

NAMOUnc: Navigation Among Movable Obstacles with Decision Making on Uncertainty Interval

Kai Zhang^{1,2,3}^a, Eric Lucet¹^b, Julien Alexandre Dit Sandretto²^c,
Shoubin Chen^{3,*}^d and David Filliat²^e

¹Paris-Saclay University, CEA, List, 91120 Palaiseau, France

²U2IS, ENSTA Paris, Institut Polytechnique de Paris, 91120 Palaiseau, France

³Guangdong Laboratory of Artificial Intelligence and Digital Economy(SZ), 51800 Shenzhen, China

Keywords: Navigation Among Movable Obstacles, Planning Under Uncertainty, Decision Making.

Abstract: Navigation among movable obstacles (NAMO) is a critical task in robotics, often challenged by real-world uncertainties such as observation noise, model approximations, action failures, and partial observability. Existing solutions frequently assume ideal conditions, leading to suboptimal or risky decisions. This paper introduces NAMOUnc, a novel framework designed to address these uncertainties by integrating them into the decision-making process. We first estimate them and compare the corresponding time cost intervals for removing and bypassing obstacles, optimizing both the success rate and time efficiency, ensuring safer and more efficient navigation. We validate our method through extensive simulations and real-world experiments, demonstrating significant improvements over existing NAMO frameworks. More details can be found in our website: <https://kai-zhang-er.github.io/namo-uncertainty/>.

1 INTRODUCTION

In real applications, robots take actions with partial and noisy observation on the environment, and a given action may cause unexpected effect due to uncertainties. If the robot has over confidence on the observation and action, it can lead to some suboptimal decisions, even to dangerous actions resulting in catastrophic effects, like destroying objects in the workspace. It is therefore essential for the robot to recognize the limitation of its observations and actions, and make decisions with awareness of uncertainty and risks.

Recent studies on task and motion planning incorporate uncertainty and update plans based on the observation and action uncertainties. For example (Safronov et al., 2020; Pan et al., 2022) take success rate (SR) into account in a manipulation task: if a grasp action fails, the planner updates its estima-

tion on SR, then replans for an action sequence with a higher predicted SR. Methods such as probabilistic symbolic planning (Silver et al., 2021) or Bayes optimization (Curtis et al., 2024) plan with uncertainty but mainly focus on optimizing SR. They often neglect joint optimization with efficiency, which is crucial for navigation tasks.

Navigation among movable obstacles (NAMO) task is mainly a navigation task but the robot is able to manipulate movable obstacles (MO). Many existing solutions (Muguira-Iturralde et al., 2023; Ellis et al., 2022) consider NAMO as a manipulation task, assuming manipulation is necessary to complete the task. However, the most common case is that the task can be finished without moving the MO, by making a detour. This oversight makes probability-based optimization less useful, as bypassing with a SR close to 1 will be preferred even if it requires significantly more time. Furthermore, in partial observation condition, the invisible region in NAMO tasks is considerably larger than in manipulation tasks. Common strategies for reducing uncertainty in manipulation tasks, such as exploring all occluded regions (Curtis et al., 2024), are often inefficient or impractical for NAMO tasks due to the scene complexity and large scale.

^a <https://orcid.org/0000-0003-1129-9944>

^b <https://orcid.org/0000-0002-9702-3473>

^c <https://orcid.org/0000-0002-6185-2480>

^d <https://orcid.org/0000-0002-9071-0051>

^e <https://orcid.org/0000-0002-5739-1618>

*Corresponding author.




Figure 1: Primary sources of uncertainty in a NAMO task. The illustration highlights four key uncertainties encountered during task execution. The blue rectangle represents the MO while the green triangle with R denotes the robot. The red square labeled G indicates the goal. In sub-figure (c), X denotes the observation while y and σ represent the prediction result and the corresponding prediction uncertainty, respectively.

This paper presents NAMOUnc, a method for solving NAMO tasks by optimizing both SR and goal reaching time in real scenarios. Our main contribution lies in enhancing the robustness of the NAMO method with respect to the following uncertainties (Fig. 1): (a) observation uncertainty on MO pose estimation caused by sensor noise, (b) model approximation uncertainty, (c) action uncertainty from imperfect controllers, and (d) blockage uncertainty from partial observation. NAMOUnc estimates these uncertainties as cost intervals and makes decisions based on their utility values, balancing removal and bypass strategies to achieve efficient and successful navigation.

In summary, our contribution includes:

1. A novel approach for solving NAMO tasks, optimizing SR and efficiency under condition of partial observability.
2. Four modules to systematically estimate and quantify the uncertainties described in Fig. 1.
3. A novel method to estimate the uncertainty caused by partial observation in unexplored region, which can effectively reduce the navigation risk and improve the efficiency.
4. Experiments in simulated and real environments to demonstrate the effectiveness of our method.

2 RELATED WORK

2.1 Navigation Among Movable Obstacles

The NAMO task has been extensively studied, with recent advancements focusing on end-to-end learning (Li et al., 2020) and hybrid approaches (Muguira-Iturralde et al., 2023; Kim et al., 2019; Xia et al., 2021). End-to-end learning methods typically employ hierarchical reinforcement learning (Li et al., 2020), generating high-level subgoals alongside low-level control parameters. Hybrid methods, on the other hand, leverage machine learning either to produce subgoals (Xia et al., 2021) or to assist in generating action sequences (Muguira-Iturralde et al., 2023). A comprehensive review of these techniques is provided in (Zhang et al., 2022).

Existing methods primarily focus on task completion without considering the associated costs. A related work (Zhang et al., 2023) introduced a strategy selection mechanism based on estimated costs, jointly optimizing SR and efficiency. However, this approach does not incorporate uncertainty, which limits its generalizability and applicability in real-world scenarios. Therefore, we propose NAMOUnc method in this paper as an extension to deal with real-world cases.

2.2 Planning with Uncertainty

Uncertainty in real-world environments introduces significant challenges, prompting the development of methods to address various types of uncertainties, including observation uncertainty, action uncertainty, and uncertainty arising from partial observability.

Observation uncertainty in task planning pertains to the confidence in object classification and pose estimation. Object classification uncertainty has been extensively studied in computer vision (Feng et al., 2021), with confidence scores commonly used to quantify this uncertainty. Pose uncertainty, often resulting from sensor noise, is often modeled as a Gaussian distribution and mitigated through repeated observations (Kaelbling and Lozano-Pérez, 2013; Garrett et al., 2020).

Action uncertainty has garnered increasing attention. In sampling-based methods, action uncertainty is represented as a probabilistic transition matrix, which can be approximated through frequent replanning (Garrett et al., 2020) or learned from demonstrations and datasets (Silver et al., 2021). This matrix enables the selection and execution of the most likely successful plan. Rather than treating it as an open-loop process, some approaches (Pan et al., 2022) it-




Figure 2: Overview of the NAMO method pipeline. When being blocked by MOs during navigation, the robot estimates the bypass and removal cost to choose an efficient strategy to continue its task. The green cloud symbols represent the uncertainties associated with each module: (a) Observation uncertainty; (b) Model uncertainty; (c) Action uncertainty; (d) Blockage uncertainty caused by partial observability.

eratively update transition probabilities based on observed action effects, dynamically adjusting plans to achieve task objectives.

Regarding *uncertainty from partial observation*, most methods (Muguira-Iturralde et al., 2023; Curtis et al., 2024; Wu et al., 2010) restrict actions to the sensor’s field of view. In manipulation tasks, studies such as (Curtis et al., 2024; Garrett et al., 2020) model the probability of unseen objects behind visible ones, incorporating this into planning to trigger exploration actions when necessary. For NAMO tasks, RNAMO (Wu et al., 2010) proposes pushing all obstacles in invisible regions; if obstacles are immovable, the robot updates its internal map and plans a detour. In larger workspaces, LaMB (Muguira-Iturralde et al., 2023) employs backward reasoning to eliminate environmental invisibility, though this approach is limited to small environments due to its exhaustive exploration of invisible regions. To address this limitation, we introduce a method to quantify the potential cost of navigating through unexplored areas, optimizing the trade-off between exploration (bypassing into unknown regions) and exploitation (manipulating visible obstacles).

3 NAVIGATION AMONG MOVABLE OBSTACLES WITH UNCERTAINTY

In a NAMO task, shown in Fig. 2, a robot needs to navigate to a goal while avoiding obstacles. With the environment map, the first step is to plan and follow the shortest path. If meeting a MO blocking this path, the robot can choose to bypass it or clear the path by removing it. As described in our previous work (Zhang et al., 2023), we first estimates the by-

pass and removal cost before making decision. The bypass cost is calculated based on a detour trajectory while the removal cost is computed in two steps: predicting the stock region for the MO, then estimating the time of moving the MO to this region. The motion planner outputs control parameters for the robot to follow the best alternative.

Uncertainties across different system components can significantly affect task success. In the context of MO detection and localization, observation-related uncertainties, such as recognition errors and pose estimation inaccuracies, can lead to collisions or unsuccessful removal attempts. Similarly, model uncertainties in estimating the costs of bypassing or removing obstacles may yield inaccurate evaluations. If the robot opts to remove a MO, the action may fail due to incomplete knowledge or discrepancies between the predicted and actual outcomes of the action. Furthermore, operating in partially observable environments introduces intrinsic uncertainty, which often results in a divergence between anticipated and actual conditions.

To quantify the impact of these uncertainties on the decision-making process, we use time intervals, denoted as $[]$, which reflect the range of potential outcomes generated by uncertainty, before selecting the best option based on their utility values. In the following section, we present the methods of quantifying these uncertainties using intervals.

4 UNCERTAINTY ESTIMATION METHOD

We now present four uncertainties considered in NAMO tasks and the methods to estimate them, before detailing the decision process.

4.1 Observation Uncertainty

While navigating the environment, the robot should localize MOs and detect whether the planned path is blocked. Due to the sensors' noise in object detection and localization error of the robot, the estimated MO pose is subject to uncertainty. However, multiple observations can refine the result, and given a specified confidence interval, a belief region representing the MO's position can be calculated for path blockage determination.

The obstacle pose is computed from the robot pose plus the relative pose of the MO given by the sensor. Assuming the robot pose from the localization algorithm is $X_r = (x_r, y_r, \theta_r)$ with covariance Σ_r , and the relative distance and angle of the i -th MO measured by the depth camera $Y^i = (d^i, \phi^i)$ with covariance Σ_Y , the obstacle pose X_{MO}^i can be obtained by: $X_{MO}^i = (x_r + d^i \cos(\theta_r + \phi^i), y_r + d^i \sin(\theta_r + \phi^i))^T$ and the covariance matrix Σ_{MO}^i by:

$$\Sigma_{MO}^i = J_r \Sigma_r J_r^T + J_y \Sigma_Y J_y^T; \quad J_r = \frac{\partial X_{MO}^i}{\partial X_r}; \quad J_y = \frac{\partial X_{MO}^i}{\partial Y^i} \quad (1)$$

When multiple observations of the same MO are received, a Kalman filter (Kalman, 1960) is applied to fuse the repeated observations and obtain the estimated MO pose and its covariance.

Given a confidence score T_{conf} , set to 95% in our experiments, though other values are possible, the belief region of the MO pose is represented as an ellipse computed from the covariance matrix. We add the size of the MO, expressed as its radius, to obtain the ellipse region where a path would lead to potential collision. In this case, the robot stops and chooses the best strategy as described in Sec. 4.5.

4.2 Bypass Cost Model Uncertainty

The bypass cost, defined as the estimated travel time to reach the goal when the robot follows a planned detour, varies depending on the trajectory and the robot's moving speed. A deterministic predictor of navigation time inherently involves uncertainty, since in practice the robot may deviate from the planned trajectory. To capture this prediction uncertainty, we express the estimated time as an interval rather than a single value.

To plan the detour, a MO and its uncertainty region (an ellipse with 95% confidence as described in Sec. 4.1) is temporarily added to the obstacle map, then the shortest path planner searches a path to bypass the MO. If no path is found, the bypass time cost is set to $C_{nav}^{by} = Inf$. On the contrary, if a path is found,

we need to estimate the navigation time from trajectory features. While the average speed was simply used in (Zhang et al., 2023), we use a Gaussian linear regressor (GLR) (Williams and Rasmussen, 1995) for better prediction and uncertainty evaluation (see Sec. 5.3).

In practice, as rotation takes more time than following straight lines, we need to take the orientation change of the trajectory into account during bypass time estimation. In addition to the trajectory length F_l , we therefore calculate the trajectory smoothness F_s and variance of direction change F_v as features to estimate the navigation time. Assuming a trajectory consists of N waypoints $pt_i, i = 0, 1, \dots, N$, each pt_i characterized by position p_i and orientation α_i , smoothness and variance are calculated using:

$$F_s = \frac{\sum_{i=1}^N |\alpha_i - \alpha_{i-1}|}{N-1}; F_v = \text{var}(|\alpha_i - \alpha_{i-1}|) \quad (2)$$

A trajectory is therefore characterized by $X = \{F_l, F_s, F_v\}$. Then, the bypass time cost and variance are predicted by a GLR:

$$T_{by}, \sigma_{by} = \text{GLR}(X) \quad (3)$$

The navigation time interval with $T_{conf} = 95\%$ confidence is: $[C_{nav}^{by}] = [T_{by} - 2\sigma_{by}, T_{by} + 2\sigma_{by}]$.

To train the regression model, we collect a set of trajectories by controlling the robot to navigate in a warehouse environment. The pose and time stamp are recorded along these trajectories, and we create a large and varied dataset for the model by sampling random start and goal points in these trajectories and computing the corresponding features and duration.

4.3 Removal Action Uncertainty

The action uncertainty relates to the uncertain outcome of loading a MO for displacement, which can be either success or failure.

For a given SR, p_a , the expected removal cost C_{MO} is

$$C_{MO} = T_{MO} \sum_{i=1}^M i p_a (1 - p_a)^{i-1} + (M T_{MO} + C_{by}) (1 - p_a)^M \quad (4)$$

where M defines the maximum attempts when the robot continuously fails to load a MO (according to the psychological view on learned helplessness (Maier and Seligman, 1976)). C_{by} denotes the obstacle bypass cost, and T_{MO} is the removal cost of a MO that can be estimated by the method proposed in (Zhang et al., 2023), where a stock region predictor and a regression model are employed to predict the placing pose and the removal time.

In practice, the SR may not be fixed and can evolve during robot operation, we therefore start by an initial estimate and update it after each trial. Similar to (Curtis et al., 2024), we model the SR of an action and the uncertainty of its estimation after t trials, p_a^t , using a Beta distribution. To obtain the initial SR, p_a^0 , we control the robot to load the MO in several trials (10 in our experiments) and record the action results. Assuming there are α successful trials and β failure cases, the initial knowledge on the trial results can be described as $p_a^0 \sim \text{Beta}(\alpha, \beta)$. During operation, when new trials are performed, the updated posterior is $p_a^t \sim \text{Beta}(\alpha + s, \beta + f)$ where s and f are the number of successful and failed object manipulation respectively. With a confidence score T_{conf} , we can obtain the SR interval:

$$p_a^t \in [\text{Beta.ppf}(\frac{1 - T_{conf}}{2}), \text{Beta.ppf}(\frac{1 + T_{conf}}{2})] \quad (5)$$

where the *Beta.ppf* is the point percent function to obtain the confidence interval of the beta distribution given a confidence score.

We compute the interval of the expected removal cost $[C_{MO}]$ by using Eq. 4 with the minimum and maximum p_a^t of the SR interval.

4.4 Blockage Uncertainty

Blockage uncertainty comes from the partial observation condition and results in a blocking probability of the robot by some unseen MOs in unexplored region. It is related to the passage width and robot size as it is more risky if the planned navigation path passes a narrow passage rather than an open space.




Figure 3: Blocking case. The blue circle represents the MO in a corridor with width W_i . The green dash curve is the planned trajectory of the robot while the red dash line is the traversal line at waypoint pt_i .

To calculate the blockage probability of a trajectory T , we take T as a set of way points $pt_i, i = 1, 2, \dots, N$ and compute the blockage probability of each way point.

We model the blockage probability based on several assumptions (shown in Fig. 3):

- (i) The width of the passage W_i at pt_i ,
- (ii) The radius of the robot r ,
- (iii) The diameter of the MO, l_{MO} , modeled as a Gaussian distribution $G(\mu, \sigma)$ where μ is the average radius of the MOs and σ is the standard deviation.
- (iv) The distance between the MO center and the wall d , modeled as a uniform distribution $d \sim U(\frac{l_{MO}}{2}, W_i - \frac{l_{MO}}{2})$, which leads to $p(d|l_{MO}, W_i) = \frac{1}{W_i - l_{MO}}$.

Given d, l_{MO}, r and W_i , the blockage probability is deterministic: $p(b|d, l_{MO}, W_i, r) = 1$ when the size of robot is larger than the space of both sides of the MO, represented by $2r > \max(d - \frac{l_{MO}}{2}, W_i - d - \frac{l_{MO}}{2})$; otherwise, $p(b|d, l_{MO}, W_i, r) = 0$.

To obtain the blockage probability conditioned only on the passage width and the robot size $p(b|W_i, r)$, we first eliminate the dependence on d through marginalization:

$$p(b|l_{MO}, W_i, r) = \int_{l_{MO}/2}^{W_i - l_{MO}/2} p(b|d, l_{MO}, W_i, r) p(d|l_{MO}, W_i, r) dd \quad (6)$$

After simplification, we get:

$$p(b|l_{MO}, W_i, r) = \begin{cases} 0, & l_{MO} < W_i - 4r \\ \frac{4r}{W_i - l_{MO}} - 1, & W_i - 4r < l_{MO} < W_i - 2r \\ 1, & W_i - 2r < l_{MO} < W_i \\ 0, & W_i < l_{MO} \end{cases} \quad (7)$$

The blockage probability in a corridor is then:

$$p(b|W_i, r) = \int p(b|l_{MO}, W_i, r) p(l_{MO}) dl_{MO} \quad (8)$$

Considering that l_{MO} satisfies a Gaussian distribution and $p(b|l_{MO}, W_i, r)$ is piecewise constant, we approximate this integral by using a sampling method.

The previous $p(b|W_i, r)$ is calculated with the assumption of a MO being on the line perpendicular to the corridor passing through pt_i (the red dash line in Fig. 3). Because W_i can be obtained from pt_i using a ray casting algorithm, and both are independent of r , the blockage probability at pt_i can be expressed as $p(b|pt_i, r) = p(b|W_i, r)$.

Assuming the MO is uniformly distributed in the space with a free area A , the probability that the MO is in the traversal line at pt_i is $p(pt_i) = \frac{W_i K}{A}$. Here K is a parameter characterising the obstacle appearance probability. It's worth noting that the assumption on uniform distribution is based on the absence of prior

knowledge about the object organization in the environment. However, if prior information is available, the distribution should be adjusted accordingly.

Therefore, the probability that the robot is blocked at pt_i is $p(b|pt_i, r) \times p(pt_i)$. Then, for a trajectory T , the blockage probability $p(b|T, r)$ can be computed by:

$$p(b|T, r) = 1 - \prod_{pt_i}^T (1 - p(b|pt_i, r) \times p(pt_i)) \quad (9)$$

The estimated cost of blockage when passing the invisible region is then :

$$[C_{blocked}] = p(b|T, r) \times [C_{MO}] \quad (10)$$

4.5 Decision Making with Uncertainty Interval

The decision-making module aims to compare the costs of bypassing and removal, then to choose the one with the lower cost.

With the uncertainties considered, the final cost of each option can be calculated by

$$[C_{by}] = [C_{nav}^{by}] + [C_{blocked}^{by}] \quad (11)$$

$$[C_{re}] = [C_{MO}] + [C_{nav}^{re}] + [C_{blocked}^{re}] \quad (12)$$

where $[C_{blocked}^{by}]$ and $[C_{blocked}^{re}]$ are the blockage costs of bypass and removal trajectories.

For the decision making between the cost intervals, we apply the Laplace criterion, described in (Denoeux, 2019), to compute the average utility of the consequences of each option. Assuming the cost satisfies the uniform distribution, the utility function can be expressed as:

$$U = \int_{\min([C])}^{\max([C])} xp(x) dx = \frac{\max([C]) + \min([C])}{2} \quad (13)$$

where $[C]$ is either $[C_{re}]$ or $[C_{by}]$ to calculate the utility while x and $p(x)$ are samples in the cost interval and its probability. Finally, the option with a smaller U is chosen as the navigation strategy.

5 SIMULATION EXPERIMENTS

We first conduct individual modules evaluation in simulation, and then compare our method with the state of the art before demonstrating a real robot application.

5.1 Simulation Environment

We implement our method in two simulated environments, a simple room and a large warehouse, as shown in Fig. 4. The room environment is built on PyBullet (Coumans and Bai, 2021) while the warehouse is based on Gazebo (Koenig and Howard, 2004). There is one MO in the room while multiple MOs are randomly put in the blue regions in the warehouse to simulate the variety of MO positions. A wheeled mobile robot with an arm needs to complete navigation tasks to reach the red goal. A LiDAR and a stereo camera are used to localize the robot and detect MOs respectively.




Figure 4: Simulation environments. Two environments are used including a simple room and a complex warehouse. The blue regions are possible places for the MOs and the red region is the goal.

5.2 Implementation Details

We use ROS Noetic (Stanford Artificial Intelligence Laboratory et al., 2018) with Movebase as the navigation framework. The environment map (including only static obstacles) is generated using GMapping (Grisetti et al., 2007) and provided to the robot as prior knowledge. We use A* (Hart et al., 1968) as the global path planner and TEB (Rösmann et al., 2015) as the local path planner. The AMCL package is employed to localize the robot from LiDAR data. The robot detects MOs by Aruco Marker (Romero-Ramirez et al., 2018) to reduce the classification uncertainty.

All the experiments are implemented in Python

with PyTorch (Paszke et al., 2017). We use an Intel i7-12700H CPU with 16G memory for the quantitative results.

5.3 Bypass Time Regression Results

To demonstrate the improved bypass time prediction, we compare the GLR regression model with the average speed method from (Zhang et al., 2023) and a trapezoid method that considers acceleration and deceleration. We collect a dataset manually by controlling the robot using a joystick and recording the planned path and corresponding time, including 1500 trajectory segments as training and 600 as testing. The average speed method calculates the speed in the training dataset, then applies it in the test set. For the GLR method, the model is fitted on the training set to predict on the test set. We calculate the absolute error between the predicted and actual navigation time. The results in Fig. 5 show that the applied GLR method outputs more accurate estimation with the lowest median absolute error (1.59s), compared to the average speed (3.43s) and trapezoid methods (3.31s). Additionally, the GLR method is more stable, with an interquartile range (IQR) of 0.69s, versus 1.35s and 1.91s for the other methods. In the case study, we found that the large error in the average speed method arises from its neglect of acceleration and deceleration during navigation. In contrast, the trapezoid method accounts for the time spent in speed changes and therefore yields smaller errors than the average speed method. However, for paths with multiple angle changes, the trapezoid method’s initial assumption, allowing only a fixed number of accelerations and decelerations, limits its adaptability. By comparison, our GLR method incorporates trajectory smoothness, enabling it to capture the key factors influencing navigation time and thus achieve the lowest prediction error.

5.4 Action Uncertainty Module Evaluation

To evaluate the effectiveness of modeling action uncertainty, we compare the task completion time with (w/) and without (w/o) the action uncertainty module in two cases: one with easy MOs (90% loading SR) and one with hard MOs (20% SR). We test the methods in three setups: ABC, AB, BC. ABC (resp. AB and BC) indicates three MOs are in regions A, B and C (resp. 2 at AB and BC) in Fig. 4.

The results in Fig. 6 show that with easy MOs (SR=90%), the w/o method takes less time in planning and making decision as it removes the MO at B




Figure 5: Boxplot of the absolute error of prediction results for the three bypass time prediction methods. The box represents the interquartile range (IQR), with the lower and upper edges indicating the 25th (Q1) and 75th (Q3) percentiles, respectively. The notch and orange line in the box marks the median value of the absolute error. Whiskers extend to the smallest and largest values within 1.5 times the IQR from Q1 and Q3.




Figure 6: Task completion time comparison on methods with/without considering action uncertainty in three environments, ABC, AB, BC. The top figure shows the case with easy MO (high SR) while the bottom one shows the hard MO (low SR).

directly. Conversely, the method w/ has the robot bypass B first to check MOs in A or C. If A and C are also blocked, the robot returns and removes the MO in B, requiring more planning and navigation time.

When MOs are hard to manipulate (SR=20%), the method w/ bypasses all MOs due to the high potential cost of the removal action, leading to faster navigation compared to w/o method that repeatedly attempts to remove MOs regardless of the cost. As for the stability, the w/ method demonstrates much lower IQR

with 2.53s for easy MOs and 1.35s for hard MOs, compared to the w/o method's 6.16s and 34.14s, respectively.

5.5 Ablation Study on Bias Between Estimated SR and Real SR

To obtain the initial SR value, we conduct prior experiments, as explained in Sec. 4.3. However, the estimated and actual SR may differ, especially when assuming all MOs share the same SR. Since SR relates to the estimated removal cost and affects the navigation strategy, we analyze the impact of errors in estimating the SR. The results on the cases with and without estimation bias in environment ABC are shown in Table 1. All the values of running time in the table are the average of 5 trials. In the method with action uncertainty (w/), an unbiased SR estimation gives the best navigation strategy with minimal time. Even with biased estimation, the method w/ still outperforms the method w/o, proving the effectiveness of the proposed action uncertainty module.

Table 1: Average running time of methods with unbiased/biased estimation on SR. The cell marked bold means better result.

	Estimated SR	Real SR	w/o (s)	w/ (s)
Unbiased	0.90	0.90	68.99±4.42	93.80±4.53
	0.50	0.50	98.31±21.81	94.19±12.86
	0.20	0.20	164.15±80.89	85.14±1.92
	Overall		110.48	91.04
Biased	0.90	0.20	164.15±80.89	168.19±39.50
		0.50	98.31±21.81	109.24±15.91
	0.50	0.20	164.15±80.89	113.46±22.17
		0.90	68.99±4.42	91.03±6.95
	0.20	0.50	98.31±21.81	85.57±2.31
		0.90	68.99±4.42	84.45±2.42
	Overall		110.48	108.66

Table 2: Average running time of method with/without blockage uncertainty in two environments.

Env	Methods	
	w/o (s)	w/ (s)
AB	67.04±2.54	77.77±2.63
ABE	141.02±16.81	90.08±2.25
Overall	104.03	83.92

5.6 Blockage Uncertainty Module Evaluation

To compare the impact of introducing the blockage uncertainty module, we design two environments AB and ABE for evaluation (with obstacles at the corresponding positions shown in Fig. 4). Environment AB demonstrates the difference of navigation strategy due

to blockage uncertainty while ABE illustrates the advantage of the blockage uncertainty module when an unexpected MO appears on the detour.

The evaluation results in Table 2 report the mean runtime of 5 trials. In environment AB, without considering the blockage uncertainty (w/o), the robot bypasses all the MOs. In contrast, the w/ method chooses to remove the MO in region B considering the potential blockage risk of the detour in narrow passage. In AB where no surprising MO appears, the bypass takes less time than the removal. However, in ABE, where an unexpected MO blocks the detour, the method w/o bypasses the MO in region B, and the MO in E (failing to find a suitable stock region for removing E), finally it returns to remove the MO at B, taking much longer time than the method w/ that removes B initially. From the overall performance, the method w/ is more efficient comparing to the method w/o.

5.7 Overall Comparison

We evaluate the overall performance in the two simulation environments (Fig. 4), with configurations of ABC, AB, BC, ABE, ABD, BCE in the warehouse. The starting points and goal points are randomly selected to create different navigation tasks. Then, we compare our method with some baseline methods, priority bypass (Rösmann et al., 2015; Hart et al., 1968), priority removal, random choice method and LaMB (Muguira-Iturralde et al., 2023).

The priority bypass methods including TEB (Rösmann et al., 2015) and A* (Hart et al., 1968), bypass all the obstacles but fail if there is no alternative way to the goal. The priority removal method refers to a category of NAMO methods (Ellis et al., 2022; Wang et al., 2020), which removes the MOs that block the path without considering the removal cost. The random choice method chooses to remove or bypass with a probability (0.5 in our experiments). The LaMB method (Muguira-Iturralde et al., 2023) is one of the latest methods considering the partial observation constraints in NAMO tasks but limited to small scale environments.

We record the running time of each method. Tasks are marked as failed if the goal is not reached within 300 seconds. Results are shown in Fig. 7. In the room environment, the only path to the goal is blocked by a MO. Therefore, the bypass method (A*) always fails. The NAMOUnc and priority removal methods complete the task more quickly, with the priority removal method slightly faster (77.68s vs. 79.56s) since it takes less time to plan bypass and make decision. In the warehouse environment, the NAMOUnc and the

Table 3: Overall performance comparison in the real environment. The columns marked bold represent the best results.

Env	TEB(Rösmann et al., 2015)	Priority Removal	Random Choice	NAMOUnc (Ours)
ABC	300.00±0.00	96.33±9.01	115.34±14.58	137.69±13.19
BC	46.66±4.56	94.78±5.08	72.73±19.94	50.41±2.62
AB	61.11±4.90	97.94±5.72	107.44±22.40	72.56±2.09
Overall	135.93	96.35	98.5	86.88




Figure 7: Running time in two simulation environments (around 70 trials for each environment), the room (left) and the warehouse (right). The outliers (circle around 300 in the figures) represent the failure cases. The orange line in the box marks the median value.




Figure 8: Real environment setup. The blue regions marked as A, B, C are possible positions of MOs. The red arrow with G is the goal.

TEB methods finish the task with comparable time cost but NAMOUnc has no failures.

6 REAL EXPERIMENTS

6.1 Environment Description

To evaluate the performance of our method in real applications, we use a Jackal differential mobile platform in a small warehouse-like environment. As shown in Fig. 8, there are maximum 3 MOs and the robot should navigate to a goal G. It is equipped with a LiDAR and a RealSense camera to observe the environment, an arm to lift MOs and the control software described in Sec. 5.2.

6.2 Experiment Results

We randomly pick goal points to create different NAMO tasks and record the running time with different MO setups, including environments ABC (all paths to the goal are blocked), AB and BC (at least one path to the goal is feasible). The quantitative results are shown in Table 3, where each cell indicates the average running time and corresponding standard deviation on 5 repeated trials. The priority bypass method (TEB) fails in the ABC environment because it finds no feasible path to the goal. Although the proposed NAMOUnc method does not achieve the best result in every individual environment, it attains the best overall performance across the three setups in terms of average running time. This shows that it achieves a good trade-off between completeness and efficiency in the search for a solution.

7 CONCLUSION AND FUTURE WORK

We have presented a NAMO framework capable of planning task and motion under four kinds of uncertainties related to observation, action, model and observability. Our planner jointly optimizes success rate and running time. Experimental results in both simulation and real environments demonstrate its ability to balance these objectives, suggesting potential extensions to optimize additional objectives like energy and safety.

A relevant perspective is to overcome simplified assumptions. For instance, observation noise may not always follow a Gaussian distribution, and action failures can stem from various factors, such as mechanical constraints that can depend on the obstacles and the environment. We believe that with more data collection, a more accurate model can be developed, such as a planner based on large language models (Honerkamp et al., 2024), enabling the proposed method to provide more robust and effective solutions for NAMO tasks.

ACKNOWLEDGEMENTS

The publication of this research was supported by the National Natural Science Foundation of China [Grant 42101445] and the Director Foundation of Guangdong Laboratory of Artificial Intelligence and Digital Economy(SZ) [Grant 25420001 and 24420004].

REFERENCES

- Coumans, E. and Bai, Y. (2016–2021). Pybullet, a python module for physics simulation for games, robotics and machine learning. <http://pybullet.org>.
- Curtis, A., Matheos, G., Gothoskar, N., Mansinghka, V., Tenenbaum, J., Lozano-Pérez, T., and Kaelbling, L. P. (2024). Partially observable task and motion planning with uncertainty and risk awareness. *arXiv preprint arXiv:2403.10454*.
- Denoeux, T. (2019). Decision-making with belief functions: A review. *International Journal of Approximate Reasoning*, 109:87–110.
- Ellis, K., Zhang, H., Stoyanov, D., and Kanoulas, D. (2022). Navigation among movable obstacles with object localization using photorealistic simulation. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1711–1716. IEEE.
- Feng, D., Harakeh, A., Waslander, S. L., and Dietmayer, K. (2021). A review and comparative study on probabilistic object detection in autonomous driving. *IEEE Transactions on Intelligent Transportation Systems*, 23(8):9961–9980.
- Garrett, C. R., Paxton, C., Lozano-Pérez, T., Kaelbling, L. P., and Fox, D. (2020). Online replanning in belief space for partially observable task and motion problems. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5678–5684. IEEE.
- Grissetti, G., Stachniss, C., and Burgard, W. (2007). Improved techniques for grid mapping with rao-blackwellized particle filters. *IEEE transactions on Robotics*, 23(1):34–46.
- Hart, P. E., Nilsson, N. J., and Raphael, B. (1968). A formal basis for the heuristic determination of minimum cost paths. *IEEE transactions on Systems Science and Cybernetics*, 4(2):100–107.
- Honerkamp, D., Büchner, M., Despinoy, F., Welschehold, T., and Valada, A. (2024). Language-grounded dynamic scene graphs for interactive object search with mobile manipulation. *IEEE Robotics and Automation Letters*.
- Kaelbling, L. P. and Lozano-Pérez, T. (2013). Integrated task and motion planning in belief space. *The International Journal of Robotics Research*, 32(9-10):1194–1227.
- Kalman, R. E. (1960). A new approach to linear filtering and prediction problems.
- Kim, B., Wang, Z., Kaelbling, L. P., and Lozano-Pérez, T. (2019). Learning to guide task and motion planning using score-space representation. *The International Journal of Robotics Research*, 38(7):793–812.
- Koenig, N. and Howard, A. (2004). Design and use paradigms for gazebo, an open-source multi-robot simulator. In *2004 IEEE/RSJ IROS*, volume 3, pages 2149–2154. IEEE.
- Li, C., Xia, F., Martin-Martin, R., and Savarese, S. (2020). Hrl4in: Hierarchical reinforcement learning for interactive navigation with mobile manipulators. In *Conference on Robot Learning*, pages 603–616. PMLR.
- Maier, S. F. and Seligman, M. E. (1976). Learned helplessness: theory and evidence. *Journal of experimental psychology: general*, 105(1):3.
- Muguira-Iturralde, J., Curtis, A., Du, Y., Kaelbling, L. P., and Lozano-Pérez, T. (2023). Visibility-aware navigation among movable obstacles. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 10083–10089. IEEE.
- Pan, T., Wells, A. M., Shome, R., and Kavraki, L. E. (2022). Failure is an option: task and motion planning with failing executions. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 1947–1953. IEEE.
- Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., Lin, Z., Desmaison, A., Antiga, L., and Lerer, A. (2017). Automatic differentiation in pytorch.
- Romero-Ramirez, F. J., Muñoz-Salinas, R., and Medina-Carnicer, R. (2018). Speeded up detection of squared fiducial markers. *Image and vision Computing*, 76:38–47.
- Rösmann, C., Hoffmann, F., and Bertram, T. (2015). Timed-elastic-bands for time-optimal point-to-point nonlinear model predictive control. In *2015 european control conference (ECC)*, pages 3352–3357. IEEE.
- Safronov, E., Colledanchise, M., and Natale, L. (2020). Task planning with belief behavior trees. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 6870–6877. IEEE.
- Silver, T., Chitnis, R., Tenenbaum, J., Kaelbling, L. P., and Lozano-Pérez, T. (2021). Learning symbolic operators for task and motion planning. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3182–3189. IEEE.

- Stanford Artificial Intelligence Laboratory et al. (2018). Robotic operating system.
- Wang, M., Luo, R., Önl, A. Ö., and Padir, T. (2020). Affordance-based mobile robot navigation among movable obstacles. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 2734–2740. IEEE.
- Williams, C. and Rasmussen, C. (1995). Gaussian processes for regression. *Advances in neural information processing systems*, 8.
- Wu, H.-n., Levihn, M., and Stilman, M. (2010). Navigation among movable obstacles in unknown environments. In *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1433–1438. IEEE.
- Xia, F., Li, C., Martín-Martín, R., Litany, O., Toshev, A., and Savarese, S. (2021). Relmogen: Integrating motion generation in reinforcement learning for mobile manipulation. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4583–4590. IEEE.
- Zhang, K., Lucet, E., dit Sandretto, J. A., and Filliat, D. (2023). Navigation among movable obstacles using machine learning based total time cost optimization. In *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 11321–11327. IEEE.
- Zhang, K., Lucet, E., Sandretto, J. A. D., Kchir, S., and Filliat, D. (2022). Task and motion planning methods: applications and limitations. In *ICINCO 2022-19th International Conference on Informatics in Control, Automation and Robotics*, pages 476–483. SCITEPRESS-Science and Technology Publications.