Uncertainty-Aware Contact-Implicit MPC via GPU-Parallelized ADMM

Herbert Wright[1], Michael Posa[1]

***Abstract*—Reasoning about different contact modes is critical for many manipulation tasks, motivating the creation of contact-implicit model predictive control (CI-MPC) approaches. Contact can also amplify uncertainty, which is common during real-world manipulation. Despite this, CI-MPC approaches often ignore any notion of uncertainty, because it would add further computational complexity. In this work, we present ARCtIC, a GPU-accelerated CI-MPC controller capable of reasoning about model uncertainty during contact-rich manipulation. Our method takes an ensemble approach and formulates an expectational MPC objective which includes state-feedback for uncertainty-awareness. We efficiently solve this optimization problem through reformulations, applying the alternating direction method of multipliers (ADMM), and GPU parallelization. Results show that our method is capable of controlling a simulated 9-DoF Trifinger under model uncertainty in real time, outperforming both robust sampling-based approaches as well as deterministic counterparts.**

I. Introduction

Contact lies at the core of robotic manipulation. Thus, it is crucial for robots to effectively reason about contact: both how to make it and the effects of their actions through it. However, rigid contacts can be difficult to optimize over due to the approximate discontinuities introduced in the underlying dynamics. For example, when a robot forcefully pushes an object, the object undergoes an impulsive jump in velocity when contact is made, creating an approximate discontinuity in state over time. The discontinuous nature of contact means that traditional control strategies can struggle when applied to contact-rich tasks. In order to better control such systems, efficient *contact-implicit* model predictive control (CI-MPC) methods have been proposed [1]–[3]. To make the optimization problem tractable, these methods often make significant relaxations or approximations to the dynamics, such as contact linearizations or smoothing.

Contact also complicates things when uncertainty is present. Imperfect perception and partial information often result in model inaccuracies, and small model errors can be amplified through contact [4]. For example, if a robot misjudges the size of an object, it may reach out thinking it will make contact and completely miss. The field of robust control [5] has a rich history of solving control problems under uncertainty, however these methods remain underexplored with respect to the online control of contact-rich systems. This is because robust approaches often add significant complexity on top of the already-complex problem of contact-rich control, and the inclusion of feedback over uncertainty is often non-trivial. For sim-to-real reinforcement learning (RL), robustness can critically be achieved via *offline domain randomization* [6] to overcome the reality gap from inaccurate models [7]. This is one of the main benefits of RL over online MPC, but RL approaches still struggle with their ability to transfer to out-of-distribution goals and sample efficiency. Thus, it is an important open challenge to efficiently perform robust and uncertainty-aware control in a contact-implicit manner. Previous approaches are either extremely computationally expensive [8], [9] or rely on highly restrictive assumptions about the type of uncertainty or task [10], [11]. This work bridges this gap and proposes the first efficient *and* effective MPC algorithm that can reason about both complementarity constraints and uncertainty. Our insight is that domain-randomization-like robustness can be achieved within C3-style [1] CI-MPC through optimization-oriented reformulations and GPU parallelization.

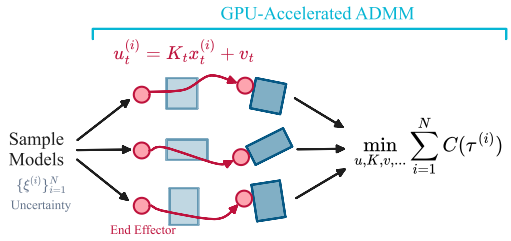

Fig. 1. Our method, ARCtIC, is a robust contact-implicit model-predictive controller. We formulate a large sampling-based optimization problem which includes feedback and solve it efficiently with GPU-accelerated ADMM-based optimization.

We present ADMM-based Robust Contact-Implicit Control (ARCtIC), an efficient MPC controller capable of controlling complementarity systems under model uncertainty (Figure 1). Our method, ARCtIC approximates an expectation-based cost through a collection of samples. Then, we solve a highly discontinuous and non-convex optimization problem by reformulating and breaking it into a few easily computable steps via the alternating direction method of multipliers (ADMM) [12]. Each of these steps is highly parallelizable and essentially becomes a short sequence of batched linear solves with positive-definite matrices. We showcase our method in three different settings, demonstrating its ability to reason about uncertainty. We show that our method is capable of reasoning about uncertainty over *object geometry*, which is a critical gap in previous robust CI-MPC strategies [13].

II. Related Works

A. Contact-Implicit Control

The highly stiff contact dynamics found in many manipulation tasks make typical control approaches difficult [13], where

[1]GRASP, University of Pennsylvania, PA USA

often the most successful approaches amount to inefficient random search (e.g. [14]). To address this, contact-implicit model predictive control (CI-MPC) methods aim to solve for a trajectory of actions through making and/or breaking contact, with contact modes implicitly chosen during the optimization. One form of these methods relaxes the contact model into a smoother version, creating artifacts in order to obtain more informative gradients [2], [3], [15], [16], to be used by first-order optimizers. Another form of contact-implicit control is to optimize for a trajectory subject to complementarity constraints [17]. This has been utilized in tandem with ADMM to build various efficient contact-implicit controllers [1], [18]–[20] ARCtIC falls into this latter category as we leverage ADMM to explicitly optimize over complementarity constraints, but differs from previous work as we additionally consider the critical role of uncertainty during manipulation.

B. Uncertainty in Manipulation

Robotic manipulation in less-structured environments inherently involves uncertainty [21], [22]. Thus, many techniques have been proposed to deal with uncertainty during manipulation. Uncertainty during planning has been tackled by operating in belief-space [23], though it is often computationally expensive. Closer to the mechanics of manipulation, early work explored synthesizing motion plans to reduce uncertainty for precise manipulation tasks [24], [25]. In robotic grasping, methods have been developed concerning robustness to *shape uncertainty* [26], [27], yet it remains under-explored in the context of contact-implicit control [13]. With respect to control for manipulation, uncertainty-awareness has been applied to sampling-based controllers [10], [28], [29], however has not seen as much success in other CI-MPC forms. In [8], optimization is performed with chance constraints over LCS uncertainty, but it was only run offline in open-loop. In contrast, our controller performs closed-loop control for high-dimensional contact-rich tasks under model uncertainty. In reinforcement learning, domain randomization over physical parameters can be viewed as enforcing robustness to model uncertainty [6]. We seek to bring this robustness to CI-MPC.

III. Preliminaries

A. Linearization of Anitescu Contact Dynamics

In this work, we use a variant of Anitescu's formulation of contact dynamics [30] to model the dynamics of a scene during optimization. We follow [1] and convert Anitescu's formulation into a linear complementarity system (LCS) [31]:

$$x_{t+1} = Ax_t + Bu_t + D\lambda_t + d \tag{1a}$$

$$0 \leq \lambda_t \perp Ex_t + F\lambda_t + Hu_t + c \geq 0, \tag{1b}$$

where $x_t = [q_t, v_t]^\top$. We refer readers to [1] for a more detailed explanation. Linearized versions of Anitescu's dynamics have been effectively leveraged to control many contact-rich manipulation tasks, including multi-object pushing [19], ball rolling [1], and a dynamic waiter task [18]. In practice, we use auto-differentiation to recover the necessary Jacobians and quantities to efficiently construct the LCS.

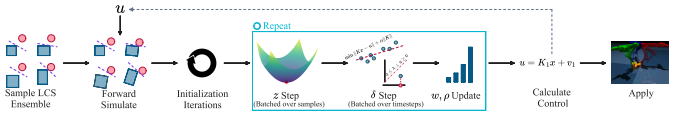

Fig. 2. Overview of the ARCtIC algorithm. We begin by sampling an ensemble of linear complementarity systems (LCS's), then we forward simulate each LCS and run a number of intialization iterations (described in Section IV-F). Next, we solve a distributional CI-MPC problem (Sections IV-A and IV-B) by applying ADMM (Section IV-C), turning it into easily computable z step (Section IV-D), δ step (Section IV-E), and ρ, w updates (Section IV-F).

B. Alternating Direction Method of Multipliers

The Alternating Direction Method of Multipliers (ADMM) is a primal-dual algorithm [12] that aims to solve:

$$\min_{z,\delta} f(z) + g(\delta) \quad \text{s.t. } h(z,\delta) = 0$$

ADMM is a dual ascent method. Specifically, it leverages the augmented Lagrangian, $\mathcal{L}_\rho(z, \delta, w)$, with w being the scaled dual variable. ADMM alternates the following updates:

$$z = \arg\min_z \mathcal{L}_\rho(z, \delta, w) \tag{2a}$$

$$\delta = \arg\min_\delta \mathcal{L}_\rho(z, \delta, w) \tag{2b}$$

$$w = w + h(z, \delta). \tag{2c}$$

In our method, we apply ADMM to a case where $h(z, \delta)$ is a bilinear equality constraint, which reduces to a linear constraint when solving (2a) or (2b) individually. ADMM is a popular technique for solving contact-implicit control problems when complementarity constraints are present [1], [18]–[20].

IV. ADMM-based Robust Contact-Implicit Control

A. The Robust Contact-Implicit Control Problem

Generally, contact-implicit control methods have assumed a known model with little uncertainty. In order to account for uncertainty, we formulate the robust contact-implicit control problem by starting from a stochastic policy optimization problem, approximated via sampling, then follow previous contact-implicit approaches by including contact forces as decision variables, subject to complementarity constraints.

We begin with the stochastic policy optimization problem:

$$\min_\pi \mathbb{E}_\xi[C_\xi(\pi)], \tag{3}$$

where $\pi : \mathcal{X} \to \mathcal{U}$ is a policy and ξ is represents parametric uncertainty. This could be the scene geometry, coefficients of friction, inertial properties, or any other variable that affects dynamics. We consider solving (3) in a receding-horizon MPC fashion, in the setting where the cost is a quadratic function. In our approach, we consider sampling $\{\xi^{(i)}\}_{i\in[N]}$ from the uncertainty distribution. Then, we define the the cost $C_\xi(\pi)$ as a quadratic cost starting from the current state and having a fixed horizon T. Finally, we linearize $\pi(x_t) = u_t = K_t x_t + v_t$ at each timestep and include x, u, K, v as decision variables with corresponding constraints.

We follow previous approaches by explicitly including contact forces, λ, as decision variables, yielding:

$$\min_{x,u,\lambda,K,v} \sum_{i=1}^{N} \left\|x_T^{(i)}\right\|_{Q_f}^2 + \sum_{t=0}^{T-1} \left\|x_t^{(i)}\right\|_Q^2 + \left\|u_t^{(i)}\right\|_R^2 \tag{4a}$$

$$\text{s.t.} \quad x_{t+1}^{(i)} = f_{\xi^{(i)}}(x_t^{(i)}, u_t^{(i)}, \lambda_t^{(i)}) \tag{4b}$$

$$0 \le \lambda_t^{(i)} \perp \Phi_{\xi^{(i)}}(x_t^{(i)}, u_t^{(i)}, \lambda_t^{(i)}) \ge 0 \tag{4c}$$

$$u_t^{(i)} = K_t x_t^{(i)} + v_t \tag{4d}$$

$$x_0^{(i)} = x_0, \tag{4e}$$

where Q_f, Q, R are cost matrices and f_ξ is the nonlinear dynamics of the system. Allowing the control input for each sample to vary based on a state-feedback controller from linearizing π acts similarly to *covariance steering*. Covariance steering helps prevent being overly conservative [8], [32].

B. Adjusted Robust CI-MPC Problem

Our method, ARCtIC, solves an adjusted version of the MPC objective in Section IV-A with an LCS model:

$$\min_{x,u,\lambda,K,v} \left[\sum_{i=1}^{N} \left\|x_{T+1}^{(i)}\right\|_{Q_f}^2 + \sum_{t=1}^{T} \left\|x_t^{(i)}\right\|_Q^2 + \left\|u_t^{(i)}\right\|_R^2\right] \tag{5a}$$

$$+ \alpha \sum_t \|K_t\|^2 \tag{5b}$$

$$\text{s.t. } x_{t+1}^{(i)} = A^{(i)} x_t^{(i)} + B^{(i)} u_t^{(i)} + D^{(i)} \lambda_t^{(t)} + d^{(i)} \tag{5c}$$

$$\eta_t^{(i)} = E^{(i)} x_t^{(i)} + F^{(i)} \lambda_t^{(i)} + H^{(i)} u_t^{(t)} + c^{(i)} \tag{5d}$$

$$0 \le \lambda_t^{(i)} \perp \eta_t^{(i)} \ge 0 \tag{5e}$$

$$u_t^{(i)} = K_t x_t^{(i)} + v_t \tag{5f}$$

$$x_1^{(i)} = \text{SIMULATELCS}^{(i)}(x_0, u_0) \tag{5g}$$

A critical difference in this problem from (4) is that the decision variables start at $t = 1$, and, because uncertainty impacts the state at $t = 1$, the first u_1 is computed from state-feedback. Additionally, a regularization cost on the K matrix has been added to penalize high-gain solutions and we follow [19], [20] by introducing η variable for the complementarity constraints.

C. ADMM Reformulation

Here, we reformulate (5) into ADMM updates, taking the form outlined in Section III-B. We rewrite the problem as:

$$\min_{x,u,\lambda,\eta,K,v} C(x,u) + \mathcal{I}_{\mathcal{D}}(x,u,\lambda,\eta) + \mathcal{I}_{\mathcal{H}}(\lambda,\eta) + \alpha \sum_t \|K_t\|^2$$

$$\text{s.t. } u_t^{(i)} = K_t x_t^{(i)} + v_t$$

where $\mathcal{I}_S(s) = 0$ if $s \in S$ and ∞ otherwise; $C(x,u)$ is the cost in (5a); $\mathcal{D}$ denotes feasible x, u, λ, η according to the constraints in (5c), (5d), and (5g); and $\mathcal{H}$ denotes feasible λ, η satisfying (5e). In practice this MPC problem/formulation contains up to $64\times$ as many decision variables and constraints compared to [1].

Next, we introduce two decision variables which contain certain copied variables:

$$z = [x, u, \lambda, \eta]^\top \qquad \delta = [x, K, v, \lambda, \eta]^\top \tag{6}$$

and formulate our objective into the following problem:

$$\min_{z,\delta} C(z) + \mathcal{I}_{\mathcal{D}}(z) + \mathcal{I}_{\mathcal{H}}(\delta) + \alpha \sum_t \left\|\delta^{K_t}\right\|^2 \tag{7a}$$

$$\text{s.t. } z^{u_t^{(i)}} = \delta^{K_t} z^{x_t^{(i)}} + \delta^{v_t} \tag{7b}$$

$$z^{x_t^{(i)}} = \delta^{x_t^{(i)}}, z^{\lambda_t^{(i)}} = \delta^{\lambda_t^{(i)}}, z^{\eta_t^{(i)}} = \delta^{\eta_t^{(i)}} \tag{7c}$$

We apply ADMM to this bilinear-constrained problem. The inclusion of multiple samples, the bilinear feedback coupling, and gain regularization differentiate this formulation from prior work [1]. We operate on the augmented Lagrangian of (7). Under this reformulation, solving the two ADMM subproblems, $\arg\min_z \mathcal{L}_\rho(z, \delta, w)$ and $\arg\min_\delta \mathcal{L}_\rho(z, \delta, w)$ becomes surprisingly parallelizable and efficient.

D. Solving for z

This section outlines how our controller efficiently solves the first ADMM subproblem: $\arg\min_z \mathcal{L}_\rho(z, \delta, w)$. We note that this optimization problem allows us to ignore the non-linear complementarity constraints and leaves us with an equality-constrained quadratic program per each particle. We can formulate the per-particle optimization problem as:

$$z^* = \arg\min_z \frac{1}{2} z^\top \tilde{Q} z + \tilde{c}^\top z \quad \text{s.t.} \quad \tilde{A} z = \tilde{b}$$

which can be solved by setting up the KKT matrix and solving the linear system. For efficiency, we use the Shur complement, $S = -\tilde{A}\tilde{Q}^{-1}\tilde{A}^\top$, which gives:

$$y^* = S^{-1}(\tilde{b} + \tilde{A}\tilde{Q}^{-1}\tilde{c}) \tag{8a}$$

$$z^* = \tilde{Q}^{-1}\tilde{A}^\top y^* - \tilde{Q}^{-1}\tilde{c} \tag{8b}$$

We leverage Cholesky solves for the computation, because $\tilde{Q} \succ 0$ and $-S \succ 0$. In practice, we apply various numerical techniques, such as Ruiz equilibration [33] for stability.

E. Solving for δ

Next, we focus on efficiently solving the second ADMM subproblem: $\arg\min_\delta \mathcal{L}_\rho(z, \delta, w)$. This can be separated into solving three independent optimization problems:

a) Solving for δ^K, δ^v: This minimization is equivalent to ridge regression across particles each timestep:

$$\delta^{K_t}, \delta^{v_t} = \arg\min_{K,v} \left\|w^{u_t^{(i)}} + z^{u_t^{(i)}} - K z^{x_t^{(i)}} + v\right\|^2 + \alpha \|K\|^2$$

This reduces to a positive definite linear solve [34], which we perform via Cholesky decompositions.

b) Solving for δ^x: Trivially, $\delta^x = w^x + z^x$.

c) Solving for $\delta^\lambda, \delta^\eta$: This can be solved analytically as the optimization problem is separable both between particles and between timesteps. It becomes a large set of 1D non-linear optimization programs, with easy analytic solutions [19], [20].

Each of these three steps can be batched per-timestep.

F. The Full ARCtIC Controller

Here, we detail the full ARCtIC algorithm, which efficiently runs ADMM to recover a solution to the adjusted problem outlined in Section IV-A. Our algorithm begins by sampling N LCS systems and forward simulates the current observed state with the known first control input for each LCS. Then, we perform a number of warm-start iterations on a simplified

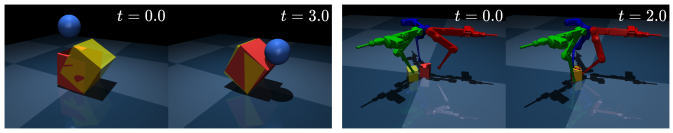

Fig. 3. Visualization of the box pivot and Trifinger reorientation tasks.

version of the problem where $K = 0$, and ρ is fixed, which keeps $\tilde{Q}, S$ constant across iterations, allowing us to cache the required Cholesky decompositions. Then, we run the full ADMM steps for a fixed set of iterations, with a predetermined ρ schedule. Finally, we return the first K_1, v_1 from our MPC plan to calculate u and apply. The full ARCtIC algorithm can be seen in Figure 2. We implement our controller in JAX, taking full advantage of GPU parallelization and JIT compilation. For LCS construction, we leverage MuJoCo MJX [35]. For solving QPs during LCS simulation, we use [36].

V. Experiments

Here, we run experiments on two contact-rich simulation tasks (Figure 3), evaluating robustness to model uncertainty.

a) Baselines: We compare our method against multiple baselines. We compare against a version of ARCtIC without feedback, referred to as **ARCtIC-NF**. We also ablate our method with a version that operates on a single non-random nominal model, instead of a distribution. We refer to this version as **ARCtIC-D**, and note that it is essentially equivalent to the C3+ [19] with a few minor differences. Additionally, we compare our method with a *deterministic* MJPC predictive sampling baseline [14], referred to as **MJPC**, as well as **EMPPI** [28], and a beefier version of EMPPI incapable of running in real-time (**EMPPI+**).

b) Box-pivot Experiment: Our first simulation experiment is a contact-rich box pivot task subject to shape uncertainty. The goal is to tip a box with randomized sides to an angle of 30° after 3 seconds. The force-controlled spherical end effector is initialized randomly above the box uniformly within bounds of 0.1 m in either direction. The box side lengths are independently uniformly sampled from the range $[0.27\text{ m}, 0.33\text{ m}]$. We compare our method (ARCtIC) with a no-feedback version (ARCtIC-NF) and a non-distributional deterministic version (ARCtIC-D), which runs our algorithm on an ensemble of a single model with a cube with 0.3 (m) sides. Results for this experiment are in the right-hand side of Figure 4. We find that both an ensembled approach and the incorporation of feedback are critical for optimal performance.

c) Trifinger Reorientation: Our second simulation task is a 9-DoF torque-controlled Trifinger [37] reorientation task. We follow [38], [39] and generate random planar goal positions $p_{\text{goal}} \in [-6\text{ cm}, 6\text{ cm}]$ and orientations $\theta \in [-0.5\text{ rad}, 0.5\text{ rad}]$, with some small initialization randomization. Each controller has 2 s to get the cube within a tolerance of 0.02 m and 0.2 rad of the goal pose. We run two experiments with this task, evaluating the Full ARCtIC, ARCtIC-NF, and ARCtIC-D baselines. Even though this task contains 40 contact variables, 31-dimensional state, and a 9-dimensional control space, our method is capable at running at a real-time rate as it takes about 70 ms per solve.

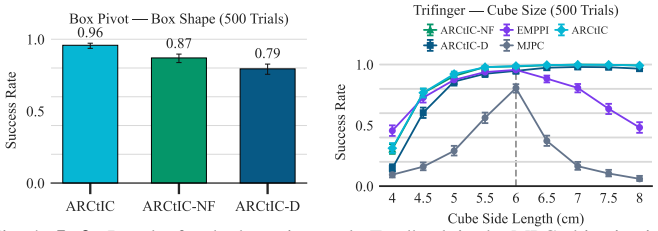

Fig. 4. **Left:** Results for the box pivot task. Feedback in the MPC objective is critical for optimal performance. **Right:** Results for the Trifinger reorientation experiment with cube scale. The ARCtIC method is the most robust controller.

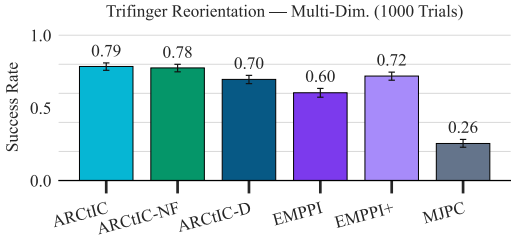

Fig. 5. Results for the 9-dimensional uncertainty distribution experiment on the Trifinger reorientation task. Error bars are 95% Wilson confidence intervals. ARCtIC and its no feedback variant (ARCtIC-NF) perform the best.

In the first Trifinger experiment, we evaluate the robustness of our method to shape uncertainty by randomizing the *scale* of the cube. We provide an uncertainty distribution over side lengths of the cube of $\mathcal{N}(6\text{ cm}, 0.5\text{ cm})$ to distributional controllers, and provide a nominal model of a 6 cm cube to deterministic controllers. Then, we evaluate each controller on different cube sizes for 500 trials each, measuring success rates. Results are shown in the left-hand side of Figure 4. We find that our method tends to outperform other methods across the board. Notably, for variants of our method, the size of the cube had a very asymmetrical effect on success rate.

For the second Trifinger experiment, we evaluate the capability of various methods under a complex, nine-dimensional uncertainty distribution. We use a distribution of uncertainty over center-of-mass, mass, fingertip coefficient of friction, box-ground coefficient of friction, and box side lengths. Results for this experiment are shown in Figure 5. We find that our method is best able to reason about the high-dimensional uncertainty distribution, however we also found that feedback wasn't particularly necessary for this task as ARCtIC and ARCtIC-NF had similar performance. We also found that real-time EMPPI was not capable of outperforming the deterministic version of our method, whereas EMPPI+ was.

VI. Conclusion

We have introduced ARCtIC, a GPU-parallelized CI-MPC algorithm that brings efficient uncertainty-awareness to contact-rich manipulation. Our method also includes a feedback term to prevent overly conservative plans. We demonstrated that our controller provided improvements on two simulated task and was capable of reasoning about shape uncertainty on a simluated 9-DoF Trifinger robot. We believe this constitutes an important step towards real-world uncertainty-awareness during contact-rich manipulation.

REFERENCES

[1] A. Aydinoglu, A. Wei, W.-C. Huang, and M. Posa, "Consensus complementarity control for multi-contact mpc," *IEEE Transactions on Robotics*, 2024.

[2] S. Le Cleac'h, T. A. Howell, S. Yang, C.-Y. Lee, J. Zhang, A. Bishop, M. Schwager, and Z. Manchester, "Fast contact-implicit model predictive control," *IEEE Transactions on Robotics*, vol. 40, pp. 1617–1629, 2024.

[3] V. Kurtz, A. Castro, A. Ö. Önol, and H. Lin, "Inverse dynamics trajectory optimization for contact-implicit model predictive control," *The International Journal of Robotics Research*, vol. 45, no. 1, pp. 23–40, 2026.

[4] W.-C. Huang, A. Aydinoglu, W. Jin, and M. Posa, "Adaptive contact-implicit model predictive control with online residual learning," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 5822–5828, IEEE, 2024.

[5] K. Zhou, J. C. Doyle, and K. Glover, *Robust and Optimal Control*, vol. 40. New Jersey: Prentice Hall, 1996.

[6] F. Muratore, F. Ramos, G. Turk, W. Yu, M. Gienger, and J. Peters, "Robot learning from randomized simulations: A review," *Frontiers in Robotics and AI*, vol. 9, p. 799893, 2022.

[7] E. Aljalbout, J. Xing, A. Romero, I. Akinola, C. R. Garrett, E. Heiden, A. Gupta, T. Hermans, Y. Narang, D. Fox, *et al.*, "The reality gap in robotics: Challenges, solutions, and best practices," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 9, 2025.

[8] Y. Shirai, D. K. Jha, and A. U. Raghunathan, "Covariance steering for uncertain contact-rich systems," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 7923–7929, IEEE, 2023.

[9] I. Mordatch, K. Lowrey, and E. Todorov, "Ensemble-cio: Full-body dynamic motion planning that transfers to physical humanoids," in *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 5307–5314, IEEE, 2015.

[10] J. Jankowski, L. Brudermüller, N. Hawes, and S. Calinon, "Robust pushing: Exploiting quasi-static belief dynamics and contact-informed optimization," *The International Journal of Robotics Research*, p. 02783649251318046, 2025.

[11] Y. Shirai, D. K. Jha, A. U. Raghunathan, and D. Romeres, "Robust pivoting: Exploiting frictional stability using bilevel optimization," in *2022 International Conference on Robotics and Automation (ICRA)*, pp. 992–998, IEEE, 2022.

[12] P. Neal, C. Eric, P. Borja, and E. Jonathan, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Foundations and Trends® in Machine learning*, vol. 3, no. 1, pp. 1–122, 2011.

[13] Y. Shirai, T. Zhao, H. T. Suh, H. Zhu, X. Ni, J. Wang, M. Simchowitz, and T. Pang, "Is linear feedback on smoothed dynamics sufficient for stabilizing contact-rich plans?," in *2025 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 11926–11932, IEEE, 2025.

[14] T. Howell, N. Gileadi, S. Tunyasuvunakool, K. Zakka, T. Erez, and Y. Tassa, "Predictive sampling: Real-time behaviour synthesis with mujoco," *arXiv preprint arXiv:2212.00541*, 2022.

[15] Y. Jiang, M. Yu, X. Zhu, M. Tomizuka, and X. Li, "Contact-implicit model predictive control for dexterous in-hand manipulation: A long-horizon and robust approach," in *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 5260–5266, IEEE, 2024.

[16] W. Jin, "Complementarity-free multi-contact modeling and optimization for dexterous manipulation," in *Proceedings of Robotics: Science and Systems (RSS)*, 2025.

[17] 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.

[18] W. Yang and M. Posa, "Dynamic on-palm manipulation via controlled sliding," in *Robotics: Science and Systems (RSS)*, July 2024.

[19] H. Bui, Y. Gao, H. Yang, E. Cui, S. Mody, B. Acosta, T. S. Felix, B. Bianchini, and M. Posa, "Push anything: Single-and multi-object pushing from first sight with contact-implicit mpc," *arXiv preprint arXiv:2510.19974*, 2025.

[20] E. Ménager, A. Bambade, W. Jallet, A. De Marchi, and J. Carpentier, "Contact-implicit inverse dynamics," 2025.

[21] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics*. MIT Press, 2005.

[22] R. Senanayake, "The role of predictive uncertainty and diversity in embodied ai and robot learning," *arXiv preprint arXiv:2405.03164*, 2024.

[23] L. P. Kaelbling and T. Lozano-Pérez, "Integrated task and motion planning in belief space," *The International Journal of Robotics Research*, vol. 32, no. 9-10, pp. 1194–1227, 2013.

[24] T. Lozano-Perez, M. T. Mason, and R. H. Taylor, "Automatic synthesis of fine-motion strategies for robots," *The International Journal of Robotics Research*, vol. 3, no. 1, pp. 3–24, 1984.

[25] M. Dogar and S. Srinivasa, "A framework for push-grasping in clutter," *Robotics: Science and systems VII*, vol. 1, pp. 65–72, 2011.

[26] M. Li, K. Hang, D. Kragic, and A. Billard, "Dexterous grasping under shape uncertainty," *Robotics and Autonomous Systems*, vol. 75, pp. 352–364, 2016.

[27] H. Wright, W. Zhi, M. Matak, M. Johnson-Roberson, and T. Hermans, "Robust bayesian scene reconstruction with retrieval-augmented priors for precise grasping and planning," *IEEE Robotics and Automation Letters*, vol. 11, no. 1, pp. 49–56, 2025.

[28] I. Abraham, A. Handa, N. Ratliff, K. Lowrey, T. D. Murphey, and D. Fox, "Model-based generalization under parameter uncertainty using path integral control," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 2864–2871, 2020.

[29] K. Chua, R. Calandra, R. McAllister, and S. Levine, "Deep reinforcement learning in a handful of trials using probabilistic dynamics models," *Advances in neural information processing systems*, vol. 31, 2018.

[30] M. Anitescu, "Optimization-based simulation of nonsmooth rigid multibody dynamics," *Mathematical Programming*, vol. 105, no. 1, pp. 113–143, 2006.

[31] W. P. Heemels, J. M. Schumacher, and S. Weiland, "Linear complementarity systems," *SIAM journal on applied mathematics*, vol. 60, no. 4, pp. 1234–1269, 2000.

[32] K. Okamoto and P. Tsiotras, "Optimal stochastic vehicle path planning using covariance steering," *IEEE Robotics and Automation Letters*, vol. 4, no. 3, pp. 2276–2281, 2019.

[33] D. Ruiz, "A scaling algorithm to equilibrate both rows and columns norms in matrices," tech. rep., CM-P00040415, 2001.

[34] T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning*. Springer Series in Statistics, New York, NY, USA: Springer New York Inc., 2001.

[35] E. Todorov, T. Erez, and Y. Tassa, "Mujoco: A physics engine for model-based control," in *2012 IEEE/RSJ international conference on intelligent robots and systems*, pp. 5026–5033, IEEE, 2012.

[36] H. Lu, Z. Peng, and J. Yang, "Mpax: Mathematical programming in jax," *arXiv preprint arXiv:2412.09734*, 2024.

[37] M. Wüthrich, F. Widmaier, F. Grimminger, J. Akpo, S. Joshi, V. Agrawal, B. Hammoud, M. Khadiv, M. Bogdanovic, V. Berenz, *et al.*, "Trifinger: An open-source robot for learning dexterity," *arXiv preprint arXiv:2008.03596*, 2020.

[38] H. Bui and M. Posa, "Enhancing task performance of learned simplified models via reinforcement learning," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 9212–9219, IEEE, 2024.

[39] W. Jin and M. Posa, "Task-driven hybrid model reduction for dexterous manipulation," *IEEE Transactions on Robotics*, vol. 40, pp. 1774–1794, 2024.