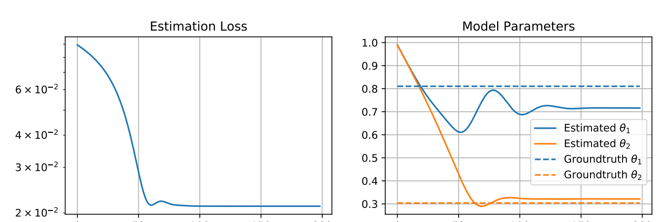

Fig. 1. Estimation of the two link lengths (in meters) of a real double pendulum modeled via spherical joints.

* Equal contribution