Title:
**Risk-Aware Path Planning for Object Transport Operations with Pose Transitions Using Point Clouds**

Authors:
Linxuan He, karinken@163.com, Hokkaido University
Hiroaki Date, hdate@ist.hokudai.ac.jp, Hokkaido University

Keywords:
Laser-based measurement, Point Clouds, Path Planning, Offset

DOI: 10.14733/cadconfP.2026.14-19

Introduction:
In indoor industrial environments, transporting large equipment through narrow, obstacle-dense spaces often requires yaw changes of the transport unit, which in this study included both the object and the worker. Therefore, this planning problem cannot be reduced to translation-only or fixed-yaw motion. Even if the transport unit is collision-free at both the initial and target yaws, intermediate rotations may still result in collisions with the environment. Moreover, collision-free motion alone is insufficient for safe transport, as a feasible path may still provide inadequate clearance from surrounding obstacles.

Existing path-planning techniques have been widely studied in the mobile robot domain [1]. Malhan and Gupta [2] studied point-cloud acquisition and reconstruction for later computation. Niwa and Masuda [3] investigated point-cloud-based collision detection for large engineering plants. Nordström et al. [4] proposed a risk-aware path finder based on A* in a time-dependent multi-agent setting. However, [2] focuses on point-cloud acquisition and reconstruction, [3] on collision detection, and [4] on risk-aware search in a different planning context. None of these approaches jointly address transport path-planning in measured point-cloud environments with yaw-dependent transport-unit occupancy, explicit validation of intermediate yaw transitions, and clearance-aware path evaluation.

To address this problem, a risk-aware path-planning framework for object transport in measured point-cloud environments was proposed in this study. The main contributions are as follows: (1) a unified voxel representation of the environment and the transport unit; (2) a pose-dependent feasibility model that validates each yaw transition throughout the intermediate rotation; and (3) an extended A*-based risk-aware path-planning framework that integrates clearance-based evaluation, enabling the planner to distinguish between merely collision-free paths and those with sufficient safety margins.

Main Idea:
The planning problem addressed in this study is as follows. Given a measured point cloud of an indoor environment and a transport unit, we plan a collision-free transport path with feasible yaw transitions and sufficient clearance. First, the environment and the transport unit are voxelized into a unified representation. Next, pose-dependent configuration-space (C-space) obstacles are precomputed by inverse offset for each discretized yaw. A risk field is then constructed that combines collision feasibility and clearance-based safety margin. Finally, an extended A* search is performed with explicit validation of intermediate rotations during pose transitions. The resulting planner balances path length and clearance-based risk, triggering pose transitions only when necessary in constrained regions.

*Proposed Method*:
We present an object transport path planning workflow for large objects that takes a measured 3D point cloud as input and builds a quantitative risk field with a unified voxel representation for both the environment and the transport unit. Hereafter, the transported object and the worker are referred to as a single transport unit. Fig. 1 shows the workflow of the proposed method. After ground removal by plane fitting with random sample consensus (RANSAC) and voxelization, pose-dependent C-space obstacles are computed by applying an inverse offset to the transport unit voxels for each discretized yaw pose $\theta$, resulting in pose-dependent C-space obstacle voxels that allow collision checking through simple occupancy tests. On these voxels, a quantitative risk field is built, consisting of an occupancy risk term and a clearance-based risk term to model the safety margin. Finally, an extended A* search is performed in the state space $(\mathbf{v},\theta)$, using a single-step cost that combines path length and risk cost, producing a feasible path and pose sequence that minimize accumulated cost while satisfying passage constraints.

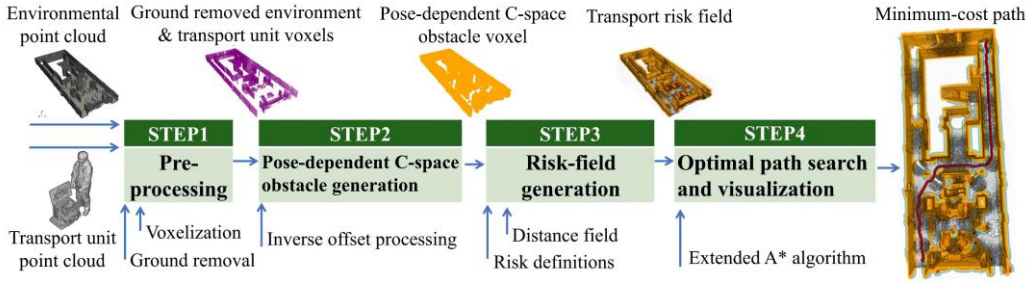


Fig. 1: Workflow of the proposed method.

*Generation of C-space Transport Unit Voxels:*
Let $V_E$ denote the set of occupied environment voxels after ground removal. Let $V_O^\theta$ denote the set of occupied voxels of the transport unit at a discretized pose $\theta$. The pose set is $\Theta = \{0^\circ, 10^\circ, 20^\circ, \ldots, 350^\circ\}$ (36 yaw angles in total). For each $\theta \in \Theta$, the configuration-space (C-space) obstacle voxel set for that pose is constructed using an inverse offset operation:

$$V_E^{(\theta,Offset)} = V_E \oplus (V_O^\theta)^{-1} \tag{1}$$

Here, $(V_O^\theta)^{-1}$ denotes the reflected (inverse) set of $V_O^\theta$, and $\oplus$ denotes the Minkowski sum. For any state $(\mathbf{v},\theta)$, if the position voxel $\mathbf{v}$ is included in $V_E^{(\theta,Offset)}$, the configuration is in collision. Thus, collision checking is reduced to an occupancy-membership query on the voxel grid.
*Object Transport Risk Field Construction*:
For each discrete pose $\theta$, we compute a voxel risk field $R^\theta(\mathbf{v})$ from the C-space obstacle set $V_E^{(\theta,Offset)}$ consisting of a binary collision indicator $R_c^\theta(\mathbf{v})$ and a clearance-based risk term $R_d^\theta(\mathbf{v})$.

The occupancy risk $R_c^\theta(\mathbf{v})$ encodes a hard feasibility constraint as a binary indicator. It is set to 1 if the voxel $\mathbf{v}$ lies inside $V_E^{(\theta,Offset)}$ (i.e., a colliding configuration), and 0 otherwise. This reduces collision checking to a lightweight voxel-occupancy membership query, avoiding expensive geometric interference tests. The clearance-based risk $R_d^\theta(\mathbf{v})$ quantifies clearance as a continuous safety margin. Let $d^\theta(\mathbf{v})$ denote the Euclidean distance from voxel $\mathbf{v}$ to the nearest occupied voxel. $R_d^\theta(\mathbf{v})$ is defined as follows:

$$R_d^\theta(\mathbf{v}) = max(0, 1 - d^\theta(\mathbf{v}) / d_{\max}). \tag{2}$$

This definition yields $R_d^\theta(\mathbf{v}) = 1$ at zero clearance (on the obstacle surface), decreases monotonically as clearance increases, and becomes 0 once the distance exceeds $d_{\max}$. We set $d_{\max} = 1$ m, rounded from the recommended human working-reach value of 1.07 m [5] for practical implementation. Overall, the risk field provides a unified input that combines hard collision constraints $R_c^\theta(\mathbf{v})$ with a soft clearance

penalty $R_d^{\theta}(\mathbf{v})$, which is later queried in the state space $(\mathbf{v},\theta)$ during risk-aware path search. The binary collision indicator is used as a hard feasibility constraint, whereas only the clearance-based risk term is used in the search cost.

*Optimal Path Search*:

To satisfy both collision constraints and pose transition constraints, an extended A* search is performed in the discrete state space $x=(\mathbf{v},\theta)$, where $\mathbf{v}=(i,j,k_{fix})$ is the voxel coordinate and $\theta$ is the discretized pose of the transport unit. In this study, $k_{fix}$ is a fixed height, so translational moves are restricted to the $xy$ plane. Each step uses an action set consisting of eight translation moves (8-neighborhood in the $xy$ plane, with $\theta$ unchanged) and two pose transitions ($\pm 10^{\circ}$, with $\mathbf{v}$ unchanged). For a candidate successor state $x'$, the single-step cost uses only the clearance-based risk term and is defined as

$$\Delta g(x,x') = w_l s\sqrt{(\Delta i)^2+(\Delta j)^2} + w_d R_d^{\theta'}(\mathbf{v}') \tag{3}$$

where $w_l$ and $w_d$ are weighting factors. Here, if $\mathbf{v}=(i,j,k_{fix})$ and $\mathbf{v}'=(i',j',k_{fix})$, then $\Delta i = i'-i, \Delta j = j'-j$, where $s$ is the voxel size. To keep the branching factor bounded while maintaining safety, we adapt the action set based on the clearance-based risk evaluated at the candidate next state $x'$. If the clearance-based risk at the candidate next voxel $\mathbf{v}'$, evaluated under the discretized yaw aligned with the translational direction from $\mathbf{v}$ to $\mathbf{v}'$, is zero, the successor is accepted and the next yaw is set to that direction-aligned yaw without additional local pose search. If the clearance-based risk at the candidate next voxel $\mathbf{v}'$, evaluated under the discretized yaw aligned with the translational direction from $\mathbf{v}$ to $\mathbf{v}'$, is greater than zero, local pose transitions are additionally allowed so the transport unit can adjust its yaw when passing through narrow sections.

When pose transitions are enabled, the next yaw is selected by choosing a feasible candidate that reduces the clearance-based risk at the next position. A feasible candidate yaw remains collision-free throughout the entire rotation and preserves next-step feasibility in the search. Thus, the candidate set is determined by collision and rotation-feasibility constraints. The intended translational direction is defined as the planar translation direction from the current voxel to the candidate next voxel before any local yaw adjustment is applied. Among feasible candidates, directional consistency with this intended translational direction serves as a secondary tie-breaking criterion.

For each candidate pose transition, the yaw change is discretized into intermediate yaws. For every intermediate yaw, occupancy is checked at both the current voxel and the candidate next voxel using occupancy-based collision tests on the voxel grid. The transition is accepted only if all intermediate poses remain collision-free at both positions. To reduce repeated checks during A* expansion, the feasibility result of each evaluated transition is stored and reused when the same transition is encountered again.

Results:

The proposed method was evaluated in a real scanned indoor environment measured with a Leica BLK360 G2 (Environment 1; 26.6 m × 10.2 m × 1.9 m) and a synthetic narrow-passage environment modeled in Unreal Engine (Environment 2; 50.0 m × 50.0 m × 2.1 m). The transport unit, consisting of the transported object and the worker, was digitized by smartphone-based LiDAR scanning using the Scaniverse app and then manually denoised; its overall dimensions are 1.83 m × 1.21 m × 0.63 m.

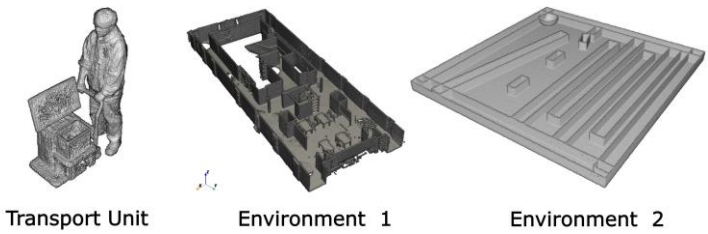


Fig. 2: Point-cloud datasets used for evaluation.

Fig. 2 shows the transport unit used in the experiments and the point cloud datasets of the two evaluated environments. Tab. 1 lists point cloud inputs, voxel statistics, and computation times. All experiments were conducted on a 64-bit Windows workstation equipped with an Intel Core i9-14900K CPU and an NVIDIA GeForce RTX 4090 GPU. The implementation was developed in Python 3.11. In both evaluated environments, inverse offset dominates the computation, taking 22.5 s in Environment 1 and 346.7 s in Environment 2, while path computation takes 12.8 s and 192.8 s, respectively. The total computation times are 49.8 s and 606.4 s, respectively.

| **Item** | | **Env.1** | **Env.2** |
|---|---|---|---|
| Number of points in environment point cloud $P_E$ | | 1,771,669 | 358,785 |
| Number of environment voxels $V_E$ | | 133,049 | 110,044 |
| Total number of voxels in all C-space obstacle sets $V_E^{(\theta,Offset)}$ | | 199,770,444 | 964,309,388 |
| Compu-tation time | Voxelization time [s] | 6.8 | 19.5 |
| | Inverse offset processing time [s] | 22.5 | 346.7 |
| | Risk-field construction time [s] | 7.8 | 47.3 |
| | Path computation time [s] | 12.8 | 192.8 |
| | Total computation time [s] | 49.8 | 606.4 |

Tab. 1: Point-cloud inputs, voxel statistics, and computation times for Environments 1 and 2.

The high-risk rate is defined as the percentage of sampled path states with a clearance-based risk value exceeding the threshold of 0.75. This threshold corresponds to approximately 25 cm clearance in the real world, indicating a limited safety margin during transport. Therefore, a lower high-risk rate means fewer path states have limited clearance for the worker, indicating less exposure to potentially unsafe segments during transport. To isolate the effects of the clearance-based risk term and pose transitions, we compare three cases, as summarized in Tab. 2. Case 1 (C1) uses neither the clearance-based risk term nor pose transitions. Case 2 (C2) uses the clearance-based risk term but does not allow pose transitions. Case 3 (C3) uses both the clearance-based risk term and pose transitions and corresponds to the proposed method. Fig. 3 visualizes the resulting transport paths in the two environments under C1–C3, together with pose transitions where they occur.

| | | **Env.1** | | | **Env.2** | | |
|---|---|---|---|---|---|---|---|
| | | C1 | C2 | C3 | C1 | C2 | C3 |
| Case setting | Clearance-based risk term | No | Yes | Yes | No | Yes | Yes |
| | Pose transitions | No | No | Yes | No | No | Yes |
| Performance metrics | Path Length [m] | 33.45 | 36.83 | 35.29 | 228.0 | 241.7 | 238.1 |
| | Minimum Clearance [m] | 0.00 | 0.14 | 0.35 | 0.00 | 0.49 | 0.49 |
| | Average Clearance [m] | 0.30 | 0.53 | 0.57 | 0.30 | 0.74 | 0.73 |
| | High-risk Rate [%] | 53.6 | 4.3 | 0.0 | 64.3 | 0.0 | 0.0 |
| | Total Cost | 173.6 | 140.6 | 125.2 | 1245.2 | 689.3 | 662.8 |

Tab. 2: Quantitative comparison of the three cases in Environments 1 and 2.

Tab. 2 shows a consistent trend across Environments 1 and 2. Compared with C1, introducing the clearance-based risk term in C2 substantially reduces the high-risk rate, from 53.6% to 4.3% in Environment 1 and from 64.3% to 0.0% in Environment 2. At the same time, the minimum clearance increases from 0.00 m to 0.14 m in Environment 1 and from 0.00 m to 0.49 m in Environment 2, while the average clearance increases from 0.30 m to 0.53 m and 0.74 m, respectively. However, this also increases path length by 10.1% and 6.0% because the fixed-pose planner must detour to maintain safety margins.

Allowing pose transitions in C3 preserves these safety benefits while partially recovering path efficiency. In both Environments 1 and 2, the path length is reduced compared with C2 (from 36.83 m to 35.29 m and from 241.7 m to 238.1 m, respectively), while maintaining low or zero high-risk rates (0.0% in both environments) and comparable clearance (minimum clearance of 0.35 m and 0.49 m, respectively). These results indicate that local pose adjustment enables the transport unit to pass through narrow regions with less detour.

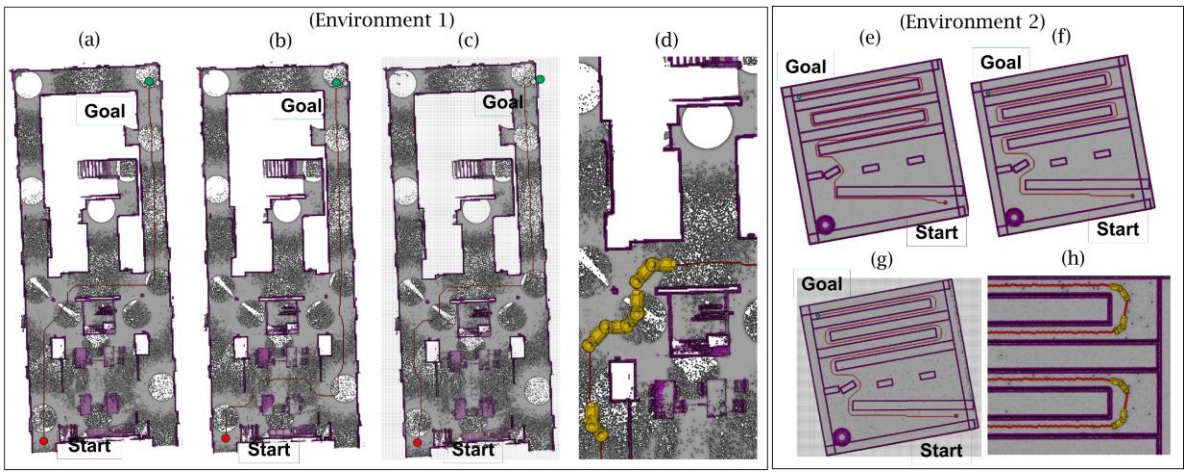


Fig. 3: Path planning results under different conditions.

Conclusions:

In this paper, a risk-aware path-planning method is proposed for object transport that combines clearance-based risk evaluation with an extended A* search in point-cloud environments. The method targets the transport unit, which includes both the transported object and the worker. Using a voxelized representation of the environment, the method jointly considers obstacle-clearance risk, pose transitions during transport, and path cost, thereby achieving joint planning of the transport path and pose sequence.

In summary, the proposed method enables path-planning for object transport in complex and narrow environments while jointly considering feasibility and safety margins. The results show that integrating pose-dependent C-space obstacles, a clearance-based risk field, and extended A* search within a unified framework is effective in the evaluated environments. Quantitatively, the proposed method reduces the high-risk rate from 53.6% to 0.0% in Environment 1 and from 64.3% to 0.0% in Environment 2. Compared with the fixed-pose risk-aware baseline C2, it also shortens the path from 36.83 m to 35.29 m in Environment 1 and from 241.7 m to 238.1 m in Environment 2. This shows that the proposed method not only reduces the proportion of path segments with insufficient safety margins, but also avoids the extra detour required by a fixed-pose risk-aware planner. These results demonstrate that explicit handling of yaw change and clearance is effective for transport planning in measured point-cloud environments.

Acknowledgements:
This work was supported by JST SPRING, Grant Number JPMJSP2119.

*Linxuan He*, https://orcid.org/0009-0009-3599-6913
*Hiroaki Date*, https://orcid.org/0000-0002-6189-2044

References:

[1] Liu, L.; Wang, X.; Yang, X.; Liu, H.; Li, J.; Wang, P.: Path planning techniques for mobile robots: Review and prospect, Expert Systems with Applications, 227, 2023, 120254. https://doi.org/10.1016/j.eswa.2023.120254

[2] Malhan, R. K.; Gupta, S. K.: Planning Algorithms for Acquiring High Fidelity Point clouds Using a Robot for Accurate and Fast 3D Reconstruction, Robotics and Computer-Integrated Manufacturing, 78, 2022, 102372. https://doi.org/10.1016/j.rcim.2022.102372

[3] Niwa, T.; Masuda, H. Interactive collision detection for engineering plants based on point-clouds. *Computer-Aided Design & Applications*, 13(4), 2016, 511–518. https://doi.org/10.1080/16864360.2015.1131546

[4] Nordström, S.; Bai, Y.; Lindqvist, B.; Nikolakopoulos, G.: A Time-dependent Risk-aware distributed Multi-Agent Path Finder based on A*, arXiv preprint, 2025, arXiv:2504.19593. https://doi.org/10.48550/arXiv.2504.19593

[5] U.S. Department of Energy: DOE Handbook: Human Factors/Ergonomics Handbook for the Design for Ease of Maintenance (DOE-HDBK-1140-2001), U.S. Department of Energy, Washington, DC, 2001. https://www.standards.doe.gov/standards-documents/1100/1140-bhdbk-2001-pt1