

GAUSSIAN PROCESS REGRESSION FOR SIM-TO-REAL TRANSFER OF HOPPING GAITS

Jeremy Krause^{1,*}, Adel Alaeddini², Pranav A. Bhounsule^{1,*}

¹Department of Mechanical and Industrial Engineering, University of Illinois at Chicago 842 W. Taylor St., Chicago, IL 60607.

²Department of Mechanical Engineering, The University of Texas at San Antonio, 1 UTSA Circle, San Antonio, TX 78249.

ABSTRACT

Simulation-based controllers are relatively easy to build and evaluate, but rarely transfer seamlessly to hardware. This is because of the reality-gap which is the discrepancy between simulations and hardware. Narrowing the reality-gap can speed up the deployment of simulated controllers to hardware, also known as sim-to-real. This paper presents sim-to-real transfer of controllers on a single leg hopping robot, a system that cycles between under-actuation during the stance phase to no-actuation during flight phase. Using simulations, we design a controller to achieve speed and height regulation once-per-step, but the controller cannot achieve accurate control on hardware. Using data from hardware, we model the mis-match between simulation and hardware using Gaussian Process Regression (GPR), recompute the controller, and redeploy it in hardware. It takes about 4 iterations to achieve accurate tracking. The results show that when GPR is used to model the step-to-step level model inaccuracy, it can lead to high accuracy sim-to-real transfer while maintaining sample efficiency. A video is here: tiny.cc/idetc2023

Keywords: Legged Robots, Gaussian Process, Sim-to-Real, Poincaré Map

1. INTRODUCTION

To accelerate the development of robotic applications in the real world, it is important to use a simulation-based methodology to develop controllers. Indeed, simulations are computationally cheap (e.g., take a fraction of the time to set up and run), can be fully automated (e.g., running test cases without supervision), and inexpensive (e.g., robots don't break in simulation) as compared to hardware. This allows one to test various controller ideas, refine them, and tune them on simulation before being deployed on hardware. However, even the best simulators have limitations. For example, it is difficult to model friction, mass, inertia, deformations, noise, etc., with high accuracy leading to the reality gap – differences between hardware and simulation. It is important

to close this gap to enable seamless deployment of simulation-based controllers to hardware, also known as sim-to-real transfer. This work presents a technique for sim-to-real transfer by modeling the discrepancy between simulation and hardware and tuning the controller based on the improved model. In particular, we demonstrate the efficacy of the approach on a single leg hopping robot.

2. BACKGROUND AND RELATED WORK

A single leg robot that can move by hopping is a good system for testing and benchmarking because the system is simple enough to model, yet offers significant control challenges that must be overcome to achieve accurate step-step control. One of the most challenging aspect of the hopping system is that it is under-actuated (less actuators than degrees of freedom) when the foot is in contact with the ground and un-actuated (no control) in the flight phase. Hence, to enable stable hopping the controller needs to provide accurate intra-step control and predict the future inter-step control, beginning at the start of the next contact phase, based solely on the state at the end of the preceding contact phase with limited data pertaining to encompassed flight phase.

The earliest hopping robot was built by Raibert [1], who showed that simple speed and height regulation by foot placement control, in the flight phase, and torque control, in the stance phase, can achieve stable hopping from step-to-step. Raibert's method of control may be formalized by using Poincaré sectioning and mapping [2]. The Poincaré section is the states at an instant in the locomotion cycle (e.g., apex for the hopper) and the Poincaré map is a function that maps the states from one Poincaré section to the next. The linearization over the Poincaré map may be used to quantify the stability of the hopping gait. The Poincaré map approach has been widely used for the control of hopping robot (for example, [3], [4], [5]).

Sim-to-real transfer has gained increased attention in recent years. Dynamic randomization has become a popular approach to achieve efficient sim-to-real transfer [6, 7]. The control policy (usually a neural network) is trained by randomizing the features

*Corresponding author: jkraus3@uic.edu, pranav@uic.edu

of the simulation (mass, inertia, friction, latency, etc.). The expectation is that the resulting control policy is robust towards variability and, hence, has a higher chance of succeeding on hardware. Such an approach is computationally demanding and, therefore, not scalable. The scalability issue has been overcome, to some extent, by seeding the learning with samples from motion capture [8] or training in a lower dimensional space, such as in the cartesian space [9]. Another approach is to use experimental data to improve sub-system model (e.g., actuator dynamics [10]).

Control improvements may also be achieved directly in hardware by using sample-efficient optimization such as Bayesian Optimization [11]. Here, the cost function (e.g., speed, energy) is mapped to parameters of the control policy using a Gaussian Process Regression (GPR). During training on hardware, the control policy parameters are directly optimized using an acquisition function, such as the probability of improvement, predicted improvement, and confidence bounds. However, GPRs do not scale well with the input space (control parameters, in this case) and data points. The first issue may be fixed by using latent space representations that map high dimensional inputs to low dimensional space for the GPRs [12, 13] and the second issue may be addressed by using sparse GPR [14]. Another approach is to use GPR to model the system dynamics and then use the GPR for control [15].

In this paper, we take a different approach. Instead of learning a control policy in hardware using GPR, we learn an error correction between the simulation-based model and the physical hardware using GPR. This approach has been used for bongo board balance [16], control of a blimp [17], control of a quadcopter [18], and control of unmanned ground vehicles [19]. The trained model is then used to develop an improved controller using existing methods (e.g., model-based control), which provides an increased correlation between the simulation-based and physical models over the original, purely simulation-based, model. In the past approaches, the corrected model is estimated over the time domain (typically 20 - 200 Hz). This leads to scalability issues with GPRs. We exploit the Poincaré map representation, which is a low-dimensional model, for control purposes. We use the GPRs to learn a correction to the Poincaré map, estimated over one or more steps (typically 2 Hz for slow systems like walking and 10 Hz for fast systems like hopping), ensuring scalability of the proposed approach. The novelty of this approach is to use Poincaré-map-based model correction, using GPRs, to ensure scalability of the approach and its demonstration on a custom-built monoped hopping robot.

The paper is organized as follows. In Sec. 3 we provide details of the physics-based model. In Sec. 4 we present the low-level and the step-level controller. In Sec. 5, we present details of the custom-built hardware. This is followed by Results in Sec. 6, Discussion in Sec 7, and Conclusion and Future work in Sec. 8.

3. MODEL

Figure 1 shows the 2D model based on the custom-built hopping robot. The hopper is confined to the sagittal plane using a boom. The coordinates of the boom-hopping robot connection are (x, y) . The hopper leg consists of a 4-link, closed chain, mechanism. The links of the 4-link mechanism are symmetrical about

the mechanism vertical axis except for the protruding second link (foot) on one side as denoted by ℓ_3 in figure 1. Only two angles are needed to define the kinematics of the chain. We define α_1 as the angle between the vertical and the line joining the boom axis to the joint between bottom two links and α_2 as half the angle between the top two links.

The links are connected to each other and to the boom through pin joints. The robot top links have length $\ell_1 = 0.11$ m and bottom links have length $\ell_2 = 0.2$ m. The top links have mass $m_1 = 0.043$ kg and inertia $I_1 = 0.0001$ kg/m² while the bottom link have mass $m_2 = 0.047$ kg and inertia $I_2 = 0.00025$ kg/m². The center of mass of the top link is at a distance of $c_1 = 0.0485$ m from the boom axis and the center of mass of the bottom link is at a distance of $c_2 = 0.1054$ m from the pin joint connecting the top and bottom link. The foot protrusion of the second link is $\ell_3 = 0.045$ m. Gravity is $g = 9.81$ m/s² and points vertically downward. The mass of the hopping robot and the inertia of the attached boom are $M = 1.466$ kg and $I = \text{diag}\{I_x, I_y, I_z\}$ where $I_x = I_z = 5$ kg/m² and $I_y = 0.0015$ kg/m².

We use Euler-Lagrange equations to describe the movement of the hopper. The hopper movement consists of a stance phase, where we assume that the foot is in contact with the ground, and a flight phase, where the hopper moves freely under the effects of gravity and momentum. There are two phase transitions, the collision transition, where the hopper transitions from a flight phase to a stance phase, and the take-off transition, where the hopper transitions from a stance phase to a flight phase. These are described in detail next.

3.1 Stance phase equations

The state variables for derivation are defined as $\mathbf{q} = [x \ y \ \alpha_1 \ \alpha_2]^T$. The Lagrangian can be written as

$$\begin{aligned} \mathcal{L} &= \mathcal{T} - \mathcal{V}, \\ &= 0.5(Mv^T v + I_x \dot{y}^2/R + I_z \dot{x}^2/R + I_y \omega^2) + \dots \\ &\quad \sum \left(m_i v_i^T v_i + I_i \omega_i^2 \right) - Mgy - \sum \left(m_i g y_i \right), \end{aligned} \quad (1)$$

where v_i , ω_i , y_i are the linear velocity, angular velocity, and relative vertical position at the center of mass of link i respectively and boom length is $R = 1.854$ m. We take the summation over the 4 links. Using the Euler-Lagrange equations gives 4 equations,

$$\mathbf{M}(\mathbf{q})\ddot{\mathbf{q}} + \mathbf{N}(\mathbf{q}, \dot{\mathbf{q}}) = \mathbf{B}\mathbf{u} + \mathbf{J}^T \mathbf{P}, \quad (2)$$

where \mathbf{M} , \mathbf{N} , \mathbf{B} are the mass matrix, accelerations due to Coriolis, centrifugal acceleration and gravity, and torque selection matrices. The control torques are $\mathbf{u} = [\tau_1 \ \tau_2]^T$ where τ_i is the torque for the two top links, \mathbf{J} is the Jacobian of the contact point with respect to the degrees of freedom and \mathbf{P} is the ground reaction force on the stance leg.

The foot does not move during stance phase (nonslip assumption). This can be computed by differentiating the position of the contact points twice, $\ddot{x}_{\text{foot}} = \ddot{y}_{\text{foot}} = 0$ to give

$$\mathbf{J}\ddot{\mathbf{q}} = -\dot{\mathbf{J}}\dot{\mathbf{q}}. \quad (3)$$

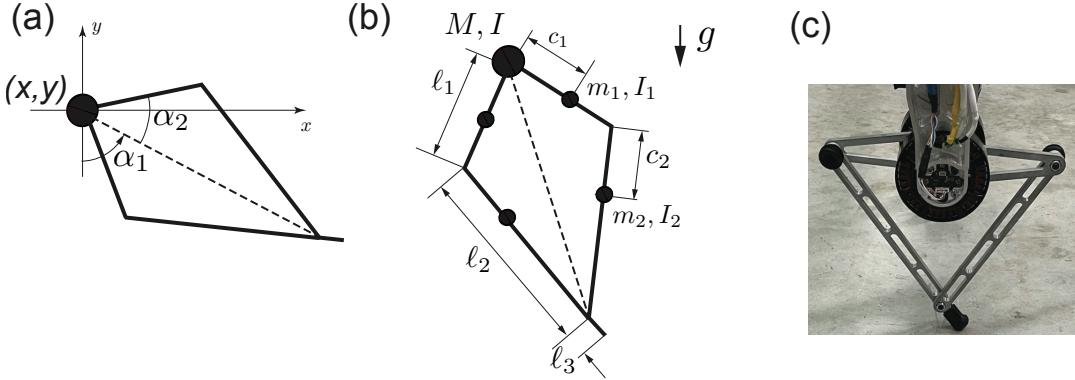


FIGURE 1: 2D MODEL: (A) CONFIGURATION VARIABLES DESCRIBING THE DEGREES OF FREEDOM, (B) MASS, CENTER OF MASS, INERTIA ABOUT CENTER OF MASS, AND LENGTH PARAMETERS, (C) HARDWARE

We can combine Eqns. 2 and 3 to get

$$\begin{bmatrix} \mathbf{M}(\mathbf{q}) & -\mathbf{J}^T \\ \mathbf{J} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \ddot{\mathbf{q}} \\ \mathbf{P} \end{bmatrix} = \begin{bmatrix} -\mathbf{N}(\mathbf{q}, \dot{\mathbf{q}}) + \mathbf{B}\mathbf{u} \\ -\mathbf{j}\dot{\mathbf{q}} \end{bmatrix}. \quad (4)$$

3.2 Flight phase equations

The swing phase equations are derived from the stance phase equations by setting $\mathbf{P} = 0$ in Eqn. 2 to get

$$\mathbf{M}(\mathbf{q})\ddot{\mathbf{q}} + \mathbf{N}(\mathbf{q}, \dot{\mathbf{q}}) = \mathbf{B}\mathbf{u}. \quad (5)$$

3.3 Flight to stance phase transition

The transition from flight to stance phase occurs when the foot makes contact with the ground. In simulation, this transition occurs when $y_{\text{foot}} = 0$. On hardware, due to deflection in the boom arm, $y_{\text{foot}} = 0$ is not a precise indicator of the flight-to-stance state transition, and thus, on hardware, transition is considered to have occurred when vertical ground reaction force \mathbf{P}_y exceeds a threshold value.

During the transition, the position variables are assumed to be the same while the velocity variables are assumed to change based on the conservation of angular momentum about the stance foot. We obtain the velocity during transition by integrating Eqn 2 and taking the limit as time goes to 0 to yield

$$\mathbf{M}(\mathbf{q}^-)(\dot{\mathbf{q}}^+ - \dot{\mathbf{q}}^-) = \mathbf{J}^T \mathbf{P}_t, \quad (6)$$

where the superscripts $-$ and $+$ indicate instants before and after collision and \mathbf{P}_t is the impulse on the stance foot. The foot comes to rest after the collision. This is given by $\dot{y}_{\text{foot}}^+ = 0$ to get

$$\mathbf{J}\dot{\mathbf{q}}^+ = \mathbf{0}. \quad (7)$$

We can combine Eqns. 6 and 7 to get

$$\begin{bmatrix} \mathbf{M}(\mathbf{q}^-) & -\mathbf{J}^T \\ \mathbf{J} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \dot{\mathbf{q}}^+ \\ \mathbf{P}_t \end{bmatrix} = \begin{bmatrix} \mathbf{M}(\mathbf{q}^-)\dot{\mathbf{q}}^- \\ \mathbf{0} \end{bmatrix}. \quad (8)$$

3.4 Stance to flight phase transition

On simulation, the stance to flight phase transition occurs when the vertical ground reaction force $\mathbf{P}_y = 0$. This reaction force is obtained from \mathbf{P} in Eqn. 4. On hardware, due to physical limitations, the stance-to-flight transition is considered to have occurred when the 4-link leg member has extended to the safe maximum value, since relying on ground reaction force would lead to a collision of the 4-link member links. During the transition from stance to flight phase the positions and velocities are assumed to remain the same.

4. CONTROLLER

4.1 Low-level controller

The low-level controller runs at 1 kHz. The controller takes in sensor data, the angles and rates, and outputs joint torque torques to the two servo drives. There are separate controllers for the stance and the flight phase.

Flight Phase controller: The flight phase controller uses a position-derivative controller to set the joint angles. In the flight phase we input a desired leg length ℓ_{des} and leg orientation θ_{des} . There is a kinematic mapping $[\ell, \theta] = \mathbf{f}(\alpha_1, \alpha_2)$. We are able to analytically invert \mathbf{f} to compute the desired angles $(\alpha_1, \alpha_2)_{\text{des}} = \mathbf{f}^{-1}(\ell_{\text{des}}, \theta_{\text{des}})$. The feedback controller can then be written as $\tau = [\tau_1 \ \tau_2] = -\mathbf{K}_p(\alpha - \alpha_{\text{des}}) - \mathbf{K}_d(\dot{\alpha} - \dot{\alpha}_{\text{des}})$, where \mathbf{K}_p and \mathbf{K}_d are the proportional and damping gain respectively. There is a feedback control on the desired landing angle, $\theta_{\text{des}} = K_1(\dot{x} - \dot{x}_{\text{des}})$, where K_1 and \dot{x}_{des} are the proportionality gain and the desired speed respectively.

Stance Phase controller: The stance phase controller regulates the force in the leg. The leg is simulated to contain a virtual spring. The leg force is $F = -K_2(\ell - \ell_0)$, where K_2 and ℓ_0 are the user specified stiffness and the rest length of the virtual spring, respectively. The force is then resolved in the x- and y-direction as follows $F_x = F \cos(\alpha_1)$ and $F_y = F \sin(\alpha_1)$. Using the Jacobian \mathbf{J} from the toe to the joints, we compute the desired torques, $\tau = [\tau_1 \ \tau_2] = \mathbf{J}^T [F_x \ F_y]$.

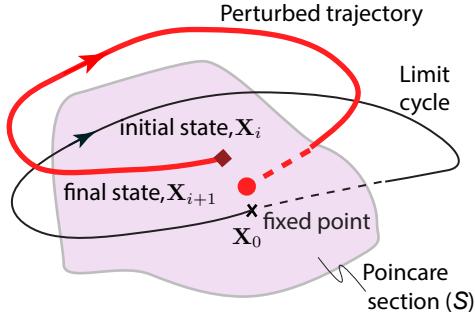


FIGURE 2: POINCARÉ MAP AND SECTION

4.2 High-level (Step-level) controller

Our goal is to track a desired apex horizontal speed \dot{x}_{des} and desired apex height y_{des} . This is achieved using a step-level, Poincaré map based, control. The step-level control attempts to regulate the speed and height once per step, hence the name. We describe the Poincaré map/section that is used as a model for the step-level controller below.

Poincaré section and map: Figure 2 shows pictorial depiction of the Poincaré section and map [2]. Consider an instant in the gait cycle (e.g., instant when the hopper is the apex). This is the Poincaré section. Let the reduced state (which is a subset of the complete state) at the Poincaré section in the current step be \mathbf{X}_i . This is shown with the red diamond. Our low-level control parameterization chooses the control \mathbf{U}_i (e.g., amplitude, gain, set-point), which takes the reduced state to \mathbf{X}_{i+1} at the Poincaré section in the next step. There is a function \mathbf{P} , known as the Poincaré map, that maps the reduced state from one step to the next. This is given by:

$$\mathbf{X}_{i+1} = \mathbf{P}(\mathbf{X}_i, \mathbf{U}_i). \quad (9)$$

Note that the control \mathbf{U}^i is set once per step and kept constant during the step. We can also find a state-control pair $(\mathbf{X}_0, \mathbf{U}_0)$ that would repeat itself at the next step, $\mathbf{X}_0 = \mathbf{P}(\mathbf{X}_0, \mathbf{U}_0)$. This state, \mathbf{X}_0 , is known as the fixed point and it gives rise to a periodic gait known as the limit cycle.

For the hopper we choose the Poincaré map at the apex where $\dot{y}_i = 0$ and $\mathbf{X}_i = [\dot{x}_i \ y_i]$ and $\mathbf{U}_i = [f\dot{x}_i \ f y_i]$ where $f\dot{x}_i$ and $f y_i$ are forces, constant for the entire stance phase, added to the stance level controller (described in Sec. 4.1) after the mid-stance phase (when $\dot{y} > 0$).

Thus, we have

$$[\dot{x}_{i+1} \ y_{i+1}] = \mathbf{P}([\dot{x}_i \ y_i \ f\dot{x}_i \ f y_i]). \quad (10)$$

Unfortunately, there is no analytical solution to the Poincaré map, \mathbf{P} . We use numerical integration to determine the state $[\dot{x}_{i+1} \ y_{i+1}]$ for the given inputs $[\dot{x}_i \ y_i \ f\dot{x}_i \ f y_i]$.

4.3 Step-level control problem

The step-level control problem is stated as follows. Given the state at the current step, $[\dot{x}_i \ y_i]$, compute the control at the

current step, $[f\dot{x}_i \ f y_i]$, such that the state at the next step is $[\dot{x}_{i+1} \ y_{i+1}] = [\dot{x}_{\text{des}} \ y_{\text{des}}]$. From Eqn. 10 we can write

$$[\dot{x}_{\text{des}} \ y_{\text{des}}] = \mathbf{P}([\dot{x}_i \ y_i \ f\dot{x}_i \ f y_i]). \quad (11)$$

These are two nonlinear equations and two control variables $f\dot{x}_i$ and $f y_i$. These equations may be solved using a nonlinear root finding algorithm, such as the MATLAB function fmincon, which is utilized here.

4.4 Sim-to-real transfer

The controller derived from the purely simulation-based Poincaré map, equation 11, leads to relatively poor tracking on hardware. This is known as the reality gap [7]; simulation-based controllers do not always work as intended on hardware. We use a Gaussian Process Regression to update the Poincaré map as follows [20].

For a given reference horizontal speed and height, we use the simulation-based Poincaré map, equation 11, to compute a control. We store the state obtained on the hardware, $[\dot{x}_{i+1}^{\text{true}} \ y_{i+1}^{\text{true}}]$. We then fit the error in the Poincaré map as follows:

$$\begin{aligned} [\dot{x}_{i+1}^{\text{true}} \ y_{i+1}^{\text{true}}] - \mathbf{P}([\dot{x}_i \ y_i \ f\dot{x}_i \ f y_i]) = \dots \\ GP([\dot{x}_i \ y_i \ f\dot{x}_i \ f y_i]). \end{aligned} \quad (12)$$

where GP is the Gaussian Process Regression model.

Then we re-compute the control using the improved Poincaré map as follows

$$\begin{aligned} [\dot{x}_{\text{des}} \ y_{\text{des}}] = \mathbf{P}([\dot{x}_i \ y_i \ f\dot{x}_i \ f y_i]) + \dots \\ GP([\dot{x}_i \ y_i \ f\dot{x}_i \ f y_i]). \end{aligned} \quad (13)$$

This is solved using a root solving program *fsolve* in MATLAB.

5. HARDWARE PLATFORM

5.1 Overview

The hardware platform utilized in this work is shown in Fig. 3. The setup consists of a custom, two degree-of-freedom, portable rotating boom with a boom-arm-attached monoped leg, linked via EtherCAT network to a real-time linux server. The boom-arm attached monoped leg is an, electrically driven, asymmetric four-link mechanism that acts to convert two opposing torques into linear force and torque. The drive section for the monoped leg mechanism consist of oppositely-mounted T-motor U10plus brushless DC motors with attached custom 14bit magnetic encoder modules, build around the IC Haus MHM encoder IC. The closed torque control loops of the monoped drive section are controlled within dual onboard Elmo Motion Control 80V/80A servo drives. Mid-level interface is provided by a real time Linux master terminal via a synchronous EtherCAT network operating at a 4kHz.

5.2 Boom

The custom portable rotating boom is a platform that facilitates the controlled motion and observation of an experimental package while constraining unwanted degrees of freedom. The rotating boom utilized in this work has two unconstrained degrees



FIGURE 3: HARDWARE SETUP

of freedom: about the vertical axis (oriented parallel to gravity) and about a rotating horizontal axis (oriented perpendicular to the vertical axis). The boom rotating axes are equipped with mechanical encoders which offer a resolution of 0.072 and 0.022 degrees in the horizontal and rotating vertical planes, respectively. Boom encoder position computation and tracking are performed by on-board Beckhoff EtherCAT terminals, operating synchronously with the EtherCAT-network-linked Elmo servo drives. The ultra high modulus carbon fiber boom arm, which couples the monoped leg mechanism utilized in this work to the body of the portable rotating boom, is adjustable in length from 1.85m to 2.75m. On-board EtherCAT communication, high-voltage low-current supply (120Vac), and low-voltage high-current supply (36Vdc) are provided via an electro-mechanical slip ring, allowing for infinite and unobstructed rotation of the boom arm about the vertical axis of the boom.

5.3 Control terminal and EtherCAT communication

High level command and control for both the monoped leg mechanism and the rotating boom platform are performed on a Dell Optiplex 9020 with an Intel i7 processor, operating Ubuntu Linux with a real-time kernel. The EtherCAT network interface adapter is 10Gtek network interface card (NIC) with an Intel 82574L chipset. EtherCAT network management and EtherCAT network interfacing between the hardware and main state machine are provided by the Acontis Technologies EtherCAT master stack. The Acontis Technologies EtherCAT master stack is a collection of libraries that provide link layer management between an NIC and user software, allowing off-the-shelf network and computer hardware to be utilized as an EtherCAT master terminal.

Through the use of the EC master stack libraries, the state machine is able to synchronously access the torque, velocity and position of each U10plus drive motor, as well as the angular position of both the horizontal and vertical boom-mounted encoders. As the torque control loop executes directly on the Elmo servo drives, the only data passed to the servo drives, synchronously,

by the state machine is the reference torque command for each servo drive. Due to the utilized real-time Linux kernel, the user software (including the state machine) is able to execute significantly faster than on a non-rt system and, due to the use of a multithreaded program architecture, numerous computational operations are able to be performed between each successive EtherCAT network cycle without blocking the operation of the EtherCAT masterstack, providing for the ability to handle substantial computational problems in real time while the hardware is in operation without blocking EtherCAT communication.

6. RESULTS

The simulation performed using MATLAB. We used the function *ode113* to integrate the equations of motion given by Eqns 4, 5, and 8. The integrator has an in-built event detection that helps to detect transitions. The Poincaré map (Eqn. 10) is obtained by integrating the equations from one apex to the next.

We generated a reference profile for \dot{x}_{des} and \dot{y}_{des} as shown in Fig. 4 (solid blue line). Using the simulation-based Poincaré map (Eqn. 10) and the MATLAB function *fsolve*, we generated a controller as a lookup table. This was then executed on the robot. This is shown in Fig. 4 (dash-dotted red line). It can be seen that there is a significant tracking error for both, \dot{x} and \dot{y} .

Next, we used the experimental data to fit the error in the Poincaré map estimation using a Gaussian Process Regression (Eqn. 12). We used the MATLAB function *fitrgp* with a constant basis function and matern52 kernel. Then we used the corrected Poincaré map to compute the control (see Eqn. 13). This was tested on the hardware and the resulting data was used to improve the estimate of the Gaussian Process. This process was repeated 3 more times, for a total of 4 training iterations and 4 hardware deployments. The 4th iteration led to acceptable results. The error in tracking using the 4th iteration GP is shown in Fig. 4 (dashed black line). It can be seen that the error has substantially reduced compared to the simulation-based controller. Thus the approach is shown to successfully improve the sim-to-real transfer

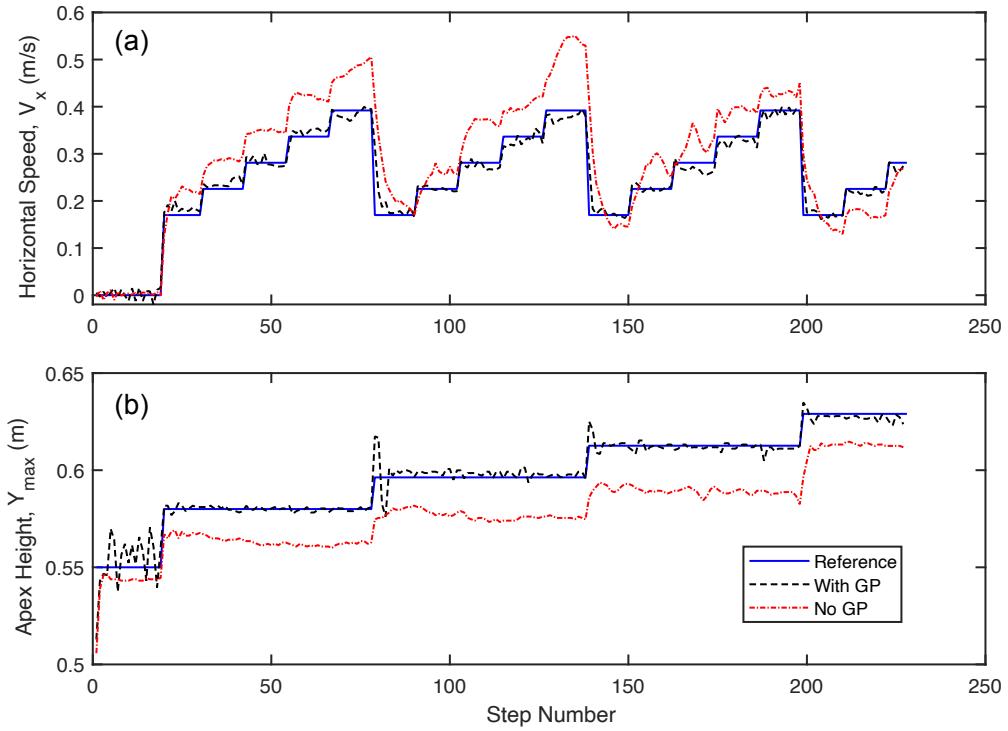


FIGURE 4: COMPARING CONTROLLERS ON HARDWARE. CONTROLLER BASED ON SIMULATION-BASED POINCARÉ MAP (NO GP) VERSUS CONTROLLER BASED ON LEARNT POINCARÉ MAP (WITH GP).

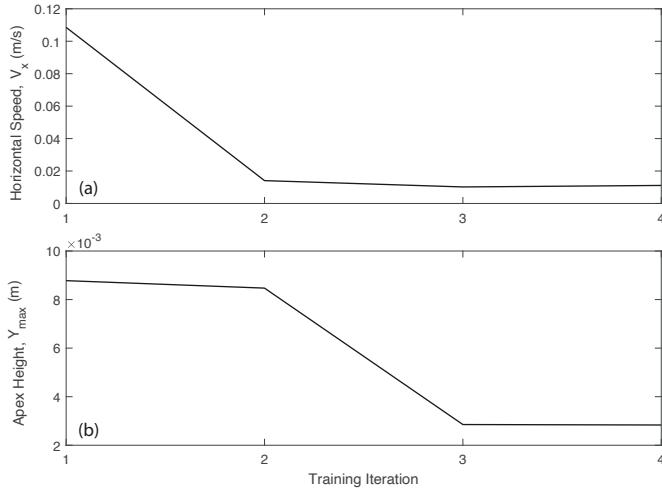


FIGURE 5: AVERAGE MEAN SQUARED ERROR OF THE POINCARÉ MAP CORRECTION VS ITERATION NUMBER.

using just 4 training iterations. A video comparing the controllers with and without GPR is here: <tiny.cc/detc2023>.

Figure 5 shows the progress in the training, the average mean squared error (MSE) vs the iteration number for both the variables. It can be seen that the MSE converges in a mere 1 iteration for velocity and about 3 iterations for apex height. Figure 6 and 7 gives more details about the mean/confidence interval for the converged iteration (top) and the standard deviation over the first and last converged iteration (bottom). From the top plots it can be seen that that 95% confidence interval is narrow for both variables. From bottom figures it can be seen that initially (iteration 1), the standard deviation is the largest when there is a step change in the reference. However, with more training (iteration 4), the standard deviation at the step change decreases. In particular, the maximum deviation in iteration 1 is about 0.03 m/s which decreases to 0.02 m/s for velocity and from 0.015 m to under 0.01 m for apex height.

7. DISCUSSION

We presented a technique for sim-to-real transfer tracking control of hopping robot. The technique consists of using a physics-based model to numerically obtain the Poincaré map. Then the map is iteratively improved by designing a controller with the Poincaré map, collecting data from hardware, and fitting a correction to the Poincaré map using Gaussian Process Regression. This process is repeated till the Gaussian Process Regression model converges, i.e., there is no improvement on the mean square error. The converged estimate of the Poincaré

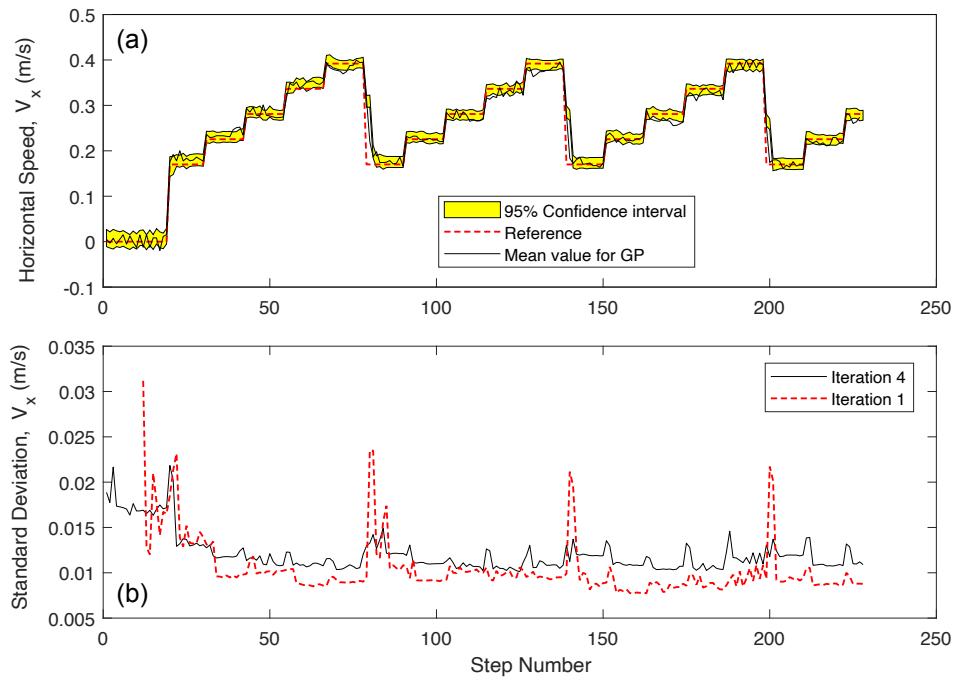


FIGURE 6: TRAINING OF POINCARÉ MAP FOR HORIZONTAL VELOCITY, (A) MEAN AND CONFIDENCE INTERVALS FOR THE TRAINING, (B) STANDARD DEVIATION FOR THE FIRST AND FOURTH (LAST) ITERATION.

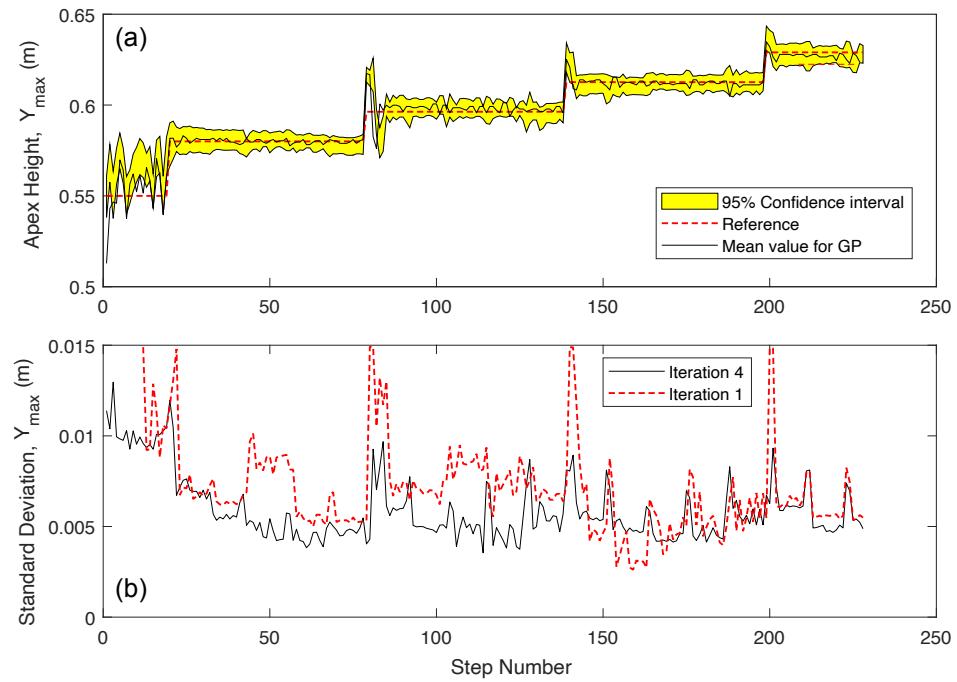


FIGURE 7: TRAINING OF POINCARÉ MAP FOR APEX HEIGHT, (A) MEAN AND CONFIDENCE INTERVALS FOR THE TRAINING, (B) STANDARD DEVIATION FOR THE FIRST AND FOURTH (LAST) ITERATION.

map provides accurate tracking compared against the simulation obtained Poincaré map.

The differences between the model and simulation were due to boom dynamics (boom deflection during motion), ground variance, and mismatch in mass, center of mass, and inertia to some extent. Some of these difference are best modeled as a bias (e.g., mass, inertia, center of mass), but others such as boom dynamics and ground variance are best modeled as a bias and an uncertainty. A Gaussian Process Regression (GPR) can effectively model both, bias and uncertainty as a mean and standard deviation respectively. The mis-modeled dynamics shows up as a steady state error (a bias) in the tracking using the simulation-based controller as shown in Fig. 4 (dash-dotted red line). The boom dynamics have the biggest effect when the hopper changes the reference speed and velocity. This can be seen in the relatively large standard deviations during these step changes (see Fig. 6 and 7).

Unlike past approaches which learn the control policy in hardware, we learn the model mismatch. The advantage of learning the model mismatch (as opposed to the control) is that the resulting model can be used in a novel scenario without taking more experimental data. Also, past approaches learn the GPR for model or control policy in the time domain. Since the GPR scales as N^3 where N is the number of data points, time domain models are challenging to scale. In contrast, we learn the GPR model at the step-level where N is step number, our GPR model can model more data sets than using time-based parameterization.

There are several limitations of our approach. Unlike past approaches that use neural networks for modeling, GPR provide local approximations. However, it is possible to use neural networks or polynomials to approximate the Poincaré map [21]. For the hopping system, there are only 4 independent variables for the GPR (see Eqn. 12) which helps in scaling the approach. However, this may not be true for high-dimensional systems. In such cases, one attractive option is to use a feature space to map the Poincaré variable to a low-dimensional space for GPR (e.g., see [12]). Finally, it is important to choose free parameters from the low-level control that have appreciable effect on the states on Poincaré map. There is, so far, no systematic way to find such parameters but human intuition and trial-and-error [22, 23].

8. CONCLUSIONS AND FUTURE WORK

We conclude that Gaussian Process Regression (GPR) can scale for sim-to-real transfer when used to model the correction to the Poincaré map rather than the low-level physics. Our results show that the convergence is fast (about 4 iterations). Thus, GPR is promising method for sim-to-real transfer.

The future work will investigate GPR in situation where there is added uncertainty (e.g., manipulate the ground profile) and the use of model-based control framework for tracking of reference signals.

ACKNOWLEDGMENTS

The authors would like to thank Acontis Technologies for the use of and assistance with their EtherCAT Master Stack software. The work is funded by US National Science Foundation through grant 2128568.

Any opinions, findings, and conclusion or recommendations expressed in this material are those of the authors and do not necessarily reflect the view of the National Science Foundation.

REFERENCES

- [1] Raibert, M.H. *Legged robots that balance*. MIT press Cambridge, MA (1986).
- [2] Strogatz, S.H. *Nonlinear dynamics and chaos*. Addison-Wesley Reading (1994).
- [3] Dadashzadeh, Behnam, Mahjoob, MJ, Nikkhah Bahrami, M and Macnab, Chris. “Stable active running of a planar biped robot using Poincare map control.” *Advanced Robotics* Vol. 28 No. 4 (2014): pp. 231–244.
- [4] Cheng, M-Y and Lin, C-S. “Measurement of robustness for biped locomotion using a linearized Poincaré map.” *Robotica* Vol. 14 No. 3 (1996): pp. 253–259.
- [5] Bhounsule, Pranav A, Zamani, Ali, Krause, Jeremy, Farra, Steven and Pusey, Jason. “Control policies for a large region of attraction for dynamically balancing legged robots: a sampling-based approach.” *Robotica* Vol. 39 No. 1 (2021): pp. 107–122.
- [6] Mordatch, Igor, Lowrey, Kendall and Todorov, Emanuel. “Ensemble-cio: Full-body dynamic motion planning that transfers to physical humanoids.” *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*: pp. 5307–5314. 2015. IEEE.
- [7] Tan, Jie, Zhang, Tingnan, Coumans, Erwin, Iscen, Atil, Bai, Yunfei, Hafner, Danijar, Bohez, Steven and Vanhoucke, Vincent. “Sim-to-real: Learning agile locomotion for quadruped robots.” *arXiv preprint arXiv:1804.10332* (2018).
- [8] Peng, Xue Bin, Coumans, Erwin, Zhang, Tingnan, Lee, Tsang-Wei, Tan, Jie and Levine, Sergey. “Learning agile robotic locomotion skills by imitating animals.” *arXiv preprint arXiv:2004.00784* (2020).
- [9] Bellegarda, Guillaume and Nguyen, Quan. “Robust high-speed running for quadruped robots via deep reinforcement learning.” *arXiv preprint arXiv:2103.06484* (2021).
- [10] Hwangbo, Jemin, Lee, Joonho, Dosovitskiy, Alexey, Belluscio, Dario, Tsounis, Vassilios, Koltun, Vladlen and Hutter, Marco. “Learning agile and dynamic motor skills for legged robots.” *Science Robotics* Vol. 4 No. 26 (2019): p. eaau5872.
- [11] Calandra, Roberto, Seyfarth, André, Peters, Jan and Deisenroth, Marc Peter. “An experimental comparison of Bayesian optimization for bipedal locomotion.” *2014 IEEE international conference on robotics and automation (ICRA)*: pp. 1951–1958. 2014. IEEE.
- [12] Rai, Akshara, Antonova, Rika, Song, Seungmoon, Martin, William, Geyer, Hartmut and Atkeson, Christopher. “Bayesian optimization using domain knowledge on the ATRIAS biped.” *2018 IEEE International Conference on Robotics and Automation (ICRA)*: pp. 1771–1778. 2018. IEEE.
- [13] Yu, Wenhao, Kumar, Visak CV, Turk, Greg and Liu, C Karen. “Sim-to-real transfer for biped locomotion.” *2019*

ieee/rsj international conference on intelligent robots and systems (iros): pp. 3503–3510. 2019. IEEE.

[14] Snelson, Edward and Ghahramani, Zoubin. “Local and global sparse Gaussian process approximations.” *Artificial Intelligence and Statistics*: pp. 524–531. 2007. PMLR.

[15] Torres, Daniel, Hernandez-Hinojosa, Ernesto and Bhounsule, Pranav A. “Control of a Bipedal Walking Using Partial Feedback Linearization and Gaussian Process Regression-Based of the Step-to-Step Map.” *ASME-International Design Engineering & Technical Conference*. 2023.

[16] Ha, Sehoon and Yamane, Katsu. “Reducing hardware experiments for model learning and policy optimization.” *2015 IEEE International Conference on Robotics and Automation (ICRA)*: pp. 2620–2626. 2015. IEEE.

[17] Ko, Jonathan, Klein, Daniel J, Fox, Dieter and Haehnel, Dirk. “Gaussian processes and reinforcement learning for identification and control of an autonomous blimp.” *Proceedings 2007 ieee international conference on robotics and automation*: pp. 742–747. 2007. IEEE.

[18] Berkenkamp, Felix, Schoellig, Angela P and Krause, Andreas. “Safe controller optimization for quadrotors with Gaussian processes.” *2016 IEEE international conference on robotics and automation (ICRA)*: pp. 491–496. 2016. IEEE.

[19] Ostafew, Chris J, Schoellig, Angela P and Barfoot, Timo-
thy D. “Learning-based nonlinear model predictive control to improve vision-based mobile robot path-tracking in challenging outdoor environments.” *2014 IEEE International Conference on Robotics and Automation (ICRA)*: pp. 4029–4036. 2014. IEEE.

[20] Williams, Christopher KI and Rasmussen, Carl Edward. *Gaussian processes for machine learning*. Vol. 2. MIT press Cambridge, MA (2006).

[21] Bhounsule, Pranav A, Kim, Myunghee and Alaeddini, Adel. “Approximation of the Step-to-step Dynamics Enables Computationally Efficient and Fast Optimal Control of Legged Robots.” *ASME 2020 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*. 2020. American Society of Mechanical Engineers.

[22] Bhounsule, Pranav A, Ruina, Andy and Stiesberg, Gregg. “Discrete-decision continuous-actuation control: balance of an inverted pendulum and pumping a pendulum swing.” *Journal of Dynamic Systems, Measurement, and Control* Vol. 137 No. 5 (2015): p. 051012.

[23] Bhounsule, Pranav A., Cortell, Jason, Grewal, Anoop, Hendriksen, Bram, KarsSEN, J.G. Daniël, Paul, Chandana and Ruina, Andy. “Low-bandwidth reflex-based control for lower power walking: 65 km on a single battery charge.” *International Journal of Robotics Research* (2014).