Proprioception and Tail Control Enable Extreme Terrain Traversal by Quadruped Robots

Yanhao Yang, Joseph Norby, Justin K. Yim, and Aaron M. Johnson

***Abstract*—Legged robotic systems leverage ground contact and the reaction forces they provide to achieve agile locomotion. However, uncertainty coupled with the discontinuous nature of contact can lead to failure in real-world environments with unexpected height variations, such as rocky hills or curbs. To enable dynamic traversal of extreme terrain, this work introduces the utilization of proprioception to estimate and react to unknown hybrid events and elevation changes and a two-degree-of-freedom tail to improve control independent of contact. Simulation results over unforeseen elevation changes show that our method can stabilize locomotion over height changes of up to 1.5 times the leg length.**

***Paper Type* – Original Work.**

I. Introduction

Quadruped platforms have a growing potential in practical mobile robotic applications. The legged design enables agile traversal and challenging maneuvers, making quadruped robots ideal for tasks in complex terrain such as environmental monitoring, last-mile deliveries, and disaster relief. However, their control strategies are currently limited in their ability to traverse extreme terrain commonly found in real-world environments. Here, "extreme terrain" refers to environments with elevation changes so drastic that they disturb the system beyond the basin of attraction of stable motion controllers and require additional planning or control to traverse.

Animals can easily traverse these extreme environments as shown in Fig. 1. Their methods are diverse, ranging from cats placing feet in repeated locations to ensure reliable contact [1] to animals using tails to reject disturbances [2] to distributed limb control to promote rapid, reactive behaviors [3]. These biological phenomena inspire us to propose new approaches for the perception and control of quadruped robots to improve robustness across extreme terrains.

One common approach to improve performance on rough terrain is to leverage perception to aid control. Exteroceptive sensors such as depth cameras or lidars are common for legged robots, which provide almost out-of-the-box mapping of the environment [4]. However, the created map may not be perfect due to factors such as sensor noise, obstacles, reflections, or lighting conditions, as shown in [5]. Proprioception, on the other hand, is another option for adapting to unstructured terrain and has proven to be powerful and trustworthy. [6] demonstrates that momentum observer-based contact sensing and event-based controllers can be robust to unknown unstructured terrain. [5] uses learning techniques for proprioception to successfully go up and down stairs blindly.

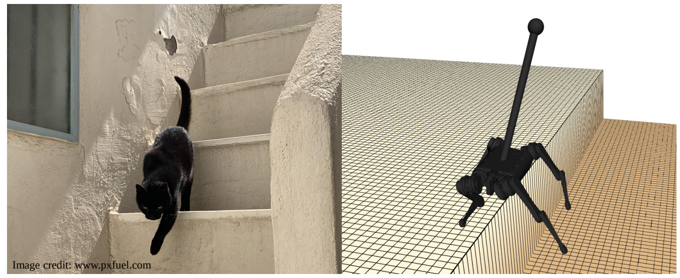

Fig. 1. A comparison of a cat and a quadruped robot with a tail walking over extreme elevation changes. Animals provide inspiration for the design, control, planning, and perception strategies of quadruped robots for traversing extreme terrains.

Another approach to improve stability over rough terrain is to bypass perception and design systems with dynamics that are stable with respect to contact errors. Many bio-inspired control approaches to improve legged locomotion stability in real-world environments, including passive stabilization and additional actuators, have been investigated. Swing-leg retraction [7], [8] observed in biology is an effective passive stabilization technique that improves gait stability on uneven terrain by changing the shape of the hybrid guard [9]. However, it cannot actively respond to extreme height changes and has a basin of attraction limited by the stroke of the leg. Some researchers have taken inspiration from other biological systems and instead added actuators that do not rely on contact such as tails or flywheels. [10], [11] equip a quadruped with a tail and use it to effectively suppress impulsive perturbations. However, given the limited angular deflection, it requires nonholonomic behavior like conic motion to maximize maneuverability [12].

In this work, we propose a proprioception-based gait planner and time-scale decoupled sequential distributed nonlinear model predictive control (NMPC) for tail control to improve extreme terrain traversal performance from both perception and control perspectives, summarized in Fig. 2. First, the gait planner updates contact, body, foothold, and terrain references

*This research was sponsored by the Army Research Office under Grant Number W911NF-19-1-0080 and the National Science Foundation under Grant Number CCF-2030859 to the Computing Research Association for the CIFellows Project. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Office, the National Science Foundation, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein.

The authors are with the Department of Mechanical Engineering, Carnegie Mellon University, Pittsburgh, PA 15213, USA, {yanhaoy, amj1}@andrew.cmu.edu

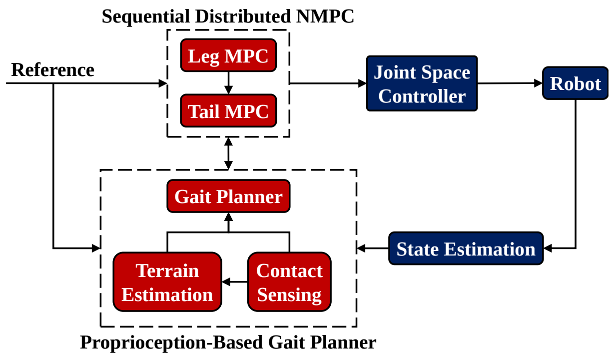

Fig. 2. Block diagram of the proposed control scheme. The system consists of three main parts: proprioception-based gait planner (bottom left red blocks), sequential distributed NMPC for leg and tail control (top left red blocks), and robot interface (blue blocks on right). Among them, the red blocks are the main works of this paper.

based on proprioceptive information to accommodate unexpected hybrid events and unknown elevation changes. Second, we show that 2 degree-of-freedom (DOF) tails with appropriate swing trajectories can complement highly underactuated legs restricted by kinematic and friction constraints caused by orientation errors to improve leg motion performance under hybrid perturbations. Specifically, time-scale decoupled tail control through sequential distributed NMPC is applied to generate non-holonomic tail behavior and efficient cooperation with existing leg controllers. Simulation results show that using the proposed proprioception-based gait planner, the quadruped robot can robustly walk down unexpected cliffs with heights up to 1.5 times the leg length. The suggested tail control method can further make such extreme terrain traversal robust to prior angular perturbations.

II. Methods

The full-stack control framework shown in Fig. 2 is based on Quad-SDK [13], which is augmented with the proposed proprioception and tail control algorithms to accommodate unexpected hybrid events and unknown elevation changes. Specifically, the proposed method can be divided into four modules: contact sensing, terrain estimation, gait planning, and tail control. These modules estimate contact and terrain and use this information to improve control to adapt to the terrain.

A. Contact Sensing Finite-State Machine

Legged robots as hybrid systems rely heavily on correct contact information to properly control and plan ground reaction forces (GRF) to traverse unstructured terrain. In this section, we discuss the implementation of a proprioception-based finite state machine and use it to perform contact sensing.

As shown in Fig. 3, the clock-based nominal contact schedule is divided into two phases - stand and swing. The finite-state machine uses leg extension and contact force to subdivide the touchdown process. As the clock transitions to the standing phase, the legs begin to extend. By examining the leg extension, contact loss can be identified and hyperextension can be avoided. If the corresponding conditions are

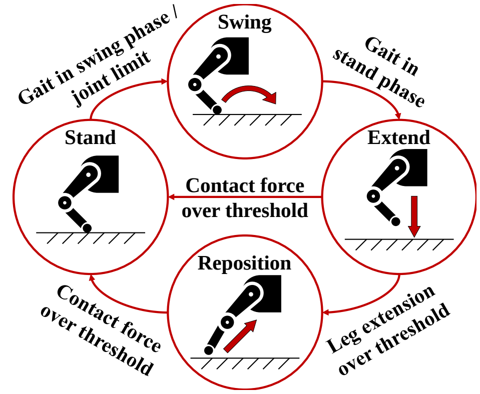

Fig. 3. Illustration of the contact sensing finite state machine. The finite state machine detects the contact state of each leg of the robot according to the listed conditions. It notifies the robot if the contacts are running on a predetermined schedule or experiencing an unexpected loss of contact.

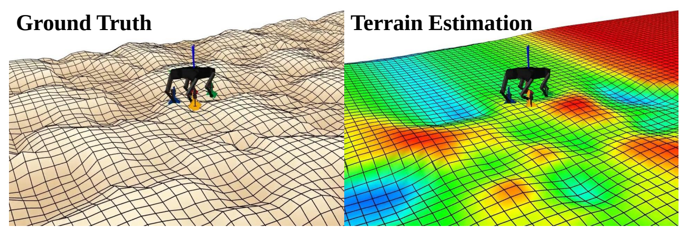

Fig. 4. Terrain estimation results obtained by simulating quadrupeds walking on rough terrain and using the proposed terrain estimation algorithm to model height variations without prior knowledge. Left: Ground truth. Right: Estimated terrain elevation.

met, the finite-state machine will claim a contact loss and switch to the repositioning phase and the corresponding gait, see Section II-C for details. On the other hand, solid support can be confirmed by examining the contact force obtained by a momentum observer or pressure sensor. Finally, early swings can be performed when the leg reaches the kinematic limits to avoid physical hard stops. Note that, unlike [6], we only consider the case of delayed contact since early contact can be handled by force control, while delayed contact would break the nominal support.

B. Terrain Estimation

Successful planning and control of legged robots on rough terrain requires not only contact information but, more importantly, a good understanding of the height variation of the terrain. Thus we propose a terrain estimation algorithm that aims to estimate unknown elevation changes by filtering the foothold history based on contact sensing information for planning the desired body trajectory and foothold. It provides terrain estimation updates for both swing and stance legs, and is designed to use limited sensors to build maps accurate enough to aid future control and planning. As shown in Fig. 4, the quadruped robot successfully estimated rough terrain and was able to use it for future planning and control.

The terrain estimation process is summarized in two steps in Algorithm 1. First, check the contact sensing result

Algorithm 1: Terrain Estimation

```
Input: terrainEst, swingSpace, contactState, footPos
Output: terrainEst, swingSpace
for legIdx ← 0 to 3 do
    if !contactState.at(legIdx) then
        if footPos.z < terrainEst.at(footPos.xy) then
            terrainEst.at(footPos.xy) ← footPos.z;
            swingSpace.at(legIdx).append(footPos.xy);
        end
    else
        terrainEst.at(footPos.xy) ← footPos.z;
        while !swingSpace.at(legIdx).empty() do
            terrainEst.at(swingSpace.at(legIdx).pop())
              ← footPos.z;
        end
    end
end
```

`contactState`. For a non-contact foot, including swing or missed contact, we check if it is below the terrain estimate `terrainEst` at the current position `footPos.xy`. If so, the current terrain estimate for this location is inaccurate, but the actual height is uncertain. We record the current height `footPos.z` temporarily, but also record that position in the list `swingSpace` and wait for the foot to next contact the ground to update it. Second, if it is in contact, we save the current foot position in the terrain estimation history. At the same time, we update all temporary values in `swingSpace` with it. Finally, we apply filters [14] to the discrete terrain estimation history to inpaint and smooth it. Note that, as shown in [15], proprioception-based terrain estimation can also be fused with exteroceptive sensing and be generated from discrete footholds by Gaussian process regression.

C. Gait Planner

Given contact sensing information and the estimated terrain, it is important to plan an appropriate gait which includes contact schedule and locations, to traverse this environment stably. In this section, we propose a gait planner that independently plans future contact and foothold sequences for each leg to reflect the actual terrain and contact from proprioception, and converge asymptotically to a central clock-based gait.

In order to reduce the impact on NMPC warm-starts, the modification of the contact schedule is designed to occur as little as possible while also capturing the entire adaptation process at once. We define three discrete events in the contact schedule: 1) miss contact, 2) landing, and 3) early swing. For each event, we generate an adaptive gait with an offset phase of a specific gait reference time t_a, a convergence reference time t_c and a convergence period T_c. The linearly changing weight w_a is calculated according to the convergence period T_c. Finally, the gait phase actually applied to the control $\phi(t)$ is calculated as a weighted sum adaptive gait and a nominal gait based on the central clock t_0 and the nominal period T_0.

$$\begin{aligned}\phi(t) &= (1-w_a)\frac{\bmod(t-t_0,T_0)}{T_0} + w_a\frac{\bmod(t-t_a,T_0)}{T_0}\\ w_a &= \max\left(1-\frac{(t-t_c)}{T_c},0\right)\end{aligned} \tag{1}$$

Specifically, for missed contact, we clear the current stand phase. When contact sensing confirms the landing, we perform a full stand phase. Finally, early swing due to kinematic constraints will assign a full swing phase. Given an adaptive gait, we can predict all future contact schedules and converge to a nominal clock-based gait.

Similar to Raibert's heuristic [16], [17], each desired footstep point p_l is determined from the NMPC-predicted trajectories as well as dynamics and kinematics heuristics as

$$\boldsymbol{p}_l = \begin{cases}\boldsymbol{p}_{\text{center}} + \boldsymbol{p}_{\text{vel}} + \boldsymbol{p}_{\text{centrifugal}} & \text{(stand)}\\ \boldsymbol{R}_v\boldsymbol{p}_0 & \text{(reposition)}\end{cases} \tag{2}$$

where

$$\begin{aligned}\boldsymbol{p}_{\text{center}} &= \underset{p}{\operatorname{argmin}}\left(\max_{i\in\text{stance}}\left\|\boldsymbol{p}-\boldsymbol{p}_{lb,i}\right\|_2^2\right)\\ \boldsymbol{p}_{\text{vel}} &= \sqrt{\frac{p_{lb_z,\text{td}}}{g}}\left(\dot{p}_{b,\text{td, ref}}-\dot{\boldsymbol{p}}_{b,\text{td}}\right)\\ \boldsymbol{p}_{\text{centrifugal}} &= \frac{p_{lb_z,\text{td}}}{g}\dot{\boldsymbol{p}}_{b,\text{td}}\times\boldsymbol{\omega}_{\text{ref}}\end{aligned} \tag{3}$$

and $\boldsymbol{p}_0$ is the foot position when the leg is extended downwards, $\boldsymbol{p}_{lb}$ and $\boldsymbol{p}_b$ are the leg base and body positions, and the subscript `td` denotes the value at the touch down time. Specifically, a minimum enclosing circle problem is formulated for the leg base positions for each stance phase and solved by the Welzl's algorithm [18] to compute $\boldsymbol{p}_{\text{center}}$, which ensures foothold reachability. Offset terms $\boldsymbol{p}_{\text{vel}}$ and $\boldsymbol{p}_{\text{centrifugal}}$ based on velocity and angular velocity [17] tracking are added to the nominal foot position to minimize undesired moments caused by GRFs during agile motion. For the repositioned foothold, $\boldsymbol{R}_v$ is the rotation matrix that rotates the nominal leg extension, which is designed to maximize the possible supporting polygons when landing. The current method simply rotates the nominally downwards leg towards the velocity direction to the joint limit. An optional approach is to parameterize the desired location by referring to the capture points discussed in [19].

D. Tail Control

The main goal of the tailed quadruped robot control is to utilize the tail to complement the kinematically restricted legs under hybrid perturbations. It requires the controller to account for the nonlinear dynamics of the tail and enables nonholonomic tail behavior in SO(3). As discussed in our previous work [20], it is not desirable to force the legs and tail to use the same controller. We can apply sequential distributed NMPC [21] to decouple tail and leg control and make it compatible with existing leg controllers.

We also propose a novel warmstart technique to decouple the time scale between NMPC update rate and finite element

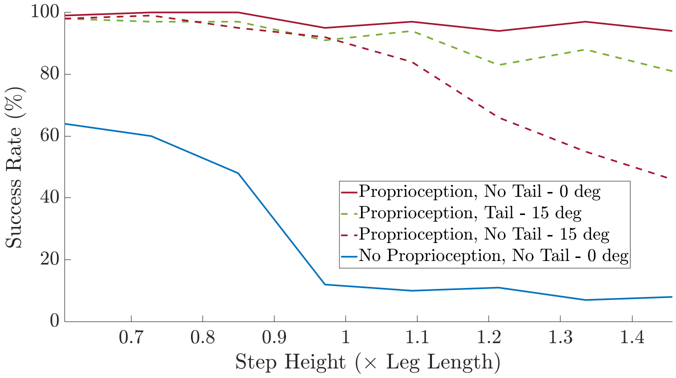

Fig. 5. Batch simulation success rate statistics under different initial conditions and environments.

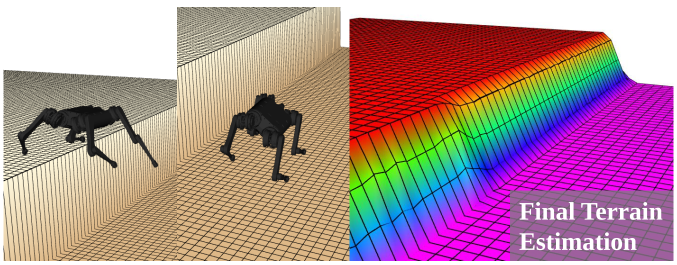

Fig. 6. Control behavior and terrain estimation results of the proposed proprioception-based gait planner when a robot walks sideways down an unforeseen cliff. Left: Proprioception recognizes the loss of contact, and the gait planner modifies the contact schedule and repositions the legs. Middle: With modified gait and controls, the robot lands in an appropriate posture. Right: The final terrain estimate successfully models stepped terrain.

discretization, and allow asynchronous solving of leg and tail NMPCs. Specifically, we make the duration of the first finite element adjustable to allow subsequent finite elements to remain aligned with the discretized collocation points. This decouples the time scale between the NMPC update rate and the finite element discretization. In this way, we can start solving the NMPC at any time. This allows both reaction to the latest state estimates or proprioception updates, and the tail and leg NMPC to run asynchronously. In addition, since the NMPC problem can be solved multiple times between the two collocation points, modeling error accumulation is greatly reduced, and the solving efficiency under warm starting is also greatly improved.

III. Experimental Evaluation

We measure performance by simulating a quadruped walking over an unknown elevation change over multiple trials (N=100) in Gazebo [22]. To measure the stability of the system under linear and angular disturbances, we initialize the robot with 0 and 15 degrees of roll and vary the elevation change from 25 to 60 cm, with success rate statistics shown in Fig. 5.

First, we compare the proposed proprioception-based gait planner with no proprioception assuming flat ground and a nominal contact schedule. The success rate without the proposed proprioception-based gait planner drops significantly

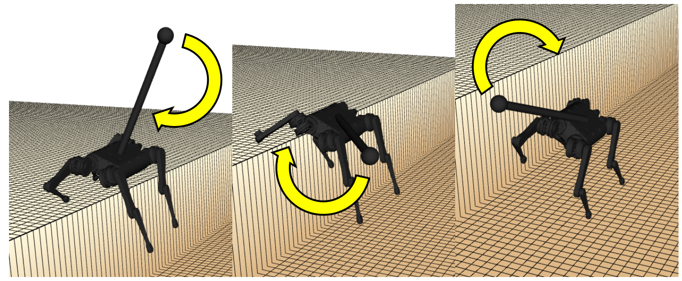

Fig. 7. Tailed robot control behavior with NMPC when falling off a cliff. The NMPC applies a conical motion that enables the tail to apply continuous torque to the body without hitting the joint limit.

with increasing cliff height. This is because the contact schedule is increasingly deviated from nominal, exiting the basin of attraction for the stability of the controller. On the other hand, using a proprioception-based gait planner, the robot can robustly walk over ledges up to 1.5 times the leg length with a success rate of over 94%. Higher altitude changes cause the motors to saturate, so even if the gait planner can maintain orientation, the robot will still fail due to too much potential energy. As shown in Fig. 6, it successfully identified contact deviation and responded by adjusting the contact schedule and repositioning the foot. The final terrain estimation also shows successful reconstruction of the step-like height drop, so the robot can be controlled and planned accordingly to adapt to the terrain.

Second, we examine the performance gains enabled by the tail. Experiments were set up with additional prior orientation errors comparing the success rates of proprioceptive-based gait planners with and without tails. The results show that even with the proposed gait planner, the success rate without the tail drops significantly with increasing stair height. This is because the initial angular disturbance rotates the robot towards the limit of the leg's abduction joint. When the tail was included, the success rate increased from 46% to 81% in the worst case. As shown in Fig. 7, it successfully augments the legs by rejecting angular errors and providing more kinematically feasible configurations for the legs. NMPC achieves this by enabling non-holonomic motion of the tail, specifically through a conical trajectory which can achieve prolonged maneuvers with limited angular deflection. This behavior has been described in [12], which has been shown to provide sustained torque and resembles a cheetah's tail motion.

IV. Conclusion and Future Work

We demonstrate extreme terrain traversal by quadrupeds through proprioception and tail control. Online updates of contact schedules, footholds, and terrain by proprioception improve robustness against unknown elevation changes. Tail control improves performance under angular perturbations by complementing the underactuated legs. Future work includes accounting for uncertainty in body state and contact estimation, extending the experiments to other types of terrain, and evaluating on hardware.

References

[1] W. Christensen, "The Physical Cat," in *Outwitting Cats: Tips, Tricks and Techniques for Persuading the Felines in Your Life That What You Want Is Also What They Want*. Globe Pequot, 2004, pp. 22–45.

[2] S. Shield, R. Jericevich, A. Patel, and A. Jusufi, "Tails, Flails, and Sails: How Appendages Improve Terrestrial Maneuverability by Improving Stability," *Integrative and Comparative Biology*, vol. 61, no. 2, pp. 506–520, 05 2021.

[3] A. Büschges, "Sensory Control and Organization of Neural Networks Mediating Coordination of Multisegmental Organs for Locomotion," *Journal of Neurophysiology*, vol. 93, no. 3, pp. 1127–1135, 2005.

[4] L. Keselman, J. Iselin Woodfill, A. Grunnet-Jepsen, and A. Bhowmik, "Intel RealSense Stereoscopic Depth Cameras," in *IEEE Conference on Computer Vision and Pattern Recognition Workshops*, July 2017.

[5] T. Miki, J. Lee, J. Hwangbo *et al.*, "Learning robust perceptive locomotion for quadrupedal robots in the wild," *Science Robotics*, vol. 7, no. 62, 2022.

[6] G. Bledt, P. M. Wensing, S. Ingersoll, and S. Kim, "Contact model fusion for event-based locomotion in unstructured terrains," in *IEEE International Conference on Robotics and Automation*, 2018, pp. 4399–4406.

[7] A. Seyfarth, H. Geyer, and H. Herr, "Swing-leg retraction: a simple control model for stable running," *Journal of Experimental Biology*, vol. 206, no. 15, pp. 2547–2555, 08 2003.

[8] C. Hubicki, J. Grimes, M. Jones *et al.*, "ATRIAS: Design and validation of a tether-free 3D-capable spring-mass bipedal robot," *The International Journal of Robotics Research*, vol. 35, no. 12, pp. 1497–1521, 2016.

[9] J. Zhu, N. J. Kong, G. Council, and A. M. Johnson, "Hybrid Event Shaping to Stabilize Periodic Hybrid Orbits," in *IEEE Intl. Conference on Robotics and Automation*, Philadelphia, PA, May 2022.

[10] R. Briggs, J. Lee, M. Haberland, and S. Kim, "Tails in biomimetic design: Analysis, simulation, and experiment," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2012, pp. 1473–1480.

[11] R. T. Fawcett, A. Pandala, J. Kim, and K. Akbari Hamed, "Real-Time Planning and Nonlinear Control for Quadrupedal Locomotion With Articulated Tails," *Journal of Dynamic Systems, Measurement, and Control*, vol. 143, no. 7, 02 2021.

[12] A. Patel and E. Boje, "On the Conical Motion of a Two-Degree-of-Freedom Tail Inspired by the Cheetah," *IEEE Transactions on Robotics*, vol. 31, no. 6, pp. 1555–1560, 2015.

[13] J. Norby, Y. Yang, A. Tajbakhsh *et al.*, "Quad-SDK," 2022. [Online]. Available: https://github.com/robomechanics/quad-software

[14] P. Fankhauser and M. Hutter, "A Universal Grid Map Library: Implementation and Use Case for Rough Terrain Navigation," in *Robot Operating System (ROS) – The Complete Reference (Volume 1)*, A. Koubaa, Ed. Springer, 2016, ch. 5.

[15] T. Homberger, L. Wellhausen, P. Fankhauser, and M. Hutter, "Support Surface Estimation for Legged Robots," in *IEEE International Conference on Robotics and Automation*, 2019, pp. 8470–8476.

[16] M. H. Raibert, *Legged robots that balance*. MIT press, 1986.

[17] D. Kim, J. Di Carlo, B. Katz *et al.*, "Highly dynamic quadruped locomotion via whole-body impulse control and model predictive control," *arXiv preprint arXiv:1909.06586*, 2019.

[18] E. Welzl, "Smallest enclosing disks (balls and ellipsoids)," in *New results and new trends in computer science*. Springer, 1991, pp. 359–370.

[19] J. Pratt, J. Carff, S. Drakunov, and A. Goswami, "Capture point: A step toward humanoid push recovery," in *IEEE-RAS International Conference on Humanoid Robots*, 2006, pp. 200–207.

[20] Y. Yang, J. Norby, J. K. Yim, and A. M. Johnson, "Improving Tail Compatibility Through Sequential Distributed Model Predictive Control," in *RSS Workshop on Software Tools for Real-Time Optimal Control*, July 2021.

[21] J. Liu, X. Chen, D. Muñoz de la Peña, and P. D. Christofides, "Sequential and iterative architectures for distributed model predictive control of nonlinear process systems," *AIChE Journal*, vol. 56, no. 8, pp. 2137–2149, 2010.

[22] N. Koenig and A. Howard, "Design and use paradigms for Gazebo, an open-source multi-robot simulator," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, vol. 3, 2004, pp. 2149–2154 vol.3.