

IMPACT: An Implicit Active-Set Augmented Lagrangian for Fast Contact-Implicit Trajectory Optimization

Author Names Omitted for Anonymous Review.

Abstract—*Contact-implicit trajectory optimization (CITO)* has attracted growing attention as a unified framework for planning and control in contact-rich robotic tasks. Recent approaches have demonstrated promising results in manipulation and locomotion without requiring a prescribed contact-mode schedule. It is well known that the underlying *mathematical programs with complementarity constraints* (MPCCs) remain numerically ill-conditioned, and systematic, scalable solution strategies for CITO remain an active area of research. More efficient and principled solvers that can handle contact constraints are therefore essential to broaden the applicability of CITO. In this work, we develop an augmented-Lagrangian approach to CITO for solving MPCC-based CITO with stationarity guarantees. The method can be interpreted as *identifying the implicit contact-mode branches on the fly* during the trajectory optimization (TO) iterations; we call this approach IMPACT (*IMPlicit contact ACTIVE-set Trajectory optimization*). We provide an efficient C++ implementation tailored to trajectory-optimization workloads and evaluate it on the open-source CITO and *contact-implicit model predictive control* (CI-MPC) benchmarks. On CITO, IMPACT achieves $2.9\times$ – $70\times$ speedups over strong baselines (geometric mean $13.8\times$). On CI-MPC, we show improved control quality for contact-rich trajectories on dexterous manipulation tasks in simulation. Finally, we demonstrate the proposed method on real robotic hardware on a T-shaped object pushing task.

I. INTRODUCTION

Optimal motion planning and control in contact-rich manipulation settings remains one of the major challenges in robotics. The planner must reason not only about continuous dynamics, but also about discrete events in which contacts are established, maintained, and broken to perform tasks such as underactuated manipulation and legged locomotion [1], [2]. A common workaround is to prescribe the contact mode sequence a priori; conditioned on this discrete choice, the resulting TO problem can often be formulated as a smooth *nonlinear programming* (NLP) and solved efficiently [3], [4]. By contrast, when the contact schedule is not fixed, the problem is typically posed as CITO. The same contact reasoning is also required in CI-MPC, where the contact mode must be inferred online under tight computation budgets. While CITO/CI-MPC is promising for autonomously inferring contact modes in complex tasks, it remains challenging due to nonsmooth contact dynamics and the combinatorics of mode changes [5], [6], [7].

CITO commonly relies on contact models with a complementarity structure to capture the on/off nature of contact and frictional interaction [8]. This yields a fundamentally nonsmooth optimization problem that is naturally cast as mathematical programs with complementarity constraints (MPCCs) [9]. At feasible points, complementarity induces




Fig. 1. IMPACT demonstrations in simulation and hardware. **Top:** Allegro Hand reorients a rubber duck in simulation. **Bottom:** a Panda robot pushes a T-shaped block to a target pose in real-robot experiments; the red curve indicates the trajectory of the block’s coordinate origin during the push.

degeneracy and violates the regularity assumptions required by standard smooth NLP constraint qualifications (CQs); consequently, standard CQs typically fail, and the associated Lagrange multipliers may become unbounded, rendering the *Karush–Kuhn–Tucker* (KKT) conditions invalid [9], [10]. As a result, off-the-shelf NLP solvers based on standard KKT/CQ theory are often numerically brittle for MPCCs: large initial complementarity violations make convergence highly initialization-sensitive and unreliable at TO scale [11].

Due to the nonsmooth and degenerate nature of contact-implicit formulations, many practical treatments modify the original complementarity structure through smoothing or relaxation, or introduce penalty terms to better fit standard solvers and their convergence assumptions [12], [13], [14]. While these treatments often improve numerical behavior, they can degrade control quality and introduce heuristic update parameters that strongly affect efficiency. Moreover, as these modifications are tightened to approximate the original complementarity conditions, degeneracy and ill-conditioning can re-emerge.

In this work, we propose IMPACT (*IMPlicit contact ACTIVE-set Trajectory optimization*), an *implicit complementarity/contact-branch selection* method for solving MPCC-based CITO, with *stationarity guarantee* for

feasible accumulation points under standard assumptions from recent augmented Lagrangian (AuLa) theory [15], [16], [17], [18]. Instead of smoothing or relaxing complementarity constraints, IMPACT retains the original nonsmooth contact constraints and handles them through the AuLa subproblems. Across these subproblems, the approximate Lagrange multipliers and penalty strength guide the selection of the active complementarity branch, enabling implicit contact-mode discovery without a prescribed homotopy schedule. This AuLa treatment also prevents the multiplier blow-up typically seen in vanilla AuLa formulations for MPCCs, providing a principled alternative to penalty-only methods. Combined with our proposed *block coordinate descent* (BCD) solver for the AuLa subproblems, IMPACT achieves substantially faster convergence than strong baselines while maintaining competitive task performance.

Our contributions are summarized as follows.

- **Algorithm.** We introduce IMPACT, an efficient CITO method based on a safeguarded AuLa scheme for MPCCs, and establish its stationarity guarantee.
- **Implementation.** We provide a C++ implementation of IMPACT with an interface designed for CITO/CI-MPC workloads.
- **Validation.** On open-source CITO benchmarks, IMPACT achieves $2.9\times-70\times$ speedups over strong baselines. On Allegro-Hand CI-MPC, it achieves comparable success rates with improved control-quality metrics, and we further demonstrate the method on real robotic hardware in a Push-T manipulation task.

II. RELATED WORK

CITO and CI-MPC methods can be broadly viewed from two perspectives: approaches that *soften* contact logic to obtain smooth gradients for planning/MPC, and approaches that *retain* nonsmooth contact logic and solve more directly. Many pipelines replace rigid complementarity with *smooth* or *relaxed* surrogates to enable gradient-based optimization, which can improve numerical behavior but may affect gradient accuracy and closed-loop performance. One line modifies the *forward contact model* itself, e.g., using compliant or differentiable contact, yielding smoother derivatives at the expense of strict rigid-contact fidelity [6], [2], [19], [13], [12]. Another line keeps *forward simulation* rigid but regularizes the sensitivity/KKT systems used to compute gradients through contact, enabling fast CI-MPC and relaxed-complementarity DDP variants [20], [21], [22]. Beyond local smoothing, global convex relaxations have also been explored for CITO, emphasizing geometric reasoning and contact sequencing [23]. In contrast, IMPACT adopts a *nonsmooth* AuLa viewpoint: we update relaxation via safeguarded multiplier/penalty rules and solve nonsmooth subproblems that explicitly enforce complementarity, aiming to preserve mode transitions and implicitly identify the contact mode.

A second class handles contact complementarity more *directly*, reflecting the inherent nonsmoothness of contact. Early CITO work introduces elastic variables and penalizes

complementarity violations so that contacts can be “discovered” during optimization [5]. Later pipelines improve reliability via continuation strategies or automatic penalty-update rules [24], [11], and related penalty reformulations also appear in *quadratic programs with complementarity constraints* (QPCC) solvers such as LCQPow [25]. Despite their widespread use, penalty/continuation methods can become numerically stiff under large penalties, and heuristic schedules may slow convergence [11]. Operator-splitting methods such as the C3 family formulate CI-MPC in a consensus ADMM form with simple, parallelizable subproblems, scaling well to many-contact settings [26], [27], [28], [29], [30]. However, these methods typically operate on linearized contact/dynamics and often require task-space tracking layers for execution [26], [30]. A recent preprint by Ménager et al. [31] also uses an AuLa and alternating-minimization/BCD scheme for contact-implicit QPCCs in one-step inverse dynamics. In contrast, IMPACT targets temporally coupled TO and CI-MPC settings with general MPCC-based CITO formulations, using an MPCC-tailored safeguarded AuLa scheme with complementarity-enforcing inner subproblems. This induces an implicit contact-mode selection behavior as slack is driven tight.

III. PROBLEM DEFINITION

We formulate both CITO and CI-MPC in a unified optimization framework in generalized coordinates. Let $X = [x_t]_{1:T}$ denote the trajectory of decision variables, e.g., states, controls, and contact forces. We use $r(x_t)$, $h(x_t, x_{t+1})$, and $g(x_t)$ to denote the residual function, equality constraints, and inequality constraints, respectively, where $h(x_t, x_{t+1})$ includes the system dynamics coupling consecutive time steps. The functions $G(x_t)$ and $H(x_t)$ denote the complementarity functions. The condition $0 \leq G(x_t) \perp H(x_t) \geq 0$ encodes the on/off nature of contact: component-wise, at least one of $G(x_t)$ or $H(x_t)$ must be zero. Algebraically, it is equivalent to $G(x_t) \geq 0$, $H(x_t) \geq 0$, and $G(x_t) \circ H(x_t) = 0$. This yields the following MPCC:

$$\min_X \frac{1}{2} \sum_{t=1}^T r(x_t)^\top r(x_t) \quad (1a)$$

$$\text{s.t. } h(x_t, x_{t+1}) = 0, \quad t = 1, \dots, T-1, \quad (1b)$$

$$g(x_t) \leq 0, \quad t = 1, \dots, T, \quad (1c)$$

$$0 \leq G(x_t) \perp H(x_t) \geq 0, \quad t = 1, \dots, T. \quad (1d)$$

Problem (1) covers different discretizations and contact models as special cases. With a polyhedral approximation of the friction cone, the per-step contact complementarity conditions can be written as an LCP; with conic friction models, the same template leads to nonlinear complementarity constraints. We focus on nonlinear sum-of-squares objectives, which capture many task-space motion-planning costs and are convenient under the augmented-Lagrangian formulation, since both penalty and Lagrangian terms can be absorbed into a nonlinear least-squares residual.

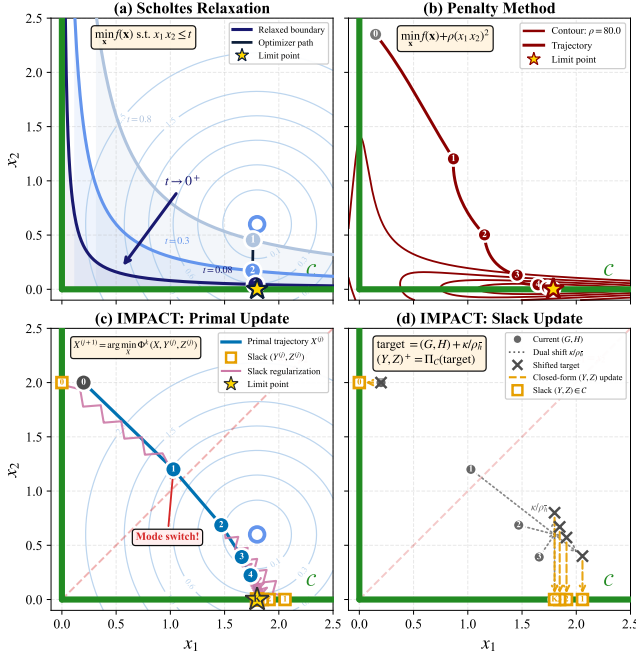


Fig. 2. Comparison of complementarity-handling methods on a 2D toy objective. Top row: Scholtes relaxation and squared-penalty method. Bottom row: IMPACT, split into the primal update and the slack update. Objective contours are shown in blue and the complementarity-feasible set is shown in green. In IMPACT, the slack variables act as anchors that regularize the next primal update, while the slack update is obtained from a dual-shifted target projected onto the complementarity set.

IV. METHOD

IMPACT solves (1) by combining a safeguarded augmented Lagrangian (AuLa) outer loop with a structured block coordinate descent (BCD) inner solver. Instead of smoothing or relaxing complementarity constraints, IMPACT keeps them explicit in each inner subproblem. Across the AuLa iterations, the approximate multipliers and penalty weights guide the active complementarity branch selection, enabling implicit contact-mode discovery without a prescribed continuation schedule. Fig. 2 illustrates this difference from Scholtes relaxation and penalty methods.

For algorithmic convenience, we rewrite (1) in *vertical form* [10] by introducing auxiliary variables Y and Z . We enforce $G(x_t) = y_t$ and $H(x_t) = z_t$, and impose complementarity directly on the variables via $0 \leq y_t \perp z_t \geq 0$. This reformulation is equivalent to the original MPCC at any feasible point, but yields a simple complementarity set that is well suited to closed-form inner updates. Define

$$\bar{h}_t(x_{t+1}, x_t, y_t, z_t) := [h(x_t, x_{t+1}); G(x_t) - y_t; H(x_t) - z_t],$$

and let

$$\mathcal{C} = \{(Y, Z) \mid 0 \leq Y \perp Z \geq 0\}.$$

The vertical-form problem can then be written compactly as

$$\begin{aligned} \min_{X, Y, Z} \quad & \frac{1}{2} \sum_{t=1}^T r(x_t)^\top r(x_t) \\ \text{s.t.} \quad & \bar{h}_t(x_{t+1}, x_t, y_t, z_t) = 0, \quad g(x_t) \leq 0, \quad (Y, Z) \in \mathcal{C}. \end{aligned} \quad (2)$$

We follow a safeguarded AuLa scheme for MPCCs in which complementarity constraints are not absorbed into the augmented Lagrangian; instead, they are enforced explicitly in the inner solve. With equality multiplier κ_t , inequality multiplier μ_t , and penalty factors $\rho_{\bar{h}}$ and ρ_g , the AuLa subproblem is

$$\begin{aligned} \min_{X, Y, Z} \quad & \Phi(X, Y, Z) := \frac{1}{2} \sum_{t=1}^T r_t^\top r_t + \frac{\rho_{\bar{h}}}{2} \sum_{t=1}^T \left\| \bar{h}_t + \frac{\kappa_t}{\rho_{\bar{h}}} \right\|^2 \\ & + \frac{\rho_g}{2} \sum_{t=1}^T \left\| \left(g_t + \frac{\mu_t}{\rho_g} \right)_+ \right\|^2 \end{aligned} \quad (3a)$$

$$\text{s.t.} \quad (Y, Z) \in \mathcal{C}. \quad (3b)$$

The outer loop clips multipliers to fixed safeguard bounds and updates them as

$$\bar{\kappa}_t^{k+1} = \bar{\kappa}_t^k + \rho_{\bar{h}}^k \bar{h}_t, \quad \mu_t^{k+1} = \left(\mu_t^k + \rho_g^k g_t \right)_+.$$

Penalty parameters are increased only when the equality feasibility residual and the inequality KKT residual fail to decrease sufficiently. This separation preserves the MPCC structure and avoids the pathological multiplier growth that can arise when complementarity is penalized directly [16], [15]. Under standard MPCC regularity assumptions, if the inner subproblems are solved with a vanishing stationarity residual, every feasible accumulation point of the safeguarded AuLa sequence is a first-order stationary point.

We solve (3) using a two-block BCD scheme. Each iteration alternates between updating the trajectory variables X and updating the auxiliary complementarity variables (Y, Z) . With (Y, Z) fixed, the X -update reduces to a smooth nonlinear least-squares problem, which we solve using a damped Gauss–Newton step with Armijo backtracking to ensure sufficient decrease of Φ .

With X fixed, the (Y, Z) -subproblem decouples across time steps and complementarity pairs. For each scalar pair $(y_{t,i}, z_{t,i})$ with $0 \leq y_{t,i} \perp z_{t,i} \geq 0$, the feasible set is the union of two convex cones,

$$\mathcal{C}_i = \{(y, 0) \mid y \geq 0\} \cup \{(0, z) \mid z \geq 0\}.$$

Thus the update is obtained by evaluating two closed-form candidates and selecting the one with smaller quadratic penalty contribution. With the safeguarded multipliers at outer iteration k and current penalty factor $\rho_{\bar{h}}$, the two candidates are

$$y_{t,i}^* = \max\left(0, G_i(x_t) + \frac{1}{\rho_{\bar{h}}} \kappa_{G,t,i}\right), \quad z_{t,i}^* = 0, \quad (4)$$

$$z_{t,i}^* = \max\left(0, H_i(x_t) + \frac{1}{\rho_{\bar{h}}} \kappa_{H,t,i}\right), \quad y_{t,i}^* = 0. \quad (5)$$




Fig. 3. IMPACT planning demos on three CITO tasks. The green dashed box marks the start pose and the orange dashed box marks the goal pose. Purple arrows visualize contact forces, and the solid green line traces the object-origin trajectory. Color indicates time from early (light blue) to late (dark orange).

We then compare the corresponding objective values and keep the lower one. This step is an exact minimization over the complementarity set for fixed X , and hence never increases the inner objective. Importantly, the update is not a pure feasibility projection: the current multiplier and penalty shift the target before projection. Therefore, the slack variables act as anchors that regularize the next primal update, as illustrated in Fig. 2. This induces the implicit active-set behavior of IMPACT: the optimizer selects the contact or no-contact branch through the AuLa subproblem rather than through an externally prescribed contact schedule or homotopy parameter.

The resulting inner solver is not intended to certify global optimality of the nonconvex subproblem. Instead, it provides an attainability interface required by the safeguarded AuLa theory. For a fixed outer iteration, assume that Φ is smooth with Lipschitz gradient on the bounded set visited by the iterates, the X -update achieves sufficient decrease, and the (Y, Z) -update is the exact minimization described above. Then, for any tolerance $\varepsilon > 0$, the BCD inner loop can reach an ε -stationary point of (3) in finitely many iterations. In our implementation, we use an inexact/budgeted variant and stop the inner loop when the augmented objective stagnates, $|\Phi(w^{j+1}) - \Phi(w^j)| \leq \tau_k$, which empirically yields stable outer-loop progress on all benchmarks.

V. EXPERIMENTS

We evaluate IMPACT on two open-source benchmarks and a real-robot Push-T task. For long-horizon planning, we use the CRISP CITO benchmark [11]; for dexterous manipulation under multi-contact dynamics, we use the Allegro-Hand CI-MPC benchmark from [13]. All CITO methods are evaluated under the same task objectives and feasibility requirements, including complementarity violation below 10^{-5} .

a) Long-horizon CITO.: We compare IMPACT against Scholtes relaxation (SR), the square penalty method (PM), and CRISP on Push Box, Push T, and Cart Transport, shown in Fig. 3. IMPACT succeeds on all tested instances and maintains competitive task-level tracking performance. In terms of tracking error, SR is best on Push Box and Push T, while CRISP is best on Cart Transport; IMPACT is second-best on Push Box and Cart Transport and remains close to CRISP on Push T. In contrast, PM yields larger tracking errors and fails on some instances of Push Box and Push T.

The main advantage of IMPACT is runtime. Compared with CRISP, IMPACT is $16.8\times$, $25.0\times$, and $34.0\times$ faster on Push Box, Push T, and Cart Transport, respectively, giving a geometric mean speedup of $24.3\times$. Even compared with the fastest baseline in wall-clock time, IMPACT achieves $5.7\times$, $2.9\times$, and $25.0\times$ speedups on the three tasks. Across all baseline comparisons, this corresponds to a $2.9\times$ – $70\times$ speedup range with a geometric mean of $13.8\times$.

b) Multi-contact CI-MPC.: We further evaluate IMPACT on the Allegro-Hand meshable-object reorientation benchmark [13], comparing against the smoothing-based *cfree* baseline with two control bounds. IMPACT achieves a comparable success rate: $91.8\% \pm 4.1$ versus $91.2\% \pm 4.2$ for *cfree*(0.1) and $92.9\% \pm 3.9$ for *cfree*(0.2). More importantly, it improves control quality, with lower control variance (1.55 ± 0.13 vs. $2.16 \pm 0.21/4.44 \pm 0.49$), smoothness (0.029 ± 0.004 vs. $0.044 \pm 0.008/0.088 \pm 0.012$), and effort (3.0 ± 1.0 vs. $4.8 \pm 2.0/8.2 \pm 4.8$). IMPACT runs at 9.53 ± 0.36 Hz, lower than *cfree*'s 50–60 Hz rate, but directly solves a multi-contact, nonsmooth MPCC-based MPC problem without smoothing or relaxation. The stick object remains the main outlier, where the long slender geometry makes contact maintenance sensitive to LCP-based mode switching.

c) Real-robot Push-T.: Finally, we deploy IMPACT on a real robotic Push-T setup. We solve a full-horizon TO problem to obtain a pushing trajectory and replan online when the object deviates from the nominal trajectory. Across 10 trials with different initial object configurations, IMPACT reaches the target in all trials, achieving a 100% success rate.

VI. CONCLUSION

We presented IMPACT, an implicit contact active-set augmented Lagrangian method for fast contact-implicit trajectory optimization. IMPACT combines a safeguarded AuLa outer loop for MPCCs with a structured BCD inner solver, enabling contact-mode discovery without a prescribed continuation schedule while retaining the original nonsmooth complementarity structure. Experiments show substantial speedups on long-horizon CITO, improved control quality on multi-contact CI-MPC, and successful transfer to real robotic hardware. These results suggest that MPCC-tailored AuLa methods are a promising alternative to smoothing, relaxation, and penalty-only approaches for contact-rich planning and control.

REFERENCES

- [1] K. Bouyarmane, S. Caron, A. Escande, and A. Kheddar, "Multi-contact motion planning and control," in *Humanoid Robotics: A Reference*, A. Goswami and P. Vadakkepat, Eds. Springer Nature, 2019.
- [2] T. Pang, H. T. Suh, L. Yang, and R. Tedrake, "Global planning for contact-rich manipulation via local smoothing of quasi-dynamic contact models," *IEEE Transactions on Robotics*, vol. 39, no. 6, pp. 4691–4711, 2023.
- [3] W. Jallet, A. Bambade, E. Arlaud, S. El-Kazdadi, N. Mansard, and J. Carpentier, "Proxddp: Proximal constrained trajectory optimization," *IEEE Transactions on Robotics*, 2025.
- [4] C. Mastalli, R. Budhiraja, W. Merkt, G. Saurel, B. Hammoud, M. Naveau, J. Carpentier, L. Righetti, S. Vijayakumar, and N. Mansard, "Crocodyl: An efficient and versatile framework for multi-contact optimal control," in *Proc. IEEE International Conference on Robotics and Automation (ICRA)*, 2020.
- [5] 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.
- [6] A. Ö. Önel, P. Long, and T. Padiş, "Contact-implicit trajectory optimization based on a variable smooth contact model and successive convexification," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 2447–2453.
- [7] J.-P. Sleiman, F. Farshidian, and M. Hutter, "Versatile multicontact planning and control for legged loco-manipulation," *Science Robotics*, vol. 8, no. 81, p. eadg5014, 2023.
- [8] Q. Le Lidec, W. Jallet, L. Montaut, I. Laptev, C. Schmid, and J. Carpentier, "Contact models in robotics: a comparative analysis," *IEEE Transactions on Robotics*, 2024.
- [9] H. Scheel and S. Scholtes, "Mathematical programs with complementarity constraints: Stationarity, optimality, and sensitivity," *Mathematics of Operations Research*, vol. 25, no. 1, pp. 1–22, 2000.
- [10] A. Nurkanović, A. Pozharskiy, and M. Diehl, "Solving mathematical programs with complementarity constraints arising in nonsmooth optimal control: A. nurkanović et al." *Vietnam Journal of Mathematics*, vol. 53, no. 3, pp. 659–697, 2025.
- [11] Y. Li, H. Han, S. Kang, J. Ma, and H. Yang, "On the Surprising Robustness of Sequential Convex Optimization for Contact-Implicit Motion Planning," in *Proceedings of Robotics: Science and Systems*, Los Angeles, CA, USA, June 2025.
- [12] V. Kurtz, A. Castro, A. Ö. Önel, 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.
- [13] W. Jin, "Complementarity-Free Multi-Contact Modeling and Optimization for Dexterous Manipulation," in *Proceedings of Robotics: Science and Systems*, Los Angeles, CA, USA, June 2025.
- [14] S. Katayama, T. Tani, and K. Tanaka, "Quasistatic contact-rich manipulation via linear complementarity quadratic programming," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 203–210.
- [15] X. Jia, C. Kanzow, P. Mehlitz, and G. Wachsmuth, "An augmented lagrangian method for optimization problems with structured geometric constraints," *Mathematical Programming*, vol. 199, no. 1, pp. 1365–1415, 2023.
- [16] L. Guo and Z. Deng, "A new augmented lagrangian method for mpcs—theoretical and numerical comparison with existing augmented lagrangian methods," *Mathematics of Operations Research*, vol. 47, no. 2, pp. 1229–1246, 2022.
- [17] A. De Marchi, X. Jia, C. Kanzow, and P. Mehlitz, "Constrained composite optimization and augmented lagrangian methods," *Mathematical Programming*, vol. 201, no. 1, pp. 863–896, 2023.
- [18] C. Kanzow and P. Mehlitz, "Convergence properties of monotone and nonmonotone proximal gradient methods revisited," *Journal of Optimization Theory and Applications*, vol. 195, no. 2, pp. 624–646, 2022.
- [19] H. J. T. Suh, T. Pang, and R. Tedrake, "Bundled gradients through contact via randomized smoothing," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 4000–4007, 2022.
- [20] 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.
- [21] G. Kim, D. Kang, J.-H. Kim, and H.-W. Park, "Contact-implicit differential dynamic programming for model predictive control with relaxed complementarity constraints," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 11 978–11 985.
- [22] G. Kim, D. Kang, J.-H. Kim, S. Hong, and H.-W. Park, "Contact-implicit model predictive control: Controlling diverse quadruped motions without pre-planned contact modes or trajectories," *The International Journal of Robotics Research*, vol. 44, no. 3, pp. 486–510, 2025.
- [23] B. P. Graesdal, S. Y. C. Chia, T. Marcucci, S. Morozov, A. Amice, P. Parrilo, and R. Tedrake, "Towards Tight Convex Relaxations for Contact-Rich Manipulation," in *Proceedings of Robotics: Science and Systems*, Delft, Netherlands, July 2024.
- [24] A. Ö. Önel, R. Corcodel, P. Long, and T. Padiş, "Tuning-free contact-implicit trajectory optimization," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 1183–1189.
- [25] J. Hall, A. Nurkanović, F. Messerer, and M. Diehl, "LCQPow: a solver for linear complementarity quadratic programs," *Mathematical Programming Computation*, vol. 17, no. 1, pp. 81–109, 2025.
- [26] A. Aydinoglu and M. Posa, "Real-time multi-contact model predictive control via admm," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 3414–3421.
- [27] A. Aydinoglu, A. Wei, W.-C. Huang, and M. Posa, "Consensus complementarity control for multi-contact mpc," *IEEE Transactions on Robotics*, 2024.
- [28] 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)*. IEEE, 2024, pp. 5822–5828.
- [29] S. Venkatesh, B. Bianchini, A. Aydinoglu, W. Yang, and M. Posa, "Approximating global contact-implicit mpc via sampling and local complementarity," *IEEE Robotics and Automation Letters*, 2025.
- [30] 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.
- [31] E. Ménager, P. Fabre, A. Bambade, W. Jallet, A. de Marchi, and J. Carpentier, "Frictional Contact-Implicit Inverse Dynamics," Aug. 2025, working paper or preprint. [Online]. Available: <https://hal.science/hal-05201780>