

Learning Fast, Tool-aware Collision Avoidance for Collaborative Robots

Joonho Lee*, Yunho Kim, Seokjoon Kim, Quan Nguyen, Youngjin Heo

Abstract— Ensuring safe and efficient operation of collaborative robots in human environments is challenging, especially in dynamic settings where both obstacle motion and tasks change over time. Current robot controllers typically assume full visibility and fixed tools, which can lead to collisions or overly conservative behavior. In our work, we introduce a tool-aware collision avoidance system that adjusts in real time to different tool sizes and modes of tool-environment interaction. Using a learned perception model, our system filters out robot and tool components from the point cloud, reasons about occluded area, and predicts collision under partial observability. We then use a control policy trained via constrained reinforcement learning to produce smooth avoidance maneuvers in under 10 milliseconds. In simulated and real-world tests, our approach outperforms traditional approaches (APF, MPPI) in dynamic environments, while maintaining sub-millimeter accuracy. Moreover, our system operates with approximately 60 % lower computational cost compared to a state-of-the-art GPU-based planner. Our approach provides modular, efficient, and effective collision avoidance for robots operating in dynamic environments. We integrate our method into a collaborative robot application and demonstrate its practical use for safe and responsive operation.

I. INTRODUCTION

THE use of robotic manipulators in contact-rich, real-world applications is rapidly expanding, driven by recent advances in control strategies and hardware [1], [2]. These developments have broadened robot deployment into a wide range of settings, from industrial sites to in-home service applications. As robots become more integrated into everyday environments, ensuring safety is becoming critical—especially for collaborative and service robots that operate alongside humans and cannot be isolated behind protective barriers like industrial machines.

Most industrial systems rely on collision detection [3], [4] and compliance control [5], [6], responding only after a collision has occurred. While widely adopted [7], these approaches fall short in dynamic environments populated by humans. A common preventive measure is to define protective zones around the robot during the pre-impact phase [8]. Sensors like LiDAR or safety-rated scanners create “red” zones to trigger safety stops and “yellow” zones to reduce speed. However, these methods do not change trajectories

Manuscript received: January, 21, 2025; Revised March, 21, 2025; Accepted May, 20, 2025.

This paper was recommended for publication by Editor Olivier Stasse upon evaluation of the Associate Editor and Reviewers’ comments.

This research has been funded by the Industrial Technology Innovation Program (P0028404) of the Ministry of Industry, Trade and Energy of Korea.

All authors are with Neuromeka Co., Ltd.

*Senior and Corresponding Author (joonho.lee@neuromeka.com)

Digital Object Identifier (DOI): see top of this page.

©2026 IEEE




Fig. 1. **Pick-and-place example.** Our system manages robot-environment interactions by taking an *interaction mode* input that determines contact permissions. (Left) The tool area is allowed to make contact with the environment during active use. (Right) At other times, contact is prohibited for both the tool and the robot.

according to obstacles and are limited because robots with high kinetic energy require large safety zones. Recent works have introduced more advanced perception systems capable of adapting trajectories in real time by measuring distances and maintaining safe separation [9], [10], [11], [12], [13], [14].

Existing approaches still face limitations in real-world applications. They often rely on fixed geometric parameters and cannot adapt to dynamic environments. They also overlook *controlled contact*, reducing effectiveness in applications like assembly and pick-and-place. Furthermore, their lack of robustness to occlusions and reliance on heuristics often increase system complexity.

In this work, we propose an adaptive collision avoidance system that provides two command modes for manipulator-environment interaction (see Fig. 1). *Engage mode* allows the tool (e.g., gripper) to make contact during active tasks, while *Protective mode* prevents all contact during transport phases. This tool-aware design enhances adaptability and enables safe and efficient operation in dynamic contact-rich settings.

A. System Requirements

We define three key requirements for an effective collision avoidance system in real-world applications:

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

- **Flexibility to Tool Settings:** Robots use a wide range of End-effector (EE)s, so the system should adapt to varying tool shapes and operational modes in real time.
- **Precision:** The system should achieve high precision in nominal operations, ensuring consistent performance.
- **Responsiveness and Reliability:** The system must operate with minimal delays to avoid unsafe conditions, remain interpretable, and remain stable under challenging conditions such as occlusions.

B. System Overview

Considering the requirements above, we developed a collision avoidance system that combines the robustness of learning-based methods with the precision of a classical controller. An overview of our system is shown in Fig. 2. The user provides three inputs: the desired EE pose, the geometry of the tool attached to the robot and the interaction mode (Engage or Protective mode). In nominal conditions, the system solves Inverse Kinematics (IK) using a classical method. A learned safety critic continuously monitors the environment and evaluates the risk of future collisions. When the risk exceeds a defined threshold, the system switches to a Reinforcement Learning (RL) policy for reactive motion generation.

The first part of our system is an encoder-decoder model that processes point cloud data along with its previous occupancy prediction. The module has two outputs: occupancy grid and safety critic values [15] for both body and tool collision constraints (explained later).

The safety critic value estimates the likelihood of future constraint violations (collisions), based on Srinivasan et al. [15]. It outputs a scalar value between 0 and 1. When the safety critic value is below a threshold (0.8 in our implementation), the RL policy output is used as an initial guess for an iterative IK method for high-precision. Conversely, when the critic indicates a high risk, the policy output is directly used for reactive collision avoidance.

The policy uses the latent space of the encoder as its input. This state encapsulates the spatial information necessary for effective collision avoidance [16]. The collision avoidance policy operates at a frequency of 50 Hz and generates joint position target residuals.

C. Contributions

Our key advancements are:

- **Fast and Adaptive Collision Avoidance:** Our system adapts in real time to user inputs and robustly handles dynamic obstacles. Our learned models adapt their outputs based on tool geometry and modes. Its low computational cost enables fast, safe execution of contact-rich tasks.
- **Seamless Integration with High Accuracy:** To balance responsiveness and precision, a fast but less accurate RL policy is activated only under high collision risk.
- **Robust Perception:** Our learned perception model removes the need for explicit robot filtering [11], mapping [17], and hand-crafted occlusion heuristics. This

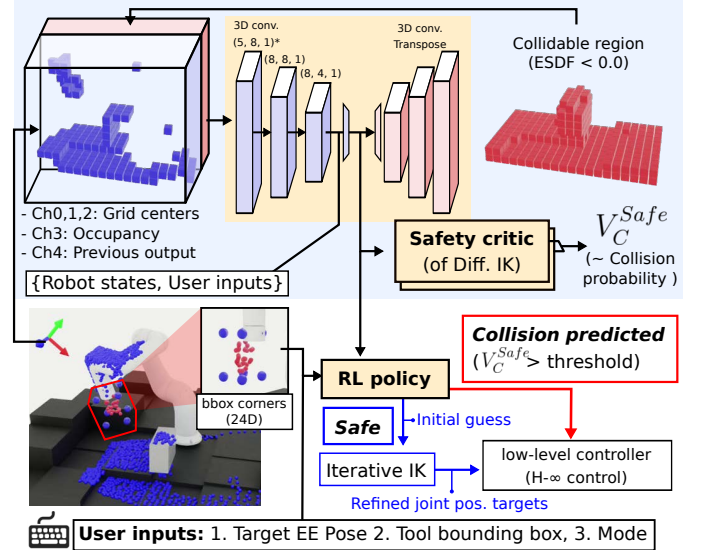


Fig. 2. **System Overview.** The RL policy generates target joint positions, used as initial guess for IK when safe. A 3D CNN encoder processes workspace occupancy and proprioceptive inputs to estimate collision probabilities via a learned safety critic. When the risk is high, RL policy is triggered. All components are trained in simulation.

results in a simpler implementation and enables reliable operation in dynamic, partially observable environments.

We validate our system through real-world experiments and qualitative evaluations. Our approach achieves high precision and reliability while operating an order of magnitude faster than state-of-the-art methods [13], [18], [19].

II. RELATED WORK

We summarize related works in the areas relevant to our research.

A. Learned Scene Representation

Our approach aligns with recent works that employ neural networks to process raw sensor data for scene representation [16], [20], [21], [22], [23]. These methods utilize encoder-decoder architectures to process exteroceptive measurements and use the resulting latent vectors as inputs for policy networks. Hoeller et al.[16] used a 3D Convolutional Neural Network (CNN) to handle noisy point clouds, producing a latent representation for navigation and terrain reconstruction. Similarly, Kulkarni et al.[20] developed a Deep Collision Encoder (DCE) to extract collision information from noisy depth images for quadrotors using an architecture inspired by variational autoencoders.

These works demonstrate that processing sensor measurements with neural networks to build implicit environmental maps, rather than relying on conventional mapping, leads to highly responsive and robust systems.

B. Collision Avoidance

Traditional approaches to safe motion planning focus on finding collision-free geometric paths using methods such as

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.




Fig. 3. Simulation Environment.

Rapidly-exploring Random Tree (RRT) [24] and Artificial Potential Field (APF) [25]. Subsequent research has explored Trajectory Optimization (TO) techniques, like CHOMP [26].

Recent research frames robot navigation as a global optimization problem. To handle complex environments with multiple possible solutions, sampling-based approaches are employed. For example, Curobo [13] combines sampling-based planning and TO for robotic manipulators.

These model-based methods generally perform well in fully observable environments where obstacles and dynamics are predictable. However, they struggle in dynamic, partially observable environments. Occlusions and the need for fast replanning make reliable collision avoidance challenging in real-world settings.

Recently, data-driven approaches, particularly those using RL, have emerged as alternatives. These methods allow robots to learn safe maneuvers by quickly responding to dynamic environment [27], [28], [20], [29].

C. Safe Reinforcement Learning

Safe RL is a well-established research area, with various approaches differing in their architectural frameworks [30]. One widely used approach is the Conditional Value at Risk (CVaR) formulation [31], [32], [33]. This method focuses on maximizing the expected return in the lower tail of the return distribution, making it effective for addressing rare but critical events. However, we opted not to adopt this approach due to the complexity involved in implementing distributional critics and tuning risk parameters [34], leaving it as a consideration for future work.

In this work, we adopt the Constrained Markov Decision Process (CMDP) framework for its simplicity in implementation and deployment, leveraging its similarity to the widely used PPO algorithm. The CMDP framework [35] offers a structured and efficient approach to incorporating constraints into RL. By defining constraints as cost functions, CMDPs allow the optimization of policies that balance safety and performance [36], [37]. These methods efficiently manage multiple constraints. Kim et al.[38] and Lee et al.[39] demonstrated the utility of CMDPs in training safe locomotion policies for quadrupedal robots on rough terrain.

III. METHOD

We provide an overview of our system and describe each component in detail.

A. Problem Definition

Figure 3A illustrates the target environment. We focus on 2.5D environments, where obstacles are primarily oriented in the vertical (z) direction with varying heights. We consider random static and dynamic obstacles, replicating typical scenarios where manipulators move around machinery, workpieces, and human workers.

Scenarios involving deep cavity insertion or complex, occluded shapes are excluded, as they exceed the sensing capabilities of fixed cameras and are beyond the scope of this work.

We consider various tool shapes attached to the EE by defining a tool region and simulating point cloud within the region as illustrated in Fig. 3B.

B. Safety Critic

We follow the approach by Srinivasan et al. [15], which defines the discounted probability of future constraint violations. The optimal safety critic for the i -th constraint estimates the following expectation of the nominal controller:

$$V_{C_i}^{safe}(s_t) = c_i(s_t) + (1 - c_i(s_t)) \sum_{t' > t}^T \mathbb{E} \left[\gamma^{t'-t} c_i(s_{t'}) \right]. \quad (1)$$

As specified in Table. I, the cost function c_i is an indicator function in our setup. The cumulative discounted probability of violations is estimated using Bellman equation:

$$V_{C_i}^{safe}(s) = c_i(s) + (1 - c_i(s)) \mathbb{E} \left[\gamma V_{C_i}^{safe}(s') \right] \in [0, 1]. \quad (2)$$

We train it via value iteration during the encoder-decoder model training.

During the deployment, we aggregate safety critics to evaluate the overall collision risk. The maximum safety value between the robot body and the tool area is used as

$$V_C^{safe}(s_t) = \max(V_{C_{body}}^{safe}(s_t), V_{C_{tool}}^{safe}(s_t)) \quad (3)$$

where $V_{C_{body}}^{safe}$ represents the value of the robot's body, and $V_{C_{tool}}^{safe}$ represents that of the tool constraint (see III-C3).

C. Constrained Policy Optimization

1) *Background:* A control problem is typically modeled as a Markov Decision Process (MDP) in RL framework, which is defined by a tuple (S, A, r, p, μ) . S is the set of states, A is the set of actions, $r : S \times A \times S \rightarrow \mathbb{R}$ is the reward function, $p : S \times A \times S \rightarrow [0, 1]$ is the state transition probability and μ is the initial state distribution. The aim of RL is to find a policy $\pi : S \mapsto \mathcal{P}(A)$ that maximizes

$$J_R(\pi) = \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t, s_{t+1}) \right], \quad (4)$$

where $\gamma \in [0, 1)$ is the discount factor. Here, $\mathbb{E}[\dots]$ represents the average over a finite batch of samples. s_0 is sampled from an initial state distribution μ and trajectories are sampled using π .

This is extended into CMDP to handle physical constraints, augmented with a set C of cost functions that capture

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

Rewards	
1. Position tracking	$1.0 - \text{clip}(p_{EE}^{target} - p_{EE} ^2, 0.0, 1.0)$
2. Orientation tracking	$1.0 - 2.0 * \arccos((q_{EE} \cdot (q_{EE}^{target})^{-1})_w)$
3. Command smoothness 1	$-1.0 \times 10^{-3} \ q_t^{des} - q_{t-1}^{des}\ ^2$
4. Command smoothness 2	$-1.0 \times 10^{-4} \ q_t^{des} - 2q_{t-1}^{des} + q_{t-2}^{des}\ ^2$
Costs	
1. Collision	1.0 if the robot body collides
2. Tool region violation (Protective mode only)	1.0 if any ESDF (corner point _i) < 0.025
3. End-effector speed limit	1.0 if EE speed exceeds 0.5 m/s
4. Joint speed limit	1.0 if $ \dot{q}_{i,t} > 0.8 \cdot \dot{q}_{i,lim}$.

TABLE I
REWARD AND COST FUNCTIONS.

constraint violations $\{c_1, \dots, c_n\}$ and corresponding limits $E = \{\epsilon_1, \epsilon_2, \dots, \epsilon_n\}$ [40]. Each $c_i : S \times A \times S \rightarrow \mathbb{R}$ maps (s_t, a_t, s_{t+1}) to the cost of the transition. An optimal policy maximizes J_R , while keeping the discounted sum of future costs c_i below their respective threshold ϵ_i , yielding the constrained optimization problem:

$$\max_{\pi} J_R(\pi) \quad \text{s.t.} \quad \forall i \in \{1, \dots, n\}, J_{C_i}(\pi) \leq \epsilon_i, \quad (5)$$

where $J_{C_i}(\pi) = \mathbb{E}[\sum_{t=0}^{\infty} \gamma^t c_i(s_t, a_t, s_{t+1})]$.

This framework has been shown to be effective for high-dimensional control problems in robotics in recent works [39], [38]. Following the previous works, we employ the Penalized Proximal Policy Optimization (P3O) algorithm.

2) *CMDP for collision avoidance*: Here we define the action and observation spaces, as well as the reward and cost functions for our CMDP. The state transition probabilities follow the rigid body dynamics of the robot.

The action space is defined by the residual changes to the robot's joint positions: $q_t^{des} = q_t + a_t$ where a_t represents the action at time t and q_t represents current joint angles. Action distributions are modeled using a Gaussian distribution. The observation space includes the previous joint position history, the latent state from the voxel encoder, and the target EE pose, represented in Cartesian coordinates and euler angles in the base frame.

The reward and cost functions are outlined in Table I. The reward function encourages EE pose tracking while also regularizing the smoothness of the robot's movements. The cost functions are defined as indicator functions.

3) *Variable Tool Region Constraint*: The tool region, defined as a bounding box attached to the EE (Fig. 2A-ii), is randomized in each episode. We evaluate the Euclidean Signed Distance Function (ESDF) values at the eight corner points and flag a constraint violation if any value is less than half the voxel width.

D. Model Implementation

1) *Architecture*: The architecture is shown in Fig. 3. A voxelized point cloud is processed by three 3D CNN layers and concatenated with proprioceptive inputs—including joint state history and user commands—before passing through a Multi Layer Perceptron (MLP) to produce a latent vector. Joint states help filter out the robot's own body and assess collision risk.

2) *Scenario Setup*: All training data was generated in simulation using approximate robot and workspace geometry (Fig. 3C), without relying on precise CAD models.

To define a new setup, users have to specify the workspace dimensions (x, y, z), grid resolution, sampling ranges for target EE poses, and a range of tool region parameters, including size and offsets. Then the obstacles and tool bounding boxes are randomized each episode.

3) *Training*: We first collect data for training the encoder–decoder and safety critic by running a nominal controller (differential IK) in simulation. Occlusion labels are assigned to grid cells where the ESDF value is less than half the grid resolution. These components are trained using supervised learning and value iteration.

The policy is then optimized using P3O [36], with the encoder–decoder and safety critic frozen to enable stable and efficient learning based on fixed perception and safety estimates in simulation.

IV. EXPERIMENTAL RESULTS

We present real-world experiments and qualitative evaluations to validate each component of our system. See Supplementary video for our main result.

A. Experimental Setup

The collision avoidance behavior and task-space position tracking performance of our system are evaluated using the setup illustrated in Fig. 4A. We used the Indy7 robot, which is a collaborative robotic arm with six degrees of freedom and is equipped with a two-finger gripper. The depth camera, an Intel RealSense D435, was mounted on the side of the table, as depicted in the figure. The wrist camera in the picture is not used for collision avoidance.

To assess the system's robustness in the presence of occlusions, we use a single point cloud source with a limited field of view (Realsense D435). This allows us to evaluate the system under realistic sensing constraints.

The robot operates within a workspace measuring 0.6 m (x -axis) by 1.0 m (y -axis) by 0.6 m (z -axis). Within this workspace, we construct a voxel grid with 0.05 m resolution.

B. Use-case Example: Pick-and-Place


Fig. 1 illustrates a pick-and-place system where our collision avoidance system is integrated. The supplementary video demonstrates this behavior in action. We use a finite state machine to switch between picking and placing operations. For grasp planning, we used ANY-grasp [41]. The Engage mode is activated when a grasp pose is detected by ANY-grasp. Protective mode is used otherwise. This setup demonstrates how our approach can be incorporated into existing pipelines.

C. Computation Time

Our system operates at 50 Hz. The average update time for each action is approximately 6 ms, enabling the system to scale for deployment at frequencies exceeding 100 Hz. All tests were performed on a desktop PC equipped with an RTX

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

A. Experimental setup




B. Engage mode




C. Protective mode



D. Tool center trajectory



E. Safety critic values



— Engage mode — Protective mode

Fig. 4. **Tool-Adaptive Collision Avoidance.** (A) Experimental setup. (B) Motion in Engage mode, where the tool area is permitted to make contact with the environment. (C) Motion in Protective mode, where the policy and the safety critic treat the tool region as part of the robot’s body and avoid contact. (D) Tool center trajectories for each mode. (E) Corresponding safety critic values.

3060 GPU (12 GB), an Intel i7-12700F CPU, and 32 GB of RAM.

Model inference averages 1.3 ms (SD 0.5 ms, max 4.0 ms), performing a single-pass evaluation of the encoder-decoder, policy, and safety critic. Other processes, including solving kinematics, communication, and observation computation, take 3.9 ms (SD 3.9 ms, max 19.0 ms). Point cloud processing—implemented with Warp [42]—runs in parallel and averages 7.9 ms (SD: 3.7 ms, max: 72.3 ms) for voxelization, occupancy computation, and raycasting.

D. Analyzing Adaptive Behavior

A static obstacle was placed at the center of the workspace to analyze obstacle avoidance behavior, following a setup similar to [43]. The robot was tasked with moving its EE from the left to an endpoint on the right.

1) *Interaction Mode Behaviors:* Fig. 4B–D illustrate the trajectories executed by the robot in each mode. In Protective mode, the robot accounts for the tool region, resulting in longer detours to avoid potential collisions.

A. raw point cloud



B. estimated collidable region



Fig. 5. **Visualization of Point Cloud and Collidable Region Estimation.** Front view snapshots at the initial configuration, at the start (i) and at a mid (ii) point of the trajectory. (A) Raw point cloud data captured during the experiment. (B) Corresponding estimated collidable regions.

2) *Safety Critic Values:* Fig. 4E shows the cost values by the safety critic, calculated as the maximum of the body cost and the tool cost value. The controller switches to the RL policy when the cost value exceeds a predefined threshold (0.8) and reverts to the nominal controller once the cost decreases.

3) *Estimated Collidable Region:* Fig. 5 shows the raw point cloud and the estimated occupancy during the experiment in Fig. 4. When the gripper is at the left side, the table is significantly occluded (Fig. 5A,B-i). Our model remembers the measurements at the initial configuration and preserves the occupancy prediction.

Our model filters out measurements corresponding to the robot’s body and tool components, such as the gripper, camera, and cables. This capability is achieved by incorporating randomly simulated tool geometry during training.

4) *Adaptation to Tool Shape:* Fig. 6A and Fig. 6B depict the task-space position tracking performance in the yz plane for three different EE configurations. Each target position was maintained for three seconds after convergence, and the average tracking error was computed. The experiment follows the setup shown in Fig. 4A, with errors measured via forward kinematics.


As collision avoidance is activated, higher tracking errors are observed near the table and obstacle. In contrast, open areas consistently exhibit errors around 0.01 mm. As larger tool areas were commanded, the high-error region expands (Fig. 6B). This demonstrates the system’s ability to adaptively predict potential collisions based on tool dimensions.

E. Validation of the Control Strategy


Fig. 6C illustrates the performance of RL alone, where tracking errors in the safe area are significantly higher compared to ours (Fig. 6A). This highlights the inherent limitations of RL in achieving the precision offered by traditional IK methods.

We also compared our system to Curobo [13], a state-of-the-art GPU-accelerated TO framework (Fig. 6D). The results are similar: Curobo also reports no feasible solution (marked as N/A) in obstacle regions. In collision-free regions, our method achieves lower tracking error because of the refinement using iterative IK.


A. Engage mode



B. Protective mode



C. RL Only. No tool



D. Trajectory Optimization (Curobo). No tool




Fig. 6. **Tracking Errors Across Different Tool Conditions and Control Approaches.** Task-space position tracking error in the yz plane, with target positions in a 0.1 m grid at a fixed x -coordinate of 0.5 m. (A) Engage mode: only robot body collisions are considered. (B) Protective mode: results for different tool sizes. (C) RL-only approach. (D) Trajectory optimization.

Our system requires fewer computational resources overall. The average planning time for Curobo is 468 ms (SD 244 ms, max 1.36 s), whereas our method operates within 10 ms. In terms of memory usage, ours requires only 0.45 GB of GPU memory, compared to approximately 4.5 GB for Curobo when running with NvBlox. It is important to note that this does not include voxel grid mapping and other system components.

While global optimizers [13], [44] excel at long-horizon planning under full observability, our lightweight system is better suited for reactive, short-horizon collision avoidance—especially when computational resources are shared with high-level processes.

F. Evaluation in Dynamic Settings

To assess the reliability of our system, we evaluate its performance in dynamic environments with moving obstacles




Fig. 7. **Dynamic obstacle avoidance experiment** (A) Experimental setup. The obstacle and waypoint move in opposite direction. The waypoint moves at 0.25 m/s. (B) Visualization of point cloud on the obstacle

TABLE II
COLLISION RATE (COLL.) AND SUCCESS RATE (SUCC.) FOR DIFFERENT METHODS UNDER STATIC AND DYNAMIC OBSTACLE CONDITIONS. ALL VALUES IN %

Method	Control	Static		0.2 m/s		0.4 m/s	
		Succ.	Coll.	Coll.	Coll.		
Perception	GT Pose	100	0	2	18		
	MPPI	100	0	6	56		
Convex hull [11], [45]	APF	54	16	78	82		
	MPPI	72	0	50	72		
Learned SDF [18]	APF	96	4	8	38		
	MPPI	100	0	16	65		
Ours (Occupancy grid)	RL	100	0	6	36		

and compare it against widely used baseline methods.

1) *Experimental Setup:* As shown in Fig. 7A, the EE is commanded to move from right to left at a speed of 0.25 m/s. For the dynamic case, an obstacle is spawned on the left side of the workspace and moves at a fixed linear velocity toward the center of the EE target trajectory. In the static case, the obstacle is fixed at the center. Each method is evaluated over 50 trials.

All controllers operate at 50 Hz; to ensure consistent timing, the simulation pauses each step until computation is complete.

2) *Baselines:* We validated our control strategy against two baseline approaches with different perception methods:

- **APF:** A classical reactive path planning method [25].
- **Model Predictive Path Integral (MPPI):** A sampling-based Model Predictive Control (MPC) method [13], [18], [46]. We followed STORM [19] implementation with a 0.5 s horizon and 4000 samples.

Traditional approaches often preprocess point clouds by clustering and approximating obstacles as ellipsoids or convex hulls [47], [11]. More recent work has adopted learned ESDF or collision check for improved robustness and reactivity [18], [48], which aligns with our learning-based approach.

For evaluation, we implemented two perception variants for the baselines: (i) point cloud obstacle approximation using convex hulls [11], [45] combined with Kalman filter-based centroid tracking, and (ii) a learned SDF model [18].

3) *Result:* Table II summarizes the performance of our method compared to baseline approaches under varying obstacle speeds. **GT** denotes access to ground-truth obstacle position and shape.

Our method achieves performance comparable to APF with learned SDF [18]. Across all scenarios, approaches using

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

TABLE III
COMPARISON ACROSS DIFFERENT SCENE REPRESENTATIONS.
ALL VALUES IN %

Method	Protective Mode			Engage Mode	
	Succ.	Coll.	Tool Coll.	Succ.	Coll.
End-to-End	10.0	2.51	5.03	11.0	6.48
Autoencoder (AE)	47.7	0.86	1.27	59.3	1.68
Mapping + AE	54.1	1.06	2.26	72.8	2.44
Ours	67.3	0.42	1.31	82.5	0.82

learned perception demonstrate higher robustness to dynamic changes and partial observability.

APF occasionally fails in static settings due to local minima and may exhibit oscillatory behavior caused by sudden, high repulsive forces. MPPI shows strong performance in static scenarios but struggles with dynamic obstacles. Despite running at high frequency without delay, its trajectory samples requires time to converge to collision-free trajectories, which limits responsiveness.

The convex hull approximation approach suffers significantly from partial occlusions, as illustrated in Fig. 7B. In such cases, the obstacle is underestimated in size, leading to increased collision rates.

4) *Scalability*: Traditional methods involving point cloud clustering and filtering [47], [11] have a per-cycle runtime that grows with the number of segmented obstacles. By contrast, our method and learned approaches using the whole scene [18], [48] maintain a constant runtime.

G. Effectiveness of the Learned Perception Model for RL

We evaluated our perception approach, which uses an encoder-decoder model to reconstruct collision regions, filter out noise, and remove robot and tool areas from the voxel grid.

1) *Baselines*: We compared three methods:

- **End-to-End Training**: Training directly from scratch using the same RL algorithm without pretraining.
- **Pretrained Autoencoder (AE)**: A 3D CNN is pretrained to reconstruct the input voxel grid, and its latent representation is used as input to the policy.
- **Mapping with Autoencoder (Mapping+AE)**: This approach explicitly filters out robot and tool points from the raw point cloud, constructs a voxel-based occupancy grid, and uses it as input to the policy network. We integrated Voxblox [17] into the simulation, allowing the policy to use the voxel map as external memory.

2) *Results*: The results, summarized in Table III, are evaluated in both control modes in simulation. In Protective mode, tool shape is randomized. The tool violation rate reflects the percentage of episodes where tool constraints were violated. Metrics are computed from 20,000 episodes, with waypoints sampled from free space ($ESDF > 0.1$) in the workspace. Each episode runs for 10 seconds.

End-to-end training with P3O yields suboptimal performance, resulting in overly conservative behavior that reduces success rates and increases collision frequency.

While the pretrained encoder (AE) improves over end-to-end training, it still leads to lower success rates and higher collision rates compared to our method.

Although the Mapping+AE approach achieves a high success rate, it struggles with perception uncertainties such as noise and occlusions. A key failure mode we observed is that robot and tool filtering occasionally removes nearby obstacle points when they are too close to the robot, resulting in unexpected collisions.

V. CONCLUSION

We presented a fast and tool-aware collision avoidance system for collaborative robots. A use case example, shown in Fig. 1, highlights its applicability to realistic scenarios involving collaborative robots under practical sensing constraints. By integrating constrained RL with a traditional differential IK approach, our system achieves both robustness and high accuracy. It seamlessly adapts to diverse tool shapes and interaction modes, ensuring safe and reliable operation while supporting real-time performance exceeding 100 Hz

Compared to baseline methods such as APF and MPPI, our system consistently achieves lower collision rates and comparable or higher success rates in dynamic environments. Experimental results further show robust performance with control latency under 10 ms and GPU memory usage of only 0.45 GB—over 10 more memory-efficient and significantly faster than Curobo’s TO [13].

Future work will focus on integrating the system into contact-rich and complex manipulation frameworks such as ALOHA [49] to study high-level intelligence. Additionally, we aim to extend the system to mobile manipulators and multi-robot scenarios operating in variable workspaces.

ACKNOWLEDGMENT

We appreciate Dayoung Kim for helping camera setup and calibration.

REFERENCES

- [1] M. Zhang, Y. Ma, T. Miki, and M. Hutter, “Learning to open and traverse doors with a legged manipulator,” *arXiv preprint arXiv:2409.04882*, 2024.
- [2] Z. Fu, T. Z. Zhao, and C. Finn, “Mobile aloha: Learning bimanual mobile manipulation with low-cost whole-body teleoperation,” *arXiv preprint arXiv:2401.02117*, 2024.
- [3] S. Haddadin, A. De Luca, and A. Albu-Schäffer, “Robot collisions: A survey on detection, isolation, and identification,” *IEEE Transactions on Robotics*, vol. 33, no. 6, pp. 1292–1312, 2017.
- [4] Y. J. Heo, D. Kim, W. Lee, H. Kim, J. Park, and W. K. Chung, “Collision detection for industrial collaborative robots: A deep learning approach,” *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 740–746, 2019.
- [5] D. Ko, D. Lee, W. Lee, K. Kim, and W. K. Chung, “Compensation modulation for tracking accuracy in free motion and compliance during interaction,” *IEEE/ASME Transactions on Mechatronics*, 2024.
- [6] L. Han, W. Xu, B. Li, and P. Kang, “Collision detection and coordinated compliance control for a dual-arm robot without force/torque sensing based on momentum observer,” *IEEE/ASME Transactions on Mechatronics*, vol. 24, no. 5, pp. 2261–2272, 2019.
- [7] S. Proia, R. Carli, G. Cavone, and M. Dotoli, “Control techniques for safe, ergonomic, and efficient human-robot collaboration in the digital industry: A survey,” *IEEE Transactions on Automation Science and Engineering*, vol. 19, no. 3, pp. 1798–1819, 2021.
- [8] O. I. de Normalización, *ISO-TS 15066: Robots and robotic devices: Collaborative robots*. ISO, 2016.

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

- [9] E. Magrini and A. De Luca, "Human-robot coexistence and contact handling with redundant robots," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2017, pp. 4611–4617.
- [10] A. Casalino, D. Bazzi, A. M. Zanchettin, and P. Rocco, "Optimal proactive path planning for collaborative robots in industrial contexts," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 6540–6546.
- [11] J. de Gea Fernández, D. Mronga, M. Günther, T. Knobloch, M. Wirkus, M. Schröer, M. Trampler, S. Stiene, E. Kirchner, V. Bargsten, et al., "Multimodal sensor-based whole-body control for human-robot collaboration in industrial settings," *Robotics and Autonomous Systems*, vol. 94, pp. 102–119, 2017.
- [12] C. Sloth and H. G. Petersen, "Computation of safe path velocity for collaborative robots," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 6142–6148.
- [13] B. Sundaralingam, S. K. S. Hari, A. Fishman, C. Garrett, K. Van Wyk, V. Blukis, A. Millane, H. Oleynikova, A. Handa, F. Ramos, et al., "Curobo: Parallelized collision-free robot motion generation," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 8112–8119.
- [14] M. Faroni, M. Beschi, and N. Pedrocchi, "Safety-aware time-optimal motion planning with uncertain human state estimation," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 12 219–12 226, 2022.
- [15] K. Srinivasan, B. Eysenbach, S. Ha, J. Tan, and C. Finn, "Learning to be safe: Deep rl with a safety critic," *arXiv preprint arXiv:2010.14603*, 2020.
- [16] D. Hoeller, N. Rudin, D. Sako, and M. Hutter, "Anymal parkour: Learning agile navigation for quadrupedal robots," *Science Robotics*, vol. 9, no. 88, p. eadi7566, 2024.
- [17] H. Oleynikova, Z. Taylor, M. Fehr, R. Siegwart, and J. Nieto, "Voxblox: Incremental 3d euclidean signed distance fields for on-board mav planning," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2017, pp. 1366–1373.
- [18] P. Liu, K. Zhang, D. Tateo, S. Jauhri, J. Peters, and G. Chalkatzaki, "Regularized deep signed distance fields for reactive motion generation," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 6673–6680.
- [19] M. Bhardwaj, B. Sundaralingam, A. Mousavian, N. D. Ratliff, D. Fox, F. Ramos, and B. Boots, "STORM: An integrated framework for fast joint-space model-predictive control for reactive manipulation," 2021.
- [20] M. Kulkarni and K. Alexis, "Reinforcement learning for collision-free flight exploiting deep collision encoding," *arXiv preprint arXiv:2402.03947*, 2024.
- [21] T. Miki, J. Lee, L. Wellhausen, and M. Hutter, "Learning to walk in confined spaces using 3d representation," *arXiv preprint arXiv:2403.00187*, 2024.
- [22] B. Yang, Q. Zhang, R. Geng, L. Wang, and M. Liu, "Real-time neural dense elevation mapping for urban terrain with uncertainty estimations," *IEEE Robotics and Automation Letters*, vol. 8, no. 2, pp. 696–703, 2022.
- [23] S. Luo, S. Li, R. Yu, Z. Wang, J. Wu, and Q. Zhu, "Pie: Parkour with implicit-explicit learning framework for legged robots," *IEEE Robotics and Automation Letters*, 2024.
- [24] S. LaValle, "Rapidly-exploring random trees: A new tool for path planning," *Research Report 9811*, 1998.
- [25] O. Khatib, "Real-time obstacle avoidance for manipulators and mobile robots," *The international journal of robotics research*, vol. 5, no. 1, pp. 90–98, 1986.
- [26] N. Ratliff, M. Zucker, J. A. Bagnell, and S. Srinivasa, "Chomp: Gradient optimization techniques for efficient motion planning," in *2009 IEEE international conference on robotics and automation*. IEEE, 2009, pp. 489–494.
- [27] J. Lee, M. Bjelonic, A. Reske, L. Wellhausen, T. Miki, and M. Hutter, "Learning robust autonomous navigation and locomotion for wheeled-legged robots," *Science Robotics*, vol. 9, no. 89, p. eadi9641, 2024.
- [28] J. Thumm and M. Althoff, "Provably safe deep reinforcement learning for robotic manipulation in human environments," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 6344–6350.
- [29] D. Kim, M. Srouji, C. Chen, and J. Zhang, "Armor: Egocentric perception for humanoid robot collision avoidance and motion planning," *arXiv preprint arXiv:2412.00396*, 2024.
- [30] J. Garcia and F. Fernández, "A comprehensive survey on safe reinforcement learning," *Journal of Machine Learning Research*, vol. 16, no. 1, pp. 1437–1480, 2015.
- [31] M. Heger, "Consideration of risk in reinforcement learning," in *Machine Learning Proceedings 1994*. Elsevier, 1994, pp. 105–111.
- [32] P. Geibel and F. Wyszotki, "Risk-sensitive reinforcement learning applied to control under constraints," *Journal of Artificial Intelligence Research*, vol. 24, pp. 81–108, 2005.
- [33] Y. C. Tang, J. Zhang, and R. Salakhutdinov, "Worst cases policy gradients," *arXiv preprint arXiv:1911.03618*, 2019.
- [34] L. Schneider, J. Frey, T. Miki, and M. Hutter, "Learning risk-aware quadrupedal locomotion using distributional reinforcement learning," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 11 451–11 458.
- [35] E. Altman, *Constrained Markov decision processes*. Routledge, 2021.
- [36] L. Shen, L. Yang, S. Chen, B. Yuan, X. Wang, D. Tao, et al., "Penalized proximal policy optimization for safe reinforcement learning," *arXiv preprint arXiv:2205.11814*, 2022.
- [37] Y. Liu, J. Ding, and X. Liu, "Ipo: Interior-point policy optimization under constraints," in *Proceedings of the AAAI conference on artificial intelligence*, vol. 34, no. 04, 2020, pp. 4940–4947.
- [38] Y. Kim, H. Oh, J. Lee, J. Choi, G. Ji, M. Jung, D. Youm, and J. Hwangbo, "Not only rewards but also constraints: Applications on legged robot locomotion," *IEEE Transactions on Robotics*, 2024.
- [39] J. Lee, L. Schroth, V. Klemm, M. Bjelonic, A. Reske, and M. Hutter, "Evaluation of constrained reinforcement learning algorithms for legged locomotion," *arXiv preprint arXiv:2309.15430*, 2023.
- [40] J. Achiam, D. Held, A. Tamar, and P. Abbeel, "Constrained policy optimization," in *International conference on machine learning*. PMLR, 2017, pp. 22–31.
- [41] H.-S. Fang, C. Wang, H. Fang, M. Gou, J. Liu, H. Yan, W. Liu, Y. Xie, and C. Lu, "Anygrasp: Robust and efficient grasp perception in spatial and temporal domains," *IEEE Transactions on Robotics*, 2023.
- [42] M. Macklin, "Warp: A high-performance python framework for gpu simulation and graphics," <https://github.com/nvidia/warp>, March 2022, nVIDIA GPU Technology Conference (GTC).
- [43] S. M. Khansari-Zadeh and A. Billard, "A dynamical system approach to realtime obstacle avoidance," *Autonomous Robots*, vol. 32, pp. 433–454, 2012.
- [44] A. Fishman, A. Walsman, M. Bhardwaj, W. Yuan, B. Sundaralingam, B. Boots, and D. Fox, "Avoid everything: Model-free collision avoidance with expert-guided fine-tuning," in *CoRL Workshop on Safe and Robust Robot Learning for Operation in the Real World*, 2024.
- [45] D. Rakita, H. Shi, B. Mutlu, and M. Gleicher, "Collisionik: A per-instant pose optimization method for generating robot motions with environment collision avoidance," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 9995–10001.
- [46] V. Vasilopoulos, S. Garg, P. Piacenza, J. Huh, and V. Isler, "Ramp: Hierarchical reactive motion planning for manipulation tasks using implicit signed distance functions," in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2023, pp. 10 551–10 558.
- [47] M. Becker, P. Caspers, T. Lilje, S. Haddadin, and M. A. Müller, "Informed circular fields: a global reactive obstacle avoidance framework for robotic manipulators," *Frontiers in Robotics and AI*, vol. 11, p. 1447351, 2025.
- [48] M. Danielczuk, A. Mousavian, C. Eppner, and D. Fox, "Object rearrangement using learned implicit collision functions," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 6010–6017.
- [49] T. Z. Zhao, J. Tompson, D. Driess, P. Florence, K. Ghasemipour, C. Finn, and A. Wahid, "Aloha unleashed: A simple recipe for robot dexterity," *arXiv preprint arXiv:2410.13126*, 2024.