Robust Trajectory Optimization Over Uncertain Terrain With Stochastic Complementarity

Luke Drnach and Ye Zhao

***Abstract*—Trajectory optimization with contact-rich behaviors has recently gained attention for generating diverse locomotion behaviors without pre-specified ground contact sequences. However, these approaches rely on precise models of robot dynamics and the terrain and are susceptible to uncertainty. Recent works have attempted to handle uncertainties in the system model, but few have investigated uncertainty in contact dynamics. In this letter, we model uncertainty stemming from the terrain and design corresponding risk-sensitive objectives for contact-implicit trajectory optimization. In particular, we parameterize uncertainties from the terrain contact distance and friction coefficients using probability distributions and propose a corresponding expected residual minimization cost approach. We evaluate our method in three simple robotic examples, including a legged hopping robot, and we benchmark one of our examples in simulation against a robust worst-case solution. We show that our risk-sensitive method produces contact-averse trajectories that are robust to terrain perturbations. Moreover, we demonstrate that the resulting trajectories converge to those generated by a traditional, non-robust method as the terrain model becomes more certain. Our study marks an important step towards a fully robust, contact-implicit approach suitable for deploying robots on real-world terrain.**

***Index Terms*—Optimization and optimal control, motion and path planning, contact modeling.**

I. Introduction

TRAJECTORY optimization has become a powerful tool for designing dynamic motions for robots with nonlinear, hybrid, under-actuated dynamics and constraints [1]–[6]. Although impressive locomotion applications abound in the literature and transfer to hardware has been reliably demonstrated in laboratory conditions [7], successful application of trajectory optimization to locomotion in general, unstructured environments is still challenging: success depends critically on multiple factors including model fidelity, environmental uncertainty, and the ability to design effective closed-loop strategies for executing planned motions [3], [5]. As optimal strategies often lie on the boundary of the feasible region, errors in the dynamic model could result in the planned trajectory becoming dynamically infeasible. Additionally, unmodeled disturbances from the environment can introduce deviations from the nominal trajectory, which propagate through the dynamics and can result in large errors over time. While fast online re-planning and robust low-level control can aid in recovering from local disturbances, reasoning about robustness in the high-level planning could fundamentally improve the overall system performance.

Within trajectory optimization, designing reliable behaviors for dynamic robot locomotion tasks that require intermittent frictional contact has been a persistent challenge over the past few decades. Contact sequences and forces can be calculated during trajectory planning using contact-implicit trajectory optimization [4]; however, this method requires exact knowledge of the terrain geometry and friction coefficients beforehand. As friction coefficients require specialized sensors to estimate and real-world terrain geometry can be intractable to model outside the laboratory, the contact-implicit method becomes highly prone to errors and failures. Errors in modeling the friction characteristics could cause a robot to slip, and errors in modeling terrain geometry could cause the robot to trip, both of which could result in a fall. We hypothesize that the failure to explicitly account for uncertainties and feedback during trajectory design is a key contributor to slow progress in translating trajectory optimization, and in particular contact-implicit trajectory optimization, research results into a depolyable technology.

Our study takes one step toward addressing this problem by deriving a risk-sensitive variant of contact-implicit trajectory optimization. We develop objective functions derived from statistics related to the traditional complementarity constraints for contact and reason about the robustness by comparing trajectories generated by our robust method to those generated using the conventional complementarity constraint method. To contribute specifically to the field, we:

- Include parametric models of uncertainty in the friction coefficient and in the contact distance into contact-implicit trajectory optimization.
- Develop risk-sensitive objectives that produce contact-averse trajectories which are robust to perturbations in terrain parameters when the terrain model is uncertain.
- Demonstrate that our method represents a smooth generalization of the traditional complementarity constraints in

Manuscript received September 24, 2020; accepted January 13, 2021. Date of publication February 1, 2021; date of current version February 17, 2021. This letter was recommended for publication by Associate Editor A. Quattrini Li and Editor L. Pallottino upon evaluation of the reviewers' comments. This work was supported by National Science Foundation under Grant DGE-1650044. *(Corresponding author: Ye Zhao.)*

Luke Drnach is with the School of Electrical, and Computer Engineering, Institute for Robotics, and Intelligent Machines, Georgia Institute of Technology, Atlanta, GA 30318 USA, and also with the Laboratory for Intelligent Decision and Autonomous Robots, G. W. Woodruff School of Mechanical Engineering, Georgia Institute of Technology, Atlanta, GA 30332 USA (e-mail: ldrnach3@gatech.edu).

Ye Zhao is with the Laboratory for Intelligent Decision and Autonomous Robots, G. W. Woodruff School of Mechanical Engineering, Georgia Institute of Technology, Atlanta, GA 30332 USA (e-mail: ye.zhao@me.gatech.edu).

This letter has supplementary downloadable material available at https://doi.org/10.1109/LRA.2021.3056064, provided by the authors.

Digital Object Identifier 10.1109/LRA.2021.3056064

that it converges to a relaxed complementarity method as the contact parameters become certain.

We evaluate our framework in three examples and benchmark one of our examples against a worst-case robust approach in simulation. We show that the control trajectories resulting from our approach are robust to perturbations in the contact parameters, since uncertainty in the contact constraints is explicitly modeled. Although our work assumes the contact constraints are uncertain, and thus we cannot enforce the exact physical constraints, we show that the solution sets of our robust objective correspond with the solution of the complementarity constraints, and we prove this correspondence in limiting cases. Thus, trajectories generated under our robust objective may facilitate implementing robust motion plans on physical robots.

II. Related Work

A. Contact-Implicit Trajectory Optimization

Contact-implicit trajectory optimization includes contact forces as decision variables in an optimal control problem [4], [8]–[12]. The contact forces are governed by set of complementarity constraints [13], and the resulting direct transcription problem is solved through a large-scale nonlinear program. Compared to approaches with predefined contact sequences [14], [15], a remarkable advantage of this contact-implicit method lies in avoiding an exhaustive search of combinatorial contact mode possibilities, which are computationally prohibitive for contact-rich robotic systems. Building on top of this contact-implicit approach, our study focuses on reasoning about robustness to uncertainties with respect to contact surface geometry and friction properties.

B. Robust Trajectory Optimization

Reasoning about the robustness of trajectory optimization has been extensively explored in robotics [16]–[18]. One robust approach is ensemble contact-invariant optimization [3] which samples uncertain physical model parameters and generates a collection of specific model instances. Trajectories associated with each model instance are coupled via a penalty cost and a single nominal trajectory is generated with a notion of robustness. In more recent works [16], [19], uncertainty in friction coefficients has been addressed by updating model parameters from errors between planned motions and simulated or experimental motions; however, these methods require multiple physical interactions to improve the estimate of the friction coefficient, and early interactions can fail due to a lack of robustness. Differing from modeling parameter uncertainties or learning friction parameters, our study reasons about robustness to contact uncertainties, which is critical for safe contact-rich planning.

Risk-sensitive optimal control, a powerful approach to reason about robustness, employs high-order statistics in the cost function design [20]–[23]. A seminal work in [24] proposed a Linear-Exponential Gaussian algorithm which includes the high-order statistics by using the expectation of the exponential transformation of a performance index as the cost. In these works, uncertainty is assumed to enter through either the state estimation or control actuation, and the cost function is transformed to produce risk-sensitive behaviors. However, these works have yet to address uncertainty from constraints dealing with contact. Here we consider uncertainty arising from the contact model, which is normally included in trajectory optimization as complementarity constraints, and we derive additional cost terms to produce risk-sensitive behaviors.

C. Stochastic Complementarity Problems

One approach to handling uncertainty in complementarity constraints is to recast them as expected residual minimization (ERM) problems. The ERM formulation, which is a smooth alternative for both linear and nonlinear complementarity constraints (LCPs and NCPs), has been extensively investigated in the context of stochastic complementarity problems (SCPs) [25]. Smoothed residual functions are often introduced as approximations of the original constraints [26], and solutions to the ERM problem are robust in the sense that they have minimum sensitivity to random SCP parameter variations. Another approach is to cast the complementarity problem as a worst-case robust optimization, as in [27]. In the case of LCPs, the worst-case variant can be solved by a single convex program. Nevertheless, the convexity assumption is conservative since many robotic problems are inherently non-convex and nonlinear. Moreover, application of both the ERM and the worst-case methods to trajectory optimization has been largely under-explored. An initial effort applied the ERM framework to solve robotic problems with stochastic complementarity [28], where uncertainty is assumed to be derived from errors in state estimation. However, that work mainly applies the ERM method as a smoothing technique so the complementarity constraints could be included in indirect trajectory optimization. In our study, we further explore the ERM technique as a method for encoding uncertainty about the terrain model into direct trajectory optimization and explicitly analyze the robustness of the resulting trajectories.

III. Problem Formulation

A. Contact-Implicit Trajectory Optimization

Contact-implicit trajectory optimization solves for the states, controls, and contact force of a robot's motion with intermittent contact through the optimal control problem:

$$\min_{x,u,\lambda} \int_0^T L(x,u,\lambda)dt + L_F(x(T)) \tag{1a}$$

$$\text{s.t.} \begin{cases} M(q)\ddot{q} + C(\dot{q},q) = Bu + J_c^\top(q)\lambda & \text{(1b)} \\ x(0) = x_0, x(T) = x_f & \\ 0 \le \lambda_N \perp \phi(q) \ge 0 & \text{(1c)} \\ 0 \le \lambda_T \perp \gamma + J_T\dot{q} \ge 0 & \text{(1d)} \\ 0 \le \gamma \perp \mu\lambda_N - e^\top\lambda_T \ge 0 & \text{(1e)} \end{cases}$$

where $x = (q, \dot{q})$ is the state, q is t he system configuration, x_0 and x_f are the initial and final states respectively, and L and L_F are the running and terminal costs respectively. Eq. (1b) represents the rigid-body dynamics with mass matrix M, Coriolis and

conservative forces C, control selection matrix B, and contact Jacobian J_c. Eqs. (1c)–(1e) are the nonlinear complementarity constraints encoding the contact conditions. Eq. (1c) encodes a normal distance constraint, where λ_N is the normal force and $\phi(q)$ is the normal distance. Eq. (1d) encodes a constraint on the sliding velocity, where γ is a slack variable related to the magnitude of the sliding velocity, J_T is the tangential part of the contact Jacobian, and λ_T is the tangential contact force. Eq. (1e) encodes a linearized friction cone constraint, where μ is the coefficient of friction and e is a vector of 1 s. The shorthand $0 \leq a \perp b \geq 0$ denotes a complementarity constraint: $a \geq 0, b \geq 0, a^\top b = 0$.

Numerical methods have already been developed to solve the problem (1a)–(1e) using either direct [4], [8], [9] or indirect [28], [29] methods. In this work, instead of developing a more computationally efficient or more accurate high-order method, as was the case in [8], [29], and [9], our goal is to develop and evaluate a framework for including contact uncertainties. Thus, we used a direct transcription method to convert the continuous dynamics and costs into their discrete analogs. We evaluated the dynamics using backward Euler integration, enforced the contact constraints at the end of each interval, and used the following quadratic cost:

$$L(x, u, \lambda) = \frac{1}{2}\left((x - x_f)^\top Q(x - x_f) + u^\top R u\right).$$

B. Stochastic Complementarity Constraints

The preceding formulation assumes perfect knowledge of the contact parameters. If any of the terms in (1c)–(1e) are uncertain or random, then resolving the contact forces becomes a stochastic complementarity problem (SCP) [30]:

$$0 \leq z \perp F(z, \omega) \geq 0, \quad \omega \in \Omega \tag{2}$$

where ω represents a random quantity on probability space $(\Omega, \mathcal{F}, \mathcal{P})$ with given probability distribution $\mathcal{P}$, z is the decision variable, and $F(\cdot)$ is a vector-valued function. Because ω is stochastic, (2) is not well-defined and in general will not have a solution for all $\omega \in \Omega$. One approach is to replace the function F with its expected value:

$$0 \leq z \perp \mathbb{E}[F(z, \omega)] \geq 0 \tag{3}$$

The expected value method is largely equivalent to solving the deterministic problem at the mean value of F, and is not expected to be robust to random variations in the parameters.

C. Expected Residual Minimization

Theoretical works have studied robust solutions to Eq. (2) in both the case when F is affine [25], [26], [31] and in the case when F is nonlinear [30]. In these works, it is common to define a residual function ψ such that the residual is zero when the complementarity conditions are satisfied:

$$\psi(a, b) = 0 \iff a \geq 0, b \geq 0, a^T b = 0 \tag{4}$$

One common choices for the residual function is the "min" function $\psi_{\min}(a, b) = \min(a, b)$. Then, the expectation of the residual can be taken to form a deterministic objective for the original SCP, which can then be minimized. This is the Expected Residual Minimization (ERM) approach that we use in this work, which is commonly formulated as:

$$\min_z \mathbb{E}[\|\psi(z, F(z, \omega))\|^2] \tag{5}$$

The ERM formulation Eq (5) has an advantage over the expected value method, in that it has minimum sensitivity to variations in the random parameters [31]:

$$\mathbb{E}[\|\psi\|^2] = \|\mathbb{E}[\psi]\|^2 + \mathbb{E}[\|\psi - \mathbb{E}[\psi]\|^2]$$

where, for vector-valued ψ, $\mathbb{E}[\|\psi - \mathbb{E}[\psi]\|^2] = \text{tr}(\text{Cov}(\psi))$ is the trace of the covariance matrix, or the total variance. Thus, the ERM approach minimizes the mean-squared residual and the total variation with respect to random parameters.

IV. Contact-Robust Trajectory Optimization

A. Stochastic Complementarity in Trajectory Optimization

Previous work using SCPs in trajectory optimization developed ERM closed-form cost functions for the special case when the elements of F are normally distributed or logistically distributed [28]. In this work, we make use of the ERM for Gaussian distributed variables:

$$F \sim \mathcal{N}(\mu_F, \sigma_F)$$

$$\mathbb{E}[\min(z, F)^2] =$$

$$z^2 - \sigma_F^2(z + \mu_F)p(z) + (\sigma_F^2 + \mu_F^2 - z^2)P(z)$$

$$p(z) = \frac{1}{\sigma_F\sqrt{2\pi}} e^{-\frac{1}{2}(\frac{z-\mu_F}{\sigma_F})^2}$$

$$P(z) = \int_{-\infty}^{z} p(t)dt = \frac{1}{2}\left(1 + \text{erf}\left(\frac{z - \mu_F}{\sigma_F\sqrt{2}}\right)\right) \tag{6}$$

where $p(z)$ and $P(z)$ are the probability density and cumulative density functions for the normal distribution, respectively, evaluated at z. However, the previous work assumed the uncertainty resulted from state estimation and propagated directly to the SCP function F [28], which may not be consistent with the true uncertainty effects.

In this work, we assume the uncertainty lies directly in the contact parameters - specifically the friction coefficient μ and normal distance ϕ (Fig. 1(a)) - and derive the corresponding distributions F. We replace contact constraints with an ERM cost which encodes uncertainty about the terrain parameters:

$$\min_{\mathbf{x}, \mathbf{u}, \lambda} \sum_{i=0}^{N-1} \left(L(x_i, u_i, \lambda_i) + \beta \mathbb{E}[\|\psi(z_i, F(z_i, \omega))\|^2]\right) \tag{7}$$

where β is a weighting scalar, N is the total number of knot points, $\mathbf{x}, \mathbf{u}$ and λ represent collections of the respective variables across knot points, and $z_i \in \{\mathbf{x}, \lambda\}$ represents the variables in the complementarity constraints. Except where noted otherwise, we used $\beta = 10^4$ to keep the ERM cost within a few orders of magnitude of the quadratic costs (Fig. 1 c). Maintaining similar orders of magnitude prevents the optimization from trading off physical feasibility for reduced control effort, especially as the uncertainty vanishes. However, we also caution against

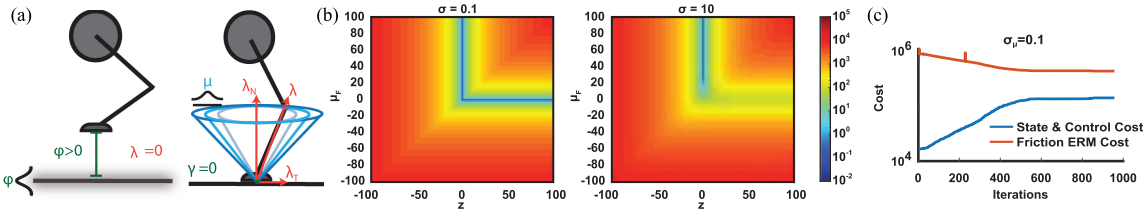

Fig. 1. (a) Contact geometry with uncertainty in terrain height (left) and friction coefficient (right). (b) ERM cost map for different values of uncertainty as a function of the decision variable z and mean of the uncertain constraint μ_F. Both subplots share the same logarithmic color axis for the ERM function value. For uncertainty less than 1, the minima of the ERM cost approaches the nonnegative axes. As uncertainty increases, the values on the positive μ_F-axis still represent local minima, but only at higher values of μ_F. (c) Convergence of the ERM and quadratic cost terms for the sliding block example shown on a logarithmic scale. A large multiplier β on the ERM cost ensures the two cost terms are within similar orders of magnitude of each other.

using arbitrarily large values of β which could cause numerical ill-conditioning.

Remark 1: A notable feature of the Gaussian ERM is that, as the uncertainty decreases, the ERM objective function approaches the residual function objective evaluated at the mean value of the uncertain variable:

$$\lim_{\sigma \to 0^+} \mathbb{E}[\min(z, F)^2] = \min(z, \mu_F)^2 \tag{8}$$

This property can be proved by using L'Hopital's rule to show that $\lim_{\sigma \to 0^+} \sigma^2 p(z) = 0$ and that $P(z) \to 1$ if $z - \mu_F > 0$ and $P(z) \to 0$ if $z - \mu_F < 0$ as $\sigma \to 0^+$.

Note that in our formulation the ERM term is added as a weighted cost. Thus, as $\sigma \to 0^+$, we recover a solution to a relaxed complementarity problem. A solution to the strict complementarity problem is recovered as $\beta \to \infty$.

B. Worst-Case Optimization

To compare against the ERM formulation, we also consider a worst-case scenario of the LCP formulation as a robust optimization [27]. A general robust counterpart (RC) of the uncertain optimization problem can be formulated as:

$$\min_{z \geq 0} \max_{\omega \in \Omega} \psi(z, F(z, \omega)) \quad \text{s.t.} \min_{\omega \in \Omega} F_i(z, \omega) \geq 0, \forall i \in \mathcal{I} \tag{9}$$

where the index set $\mathcal{I}$ comprises all the possible LCP instances. However, this RC is computationally challenging to solve in general. To derive a tractable RC, we impose an ∞-norm assumption on the uncertainty set Ω as $\Omega_\infty = \{\omega : ||\omega||_\infty \leq 1\}$. Given this set, we can express the residual value as $F(z, \omega) = F_0(z) + \sum_{k=1}^{K} \omega_k F_k(z)$. Accordingly, the robust optimization of LCP contact dynamics becomes

$$\min_{z \geq 0, \epsilon} \epsilon, \quad \text{s.t. } F_k(z) \geq 0,\ z F_k(z) \leq \epsilon,\ \forall k \in \{1, \ldots, K\}$$

where ϵ is a slack variable to minimize. Compared with the ERM formulation, this robust LCP formulation does not require a probability distribution over the uncertain parameters. Different from the expected value formulation which minimizes the mean scenario, the robust formulation version reasons about the worst-case scenario by enumerating K LCP instances in the uncertainty set Ω_∞. In practice, the worst-case scenario corresponds to a specific LCP instance; in this work we benchmark the performance of the ERM against this worst-case instance in one of our examples.

C. Characterizing Physical Contact Uncertainties

Here we explicitly parameterize uncertainties in the friction cone and in the contact geometry and develop the corresponding ERM cost functions. Specifically, we assume a normally distributed friction coefficient and a normally distributed error in the distance to the terrain and then derive the corresponding distributions used in the ERM objective.

1) Uncertainty in the Friction Coefficient: We assume the friction coefficient μ is normally distributed with mean $\bar{\mu}$ and standard deviation σ_μ : $\mu \sim \mathcal{N}(\bar{\mu}, \sigma_\mu)$. By linearity, the friction cone defect F_{FC} is also normally distributed:

$$F_{FC} = \mu \lambda_N - e\lambda_T \sim \mathcal{N}(\bar{\mu}\lambda_N - e\lambda_T, \sigma_\mu \lambda_N) \tag{10}$$

Thus, we can replace the constraint (1) with the ERM objective (6), where $\mu_F = \bar{\mu}\lambda_N - e\lambda_T$ and $\sigma_F = \sigma_\mu \lambda_N$.

2) Uncertainty in the Contact Distance: We assume the terrain is flat but that the contact distance is uncertain. The normal distance ϕ to the terrain is $\phi(q) = \eta^\top (h(q) - r)$ where η is surface normal of the terrain, and $h(q)$ and r are the Cartesian positions of the end effector and the nearest point on the terrain, respectively. We assume the normal distance ϕ is normally distributed: $\phi(q) \sim \mathcal{N}(\bar{\phi}(q), \sigma_\phi)$. In practice, the uncertainty can vary along the terrain therefore depend on configuration, i.e., $\sigma_\phi = \sigma_\phi(q)$. In either case, we can replace the normal distance constraint (1) with the ERM objective (6), where $\mu_F = \bar{\phi}(q)$ and $\sigma_F = \sigma_\phi$.

Remark 2: In theory, uncertainty in contact distance can also be expressed as uncertainty in the Cartesian coordinates of the nearest contact point, $r \sim \mathcal{N}(\bar{r}, \Sigma_r)$. However, because the terrain orientation η is known, the normal distance becomes $\phi \sim \mathcal{N}(\eta^\top (h(q) - \bar{r}), \eta^\top \Sigma_R \eta)$, which is equivalent to the preceding formulation in terms of normal distance.

Remark 3: Our formulation for uncertainty in the contact distance can be reformulated to account for uncertainty in the terrain orientation η. If we assume the contact distance is known but that the terrain orientation is normally distributed - i.e. $\eta \sim \mathcal{N}(\bar{\eta}, \Sigma_\eta)$, then the normal distance follows a normal distribution: $\phi \sim \mathcal{N}(\bar{\eta}^\top (h(q) - r), (h(q) - r)^\top \Sigma_\eta (h(q) - r))$ which is again equivalent to a distribution over the normal distance: $\phi \sim \mathcal{N}(\bar{\phi}(q), \sigma(q)_\phi)$. However, as the terrain orientation η also partly defines the contact Jacobian J_c, additional care should be taken to ensure the uncertainty effects are consistent across the normal distance, sliding velocity, and dynamics. Propagating

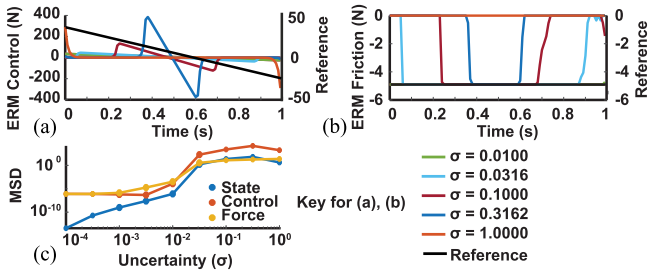

Fig. 2. Trajectories generated by the ERM method at different levels of uncertainty compared to the reference trajectory, for the same value of the friction coefficient. The control force (a) and frictional force (b) in the reference trajectory (shown on the right axis of each subplot) are linear for the entire motion. Under different values of uncertainty, the control and friction forces change and become shorter in duration. (c) As the uncertainty decreases, the ERM trajectories converge to the reference trajectories. All trajectories were generated using an expected friction coefficient $\bar{\mu} = 0.5$.

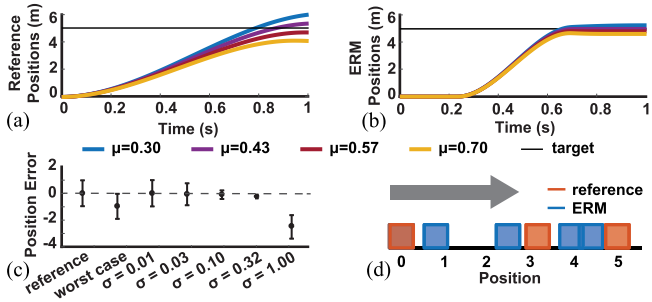

Fig. 3. Comparison of trajectories generated by the reference (a) and ERM (b) controls under terrain friction perturbations. The horizontal line represents the target position. (c) The mean and range of the final position error for different models of friction uncertainty. As uncertainty increases, the range of final positions decreases. However, if the uncertainty is too large, the planned motion is infeasible, and the simulation produces large position errors. (d) Selected frames of the block's motion from the reference and ERM trajectories. The arrow indicates the direction of motion.

uncertainty effects to the dynamics and deriving a corresponding risk-sensitive cost could be possible but is beyond the scope of the ERM framework we pursue here.

V. Simulation Experiments

Here we detail a set of simulation experiments to compare our ERM formulation to a baseline with non-stochastic nonlinear complementarity constraints, which we will refer to as the "non-robust" case. We study three examples: a block sliding over a surface with friction, a cart and double pendulum driven by contact with the ground, and a footed hopper with toe and heel contact. We compare the trajectories generated by our ERM formulation to those generated by the non-robust case for a range of uncertainty parameter values and used a mean-squared difference (MSD) criterion to assess convergence to the reference trajectory. All of our trajectory optimization examples were implemented in MATLAB using Drake [32] and solved using SNOPT [33]. Solution times are reported from unoptimized MATLAB code and are presented to evaluate the effect of the ERM cost on convergence speed; faster times could be achieved with optimized and compiled code. Unless otherwise noted, all of our examples were discretized using 101 kn points, initialized by linearly interpolating between the initial and final states, and solved to major optimality and major feasibility tolerances of 10^{-6}. Additionally, all ERM solutions were warm-started using the non-robust solution. Our code is available at https://github.com/GTLIDAR/RobustContactERM.

A. Sliding a Block Over a Surface With Unknown Friction

To benchmark the performance of the ERM method, we first study a two-dimensional 1 kg block with height 1 m sliding over a surface with uncertain friction (see Fig. 3(d)). The configuration of the block $q = [x_b, z_b]^\top$ is given by its planar center of mass (CoM) and the control is a horizontal force acting on the block. The initial and final states are $x_0 = [0, 0.5, 0, 0]^\top$ and $x_N = [5, 0.5, 0, 0]^\top$, and the total time is 1 s. The running cost has weight matrices $R = \text{diag}([100, 100])$ and $Q = \text{diag}([1, 1, 1, 1])$. We first solved to a tolerance of 10^{-6} and then solved to a tolerance of 10^{-8}. For the reference, non-robust trajectory we used a friction coefficient of $\mu = 0.5$. For the ERM trajectories, we assumed a mean friction $\bar{\mu} = 0.5$ and tested 9 values of σ logarithmically spaced between $\sigma = 0.001$ and $\sigma = 1.0$.

We compared the open-loop performance of the ERM controls to the non-robust controls in simulation using a time-stepping approach [13], [34] and 10 values of terrain friction linearly spaced between $\mu = 0.3$ and $\mu = 0.7$. All simulations started from the initial state x_0 and ran for 1 s. We also compared the simulations to a control generated using the worst-case scenario, where the friction uncertainty set was considered to be $\mu \in [0.3, 0.7]$ - in this case, the worst-case solution corresponds to using the lowest friction coefficient value, $\mu = 0.3$. We quantified the performance of the controls as the difference between target position of the block and the position achieved after 1 s.

B. A Contact Driven Cart With Unknown Terrain Height

Our second example is a double-pendulum connected to a cart which is constrained along a horizontal frictionless track (see Fig. 4(a)). The mass of the cart is 1 kg; each of the pendulums have mass 1 kg, length 1 m, and their CoMs coincide with their geometric centers. The configuration of the cart is $q = [x_c, \theta_1, \theta_2]$, where x_c is the horizontal position of the center of the cart and θ_1 and θ_2 are the angles of the pendulums. Only the joints of the pendulums are actuated; the cart must propel itself through contact with the ground.

The cart must travel from $x_{c,0} = 0$ m to $x_{c,N} = 5$ m in 1 s, starting and stopping at rest and with the end-effector in contact with the terrain. We used a $R = \text{diag}([1, 1])$ and $Q = \text{diag}([1100, 10, 1100, 10])$ in the cost function. For the reference trajectory, we first solved for a feasible trajectory which we used to warm-start the optimal solve. In the ERM experiments, we assumed a flat terrain with a mean distance of 1.5 m from the center of the cart and tested 7 values of height uncertainty logarithmically spaced between $\sigma = 0.001$ and $\sigma = 1$. Here we weighted the ERM cost by $\beta = 10^5$.

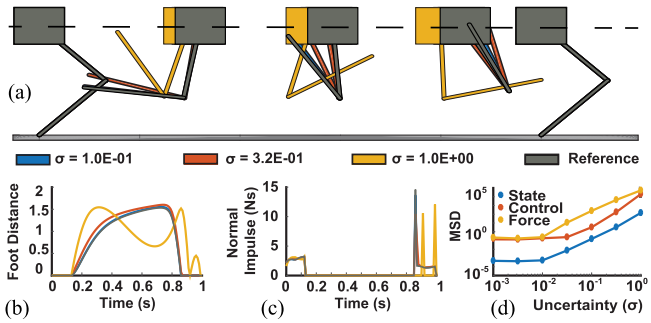

Fig. 4. Illustration of the relationship between contact uncertainty and foot clearance height. As uncertainty increases, ERM increases the distance to the terrain. (a) Selected configurations of the cart under different values of uncertainty, where the cart is constrained along a horizontal track. For $\sigma < 0.1$, the configurations are indistinguishable from the non-ERM reference trajectory. (b) The normal distance between the endpoint of the contact-driven cart and the terrain over the entire trajectory. As uncertainty increases, the distance increases until the second link flips over. (c) Normal ground reaction forces between the terrain and the endpoint of the cart. (d). Mean-squared difference between the ERM and reference solutions. As uncertainty decreases, the ERM trajectories converge to the reference trajectory.

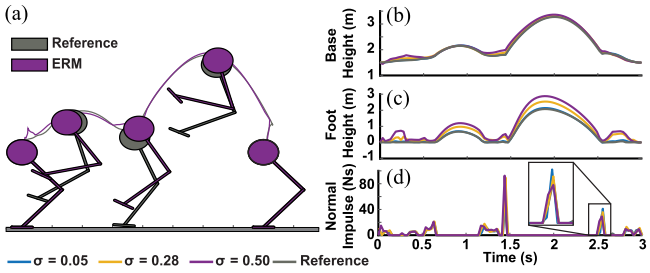

Fig. 5. Illustration of the hopper experiments with both friction and terrain uncertainty. (a) Selected configurations of the hopper in both the reference trajectory and in the ERM trajectory for height uncertainty $\sigma = 0.50$ and friction coefficient uncertainty $\sigma = 0.01$. Traces indicate the center of mass trajectories. (b) Base and (c) foot height trajectories under different uncertainties. The ERM cost increases the foot, but not the base, height. (d) Normal impulse trajectories across different height uncertainties.

C. A Single-Legged Hopper

Our final example is a single-legged hopper, free to move in the plane and with contact points at the toe and heel (see Fig. 5(a)). The configuration of the system is $q = [x_c, y_c, \theta_1, \theta_2, \theta_3]$ and the controls are the joint torques. In this example, the hopper must traverse 4 m in 3 s, starting and stopping at rest. The weights in the running cost were $R = \text{diag}([0.01, 0.01, 0.01])$ and $Q = \text{diag}([1, 10, 10\,100\,100, 1, 1, 1, 1, 1])$. As this example is more complex than our previous examples, we initialized the problem with all zeros and employed a relaxation technique [4], [35]. We first solved a series of problems in which we relaxed and then progressively tightened the complementarity constraints in powers of 10 from 10 to 10^{-4}, and then solved the strictly feasible problem. We then used the feasible solution to warm-start the optimal solve. In our ERM experiments, we used a friction ERM with mean $\bar{\mu} = 0.5$ and uncertainty of $\sigma = 0.01$ and a distance ERM with expected distance between the hopper base and terrain of 1.5 m under three uncertainties, $\sigma \in \{0.05, 0.28, 0.50\}$.

TABLE I
SELECTED SOLVE TIMES (S) FOR THE PUSHBLOCK EXAMPLE (COMPARED TO THE REFERENCE SOLVE TIME OF 101.0 S)

Uncertainty (σ)	10^{-4}	10^{-3}	10^{-2}	10^{-1}	10^{0}
ERM w/o Warmstart	524.1	756.3	343.2	350.4	45.3
ERM w/ Warmstart	182.1	185.3	200.2	293.2	109.0

VI. RESULTS

A. ERM Biases Away From Contact Interaction

To understand the effect of uncertainty on the solutions, we mapped the ERM cost landscapes for different levels of uncertainty (Fig. 1 b). At $\sigma = 0.1$, the ERM costmap has a set of low values near the nonnegative axes, which supports the claimed property that the ERM converges to the relaxed complementarity constraint when the uncertainty vanishes (Eq. (8)). However, as the uncertainty increases, the cost along the decision variable axis increases and, when the uncertainty is high enough, the cost for low values of the mean of the uncertain constraint also increases. For contact problems, a high uncertainty should bias the optimal trajectory towards reducing the friction force (therefore increasing the friction cone residuals) for uncertain friction and towards increasing ground clearance for uncertain terrain height.

B. ERM Increases Solve Time

Using the block example as a benchmark, we compared the total solve times for the ERM problem across a range of uncertainties when the ERM was and was not warm-started with the reference trajectory solution (Table I). As a note, across uncertainties the average number of iterations before the ERM converged was 772 for the ERM with warm-start (including iterations to get the warm-start) and 2357 without the warm-start; the reference trajectory converged in a total of 430 iterations. The large number of iterations is due in part to the tight tolerances used in the optimization. In most cases, we found that warm-starting the ERM with the reference resulted in a faster total solve time (including the warm-start solve time) compared to solving the ERM from the naive linear interpolation initialization, indicating the main benefit of the ERM lies in its robustness and not its smoothing properties. The remainder of this section therefore deals with solve times for ERM problems which were warm-started from the reference solution only.

For completeness, the total solve times for the ERM in the cart example in Fig. 4(a) were 291.0 s, 218.9 s, and 414.7 s respectively for $\sigma = 0.1$, 0.32, and 1.0, compared to the solve time of 92.4 s for the reference. Likewise, the solve times for the footed hopper were 6705 s, 3859 s, and 5875 s respectively for height uncertainties $\sigma = 0.05$, 0.28, and 0.5 compared to the reference time of 794.8 s.

C. ERM Generates Controls Robust to Changes in Friction

For low values of uncertainty ($\sigma \leq 0.01$), the ERM method produced trajectories that were nearly indistinguishable from the non-robust reference trajectory (Fig. 2), with mean-squared deviations less than 10^{-6} for state, 10^{-4} for control, and 10^{-3}

for contact force trajectories. For moderate values of uncertainty ($0.01 < \sigma < 1.0$), the generated trajectories deviate from the nominal trajectory, and the magnitude of the deviation grows with the magnitude of the uncertainty. Specifically, the controls are more aggressive and nonzero for only part of the duration and the friction forces are also nonzero for only part of the trajectory. At the highest value for uncertainty we tested ($\sigma = 1.0$), the control approaches a bang-bang control, and the friction forces are zero for the duration, which indicates that the ERM may produce infeasible solutions if the uncertainty is too high.

In open-loop simulation, the ERM generated controls with low uncertainty ($\sigma \leq 0.01$, $\sigma < 0.01$ not shown) produced trajectories with a spread in final state similar to the reference control (Fig. 3). Under frictional perturbations, the ERM controls with $\sigma \leq 0.01$ resulted in final positions within 1 m of the target and with an average error of 0 m. As the uncertainty increased, the spread in final positions decreased from 1 m to 0.22 m, indicating the uncertainty produced controls that were more robust to frictional perturbations. However, for $\sigma = 1.0$, the performance degraded and the open-loop average position error was -2.4 m. In this case, the ERM cost landscape corresponds to that in Fig. 1(b) ($\sigma = 10$), as the uncertainty in friction is multiplied by the normal force, $\lambda_N = 9.8N$. Moreover, the ERM solution set no longer corresponds to the complementarity solution, and therefore produces a physically infeasible solution. However, this behavior is sensitive to the scaling of the normal force; if we had used a different scale, such as kilonewtons, then the uncertainty also would have been multiplied by the scaling, altering the ERM landscape to produce feasible solutions.

In contrast, the worst-case scenario always produced a feasible trajectory with respect to at least one value of the friction coefficient in the uncertainty set. However, in open-loop simulations, the worst-case control had an average error of -0.95 m and a range of -1.9 m, which was comparable to the range produced by the reference non-robust control.

D. ERM Increases Foot Ground Clearance

In the cart example, low values of uncertainty ($\sigma \leq 0.01$) in the ERM objective resulted in trajectories that were close to the optimal non-robust trajectory (Fig. 4). However, as the uncertainty increased, the ERM-generated trajectories increased the distance between the terrain and the end-effector. The average distance to the terrain was 0.84 m for $\sigma = 1.0$ and 0.79 m for $\sigma = 0.001$, compared to an average distance of 0.79 m in the reference non-robust trajectory.

For all values of terrain height uncertainty tested in the hopping example, incorporating terrain and friction uncertainties increased the foot clearance of the hopper (Fig. 5). Moreover, the increase in foot clearance trended with the uncertainty in the terrain contact distance, with increases of 4.7% ($\sigma = 0.05$), 4.32% ($\sigma = 0.28$) and 74.1% ($\sigma = 0.50$) in our experiments. In contrast, the height of the base of the hopper increases only marginally, with increases of 0.8% ($\sigma = 0.05$), 1.7% ($\sigma = 0.28$), and 3.1% ($\sigma = 0.50$). In all cases, the friction uncertainty was fixed at $\sigma_\mu = 0.01$.

VII. Discussion and Conclusion

Our ERM method for modeling terrain uncertainties is an important step towards a deployable terrain-robust contact-implicit trajectory optimization. As explored in previous work [28], the ERM approach represents a stochastic variant of a smoothed complementarity problem [36] and permits some nonzero contact force at a nonzero distance (Fig. 5(c,d)). While in principle the ERM method could also act as a soft contact model and permit ground penetration, we did not observe that behavior in our work here, although that result could be due to our choice of β. One advantage of our approach over the previous work [28] is that our approach explicitly models and is robust to uncertainty in the contact parameters. By evaluating a variety of uncertainty parameters, we demonstrated that our approach generates trajectories of varying robustness and converges to the traditional, non-robust solution as the uncertainty vanishes.

The proposed ERM method is similar to the previous ensemble approach in that both achieve robustness by introducing a cost with respect to random parameter variations [3]. However, unlike the previous approach, we did not need more than one trajectory to achieve robustness. Instead, we assumed normal distributions over the friction coefficient and terrain height and calculated the expected value analytically as in [28]. A closed-form expression for the expectation allowed us to avoid sampling-based approaches. However, as in the original work [28], we note that the ERM objective in our study has no physical meaning - the interpretation of the complementarity constraints is lost when the constraints are replaced with the residual function in Eq. (4). Moreover, the normal distribution assumption places probability density on negative friction coefficients, which are physically infeasible. Future work may improve on our work by developing terrain-robust objectives that admit a physical interpretation, or by deriving ERM methods for distributions over the positive reals, such as a truncated Gaussian or a Gamma distribution.

In the context of risk-sensitive control, our approach is analogous to the risk-averse control in [22]. For our work under friction uncertainty, the optimization incurs little additional cost from the uncertainty if the system is at rest and there are no tangential frictional forces, provided the normal force is sufficiently large. Thus, the friction ERM cost promotes the short and fast sliding motions observed in the sliding block example. In the uncertain terrain distance model, the ERM cost penalizes proximity to the expected terrain, and thus the system tends to move away from the terrain, using more control and taking higher steps to reach the goals. These behaviors can be understood as risk averse, as the ERM minimizes the interactions between the system and the uncertain terrain. In contrast to our approach, risk-seeking behaviors could reward the system for making more contact interactions with the environment and could be useful for robots to collect more terrain data for estimation. Here we focused on simple examples to demonstrate these risk-sensitive behaviors of the ERM approach; however, how these behaviors scale up to more complex robots and more challenging terrains remains an open avenue of research.

In this work, we compared our ERM approach to uncertainty in the complementarity conditions to a worst-case solution. The

worst-case solution, inspired by [27], is a robust, distribution-free method to solve complementarity problems. However, unlike our ERM method, the worst-case method assumes a discrete set of uncertain values and solves for the value in the set that maximizes the complementarity residual. This is analogous to choosing a particular value of the uncertain parameter (for example, choosing the friction coefficient), and then solving the corresponding optimization. While the worst-case may achieve a robust solution to the complementarity problem, that robustness does not translate to the generated controls, as the worst-case solutions produce open-loop trajectories with the same endpoint variation as the standard contact-implicit method. Thus, although our approach may not strictly satisfy the complementarity constraints for all values of the uncertain parameters, it does have an advantage over the worst-case method in that the control trajectories inherit robustness from the ERM solutions.

One important feature of our work is that, as the uncertainty approaches zero, the trajectories approach the solutions generated by using the mean value of the uncertain parameters. While the low uncertainty case can be interpreted as a smooth and accurate approximation to the original nonsmooth complementarity constraint, we also note that the property alone is important, as it opens an avenue for combining model-based approaches, such as contact-implicit trajectory optimization, with model-free Bayesian optimization methods [20], [37]. Future work could combine our work here with measurements from the terrain to estimate the terrain parameters during locomotion and close the loop of terrain estimation and robust trajectory optimization.

REFERENCES

[1] N. Ratliff, M. Zucker, J. A. Bagnell, and S. Srinivasa, "CHOMP: Gradient optimization techniques for efficient motion planning," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2009, pp. 489–494.

[2] J. Schulman *et al.*, "Motion planning with sequential convex optimization and convex collision checking," *Int. J. Robot. Res.*, vol. 33, no. 9, pp. 1251–1270, Aug. 2014.

[3] I. Mordatch, K. Lowrey, and E. Todorov, "Ensemble-cio: Full-body dynamic motion planning that transfers to physical humanoids," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2015, pp. 5307–5314.

[4] M. Posa, C. Cantu, and R. Tedrake, "A direct method for trajectory optimization of rigid bodies through contact," *Int. J. Robot. Res.*, vol. 33, no. 1, pp. 69–81, 2014.

[5] M. Posa, S. Kuindersma, and R. Tedrake, "Optimization and stabilization of trajectories for constrained dynamical systems," in *Proc. Int. Conf. Robot. Automat.*, 2016, pp. 1366–1373.

[6] S. Kuindersma *et al.*, "Optimization-based locomotion planning, estimation, and control design for atlas," *Auton. Robots*, vol. 40, no. 3, pp. 429–455, 2016.

[7] B. Katz, J. D. Carlo, and S. Kim, "Mini cheetah: A platform for pushing the limits of dynamic quadruped control," in *Proc. Int. Conf. Robot. Automat.*, May 2019, pp. 6295–6301.

[8] Z. Manchester, N. Doshi, R. J. Wood, and S. Kuindersma, "Contact-implicit trajectory optimization using variational integrators," *Int. J. of Robot. Res.*, vol. 38, no. 12/13, pp. 1463–1476, 2019.

[9] A. Patel, S. Shield, S. Kazi, A. M. Johnson, and L. T. Biegler, "Contact-implicit trajectory optimization using orthogonal collocation," *IEEE Robot. and Automat. Lett.*, vol. 4, no. 2, pp. 2242–2249, Apr. 2019.

[10] I. Mordatch, E. Todorov, and Z. Popović, "Discovery of complex behaviors through contact-invariant optimization," *ACM Trans. Graph.*, vol. 31, no. 4, pp. 43–43:8, Aug. 2012.

[11] J.-P. Sleiman, J. Carius, R. Grandia, M. Wermelinger, and M. Hutter, "Contact-implicit trajectory optimization for dynamic object manipulation," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2019, pp. 6814–6821.

[12] H. Dai, A. Valenzuela, and R. Tedrake, "Whole-body motion planning with centroidal dynamics and full kinematics," in *Proc. IEEE-RAS Int. Conf. Humanoid Robots*, 2014, pp. 295–302.

[13] D. E. Stewart and J. C. Trinkle, "An implicit time-stepping scheme for rigid body dynamics with inelastic collisions and coulomb friction," *Int. J. Numer. Methods Eng.*, vol. 39, no. 15, pp. 2673–2691, 1996.

[14] D. Pardo, M. Neunert, A. Winkler, R. Grandia, and J. Buchli, "Hybrid direct collocation and control in the constraint-consistent subspace for dynamic legged robot locomotion," in *Robotics: Science and Systems*, 2017.

[15] J. Carpentier, R. Budhiraja, and N. Mansard, "Learning feasibility constraints for multi-contact locomotion of legged robots," in *Robotics: Science and Systems*, p. 9p, 2017.

[16] J. Luo and K. Hauser, "Robust trajectory optimization under frictional contact with iterative learning," *Auton. Robots*, pp. 1447–1461, 2017.

[17] A. M. Johnson, J. E. King, and S. Srinivasa, "Convergent planning," *IEEE Robot. and Automat. Lett.*, vol. 1, no. 2, pp. 1044–1051, Jul. 2016.

[18] Y. Pan, E. Theodorou, and K. Bakshi, "Robust trajectory optimization: A cooperative stochastic game theoretic approach," in *Robotics: Science and Systems*, 2015.

[19] C. Song and A. Boularias, "Identifying Mechanical Models Through Differentiable Simulations," *Proc. Mach. Learn. Res.*, May 2020, pp. 1–12.

[20] S. R. Kuindersma, R. A. Grupen, and A. G. Barto, "Variable risk control via stochastic optimization," *Int. J. Robot. Res.*, vol. 32, no. 7, pp. 806–825, 2013.

[21] S. Singh, Y. Chow, A. Majumdar, and M. Pavone, "A framework for time-consistent, risk-sensitive model predictive control: Theory and algorithms," *IEEE Trans. Autom. Control*, vol. 64, no. 7, pp. 2905–2912, Jul. 2019.

[22] F. Farshidian and J. Buchli, "Risk sensitive, nonlinear optimal control: Iterative linear exponential-quadratic optimal control with gaussian noise," 2015, *arXiv:1512.07173*.

[23] B. Ponton, S. Schaal, and L. Righetti, "Risk sensitive nonlinear optimal control with measurement uncertainty," 2016, *arXiv:1605.04344*.

[24] D. Jacobson, "Optimal stochastic linear systems with exponential performance criteria and their relation to deterministic differential games," *IEEE Trans. Autom. Control*, vol. 18, no. 2, pp. 124–131, Apr. 1973.

[25] X. Chen and M. Fukushima, "Expected residual minimization method for stochastic linear complementarity problems," *Math. Oper. Res.*, vol. 30, no. 4, pp. 1022–1038, 2005.

[26] C. Chen and O. L. Mangasarian, "A class of smoothing functions for nonlinear and mixed complementarity problems," *Comput. Optim. Appl.*, vol. 5, no. 2, pp. 97–138, 1996.

[27] Y. Xie and U. V. Shanbhag, "On robust solutions to uncertain linear complementarity problems and their variants," *SIAM J. Optim.*, vol. 26, no. 4, pp. 2120–2159, 2016.

[28] Y. Tassa and E. Todorov, "Stochastic complementarity for local control of discontinuous dynamics," in *Robotics: Science and Systems*, 2010.

[29] J. Carius, R. Ranftl, V. Koltun, and M. Hutter, "Trajectory optimization with implicit hard contacts," *IEEE Robot. Automat. Lett.*, vol. 3, no. 4, pp. 3316–3323, Oct. 2018.

[30] M.-J. Luo and Y. Lu, "Properties of expected residual minimization model for a class of stochastic complementarity problems," *J. Appl. Math.*, vol. 2013, pp. 1–7, 2013.

[31] X. Chen, C. Zhang, and M. Fukushima, "Robust solution of monotone stochastic linear complementarity problems," *Math. Program.*, vol. 117, no. 1/2, pp. 51–80, Mar. 2009.

[32] R. Tedrake and the Drake Development Team, Drake: Model-based design and verification for robotics, 2019. [Online]. Available: https://drake.mit.edu

[33] P. E. Gill, W. Murray, and M. A. Saunders, "SNOPT: An SQP algorithm for large-scale constrained optimization," *SIAM Rev*, vol. 47, no. 1, pp. 99–131, 2005.

[34] M. Anitescu and F. A. Potra, "Formulating dynamic multi-rigid-body contact problems with friction as solvable linear complementarity problems," *Nonlinear Dyn.*, vol. 14, no. 3, pp. 231–247, 1997.

[35] S. Scholtes, "Convergence properties of a regularization scheme for mathematical programs with complementarity constraints," *SIAM J. Optim.*, vol. 11, no. 4, pp. 918–936, Jan. 2001.

[36] C. Chen and O. L. Mangasarian, "A class of smoothing functions for nonlinear and mixed complementarity problems," *Comput. Optim. Appl.*, vol. 5, no. 2, pp. 97–138, Mar. 1996.

[37] T. Seyde, J. Carius, R. Grandia, F. Farshidian, and M. Hutter, "Locomotion planning through a hybrid bayesian trajectory optimization," in *Proc. Int. Conf. Robot. Automat.*, 2019, pp. 5544–5550, ISSN: 1050–4729.