ICRA2026 Late Breaking Results Poster presented at
2026 IEEE International Conference on Robotics and Automation (ICRA 2026)
June 1-5, 2026. Vienna, Austria

Risk-Aware Control of Tendon-Driven Continuum Robots via CVaR-MPPI with Residual Learning for Hysteresis Compensation : A Pilot Study

DongJun Lee[1] and Dongwook Kim*
Department of Robotics and Mechatronics Engineering, DGIST
`djl3372@gmail.com, dw_kim@dgist.ac.kr`
*Corresponding author

***Abstract*—Tendon-driven Continuum Robots (TDCRs) are widely used in confined operating systems due to their thin shape, flexibility, and compliance making them easily deployable in narrow or contact-rich environments. However, real-time safe control near obstacles remains challenging. Computationally expensive dynamic models, such as the Cosserat rod model, are impractical for real-time control. Conventional model predictive control (MPC) methods require linearization of the dynamics, limiting their applicability to the complex nonlinear behavior of TDCRs, including hysteresis. In this paper, we adopt the Piecewise Constant Curvature (PCC) model, which assumes constant curvature for each link. While computationally cheap, this approximation contains modeling errors that, combined with mechanical friction, backlash, and misalignment at the rolling joints, result in unpredictable hysteresis. Also, we propose CVaR-MPPI(Conditional Value-at-Risk Model Predictive Path Integral), a controller that combines sampling based planning with probability safety under uncertainty environment, improving both worst-case risk managing and sampling efficiency. In simulation with 100 iterations, CVaR-MPPI improves the success rate from 80% to 85% and the mean safety clearance by 129%, while maintaining end-effector tracking error compared to standard MPPI, as detailed in the simulation results. The controller runs at 50Hz with 8192 samples, demonstrating real-time feasibility.**

I. INTRODUCTION

Continuum robots with tendon-driven actuation offer inherent compliance that makes them promising for tasks in confined and contact-rich environments such as minimally invasive surgery, in-pipe inspection, and industrial maintenance. However, achieving real-time safety control near obstacles is challenging due to Computationally expensive dynamic model(Cosserat rod model). In particular, hysteresis caused by mechanical friction, backlash and misalignment combined with the approximation of the widely used Piecewise Constant Curvature (PCC) kinematic model, introduces significant and unpredictable deviations between the commanded and actual poses of the robot.

Existing approaches to safety control for continuum robots face notable limitations under uncertainty environment. Fiber Bragg Grating (FBG) sensors provide accurate shape sensing but are expensive and require complex calibration. Cosserat rod models offer high-fidelity dynamics but are too computationally expensive for real-time control. Deterministic safety methods such as Control Barrier Functions (CBFs) and Hamilton-Jacobi (HJ) require exact dynamical models, which becomes unreliable under the hysteresis-induced model mismatch inherent to our rigid rolling-joint structure. We adopt rigid rolling joints for their rapid, low-cost prototyping while clearly exhibiting hysteresis from accumulated contact friction - making is a suitable testbed for validating risk-aware control. Unlike these deterministic methods, CVaR-based safety analysis does not require an exact model - it quantifies risk by focusing on the worst-case tail of the cost distribution, remaining robust under model uncertainty [2].

Standard Model Predictive Path Integral (MPPI) control [1] evaluates thousands of sampled trajectories but relies on the average or sum of obstacle penalties, which makes it vulnerable to edge-case collisions. To address this, we propose CVaR-MPPI, which incorporates Conditional Value-at-Risk (CVaR) into the MPPI framework, inspired by recent work on risk-aware MPPI [3]. Unlike standard MPPI that optimizes around the mean cost, CVaR-MPPI focuses on the worst $(1-\alpha)K$ highest-cost samples, explicitly penalizing high-risk trajectories that standard MPPI ignores.

Contributions. (1) We integrate CVaR-based risk measure into MPPI for a tendon-driven continuum robot with a PCC kinematic model. (2) We demonstrate in simulation that CVaR-MPPI significantly improves safety margins without sacrificing tracking performance. (3) We show real-time feasibility at 50 Hz on commodity GPU hardware.

A. Hardware Platform and Kinematic Model

Our platform is a tendon-driven continuum robot with rigid rolling joints. We adopt this design for its rapid, low-cost prototyping, while it clearly exhibits hysteresis from accumulated contact friction at the rolling joints — making it a suitable testbed for validating risk-aware control. We employ the PCC model, which assumes each of the n links bends uniformly at θ_d/n, approximating constant-curvature arcs. While computationally cheap, this approximation, combined with mechanical friction, backlash, and misalignment at the rolling joints, results in unpredictable hysteresis and a

Algorithm 1 CVaR-MPPI

```
Require: State $\boldsymbol{\theta}_t$, goal $\mathbf{x}_{\mathrm{goal}}$, horizon $H$, samples $K$, CVaR $\alpha$, temperature $\lambda$
Ensure: Optimal control $\mathbf{u}_0^{\mathrm{opt}}$
/* 1. Sample K trajectories & compute cost */
for $i = 1$ to $K$ do
    Sample $\boldsymbol{\epsilon}_k^{(i)} \sim \mathcal{N}(0, \Sigma)$; rollout $\tilde{\mathbf{u}}_k^{(i)} = \mathbf{u}_k + \boldsymbol{\epsilon}_k^{(i)}$
    $\mathbf{x}_k^{(i)} \leftarrow \mathrm{FK}(f_\theta(\mathbf{l}_k^{(i)}))$
    Compute trajectory cost $J_i$
end for
/* 2. CVaR Risk Score (per sample) */
for $i = 1$ to $K$ do
    $r_k^{(i)} \leftarrow \sum_n \left(\max(0, \mathrm{SafetyMargin} - \mathrm{dist}_{k,n}^{(i)})\right)^2$
    $\mathcal{S}_i \leftarrow \frac{1}{\beta} \log\left(\frac{1}{H} \sum_k \exp(\beta \cdot r_k^{(i)})\right)$
end for
/* 3. CVaR Importance Weighting & Control Update */
$\mathrm{VaR}_\alpha \leftarrow \mathrm{Quantile}_{1-\alpha}(\{\mathcal{S}_i\}_{i=1}^K)$
$\mathcal{I}_{\mathrm{worst}} \leftarrow \{i : \mathcal{S}_i \geq \mathrm{VaR}_\alpha\}$
for $i = 1$ to $K$ do
    $\hat{J}_i \leftarrow J_i + w_{\mathrm{cvar}} \frac{(\mathcal{S}_i - \mathrm{VaR}_\alpha)_+}{\alpha}$
end for
Compute weights $w_i$ from $\hat{J}_i$ for $i \in \mathcal{I}_{\mathrm{worst}}$; $w_i = 0$ otherwise
$\mathbf{U}_{\mathrm{new}} \leftarrow \mathbf{U} + \sum_{i=1}^K w_i \boldsymbol{\epsilon}^{(i)}$;    $\mathbf{u}_0^{\mathrm{opt}} \leftarrow \mathbf{U}_{\mathrm{new}}[0]$
```

growing mismatch between PCC predictions and the actual robot pose.

B. CVaR-MPPI Algorithm

The proposed controller is summarized in Algorithm 1. At each time step, we sample K trajectories through the PCC forward kinematics, compute per-sample risk scores based on obstacle proximity, and apply CVaR-based importance weighting to penalize high-risk trajectories.

II. SIMULATION RESULTS

We evaluate CVaR-MPPI against standard MPPI over 100 iterations with parameters $K = 8192$, $H = 10$ (0.5 s), and $\alpha = 0.05$. Results are summarized in Table I and Fig. 1.

TABLE I
COMPARISON OF STANDARD MPPI VS. CVAR-MPPI (100 ITER).

Metric	**CVaR OFF**	**CVaR ON**
Success rate	80%	85%
Mean clearance	3.82 mm	8.76 mm
Worst 5% clearance	1.07 mm	3.61 mm
EE tracking error	90.8 mm	93.4 mm

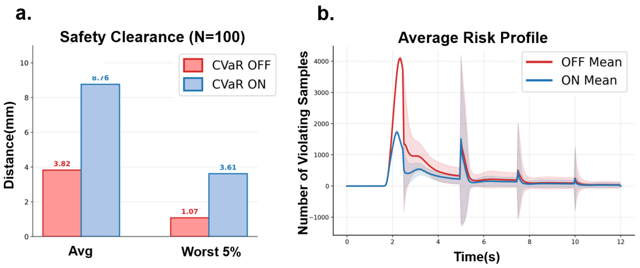

Fig. 1. (a) Safety clearance comparison (mean and worst 5%). CVaR-MPPI improves both metrics significantly. (b) Average risk profile over time. CVaR-MPPI suppresses the peak number of violating samples from ∼4,000 to ∼1,750.

CVaR-MPPI improves the success rate from 80% to 85% and the mean safety clearance by 129% (from 3.82 mm to 8.76 mm). The worst 5% clearance increases from 1.07 mm to 3.61 mm, providing a buffer against sensor noise and model mismatch. Peak violating samples are reduced from ∼4,000 to ∼1,750. End-effector tracking error remains comparable (90.8 mm vs. 93.4 mm). The controller runs at 50 Hz with $K = 8192$ samples and $H = 10$ steps. Remaining failures are attributed to the limited 0.5 s planning horizon near obstacles. Extending the horizon involves a trade-off with real-time performance, which can be addressed through GPU parallelization in future work.

III. CONCLUSION AND FUTURE WORK

We presented CVaR-MPPI, a risk-aware sampling-based controller for tendon-driven continuum robots that improves obstacle avoidance safety under hysteresis-induced model uncertainty. By focusing on the worst-case tail of sampled trajectory costs, CVaR-MPPI provides significantly improved safety margins while maintaining real-time performance.

Ongoing and future work includes: (1) *Visual state estimation* using stereo cameras and load cells with Bayesian shape estimation, to replace simulation ground-truth state with real sensor feedback; (2) *Dynamic obstacle avoidance* extending CVaR-MPPI to moving obstacles via time-varying risk constraints; (3) *Residual learning* for hysteresis compensation, using a data-driven neural network trained on real-world data to learn the nonlinear hysteresis behavior that cannot be modeled in simulation, compensating for PCC model mismatch in closed-loop.

REFERENCES

[1] G. Williams, N. Wagener, B. Goldfain, P. Drews, J. M. Rehg, B. Boots, and E. A. Theodorou, “Information theoretic MPC for model-based reinforcement learning,” in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, 2017, pp. 1714–1721.

[2] M. P. Chapman, R. Bonalli, K. M. Smith, I. Yang, M. Pavone, and C. J. Tomlin, “Risk-sensitive safety analysis using conditional value-at-risk,” *IEEE Trans. Autom. Control*, vol. 67, no. 12, pp. 6521–6536, 2022.

[3] J. Yin, Z. Zhang, and P. Tsiotras, “Risk-aware model predictive path integral control using conditional value-at-risk,” in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, 2023, pp. 3702–3708.

ACKNOWLEDGMENTS

we used AI-assisted tools for coding support during the development of this work.