

Auto-Tuning of Controller and Online Trajectory Planner for Legged Robots

Schperberg, Alexander; Di Cairano, Stefano; Menner, Marcel

TR2022-085 August 06, 2022

Abstract

This letter presents an approach for auto-tuning feedback controllers and online trajectory planners to achieve robust locomotion of a legged robot. The auto-tuning approach uses an Unscented Kalman Filter (UKF) formulation, which adapts/calibrates control parameters online using a recursive implementation. In particular, this letter shows how to use the auto-tuning approach to calibrate cost function weights of a Model Predictive Control (MPC) stance controller and feedback gains of a swing controller for a quadruped robot. Furthermore, this letter extends the auto-tuning approach to calibrating parameters of an online trajectory planner, where the height of a swing leg and the robot's walking speed are optimized, while minimizing its energy consumption and foot slippage. This allows us to generate stable reference trajectories online and in real time. Results using a high-fidelity Unitree A1 robot simulator in Gazebo provided by the robot manufacturer show the advantages of using auto-tuning for calibrating feedback controllers and for computing reference trajectories online for reduced development time and improved tracking performance

IEEE Robotics and Automation Letters 2022

Auto-Tuning of Controller and Online Trajectory Planner for Legged Robots

Alexander Schperberg, Stefano Di Cairano, and Marcel Menner

Abstract—This letter presents an approach for auto-tuning feedback controllers and online trajectory planners to achieve robust locomotion of a legged robot. The auto-tuning approach uses an Unscented Kalman Filter (UKF) formulation, which adapts/calibrates control parameters online using a recursive implementation. In particular, this letter shows how to use the auto-tuning approach to calibrate cost function weights of a Model Predictive Control (MPC) stance controller and feedback gains of a swing controller for a quadruped robot. Furthermore, this letter extends the auto-tuning approach to calibrating parameters of an online trajectory planner, where the height of a swing leg and the robot’s walking speed are optimized, while minimizing its energy consumption and foot slippage. This allows us to generate stable reference trajectories online and in real time. Results using a high-fidelity Unitree A1 robot simulator in Gazebo provided by the robot manufacturer show the advantages of using auto-tuning for calibrating feedback controllers and for computing reference trajectories online for reduced development time and improved tracking performance.

Index Terms—Legged Robots; Machine Learning for Robot Control; Calibration and Identification; Integrated Planning and Learning; Probability and Statistical Methods

I. INTRODUCTION

PLANNING and control of legged systems is challenging due to the tight coupling of reaction forces generated by foot contacts with the environment and the motion of the robot’s base. This problem is further complicated as legged systems for real-world deployment are expected to traverse highly unstructured environments such as debris, obstacles, and rough terrain. Traditional control designs such as those employing Raibert or heuristic controllers [1] can only guarantee physical feasibility on smooth or flat surfaces. Optimization-based approaches often use a cost function to account for kinematic, dynamic, and friction-cone constraints. E.g., the Linear Inverted Pendulum approach optimizes over the Center of Mass (CoM) position, where the footsteps must be specified *a priori* to satisfy the Zero-Moment Point [2]. This approach has some drawbacks because pre-defining the footholds’ locations on the ground may restrict the robot’s range of achievable motion. Further, whole-body controllers that do not require pre-defined footholds, e.g., in [3], may be sub-optimal as they only consider the current joint torques.

Manuscript received: February 24, 2022; Revised: May 21, 2022; Accepted: June 05, 2022.

This paper was recommended for publication by Editor Abderrahmane Kheddar upon evaluation of the Associate Editor and Reviewers’ comments. All authors are with Mitsubishi Electric Research Laboratories (MERL), Cambridge, MA, 02139, USA (e-mail: aschperberg28@g.ucla.edu; dicairano@ieee.org; menner@ieee.org).

Digital Object Identifier (DOI): see top of this page.

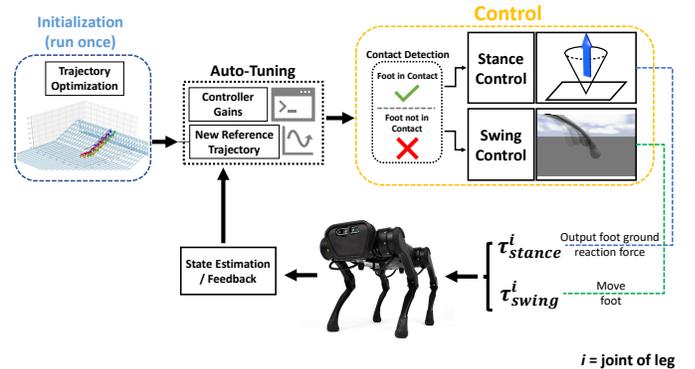


Fig. 1: **Auto-tuning framework.** We initialize the reference trajectory for the first few footsteps and CoM base using trajectory optimization. Then, auto-tuning is applied to estimate new reference trajectories which account for desired physical parameters. Further, auto-tuning is applied on cost function weights for a stance controller (MPC) and feedback gains of a swing controller (PD).

On the other hand, trajectory Optimization (TO) is often used for predictive planning of feasible motions for legged systems [4]–[7]. TO methods typically employ some form of centroidal dynamics, optimizing over contact forces [4]. However, TO is computationally expensive and often not suitable for online re-planning [8]. Further, open-loop planners are susceptible to changes in terrain.

Due to the inverse relationship between the complexity of the dynamic model used for locomotion and the computation time required for an optimization solver, a popular approach is to employ Reinforcement Learning (RL) [9]–[14]. E.g., in [9], a stochastic policy is realized with neural networks to simulate a foothold and base motion controller using Trust-Region Policy Optimization. This is done by considering a reward function that consists of several physical parameters such as the error between actual and reference footstep position, penalizing foot slippage or large swing-leg velocities, or sudden changes in base orientation. An RL approach has also been successfully demonstrated in [10], which learns a neural network that acts only on a stream of proprioceptive signals for locomotion on uneven terrain.

Although tremendous progress has been made in the RL literature to achieve legged robot locomotion, this approach is highly complex and requires significant amounts of data. For example, in [10], a two-stage training process is required that involves a teacher/student policy in addition to particle-

filtering, which maintains certain terrain parameters used to classify what is and is not traversable during training. In [9], the reward function and corresponding weighing parameters required intricate fine-tuning and assumes the user has extensive computational resources available during training. Another work can be found in [15], which uses RL to create adaptive trajectories based on following a reference trajectory calculated offline first through TO. Similar to our work, they also use step clearance and slippage as objectives to their calibration. However, different from ours, while we also use TO methods, we only rely on the reference trajectory from TO for the first few footsteps, while future footsteps are planned online. Overall, the complexity of formulating the correct reward function in RL and the effort required for tuning hyperparameters may cause significant delay in development time. This also may make RL difficult to reproduce or generalize to a wide-variety of robotic platforms without significant amounts of tuning.

This letter proposes to use an alternative method, see Fig. 1, to the above approaches for achieving robust locomotion. In [16], [17], a Kalman filter technique was used to estimate control parameters with a training objective that evaluates the performance of a closed-loop system online using a recursive implementation. This approach was successfully used to automatically tune a variety of control architectures such as the weights of a neural network, cost function weights of an LQR, or gains of a PID controller to achieve stable motion of an autonomous vehicle. Although auto-tuning methods have been applied before to tune various controller gains, typically using some form of Bayesian optimization (BO) [18], [19], the applied method differs as we do not require a trial-and-error implementation, do not need a surrogate function, and can handle disturbances due to the recursive nature of the applied method [16], [17]. One of the issues with BO is that their performance decreases in higher dimensions [20]. In [21], this problem was overcome by including domain knowledge into BO for tuning walking controllers of humanoid robots. In [22], the parameters of running and jumping motions are tuned through a policy search of generic sets of motion primitives and their cost functions. Both [21], [22] require a trial-and-error implementation which requires extensive training using simulators before application to hardware. Our method however, is able to calibrate control parameters on a single run without relying on trial-and-error.

Additionally, while other algorithms have been proposed to tune controller gains [23], [24], these algorithms have never been applied (or evaluated) to directly tune reference trajectories without having to include a predefined set of reference trajectories [25].

This letter employs the method of auto-tuning initially proposed in [16], [17] on a legged system using an unscented Kalman filter (UKF) implementation and apply it not only to controller gains, but directly on physically-meaningful parameters such as ‘step clearance’, ‘forward progress’, ‘energy efficiency’, and ‘dynamic and kinematic gait stability’, see Fig. 2. This is achieved by online tuning the robot’s future reference trajectories which uses a training objective to satisfy user-defined specifications or requirements. We demonstrate

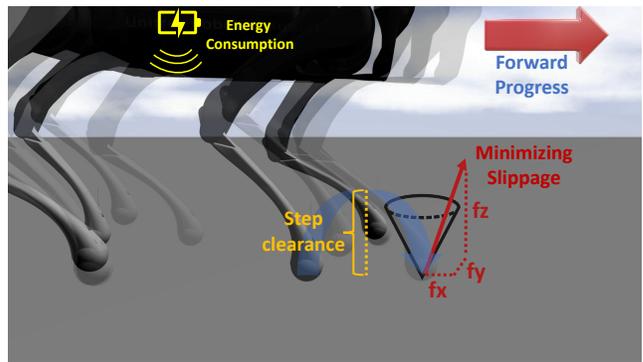


Fig. 2: **Auto-Tuning reference trajectories in Gazebo.** The robot follows a reference trajectory that is being tuned by the auto-tuning formulation. The figure shows the robot following a desired velocity (forward progress), a desired foot height (step clearance), minimizing foot slippage (ground reaction forces), and minimizing energy consumption.

the method on a quadruped robot, the Unitree A1, in a high-fidelity physics simulation provided by the robot manufacturer. We demonstrate walking on both even and uneven terrain, and present several test-cases such as tuning reference trajectories that satisfy the physical parameters mentioned previously, and controller weights of a swing and stance controller. The presented method for automating the controller calibration and the trajectory planner can save development/deployment time as manually tuning control parameters can be tedious and time consuming, facilitates in abstracting the tuning problem into physically meaningful parameters, and enables greater autonomy of robotic systems.

II. PRELIMINARIES

A. Notation

Given two integer indices n, m with $m < n$ and $\mathbf{x}_i \in \mathbf{R}^{n_x}$, we define $\mathbf{x}_{m|n} \in \mathbf{R}^{n_x(n-m+1)}$ as the vectorized sequence that comprises \mathbf{x}_i from $i = m$ through $i = n$,

$$\mathbf{x}_{m|n} := \begin{bmatrix} \mathbf{x}_m \\ \vdots \\ \mathbf{x}_n \end{bmatrix}.$$

We define $\|\mathbf{x}\|_{\Sigma} := \mathbf{x}^T \Sigma \mathbf{x}$ and $\mathcal{N}(\boldsymbol{\mu}, \Sigma)$ as the Gaussian distribution with mean vector $\boldsymbol{\mu}$ and covariance matrix Σ . The notation $\mathbf{x} \sim \mathcal{N}(\boldsymbol{\mu}, \Sigma)$ means \mathbf{x} sampled from $\mathcal{N}(\boldsymbol{\mu}, \Sigma)$. Further, $\text{diag}(\boldsymbol{\lambda}) \in \mathbf{R}^{n_\lambda \times n_\lambda}$ is a matrix, whose diagonal entries are the entries of a vector $\boldsymbol{\lambda} \in \mathbf{R}^{n_\lambda}$.

B. Unscented Kalman Filter for Control Parameter Tuning

This letter applies the auto-tuning method proposed in [16], [17] to calibrate control parameters of the Model Predictive Control (MPC) stance controller, the PD swing controller, as well as the reference trajectory. The auto-tuning method is based on an UKF, which ‘‘estimates’’ the optimal control parameters measured with respect to a training objective. The UKF uses deterministic samples (called sigma points) around the mean, which are propagated and used to update the mean

and covariance estimates [26]. Further, it uses a model of the system dynamics in order to obtain evaluations of the sigma points, which are then used to update the control parameters.

The method is model-based, i.e., it uses a model of a dynamical system (denoted by $\text{dyn}(\mathbf{x}_k, \mathbf{u}_k)$),

$$\mathbf{x}_{k+1} = \text{dyn}(\mathbf{x}_k, \mathbf{u}_k) + \mathbf{w}_k \quad (1)$$

with the state $\mathbf{x}_k \in \mathbb{R}^n$, the control input $\mathbf{u}_k \in \mathbb{R}^m$, and process noise or model mismatch \mathbf{w}_k at time k . The method calibrates control parameters, $\boldsymbol{\theta} \in \mathbb{R}^L$, of a generic controller, $\mathbf{u}_k = \kappa_\theta(\mathbf{x}_k)$ to minimize the training objective

$$\|\mathbf{y}_k - \mathbf{h}(\boldsymbol{\theta}_k)\|_{\mathbf{C}_y^{-1}} \quad (2)$$

with a positive definite \mathbf{C}_y , desired nominal values \mathbf{y}_k , and specification function $\mathbf{h}(\boldsymbol{\theta})$, where

$$\mathbf{h}(\boldsymbol{\theta}) := \mathbf{r}(\mathbf{x}_{k-N}, \mathbf{x}_{k-N+1}, \dots, \mathbf{x}_k, \mathbf{u}_{k-N}, \dots, \mathbf{u}_{k-1}),$$

i.e., $\mathbf{y}_k = \mathbf{h}(\boldsymbol{\theta}_k)$ when the dynamical system satisfies all the specifications in the training objective, exactly.

For the legged robot, the control parameters are tuned during operation (episodically, from time step $k-N$ to k) according to

$$\boldsymbol{\theta}_k = \boldsymbol{\theta}_{k-N} + \Delta\boldsymbol{\theta}_k, \quad (3)$$

where the update $\Delta\boldsymbol{\theta}_k$ is computed based on sensor measurements, \mathbf{x}_k , \mathbf{u}_k , and the training objective in (2).

The idea of the auto-tuning method is to treat the control parameter adaptation problem as an optimal estimation problem with prior distributions $\Delta\boldsymbol{\theta}_k^{\text{prior}} \sim \mathcal{N}(\mathbf{0}, \mathbf{C}_\theta)$ and $\mathbf{y}_k \sim \mathcal{N}(\mathbf{h}(\boldsymbol{\theta}_k), \mathbf{C}_y)$. Thus, the parameter tuning law in (3) results from the corresponding posterior distribution,

$$\Delta\boldsymbol{\theta}_k = \mathbf{K}_k (\mathbf{y}_k - \hat{\mathbf{h}}_k) \quad (4a)$$

with the Kalman gain, \mathbf{K}_k , computed as

$$\hat{\boldsymbol{\theta}}_k = \sum_{j=0}^{2L} v_j^a \boldsymbol{\theta}_k^{\text{sp},j} \quad (4b)$$

$$\hat{\mathbf{h}}_k = \sum_{j=0}^{2L} v_j^a \mathbf{h}_k^{\text{sp},j} \quad (4c)$$

$$\mathbf{h}_k^{\text{sp},j} = \mathbf{h}(\boldsymbol{\theta}_k^{\text{sp},j}) \quad (4d)$$

$$\mathbf{S}_k = \mathbf{C}_v + \sum_{j=0}^{2L} v_j^c (\mathbf{h}_k^{\text{sp},j} - \hat{\mathbf{h}}_k)(\mathbf{h}_k^{\text{sp},j} - \hat{\mathbf{h}}_k)^\top \quad (4e)$$

$$\mathbf{Z}_k = \sum_{j=0}^{2L} v_j^c (\boldsymbol{\theta}_k^{\text{sp},j} - \hat{\boldsymbol{\theta}}_k)(\mathbf{h}_k^{\text{sp},j} - \hat{\mathbf{h}}_k)^\top \quad (4f)$$

$$\mathbf{K}_k = \mathbf{Z}_k \mathbf{S}_k^{-1} \quad (4g)$$

and the posterior covariance, which is used to generate the sigma points, computed as

$$\mathbf{P}_{k|k-N} = \mathbf{C}_\theta + \sum_{j=0}^{2L} v_j^c (\boldsymbol{\theta}_k^{\text{sp},j} - \hat{\boldsymbol{\theta}}_k)(\boldsymbol{\theta}_k^{\text{sp},j} - \hat{\boldsymbol{\theta}}_k)^\top \quad (4h)$$

$$\mathbf{P}_{k|k} = \mathbf{P}_{k|k-N} - \mathbf{K}_k \mathbf{S}_k \mathbf{K}_k^\top, \quad (4i)$$

where $\boldsymbol{\theta}_k^{\text{sp},j}$ is the j th sigma point, v_j^c and v_j^a denote weights associated with the sigma points, \mathbf{Z}_k is the cross-covariance matrix, \mathbf{S}_k is the innovation covariance, and $\mathbf{P}_{k|k}$ is the estimate covariance. Hence, the UKF implementation uses $2L+1$ sigma points. The reader is referred to [16], [17] for more details.

Remark 1: For each sigma point evaluation (4d), the system dynamics (1) is simulated, where the model mismatch, \mathbf{w}_k , is calculated using measured data, \mathbf{x}_{k-N} through \mathbf{x}_k .

Remark 2: We choose the weights $v_0^a = v_0^c = 0$, $v_i^a = v_i^c = (1 - v_0^a)/(2L)$ and the sigma points as $\boldsymbol{\theta}_k^{\text{sp},0} = \boldsymbol{\theta}_k$, $\boldsymbol{\theta}_k^{\text{sp},j} = \boldsymbol{\theta}_k + \sqrt{L/(1 - v_0^a)} \boldsymbol{\Gamma}^j$ for $j = 1, \dots, L$, and $\boldsymbol{\theta}_k^{\text{sp},j} = \boldsymbol{\theta}_k - \sqrt{L/(1 - v_0^a)} \boldsymbol{\Gamma}^j$ for $j = L+1, \dots, 2L$ with $\boldsymbol{\Gamma}^j$ being the j th column of $\boldsymbol{\Gamma}$ and $\mathbf{P}_{k-N|k-N} = \boldsymbol{\Gamma} \boldsymbol{\Gamma}^\top$, i.e., $\boldsymbol{\Gamma}$ is calculated using the Cholesky decomposition.

III. ROBOT CONTROL ARCHITECTURE

A. Swing Controller

The swing controller in this letter is similar to [27], which is used to compute the torque for each foot i for all three joints of the robot as (time step k omitted for simplicity)

$$\boldsymbol{\tau}_i = \mathbf{J}_i^\top [\mathbf{K}_p (\mathbf{p}_{i,\text{ref}}^b - \mathbf{p}_i^b) + \mathbf{K}_d (\mathbf{v}_{i,\text{ref}}^b - \mathbf{v}_i^b)] + \mathbf{J}_i \boldsymbol{\Lambda}_i (\mathbf{a}_{i,\text{ref}}^b - \mathbf{J}_i^\top \dot{\mathbf{q}}_i) + \mathbf{V}_i \dot{\mathbf{q}}_i + \mathbf{G}_i \quad (5)$$

where $\boldsymbol{\tau}_i \in \mathbb{R}^3$ is the joint torque, $\mathbf{q}_i \in \mathbb{R}^3$ and $\dot{\mathbf{q}}_i \in \mathbb{R}^3$ are the current joint position and velocity of foot i , $\mathbf{J}_i \in \mathbb{R}^{3 \times 3}$ is the foot Jacobian, \mathbf{K}_p and \mathbf{K}_d are the proportional and derivative (PD) gain matrices (3×3 diagonal positive semi-definite), $\mathbf{p}_{i,\text{ref}}^b \in \mathbb{R}^3$ and $\mathbf{p}_i^b \in \mathbb{R}^3$ are the reference and current footstep positions in the body frame, $\mathbf{v}_{i,\text{ref}}^b \in \mathbb{R}^3$ and $\mathbf{v}_i^b \in \mathbb{R}^3$ are the reference and current footstep velocities in the body frame, $\mathbf{a}_{i,\text{ref}}^b \in \mathbb{R}^3$ is the reference footstep acceleration in the body frame, $\mathbf{V}_i \in \mathbb{R}^3$ is the torque due to the coriolis and centrifugal forces, $\mathbf{G}_i \in \mathbb{R}^3$ is the torque due to gravity, and $\boldsymbol{\Lambda}_i \in \mathbb{R}^{3 \times 3}$ is the operational mass matrix.

B. Stance Controller

The stance MPC calculates ground reaction forces \mathbf{f}_i for all feet i and is formulated as in [27],

$$\min_{\mathbf{x}, \mathbf{f}} \sum_{k=0}^{N_{\text{MPC}}} \|\mathbf{x}_k - \mathbf{x}_{k,\text{ref}}\|_{\mathbf{Q}} + \|\mathbf{f}_k\|_{\mathbf{R}} \quad (6a)$$

$$\begin{aligned} & f_{k,\text{min}} \leq f_{k,z} \leq f_{k,\text{max}} \\ & -\mu f_{k,z} \leq \pm f_{k,x} \leq \mu f_{k,z} \\ & -\mu f_{k,z} \leq \pm f_{k,y} \leq \mu f_{k,z} \end{aligned} \quad (6b)$$

$$\begin{aligned} \mathbf{x}_{k+1} &= \mathbf{A} \mathbf{x}_k + \mathbf{B} \mathbf{f}_k \\ \mathbf{D}_k \mathbf{f}_k &= 0 \end{aligned}$$

with the state

$$\mathbf{x}_k = \begin{bmatrix} \boldsymbol{\Theta}_k \\ \mathbf{r}_k \\ \boldsymbol{\omega}_k \\ \mathbf{v}_k \end{bmatrix}, \quad (6c)$$

where $\boldsymbol{\Theta}$ is the robot's orientation, \mathbf{r} is the CoM base position, $\boldsymbol{\omega}$ is the angular velocity, \mathbf{v} is the linear velocity, μ is the friction coefficient, \mathbf{A} and \mathbf{B} are the dynamic matrices for state propagation, \mathbf{D}_k is a force selection matrix (selecting forces that are not in contact to be equal to zero), and $\mathbf{Q} \in \mathbb{R}^{12 \times 12}$ and $\mathbf{R} \in \mathbb{R}^{12 \times 12}$ are diagonal positive semi-definite cost matrices, see [27] for further details. The joint torques, which are the input to the torque-controlled motors, are obtained using the forces \mathbf{f}_k resulting from (6) as

$$\boldsymbol{\tau}_{i,k} = \mathbf{J}_{i,k}^\top \mathbf{R}_{b,i,k}^{w,\top} (-\mathbf{f}_{i,k}), \quad (7)$$

where $\mathbf{f}_{i,k} \in \mathbb{R}^3$ is the force vector associated with leg i as subset of \mathbf{f}_k , and $\mathbf{R}_{b,i,k}^w$ is the rotation matrix from world to body frame of leg i at time step k .

C. Problem Definition and Contributions

In this letter, we implement model-based controllers and auto-tune their control parameters. Model-based control offers the advantages that the physics of a dynamical system are utilized to make informed decisions, less data are needed to improve the robot's performance, safety guarantees may be provided, and that the computational structure remains fixed. This control philosophy stands in contrast to black-box controllers and reduces the amount of required tuning to a few selected control parameters.

Problem A: Auto-tune the gains of the swing controller in (5), i.e., \mathbf{K}_p , \mathbf{K}_d . Section IV-A addresses this problem by adjusting the controller calibration method in [16], [17] to a swing controller of a legged robot. This is achieved by defining a training objective suited for a swing controller as well as deriving a motion model suited to model the swing motion of a legged robot.

Problem B: Auto-tune the cost function weights of the stance controller in (6), i.e., \mathbf{Q} , \mathbf{R} . Section IV-B addresses this problem by adjusting the controller calibration method in [16], [17] to a stance controller of a legged robot. This is achieved by defining a training objective suited for a stance controller as well as deriving a motion model suited to model the stance phase of a legged robot.

Problem C: Develop concept for online computation of reference trajectories. Section IV-C addresses this problem by continuously generating reference trajectories for continuous operation of a legged robot.

Problem D: Develop concept for adaptively refining reference trajectory to the current task and the current environment. Section IV-D addresses this problem by adjusting the method in [16], [17] to calibrate parameters of reference trajectories. This is achieved by parameterizing the reference trajectories, defining a suitable training objective for the operation of the legged robot, and deriving a motion model for the robot's movements.

Finally, Section V presents the application of the proposed algorithms to a third-party high-fidelity simulator of the Unitree A1 in Gazebo. The high-fidelity simulator is provided by Unitree [28] and it includes contact dynamics, multi-body dynamics, sensor noise, etc.

IV. AUTO-TUNING CONTROLLER AND REFERENCE TRAJECTORIES OF LEGGED ROBOT

The overall implementation of the auto-tuning method is as follows: (1) we generate a feasible motion plan for the first few footsteps using TO; (2) we use this initialization as part of the reference-generator function to estimate new trajectories without optimization; (3) the auto-tuning formulation is employed for both improved trajectory tracking (i.e., auto-tuning the swing or stance controller) and modifying the trajectories directly in order to minimize energy consumption, produce forces that minimize slippage, or satisfy a desired

step clearance (i.e., step height) or forward progress (i.e., base velocity), see Fig. 2

A. Training Objectives for Auto-Tuning Swing Controller

For auto-tuning the swing controller in (5), we parametrize the gains, i.e., $\mathbf{K}_p = \mathbf{K}_p(\boldsymbol{\theta})$, $\mathbf{K}_d = \mathbf{K}_d(\boldsymbol{\theta})$, and we use the dynamical model using the notation in (1) for leg i ,

$$\mathbf{x}_{i,k} = \begin{bmatrix} \mathbf{q}_{i,k} \\ \dot{\mathbf{q}}_{i,k} \end{bmatrix} \quad (8a)$$

$$\text{dyn}(\mathbf{x}_{i,k}, \boldsymbol{\tau}_{i,k}) = \begin{bmatrix} \mathbf{q}_{i,k} + dt\dot{\mathbf{q}}_{i,k} + dt^2\ddot{\mathbf{q}}_{i,k} \\ \dot{\mathbf{q}}_{i,k} + dt\ddot{\mathbf{q}}_{i,k} \end{bmatrix} \quad (8b)$$

$$\ddot{\mathbf{q}}_i = \mathbf{M}_i(\mathbf{q}_i)^{-1}(\boldsymbol{\tau}_i - \mathbf{V}_i(\mathbf{q}_i, \dot{\mathbf{q}}_i) - \mathbf{G}_i(\mathbf{q}_i)), \quad (8c)$$

where $\mathbf{M}_i \in \mathbb{R}^{3 \times 3}$ is the mass matrix in the joint space, and $\ddot{\mathbf{q}}_i \in \mathbb{R}^{3 \times 3}$ is the current joint acceleration of foot i . The training objective (2) is defined to improve reference trajectory tracking as

$$\mathbf{y}_k = \begin{bmatrix} \mathbf{p}_{i,k-N|k}^{\text{ref}} \\ \mathbf{v}_{i,k-N|k}^{\text{ref}} \end{bmatrix}, \quad \mathbf{h}(\boldsymbol{\theta}) = \begin{bmatrix} \mathbf{p}_{i,k-N|k} \\ \mathbf{v}_{i,k-N|k} \end{bmatrix}, \quad (8d)$$

where $\mathbf{p}_{i,k-N|k}^{\text{ref}}$ and $\mathbf{v}_{i,k-N|k}^{\text{ref}}$ are positions and velocities of foot i from time step $k-N$ through k as in (10) and (11).

The model mismatch $\mathbf{w}_{k-N|k} \in \mathbb{R}^6$ for the swing controller is computed using (8) and the measured states for the past swing execution. The UKF-based control parameter adaptation method uses (5) and (8) to "simulate" an execution of a swing using the gains defined by the sigma points, $\mathbf{K}_p(\boldsymbol{\theta}^{\text{sp},j})$ and $\mathbf{K}_d(\boldsymbol{\theta}^{\text{sp},j})$, considering the model mismatch, $\mathbf{w}_{k-N|k}$. Note that $\mathbf{p}_{i,k}^b$ and $\mathbf{v}_{i,k}^b$ in (5) can be calculated using the robot forward kinematics, which uses $\mathbf{q}_{i,k}$ as input, and the footstep Jacobian by $\mathbf{v}_{i,k}^b = \mathbf{J}_i(\mathbf{q}_{i,k})\dot{\mathbf{q}}_{i,k}$.

B. Training Objectives for Auto-Tuning Stance Controller

For auto-tuning the stance controller cost function weights in (6), i.e., $\mathbf{Q} = \mathbf{Q}(\boldsymbol{\theta})$, $\mathbf{R} = \mathbf{R}(\boldsymbol{\theta})$, we use the dynamical model (1) as

$$\text{dyn}(\mathbf{x}_k, \mathbf{f}_k) = \begin{bmatrix} \boldsymbol{\Theta}_k + \mathbf{R}_b^w \boldsymbol{\omega}_k \\ \mathbf{r}_k + dt\mathbf{v}_k \\ \boldsymbol{\omega}_k + dt(\sum_{i=1}^4 \hat{\mathbf{I}}^{-1} \begin{bmatrix} \mathbf{p}_{i,k}^b \\ \mathbf{f}_k^i \end{bmatrix} \times \mathbf{f}_k^i) \\ \mathbf{v}_k + dt(\sum_{i=1}^4 \frac{\mathbf{f}_k^i}{m} + \mathbf{g}) \end{bmatrix} \quad (9a)$$

with the state \mathbf{x}_k defined as in (6), where $\hat{\mathbf{I}}$ represents the inertia tensor in the world frame (\times indicates a skew matrix) [27], \mathbf{R}_b^w is the rotation matrix from world to body frame, and \mathbf{g} is the gravity vector. The training objective in (2) is defined to improve reference trajectory tracking as

$$\mathbf{y}_k = \begin{bmatrix} \mathbf{x}_{k-N|k}^{\text{ref}} \end{bmatrix}, \quad \mathbf{h}(\boldsymbol{\theta}) = \begin{bmatrix} \mathbf{x}_{k-N|k} \end{bmatrix} \quad (9b)$$

Here, too, the UKF-based control parameter adaptation method uses (6) and (9) to "simulate" the MPC using the cost function weights defined by the sigma points $\mathbf{Q}(\boldsymbol{\theta}^{\text{sp},j})$, considering the model mismatch $\mathbf{w}_{k-N|k}$.

C. Generating Reference Trajectories

To ensure kinematic and dynamically feasible motion for the first few footsteps, we apply the methods proposed in [4], which uses trajectory optimization.

1) *Reference Trajectory for Feet*: In order to continuously generate new reference trajectories, we use

$$\mathbf{p}_{i,k}^{\text{ref}} = \mathbf{p}_{i,k-N}^{\text{ref}} + \mathbf{v}_{i,k-N}^{\text{ref}} \Delta T \quad (10a)$$

$$\mathbf{p}_{i,k}^{\text{ref}} = \begin{bmatrix} p_{i,k}^{x,\text{ref}} \\ p_{i,k}^{y,\text{ref}} \\ p_{i,k}^{z,\text{ref}} \end{bmatrix}, \quad \mathbf{v}_{i,k}^{\text{ref}} = \begin{bmatrix} v_{i,k}^{x,\text{ref}} \\ v_{i,k}^{y,\text{ref}} \\ v_{i,k}^{z,\text{ref}} \end{bmatrix} \quad (10b)$$

where $\mathbf{p}_{i,k-N}^{\text{ref}}$ denotes the position of foot i on the ground at time step $k-N$, $\mathbf{p}_{i,k}^{\text{ref}}$ is the next position of foot i on the ground at time step k , which are determined by a desired CoM body velocity, $\mathbf{v}_{i,k}^{\text{ref}}$, and the phase time ΔT , i.e., the time that the foot is in stance or swing phase. Throughout, we assume $v_{i,k}^{z,\text{ref}} = 0$.

Then, we use $s(t) = 2\pi \frac{N-k+t}{N}$ with $t = k-N, k-N+1, \dots, k$ to create cycloidal footstep reference trajectories in between the ground positions in (10) with

$$p_{i,t}^{x,\text{ref}} = p_{i,k-N}^{x,\text{ref}} + \left(p_{i,k}^{x,\text{ref}} - p_{i,k-N}^{x,\text{ref}} \right) \frac{s(t) - \sin(s(t))}{2\pi} \quad (11a)$$

$$p_{i,t}^{y,\text{ref}} = p_{i,k-N}^{y,\text{ref}} + \left(p_{i,k}^{y,\text{ref}} - p_{i,k-N}^{y,\text{ref}} \right) \frac{s(t) - \sin(s(t))}{2\pi} \quad (11b)$$

and

$$p_{i,t}^{z,\text{ref}} = \begin{cases} p_{i,k-N}^{z,\text{ref}} + p_{\text{max}}^z \frac{1 - \cos(s(t))}{2} & \text{if } s(t) \leq \pi \\ p_{i,k}^{z,\text{ref}} + \left(p_{\text{max}}^z + p_{i,k-N}^{z,\text{ref}} - p_{i,k}^{z,\text{ref}} \right) \frac{1 - \cos(s(t))}{2} & \text{if } s(t) > \pi \end{cases} \quad (11c)$$

where p_{max}^z is the step clearance/apex height of the swing trajectory, i.e., the maximum p^z component of the footstep in swing. If the foot is in stance, $p_{\text{max}}^z = 0$.

2) *Reference Trajectory for Center of Mass*: The reference trajectory of the body's CoM is specified by

$$\begin{bmatrix} \mathbf{r}_k^{\text{ref}} \\ \Theta_k^{\text{ref}} \end{bmatrix} = \begin{bmatrix} \mathbf{r}_{k-N}^{\text{ref}} \\ \Theta_{k-N}^{\text{ref}} \end{bmatrix} + \begin{bmatrix} \mathbf{v}_{k-N}^{\text{ref}} \\ \boldsymbol{\omega}_{k-N}^{\text{ref}} \end{bmatrix} \Delta T \quad (12)$$

with the current body position, $\mathbf{r}_{k-N}^{\text{ref}} \in \mathbb{R}^3$, and orientation, $\Theta_{k-N}^{\text{ref}} \in \mathbb{R}^3$, and $\mathbf{r}_k^{\text{ref}}, \Theta_k^{\text{ref}}$ are composed of the next desired body position, and $\mathbf{v}_{k-N}^{\text{ref}}, \boldsymbol{\omega}_{k-N}^{\text{ref}}$ is the desired velocity (linear and angular) over the phase duration ΔT . Here, we use $\bar{s}(t) = \frac{N-k+t}{N}$ with $t = k-N, k-N+1, \dots, k$ to linearly interpolate between current and desired states with

$$\begin{bmatrix} \mathbf{r}_t^{\text{ref}} \\ \Theta_t^{\text{ref}} \end{bmatrix} = \bar{s}(t) \begin{bmatrix} \mathbf{r}_k^{\text{ref}} \\ \Theta_k^{\text{ref}} \end{bmatrix} + (1 - \bar{s}(t)) \begin{bmatrix} \mathbf{r}_{k-N}^{\text{ref}} \\ \Theta_{k-N}^{\text{ref}} \end{bmatrix}.$$

D. Training Objectives for Auto-Tuning Reference Trajectory

For auto-tuning the reference trajectory, we use the same states and dynamic model as for the swing controller in (8). The main difference is in the implementation of the training objective and control parameters, i.e., we do not auto-tune the \mathbf{K}_p and \mathbf{K}_d gain matrices. Instead, we tune two parameters

often desirable for legged robots, which are to achieve a desired step clearance in order to step over an obstacle and reach a desired velocity. Hence for leg i , we parametrize the step clearance and the desired forward velocity,

$$p_{\text{max}}^z = \theta_z, \quad v_{i,k}^{x,\text{ref}} = \theta_v, \quad \boldsymbol{\theta} = \begin{bmatrix} \theta_z \\ \theta_v \end{bmatrix}$$

and the training objective for adapting such parameters is

$$\mathbf{y}_k = \begin{bmatrix} p_{\text{des}}^z \\ v_{\text{des}}^x \\ 0 \\ 0 \end{bmatrix}, \quad \mathbf{h}(\boldsymbol{\theta}) = \begin{bmatrix} h_z(\boldsymbol{\theta}) \\ h_v(\boldsymbol{\theta}) \\ h_f(\boldsymbol{\theta}) \\ h_e(\boldsymbol{\theta}) \end{bmatrix}, \quad (13)$$

where $h_z(\boldsymbol{\theta})$ is the achieved step clearance, $h_v(\boldsymbol{\theta})$ is the achieved velocity, $h_f(\boldsymbol{\theta})$ computes lateral forces of the robot and is used to reduce slippage, and $h_e(\boldsymbol{\theta})$ is the energy consumption. Note that p_{des}^z measures the achieved step clearance in closed loop, whereas p_{max}^z is the control parameter that defines the reference trajectory. Thus, p_{des}^z and p_{max}^z may be different, e.g., due to a model mismatch of the physical robot and the dynamical system model.

1) *Step clearance optimization $h_z(\boldsymbol{\theta})$* : To propagate the states for optimizing the step clearance, we use the same method as in Section IV-A. However, here we “simulate” the various reference trajectories defined using the sigma points for the swing leg, where we do not change the controller gains of $\mathbf{K}_p, \mathbf{K}_d$. Hence, we modify the reference trajectory of the footstep using the reference generator function described in Section IV-C. Thus, $h_z(\boldsymbol{\theta}) = \max(p_{k-N|k}^z)$, i.e., the maximum value of the z component of the footstep trajectory (after propagation with $\mathbf{w}_{k-N|k}$).

2) *Forward progress optimization $h_v(\boldsymbol{\theta})$* : For forward progress, we propagate the CoM position using the desired velocity, v_{des}^x , with (12), where the initial state is the actual state.

3) *Slippage optimization $h_f(\boldsymbol{\theta})$* : Representing slippage of the foot first requires the propagation of ground reaction forces along the trajectory, which can be achieved by rearranging the MPC state-space equations in (6),

$$\mathbf{f}_k = \mathbf{B}^+(\mathbf{x}_{k+1} - \mathbf{A}\mathbf{x}_k), \quad (14)$$

where \mathbf{B}^+ is the Moore-Penrose inverse and \mathbf{x}_k is propagated using the sigma points. We chose (14) due to its similarity to (6), but other slippage formulations are possible, too, e.g., as in [29]. If the foot is not in contact, $\mathbf{f}_k = \mathbf{0}$. Then, the foot slippage is computed using similar ideas to friction cone constraints (see, e.g., [4]) as

$$h_f(\boldsymbol{\theta}) = \sum_{i=1}^4 \frac{\sqrt{f_{i,\text{avg}}^x + f_{i,\text{avg}}^y}}{f_{i,\text{avg}}^z},$$

$$\begin{bmatrix} f_{i,\text{avg}}^x \\ f_{i,\text{avg}}^y \\ f_{i,\text{avg}}^z \end{bmatrix} = \frac{1}{N} \sum_{t=k-N}^k \begin{bmatrix} |f_{i,t}^x| \\ |f_{i,t}^y| \\ |f_{i,t}^z| \end{bmatrix}.$$

4) *Energy consumption optimization $h_e(\boldsymbol{\theta})$* : The energy consumption is included as it may be desirable to find the best balance between maximizing step clearance and forward

progress, while simultaneously minimizing the energy consumption required. The energy consumption can be computed using the joint velocity $\dot{\mathbf{q}}$, torque $\boldsymbol{\tau}_t$, and phase time ΔT , as

$$h_e(\boldsymbol{\theta}) = \sum_{t=k-N}^k |\dot{\mathbf{q}}_t|^\top |\boldsymbol{\tau}_t| \Delta T.$$

Remark 3: To ensure we do not violate kinematic or dynamic constraints as we simulate new trajectories using $\boldsymbol{\theta}$, we can use $|\mathbf{R}_{w}^b [\mathbf{p}_{i,k} - \mathbf{r}_k] - \bar{\mathbf{p}}_i^b| < \mathbf{b}$, where \mathbf{R}_{w}^b is the rotation matrix from base to world frame, and $\bar{\mathbf{p}}_i^b$ is the nominal footstep position centered in the kinematic bounding box specified by $\mathbf{b} \in \mathbb{R}^3$. If the bounding box constraint does not hold, we can impose a cost for that particular simulation guiding the auto-tuner away from such control parameters.

V. RESULTS USING PHYSICS-BASED SIMULATOR

We implemented the auto-tuning algorithms for the Unitree A1 robot [28] within a realistic simulation using Gazebo with the bullet physics engine. The Unitree A1 simulator is made available by the robot manufacturer. The simulation considers not only localization noise but also friction within each torque-controlled motor, where torque commands are sent using the robotic operating system (ROS). Further, the Gazebo simulation environment is executed in a parallel thread, and hence the controller and the auto-tuning algorithm need to be executed during run-time of the simulation. Hence, being able to execute the applied algorithms using this high-fidelity simulator indicates their applicability to the physical robot. The A1 Unitree is a 12.7 kg quadruped robot with 3 degrees of freedom per leg. For obtaining an initial trajectory as in [4], we use the nonlinear programming solver IPOPT [30]. Optimization of the MPC stance controller was done using the quadratic programming solver qpOASES [31]. Constraints and overall problem formulation were setup using the CasADi [32] software. The phase time was chosen as 0.2 s, which is also the update frequency of the parameters. The swing and stance controller are executed at sampling frequency at $dt = 0.01$ s, and $N_{\text{MPC}} = 5$ in (6).

A. Auto-Tuning for Optimizing Step Clearance and Forward Progress

The first test case is to auto-tune reference trajectories aiming at providing a desired forward progress and step clearance. Here, we consider only the first and second element in (13), where we use $p_{\text{des}}^z = 0.15$ m and $v_{\text{des}}^x = 0.45$ m/s. The results are shown in Fig. 3-A), B), C). In A), we show the cost defined as $\|\mathbf{y}_k - \mathbf{h}(\boldsymbol{\theta})\|_{C_y^{-1}} = \left\| \begin{bmatrix} p_{\text{des}}^z - h_z(\boldsymbol{\theta}) & v_{\text{des}}^x - h_v(\boldsymbol{\theta}) \end{bmatrix} \right\|_{C_y^{-1}}$ as in (2). Thus, as the cost decreases, the closer the actual step clearance and forward progress get to the desired value. We decompose the results in A) and show how the actual velocity reaches the desired velocity in B), and how the actual step clearance reaches the desired step clearance in C).

B. Auto-Tuning for Optimizing Step Clearance, Forward Progress, Slippage, and Energy Consumption

In D), we show the cost when setting a desired step clearance and forward progress with $p_{\text{des}}^z = 0.1$ m and

$v_{\text{des}}^x = 0.4$ m/s, while minimizing energy consumption and foot slippage as in (13). The cost decreases quickly within a few time steps, which indicates that the auto-tuning method was able to find a good balance between reaching its desired step clearance and forward progress, while ensuring a minimization of both energy consumption and foot slippage.

C. Auto-Tuning Stance Controller

Next, we demonstrate the auto-tuning method to calibrate the gains of the stance controller. For simplicity, here we auto-tune only the diagonal elements of $\mathbf{Q} = \mathbf{Q}(\boldsymbol{\theta})$ in (6), but \mathbf{R} can similarly be tuned. We initialize $\mathbf{Q} = \text{diag}([1, 1, Q_3, Q_4, Q_5, Q_6, 1000, 1, 1, 1, 1, 1])$, where Q_3, Q_4, Q_5, Q_6 are initialized to 300, and will be further calibrated, i.e., $\mathbf{Q}(\boldsymbol{\theta})$ with $\boldsymbol{\theta} \in \mathbb{R}^4$. The cost function is tuned online for 20 s of locomotion using trot gait. In Fig. 3-E), we present the cost difference between the auto-tuning case and the no auto-tuning case. The difference decreases over time, which shows that the auto-tuning improves the cost using auto-tuning decreased significantly compared to the cost without using auto-tuning. Additionally, without auto-tuning, the robot eventually falls after about 8 seconds of locomotion, because the initial gains of the MPC controller have not been properly hand-tuned. Note that some hand-tuning was necessary to ensure initial gains that at least allow the robot to start moving a few steps without falling. After auto-tuning the stance controller for 20 s, $Q_3 = 5245.53$, $Q_4 = 1172.34$, $Q_5 = 662.45$, and $Q_6 = 1172.25$. Further, we applied the tuned controller along with the tuned swing controller in Section V-D and demonstrated robust locomotion on uneven terrain, see Fig. 4.

D. Auto-Tuning Swing Controller

Lastly, we apply the auto-tuning method to calibrate the gains of the swing controller, \mathbf{K}_p and \mathbf{K}_d . We choose $\boldsymbol{\theta} \in \mathbb{R}^6$ consisting of the three diagonal elements of the proportional gain matrix, \mathbf{K}_p , and the three diagonal elements of the derivative gain matrix, \mathbf{K}_d . We initialized $\mathbf{K}_d = \text{diag}([0.1, 10, 10])$ and $\mathbf{K}_p = \text{diag}([150.11, 16.11, 10.11])$ to be equal to the leg's natural frequency using the inverted pendulum model multiplied by the operational mass matrix, see [27] for more detail. Fig. 3-F), shows the difference of the cost when not auto-tuning with the cost when auto-tuning for over 20 seconds of locomotion. The cost curve follows a downward trend, demonstrating that the cost when auto-tuning is smaller than without auto-tuning. The gains using auto-tuning outperformed the gain selection without auto-tuning, which was based on the current natural frequency calculation and was shown to perform well in [27]. The final gains after tuning was $\mathbf{K}_d = \text{diag}([13.82, 12.42, 17.80])$ and $\mathbf{K}_p = \text{diag}([166.26, 22.75, 14.03])$.

Fig. 4 shows the robot walking on uneven ground using auto-tuning of the swing controller gains and the stance controller cost function.

Remark 4: While this letter presents one simulation run, we have obtained similar results for repeated test cases, which is expected as the auto-tuning formulation is robust to noise/disturbances due to its filter-based design.

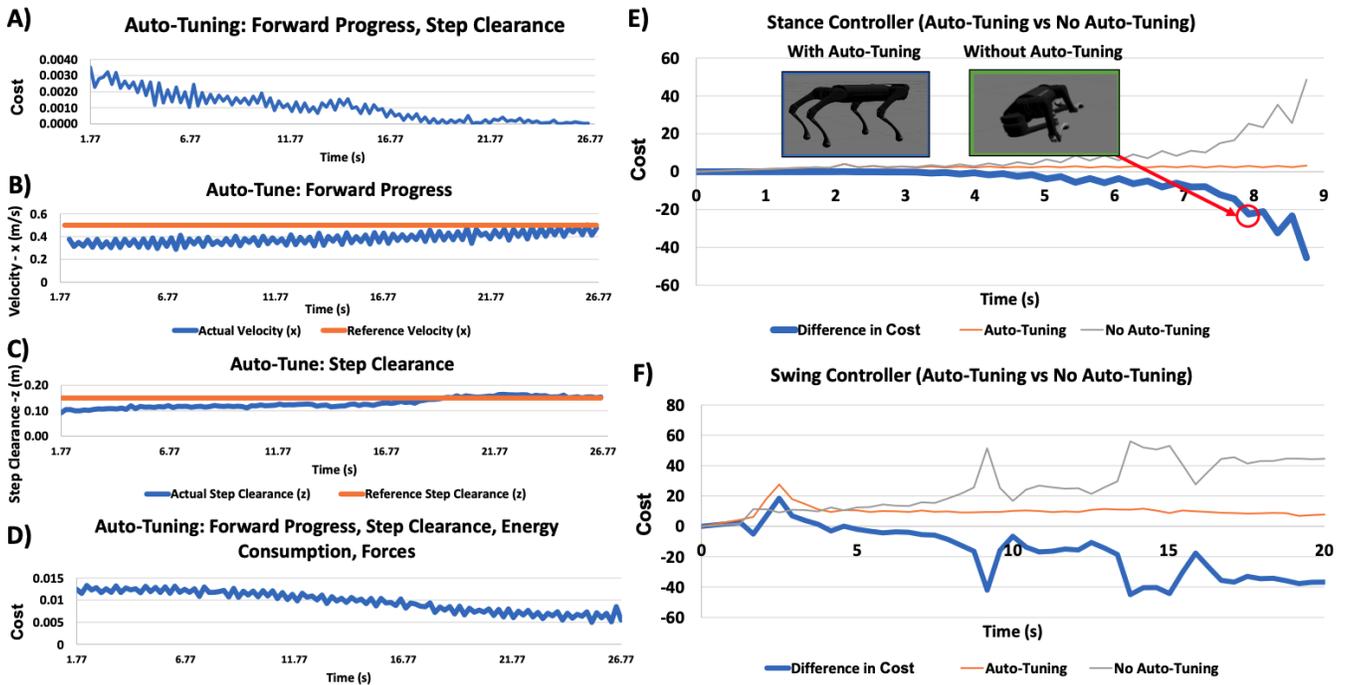


Fig. 3: **Auto-Tuning results.** The auto-tuning method successfully calibrated control parameters that generate reference trajectories (A-D), and controller gains (E-F). To make this evaluation, we used the cost as calculated using (2). The cost decreases (A and D) as the control parameters (e.g., forward progress and step clearance) get to their desired reference values. In E-F, we show the difference between the cost when not auto-tuning with the cost when auto-tuning.

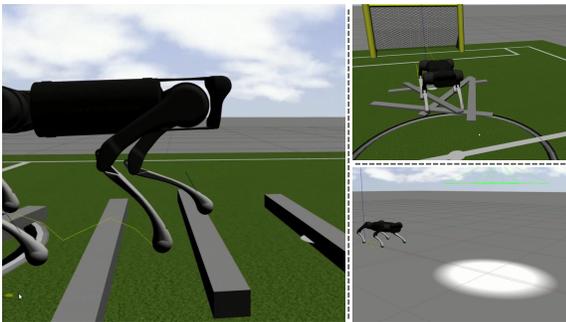


Fig. 4: **Results for robust locomotion on uneven terrain.** After auto-tuning the stance and swing controllers, we employ the robot on several test cases to demonstrate robust locomotion. We show the robot traversing over and on large beams and planks in the left and top right with a trot gait, respectively, and demonstrate a successful jumping gait in the bottom right. The robot is not aware of the obstacles and must overcome them by using the auto-tuned controllers.

E. Computation Times

The computation times of auto-tuning the swing and stance controllers, as well as the reference trajectory are shown in Table I. Table I lists the maximum, minimum, and median computation times of the auto-tuning algorithm for tuning the gains of the swing controller, the MPC weights of the stance controllers and the parameters of the reference trajectory (Reference Traj. A – D in Table I refers to the training objectives specified by A – D in Fig. 3). The computation times were

TABLE I: Auto-Tuning Computation Times

Auto-Tune	Min. Time	Median Time	Max. Time
Reference Traj. A)	0.0108s	0.0143s	0.0370s
Reference Traj. B)	0.0072s	0.0089s	0.0121s
Reference Traj. C)	0.0086s	0.0093s	0.0235s
Reference Traj. D)	0.0632s	0.0662s	0.0667s
Stance Controller E)	0.0409s	0.0465s	0.0807s
Swing Controller F)	0.0078s	0.0569s	0.1006s

obtained while running on a laptop using 4 CPU cores (Intel core i7-8850H CPU at 2.60 Ghz) with a Quadro P3200 GPU. Further, the controller ran in parallel with a Gazebo simulation, which requires more computational resources than an implementation on hardware would. As the computation times are below the update rate of the auto-tuner updating at the end of the phase time with 0.2 s, we can conclude that the algorithms in this letter can be executed in real time on comparable computational resources.

Remark 5: The computation times scale linearly with the amount of parameters to be tuned. Although the applied algorithms can be executed online, depending on user specifications, there is the option to execute the algorithms offline or in parallel. E.g., the auto-tuner for the swing and stance controllers can update the gains on a different thread, which runs in parallel to the low-level/high-level controllers of the robot. Additionally, the methods could also be implemented offline if computational resources are limited.

VI. CONCLUSIONS

In this work we successfully demonstrated robust locomotion. An auto-tuning method was implemented, which is based on an unscented Kalman filter formulation, to calibrate the gains of the swing controller and the cost function weights of a stance controller. We also demonstrated that the method can be applied for directly auto-tuning the reference trajectory, i.e., reference trajectories are calibrated to make the robot achieve a desired forward progress and step clearance, while also minimizing energy consumption and foot slippage. We showed that the method can be easily generalized to consider diverse control parameter by demonstrating the auto-tuning on both controller gains and on physically meaningful parameters of a reference trajectory. Future work may include experimental validation on the robot platform and extending the method to consider tuning for optimal footstep timings of various gait sequences, consider dynamic environments that require the auto-tuner to calibrate changing step heights and base velocities, and also adapt the aggressiveness of the tuning through the unscented Kalman filter's covariance matrices based the robot's environment or desired task. Future work may also include more sudden changes in the robot's motion, e.g., by leveraging the "prediction model" of the Kalman filter in addition to the "measurement model" used in this letter.

REFERENCES

- [1] M. H. Raibert, *Legged Robots That Balance*. USA: Massachusetts Institute of Technology, 1986.
- [2] S. Kajita, F. Kanehiro, K. Kaneko, K. Fujiwara, K. Harada, K. Yokoi, and H. Hirukawa, "Biped walking pattern generation by using preview control of zero-moment point," in *2003 IEEE International Conference on Robotics and Automation*, vol. 2, pp. 1620–1626, 2003.
- [3] S. Kuindersma, F. Permenter, and R. Tedrake, "An efficiently solvable quadratic program for stabilizing dynamic locomotion," in *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 2589–2594, 2014.
- [4] A. W. Winkler, C. D. Bellicoso, M. Hutter, and J. Buchli, "Gait and trajectory optimization for legged systems through phase-based end-effector parameterization," *IEEE Robotics and Automation Letters*, vol. 3, no. 3, pp. 1560–1567, 2018.
- [5] Y. Zhao, H.-C. Lin, and M. Tomizuka, "Efficient Trajectory Optimization for Robot Motion Planning," in *2018 15th International Conference on Control, Automation, Robotics and Vision (ICARCV)*, pp. 260–265, 2018.
- [6] J. Norby and A. M. Johnson, "Fast global motion planning for dynamic legged robots," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 3829–3836, 2020.
- [7] A. Bratta, R. Orsolino, M. Focchi, V. Barasuol, G. G. Muscolo, and C. Semini, "On the hardware feasibility of nonlinear trajectory optimization for legged locomotion based on a simplified dynamics," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 1417–1423, 2020.
- [8] N. Rathod, A. Bratta, M. Focchi, M. Zanon, O. Villarreal, C. Semini, and A. Bemporad, "Mobility-enhanced MPC for legged locomotion on rough terrain," *arXiv preprint arXiv:2105.05998*, 2021.
- [9] V. Tsounis, M. Alge, J. Lee, F. Farshidian, and M. Hutter, "Deepgait: Planning and control of quadrupedal gaits using deep reinforcement learning," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 3699–3706, 2020.
- [10] J. Lee, J. Hwangbo, L. Wellhausen, V. Koltun, and M. Hutter, "Learning quadrupedal locomotion over challenging terrain," *Science Robotics*, vol. 5, no. 47, 2020.
- [11] K. Ito and F. Matsuno, "A study of reinforcement learning for the robot with many degrees of freedom - acquisition of locomotion patterns for multi-legged robot," in *2002 IEEE International Conference on Robotics and Automation*, vol. 4, pp. 3392–3397, 2002.
- [12] Z. Liang, S. Zhu, and X. Jin, "Walking parameters design of biped robots based on reinforcement learning," in *30th Chinese Control Conference*, pp. 4017–4022, 2011.
- [13] H. Kimura, T. Yamashita, and S. Kobayashi, "Reinforcement learning of walking behavior for a four-legged robot," in *40th IEEE Conference on Decision and Control*, vol. 1, pp. 411–416, 2001.
- [14] A. Schperberg, S. Tsuei, S. Soatto, and D. Hong, "SABER: Data-driven motion planner for autonomously navigating heterogeneous robots," *IEEE Robotics and Automation Letters*, vol. 6, no. 4, pp. 8086–8093, 2021.
- [15] S. Gangapurwala, M. Geisert, R. Orsolino, M. Fallon, and I. Havoutis, "Real-time trajectory adaptation for quadrupedal locomotion using deep reinforcement learning," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 5973–5979, 2021.
- [16] M. Menner, K. Berntorp, and S. Di Cairano, "A Kalman filter for online calibration of optimal controllers," in *2021 IEEE Conference on Control Technology and Applications (CCTA)*, pp. 441–446, 2021.
- [17] M. Menner, K. Berntorp, and S. Di Cairano, "Automated controller calibration by Kalman filtering," *arXiv preprint arXiv:2111.10832*, 2021.
- [18] A. Marco, P. Hennig, J. Bohg, S. Schaal, and S. Trimpe, "Automatic LQR tuning based on gaussian process global optimization," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 270–277, 2016.
- [19] M. Neumann-Brosig, A. Marco, D. Schwarzmann, and S. Trimpe, "Data-efficient autotuning with Bayesian Optimization: An industrial control study," *IEEE Transactions on Control Systems Technology*, vol. 28, p. 730–740, May 2020.
- [20] R. Calandra, A. Seyfarth, J. Peters, and M. P. Deisenroth, "Bayesian optimization for learning gaits under uncertainty," *Annals of Mathematics and Artificial Intelligence*, vol. 76, no. 1, pp. 5–23, 2016.
- [21] A. Rai, R. Antonova, S. Song, W. Martin, H. Geyer, and C. Atkeson, "Bayesian optimization using domain knowledge on the atrias biped," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 1771–1778, 2018.
- [22] C. Gehring, S. Coros, M. Hutter, C. Dario Bellicoso, H. Heijnen, R. Diethelm, M. Bloesch, P. Fankhauser, J. Hwangbo, M. Hoepflinger, and R. Siegwart, "Practice makes perfect: An optimization-based approach to controlling agile motions for a quadruped robot," *IEEE Robotics Automation Magazine*, vol. 23, no. 1, pp. 34–43, 2016.
- [23] Y. Rahman, A. Xie, J. B. Hoagg, and D. S. Bernstein, "A tutorial and overview of retrospective cost adaptive control," in *2016 American Control Conference (ACC)*, pp. 3386–3409, 2016.
- [24] D. Clever, R. Malin Schemschat, M. L. Felis, and K. Mombaur, "Inverse optimal control based identification of optimality criteria in whole-body human walking on level ground," in *2016 6th IEEE International Conference on Biomedical Robotics and Biomechanics (BioRob)*, pp. 1192–1199, 2016.
- [25] F. L. Haufe, S. Maggioni, and A. Melendez-Calderon, "Reference trajectory adaptation to improve human-robot interaction: A database-driven approach," in *2018 40th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pp. 1727–1730, 2018.
- [26] G. A. Terejanu, "Unscented Kalman filter tutorial," *University at Buffalo, Buffalo*, 2011.
- [27] J. Di Carlo, P. M. Wensing, B. Katz, G. Bledt, and S. Kim, "Dynamic locomotion in the MIT Cheetah 3 through convex model-predictive control," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 1–9, 2018.
- [28] Unitree Robotics, "GitHub repository for Unitree A1," <https://github.com/unitreerobotics>, accessed 2021-09.
- [29] G. Bledt, P. M. Wensing, S. Ingersoll, and S. Kim, "Contact model fusion for event-based locomotion in unstructured terrains," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 4399–4406, 2018.
- [30] A. Wächter and L. T. Biegler, "On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming," *Math. Program.*, vol. 106, no. 1, pp. 25–57, 2006.
- [31] H. J. Ferreau, C. Kirches, A. Potschka, H. G. Bock, and M. Diehl, "qpOASES: A parametric active-set algorithm for quadratic programming," *Math. Program. Computation*, vol. 6, no. 4, pp. 327–363, 2014.
- [32] J. Andersson, J. Gillis, G. Horn, J. Rawlings, and M. Diehl, "CasADi – A software framework for nonlinear optimization and optimal control," *Math. Program. Computation*, vol. 11, no. 1, pp. 1–36, 2019.