

Direct Contact-Tolerant Motion Planning With Vision Language Models

He Li¹, Jian Sun¹, Chengyang Li³, Guoliang Li¹, Qiyu Ruan¹, Shuai Wang^{2,†}, and Chengzhong Xu^{1,†}

Abstract—Navigation in cluttered environments often requires robots to tolerate contact with movable or deformable objects to maintain efficiency. Existing contact-tolerant motion planning (CTMP) methods rely on indirect spatial representations (e.g., prebuilt map, obstacle set), resulting in inaccuracies and a lack of adaptiveness to environmental uncertainties. To address this issue, we propose a direct contact-tolerant (DCT) planner, which integrates vision-language models (VLMs) into direct point perception and navigation, including two key components. The first one is VLM point cloud partitioner (VPP), which performs contact-tolerance reasoning in image space using VLM, caches inference masks, propagates them across frames using odometry, and projects them onto the current scan to generate a contact-aware point cloud. The second innovation is VPP guided navigation (VGN), which formulates CTMP as a perception-to-control optimization problem under direct contact-aware point cloud constraints, which is further solved by a specialized deep neural network (DNN). We implement DCT in Isaac Sim and a real car-like robot, demonstrating that DCT achieves robust and efficient navigation in cluttered environments with movable obstacles, outperforming representative baselines across diverse metrics. The code is available at: <https://github.com/ChrisLeeUM/DCT>.

I. INTRODUCTION

Navigation in complex environments is a fundamental task for autonomous mobile robots. Traditional navigation algorithms assume strict collision avoidance, treating each obstacle as a rigid body that must be completely avoided [1]–[3]. However, they can fail in cases when no collision-free paths exist, especially in highly cluttered scenarios.

In practice, many obstacles (e.g., curtains, empty boxes) are movable or deformable and can be safely contacted. Therefore, by allowing controlled contact with such objects, a robot can navigate through spaces otherwise blocked, achieving higher efficiency [4], [5]. This navigation problem, known as contact-tolerant motion planning (CTMP), centers on reasoning whether obstacles can be safely contacted [6], [7] and planning paths among such obstacles [8], [9].

Existing CTMP methods can be classified into three paradigms: search-based, reinforcement learning (RL) based,

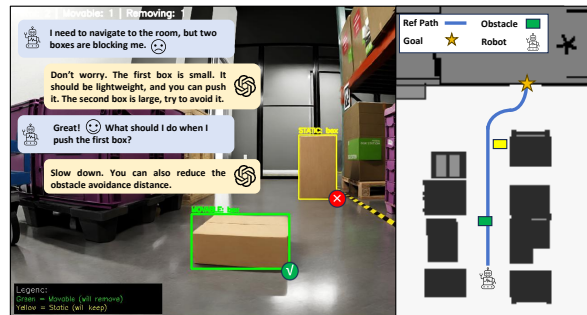
This work was jointly funded by MOST and The Science and Technology Development Fund (FDCT), Macau SAR (File no. 0074/2025/AMJ), National Natural Science Foundation of China (Grant No. 62371444), Shenzhen Science and Technology Program (Grant No. RCYX20231211090206005), TUBITAK-CAS Bilateral Cooperation Program (Grant No. 321GJHZ2025118MI), and Shenzhen Institutes of Advanced Technology Original Exploration Program.

¹He Li, Jian Sun, Guoliang Li, Qiyu Ruan, and Chengzhong Xu are with the State Key Laboratory of Internet of Things for Smart City (SKL-IOTSC), University of Macau, Macau, China.

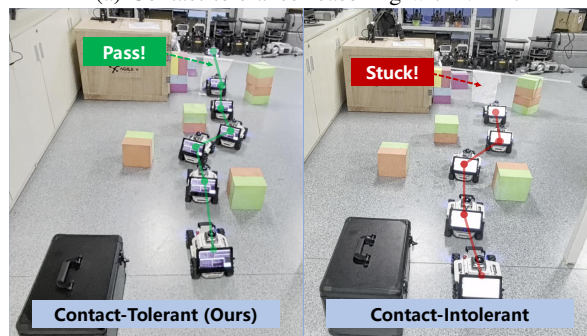
²Shuai Wang is with the Shenzhen Institutes of Advanced Technology (SIAT), Chinese Academy of Sciences, Shenzhen, China.

³Chengyang Li is with the Department of Electrical and Computer Engineering, The University of Hong Kong, Hong Kong, China.

Corresponding author: Chengzhong Xu (czxu@um.edu.mo) and Shuai Wang (s.wang@siat.ac.cn).



(a) Contact-tolerance reasoning with VLM.



(b) Contact-tolerant vs. contact-intolerant.

Fig. 1: DCT motion planning with VLM.

and hybrid planning approaches. Search-based methods [10]–[12] formulate CTMP as a combinatorial search over obstacle positions and movability cases under all possible interaction sequences prior to execution. These approaches require prior environmental knowledge and scale poorly as the number of obstacles increases. RL-based methods [7], [13], [14] formulate CTMP as a policy-learning problem, where neural networks optimize action selection through interactions with fixed and movable obstacles. These approaches are sensitive to training distribution and often face challenges in real-world deployment. Furthermore, how to set appropriate penalty functions for various movable obstacles remains challenging. Hybrid planning approaches address these limitations by integrating global search with local perception and planning. In this paradigm, movable obstacles are typically treated as free or low-cost regions during global search, with their movability assessed dynamically at runtime using visual models [6], robot-obstacle contacts [15], or receding-horizon optimization [9]. These strategies reduce exhaustive global search and improve flexibility in uncertain environments.

Despite these advancements in the development of CTMP, existing approaches often rely on indirect spatial representations (e.g., prebuilt map, obstacle set). This would inevitably introduce errors (e.g., a convex set cannot match an arbitrary shape) to the planner, leading to conservative policies, or increasing the risk of collisions. Using a map may also

result in a lack of adaptiveness to environmental changes. In addition, existing methods struggle with movability reasoning, which is jointly dependent on robot capabilities, obstacle properties, and task-specific requirements. These issues reduce their robustness and generalizability under real-world uncertainties. To fill the gap, we propose DCT, a direct contact-tolerant (DCT) motion planning system that leverages vision–language models (VLMs) for contact-tolerant point-to-action navigation, as shown in Fig. 1. The DCT consists of two key modules: the *VLM point cloud partitioner (VPP)* and *VPP guided navigation (VGN)*.

The VPP module identifies candidate movable obstacles from RGB images using an open-set detector guided by language prompts, and applies a task-conditioned VLM for mask filtering. Since such VLM reasoning cannot be triggered for each point frame due to its inference latency, we propose to cache the (prompt, caption, mask) pair, together with the robot pose obtained using lidar-inertial odometry, into a temporal memory list. When the subsequent scan arrives, VPP will retrieve the mask by calling memory from previous VLM executions, propagate the mask across frames through planar homography, and associate the mask with the current lidar scans. This VLM temporal memory design enables the robot to partition each scan into contact-tolerant and contact-intolerant point sets at high frequency.

The VGN module will adopt the contact-partitioned point cloud to formulate the *direct point distance constraints*. This is in contrast to existing CTMP that adopts indirect distance constraints. However, this would result in a large-scale model predictive control (LMPC) problem with thousands of constraints, which cannot be solved by commercial solvers in real time. Therefore, we propose a specialized deep neural network (DNN), which trains an agent to solve this LMPC by imitating the optimization process, such that time-consuming iterative computations are converted into real-time feed-forward inference. In addition, VGN also includes a correcting mode: when pushing fails, contact points will be relabeled as contact-intolerant, and the robot reverses to a safe state for replanning.

Our contributions are summarized as follows:

- We propose VPP, a real-time point partitioner that identifies movable objects using VLM contact-tolerance reasoning and memory-based mask propagation.
- We propose VGN, a fast learned planner that directly operates on contact-partitioned point cloud and supports real-time robot planning and control.
- We implement DCT in both Isaac Sim and real-world environments, with extensive evaluations showing superior navigation performance compared to various baselines in cluttered scenarios.

II. RELATED WORK

Search-based CTMP treats the problem as a global search task. The planner explicitly reasons over the positions and movability of obstacles to construct a full interaction sequence before execution. Such formulations typically require deciding which obstacles to move, in what order, and how far,

yielding an NP-hard combinatorial problem [16]. To solve this problem, existing search-based CTMP methods mostly focus on how to reduce the search or sample complexity. For instance, by constraining each object to a single interaction, it is possible to employ a reverse-order search to reduce the search space [10]–[12]. Moreover, large language models can be adopted to guide non-uniform sampling [12], enabling robots to efficiently decide which obstacles to move and in what order to push. However, these methods often rely on a fully observable environment (e.g., map), which may limit their applicability in real-world settings.

RL-based CTMP formulates CTMP as a policy-learning problem, where the neural network learns a policy through interactions with contact-tolerant and contact-intolerant obstacles. Given designed reward functions, the neural network aims to maximize cumulative action rewards of the CTMP task [7], [14], [17]. For instance, a curriculum RL framework is utilized to gradually learn from simple collision avoidance tasks to more complex CTMP tasks [7]. The designed reward functions enable the policy to make contact with movable obstacles when necessary, while discouraging indiscriminate contact. Moreover, domain randomization enhances policy robustness, enabling effective adaptation to the uncertainties in real-world scenarios [14]. Recently, a hierarchical RL framework has been proposed [13], which enables a mobile manipulator to realize CTMP by combining high-level pushing strategy generation with low-level whole-body control. Nonetheless, RL-based CTMP approaches suffer from reward sparsity issues and the difficulty of training in real-world environments.

Hybrid CTMP integrates offline global search with online local planning. In this paradigm, movable obstacles are treated as free space or low-cost regions, and their movability is determined during execution. As such, hybrid CTMP alleviates the computational overhead of exhaustive global search [6], [9], [15]. For instance, extended path planning with pushing interactions has been developed in [6], by combining object localization and movability reasoning. A directed visibility graph (DV-graph) is introduced [15], where edges encode both traversability and obstacle manipulation costs, and interaction planning updates obstacle affordances for future use. In [9], an optimization framework is proposed, which models movable objects as interactive elements and incorporates contact dynamics into trajectory optimization, allowing robots to generate smooth, feasible, and predictable paths in cluttered environments. Nonetheless, existing hybrid methods often rely on indirect spatial representations, which would inevitably introduce errors to the robot actions.

III. PROBLEM STATEMENT

A. Task Description

We consider a CTMP system, which consists of 1 autonomous robot, contact-intolerant fixed obstacles \mathcal{M}^{fix} , and contact-tolerant obstacles \mathcal{M}^{mov} . Given a target s^* , and robot sensor inputs $\{\mathcal{P}, \mathcal{C}, \mathcal{I}\}$, where \mathcal{P} , \mathcal{C} , and \mathcal{I} denote lidar point clouds, camera images, and inertial measurements, respectively, the objective of CTMP is to generate a sequence

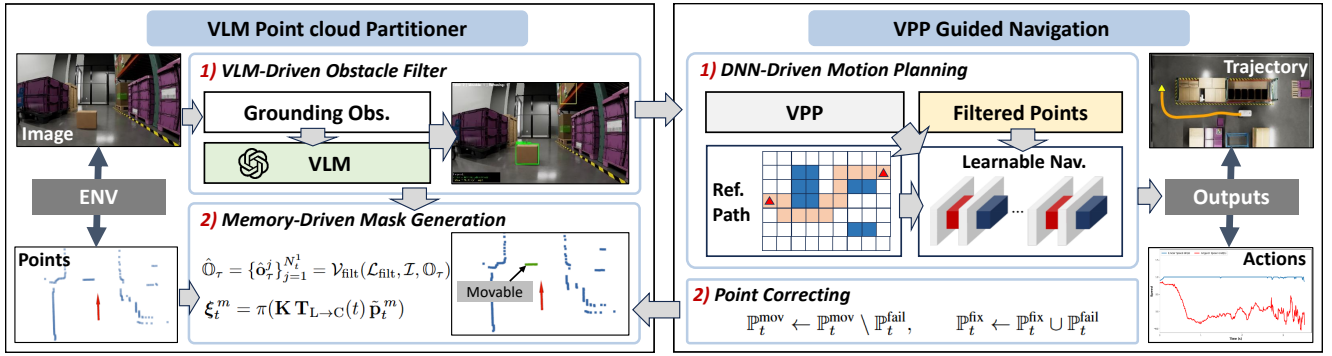


Fig. 2: System architecture of DCT, which consists of VPP and VGN modules.

of robot states \mathcal{S} and executable actions \mathcal{U} reaching the target. This can be realized with a repeated receding-horizon planning, where the system only plans a finite sequence of H discretized future steps $\mathcal{H}_t \triangleq \{t, \dots, t + H\}$, where Δt is the time length between two states. Specifically, at time t , the robot state is denoted as $\mathbf{s}_t = (x_t, y_t, \theta_t)$, where (x_t, y_t) and θ_t are position and orientation. The robot action is $\mathbf{u}_t = (v_t, \psi_t)$, where v_t and ψ_t are the linear and angular velocities. The CTMP task at the current horizon needs to generate state sequence $\mathcal{S}_t = \{\mathbf{s}_t, \dots, \mathbf{s}_{t+H}\}$ and control sequence $\mathcal{U}_t = \{\mathbf{u}_t, \dots, \mathbf{u}_{t+H}\}$ under kinetic constraints $\{\mathcal{S}_t, \mathcal{U}_t\} \in \mathcal{F}_t$, with \mathcal{F}_t detailed in Section IV-C.

B. Contact-Tolerant Collision Avoidance

In CTMP, we need to handle distinct obstacles \mathcal{M}^{fix} and \mathcal{M}^{mov} , and collision avoidance becomes different from conventional motion planning. Particularly, given the current point set $\mathbb{P}_t = \{\mathbf{p}_t^1, \dots, \mathbf{p}_t^M\}$ obtained from lidar scans, where $\mathbf{p}_t^i \in \mathbb{R}^3$ denotes the position of the i -th point and M is the total number of obstacle points, we partition \mathbb{P}_t as $\mathbb{P}_t = \mathbb{P}_t^{\text{mov}} \cup \mathbb{P}_t^{\text{fix}}$, where $\mathbb{P}_t^{\text{mov}}$ denotes contactable points and $\mathbb{P}_t^{\text{fix}}$ denotes contact-intolerant points. For CTMP, contact with contactable points $\mathbb{P}_t^{\text{mov}}$ is permitted. Therefore, hard collision avoidance should only be applied to $\mathbb{P}_t^{\text{fix}} = \mathbb{P}_t \setminus \mathbb{P}_t^{\text{mov}}$:

$$\text{dist}(\mathbb{G}_t(\mathbf{s}_t), \mathbb{P}_t^{\text{fix}}) \geq d_{\min}, \quad t \in \mathcal{H}_t, \quad (1)$$

where \mathbb{G}_t denotes the robot set and d_{\min} is a predefined safety margin. To determine the distance dist , it is equivalent to compute the distances between any two points within \mathbb{G}_t and $\mathbb{P}_t^{\text{fix}}$, and then take the minimum:

$$\text{dist}(\mathbb{G}_t(\mathbf{s}_t), \mathbb{P}_t^{\text{fix}}) = \min\{\|\mathbf{x} - \mathbf{y}\| \mid \mathbf{x} \in \mathbb{G}_t, \mathbf{y} \in \mathbb{P}_t^{\text{fix}}\}. \quad (2)$$

IV. DIRECT CONTACT-TOLERANT MOTION PLANNING

A. System Overview

The challenges in tackling CTMP constraint (1) are two-fold. First, the point partition process will rely on other sensor inputs $\{\mathcal{C}, \mathcal{I}\}$, and fusing different data modalities with distinct frequencies is non-trivial. Second, the number of points in $\mathbb{P}_t^{\text{fix}}$ is large, making direct computation of (1) computationally intensive. Existing indirect approaches convert $\mathbb{P}_t^{\text{fix}}$ to sets, but this would lead to degradation of solution accuracy.

To address the above challenges, we propose DCT illustrated in Fig. 2, which consists of VPP and VGN modules. The VPP module identifies movable objects using VLM contact-tolerance reasoning and memory-based mask propagation. It processes raw visual and point cloud inputs, grounds object detections, figures out movable elements, and publishes filtered point representations $\mathbb{P}_t^{\text{fix}}$ for downstream planner. The VGN module leverages a fast learned planner to map filtered point sets $\mathbb{P}_t^{\text{fix}}$ into control signals \mathbf{u}_t . It also consists of a point correction block that updates obstacle movability properties upon pushing failure.

B. Movable Obstacle Grounding with VPP

The goal of VPP is to identify lidar points corresponding to *movable* obstacles. VLM inference is slow, whereas point cloud filtering requires real-time updates. To address this challenge, we introduce two complementary strategies: VLM-driven filtering, which leverages the reasoning capability of VLMs to accurately categorize obstacles, and real-time mask generation, which enables efficient, high-frequency filtering of lidar data.

1) **VLM-Driven Obstacle Filter**: At time t , given a grounding prompt \mathcal{L}_g and the image \mathcal{I}_t , we use a grounding model [18] GRD to localize candidate obstacle regions $\mathbb{O}_t = \{\mathbf{o}_t^i\}_{i=1}^{N_t^0}$, where N_t^0 is the number of detected objects. GRD is an open-set detector that couples vision and language to find objects specified by \mathcal{L}_g . Since “movable” depends on task and context (obstacle type, environment, material, and robot capability), we introduce a VLM $\mathcal{V}_{\text{filt}}$ to filter candidates. Conditioned on a task prompt $\mathcal{L}_{\text{filt}}$ together with $\mathcal{I}, \{\mathbf{o}_t^{N_0}\}$, the VLM returns a movable-obstacle mask list

$$\hat{\mathbb{O}}_t = \{\hat{\mathbf{o}}_t^j\}_{j=1}^{N_t^1} = \mathcal{V}_{\text{filt}}(\mathcal{L}_{\text{filt}}, \mathcal{I}, \mathbb{O}_t), \quad (3)$$

where N_t^1 is the number of obstacles after filtering. Then we record time t as refresh time τ , the obstacle list $\hat{\mathbb{O}}_\tau$, and the corresponding robot position as \mathbf{s}_τ .

2) **Memory-Driven Mask Generation**: Due to limited onboard compute, VLM cannot always be triggered for each point frame. In DCT, VLM detection is triggered only when the moving distance exceeds d_{thres} or the time since the last detection exceeds t_{thres} . Between two visual inferences, we therefore *reuse and propagate* the last trusted masks and reconcile them when a new detection arrives. This design will bring challenges to multi-modal data alignment due to

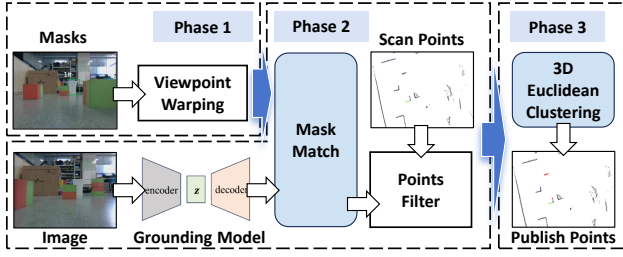


Fig. 3: Memory-driven mask generation.

the movement of the robot. We address it with a three-phase procedure shown in Fig. 3: (i) viewpoint warping; (ii) detection-driven reconciliation; and (iii) point-level refinement and publishing.

a) *Phase I: Viewpoint warping*: Given the robot state at time t and calibrated intrinsics/extrinsics obtained from lidar-inertial odometry, we compute the inter-frame homography $\mathbf{H}_{\tau \rightarrow t}$ and warp each mask:

$$\hat{\mathbf{o}}_t^n = \mathcal{W}(\hat{\mathbf{o}}_\tau^n, \mathbf{H}_{\tau \rightarrow t}), \quad \hat{\mathcal{O}}_t = \{\hat{\mathbf{o}}_t^j\}_{j=1}^{N_\tau^1}, \quad (4)$$

where $\mathcal{W}(\cdot)$ applies homogeneous projection to every pixel $(x, y) \in \hat{\mathbf{o}}_\tau^n$. Because of inevitable state-estimation error, $\hat{\mathcal{O}}_t$ may still be misaligned with the current scene.

b) *Phase II: Detection-driven reconciliation*: Given the new image \mathcal{I}_t , the grounding model yields candidate movable regions $\mathcal{O}_t = \{\mathbf{o}_t^i\}_{i=1}^{N_t^0}$. We match these to the propagated set $\hat{\mathcal{O}}_t$ using the intersection-over-union

$$\text{IoU}(\mathbb{A}, \mathbb{B}) = \frac{|\mathbb{A} \cap \mathbb{B}|}{|\mathbb{A} \cup \mathbb{B}|}. \quad (5)$$

For each $\hat{\mathbf{o}}_t^j$ in $\hat{\mathcal{O}}_t$, define $i^*(j) = \arg \max_i \text{IoU}(\hat{\mathbf{o}}_t^j, \mathbf{o}_t^i)$, where $i^*(j)$ denotes the index of the detection in \mathcal{O}_t that attains the highest IoU with the propagated mask $\hat{\mathbf{o}}_t^j$. If $\text{IoU}(\hat{\mathbf{o}}_t^j, \mathbf{o}_t^{i^*(j)}) \geq \sigma_{\text{IoU}}$ (e.g., $\sigma_{\text{IoU}} = 0.90$), we regard the pair as a successful match and *update* the trusted mask geometry $\hat{\mathbf{o}}_t^j \leftarrow \mathbf{o}_t^{i^*(j)}$. Propagated masks without a valid match are discarded for safety. Conversely, unmatched detections in \mathcal{O}_t are *inserted* as new trusted masks. The result is the reconciled set $\hat{\mathcal{O}}_t$.

c) *Phase III: Point-level refinement and publishing*: Let the lidar scan at time t be

$$\mathbb{P}_t = \{\mathbf{p}_t^m\}_{m=1}^M, \quad \mathbf{p}_t^m = [x_t^m, y_t^m, z_t^m]^\top \in \mathbb{R}^3. \quad (6)$$

Each point is augmented to $\tilde{\mathbf{p}}_t^m = [x_t^m, y_t^m, z_t^m, 1]^\top$ and projected to the image as

$$\boldsymbol{\xi}_t^m = \pi(\mathbf{K} \mathbf{T}_{L \rightarrow C}(t) \tilde{\mathbf{p}}_t^m), \quad \pi([a, b, c]^\top) = [a/c, b/c]^\top, \quad (7)$$

with $\boldsymbol{\xi}_t^m \in \mathbb{R}^2$ the pixel coordinates. Using the reconciled masks $\hat{\mathcal{O}}_t$ (let $\hat{\mathcal{U}}_t = \bigcup \hat{\mathbf{o}}_t^i$), the scan is partitioned as

$$\mathbb{P}_t^{\text{mov}} = \{\mathbf{p}_t^m \in \mathbb{P}_t \mid \boldsymbol{\xi}_t^m \in \hat{\mathcal{U}}_t\}, \quad (8)$$

$$\mathbb{P}_t^{\text{fix}} = \mathbb{P}_t \setminus \mathbb{P}_t^{\text{mov}}. \quad (9)$$

While mask gating is effective, image-level decisions ignore 3D continuity and may include isolated speckles or miss nearby returns from the same object. We therefore perform 3D Euclidean clustering (e.g., DBSCAN with radius

$\varepsilon_{\text{cluster}}$ and minimum size N_{min}) on $\mathbb{P}_t^{\text{mov}}$ to obtain clusters $\mathcal{C}_t^{\text{mov}} = \{\mathcal{C}_k\}$. We then apply two symmetric corrections:

- 1) *Outlier suppression*: points selected by the mask but not belonging to any cluster \mathcal{C}_k with $|\mathcal{C}_k| \geq N_{\text{min}}$ are *removed* from movable: $\mathbb{P}_t^{\text{mov}} \leftarrow \bigcup_{k: |\mathcal{C}_k| \geq N_{\text{min}}} \mathcal{C}_k$.
- 2) *Cluster completion*: points in $\mathbb{P}_t^{\text{fix}}$ that lie within distance $\varepsilon_{\text{cluster}}$ of some movable cluster \mathcal{C}_k (i.e., $\text{dist}(\mathbf{p}, \mathcal{C}_k) \leq \varepsilon_{\text{cluster}}$) are *added* to movable: $\mathbb{P}_t^{\text{mov}} \leftarrow \mathbb{P}_t^{\text{mov}} \cup \{\mathbf{p} \in \mathbb{P}_t^{\text{fix}} \mid \text{dist}(\mathbf{p}, \mathcal{C}_k) \leq \varepsilon_{\text{cluster}}, \exists k\}$.

Finally, we publish the refined streams $(\mathbb{P}_t^{\text{mov}}, \mathbb{P}_t^{\text{fix}})$ at the lidar rate. This pipeline preserves temporal consistency under sparse visual updates while producing spatially coherent point sets for downstream planning and control.

C. Fast Motion Planning with VGN

1) *Learned Distance*: Given $\mathbb{P}_t^{\text{fix}}$ generated by VPP, the computation of distance dist is a convex optimization problem given by [1]:

$$\forall \mathbf{p}_t^i \in \mathbb{P}_t^{\text{fix}} : \min_{\mathbf{x}} \|\mathbf{R}_t(\mathbf{s}_t)\mathbf{x} + \mathbf{t}_t(\mathbf{s}_t) - \mathbf{p}_t^i\|_2^2 \quad (10)$$

s.t. $\mathbf{G}\mathbf{x} \leq \mathbf{g}$,

where $\mathbf{R}_t(\mathbf{s}_t)$ is the rotation matrix representing the orientation of the robot, $\mathbf{t}_t(\mathbf{s}_t)$ is the translation vector denoting the position of the robot, and matrices (\mathbf{G}, \mathbf{g}) define the robot edges. However, the complexity is high when the number of points in $\mathbb{P}_t^{\text{fix}}$ is large, resulting in unacceptable latency.

To address this issue, we propose a DNN method for faster distance computation following [19]. The DNN trains an agent to solve (10) by imitating the optimal algorithm, such that time-consuming iterative computations are converted into real-time feed-forward inference. The system architecture of DNN is shown in Fig. 4, which consists of three phases.

a) *Phase I: Dual Transformation*: We first transform problem (10) to its dual formulation under the strong duality property [3], [20]:

$$\forall \mathbf{p}_t^i \in \mathbb{P}_t^{\text{fix}} : \max_{\boldsymbol{\lambda}_t^i, \boldsymbol{\mu}_t^i} \boldsymbol{\mu}_t^{i\top} (\mathbf{G}\mathbf{R}_t(\mathbf{s}_t)^\top [\mathbf{p}_t^i - \mathbf{t}_t(\mathbf{s}_t)] - \mathbf{g}) \quad (11)$$

s.t. $\boldsymbol{\mu}_t^i \geq \mathbf{0}$, $\|\boldsymbol{\lambda}_t^i\| \leq 1$, $\boldsymbol{\mu}_t^{i\top} \mathbf{G} + \boldsymbol{\lambda}_t^{i\top} \mathbf{R}_t(\mathbf{s}_t) = \mathbf{0}$,

where $\{\boldsymbol{\mu}_t^i, \boldsymbol{\lambda}_t^i\}$ are dual variables.

b) *Phase II: Demonstration Data Generation*: We solve the dual problem (11) to generate the demonstration dataset $\mathcal{D} = \{\mathcal{A}^{(j)}\}_{j=1}^J$. The j -th sample is given by $\mathcal{A}^{(j)} = \{\mathcal{X}^{(j)}, \mathcal{Y}^{(j)}\}$, where the input tuple is given by

$$\mathcal{X}^{(j)} = \left\{ \mathbf{G}, \mathbf{g}, \mathbf{s}_t^{(j)}, \mathbf{p}_t^{(j)} \right\},$$

and the output tuple $\mathcal{Y}^{(j)} = \{\boldsymbol{\mu}_t^{(j)}, \boldsymbol{\lambda}_t^{(j)}\}$.

c) *Phase III: DNN Training and Inference*: The DNN consists of 6 fully connected layers with 32 units, with ReLU, normalization, and tanh activation functions. For model training, we use the mean squared error as the loss function and optimize the network parameters using stochastic gradient descent. The pre-trained DNN model is deployed at the robot to make real-time inference.




Fig. 4: Deep unfolded neural network.

With DNN, the optimal solution of $\{\lambda_t^i, \mu_t^i\}$ to (11) is obtained in microseconds. The distance in (2) is thus $\text{dist}(\mathbb{G}_h(\mathbf{s}_h), \mathbb{P}_h^{\text{fix}}) = \text{DNN}(\mathbf{G}, \mathbf{g}, \mathbf{s}_h, \mathbb{P}_h^{\text{fix}})$, where $\text{DNN}(\cdot)$ takes the minimum of the objective of (11) for all $\mathbf{p}_h^i \in \mathbb{P}_h^{\text{fix}}$.

2) **Motion Planning**: the VGN problem can be cast as an optimization problem.

$$\mathbb{P}_t : \min_{\mathcal{U}_t, \mathcal{S}_t} C_t(\mathcal{S}_t) \quad (12a)$$

$$\text{s.t. } \mathbf{s}_{h+1} = \mathbf{s}_h + f(\mathbf{s}_h, \mathbf{u}_h)\Delta t, \quad h \in \mathcal{H}_t \quad (12b)$$

$$\mathbf{u}_{\min} \preceq \mathbf{u}_h \preceq \mathbf{u}_{\max}, \quad h \in \mathcal{H}_t \quad (12c)$$

$$\text{DNN}(\mathbf{G}, \mathbf{g}, \mathbf{s}_h, \mathbb{P}_h^{\text{fix}}) \geq d_{\min}, \quad h \in \mathcal{H}_t, \quad (12d)$$

where $C_t(\mathcal{S}_t)$ is the system utility encouraging the robot to approach target \mathbf{s}^* . It measures the distance deviation from a given reference path:

$$C_t(\mathcal{S}_t) = \sum_{h \in \mathcal{H}_t} \|\mathbf{s}_{h+1} - \mathbf{s}_{h+1}^\diamond\|_2^2, \quad (13)$$

where $\mathbf{s}_{h+1}^\diamond = \{\mathbf{s}_{t+1}^\diamond, \dots, \mathbf{s}_{t+H}^\diamond\}$ is the waypoints that are computed using a global path planner (e.g., A*) and \mathbf{s}^* . Note that the kinetic set \mathcal{F}_t in Section III now explicitly becomes $\mathcal{F}_t = \{(12b), (12c), (12d)\}$, where (12b) represents the state evolution ($f(\cdot)$ defines vehicle dynamics) and (12c) represents the physical constraint [1] ($\mathbf{u}_{\min}, \mathbf{u}_{\max}$ are control limits). We collectively define these time-invariant parameters as $\mathcal{P} = \{f, \mathbf{u}_{\min}, \mathbf{u}_{\max}, d_{\min}, \mathbf{G}, \mathbf{g}\}$.

This problem can be solved by alternately optimizing $\{\mathcal{U}_t, \mathcal{S}_t\}$ and executing DNN, as shown in [19]. The entire DCT procedure is summarized in Algorithm 1.

3) **Point Correcting**: VGN enters correcting mode when the robot fails to push through or cross an object. We also treat the robot as stuck if the linear speed is not zero but the pose change between two time steps is very small. Let $\Delta d_t = \|\mathbf{s}_t - \mathbf{s}_{t-1}\|_2$. If $v_t \geq v_{\min}$ and $\Delta d_t \leq d_{\text{stuck}}$ for τ_{stuck} consecutive steps, the correction is triggered.

First, we update the point labels. We find the points that belong to obstacles that did not move during the failed push and denote them by $\mathbb{P}_t^{\text{fail}}$. These points are re-labeled as non-movable. The point sets are then updated as

$$\mathbb{P}_t^{\text{mov}} \leftarrow \mathbb{P}_t^{\text{mov}} \setminus \mathbb{P}_t^{\text{fail}}, \quad \mathbb{P}_t^{\text{fix}} \leftarrow \mathbb{P}_t^{\text{fix}} \cup \mathbb{P}_t^{\text{fail}}. \quad (14)$$

Next, the robot moves backward to a previous safe pose. We keep a short history of states $\mathbb{S}^* = \{\mathbf{s}_t, \dots, \mathbf{s}_{t-F}\}$, where F is the number of traceback steps and \mathbf{s}_{t-F} is the last feasible checkpoint. We temporarily replace the reference path with \mathbb{S}^* (traversed in reverse) and execute reverse navigation with speed limits $v_t \in [-v_{\text{rev}}, 0], v_{\text{rev}} > 0$. The robot reverses along the historical states until it reaches \mathbf{s}_{t-F} . Finally, the global planner computes a new reference path \mathbb{S}^{new} that avoids the newly fixed points, and the robot resumes normal navigation along this path.

Algorithm 1: DCT Motion Planner

Input: RGB images $\{\mathcal{I}_t\}$, point clouds $\{\mathbb{P}_t\}$, robot poses $\{\mathbf{s}_t\}$, target \mathbf{s}^* , grounding model GRD, task prompt $\mathcal{L}_{\text{filt}}$, model $\mathcal{V}_{\text{filt}}$, thresholds $d_{\text{thres}}, t_{\text{thres}}, \sigma_{\text{IoU}}$, parameters \mathcal{P}

Output: Planned trajectory \mathbf{s}_t , control sequence \mathbf{u}_t

```

1 Memory list  $\hat{\mathbb{O}}_\tau$ ; last refresh  $(\tau, \mathbf{s}_\tau)$ ;
2 for each time step  $t$  do
   // VLM-Driven Obstacle Filter
3 if  $\text{dist}(\mathbf{s}_t, \mathbf{s}_\tau) > d_{\text{thres}}$  or  $(t - \tau) > t_{\text{thres}}$  or
    $\hat{\mathbb{O}}_\tau = \emptyset$  then
4    $\mathbb{O}_t \leftarrow \text{GRD}(\mathcal{I}_t)$ ;
5    $\hat{\mathbb{O}}_t \leftarrow \mathcal{V}_{\text{filt}}(\mathcal{L}_{\text{filt}}, \mathcal{I}_t, \mathbb{O}_t)$ ;
6    $(\tau, \mathbf{s}_\tau) \leftarrow (t, \mathbf{s}_t), \hat{\mathbb{O}}_\tau \leftarrow \hat{\mathbb{O}}_t$ ;
   // Mem-Driven Mask Generation
7 Warp  $\hat{\mathbb{O}}_\tau$  to  $t$  using pose change;
8 Fresh detection  $\mathbb{O}_t \leftarrow \text{GRD}(\mathcal{I}_t)$ ;
9 Reconcile masks by IoU ( $\geq \sigma_{\text{IoU}}$ );
10 Drop unmatched, obtain  $\hat{\mathbb{O}}_t$ , and update:
     $\mathbb{P}_t^{\text{mov}} \leftarrow \text{PointsInside}(\hat{\mathbb{O}}_t), \mathbb{P}_t^{\text{fix}} \leftarrow \mathbb{P}_t \setminus \mathbb{P}_t^{\text{mov}}$ ;
11 Refine  $\mathbb{P}_t^{\text{mov}}, \mathbb{P}_t^{\text{fix}}$  via clustering;
12 Publish  $\mathbb{P}_t^{\text{mov}}, \mathbb{P}_t^{\text{fix}}$ ;
   // VPP guided navigation
13  $(\mathcal{S}_t, \mathcal{U}_t) \leftarrow \text{VGN}(\mathbf{s}_t, \mathbb{P}_t^{\text{fix}}; \mathbf{s}^*, \text{DNN}, \mathcal{P})$ ;

```

V. EXPERIMENTS


We implemented the proposed DCT in Python using ROS Noetic and ROS2. We used Isaac Sim [21], a high-fidelity simulation platform that provides accurate physical modeling and realistic visual simulation for robot testing and training. The simulation experiment is based on Nova Carter equipped with a simulated lidar and an RGB camera, as shown in Fig. 5. The simulation and ROS packages are implemented on a Linux workstation with an Intel i9-11900 CPU and an NVIDIA RTX 4090 GPU.

We also deployed DCT on a real robot in Fig. 6a, with its length, width, longitudinal wheelbase, and lateral wheelbase being 32.2 cm, 22.0 cm, 20.0 cm, and 17.5 cm, respectively. The robot is equipped with a Livox LiDAR, an RGB-D camera, and an onboard Intel NUC computer, and operates in differential-drive mode.

A. Validation of VPP Under Different VLMs


In this experiment, we assess the effectiveness of VPP. We construct an evaluation dataset, which consists of image samples annotated with grounding information. The VPP is required to identify all movable or passable obstacles. We selected four state-of-the-art language models for evaluation: GPT-5 [22], Gemini 2.5 [23], Qwen-vl [24], and Llama 4 [25], and evaluated their performance in terms of Precision, Recall, F1-score, and Accuracy.

The experimental results of different VLMs are summarized in Fig. 7. First, GPT-5 achieved perfect Precision



(a) Nova Cater. (b) Camera and lidar data.

Fig. 5: Simulation setup.



(a) Robot. (b) Case 1. (c) Case 2.

Fig. 6: Experimental setup.

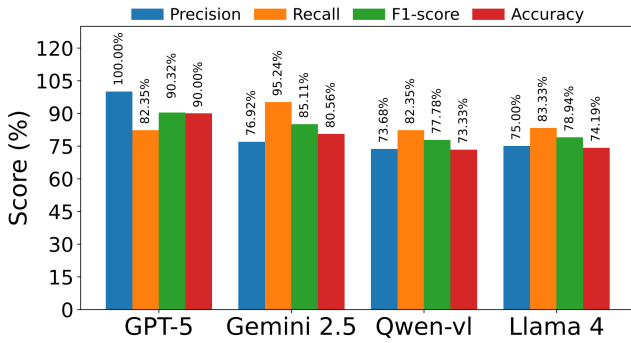


Fig. 7: Evaluation of VPP with various VLMs.

(100%), demonstrating that all identified pushable obstacles were correct. The remaining three models achieved a Precision of about 75%, which introduces a potential risk of collision when robots attempt to push incorrectly identified obstacles. This indicates that GPT-5 offers a more reliable choice for robots operating under strict safety requirements. Second, Gemini 2.5 exhibited the highest Recall (95.24%), more than 10% higher than the other models. This suggests that Gemini 2.5 adopts a more aggressive selection strategy, achieving higher recall at the cost of lower precision. This may benefit exploration-oriented scenarios where occasional failures are acceptable. Third, GPT-5 demonstrated the most balanced overall performance. Its F1-score (90.32%) and Accuracy (90.00%) significantly exceeded those of Gemini 2.5 (85.11%, 80.56%), Qwen-vl (77.78%, 73.33%), and Llama 4 (78.94%, 74.19%). These results confirm that GPT-5 achieves a balanced trade-off across all metrics, delivering both safety and coverage with high reliability. Thus, GPT-5 was selected for the subsequent experiments.

B. Evaluation in Different Obstacle Scenarios

To verify the performance gain brought by DCT, we construct three typical obstacle scenarios in Isaac Sim:

- **Case 1:** Movable obstacle with wide side path;
- **Case 2:** Movable obstacle with narrow side path;
- **Case 3:** Fixed obstacle with wide side path.

The robot is expected to navigate to the given destination

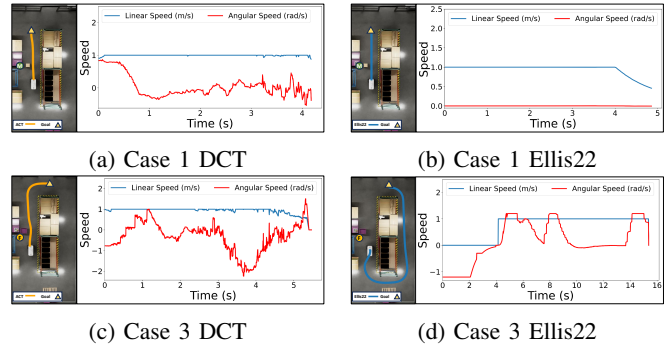


Fig. 8: Evaluation across different obstacle scenarios.

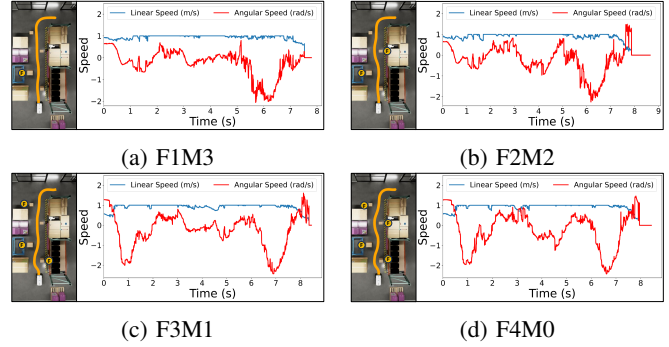


Fig. 9: Evaluation in mixed cluttered environments.

with the obstacle along the route. The target speed is set as 1 m/s for all schemes. We compare the proposed DCT with two representative approaches: NeuPAN [19] and Ellis22 [6]. NeuPAN represents a state-of-the-art direct point navigation method that utilizes lidar measurements to generate a local planned path and accurate control actions. Ellis22 is a CTMP approach, which incorporates object localization, movability reasoning, and path planning into a unified framework. Notably, Ellis22 requires an occupancy grid map and a clear obstacle set. The experiment results are shown in Table I.

In case 1, all three methods successfully reached the goal. DCT achieved the lowest navigation time (4.22s) and the highest average speed (0.915 m/s). In contrast, Ellis22 took the longest navigation time (4.917s) and the lowest average speed (0.793 m/s). This is because Ellis22 adopts the Pure Pursuit for robot control, which prioritizes geometric stability and smooth path tracking, leading to conservative speed control. In contrast, DCT and NeuPAN are trained with efficiency objectives, enabling faster navigation. This can be seen in Fig. 8a, where the speed of DCT remains almost constant at 1 m/s, whereas the speed of Ellis22 exhibits a noticeable decrease before reaching the goal position.

In case 2, the side path was intentionally designed to be too narrow to pass without contact. Both DCT and Ellis22 successfully complete the task, while NeuPAN failed to complete the task, as it modeled the movable obstacle as hard constraints and thus could not generate a feasible trajectory. These results underscore the importance of controllable contact in constrained environments and highlight the efficiency advantage of DCT.

In case 3, when facing a contact-intolerant obstacle,

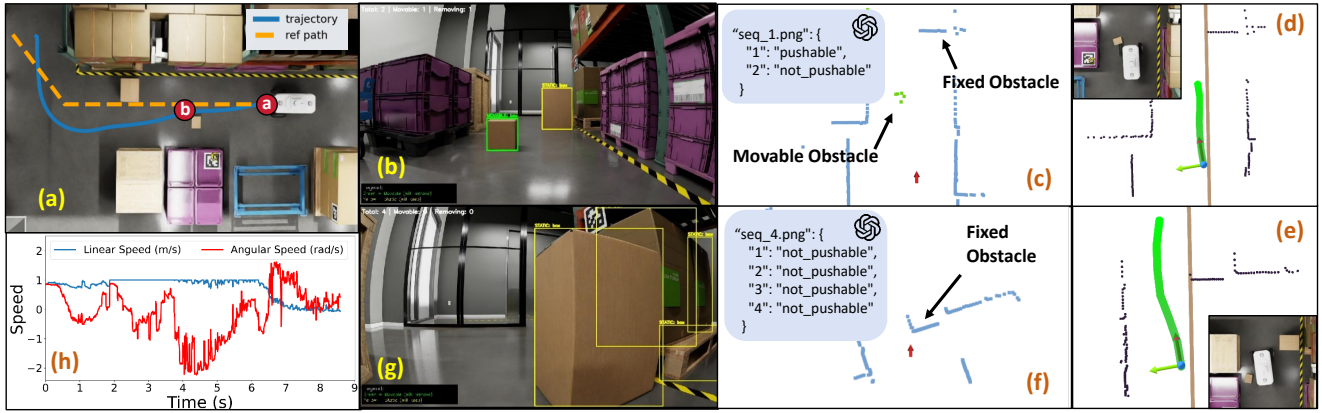


Fig. 10: Contact-tolerance reasoning and trajectory-control profiles of DCT in Isaac Sim.

TABLE I: Evaluations in Different Obstacle Scenarios

	Success			Navigation Time (s)			Average Speed (m/s)			Navigation Distance (m)		
	DCT	NeuPAN	Ellis22	DCT	NeuPAN	Ellis22	DCT	NeuPAN	Ellis22	DCT	NeuPAN	Ellis22
Case 1	✓	✓	✓	4.22	4.37	4.917	0.915	0.889	0.793	3.86	3.89	3.90
Case 2	✓	✗	✓	3.54	-	4.42	0.938	-	0.77	3.22	-	3.40
Case 3	✓	✓	✓	5.72	5.62	15.42	0.933	0.950	0.749	5.33	5.34	11.24

TABLE II: Results in Mixed Cluttered Environments

	SR	Nav. Time (s)	Avg. Speed (m/s)	Nav. Dist. (m)
F4M0	0.7	8.87	0.82	7.27
F3M1	0.9	8.53	0.83	7.08
F2M2	1.0	7.99	0.87	6.95
F1M3	1.0	7.69	0.90	6.92

DCT, NeuPAN, and Ellis22 showed significantly distinct performance. Ellis22 planned a longer route to avoid the obstacle, resulting in the longest navigation time (15.42s) and distance (11.24m). This can be seen in Fig. 8d, where the robot first turns around from 0 to 4.2 s, and then goes around the shelf to reach the end point from 4.2 s to 16 s. This is because the occupancy grid maps of Ellis22 require conservative obstacle inflation, thereby reducing accuracy and limiting applicability in high-precision tasks. In contrast, DCT accomplished the navigation task with significantly shorter navigation time (5.72 s) and distance (5.33 m), as shown in Fig. 8c. These results confirm that DCT adapts well even when obstacles cannot be pushed, maintaining efficiency while being more robust across diverse scenarios.

C. Evaluation in Mixed Cluttered Environments

To further assess the navigation performance of DCT under mixed cluttered environments, we design scenarios with varying numbers of fixed and movable obstacles. The experiment results are presented in Table II.

The results reveal that in highly cluttered environments composed only of fixed obstacles (F4M0), navigation becomes challenging, with the success rate reduced to 70% and both navigation time and distance increased. Changing one obstacle to movable (F3M1) yields a substantial performance gain, achieving a 90% success rate while lowering navigation cost. When two or more movable obstacles are available (F2M2 and F1M3), the robot consistently achieves 100% success while also attaining shorter navigation times, higher average speeds, and reduced path lengths. In particular,

F1M3 reaches the best performance, with the shortest navigation time (7.69 s), highest average speed (0.90 m/s), and shortest path length (6.92 m). The results are corroborated by Fig. 9. When all obstacles are fixed, DCT avoids each one, producing the longest trajectory (Fig. 9d). Allowing contact with movable obstacles reduces avoidance maneuvers and shortens the navigation distance. In particular, when three obstacles are designated as movable, the navigation distance is markedly reduced, as shown in Fig. 9a.

Finally, an illustration of DCT in Isaac Sim is shown in Fig. 10. The result demonstrates that the ability to detect and exploit movable obstacles is crucial in dense environments, as controlled contact improves goal reachability and lowers navigation cost.

D. Verification of DCT on Real Mobile Robot

In this experiment, we evaluate the real-world performance of DCT by deploying it on a ground mobile robot in two representative scenarios: (i) traversing a curtain (shown in Fig. 6b) and (ii) navigating among movable boxes (shown in Fig. 6c). As shown in Fig. 11, in the curtain scenario, DCT identifies the curtain as passable, projects the predicted mask into the scan data, maintains this memory throughout the run, and passes through the curtain to reach the goal position. As shown in Fig. 12, in the box scenario, DCT classifies a small box as movable and navigates to the goal position with harmless contact with the small box (Fig. 12), while maintaining hard collision constraints for all other obstacles. Interestingly, the robot passes under a chair and avoids the chair legs with agility, which demonstrates its capability to handle arbitrary-shape obstacles under the direct point paradigm. Across both experiments, DCT restricts physical interaction strictly to obstacles recognized as movable and enforces strict collision avoidance elsewhere, showing its excellent balance between navigation efficiency and safety.

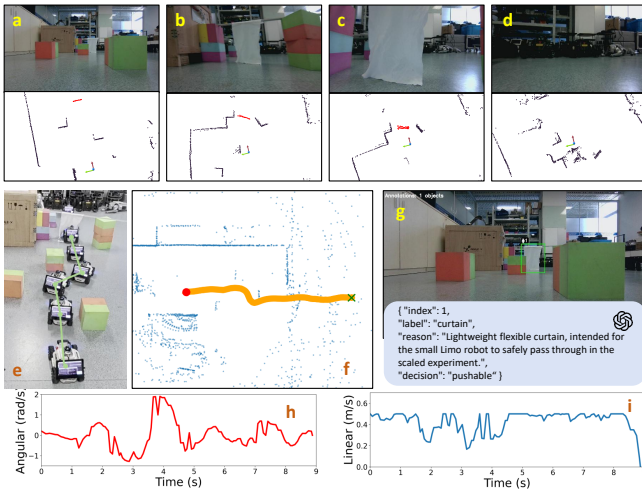


Fig. 11: Navigation with deformable obstacle.

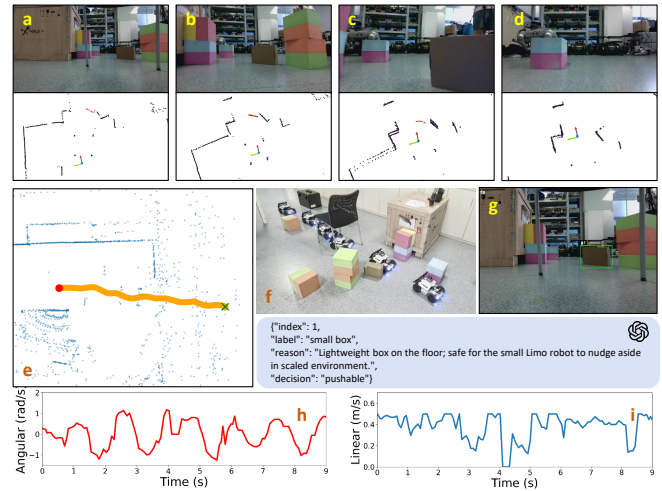


Fig. 12: Navigation with movable obstacle.

VI. CONCLUSION

This paper proposed DCT, a direct contact-tolerant planner for robot navigation among movable obstacles in cluttered environments. Through integration of VPP for contact-tolerant reasoning and VGN for contact-tolerant execution, DCT enables the robot to recall short-term memory, make informed contact decisions, act with high accuracy, and maintain efficiency under uncertainty. Extensive simulation and real-world evaluations confirm its robustness and superiority over representative baselines. Future work will extend DCT toward richer multi-modal reasoning and large-scale real-world deployments.

REFERENCES

- [1] X. Zhang, A. Liniger, and F. Borrelli, "Optimization-based collision avoidance," *IEEE Transactions on Control Systems Technology*, vol. 29, no. 3, pp. 972–983, 2020.
- [2] S. Kousik, B. Zhang, P. Zhao, and R. Vasudevan, "Safe, optimal, real-time trajectory planning with a parallel constrained bernstein algorithm," *IEEE Transactions on Robotics*, vol. 37, no. 3, pp. 815–830, 2020.
- [3] R. Han, S. Wang, S. Wang, Z. Zhang, Q. Zhang, Y. C. Eldar, Q. Hao, and J. Pan, "Rda: An accelerated collision free motion planner for autonomous navigation in cluttered environments," *IEEE Robotics and Automation Letters*, vol. 8, no. 3, pp. 1715–1722, 2023.
- [4] B. Frank, C. Stachniss, R. Schmedding, M. Teschner, and W. Burgard, "Real-world robot navigation amongst deformable obstacles," in *2009 IEEE International Conference on Robotics and Automation*, pp. 1649–1654, IEEE, 2009.
- [5] M. Stilman and J. J. Kuffner, "Navigation among movable obstacles: Real-time reasoning in complex environments," *International Journal of Humanoid Robotics*, vol. 2, no. 04, pp. 479–503, 2005.
- [6] K. Ellis, H. Zhang, D. Stoyanov, and D. Kanoulas, "Navigation among movable obstacles with object localization using photorealistic simulation," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 1711–1716, IEEE, 2022.
- [7] H.-C. Wang, S.-C. Huang, P.-J. Huang, K.-L. Wang, Y.-C. Teng, Y.-T. Ko, D. Jeon, and L.-C. Wu, "Curriculum reinforcement learning from avoiding collisions to navigating among movable obstacles in diverse environments," *IEEE Robotics and Automation Letters*, vol. 8, no. 5, pp. 2740–2747, 2023.
- [8] B. Frank, C. Stachniss, R. Schmedding, M. Teschner, and W. Burgard, "Learning object deformation models for robot motion planning," *Robotics and Autonomous Systems*, vol. 62, no. 8, pp. 1153–1174, 2014.
- [9] H. Wang, Q. Wang, F. Gao, and S. Shen, "Contact-aware motion planning among movable objects," *arXiv preprint arXiv:2502.03317*, 2025.
- [10] M. Stilman and J. J. Kuffner, *Planning Among Movable Obstacles with Artificial Constraints*, pp. 119–135. Berlin, Heidelberg: Springer Berlin Heidelberg, 2008.
- [11] M. Stilman and J. Kuffner, "Planning among movable obstacles with artificial constraints," *The International Journal of Robotics Research*, vol. 27, no. 11-12, pp. 1295–1307, 2008.
- [12] Y. Zhang and Y. Kantaros, "Namo-llm: Efficient navigation among movable obstacles with large language model guidance," *IEEE Robotics and Automation Letters*, vol. 10, no. 12, pp. 13026–13033, 2025.
- [13] T. Yang, J. Hwang, J. Jeong, M. Yoon, and S.-E. Yoon, "Efficient navigation among movable obstacles using a mobile manipulator via hierarchical policy learning," *arXiv preprint arXiv:2506.15380*, 2025.
- [14] L. Yao, V. Modugno, A. M. Delfaki, Y. Liu, D. Stoyanov, and D. Kanoulas, "Local path planning among pushable objects based on reinforcement learning," in *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 3062–3068, IEEE, 2024.
- [15] B. He, G. Chen, W. Wang, J. Zhang, C. Fermuller, and Y. Aloimonos, "Interactive-far: interactive, fast and adaptable routing for navigation among movable obstacles in complex unknown environments," in *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 5402–5409, 2024.
- [16] G. Wilfong, "Motion planning in the presence of movable obstacles," in *Proceedings of the Fourth Annual Symposium on Computational Geometry*, SCG '88, (New York, NY, USA), p. 279–288, Association for Computing Machinery, 1988.
- [17] C. Li, F. Xia, R. Martin-Martin, and S. Savarese, "Hrl4in: Hierarchical reinforcement learning for interactive navigation with mobile manipulators," in *Conference on Robot Learning*, pp. 603–616, PMLR, 2020.
- [18] S. Liu, Z. Zeng, T. Ren, F. Li, H. Zhang, J. Yang, Q. Jiang, C. Li, J. Yang, H. Su, et al., "Grounding dino with grounded pre-training for open-set object detection," in *European Conference on Computer Vision*, pp. 38–55, Springer, 2024.
- [19] R. Han, S. Wang, S. Wang, Z. Zhang, J. Chen, S. Lin, C. Li, C. Xu, Y. C. Eldar, Q. Hao, and J. Pan, "Neupan: Direct point robot navigation with end-to-end model-based learning," *IEEE Transactions on Robotics*, vol. 41, pp. 2804–2824, 2025.
- [20] Z. Zhang, Y. Zhang, R. Han, L. Zhang, and J. Pan, "A generalized continuous collision detection framework of polynomial trajectory for mobile robots in cluttered environments," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 9810–9817, 2022.
- [21] NVIDIA, "Nvidia isaac sim." <https://developer.nvidia.com/isaac/sim>. Version 4.5, robotics simulation platform.
- [22] OpenAI, "Gpt-5." <https://platform.openai.com>, 2025. Accessed via API, 2025.
- [23] Google, "Gemini." <https://ai.google.dev>, 2025. Accessed via API, 2025.
- [24] J. Bai et al., "Qwen-vl: A versatile vision-language model for understanding, localization, text reading, and beyond," *arXiv preprint arXiv:2308.12966*, 2023.
- [25] H. Touvron et al., "Llama: Open and efficient foundation language models," *arXiv preprint arXiv:2302.13971*, 2023.