

MIT Open Access Articles

Online Optimal Landing Control of the MIT Mini Cheetah

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: S. H. Jeon, S. Kim and D. Kim, "Online Optimal Landing Control of the MIT Mini Cheetah," 2022 International Conference on Robotics and Automation (ICRA), Philadelphia, PA, USA, 2022, pp. 178-184.

As Published: 10.1109/icra46639.2022.9811796

Publisher: IEEE

Persistent URL: <https://hdl.handle.net/1721.1/153616>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike



Real-time Optimal Landing Control of the MIT Mini Cheetah

Se Hwan Jeon¹, Sangbae Kim¹, and Donghyun Kim²

Abstract—Quadrupedal landing is a complex process involving large impacts, elaborate contact transitions, and is a crucial recovery behavior observed in many biological animals. This work presents a real-time, optimal landing controller that is free of pre-specified contact schedules. The controller determines optimal touchdown postures and reaction force profiles and is able to recover from a variety of falling configurations. The quadrupedal platform used, the MIT Mini Cheetah, recovered safely from drops of up to 8 m in simulation, as well as from a range of orientations and planar velocities. The controller is also tested on hardware, successfully recovering from drops of up to 2 m.

I. INTRODUCTION

Safe recovery from planned drops or unexpected falls is one of the most crucial features of animals and is beneficial both in navigating challenging terrain and preventing significant damage in the case of unexpected drops. Furthermore, a robust landing controller opens up the possibility of deploying quadruped robots directly into harsh environments where supervision can be difficult. However, in contrast to the progress in quadrupedal locomotion on robotic platforms [1], [2], there has been relatively little work exploring how robots address significant changes in elevation or react to falling. Several works have explored controlling a body’s orientation while rotating, but do not focus on the reaction force profiles and poses needed to safely recover from the fall [3], [4].

A common approach for landing actuated systems is to control contact point impedances with some additional feedforward reaction forces. Lynch et al. describes the so-called ‘soft landing problem’, exploring control methods to minimize foot penetration depth into a surface from a fall with prismatic actuation, but discussion was limited to single impacts on a body without rigid contacts [5]. Similar to quadrupeds, [6] explores actuated landing gear for rotorcraft and finding optimal impedances to mitigate impacts, but considers only horizontal landings. Although previous works demonstrated the effectiveness of impedance control, they did not capture all detailed motion included in landing behavior such as landing posture, reaction force profile, contact location, and sequence. Without taking into account these complexities, we cannot fully utilize the robot hardware’s capability to achieve smooth and safe landings.

For motion planning problems consisting of multiple, complex decision processes, trajectory optimization methods

¹Se Hwan Jeon and Sangbae Kim are with the Department of Mechanical Engineering, Massachusetts Institute of Technology, Cambridge, MA 02139, USA sehwan@mit.edu, sangbae@mit.edu

²Donghyun Kim is with the Department of Mechanical Engineering, University of Massachusetts Amherst, Amherst, MA 01003, USA donghyunkim@cs.umass.edu




Fig. 1. **High speed landing of quadruped robot.** The proposed controller enables the robot to safely land after falling from a high place. (a) In the simulation, we accomplished safe landing from 2.5m-high free drop. (b) In hardware experiments, we demonstrated 2m-high free drop.

have been widely used. It is worth noting that automatic contact sequence selection is crucial to controlling an landing because feasible contact sequences could be entirely different depending on the orientation of the body at touchdown. This timing becomes especially important at higher velocities. Previous work by Winkler et al. parameterizes the gait to optimize for the timings of the contact schedule, but the complexity of the formulation makes it difficult to solve for real-time applications [7]. Similarly, the work of [8] shows realistic, dynamic behaviors independent of prescribed contacts, but is likewise intractable for real-time performance.

Several other approaches such as reinforcement learning or model-predictive control (MPC) present interesting landing behaviors also demonstrated on hardware. [9] demonstrated planar landing and airborne orientation control on the SpaceBok quadruped in a ‘‘low-gravity’’ environment, but it is unclear how easily the algorithm could be applied for real-world, 3D conditions with more dramatic impacts and inertial effects. While planar landing and jumping was also demonstrated in [10] and [11], touchdown was made without considering optimal touchdown positions or timings, and were from relatively low heights with little pitch or roll of the body. With a heuristics-guided MPC, [12] demonstrated dropping the MIT Mini Cheetah from around 1.5 m, but did not investigate landing with changes in orientation or body velocities. Simple joint or Cartesian impedance control

can often be enough to stabilize a robot’s motion from small drops, but fail consistently at higher velocities or with significant orientations away from the horizontal.

In this work, we use a nonlinear trajectory optimization including contact complementary constraints to find optimal landing postures and reaction force profiles. The formulation is based on [13], and we modified the algorithm to accommodate actuator torque and leg kinematic limits. The optimization is solved while the robot is airborne in a MPC fashion at approximately 10 Hz. Once touchdown is detected, the optimized trajectory is then tracked with a whole-body impulse controller [14]. In our high-fidelity physics simulation environment, the Mini Cheetah robot recovered from drops of up to 8 m high, angular velocities between -0.5 and 0.5 rad/s, horizontal velocities ranging from -1.5 to 1.5 m/s, pitch orientations ranging from $-\frac{\pi}{3}$ to $\frac{\pi}{3}$, and roll orientations ranging from $-\frac{\pi}{6}$ to $\frac{\pi}{6}$. In hardware tests, the Mini Cheetah was able to land successfully from a height of 2 m.

The contribution of this work is two-fold: 1) we formulate a real-time landing control framework that finds optimal contact locations and timings from various orientations and body velocities, and 2) we demonstrate successful landings in simulation from significant heights and orientations and in hardware from heights of up to 2 m, as shown in Figure 1.

II. SYSTEM OVERVIEW

The MIT Mini Cheetah has a mass of approximately 9 kg with 12 modular actuators (ab/ad, hip, and knee for each of its four legs). Each actuator is capable of producing a maximum torque of 17 N m and a continuous torque of 6.9 N m [15]. Because the legs make up less than 10% of its mass, the dynamics of the Mini Cheetah can be approximated as a single rigid body model (SRBM) given as

$$m\ddot{\mathbf{p}} = \sum_{i=1}^{n_c} \boldsymbol{\lambda}_i - \mathbf{f}_g, \quad (1)$$

$$\frac{d}{dt}(\mathbf{I}\boldsymbol{\omega}) = \sum_{i=1}^{n_c} (\mathbf{r}_i - \mathbf{p}) \times \boldsymbol{\lambda}_i, \quad (2)$$

where m is the mass, n_c is the number of contacts, \mathbf{p} , \mathbf{r}_i , $\boldsymbol{\lambda}_i$, and \mathbf{f}_g are the vectors for the body position, foot positions, reaction forces, and gravitational force, expressed in the world frame. $\mathbf{I} \in \mathbb{R}^{3 \times 3}$ and $\boldsymbol{\omega}$ are the rotational inertia tensor of the body and angular velocity of the body respectively, expressed in the body frame.

III. FRAMEWORK

The landing controller in this work consists of three major components: nonlinear model predictive control, a trained neural network that outputs initial guesses for trajectories, and a whole-body controller. Based on its current state estimate, a nonlinear program (NLP) solves for dynamically feasible landing trajectories. The optimization is “warm started” with a guess for the trajectories output by the trained neural network to improve convergence speed. Finally, when

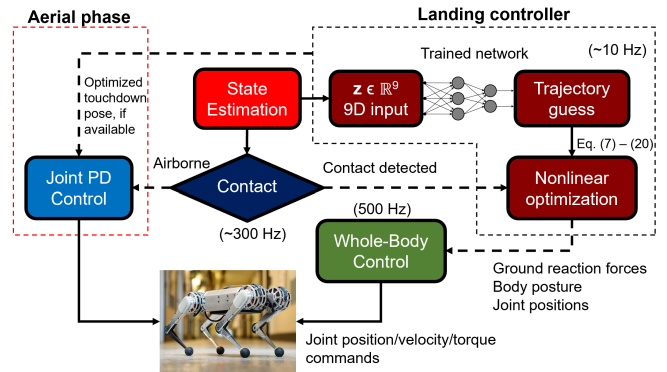


Fig. 2. Overview of the components of the landing controller. If contact is detected, the latest optimization solution is tracked with whole-body control. Otherwise, the robot maintains a nominal or optimized pose with joint PD control while airborne.

touchdown is detected, the optimized trajectory is used as a reference and tracked by a whole-body controller. An overview of this framework is shown in Figure 2.

A. Optimization Formulation

The controller is formulated as a direct transcription, kinodynamic trajectory optimization over N timesteps, similar to [16] and [17] with three key modifications for real-time performance: the use of a simpler model, assuming no-slip conditions in the CCCs, and reducing the number of decision variables necessary for the optimization.

The centroidal dynamics of a system considers the aggregate effects of each rigid link projected onto its center of mass frame, as detailed in [18]. However, given the low inertia limbs of many quadrupedal platforms and relatively small swinging movements during, this contribution is ignored, and the SRBM is used instead. This simplification significantly simplifies the dynamic constraints on the optimization.

With significant roll or pitch, it is difficult to determine how to prescribe appropriate contact schedules or sequences for landing, especially with the short timeframes high-speed landings would involve. However, the contact complementarity constraints allow these contact schedules to be directly optimized over without using mixed-integer formulations to characterize making and breaking contact. They are a set of conditions that require *either* the vertical ground reaction forces *or* the foot’s height in the world frame to be zero when the other is positive. While the CCCs outlined in [13] consider the full Coulomb friction model with sliding, this formulation only considers no-slip constraints, expressed as

$$\phi_i(\mathbf{q}_b, \mathbf{q}_j) \geq 0 \quad (3)$$

$$\lambda_z \geq 0 \quad (4)$$

$$\lambda_z \phi_i(\mathbf{q}_b, \mathbf{q}_j) \leq \epsilon, \quad (5)$$

where $\phi_i(\mathbf{q}_b, \mathbf{q}_j)$ is the signed distance to the ground from the i^{th} foot, \mathbf{q}_b and \mathbf{q}_j are the floating base and joint generalized coordinates respectively, and ϵ is a slack parameter to encourage convergence. The no-slip condition is enforced

with

$$(\mathbf{r}_{k+1} - \mathbf{r}_k)\lambda_{z,k} = 0, \quad (6)$$

where k is a timestep index in the optimization. This modification of the CCCs allows for more consistent and faster convergence to solutions, a key requirement for real-time performance.

To further reduce the complexity of the problem, the joint velocities are also excluded as decision variables. We post-process the optimized joint angle trajectories to find reasonable approximations for the joint velocities. This likewise improves the convergence and solve times of the NLP.

The full optimization formulation can be then be given as

$$\min_{\mathbf{x}, \mathbf{u}, \mathbf{q}_j} \|(\mathbf{x}_N - \mathbf{x}_{ref})\|_Q^2 \quad \text{s.t.} \quad (7)$$

$$\text{(Dynamics)} \quad \dot{\mathbf{x}}_{k+1} = f(\mathbf{x}_k, \mathbf{u}_k), \quad (8)$$

$$\text{(Initial conditions)} \quad \mathbf{x}(0) = \mathbf{x}_0, \quad (9)$$

$$\mathbf{r}(0) = \mathbf{r}_0, \quad (10)$$

$$\text{(Contact constraints)} \quad CCC(\lambda_i, \mathbf{r}_{i,k}) \quad (11)$$

$$(\mathbf{r}_{k+1} - \mathbf{r}_k)\lambda_{z,k} = 0, \quad (12)$$

$$\lambda_{x,y} \in \mathcal{F}(\mu, \lambda_z), \quad (13)$$

$$\text{(Kin. constraints)} \quad \mathbf{r}_{i,k}^T \mathbf{r}_{i,k} \leq l_{max}^2, \quad (14)$$

$$\mathbf{r}_{i,k} - g(\mathbf{q}_j) = 0, \quad (15)$$

$$\text{(Force limits)} \quad |J^T \boldsymbol{\lambda}_k| \leq \boldsymbol{\tau}_{max}, \quad (16)$$

$$\text{(Bounds)} \quad \mathbf{r} \in \mathcal{R}_{x,y,z}(\mathbf{x}_i), \quad (17)$$

$$\mathbf{x}_i \in \mathcal{X}, \quad (18)$$

$$\mathbf{x}(t_f) \in \mathcal{X}_{term}, \quad (19)$$

$$\mathbf{q}_{j,i} \in \mathcal{Q}, \quad (20)$$

where the decision variables $X = [\Theta^T \mathbf{p}^T \boldsymbol{\omega}^T \dot{\mathbf{p}}^T]^T \in \mathbb{R}^{12}$ is the state of the SRBM with roll-pitch-yaw orientation, $U = [\mathbf{r}^T \boldsymbol{\lambda}^T]^T \in \mathbb{R}^{24}$ is the inputs, and $\mathbf{q}_j \in \mathbb{R}^{12}$ is the joint positions of the legs. The foot Jacobians are represented by J , Q is a weighting matrix, μ is the friction coefficient, $CCC(\lambda_i, \mathbf{r}_{i,k})$ is the set of contact complementarity conditions, and $f(\mathbf{x}_k, \mathbf{u}_k)$, $g(\mathbf{q}_j)$, $\mathcal{F}(\mu, \lambda_z)$, and $\mathcal{R}(\mathbf{x}_i)$ are functions for the dynamics of the SRBM, the forward kinematics of the feet, friction pyramid constraints, and foot kinematic box constraints respectively. The optimization variables are bounded by l_{max} , $\boldsymbol{\tau}_{max}$, \mathcal{X} , and \mathcal{Q} , which are the maximum leg length, torque limits, state limits, and joint limits respectively. The dynamics of the system are propagated through the N timesteps with Euler forward integration. An example of an optimized trajectory in the MATLAB environment can be seen in Figure 3. Additionally, instantaneous changes in velocity due to impact are not considered. Due to the light, low-inertia limbs of the Mini Cheetah, it was assumed that impact forces could be ignored at touchdown.

Discretization and Initial Conditions: To reduce solve times, it was important to keep the number of timesteps N small. However, certain phases of the trajectory, such as directly after impact, requires high resolution for faithful tracking and smoother force profiles. This tradeoff was




Fig. 3. **Landing trajectory optimizations visualized in MATLAB.** The optimization scheme is able to find solutions involving a variety of different contact timings to stabilize the body to a nominal resting height.

approached by implementing unevenly discretized timesteps for the phases of the trajectory. The timesteps nears the final phase of the landing trajectory are larger in comparison to the period directly after impact. This allows the optimization to find smaller forces over longer time periods to stabilize the robot's state.

The final landing height of the quadruped is unknown, so its initial height is set to "expect" a touchdown immediately at all times. With constant updates from state estimation, this makes the controller independent of the global height of the robot relative to the ground plane it is landing on. The initial height of the robot is calculated to be

$$p_{z,hmin} = \min(S_z^B R_0 \mathbf{p}_{h,i} \forall i) \quad (21)$$

$$p_{z,0} = l_{max} + |p_{z,hmin}| + |v_{z,0} \Delta t_0|, \quad (22)$$

where ${}^B R_0$ is the rotation matrix from the initial roll-pitch-yaw orientation of the SRBM, $\mathbf{p}_{h,i}$ is the vector from the COM to the i^{th} hip, S_z is a selection matrix to find the z-component of a position vector, $p_{z,hmin}$ is the lowest height of the hips of the quadruped, and Δt_0 is the duration of the first timestep. While the initial angular velocity of the robot would also contribute to the orientation of the robot at touchdown, small angular velocities are assumed for simplicity.

This offset ensures that the lowest foot is able to reach the ground only *after* the first timestep from the forward Euler integration $p_{z,1} = p_{z,0} + \Delta t_0 v_{z,0}$, as shown in Figure 4. While the height of the SRBM could be positioned so that feet begin in contact with the ground as in [19], this prevents the NLP from finding optimal footstep locations relative to the body. With this formulation, the trajectory is always "expecting" to touchdown immediately after the first timestep based on its current velocity, and by solving this problem in a model-predictive fashion, becomes independent of the global coordinates of the body.

Cost: The cost in Eqn (7) consists only of the quadratic penalty between the final state of the optimized trajectory, X_N , and a desired final position, X_{ref} . The weighting matrix Q was only non-zero for entries corresponding to velocities, orientation, and a final desired resting height. It was found that penalizing only the terminal state significantly improved convergence speed, and, in some cases, solution quality. Introducing running costs on applied forces or states often




Fig. 4. **The initial height setting given by Eq. 22.** This guarantees that the robot is unable to make contact during the first timestep and to touchdown in the second, allowing for contact locations and poses to be optimized over. If the optimal pose from a previous solution is unavailable, it is initialized with some nominal pose.

generated impractical trajectories, where the knees of the Mini Cheetah would penetrate the ground, or significant forces would be commanded before touchdown. In [7], costs similarly seemed to increase solve times and the overall complexity of the NLP.

Without thorough biomechanical study, it is difficult to ascertain which quantities, if any, are best to optimize for landing. This cost formulation allows for more flexibility in finding landing trajectories instead of imposing costs that may guide the optimization to infeasible solutions. By only penalizing the final desired state, feasible landing trajectories can be generated quickly and reliably.

Kinematic Box Limits: From initial trajectory optimizations, it was observed that the touchdown angle of the feet correlated strongly with the direction of the COM velocity in the world frame, as shown in Figure [FOOT ANGLE AND VEL PLOT]. Intuitively, it follows that in order to generate force to oppose some velocity, it would be efficient to align the moment arm of the contact point with the direction of that velocity as well. With this observation, a simple heuristic constraint was added to the optimization that would adjust the kinematic box limits of the quadruped’s feet based on its current velocity similar to [7], as shown in Figure 5. The equations to govern the size of the box limits are given as

$$\mathcal{R}_{x,y,z}(\mathbf{x}_i) = \{r | r \leq d \frac{v}{v_{max}}\}, \quad (23)$$

where d is the maximum limit of the box in x , y , and z directions, v is the x , y , or z component of the body velocity expressed in the body frame, and v_{max} determines the speed limit at which the box will extend to the kinematic limits of the leg. The values for d and v_{max} were determined experimentally and adjusted based on the solutions of the trajectory optimizations. By limiting the range of locations the feet could be positioned relative to the body, kinematic limits can be implicitly enforced, improving the convergence of the optimization.

B. Supervised Learning Framework

As studied in [20], the initial guess provided to a nonlinear optimization is crucial. Depending on its quality, the

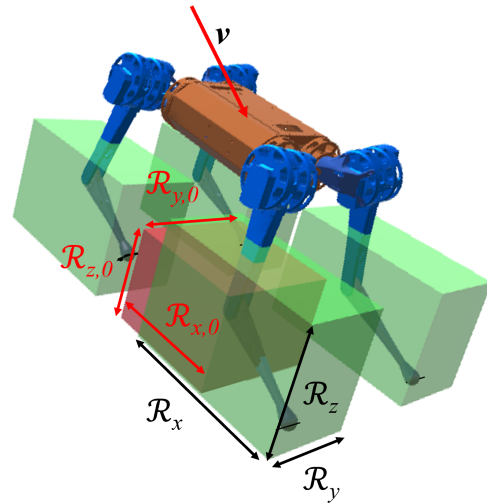


Fig. 5. **Visualized kinematic box limits.** The red box marks the *default* foot position bounds, the dimensions of which are determined by \mathcal{R}_0 . The kinematic bounds in the optimization (green) adapt to the body velocity v expressed in the body frame and falling conditions as necessary to encourage convergence, the lengths of which are denoted with \mathcal{R} .

optimization could be attracted to undesirable local minima or fail to find a feasible solution entirely. To improve convergence speed as well as solution quality, a neural network is trained in a supervised learning fashion to warm start the nonlinear trajectory optimization outlined in Section III. The inputs to the optimization are only the initial orientation, velocity, and angular velocity of the SRBM, which will be denoted $\mathbf{z} = [\Theta^T \ \omega^T \ \dot{\mathbf{p}}^T]^T \in \mathbb{R}^9$. Random samples are taken from selected ranges of \mathbf{z} and optimal trajectories are generated for each initial falling condition to train the network. Table 1 details the selected ranges for each element in \mathbf{z} .

TABLE I
RANGE OF SAMPLED INPUT SPACE FOR SUPERVISED LEARNING

\mathbf{z}	Θ_x	Θ_y	Θ_z	ω (r/s)	$v_{x,y}$ (m/s)	v_z (m/s)
Min.	$-\pi/4$	$-\pi/3$	$-\pi/2$	-1	-1.5	-6
Max.	$\pi/4$	$\pi/3$	$\pi/2$	1	1.5	-2

The generated solutions were post-processed so that trajectories with ground penetration or resteps were removed. The data was normalized with respect to the mean and standard deviation of each variable in the output trajectory. A neural network with 2 hidden layers with 128 nodes each was trained to generate output trajectories to serve as a warm start to the optimization.

C. Whole-body Controller

The whole-body controller used in this work is a quadratic program (QP) detailed in [14] with a control bandwidth of 500 Hz. The fast update rates allow for reference trajectories to be tracked closely and stabilized with small deviations in expected states. An optimized landing trajectory is tracked by the WBC when touchdown is detected. Because timesteps

are discretized unevenly, interpolated values of the trajectory are calculated to be tracked by the WBC.

IV. IMPLEMENTATION

The controller was developed in MATLAB with the CasADi symbolic framework and the `spatial v2` package from Roy Featherstone [21], [22]. With Casadi’s internal code generation features, the optimization was exported to C++ where the Lightweight Communications and Marshalling package (LCM) was used to communicate between the optimization and the simulation environment and the robot hardware asynchronously [23]. The commercial nonlinear optimization solver Artelys Knitro was used to solve the problem online [24]. While the solver IPOPT demonstrated similar solve speeds, Knitro was far more consistent in finding solutions from the range of initial falling conditions given.

The neural network was developed with PyTorch and similarly exported to C++ with the LibTorch packages. Without the overhead of the MATLAB interface, the optimizations were able to be solved faster as well.

The optimization is run as a process separate from the main body of the software on the Mini Cheetah. Requests for optimizations are sent via LCM from the robot’s state estimator and received when a solution is found. If a solution is not found within 300 ms, a new request is made and the optimization is restarted. Decoupling the optimization from the main processes on Mini Cheetah in this manner allows for asynchronous, model-predictive control.

Touchdown is detected when a threshold on the joint velocities of the legs is exceeded. The latest feasible trajectory is then interpolated and tracked with the QP-based whole-body controller from [14]. While there have been learned and probabilistic approaches to detecting contact proprioceptively [25], [26], joint encoder feedback is used instead. Due to its low response time of around 3 ms and the impulsive nature of landing that could be difficult to learn, encoder data was used as a touchdown trigger for its simplicity.

V. RESULTS

A. Trained Trajectory Generation

The network was trained with roughly 1500 samples of trajectory data over a range of falling conditions. The state space was uniformly sampled over, but the majority of the successful landings tended to be near the horizontal body posture, as landing in a near-horizontal position is the simplest landing to stabilize. Samples from the space of successful landing conditions are shown in Figure 6.

B. Real-time Performance

By warm starting the optimization with the trajectories generated by the neural net, there is a significant decrease in computation time that allows for real-time performance to be possible. On a typical desktop computer in the MATLAB environment, \mathbf{z} was randomly sampled across the same ranges as in Table 1 and optimized with the warm-start from the neural




Fig. 6. **Training data sampling and warm-start evaluation.** A: sampled space of falling conditions for training data. B: comparisons of solve times for neural net warm-started optimization, cold-started optimization, and two-stage warm-started optimization, from left to right.

net. This was compared with the solve times from cold-starts and a two-stage optimization process. The two-stage optimization involved solving the NLP outlined in Section III *without* kinematic decision variables, and using its solution as a warm start to the full kino-dynamic optimization. As shown in Figure 6, the neural-net warm started optimization showed speedups of roughly 5-10 times over the range of falling conditions. Even with the overhead of MATLAB, solve times were in the range of 5-8 Hz, allowing for real-time performance onboard the Mini Cheetah.

C. Simulation

The Mini Cheetah was initialized in various falling configurations (orientation and velocities) in the Robot-Software environment and simulated to test recovery behavior. The desired final state for $X_{ref} = [\Theta^T \mathbf{p}^T \boldsymbol{\omega}^T \dot{\mathbf{p}}^T]^T$ was set to be all zero except the height, which was set to 0.25 m. Zero entries in the weighting matrix Q , ensures no penalty was applied on the final yaw, p_x , or p_y .

The Robot-Software simulation evaluates its dynamics at 1000 Hz and includes motor constraints, rotor inertias, and impact calculations in its environment. A friction coefficient of $\mu = 0.75$ was used, and it was assumed that there was no state estimation error in the simulation. Torque and force plots for a landing trajectory from a 2.5 m high pitched landing with small lateral velocities are shown in Figure 7. While the impact on touchdown causes a high initial tracking error, this quickly goes to zero within the first 0.1 s, and both torque and force profiles are tracked closely. With the adjustments from the whole-body controller, the trajectories are well followed even with significant disturbances. Deliberate errors in velocity, orientation, or angular velocity were used as inputs to the optimization, but the controller was often able to stabilize to a nominal resting position with small errors in the expected falling conditions \mathbf{z} . The authors plan to more fully characterize the robustness of this controller in future work.

To test the limits of the hardware in simulation, the Mini Cheetah was dropped horizontally from a height of 8 m, corresponding to vertical velocities of roughly -12 m s^{-1} . The controller was able to stabilize the fall, but further work




Fig. 7. **Actuation limits during simulated landing.** A: desired (τ^*) and actual (τ) torques from the simulation environment. The black dotted lines represent the absolute torque limits of the individual motors. B: desired (λ^*) and actual (λ) ground reaction forces commanded by the whole-body controller. The Mini Cheetah is able to stabilize itself from significant drops while respecting the torque limits of the system.

must be done to investigate the effect of the impacts and post-impact velocities at these higher speeds.

Recovering from landing is far more sensitive to roll and lateral velocity than pitch and longitudinal velocity. This is intuitive given the design of the Mini Cheetah, where only the single ab/ad motor is dedicated to stabilizing significant disturbances in body roll and lateral velocity.

D. Hardware

To verify the controller on hardware and ensure safe landings, it was critical that the state estimation of the Mini Cheetah was as accurate as possible. However, there was severe and constant drift in state estimates once the robot is in the air because kinematics based estimation is not possible while airborne. With unreliable state estimation, falling conditions were instead hard coded to be approximately similar to the state of the robot at touchdown instead.

Several horizontal drops were performed from heights of up to 2 m successfully, as shown in Figure 1. Although hard-coded conditions were used as inputs to the optimization, the Mini Cheetah was released by hand, causing deviations from the expected falling state. Despite these errors, the Mini Cheetah was able to recover to a nominal resting height, as shown in the plots in Figures 8. The state estimator was paused in a resting state until touchdown was detected because of its drift, causing a significant difference between the estimated and actual velocities of the quadruped at touchdown, but the WBC is able to stabilize the robot about the planned trajectory.

Because of the unreliable state estimation, drops with significant pitch, roll, or lateral velocities were not tested to prevent damage to the robot. While simulation results suggest these landings would also be possible, we found that significant orientation errors in the state estimator quickly caused the controller to become unstable in the Robot-




Fig. 8. **Trajectory tracking on hardware.** Desired body posture (p^* , v^*) and body posture from state estimation (p , v). Note that the velocity discrepancy is due to pausing the state estimation until touchdown.

Software environment.

VI. CONCLUSIONS

This paper presents a control framework that is able to reason about optimal contact locations and timings in real-time, and demonstrates successful landings in both simulation and hardware. By modifying the contact constraints and model considered, dynamic trajectories can be planned independently from prescribed contact modes at real-time rates.

Future work will involve hardware upgrades to the Mini Cheetah, implementing onboard vision, and characterizing the robustness of the controller. Processor and power board upgrades to the Mini Cheetah will enable it to solve the presented formulation completely untethered and onboard, and with greater margin for motor failure. The authors also plan to integrate an event camera into the Mini Cheetah for more accurate localization and state estimation. Additionally, while it was shown that the controller was able to stabilize itself from small errors in the falling conditions, we have yet to characterize this. In the future, the robustness of the controller and its performance over uneven terrain will be explored in greater detail.

ACKNOWLEDGMENT

This work was supported by Naver Labs and the Centers for ME Research and Education at MIT and SUSTech. The authors would like to thank Elijah Stanger-Jones and Aditya Mehrotra for their assistance in hardware experimentation, and the other members of the Biomimetic Robotics Lab for their insight and support in this work.

REFERENCES

- [1] T. Dudzik, M. Chignoli, G. Bledt, B. Lim, A. Miller, D. Kim, and S. Kim, "Robust autonomous navigation of a small-scale quadruped robot in real-world environments," Oct. 2020. DOI: 10.1109/IROS45743.2020.9340701.
- [2] J. Lee, J. Hwangbo, L. Wellhausen, V. Koltun, and M. Hutter, "Learning quadrupedal locomotion over challenging terrain," *Science Robotics*, vol. 5, no. 47, eabc5986, 2020. DOI: 10.1126/scirobotics.abc5986.
- [3] J. T. Bingham, J. Lee, R. N. Haksar, J. Ueda, and C. K. Liu, "Orienting in mid-air through configuration changes to achieve a rolling landing for reducing impact after a fall," in *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2014, pp. 3610–3617. DOI: 10.1109/IROS.2014.6943068.

- [4] V. Kurtz, H. Li, P. M. Wensing, and H. Lin, *Mini cheetah, the falling cat: A case study in machine learning and trajectory optimization for robot acrobatics*, 2021. arXiv: 2109.04424 [cs.RO].
- [5] D. J. Lynch, K. M. Lynch, and P. B. Umbanhowar, "The soft-landing problem: Minimizing energy loss by a legged robot impacting yielding terrain," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 3658–3665, 2020. DOI: 10.1109/LRA.2020.2977260.
- [6] J. Kiefer, M. Ward, and M. Costello, "Rotorcraft Hard Landing Mitigation Using Robotic Landing Gear," *Journal of Dynamic Systems, Measurement, and Control*, vol. 138, no. 3, Jan. 2016, 031003. DOI: 10.1115/1.4032286. eprint: https://asmedigitalcollection.asme.org/dynamicsystems/article-pdf/138/3/031003/6121589/ds_138_03_031003.pdf. [Online]. Available: <https://doi.org/10.1115/1.4032286>.
- [7] A. Winkler, 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. PP, May 2018. DOI: 10.1109/LRA.2018.2798285.
- [8] I. Mordatch, E. Todorov, and Z. Popović, "Discovery of complex behaviors through contact-invariant optimization," *ACM Transactions on Graphics - TOG*, vol. 31, Jul. 2012. DOI: 10.1145/2185520.2185539.
- [9] N. Rudin, H. Kolvenbach, V. Tsounis, and M. Hutter, "Cat-like jumping and landing of legged robots in low gravity using deep reinforcement learning," *IEEE Transactions on Robotics*, pp. 1–12, 2021. DOI: 10.1109/TRO.2021.3084374.
- [10] Q. Nguyen, M. J. Powell, B. Katz, J. D. Carlo, and S. Kim, "Optimized jumping on the mit cheetah 3 robot," in *2019 International Conference on Robotics and Automation (ICRA)*, 2019, pp. 7448–7454. DOI: 10.1109/ICRA.2019.8794449.
- [11] M. Hutter, C. Gehring, M. Bloesch, M. Hoepflinger, C. Remy, and R. Siegwart, "Starleth: A compliant quadrupedal robot for fast, efficient, and versatile locomotion," Sep. 2012, pp. 483–490. DOI: 10.1142/9789814415958_0062.
- [12] G. Bledt, "Regularized predictive control framework for robust dynamic legged locomotion," Ph.D. dissertation, Jan. 2020. DOI: 10.13140/RG.2.2.24766.02883.
- [13] M. Posa, C. Cantu, and R. Tedrake, "A direct method for trajectory optimization of rigid bodies through contact," *The International Journal of Robotics Research*, vol. 33, no. 1, pp. 69–81, 2014. DOI: 10.1177/0278364913506757. eprint: <https://doi.org/10.1177/0278364913506757>. [Online]. Available: <https://doi.org/10.1177/0278364913506757>.
- [14] D. Kim, J. D. Carlo, B. Katz, G. Bledt, and S. Kim, "Highly dynamic quadruped locomotion via whole-body impulse control and model predictive control," *arXiv: 1909.06586*, 2019.
- [15] B. Katz, J. D. Carlo, and S. Kim, "Mini cheetah: A platform for pushing the limits of dynamic quadruped control," in *2019 International Conference on Robotics and Automation (ICRA)*, 2019, pp. 6295–6301. DOI: 10.1109/ICRA.2019.8793865.
- [16] H. Dai, A. Valenzuela, and R. Tedrake, "Whole-body motion planning with centroidal dynamics and full kinematics," in *2014 IEEE-RAS International Conference on Humanoid Robots*, 2014, pp. 295–302. DOI: 10.1109/HUMANOIDS.2014.7041375.
- [17] M. Chignoli, D. Kim, E. Stanger-Jones, and S. Kim, "The MIT humanoid robot: Design, motion planning, and control for acrobatic behaviors," *CoRR*, vol. abs/2104.09025, 2021. arXiv: 2104.09025. [Online]. Available: <https://arxiv.org/abs/2104.09025>.
- [18] D. Orin, A. Goswami, and S.-H. Lee, "Centroidal dynamics of a humanoid robot," *Autonomous Robots*, vol. 35, Oct. 2013. DOI: 10.1007/s10514-013-9341-4.
- [19] V. Samy, S. Caron, K. Bouyarmane, and A. Kheddar, "Post-impact adaptive compliance for humanoid falls using predictive control of a reduced model," in *2017 IEEE-RAS 17th International Conference on Humanoid Robotics (Humanoids)*, 2017, pp. 655–660. DOI: 10.1109/HUMANOIDS.2017.8246942.
- [20] N. Mansard, A. DelPrete, M. Geisert, S. Tonneau, and O. Stasse, "Using a memory of motion to efficiently warm-start a nonlinear predictive controller," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*, 2018, pp. 2986–2993. DOI: 10.1109/ICRA.2018.8463154.
- [21] J. A. E. Andersson, J. Gillis, G. Horn, J. B. Rawlings, and M. Diehl, "CasADi – A software framework for nonlinear optimization and optimal control," *Mathematical Programming Computation*, vol. 11, no. 1, pp. 1–36, 2019. DOI: 10.1007/s12532-018-0139-4.
- [22] R. Featherstone. (2021). "Spatial-v2 spatial vector and rigid-body dynamics software," [Online]. Available: <http://royfeatherstone.org/spatial/v2/> (visited on 09/12/2021).
- [23] A. S. Huang, E. Olson, and D. C. Moore, "Lcm: Lightweight communications and marshalling," in *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2010, pp. 4057–4062. DOI: 10.1109/IROS.2010.5649358.
- [24] R. H. Byrd, J. Nocedal, and R. A. Waltz, "Knitro: An integrated package for nonlinear optimization," in *Large-Scale Nonlinear Optimization*, G. Di Pillo and M. Roma, Eds. Boston, MA: Springer US, 2006, pp. 35–59. DOI: 10.1007/0-387-30065-1_4. [Online]. Available: https://doi.org/10.1007/0-387-30065-1_4.
- [25] G. Bledt, P. Wensing, S. Ingersoll, and S. Kim, "Contact model fusion for event-based locomotion in unstructured terrains," May 2018. DOI: 10.1109/ICRA.2018.8460904.
- [26] T.-Y. Lin, R. Zhang, J. Yu, and M. Ghaffari, "Deep multi-modal contact estimation for invariant observer design on quadruped robots," *CoRR*, vol. abs/2106.15713, 2021. arXiv: 2106.15713. [Online]. Available: <https://arxiv.org/abs/2106.15713>.