Robustness without Wrinkles: Parallel Simulation & Robust MPC for Certified Deformable Manipulation

Wei-Chen Li*, Jeffrey Fang*, Sasanka Polisetti, Yuexi Song, Glen Chou
Georgia Institute of Technology, *Equal Contribution
{wli777, jfang301, spolisetti6, ysong644, chou}@gatech.edu
Project Website: (Link) **Code**: (Github) **Video**: (YouTube)

***Abstract*—We present CORD-SLS, a real-time control method for safe deformable object manipulation, with a focus on ropes and cloth. CORD-SLS combines a GPU-parallel differentiable simulator with contact smoothing and a GPU-parallel output-feedback robust MPC controller to enable efficient planning under intermittent contact, model uncertainty, and perception error. We show that the simulator accelerates model-based RL training for manipulation policies. To improve real-world robustness, we use conformal prediction to calibrate perception-error bounds for MPC, producing reachable tubes that enable high-probability safe control. We evaluate CORD-SLS on high-dimensional, contact-rich rope and cloth manipulation tasks in simulation and hardware, achieving millisecond-speed planning and exceeding baselines in safety, speed, and task success.**

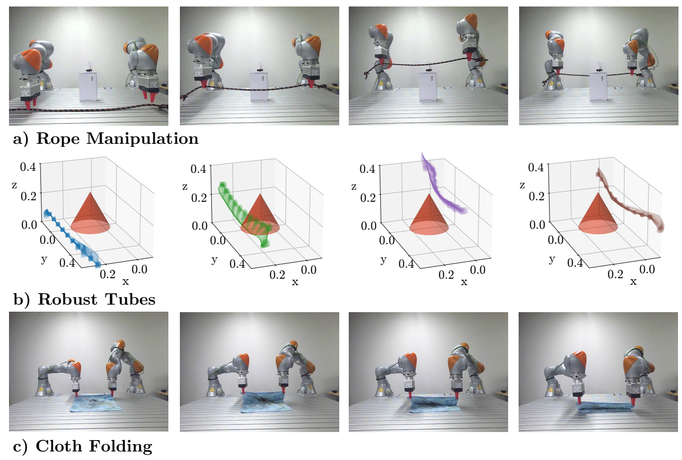

Fig. 1: Hardware results showcase CORD-SLS's real-world efficacy. In **a)**, CORD-SLS solves a constrained rope manipulation task with reachable tubes ensuring robust constraint satisfaction **b)** when avoiding the obstacle (red). In **c)**, CORD-SLS extends to cloth folding.

I. Introduction

Manipulation of deformable objects, like ropes and cloth, is important for safety-critical applications in home robots, logistics, and surgery. Yet it requires planning over high-dimensional nonlinear dynamics while reasoning about contact and uncertainty. Prior work leaves a gap between scalable planning, deformable contact dynamics, and robustness. To close these gaps, we propose CORD-SLS (Contact-smoothed Output-feedback Robust control of Deformables via System Level Synthesis), a framework for real-time differentiable simulation and robust perception-based manipulation planning for deformable objects. Our contributions are:

- A GPU-parallel, differentiable, real-time, contact-smoothed deformable object simulator, allowing for planning-friendly contact discovery.
- A GPU-parallel method for real-time synthesis of robust output-feedback control policies via reachability-constrained SLS, and gradient-accelerated model-based reinforcement learning using analytical policy gradients.
- A real-world vision-based implementation using keypoint tracking with conformal uncertainty calibration, enabling probabilistic robustness guarantees under perception and modeling uncertainty.
- Evaluation on contact-rich rope and cloth manipulation tasks in simulation and hardware, outperforming baselines in safety, speed, and task success.

II. Related Work

A. Deformable Object Simulation

Position-based dynamics (PBD) and its extension, XPBD, simulate deformable objects by projecting particles onto geometric constraints, enabling fast and stable simulations. However, PBD-based methods achieve robustness at the expense of physical fidelity and existing differentiable variants do not fully address rigid-deformable contact and friction. While Cosserat and discrete elastic rod models capture rich deformation behavior and frictional contact [11], contact-rich planning remains challenging due to the need for explicit contact-mode discovery [5]. Existing differentiable cloth simulators can efficiently model thin-shell dynamics and contact, but planning similarly remains difficult due to contact discontinuities [12].

B. Planning

Deformable-object planning is difficult because perception, dynamics, and control are tightly coupled and models must trade off between fidelity and speed [2]. Classical planners use reduced geometric abstractions or hybrid planning-control schemes for tractability [13, 16], but often produce open-loop plans with limited robustness. Learned dynamics models [14] enable planning but require extensive training and provide limited robustness guarantees. Differentiable simulators enable gradient-based planning [4], but discontinuous gradients impede contact discovery for planning. Our method instead uses contact smoothing to plan through contact without bespoke skills. Safety and robustness are less-studied. Some methods plan nominally-safe motions [8] or use barrier functions [15], but provide only local safety guarantees and do not address contact discovery. While safe output-feedback control has been considered in recent work [6], these methods struggle to scale to high-dimensional deformable objects.

III. Problem Statement

We consider a known state representation $q \in \mathcal{Q} \subseteq \mathbb{R}^{n_q}$ for a deformable object of interest (e.g., 3D positions of links for rope, or of mesh node positions for cloth) actuated by control input $u \in \mathcal{U} \subseteq \mathbb{R}^{n_u}$ and governed by uncertain dynamics (1a), where $f : \mathcal{Q} \times \mathcal{U} \to \mathcal{Q}$. To account for perception latency, we model the system as being observed with uncertain time-delayed measurements $y \in \mathcal{Y} \subseteq \mathbb{R}^{n_y}$ according to an observation map $h : \mathcal{Q} \to \mathcal{Y}$ (1b),

$$q_{k+1} = f(q_k, u_k) + E(q_k) w_k, \tag{1a}$$

$$y_{k+1} = h(q_k) + F(q_k) e_k, \tag{1b}$$

where $E : \mathcal{Q} \to \mathbb{R}^{n_q \times n_q}$ and $F : \mathcal{Q} \to \mathbb{R}^{n_y \times n_y}$ and $w \in \mathcal{B}^{n_q} := \{w \in \mathbb{R}^{n_q} \mid \|w\|_2 \leq 1\}$ and $e \in \mathcal{B}^{n_y}$ are bounded disturbances. We consider two specific problems:

Problem 1. *Differentiable contact-smoothed deformable simulator.* Develop a simulator for deformable linear objects and cloth that 1) is differentiable, 2) GPU-parallel, and 3) provides informative gradients for planning and policy learning through contact, defining the planner dynamics

$$q_{k+1} = f(q_k, u_k). \tag{2}$$

Problem 2. *Reachability-informed MPC for robust deformable manipulation.* Optimize a nominal length-N state $\mathbf{z} := \{z_k\}_{k=0}^{N}$ and control $\mathbf{v} := \{v_k\}_{k=0}^{N-1}$ trajectory that is feasible for the simulator dynamics (2), together with a causal output-feedback controller $\pi := (\pi_0, \ldots, \pi_{N-1})$, where $\pi_t : \mathcal{Y}^{t+1} \to \mathcal{U}$ stabilizes the true deformable dynamics (1a) about $(\mathbf{z}, \mathbf{v})$ and ensures closed-loop satisfaction of constraints $\mathcal{S} \subseteq \mathcal{Q} \times \mathcal{U}$.

IV. Method

We develop a differentiable deformable-object simulator with contact smoothing for gradient-based planning (Sec. IV-A), a GPU-parallel output-feedback SLS framework (Sec. IV-B), a sample-efficient model-based RL strategy using APG (Sec. IV-C), and a CP-based calibration of perception errors into reachable-tube guarantees (Sec. IV-D).

A. Differentiable, Smoothed Deformable Object Simulator

Existing deformable simulators often assume that deformable objects remain attached to the robot. We introduce a differentiable simulator that captures switching between inactive and active constraints. Consider a nodal system with n_v nodes, whose 3D coordinates, together with the n_u actuated degrees of freedom (DoFs), form the configuration vector $q \in \mathbb{R}^{n_q}$, where $n_q = 3n_v + n_u$. For an elastic energy $E(q)$, the corresponding internal elastic force is given by $-\partial E(q)/\partial q$. Under a quasi-static assumption, the system satisfies

$$M \frac{q^+ - q}{\delta t^2} + \frac{\partial E(q^+)}{\partial q^+} = \tau(u) + \sum_{i=1}^{n_c} J_i(q)^\top \lambda_i, \tag{3}$$

where M is the mass matrix, $\tau(u)$ are the external generalized forces, including gravity acting on the nodal system and actuation forces applied to the actuated DoFs, $\lambda_i \in \mathbb{R}^3$ is the ith constraint force, δt is the discrete timestep, and $J_i \in \mathbb{R}^{3 \times n_q}$ is the corresponding constraint Jacobian. For holonomic constraints, such as a gripper rigidly interacting with the nodal system, we suppose an active constraint $c_i(q) = 0 \in \mathbb{R}^3$ enforces the gripper attachment. The next configuration q^+ and constraint force λ_i must satisfy the complementarity condition

$$0 = J_i(q)\,(q^+ - q) + c_i(q) \ \perp\ \lambda_i \in \mathbb{R}^3, \tag{4}$$

where the constraint Jacobian is defined as $J_i(q) := \partial c_i(q)/\partial q$. For contact constraints, we adopt Anitescu's convex relaxation [1]. Here, the next step configuration q^+ and contact force λ_i satisfy

$$\mathcal{F}_\mu^* \ni J_i(q)\,(q^+ - q) + \begin{bmatrix}\phi_i(q) & 0 & 0\end{bmatrix}^\top \ \perp\ \lambda_i \in \mathcal{F}_\mu, \tag{5}$$

where $\phi_i(q)$ is the signed distance of contact pair i, $\mathcal{F}_\mu$ is the friction cone, and $\mathcal{F}_\mu^*$ is its dual cone.

Simulation Solver We solve the dynamics in (3) subject to the holonomic and contact constraints in (4) and (5) using a two-stage procedure. In the first stage, we solve for an intermediate configuration q^* satisfying

$$\begin{aligned} F(q^*; q, u) &:= M \frac{q^* - q}{\delta t^2} + \frac{\partial E(q^*)}{\partial q^*} - \tau(u) \\ &\quad + \sum_{i=1}^{n_h} J_i(q)^\top D_i \left(J_i(q)\,(q^* - q) + c_i(q)\right) = 0, \end{aligned} \tag{6}$$

where n_h denotes the number of holonomic constraints, and $D_i \in \mathbb{R}^{3 \times 3}$ is a diagonal stiffness matrix associated with the ith holonomic constraint. When the constraint is active, D_i is chosen sufficiently large to enforce (4); otherwise, $D_i = 0$. In the second stage, we solve the optimization problem

$$\begin{aligned} \min_{q^+} \quad & \tfrac{1}{2}\|q^+ - q^*\|_P^2 \\ \text{s.t.} \quad & J_i(q^*)(q^+ - q^*) + \begin{bmatrix}\phi_i(q^*) & 0 & 0\end{bmatrix} \in \mathcal{F}_\mu^*, \\ & \forall i = n_h + 1, \ldots, n_c. \end{aligned} \tag{7}$$

where $P := M + \frac{\partial^2 E(q^*)}{\partial (q^*)^2} + \sum_{i=1}^{n_h} J_i(q)^\top D_i J_i(q)$. The KKT conditions of (7) recover the contact constraint in (5) together with $P\,(q^+ - q^*) = \sum_{i=n_h+1}^{n_c} J_i(q)^\top \lambda_i$, which corresponds to a linearization of the dynamics about the intermediate configuration q^*. Given a (q, u) pair, solving (6)–(7) gives the next configuration q^+. The solvers for (6)–(7) are implemented in JAX to leverage GPU-accelerated matrix factorizations. Their derivatives are computed using the implicit function theorem, i.e., by implicitly differentiating the KKT conditions of the associated optimization problems.

Contact Smoothing Contact dynamics often exhibit a vanishing-gradient issue: prior to contact activation, the control inputs associated with the actuated DoFs have no influence on the state of the deformable object. As a result, gradient-based optimizers may stagnate due to the absence of informative gradient signals, failing to reach the goal.

To address this issue, we adopt the idea of surrogate gradients, in which the gradient used for planning is replaced by a surrogate estimator that does not correspond to the true gradient obtained by differentiating through the physical rollout. This introduces a controlled bias in exchange for improved gradient information in regimes where the true gradient is uninformative. The holonomic constraint in (4) is active only when $c_i(q) = 0$, at which point a nonzero stiffness matrix D_i is used in (6) to enforce the constraint. This switching between inactive and active constraints introduces both discontinuities and vanishing gradients. To smooth this transition, we redefine D_i as $D_i = D_{i,\max} \exp(-\|c_i(q)\|_2/\kappa)$, where κ is a smoothing parameter. As $\kappa \to 0$, this formulation approaches the original indicator-like activation behavior.

The contact constraint in (5), which encodes non-contact, stiction, and sliding regimes, similarly introduces discontinuous and vanishing gradients. We smooth this by relaxing the complementarity condition in (5) to $(J_i(q)\,(q^+ - q) + \begin{bmatrix}\phi_i(q) & 0 & 0\end{bmatrix})^\top \lambda_i = \kappa$, which recovers the original complementarity constraint when $\kappa = 0$. The smoothed dynamics is used solely for gradient computation and the nonsmooth dynamics for forward rollouts, preserving model accuracy.

B. Simulator-Informed GPU-Parallel Output Feedback SLS

To plan robust output feedback policies for the deformable object, we exploit the differentiability of the simulator to obtain the linearized dynamics, $\mathbf{A}, \mathbf{B}$, and observation model, $\mathbf{C}$, about a nominal trajectory and solve the output-feedback SLS problem of [10],

$$
\begin{aligned}
\min_{\substack{\mathbf{z},\mathbf{v},\\ \boldsymbol{\Phi},\boldsymbol{\tau}}} \quad & J(\mathbf{z}, \mathbf{v}, \boldsymbol{\tau}, \boldsymbol{\Phi}) && \text{(8a)}\\
\text{s.t.} \quad & z_{k+1} = f(z_k, v_k), \qquad \forall k \in [N] && \text{(8b)}\\
& \begin{bmatrix}\mathbf{I} - \mathbf{Z}\mathbf{A}(\mathbf{z},\mathbf{v}), & -\mathbf{Z}\mathbf{B}(\mathbf{z},\mathbf{v})\end{bmatrix} \boldsymbol{\Phi} = \begin{bmatrix}\mathbf{I}, & \mathbf{0}\end{bmatrix}, && \text{(8c)}\\
& \boldsymbol{\Phi} \begin{bmatrix}(\mathbf{I} - \mathbf{Z}\mathbf{A}(\mathbf{z},\mathbf{v}))^\top, & (-\mathbf{Z}\mathbf{C}(\mathbf{z},\mathbf{v}))^\top\end{bmatrix}^\top = \begin{bmatrix}\mathbf{I}, & \mathbf{0}\end{bmatrix}^\top, && \text{(8d)}\\
& \sum_{j=0}^{k-1} \left\| c_i^\top \boldsymbol{\Phi}^{\mathrm{w}}_{k,j} E_j \right\|_2 + \left\| c_i^\top \boldsymbol{\Phi}^{\mathrm{e}}_{k,j} F_j \right\|_2 + c_i^\top \begin{bmatrix} z_k \\ v_k \end{bmatrix} \le -b_i && \text{(8e)}\\
& \qquad \forall i \in [n_c],\ \forall k \in [N], && \text{(8f)}\\
& \textstyle\sum_{j=0}^{k-1} \left\|\boldsymbol{\Phi}^{\mathrm{w}}_{k,j} E_j\right\|_2 + \left\|\boldsymbol{\Phi}^{\mathrm{e}}_{k,j} F_j\right\|_2 \le \tau_k, \quad \forall k \in [N]. && \text{(8g)}
\end{aligned}
$$

Here, $\boldsymbol{\Phi}^{\mathrm{w}}$ and $\boldsymbol{\Phi}^{e}$ denote the closed-loop responses from process disturbances and observation disturbances respectively. Constraints (8f)–(8g) use these responses to compute reachable tubes and tighten constraints such that all trajectories remain within the safe set despite disturbances. GPU-SLS [7] directly computes the state-feedback responses $\boldsymbol{\Phi}^{\mathrm{w}}$. Following [10], the observer responses $\boldsymbol{\Phi}^{\mathrm{e}}$ can be recovered through the time-varying Kalman gains L_k from the Kalman recursions,

$$
\begin{aligned}
L_{j+1} &= -\left(F_j F_j^\top + C_j \Pi_j C_j^\top\right)^{-1} C_j \Pi_j A_j^\top, \\
\Pi_{j+1} &= E_j E_j^\top + A_j \Pi_j A_j^\top + A_j \Pi_j C_j^\top L_{j+1}.
\end{aligned} \tag{9}
$$

[10] calculates these recursions sequentially, creating a bottleneck for long-horizon deformable-object planning. Similar to GPU-SLS, we reformulate the Kalman recursions as associative operators and evaluate them using parallel prefix scans on the GPU. The resulting gains are used to construct $\boldsymbol{\Phi}^{\mathrm{e}}$, yielding output-feedback reachable tubes that jointly account for dynamic and perception disturbance in real-time.

C. Accelerating MBRL via Analytical Policy Gradients (APG)

The differentiable simulator proposed in Sec. IV-A can be used to support model-based reinforcement learning (MBRL) in multiple ways; we point out two ways here. First, it can generate expert trajectories to guide policy learning, improving sample efficiency in contact-rich manipulation where intermittent contact makes exploration difficult. Second, by incorporating the contact-smoothed simulator as a differentiable layer, direct policy gradients can guide policies toward discovering contact with reduced reliance on stochastic exploration, improving the sample efficiency of policy learning. We denote this MBRL variant of our method as APG, as it leverages analytical policy gradients to accelerate policy learning.

Task	Method	# Sim Eval.	Time Per-Step [ms]	Safety Rate [%]	Goal Error [m]
Lift Rope	Ours	$5.4 \cdot 10^3$	50	100.0	$5.48 \cdot 10^{-2}$
	MPPI	$1.0 \cdot 10^6$	180	100.0	$6.81 \cdot 10^{-2}$
	PPO	$1.6 \cdot 10^7$	0.4	96.5	$5.87 \cdot 10^{-2}$
	APG	$1.6 \cdot 10^7$	0.4	100.0	$5.73 \cdot 10^{-2}$
Drag Rope	Ours	$1.4 \cdot 10^4$	70	99.9	$3.60 \cdot 10^{-2}$
	MPPI	$2.6 \cdot 10^6$	230	1.0	$8.14 \cdot 10^{-2}$
	PPO	$4.0 \cdot 10^7$	0.4	97.0	$3.20 \cdot 10^{-2}$
	APG	$4.0 \cdot 10^7$	0.4	93.0	$3.02 \cdot 10^{-2}$
Fold cloth	Ours	$3.4 \cdot 10^3$	760	100.0	$6.11 \cdot 10^{-2}$
	MPPI	$6.4 \cdot 10^4$	1900	64.0	$7.38 \cdot 10^{-2}$
Flatten cloth	Ours	$3.4 \cdot 10^3$	710	100.0	$2.34 \cdot 10^{-2}$
	MPPI	$6.4 \cdot 10^4$	1100	80.0	$5.26 \cdot 10^{-2}$

TABLE I: Runtime and performance metrics across tasks.

D. Calibrated Reachability Guarantees

To account for sim-to-real mismatch and perception error, we use conformal prediction (CP) to calibrate an uncertainty bound using one-step hardware rollouts. On hardware, the true dynamics error and state-estimation error are not independently observable, as we only have access to the perceived state returned by perception. We therefore calibrate the aggregate one-step prediction residual between the model-predicted next state and the next perceived state measured by perception. This residual contains the combined effects of model mismatch and perception error. We then use CP to obtain a high-probability disturbance bound, which we incorporate into our output-feedback SLS formulation to compute reachable tubes with probabilistic safety guarantees.

V. RESULTS

We compare CORD-SLS and APG against MPPI and model-free RL methods (namely, PPO) on deformable manipulation (Sec. V-A) and evaluate the efficacy of our method on hardware (Sec. V-B).

A. Deformable Manipulation Benchmarks

We benchmark our approach on four deformable manipulation tasks; we show two in Fig. 2. In the lift rope task, a pair of grippers must establish contact with the rope to activate the holonomic constraint required to lift it. Our method leverages gradients from the dynamics to guide actions toward making contact. In contrast, MPPI requires substantially more dynamics evaluations to identify contact-making actions, resulting

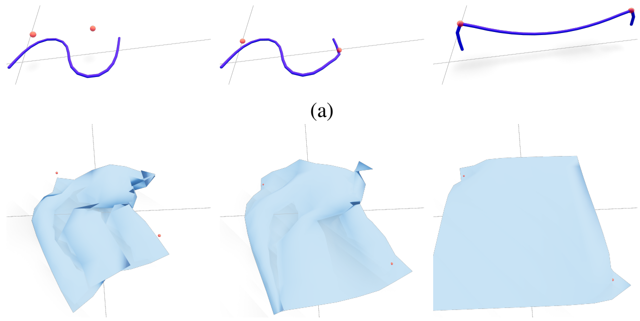

(a)

(b)

Fig. 2: Our method on (a) Rope lifting and (b) Cloth flattening.

in higher execution times. Among the RL methods, APG leverages differentiable dynamics to compute policy gradients, whereas PPO is warm-started with 10 expert demonstrations generated by CORD-SLS to enable it to solve the task within the same interaction budget. Without this warm start, PPO fails to make contact with the rope after the same number of environment interactions, and fails to complete the task.

In the drag rope task, the gripper must establish contact with the rope and transport it to a goal location while avoiding an obstacle. CORD-SLS explicitly accounts for disturbances and computes trajectory tubes, enabling robust obstacle avoidance. Although obstacle avoidance is encoded as a penalty in the MPPI cost function, the best sampled trajectories often pass through the obstacle, leading to a very low safety rate. Increasing the obstacle penalty further causes the rope to remain near its initial position. APG finds a lower-cost policy than PPO with fewer simulator samples, but both RL methods exhibit constraint violations because they do not explicitly account for disturbances during execution.

For the cloth folding and flattening tasks, despite the high-dimensional state space (300 DoF), the GPU-parallelized simulator and SLS solver allow each MPC update to be computed in under one second. MPPI requires significantly more dynamics evaluations to achieve comparable performance, resulting in longer execution times as well as higher goal error and constraint violation rates. We omit RL baselines for the 300-DoF cloth because while it remains tractable for optimization, learning a neural policy is prohibitively sample-intensive.

B. Hardware

Lastly, we validate our approach on hardware. We consider two deformable manipulation tasks: obstacle rope manipulation and cloth folding. We estimate the state using RGB-D observations, where SAM [3] grounds keypoints that are tracked in real time by TAPNext++ [9]. We then use these keypoint poses to reconstruct the object state. In Fig. 1**(a)**, we show that, due to the contact-smoothness and differentiability of our simulator, the manipulators are able to make contact with the rope and navigate it over the obstacle. In Fig. 1**(b)** we visualize the robust tubes generated at selected timestamps which account for both dynamical and measurement uncertainty, showing that we remain safe despite uncertainty. In Fig. 1**(c)** we show the manipulators successfully completing a cloth folding task, demonstrating that our formulation extends to high-dimensional deformable object tasks.

VI. Conclusion

We presented CORD-SLS, a framework for safe, contact-rich deformable object manipulation that integrates differentiable simulation, robust GPU-accelerated output-feedback, and data-driven uncertainty calibration, enabling real-time contact-rich planning with high-probability safety guarantees.

References

[1] Mihai Anitescu. Optimization-based simulation of nonsmooth rigid multibody dynamics. *Mathematical Programming*, 105(1):113–143, 2006.

[2] Veronica E Arriola-Rios, Puren Guler, Fanny Ficuciello, Danica Kragic, Bruno Siciliano, and Jeremy L Wyatt. Modeling of deformable objects for robotic manipulation: A tutorial and review. *Frontiers in Robotics and AI*, 7:82, 2020.

[3] Nicolas Carion, Laura Gustafson, Yuan-Ting Hu, Shoubhik Debnath, Ronghang Hu, Didac Suris, Chaitanya Ryali, Kalyan Vasudev Alwala, Haitham Khedr, Andrew Huang, et al. Sam 3: Segment anything with concepts. *arXiv preprint arXiv:2511.16719*, 2025.

[4] Siwei Chen, Yiqing Xu, Cunjun Yu, Linfeng Li, Xiao Ma, Zhongwen Xu, and David Hsu. Daxbench: Benchmarking deformable object manipulation with differentiable physics. *arXiv preprint arXiv:2210.13066*, 2022.

[5] Yizhou Chen, Yiting Zhang, Zachary Brei, Tiancheng Zhang, Yuzhen Chen, Julie Wu, and Ram Vasudevan. Differentiable discrete elastic rods for real-time modeling of deformable linear objects. *arXiv preprint arXiv:2406.05931*, 2024.

[6] Glen Chou, Necmiye Ozay, and Dmitry Berenson. Safe output feedback motion planning from images via learned perception modules and contraction theory. In *International Workshop on the Algorithmic Foundations of Robotics*, pages 349–367. Springer, 2022.

[7] Jeffrey Fang and Glen Chou. Safe large-scale robust nonlinear MPC in milliseconds via reachability-constrained system level synthesis on the GPU. *arXiv preprint arXiv:2604.07644*, 2026.

[8] Jing Huang, Xiangyu Chu, Xin Ma, and Kwok Wai Samuel Au. Deformable object manipulation with constraints using path set planning and tracking. *IEEE Transactions on Robotics*, 39(6):4671–4690, 2023.

[9] Sebastian Jung, Artem Zholus, Martin Sundermeyer, Carl Doersch, Ross Goroshin, David Joseph Tan, Sarath Chandar, Rudolph Triebel, and Federico Tombari. Tapnext++: What's next for tracking any point (tap)? *arXiv preprint arXiv:2604.10582*, 2026.

[10] Antoine P Leeman, Shuyu Zhan, Melanie N Zeilinger, and Glen Chou. Vision-SLS: Safe perception-based control from learned visual representations via system level synthesis. *arXiv preprint arXiv:2604.24894*, 2026.

[11] Wei-Chen Li and Glen Chou. A convex formulation of compliant contact between filaments and rigid bodies. *arXiv preprint arXiv:2509.13434*, 2025.

[12] Junbang Liang, Ming Lin, and Vladlen Koltun. Differentiable cloth simulation for inverse problems. In *Advances in Neural Information Processing Systems*, volume 32, 2019.

[13] Mitul Saha and Pekka Isto. Motion planning for robotic manipulation of deformable linear objects. In *Proceedings 2006 IEEE International Conference on Robotics and Automation, 2006. ICRA 2006.*, pages 2478–2484. IEEE, 2006.

[14] Bokui Shen, Zhenyu Jiang, Christopher Choy, Silvio Savarese, Leonidas J Guibas, Anima Anandkumar, and Yuke Zhu. Action-conditional implicit visual dynamics for deformable object manipulation. *The International Journal of Robotics Research*, 43(4):437–455, 2024.

[15] Yunxi Tang, Xiangyu Chu, Jing Huang, and KW Samuel Au. Learning-based mpc with safety filter for constrained deformable linear object manipulation. *IEEE Robotics and Automation Letters*, 9(3):2877–2884, 2024.

[16] Yutong Zhang, Fei Liu, Xiao Liang, and Michael Yip. Achieving autonomous cloth manipulation with optimal control via differentiable physics-aware regularization and safety constraints. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 9931–9938. IEEE, 2024.