

# Uncertainty-Aware Non-Prehensile Manipulation with Mobile Manipulators under Object-Induced Occlusion

Jiwoo Hwang<sup>1</sup>, Taegeun Yang<sup>2</sup>, Jeil Jeong<sup>1</sup>, Minsung Yoon<sup>2</sup> and Sung-Eui Yoon<sup>2†</sup>

**Abstract**—Non-prehensile manipulation using onboard sensing presents a fundamental challenge: the manipulated object occludes the sensor’s field of view, creating occluded regions that can lead to collisions. We propose *CURA-PPO*, a reinforcement learning framework that addresses this challenge by explicitly modeling uncertainty under partial observability. By predicting collision possibility as a distribution, we extract both risk and uncertainty to guide the robot’s actions. The uncertainty term encourages active perception, enabling simultaneous manipulation and information gathering to resolve occlusions. When combined with confidence maps that capture observation reliability, our approach enables safe navigation despite severe sensor occlusion. Extensive experiments across varying object sizes and obstacle configurations demonstrate that *CURA-PPO* achieves up to  $3\times$  higher success rates than the baselines, with learned behaviors that handle occlusions. Our method provides a practical solution for autonomous manipulation in cluttered environments using only onboard sensing.

## I. INTRODUCTION

Non-prehensile manipulation, which controls objects without grasping, is commonly used to handle items that exceed grasping capabilities due to size, weight, or surface properties [1]–[9]. Mobile platforms—including quadrupeds [5], wheeled robots [7], and mobile manipulators [9]—extend these capabilities to larger workspaces, enabling tasks such as transporting boxes in warehouses or repositioning furniture. However, a fundamental challenge arises when these platforms rely solely on onboard sensors (e.g., LiDAR): the manipulated object itself occludes the sensors’ fields of view, creating dangerous blind regions during manipulation. This occlusion-induced uncertainty becomes critical when unexpected obstacles appear—such as items falling from shelves or personnel entering the workspace—thereby requiring real-time perception and avoidance. Object-induced occlusion in non-prehensile tasks remains underexplored in the literature, despite its practical importance for safe autonomous manipulation in cluttered environments with only local sensing.

To address the perceptual challenges posed by object-induced occlusion, we leverage uncertainty-aware decision-making strategies. Recent work has shown that when agents explicitly quantify uncertainty and incorporate its reduction as an objective, they exhibit information-seeking behaviors—actively exploring ambiguous regions rather than relying on incomplete observations [10]–[14]. This principle

<sup>1</sup>J. Hwang and J. Jeong are with the Robotics Program at the Korea Advanced Institute of Science and Technology (KAIST), Daejeon, 34141, Republic of Korea.

<sup>2</sup>T. Yang, M. Yoon and S. Yoon are with the School of Computing at the same institute, KAIST.

<sup>†</sup>S. Yoon is a corresponding author. [sungeui@kaist.edu](mailto:sungeui@kaist.edu)

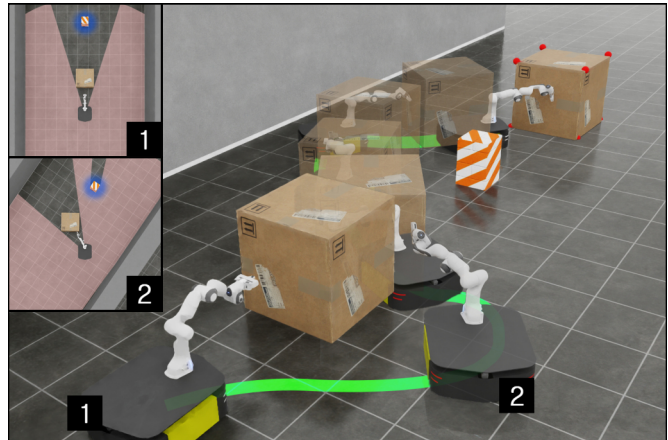


Fig. 1: **Uncertainty-Aware Non-prehensile Manipulation under Object-Induced Occlusion.** A mobile manipulator pushes an object to the target using local onboard sensing. (1) The object induces occlusion, hiding an unforeseen obstacle. (2) The mobile manipulator adjusts its behavior to actively reduce uncertainty from occlusion, revealing the hidden obstacle for safe non-prehensile manipulation.

directly applies to our setting: by incorporating uncertainty reduction alongside the manipulation objective, the robot proactively resolves occlusions while maintaining task progress. Such uncertainty-driven exploration proves essential when the manipulated object blocks sensor views, as it encourages the robot to strategically maneuver for better observability without neglecting its manipulation goal.

Building on these principles, we propose *CURA-PPO* (Collision Uncertainty-Risk Aware PPO), a reinforcement learning (RL) framework that enables safe non-prehensile manipulation with local sensing under occlusions.

In summary, our main contributions are *threefold*:

- We formulate non-prehensile manipulation under object-induced occlusion as an uncertainty-aware decision-making problem, where a mobile manipulator manages cluttered environments using only an onboard 2D LiDAR while the manipulated object creates occluded regions.
- We propose *CURA-PPO*, an RL framework incorporating a Distributional Collision Estimator (*DCE*) that predicts collision possibility as a distribution. By integrating the distribution’s mean (*risk*) and variance (*uncertainty*) as intrinsic costs, our approach jointly optimizes task completion, collision avoidance, and uncertainty reduction.
- We demonstrate through extensive simulation experiments that *CURA-PPO* significantly outperforms the baselines across varying object sizes and obstacle configurations, with qualitative analysis revealing learned active perception behaviors that proactively resolve occlusions.

## II. RELATED WORK

### A. Non-Prehensile Manipulation

*Non-prehensile* manipulation refers to a class of interactions that control an object by applying unilateral constraints (e.g., pushing), without the force closure provided by a grasp’s bilateral constraints [1]. This approach is effective for manipulating objects beyond a robot’s physical limits, such as its payload capacity or grasp range [2]–[9]. Building on these advantages, research in non-prehensile manipulation spans tasks at varying scales, from precise object handling on tabletops with fixed-base arms [2]–[4] to long-range manipulation in extended environments with mobile robots [5]–[9].

For long-range manipulation, prior work has shown promising results, yet it assumes complete knowledge of all obstacles in the environment, including their locations and geometries. In practical scenarios where such information is not readily available, robots rely on onboard sensors. This setting introduces challenges from occlusions, as the manipulated object blocks the sensors’ views and limits environmental awareness. To address these challenges, we propose an uncertainty-aware manipulation policy under partial observability from onboard sensors. Specifically, our approach incorporates active sensing—purposeful movements to gather information—to handle uncertainty from occlusions while balancing uncertainty reduction with task execution.

### B. Uncertainty-Aware Decision-Making

Decision making in partially observable environments requires handling uncertainty from incomplete sensing [15]. Such uncertainty has been quantified as the variance of model predictions, either through distributional representation learning [16]–[18], or sampling-based approximations [19]–[21]. To enhance robustness under partial observability, recent work incorporates uncertainty estimates into decision making [22]–[25]. Moreover, using uncertainty reduction as a learning signal encourages information-seeking behaviors, enabling agents to recognize and resolve perceptual ambiguities [10]–[14]. We apply this principle to non-prehensile manipulation with local sensing, where occlusions reduce observability. Our approach quantifies uncertainty through variance in collision possibility prediction: pushing toward observable regions yields low variance, while pushing into occluded areas produces high variance. Incorporating this variance into policy objective, we encourage active perception behaviors that reduce occlusions during manipulation.

## III. VARIABLE NOTATIONS

This section summarizes the key notations and variables used throughout this paper. Cartesian position, linear velocity, and linear acceleration are denoted as  $\mathbf{p}$ ,  $\mathbf{v}$ , and  $\dot{\mathbf{v}} \in \mathbb{R}^3$ . XYZ Euler angles  $\boldsymbol{\theta}$  represent the orientation, with angular velocity  $\boldsymbol{\omega} \in \mathbb{R}^3$ . For the seven-DOF manipulator, joint position, velocity, and acceleration are  $\mathbf{q}$ ,  $\dot{\mathbf{q}}$ , and  $\ddot{\mathbf{q}} \in \mathbb{R}^7$ . Superscripts denote reference frames, while subscripts indicate entities, axis components, and time steps. For example,  $p_{g,x,t}^b$  represents the  $x$ -component of the goal position in the base frame at time step  $t$ , where  $w, b, e, o$ , and  $g$  denote the world, base, end-effector, object, and goal frames, respectively.

## IV. PRELIMINARIES

This section provides the necessary background on Partially Observable Markov Decision Process (POMDP) and Proximal Policy Optimization (PPO) for understanding the subsequent methodological contributions. In contrast to fully observable MDPs, a POMDP is defined by the tuple  $(\mathcal{S}, \mathcal{A}, \mathcal{T}, \mathcal{R}, \Omega, \mathcal{O}, \gamma)$ , where agents receive observations  $\mathbf{o} \in \Omega$  rather than direct access to states  $\mathbf{s} \in \mathcal{S}$ . Here,  $\mathcal{A}$  is the action space,  $\mathcal{T}$  the transition dynamics,  $\mathcal{R}$  the reward function,  $\mathcal{O}$  the observation function mapping states into the observation space  $\Omega$ , and  $\gamma$  the discount factor.

The goal of reinforcement learning is to learn a policy  $\pi_\phi(\mathbf{a}_t|\mathbf{o}_t)$  that maximizes the expected cumulative reward. PPO [26] optimizes this objective through iterative policy improvement using a surrogate objective:

$$\mathcal{J}^{PPO}(\phi) = \mathbb{E}_{(\mathbf{o}_t, \mathbf{a}_t) \sim \pi_{\phi_{\text{old}}}} \left[ \frac{\pi_\phi(\mathbf{a}_t|\mathbf{o}_t)}{\pi_{\phi_{\text{old}}}(\mathbf{a}_t|\mathbf{o}_t)} \cdot \hat{A}_t \right] \quad (1)$$

where the probability ratio  $r(\phi) = \frac{\pi_\phi(\mathbf{a}_t|\mathbf{o}_t)}{\pi_{\phi_{\text{old}}}(\mathbf{a}_t|\mathbf{o}_t)}$  measures policy change. In practice, the probability ratio is clipped to encourage stable policy updates. The advantage function  $A(\mathbf{s}_t, \mathbf{a}_t) = Q(\mathbf{s}_t, \mathbf{a}_t) - V(\mathbf{s}_t)$  compares the value of taking a specific action to the average value at that state, where  $Q$  and  $V$  represent action-value and state-value functions, respectively. During training, the estimated advantage  $\hat{A}_t$  serves as a learning signal in policy optimization [27]: positive values encourage the policy to increase the probability of the chosen action, whereas negative values decrease it.

## V. METHOD

We present *Collision Uncertainty-Risk Aware PPO (CURA-PPO)*, a modified reinforcement learning algorithm for long-range non-prehensile manipulation under occlusions caused by the manipulated object itself. Our approach is composed of two key components: a *Distributional Collision Estimator (DCE)* that predicts the distribution of collision possibility, and a modified PPO objective that enhances safety by reducing uncertainty and minimizing collision risk. Subsequent sections present the task formulation, policy architecture, DCE implementation, and training procedure.

### A. Task and Environment Formulation

We formulate the task as a goal-conditioned non-prehensile manipulation problem in the presence of obstacles, using a mobile manipulator that relies solely on a base-mounted 2D LiDAR. The sensor provides partial observations, as the manipulated object occludes portions of its field of view—a fundamental constraint that introduces perceptual uncertainty during manipulation. Fig. 3 illustrates our simulation environment, detailing the spatial layout for the task, including start, goal, and obstacle regions. The object and obstacles vary in size, with obstacles appearing at different locations and times during the episode, simulating scenarios such as boxes suddenly falling from shelves or tools left in the workspace. This variability motivates active perception during manipulation, as reducing uncertainty in occluded regions improves safety.

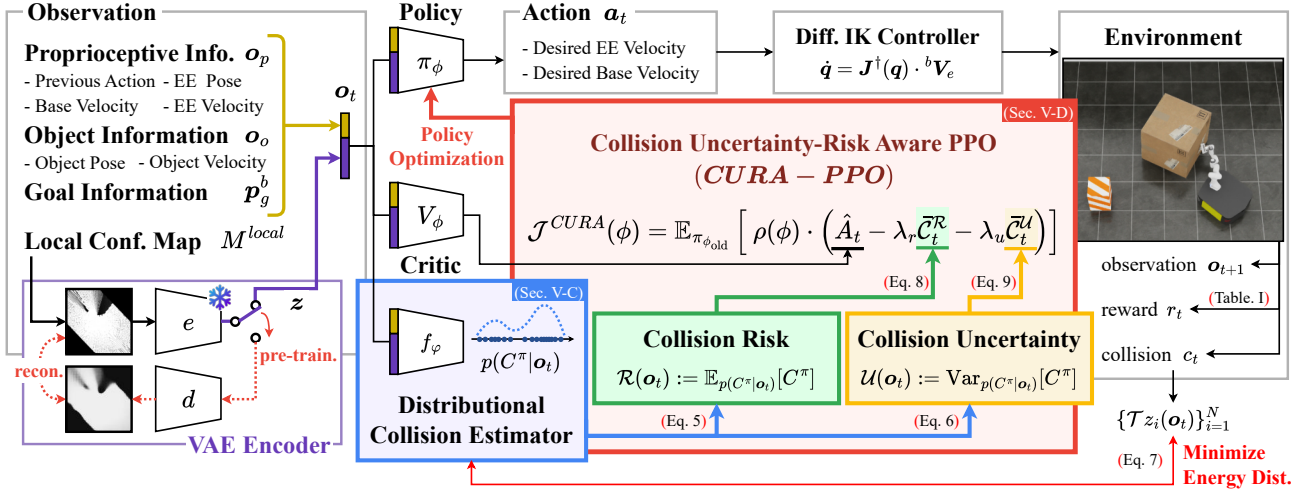


Fig. 2: **Overall Framework.** Local observations including proprioceptive, object, goal, and confidence map features are processed by the policy network  $\pi_\phi$  (see Sec. V-B), which outputs velocity commands executed through differential inverse kinematics. The Distributional Collision Estimator (DCE,  $f_\phi$ ) predicts collision possibility as a distribution, extracting risk and uncertainty signals (see Sec. V-C) that augment the PPO objective in *CURA-PPO*, encouraging behaviors that reduce both collision risk and perceptual uncertainty (see Sec. V-D).

### B. Non-prehensile Manipulation Policy under Uncertainty

As illustrated in Fig. 2, our manipulation policy  $\pi_\phi$  maps partial LiDAR observations and the goal pose to robot actions for collision-free object manipulation. We formulate this as a Partially Observable Markov Decision Process (POMDP) with a policy  $\pi_\phi(\mathbf{a}_t|\mathbf{o}_t)$  that maximizes the expected return.

1) *Sensor Encoding:* We encode the robot’s LiDAR observations into a compact spatiotemporal representation that accounts for both current measurements and their reliability over time. Simply accumulating past LiDAR scans expands spatial coverage but lacks temporal distinction. This becomes problematic when unforeseen obstacles appear in previously observed but currently occluded regions, treating potentially dangerous areas as safe based on outdated observations. We address this issue using confidence maps [28], where sensor readings decay exponentially over time—keeping recent measurements reliable while fading outdated ones:

$$M_t(x, y) = \begin{cases} 1, & \text{if } (x, y) \text{ observed at } t \\ \alpha \cdot M_{t-1}(x, y), & \text{otherwise} \end{cases} \quad (2)$$

where  $M_t$  represents the global confidence map at time step  $t$ , with values  $M_t(x, y) \in [0, 1]$  at each pixel (0.1m resolution) and a decay factor  $\alpha \in (0, 1)$ . For encoding, we extract a local window  $M_t^{\text{local}} \in \mathbb{R}^{100 \times 100}$  centered at the robot’s base from this global map and encode it into a latent vector  $z_t$  using a Variational Autoencoder (VAE) pretrained on random trajectories and frozen during policy learning.

2) *Policy Architecture:* Our policy  $\pi_\phi : \Omega \rightarrow \mathcal{A}$  maps the 4-tuple observation  $\mathbf{o} = (\mathbf{o}_p, \mathbf{o}_o, \mathbf{p}_g^b, z)$  to an action. The observation, expressed in the base frame, consists of: (1) proprioceptive information  $\mathbf{o}_p = [\mathbf{p}_e^b, \mathbf{v}_e^b, \mathbf{v}_b^b, \mathbf{a}_{t-1}]$  including the end-effector pose and velocity, velocity, base velocity, and previous action; (2) object information  $\mathbf{o}_o = [\mathbf{p}_o^b, \mathbf{v}_o^b]$  with current pose and velocity; (3) goal pose  $\mathbf{p}_g^b$ ; and (4) an encoded confidence map  $z$  providing partial environmental information. The policy outputs velocity commands  $\mathbf{a} = [\mathbf{v}_b^{b, \text{des}}, \mathbf{v}_e^{b, \text{des}}]$  for the

TABLE I: Manipulation Policy Reward  $\mathcal{R} = \mathcal{R}_{\text{task}} + \mathcal{R}_{\text{reg}}$

Reward Term	Expression
<b>Task Rewards:</b> $\mathcal{R}_{\text{task}} = \sum_{k=1}^4 r_k$	
Keypoint ( $r_1$ )	$w_1 \cdot \exp(-0.25 \ \mathbf{G}_g^w - \mathbf{G}_o^w\ _2)$
Progress ( $r_2$ )	$w_2 \cdot \exp(5.0 (\mathbf{v}_o^w \cdot \hat{\mathbf{p}}_{o \rightarrow g}^w - 1))$
Velocity ( $r_3$ )	$w_3 \cdot \ \mathbf{v}_o^w\ _2 / v_{\text{max}}$
Contact Point ( $r_4$ )	$w_4 \cdot \exp(-0.1 \ \mathbf{p}_r^w - \mathbf{p}_e^w\ _2)$
<b>Regularization Rewards:</b> $\mathcal{R}_{\text{reg}} = \sum_{k=5}^6 r_k$	
Action Rate ( $r_5$ )	$w_5 \cdot \ \mathbf{a}_t - \mathbf{a}_{t-1}\ _2$
Smooth Action ( $r_6$ )	$w_6 \cdot \ \ddot{\mathbf{q}}\ _2$
<b>Reward Weights</b>	
$w_1=4.0$	$w_2=2.0$ $w_3=1.0$ $w_4=1.0$ $w_5=-0.1$ $w_6=-0.1$

- $\hat{\cdot}$  denotes a unit vector.
- $v_{\text{max}}$  indicates the maximum speed of mobile base.
- $\mathbf{G} \in \mathbb{R}^{24}$  is the keypoint vector proposed in [5], consisting of the 3D coordinates of the eight object corners.
- $\mathbf{p}_r$  is a randomly sampled point on the object surface following [9].

omnidirectional base and end-effector, which are converted to joint-space targets via differential inverse kinematics and executed through joint impedance control.

3) *Implementation Details:* We use a ClearPath Ridgeback with a 7-DOF Franka arm and a 2D LiDAR with a 360° field of view, operating at a maximum base speed  $v_{\text{max}}$  of 0.8 m/s. Training was conducted in Isaac Sim [29] with 2048 parallel environments on an NVIDIA RTX 4090 GPU, using Proximal Policy Optimization (PPO) [26] to optimize the policy. The policy operates at 5 Hz, generating velocity commands, while joint impedance control runs at 200 Hz. The differential IK employs the Damped Least Squares (DLS) method with a damping factor  $\lambda = 0.02$ . The confidence map uses an exponential update with  $\alpha = 0.9$ . Table I summarizes the reward terms used for training.

The keypoint reward minimizes the distance between corresponding object and goal keypoints [5], while the contact point reward guides the end-effector toward multiple points sampled on the target surface, encouraging diverse pushing behaviors [9].

### C. Distributional Collision Estimator

While our manipulation policy can execute pushing actions based on partial observations, it struggles to generate movements that actively reduce occlusion-induced uncertainty. To address this challenge, we introduce a *Distributional Collision Estimator (DCE)* that models collision possibility as a distribution, enabling optimization of both its expectation (*risk*) and variance (*uncertainty*) during policy learning.

1) *Distribution of Collision Possibility*: We define collision possibility  $C^\pi$  as a random variable representing the discounted accumulation of collision indicators while following policy  $\pi_\phi$  to capture the urgency of future collisions:

$$C^\pi := \sum_{k=0}^{\infty} \gamma_c^k c_{t+k}, \quad c_t = \mathbb{1}_{\text{collision at time } t} \in \{0, 1\} \quad (3)$$

where  $\gamma_c \in (0, 1)$  is a discount factor, set to 0.9, and  $c_t$  indicates collision at time  $t$ . To obtain the inherent uncertainty arising from partial observability, we approximate the conditional distribution  $p(C^\pi | \mathbf{o}_t)$  using quantile regression [30]. Our DCE network  $f_\varphi : \Omega \rightarrow \mathbb{R}^N$  predicts  $N$  quantiles:

$$f_\varphi(\mathbf{o}_t) = \{z_i(\mathbf{o}_t)\}_{i=1}^N \approx \{q_{\tau_i}(C^\pi | \mathbf{o}_t)\}_{i=1}^N, \quad \tau_i = \frac{i - 0.5}{N} \quad (4)$$

where  $q_{\tau_i}(C^\pi | \mathbf{o}_t)$  denotes the  $i$ -th quantile of  $p(C^\pi | \mathbf{o}_t)$ .

2) *Risk and Uncertainty Estimation*: We define *risk* and *uncertainty* as the conditional expectation and variance of  $C^\pi$  given the observation  $\mathbf{o}_t$ . From the DCE's quantile predictions  $\{z_i(\mathbf{o}_t; \varphi)\}_{i=1}^N$ , we estimate the risk and uncertainty:

$$\mathcal{R}(\mathbf{o}_t) := \mathbb{E}_{p(C^\pi | \mathbf{o}_t)}[C^\pi] \approx \frac{1}{N} \sum_{i=1}^N z_i(\mathbf{o}_t; \varphi) \quad (5)$$

$$\mathcal{U}(\mathbf{o}_t) := \text{Var}_{p(C^\pi | \mathbf{o}_t)}[C^\pi] \approx \frac{1}{N} \sum_{i=1}^N (z_i(\mathbf{o}_t; \varphi) - \mathcal{R}_\varphi(\mathbf{o}_t))^2 \quad (6)$$

We denote these empirical estimates as  $\mathcal{R}_\varphi$  and  $\mathcal{U}_\varphi$  hereafter. Risk and uncertainty capture different aspects of collision prediction: high risk indicates likely collisions, whereas high uncertainty reflects unreliable predictions from partial observations. These complementary metrics guide distinct behaviors of the policy—risk drives collision avoidance while uncertainty promotes active perception to resolve ambiguity.

3) *Training the Distributional Collision Estimator*: We train the DCE network  $f_\varphi$  by minimizing the energy distance between predicted and target quantile distributions [31]:

$$\mathcal{L}_{\text{DCE}}(\varphi) = 2\mathbb{E}_{i,j}[z_i - \mathcal{T}z_j] - \mathbb{E}_{i,j}[\mathcal{T}z_i - \mathcal{T}z_j] - \mathbb{E}_{i,j}[z_i - z_j] \quad (7)$$

where  $\mathcal{T}$  denotes the distributional Bellman operator computing target quantiles from observed trajectories. The energy distance loss promotes distributional consistency: the first term aligns predictions with targets while the remaining terms preserve distributional shape by preventing collapse to point estimates. During training, DCE parameters  $\varphi$  are updated jointly with the policy  $\pi_\phi$  using  $N = 50$  quantiles to balance resolution and computational efficiency.

---

### Algorithm 1: CURA-PPO

---

```

1 for iteration = 1, 2, ... do
2   for env = 1, 2, ..., N do
3     Run policy  $\pi_{\phi_{\text{old}}}$  for  $T$  timesteps;
4     Compute advantage estimates  $\hat{A}_{1..T}$ ;
5     Compute risk cost  $\mathcal{C}_{1..T}^{\mathcal{R}}$  w.r.t  $f_{\varphi_{\text{old}}}$ 
6     (Eq. 8);
7     Compute uncertainty cost  $\mathcal{C}_{1..T}^{\mathcal{U}}$  w.r.t  $f_{\varphi_{\text{old}}}$ 
8     (Eq. 9);
9     Compute CURA policy objective  $\mathcal{J}^{\text{CURA}}$  (Eq. 10)
       w.r.t  $\pi_\phi$ , with  $K$  epochs and minibatch size  $M$ ;
10    Update actor, critic and DCE:  $\phi_{\text{old}} \leftarrow \phi$ ,  $\varphi_{\text{old}} \leftarrow \varphi$ ;

```

---

### D. CURA-PPO: Collision Uncertainty-Risk Aware PPO

Building upon the distributional predictions from the DCE, we introduce *CURA-PPO*, augmenting standard PPO with a policy objective that increases the probability of selecting actions that reduce the DCE-estimated risk and uncertainty. This integration transforms a single-objective task-completion formulation into a multi-objective problem that also targets collision avoidance and uncertainty reduction.

1) *Cost Formulation from Distributional Statistics*: We formulate two one-step temporal-difference (TD) cost signals from the DCE's predictions. The *risk cost* is the TD residual of the expected collision possibility, quantifying how much the current action increases expected future collision beyond the immediate collision. The *uncertainty cost* is the change in predictive variance, quantifying whether the next observation becomes more or less uncertain than the current one:

$$\mathcal{C}_t^{\mathcal{R}} = (c_t + \gamma_c \mathcal{R}_\varphi(\mathbf{o}_{t+1}) - \mathcal{R}_\varphi(\mathbf{o}_t)) \quad (8)$$

$$\mathcal{C}_t^{\mathcal{U}} = (\mathcal{U}_\varphi(\mathbf{o}_{t+1}) - \mathcal{U}_\varphi(\mathbf{o}_t)) \quad (9)$$

where  $\mathbf{o}_{t+1}$  is the resulting observation after executing action  $\mathbf{a}_t$  given observation  $\mathbf{o}_t$ . In the risk cost (Eq. 8),  $c_t$  denotes the immediate collision indicator and  $\gamma_c$  the discount factor from Eq. 3. A negative risk cost means the current action reduces the expected collision possibility. A negative uncertainty cost (Eq. 9) indicates reduced predictive uncertainty.

2) *Policy Objective*: We integrate these costs into PPO by augmenting the advantage with penalty terms. Building on Equation 1, the *CURA-PPO* objective becomes:

$$\mathcal{J}^{\text{CURA}}(\phi) = \mathbb{E}_{\pi_{\phi_{\text{old}}}} \left[ \rho(\phi) \cdot \left( \hat{A}_t - \lambda_r \bar{\mathcal{C}}_t^{\mathcal{R}} - \lambda_u \bar{\mathcal{C}}_t^{\mathcal{U}} \right) \right] \quad (10)$$

where  $\bar{\mathcal{C}}_t^{\mathcal{R}}$  and  $\bar{\mathcal{C}}_t^{\mathcal{U}}$  denote normalized costs computed across the batch, weighted by  $\lambda_r$  and  $\lambda_u$  respectively, with  $\lambda_r = 0.25$  and  $\lambda_u = 1.0$ . The probability ratio  $\rho(\phi) = \frac{\pi_\phi(\mathbf{a}_t | \mathbf{o}_t)}{\pi_{\phi_{\text{old}}}(\mathbf{a}_t | \mathbf{o}_t)}$  and its clipping mechanism are identical to standard PPO. The subtracted cost terms act as intrinsic penalties, increasing the probability of actions that reduce the expected collision possibility and predictive uncertainty. Algorithm 1 presents the complete training procedure: rollouts and optimization.

TABLE II: **Quantitative Manipulation Performance.** Success rates ( $\uparrow$ ) on 10,000 test cases.

Group	Method	Uniform			Adversarial		
		0.5 <sup>3</sup>	0.75 <sup>3</sup>	1.0 <sup>3</sup>	0.5 <sup>3</sup>	0.75 <sup>3</sup>	1.0 <sup>3</sup>
G1	Push w/o Occlusion	92.04	88.10	89.49	86.57	79.41	74.19
	Push w/ Occlusion (Baseline)	74.20	66.39	42.34	65.25	48.86	22.90
	<i>CURA-PPO</i> (Ours)	<b>88.73</b>	<b>81.61</b>	<b>80.13</b>	<b>77.12</b>	<b>74.59</b>	<b>72.42</b>
G2	Baseline w/ Conf.	80.63	68.77	63.74	69.65	63.68	54.13
	<i>CURA</i> w/o Uncertainty	84.57	69.35	64.27	70.55	64.08	57.80
	<i>CURA</i> w/o Risk	86.09	79.93	78.45	78.20	70.06	67.82
G3	Push w/ Base	66.33	47.45	40.18	56.97	30.94	17.64
	<i>CURA-PPO</i> w/ Base	86.75	72.20	50.08	76.41	57.31	31.03

- 0.5<sup>3</sup>, 0.75<sup>3</sup>, 1.0<sup>3</sup> denote the size of manipulated cuboid object.
- Each group (G1, G2, and G3) is distinguished for clarify the explanation.

## VI. EXPERIMENTAL RESULTS

We evaluate *CURA-PPO* through extensive experiments assessing its effectiveness in handling object-induced occlusions during non-prehensile manipulation. Our evaluation includes comparisons with two reference approaches to quantify both the severity of occlusion-induced performance degradation and our method’s ability to mitigate this challenge. We conduct ablation studies analyzing individual component contributions, and qualitative analysis of learned behaviors. Finally, we examine the advantage of mobile manipulation over base-only pushing for uncertainty reduction.

### A. Experimental Setup

Experiments are conducted in the simulation environment described in Section V-A, with systematic variations to test performance under various occlusion conditions. We evaluate three object sizes (0.5m, 0.75m, and 1.0m), where larger objects create more severe occlusions by blocking wider sensor regions. As shown in Fig. 3, obstacles (0.4m, 0.6m, and 1.0m cuboids) spawn randomly within a region 3-11m ahead of the initial object position, providing sufficient clearance to avoid immediate collisions. We test two obstacle scenarios: (1) *Uniform*, where obstacles appear randomly across the workspace, and (2) *Adversarial*, where obstacles are positioned along the direct object-to-goal path, requiring active perception and collision avoidance. In both scenarios, obstacle positions are randomly spawned at random times to reflect conditions where unexpected obstacles appear during manipulation. Episodes terminate upon collision or after 50 seconds, with success defined as reaching the goal without collision. All results are averaged over 10,000 episodes for each combination of object size and obstacle scenario.

### B. Evaluation of Manipulation Performance

We compare *CURA-PPO* against two reference policies to assess the severity of occlusion-induced performance degradation and validate our method’s effectiveness:

- *Push w/o Occlusion*: A policy trained in a setting where all occlusions from the manipulated object are artificially removed by filtering LiDAR ray intersections, relying solely on task rewards and temporal sensor accumulation. This serves as an upper bound, representing achievable performance without occlusion constraints.




Fig. 3: **Environment used for training and evaluation.** The initial pose for the object is sampled from the *red* region. Obstacles are randomly generated within the *yellow* region, and the goal pose is sampled from the *green* region. The object and obstacles are cuboids of varying sizes, with 0 to 6 obstacles present in each episode. The top-left shows an obstacle, and the top-right shows the object.

- *Push w/ Occlusion*: A baseline policy trained under realistic occlusion conditions, relying solely on task rewards without confidence maps or explicit uncertainty handling.

Table II (Group G1) clearly demonstrates the impact of occlusion on manipulation performance and validates our method’s effectiveness. *Push w/o Occlusion* consistently achieves high performance across all object sizes and obstacle scenarios, confirming that successful manipulation is feasible when sensing is unobstructed. In contrast, *Push w/ Occlusion* exhibits substantial performance degradation, with success rates dropping significantly as object size increases. This degradation is particularly severe in the *Adversarial* scenario where occlusion severely hinders obstacle detection.

*CURA-PPO* recovers much of the performance lost to occlusion, achieving substantial improvements over the baseline particularly in challenging scenarios. When occlusion occurs, our method actively reduces uncertainty in occluded regions while minimizing collision risk through distributional estimates from the DCE. As a result, *CURA-PPO* nearly matches the performance of *Push w/o Occlusion*, demonstrating that explicit risk and uncertainty modeling effectively mitigates manipulation under partial observability.




Fig. 4: **Qualitative Manipulation Behavior under Occlusion.** Trajectories shown with robot paths in *green* and object paths in *magenta*. (Left) *CURA-PPO* with corresponding confidence maps. The robot actively repositions laterally (a→b, c→d) to reveal occluded regions, successfully detecting and avoiding obstacles during pushing the object to the goal. (Right) The baseline policy, which lacks both confidence maps and DCE-based costs, fails to explore occluded regions, resulting in collision with an undetected obstacle that was hidden behind the manipulated object (e→f). Without these uncertainty-aware mechanisms, the robot maintains a direct path that leads to task failure.

### C. Analysis of Components within *CURA-PPO*

We conduct ablation studies to analyze the individual contributions of three key components—confidence maps, risk cost, and uncertainty cost—with results in Table II (G2).

1) *Evaluation of Confidence Maps:* We evaluate *Baseline w/ Conf.*, a variant that incorporates confidence maps without DCE-based costs.

Confidence maps accumulate past sensor readings while applying temporal decay, enabling it to distinguish recently observed regions from outdated information. This temporal awareness improves performance over the standard baseline (i.e., *Push w/ Occlusion*), particularly for larger objects where occlusion affects wider regions. Results indicate that capturing observation reliability provides guidance for decision-making, though the improvement remains limited without explicit uncertainty-driven behaviors.

2) *Evaluation of Risk and Uncertainty Costs:* We examine the individual effects of risk and uncertainty costs by testing variants that exclude each component while retaining confidence maps. *CURA w/o Uncertainty* improves upon *Baseline w/ Conf.* but shows limited enhancement compared to *CURA-PPO*, lacking guidance for active perception to resolve occlusions. *CURA w/o Risk* yields higher success rates, as uncertainty reduction promotes information-seeking behaviors addressing partial observability challenges.

These results reveal a hierarchy in component importance. Uncertainty cost proves more critical, directly motivating the robot to reduce occlusion through active sensing—a fundamental requirement for safe navigation under partial observability. Risk cost provides complementary collision avoidance signals that enhance safety margins beyond the task rewards alone. While task rewards implicitly penalize collisions through episode termination, risk cost offers continuous gradient signals that guide the robot away from potential collisions before critical situations arise.

Combining both costs (i.e., *CURA-PPO*) achieves the best performance, with uncertainty cost driving active perception and risk reducing expected collision possibility. This complementary effect proves particularly valuable in challenging scenarios with large objects and adversarial obstacles, where both proactive information gathering and collision avoidance are essential for successful manipulation under occlusion.




Fig. 5: **Visualization of DCE predictions** at time steps from Fig. 4. (Left) *CURA-PPO* reduces variance from c[red] to d[blue] through active perception. (Right) Baseline maintains high variance from e[red] to f[blue], leading to collision.

### D. Qualitative Analysis of Manipulation Behavior

We analyze manipulation trajectories and corresponding DCE predictions (Figs. 4 and 5) to examine how our method develops active perception strategies under occlusion, comparing *CURA-PPO* against the baseline policy (i.e., *Push w/ Occlusion*) to reveal distinct manipulation behavior.

Fig. 4 illustrates the contrasting manipulation strategies between *CURA-PPO* and the baseline. In sequences (a→b) and (c→d), *CURA-PPO* responds to low-confidence regions in its perception map by executing lateral movements deviating from the direct path, revealing undetected obstacles while maintaining object control. These uncertainty-reduction maneuvers, though seemingly inefficient for immediate progress, enable obstacle avoidance during manipulation. The baseline policy (e→f), lacking uncertainty signals, maintains direct pushing toward the goal and collides with the unexpected obstacle hidden behind the manipulated object. This behavioral difference highlights how uncertainty-aware decision-making influences manipulation strategy.

Complementing Fig. 4, the DCE predictions in Fig. 5 provide quantitative evidence for uncertainty-aware active perception. At (c), the estimated distribution shows higher variance due to limited information. Our approach reduces variance at (d) via uncertainty-reducing actions, indicating active perception lowers DCE uncertainty in collision possibility. In contrast, the baseline faces a similar situation at (e) but proceeds without DCE guidance. This single-objective strategy keeps variance high through (f), leading to a collision. These observations suggest DCE predictions provide reliable guidance for safe manipulation under occlusion.




Fig. 6: **Manipulation Behavior in Adversarial Scenario.** Sequential snapshots (left to right) demonstrate *CURA-PPO*'s response to adversarial obstacles that suddenly appear in previously clear regions. *Green* and *magenta* paths show robot and object trajectories.

The active perception capabilities of *CURA-PPO* prove particularly important in the *Adversarial* scenario, where obstacles appear along the direct pushing path. Fig. 6 shows how obstacles emerging in previously observed regions create a critical challenge—they become hidden behind the manipulated object during pushing. *CURA-PPO* successfully addresses this challenge through its uncertainty-aware strategy: confidence maps signal when spatial information becomes unreliable, triggering lateral movements that reveal the occluded obstacle. This proactive detection enables the robot to adjust its pushing direction in real-time, guiding the object safely around the obstacle to the goal.

#### E. Role of Kinematic Redundancy in Manipulation

We conducted an additional ablation study using a base-only platform (without a manipulator arm) to evaluate how kinematic redundancy affects manipulation performance under occlusion. Table II (G3) presents two configurations trained under occlusion: *Push w/ Base*, a base-only agent with task rewards, and *CURA-PPO w/ Base*, which applies our uncertainty-aware framework to the same platform.

The results show that base-only manipulation consistently underperforms compared to configurations with a manipulator arm. In the absence of kinematic redundancy, the robot has limited ability to reduce object-induced occlusion while maintaining contact and therefore often advances while the object's forward path in or near unobserved regions. Although *CURA-PPO w/ Base* improves over the base-only baseline, these benefits decline as object size increases and occlusion becomes more severe. As shown in Fig. 7, this example illustrates a failure of the base-only approach. By contrast, leveraging kinematic redundancy enables the mobile manipulator to maintain contact while improving visibility, which leads to successful obstacle detection and avoidance.

For a more intuitive understanding of the experimental results, please refer to the supplementary video.

#### ACKNOWLEDGMENTS

This work was supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (RS-2025-25443318 and RS-2023-00237965).

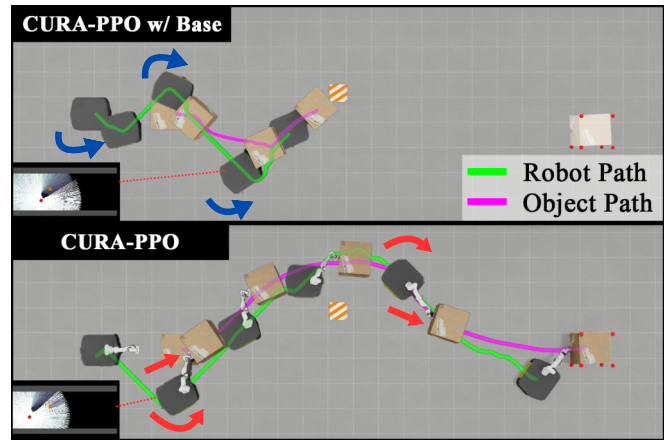


Fig. 7: **Behavioral Comparison of Base-Only and Mobile Manipulator Systems.** (Top) Base-only platform fails due to restricted viewing angle, resulting in insufficient visual coverage while pushing the object to the target. (Bottom) Mobile manipulator successfully leverages kinematic redundancy to maintain object contact while adjusting viewing angles for effective perception.

## VII. CONCLUSION

We presented *CURA-PPO*, a novel reinforcement learning framework that addresses the challenge of object-induced occlusion in non-prehensile manipulation using only onboard sensing. By modeling collision possibility as a distribution through the *Distributional Collision Estimator (DCE)*, our approach extracts both *risk* and *uncertainty* signals to augment standard policy optimization. These complementary costs encourage distinct yet synergistic behaviors: *risk* drives collision avoidance while *uncertainty* promotes active perception to resolve occlusions during manipulation. Experiments across varying object sizes and obstacle configurations demonstrate that *CURA-PPO* achieves robust performance in both *Uniform* and *Adversarial* scenarios. Qualitative analysis reveals that the learned policy strategically maneuvers to gather environmental information while maintaining object control, mitigating the occlusion challenge. Future work will focus on deploying this approach on real robot systems and extending the framework to incorporate additional sensor modalities such as RGB-D cameras for enhanced environmental perception in complex real-world scenarios.

## REFERENCES

- [1] Fabio Ruggiero, Vincenzo Lippiello, and Bruno Siciliano, "Non-prehensile dynamic manipulation: A survey", *IEEE Robotics and Automation Letters*, vol. 3, no. 3, pp. 1711–1718, 2018.
- [2] Nils Dengler, David Großklaus, and Maren Bennewitz, "Learning goal-oriented non-prehensile pushing in cluttered scenes", in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 1116–1122.
- [3] Nils Dengler, Juan Del Aguila Ferrandis, João Moura, Sethu Vijayakumar, and Maren Bennewitz, "Learning goal-directed object pushing in cluttered scenes with location-based attention", *arXiv preprint arXiv:2403.17667*, 2024.
- [4] Juan Del Aguila Ferrandis, João Moura, and Sethu Vijayakumar, "Nonprehensile planar manipulation through reinforcement learning with multimodal categorical exploration", in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2023, pp. 5606–5613.

- [5] Seunghun Jeon, Moonkyu Jung, Suyoung Choi, Beomjoon Kim, and Jemin Hwangbo, “Learning whole-body manipulation for quadrupedal robot”, *IEEE Robotics and Automation Letters*, vol. 9, no. 1, pp. 699–706, 2023.
- [6] Yuming Feng, Chuye Hong, Yaru Niu, Shiqi Liu, Yuxiang Yang, Wenhao Yu, Tingnan Zhang, Jie Tan, and Ding Zhao, “Learning multi-agent loco-manipulation for long-horizon quadrupedal pushing”, *arXiv preprint arXiv:2411.07104*, 2024.
- [7] Yujie Tang, Hai Zhu, Susan Potters, Martijn Wisse, and Wei Pan, “Unwieldy object delivery with nonholonomic mobile base: A stable pushing approach”, *IEEE Robotics and Automation Letters*, vol. 8, no. 11, pp. 7727–7734, 2023.
- [8] Idil Ozdamar, Doganay Sirintuna, Robin Arbaud, and Arash Ajoudani, “Pushing in the dark: A reactive pushing strategy for mobile robots using tactile feedback”, *IEEE Robotics and Automation Letters*, 2024.
- [9] Ioannis Dadiotis, Mayank Mittal, Nikos Tsagarakis, and Marco Hutter, “Dynamic object goal pushing with mobile manipulators through model-free constrained reinforcement learning”, *arXiv preprint arXiv:2502.01546*, 2025.
- [10] Xuran Pan, Zihang Lai, Shiji Song, and Gao Huang, “Activenerf: Learning where to see with uncertainty estimation”, in *European Conference on Computer Vision*. Springer, 2022, pp. 230–246.
- [11] Georgios Georgakis, Bernadette Bucher, Anton Arapin, Karl Schmeckpeper, Nikolai Matni, and Kostas Daniilidis, “Uncertainty-driven planner for exploration and navigation”, in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 11295–11302.
- [12] Liren Jin, Xieyuanli Chen, Julius Rückin, and Marija Popović, “Neubv: Next best view planning using uncertainty estimation in image-based neural rendering”, in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2023, pp. 11305–11312.
- [13] Nils Dengler, Jesper MÅžcke, Rohit Menon, and Maren Bennewitz, “Efficient manipulation-enhanced semantic mapping with uncertainty-informed action selection”, *arXiv preprint arXiv:2506.02286*, 2025.
- [14] Shengmiao Jin, Yuchen Mo, and Wenzhen Yuan, “Learning to double guess: An active perception approach for estimating the center of mass of arbitrary objects”, *arXiv preprint arXiv:2502.02663*, 2025.
- [15] Mykel J Kochenderfer, *Decision making under uncertainty: theory and application*, MIT press, 2015.
- [16] Marc G Bellemare, Will Dabney, and Rémi Munos, “A distributional perspective on reinforcement learning”, in *International conference on machine learning*. PMLR, 2017, pp. 449–458.
- [17] Will Dabney, Mark Rowland, Marc Bellemare, and Rémi Munos, “Distributional reinforcement learning with quantile regression”, in *Proceedings of the AAAI conference on artificial intelligence*, 2018, vol. 32.
- [18] Alexander Amini, Wilko Schwarting, Ava Soleimany, and Daniela Rus, “Deep evidential regression”, *Advances in neural information processing systems*, vol. 33, pp. 14927–14937, 2020.
- [19] Yarín Gal and Zoubin Ghahramani, “Dropout as a bayesian approximation: Representing model uncertainty in deep learning”, in *international conference on machine learning*. PMLR, 2016, pp. 1050–1059.
- [20] Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell, “Simple and scalable predictive uncertainty estimation using deep ensembles”, *Advances in neural information processing systems*, vol. 30, 2017.
- [21] Gaon An, Seungyong Moon, Jang-Hyun Kim, and Hyun Oh Song, “Uncertainty-based offline reinforcement learning with diversified q-ensemble”, *Advances in neural information processing systems*, vol. 34, pp. 7436–7447, 2021.
- [22] Benjamin Rosman, Majd Hawasly, and Subramanian Ramamoorthy, “Bayesian policy reuse”, *Machine Learning*, vol. 104, no. 1, pp. 99–127, 2016.
- [23] Felipe Leno Da Silva, Pablo Hernandez-Leal, Bilal Kartal, and Matthew E Taylor, “Uncertainty-aware action advising for deep reinforcement learning agents”, in *Proceedings of the AAAI conference on artificial intelligence*, 2020, vol. 34, pp. 5792–5799.
- [24] Ralf Roemer, Armin Lederer, Samuel Tesfazgi, and Sandra Hirche, “Vision-based uncertainty-aware motion planning based on probabilistic semantic segmentation”, *IEEE Robotics and Automation Letters*, vol. 8, no. 11, pp. 7825–7832, 2023.
- [25] Shivin Dass, Jiaheng Hu, Ben Abbatematteo, Peter Stone, and Roberto Martín-Martín, “Learning to look: Seeking information for decision making via policy factorization”, *arXiv preprint arXiv:2410.18964*, 2024.
- [26] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov, “Proximal policy optimization algorithms”, *arXiv preprint arXiv:1707.06347*, 2017.
- [27] John Schulman, Philipp Moritz, Sergey Levine, Michael Jordan, and Pieter Abbeel, “High-dimensional continuous control using generalized advantage estimation”, *arXiv preprint arXiv:1506.02438*, 2015.
- [28] Hyeonjeol Ryu, Minsung Yoon, Daehyung Park, and Sung-Eui Yoon, “Confidence-based robot navigation under sensor occlusion with deep reinforcement learning”, in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 8231–8237.
- [29] M. Mittal et al., “Orbit: A unified simulation framework for interactive robot learning environments”, *IEEE Robotics and Automation Letters (RA-L)*, 2023.
- [30] Roger Koenker, *Quantile regression*, vol. 38, Cambridge university press, 2005.
- [31] Lukas Schneider, Jonas Frey, Takahiro Miki, and Marco Hutter, “Learning risk-aware quadrupedal locomotion using distributional reinforcement learning”, in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 11451–11458.