

Active-Perceptive Motion Generation for Mobile Manipulation

Snehal Jauhri^{1*}, Sophie Lueth^{1*}, and Georgia Chalvatzaki^{1,2,3}



Fig. 1: Real-world execution of our active perception and grasping pipeline. The mobile manipulator robot, equipped with a head-mounted camera, efficiently explores the scene to detect grasps on the target object and effectively performs the grasp by evaluating its executability. Video demonstrations are provided in the supplementary video and the project website: sites.google.com/view/actpermoma.

Abstract—Mobile Manipulation (MoMa) systems incorporate the benefits of mobility and dexterity, due to the enlarged space in which they can move and interact with their environment. However, even when equipped with onboard sensors, e.g., an embodied camera, extracting task-relevant visual information in unstructured and cluttered environments, such as households, remains challenging. In this work, we introduce an active perception pipeline for mobile manipulators to generate motions that are informative toward manipulation tasks, such as grasping in unknown, cluttered scenes. Our proposed approach, ActPerMoMa, generates robot paths in a receding horizon fashion by sampling paths and computing path-wise utilities. These utilities trade-off maximizing the visual Information Gain (IG) for scene reconstruction and the task-oriented objective, e.g., grasp success, by maximizing grasp reachability. We show the efficacy of our method in simulated experiments with a dual-arm TIAGo++ MoMa robot performing mobile grasping in cluttered scenes with obstacles. We empirically analyze the contribution of various utilities and parameters, and compare against representative baselines both with and without active perception objectives. Finally, we demonstrate the transfer of our mobile grasping strategy to the real world, indicating a promising direction for active-perceptive MoMa.

I. INTRODUCTION

We envision a near future where embodied agents, such as mobile manipulators, can operate autonomously in everyday environments like households. However, performing tasks in these environments is challenging due to the unstructured nature and unpredictability of the real world. Hence, some robots will actively gather information about their surroundings through embodied sensors, like an embodied camera whilst operating and changing their environment [1]. While current advances in AI and machine learning for

robotics have unlocked new capabilities for table-top manipulation [2]–[4], or language-driven navigation and manipulation [5]–[9], mobile manipulation in unknown (or partially known) scenes poses significant challenges [10]–[12], as the MoMa robot needs to consider both scene reconstruction and task-oriented objectives.

Active perception [13], [14] refers to the ability of an agent to “*know why it wishes to sense, and then chooses what to perceive, and determines how, when and where to achieve that perception,*” [1, p. 178]. For mobile robots, the robot’s objective is typically reconstruction, i.e., obtaining volumetric information about the scene/target object [15]–[18]. Many proposed active-perception methods use a Next-Best-View (NBV) strategy [19]–[23], primarily choosing viewpoints based on information gain (IG) [19] that minimizes uncertainty by exploring unobserved regions. A good overview and comparison of different IG formulations for NBV is provided in [24]. Notably, an active perception approach that only considers movement to the NBV with the most information gain can lead to unnecessarily large motions. Hence, the authors of [25] consider IG over paths using a graph-based approach, while in [26], a receding horizon viewpoint and path planning method is proposed since this formulation naturally benefits active perception as it adapts to newly observed information.

This work focuses on active perception that enables mobile manipulation in unknown environments, with a focus on mobile grasping. When grasping with static manipulators, recent methods adopt grasp quality metrics to choose the next robot viewpoint that minimizes uncertainty in the grasp pose estimation [2], [27]. Breyer et al. [2] exploit the fact that only negligible performance differences can be detected in the different formulations of volumetric IG [19], [24] and propose a rear-side voxel IG computation over views equally distributed on a half-sphere above the target object. The active perception is considered complete when a stable grasp quality has been found. However, these methods have several drawbacks when considering their application to MoMa.

*Authors contributed equally

This research received funding from the European Union’s Horizon program under grant agreement no. 101120823, project MANiBOT, the German Research Foundation (DFG) Emmy Noether Programme (CH 2676/1-1), and the Daimler Benz foundation.

¹Computer Science Department, Technische Universität Darmstadt, Germany ²Hessian.AI, Darmstadt, Germany ³Center for Mind, Brain and Behavior, Uni. Marburg and JLU Giessen, Germany {snehal.jauhri, georgia.chalvatzaki}@tu-darmstadt.de, sophie.lueth@stud.tu-darmstadt.de

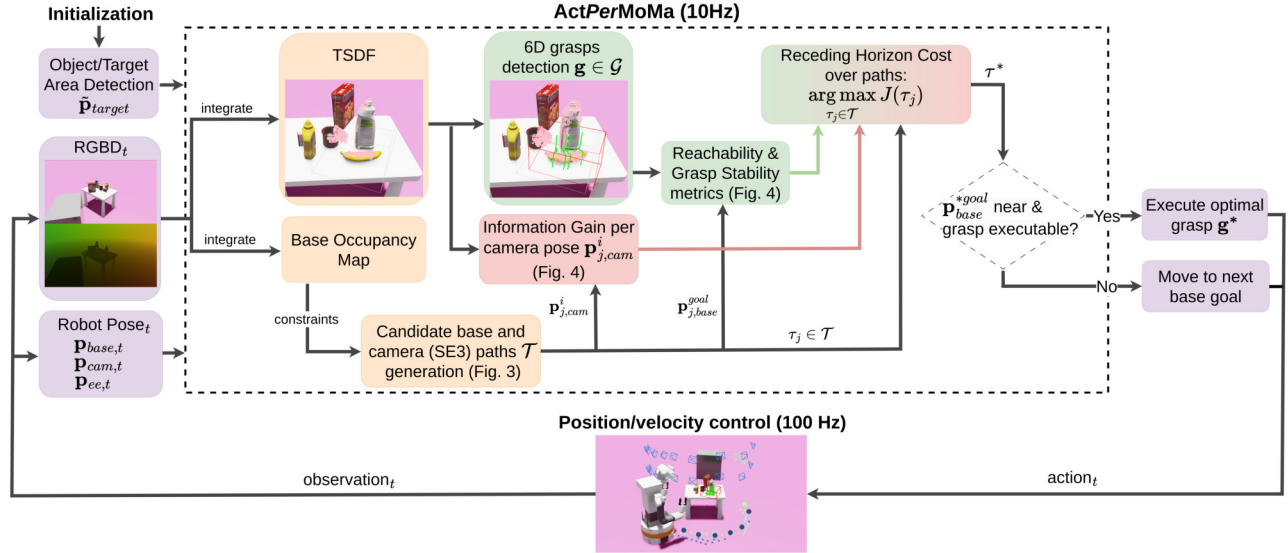


Fig. 2: ActPerMoMa pipeline. Using a rough initial knowledge about the target area or target object position $\tilde{\mathbf{p}}_{target}$, we continuously plan and execute informative motions for the mobile grasping task. At every timestep t , the RGBD information from the head-mounted embodied camera is integrated into a scene TSDF for both grasp detection and information gain computation. Using the currently known free space for movement of the robot base, we sample candidate robot paths \mathcal{T} , including both base and camera poses, towards the target. For each candidate path $\tau_j \in \mathcal{T}$, we compute the information gained from camera views $\mathbf{p}_{j,cam}^i$ in the path, and the reachability of stable detected grasps from the final base poses $\mathbf{p}_{j,base}^{goal}$ in the path. We trade-off these objectives with a receding horizon cost J_τ and take a step of the optimal path τ^* for execution at every timestep.

First, a MoMa robot can be placed *anywhere* in a scene and can reach *any* viewpoint (subject to scene restrictions). Thus, the range of movements of a MoMa robot is larger which makes wasteful motions especially costly. This also means that NBV-only approaches are sub-optimal as they do not consider information gained *during* movement from one NBV to another. Second, the reachability and feasibility of grasps and viewpoints are crucial when formulating planning for MoMa, since grasps of high quality could be challenging to reach and might cause unnecessary robot movements.

In this work, we propose an effective and efficient approach for visually informative motion generation for mobile manipulators to perform tasks in unknown, cluttered scenes. Particularly, we consider the problem of *mobile grasping*. Our method, Active-Perceptive Motion Generation for Mobile Manipulation (ActPerMoMa), resolves mobile grasping efficiently by planning over paths, collecting enough visual scene information to infer collisions and good grasps, and accounting for grasp executability. We ablate and benchmark our method against active grasping baselines, showcasing design choices that lead to performance gains toward MoMa applications.

To summarize: (i) We propose a novel formulation for active perceptive mobile manipulation, generating paths toward objects of interest without any prior scene knowledge. (ii) We calculate path-wise utilities over robot poses. Our motion generation objective balances exploration, maximizing visual information gain, and exploitation of task-specific information, such as grasp executability for mobile grasping. (iii) To ensure robot reactivity to new data, we use a receding-horizon control approach. We sample and evaluate numerous potential paths towards an approximate object location, based on feasible base goal-poses perceived in the current scene.

II. ACTIVE-PERCEPTIVE MOTION GENERATION FOR MOBILE MANIPULATION

We consider scenarios where a MoMa robot is placed in a previously unseen environment and is tasked with picking up a target object placed on a surface among clutter. To achieve this task, the mobile manipulator needs to use its multiple embodiments, i.e., a mobile base to move in the scene, an RGBD camera to perceive the scene and gather information, and an arm/end-effector to execute 6DoF grasps. Without loss of generality to different physical MoMa designs, we can simplify the description of the robot’s state as a combination of its mobile base pose $\mathbf{p}_{base} \in SE(2)$, its camera pose $\mathbf{p}_{cam} \in SE(3)$, and its end-effector pose $\mathbf{p}_{ee} \in SE(3)$. In principle, we could consider a whole-body joint representation, but in this work we decouple the poses of the different embodiments for simplicity.

Hypothesizing some rough knowledge of the area where the target object is, we obtain an approximate bounding box of the region where the target object lies, the center of which we denoted by $\tilde{\mathbf{p}}_{target}$. In practice, this can be done by either exploring and executing an RGB object detector or via information from a user instruction such as “Pickup the object from the right corner of the table”, but no prior information about the scene is assumed. Using the point cloud from the embodied camera we can build a volumetric representation of the scene, a 3D Truncated Signed Distance Function (TSDF), to effectively plan collision-free motions in the observed environment. The TSDF is also used to detect a set of 6DoF grasps $\mathcal{G} = \{\mathbf{g}_i\}_{i=0}^{N_g}$, where $\mathbf{g} \in SE(3)$ of maximum number N_g , in the object region using an $SE(3)$ grasp detection network that can process volumetric information, such as [28], [29]. Example scenarios we consider are visualized in Fig. 2 and Fig. 3.

Our overall objective is to generate efficient motions for the MoMa robot to find and execute a grasp to pick up the target object. For mobile grasping with active information accumulation, we consider the following design principles for motion generation:

- (i) Mobile manipulators operating under partial information need to trade-off between exploration (scene understanding) and exploitation (executing a task-oriented action).
- (ii) Since information is continuously gathered, the control formulation must be adaptive and reactive, and the costs must be considered and updated at every time interval.
- (iii) The key utilities to be balanced are the robot movement cost, the information gained about the target-object/scene, and the likelihood of grasp success.
- (iv) While grasp detectors [28], [29] may predict high-quality grasps, they are not necessarily feasible for the MoMa robot. Thus, grasping utility should consider metrics like the likelihood of reaching a grasp from different base locations [30].

In this section, we introduce a holistic motion generation pipeline that satisfies these principles (illustrated in Fig. 2). We sample many candidate feasible paths for the robot to execute the task while effectively using its embodiments (sec. II-A). To ensure that the robot continuously adapts to new information gathered, we choose a sampling-based receding horizon control formulation (sec. II-B). Our control formulation balances the objectives of information gain (sec. II-C) and the utility of grasp executability (sec. II-D).

A. Candidate goals & paths generation

Our objective is to move towards a target object, whose location is only roughly known in an unobserved scene, in the most informative and time/energy efficient manner to grasp it. We measure efficiency w.r.t the total distance traveled, viewpoints visited, and number of failures. At each time step, we sample candidate paths for the robot and evaluate utilities over those paths. To approach the target object, we sample N_b base poses near the approximate target object position $\tilde{\mathbf{p}}_{target}$ within a radius that affords robot reachability [30], [31]. These base poses serve as goals for our path generation $\{\mathbf{p}_{base}^{goal_i}\}_{i=0}^{N_b}$. We ensure that base goals are collision-free by performing a simple collision-check with the scene’s continuously generated TSDF grid or base occupancy map. The resampling of new base goals and paths at every time step ensures feasibility based on new scene information. We also jointly sample SE(3) camera poses at these base goal poses such that the robot always looks at the target area.

Our objective is to obtain the optimal motion of the robot toward the target object by planning to the candidate base goals, which should allow reaching the object to be grasped. We, thus, sample M candidate paths $\mathcal{T} = \{\tau_j\}_{j=0}^M$ to all the N_b candidate base goals $\{\mathbf{p}_{base}^{goal_i}\}_{i=0}^{N_b}$ using optimal path planners—in this work we plan over discretized grids with A^* . Each path $\tau \in \mathcal{T}$ consists of base poses from the current robot base to the base goal as well as sampled feasible camera poses $\mathbf{p}_{cam} \in SE(3)$ along the path, i.e., $\tau_i = \{\{\mathbf{p}_{base}^0, \mathbf{p}_{cam}^0\}, \{\{\mathbf{p}_{base}^1, \mathbf{p}_{cam}^1\} \dots \{\mathbf{p}_{base}^{goal}, \mathbf{p}_{cam}^{goal}\}\}$.




Fig. 3: Example scene with sampled candidate paths (blue) for the robot pose towards the target object (red box). The paths consist of SE(2) poses for the base and SE(3) poses for the head-mounted camera (visualized from the robot to the base goals). The current optimal path is highlighted in green.

Example candidate paths are visualized in Fig. 3. We then generate robot motion based on these paths using a receding horizon control formulation, as detailed next.

B. Receding-horizon control

We use a receding-horizon control formulation to generate the robot’s motion to find and execute a grasp on the target object. At each time step, we choose the optimal path among the sampled candidate paths \mathcal{T} and execute an action towards the first waypoint along this current optimal path. The re-computation of actions at every timestep ensures robot reactivity to newly observed scene information.

Formally, given the observation of the scene, i.e., the observed TSDF \mathbf{o}_{TSDF} , the detected set of grasps \mathcal{G} , and the sampled candidate paths \mathcal{T} , we compute the current optimal path $\tau^* \in \mathcal{T}$ based on the expected information gain J_{IG} and the utility of the grasps’ executability J_{exec} in the paths:

$$\tau^* = \arg \max_{\tau \in \mathcal{T}} J_{IG}(\mathbf{o}_{TSDF}, \tau) + J_{exec}(\mathcal{G}, \tau) \quad (1)$$

Utilities J_{IG} and J_{exec} are detailed in subsections II-C & II-D.

For movement at every time step, we use the first waypoint along the chosen optimal path, i.e., $\{\mathbf{p}_{base}^{*1}, \mathbf{p}_{cam}^{*1}\} \in \tau^*$ and run a low-level controller that executes IK-based velocities for the robot base and the camera. If the optimal path τ^* contains exactly one waypoint, i.e., the robot is close enough to the final chosen base goal $\mathbf{p}_{base}^{*goal}$, we finally consider grasp execution. If the grasp execution utility J_{exec} is above a threshold, we execute the grasp with the highest utility (sec. II-D) by activating the arm/end-effector and planning a motion to the SE(3) grasp.

C. Information gain computation & grasp detection

Information gain computation: Our perception objective is to obtain more information about the target object in order to grasp it. For this, we use an information gain (IG) formulation inspired by [24], [2]. We continuously build a voxel-based TSDF representation \mathbf{o}_{TSDF} of the scene and calculate the rear-side voxel information gain IG_{rear} [24]. In this setup, we cast rays from each candidate camera view $\mathbf{p}_{cam} \in SE(3)$ and count voxels that are on the rear side of the observed TSDF voxels and would thus be revealed by the candidate camera view. Since the approximate location of the target object is known, only voxels in an approximate




Fig. 4: Left: Example rear-side Information Gain (IG) for a candidate view. Pink voxels denote observed TSDF voxels. Blue voxels are on the rear side of the observed TSDF, which could be revealed by a candidate view. Views are colored red to green, denoting lower to higher IG. **Right:** Reachability map of the robot’s left arm, reduced from 6 dimensions (SE(3)) to 3 for visualization. Red and green points denote lower and higher reachability. Current detected 6D grasps are visualized in green on a target object.

bounding box around the target object are considered. Hence, this IG formulation rewards camera views that see more of the target object than has already been seen. More precisely per [2], for every viewpoint \mathbf{p}_{cam} , a set of rays R are generated by casting from a virtual camera placed at the respective view pose. Every ray r traverses voxels of the TSDF $v \subset \mathcal{O}_{TSDF}$ until it hits an observed surface. Therefore the rear-side IG is computed as $IG_{rear} = \sum_{r \in R} \sum_v \mathcal{I}(v)$, where $\mathcal{I}(v) = 1$ if the voxel is on the rear of an existing voxel and within the approximate target object bounding box.

Unlike [24], [2], we consider not only the Next-Best-View (NBV) but the IG *over paths* taken by the robot. Instead of only sampling a few viewpoints around the target object, we consider candidate viewpoints from each candidate path τ from our sampled paths \mathcal{T} . Moreover, we also consider the cost of reaching the viewpoints in the paths by weighting the IG by the distance to the viewpoints $dist(\mathbf{p}_{cam})$ along the path. This takes care of our requirement that *information gained sooner is better than later*. We can thus calculate the total IG over each candidate path τ as

$$J_{IG}(\mathcal{O}_{TSDF}, \tau) = \sum_{\mathbf{p}_{cam} \in \tau} \frac{IG_{rear}(\mathcal{O}_{TSDF}, \mathbf{p}_{cam})}{dist(\mathbf{p}_{cam})^2}. \quad (2)$$

An example visualization of the rear-side voxel IG is provided in Fig 4.

Grasp detection: At every time step, we query grasps in the target object region using the observed TSDF of the scene using a grasp detection network. In this work, we use the VGN [28] grasp detection network that can predict an SE(3) grasp pose for every 3D voxel of the TSDF along with a grasp quality prediction q . We use a grasp quality threshold q_{th} hyperparameter to detect good grasps with a high likelihood of success. Given that the built TSDF contains only partial/incomplete information, it is also important to consider grasp detector inaccuracy. Hence, as in [2], we also consider grasp *stability* by ensuring that a grasp predicted for the same 3D voxel on the TSDF has a high-quality score for n_{stab} steps. Finally, grasps that have a high likelihood of success but are challenging to reach can cause long, sub-optimal robot motions. To this end, we also score grasps based on their reachability from the candidate base goals around the target object, as detailed next.

D. Reachability utility & grasp selection

As we detect stable grasps over time, the motion generation should smoothly switch towards the actual execution of grasps, i.e., switching from exploratory to exploitative behavior. For this, the robot needs to be positioned at a base goal where a high-quality grasp can be executed easily. We achieve this ability by computing a grasp reachability/executability utility J_{exec} corresponding to each candidate path τ . The reachability of any SE(3) end-effector pose of a robot from a given base pose can be found by pre-computing a reachability map [32]–[34]. This map is obtained by executing forward kinematics for many different joint configurations and storing the visitations and the manipulability of the 6D voxels visited by the end-effector. We refer to [32], [34] for a full description of reachability map computation. A visualization of the reachability map used in our approach is provided in Fig 4.

In our pipeline, we pre-compute the reachability map \mathcal{R} and query it for *each grasp* $\mathbf{g} \in \mathcal{G}$, when executed *from each candidate base goal* \mathbf{p}_{base}^{goal} corresponding to our sampled paths $\tau \in \mathcal{T}$. The highest reachability over all grasps gives us a utility score J_{exec} for each candidate path τ . In the case of a dual-armed mobile manipulator (as used in our experiments) we can also compute the reachability of both arms and use the maximum for utility computation, as well as for arm selection. Moreover, as the robot reaches the optimal base goal, this utility is also used for grasp selection as we execute the most reachable grasp from the final base goal. To ensure that the proximity of the robot to the base goal of the path $\mathbf{p}_{base}^{goal} \in \tau$ is also considered, we weigh the reachability utilities by the length of the path $len(\tau)$, resulting in

$$J_{exec}(\mathcal{G}, \tau) = \frac{\max_{\mathbf{g} \in \mathcal{G}} \mathcal{R}(\mathbf{g}, \mathbf{p}_{base}^{goal})}{len(\tau)} \quad (3)$$

E. Additional hyperparameters

To smoothly switch between the two objectives in (1), we weigh the IG and grasp execution utilities by factors w_{IG} and w_{exec} . To avoid noisy grasps being used for movement and execution, we filter out unstable grasps, i.e., grasps that disappear after a few timesteps. Another problem that can appear is that the robot can oscillate between two base goals that move the robot in opposing directions if they have similar overall utility (1). Thus, we introduce a momentum term that continues to move the robot in a direction unless the utility of another direction is significantly higher.

III. EXPERIMENTS

A. Experimental setup & metrics

We conduct experiments both in simulation and the real world. Our setup contains a dual-armed TIAGo++ MoMa robot with a holonomic base and a head-mounted camera. The robot has 20 DoF in total (2 for the head, 1 for the torso, 3 for the holonomic base, and 7 for each arm), allowing 5D view sampling. In the simulated setup in Isaac Sim, we spawn the robot at a maximum distance of 2 m w.r.t. the approximate target location and consider two scenarios; a

TABLE I: Ablations & Hyperparameter study – Simple scenes

Hyperparameters					
Approach	SR (%) \uparrow	AR (%) \downarrow	GFR (%) \downarrow	d_{total} (m) \downarrow	v_{total} \downarrow
ActPerMoMa-Quality (0.7)	92.6	1.8	5.6	3.93\pm1.88	13.81\pm6.33
ActPerMoMa-Quality (0.9)	93.4	2.2	4.4	4.02 \pm 1.99	13.91 \pm 6.31
ActPerMoMa-StableGrasp (1)	94.2	1.0	4.8	4.01\pm1.99	13.87 \pm 6.22
ActPerMoMa-StableGrasp (5)	94.0	3.2	2.8	4.19 \pm 1.87	14.38 \pm 5.87
ActPerMoMa-IGweight (3.0)	91.0	5.6	3.4	4.18 \pm 2.21	14.36 \pm 6.70
ActPerMoMa-IGweight (0.2)	95.4	1.4	3.2	3.59\pm1.69	12.67 \pm 5.39
ActPerMoMa-momentum (0)	94.8	3.2	2.0	3.85 \pm 1.92	13.38 \pm 5.78
ActPerMoMa-momentum (700)	93.0	2.2	4.8	3.82 \pm 1.88	13.32 \pm 6.00
Ablation					
ActPerMoMa [†]	95.4	1.4	3.2	3.59 \pm 1.70	12.67 \pm 5.39
ActPerMoMa-IG-only	98.4	0.8	0.8	2.81\pm1.06	10.42\pm4.0
ActPerMoMa-no-weights	42.4	54.6	3.0	7.96 \pm 2.15	25.36 \pm 6.38
Hard Grasps					
ActPerMoMa [†]	66.8	23.6	9.6	5.33 \pm 2.77	18.75 \pm 9.04
ActPerMoMa-IG-only	65.2	31.6	3.2	3.59\pm2.35	17.88\pm9.52

[†] Quality=0.8, StableGrasp=1, IGweight=0.2, momentum=800

simple one, where a table is placed in free space and with 4 randomly spawned objects from the YCB dataset [35], and a more complex one with 6 objects to create more clutter and with a random obstacle sampled around the table to obstruct the path of the robot. For grasp detection, we used a pre-trained VGN [28]. The real-world setup resembles the one in simulation as shown in Fig. 1.

We conduct several experiments to ablate our method and justify the selection of key hyperparameters of our algorithm while also comparing against baselines. To measure the performance gains of the compared approaches, we employ the following metrics; **Success Rate** (SR), the percentage of episodes that finish with successful grasp execution, **Abort Rate**: (AR) the percentage of episodes that end without finding executable grasps in the given time budget, **Grasp Failure Rate** (GFR): the percentage of episodes ending in grasp failure, **total distance covered** (d_{total}), and the **total number of views** visited (v_{total}). We run every experiment for each scenario for 500 episodes and report average metrics and standard deviation when applicable. In the tables, we note superior performance of at least 0.5% in **bold**.

B. Ablations & hyperparameter study

To justify design decisions and better balance our algorithm’s exploration-exploitation dilemma, we conduct an extensive ablation study given the aforementioned metrics. In the following, we present the results for (i) ActPerMoMa-IG-only: ablation without the grasp executability objective in which case we execute a grasp as soon as we are within reach of the object, which in practice is 0.85 m away from the approximate target object; (ii) ActPerMoMa-no-weights: ablation without path-length-related scaling of the utilities (see eqs. 2, 3)). Additionally, we tune important parameters to find the best configuration for our method and present results for varying grasp quality thresholds (ActPerMoMa-Quality), grasp stability windows (ActPerMoMa-StableGrasp), IG weighting factors (ActPerMoMa-IGweight), and the momentum term that punishes oscillatory paths (ActPerMoMa-momentum).

Tables I and II present the results of our study for simple and complex scenes, respectively. Focusing first on Table I, we notice that in general for simple scenes the SR is high. That is because the robot has the freedom to explore the scene and find good grasps. From the hyperparameters, it

TABLE II: Ablations & Hyperparameter study – Complex scenes

Hyperparameters					
Approach	SR (%) \uparrow	AR (%) \downarrow	GFR (%) \downarrow	d_{total} (m) \downarrow	v_{total} \downarrow
ActPerMoMa-Quality (0.7)	91.4	0.6	8.0	4.57 \pm 2.39	16.47\pm8.78
ActPerMoMa-Quality (0.9)	88.8	5.0	6.2	4.84 \pm 2.73	17.11 \pm 9.54
ActPerMoMa-StableGrasp (1)	92.6	2.2	5.2	4.57\pm2.51	16.20\pm8.65
ActPerMoMa-StableGrasp (5)	91.2	2.8	6.0	4.92 \pm 2.48	17.47 \pm 8.80
ActPerMoMa-IGweight (3.0)	85.8	8.6	5.6	5.10 \pm 3.18	17.83 \pm 10.38
ActPerMoMa-IGweight (0.2)	92.6	1.8	5.6	4.31\pm2.27	15.60\pm8.39
ActPerMoMa-momentum (0)	83.2	13.0	3.8	5.45 \pm 3.40	18.77 \pm 11.06
ActPerMoMa-momentum (700)	90.6	3.8	5.6	4.56\pm2.53	16.28 \pm 9.11
Ablation					
ActPerMoMa [†]	92.6	1.80	5.6	4.31 \pm 2.27	15.60 \pm 8.39
ActPerMoMa-IG-only	96.8	0.8	2.4	3.52 \pm 1.61	12.87 \pm 6.31
ActPerMoMa-no-weights	48.4	42.0	10.6	9.05 \pm 3.79	13.28 \pm 10.44
Hard Grasps					
ActPerMoMa [†]	61.8	29.6	8.6	7.19 \pm 4.17	24.81\pm13.24
ActPerMoMa-IG-only	56.2	36.2	7.6	5.77\pm4.11	25.52 \pm 13.82

[†] Quality=0.8, StableGrasp=1, IGweight=0.2, momentum=800

is evident that even with a low grasp quality threshold we can achieve over 90% SR with low GFR, but the rest of the metrics are comparable. We see that the temporal window during which a detected grasp is stable does not really change the performance. However, we notice that using too large time windows (over 10 frames as in [2]) leads to poor mobile grasping results. The use of a small weight on the IG once we detect grasps seems to play a critical role in performance, as it shows that once we find a grasp, we can consider exploiting this information while maintaining some exploration. Finally, in simple scenes, the momentum may sometimes even lead to a drop in performance, as it may end in more failed attempts. Continuing with the ablations, we see that the method that does not penalize the utilities with the path-related scaling performs very poorly. We note that VGN, trained on table-top environments, mainly favors top-down grasps that are easier to execute. Conversely, when we restrict VGN to generate only side “hard” grasps (at 45 $^\circ$), we see a significant drop in performance in finding grasps (higher AR).

Table II, with complex scenes, shows a similar trend but with a noticeable drop in the overall performance, which is evident also by the generally higher grasp failure rates. Here we see the benefits of our momentum term as high momentum in complex scenes leads to better performance and reduced path lengths. Our ablation places the method without the grasp-related utility higher than the full ActPerMoMa approach, but in the “hard” grasp scenario, we notice a significant \sim 6% performance improvement when accounting for the grasps utility while planning. Overall, we propose our full objective as it has a higher chance of giving us reliable base poses for executing mobile grasps, especially when considering more realistic real-world scenarios where many hard grasps also exist, e.g., due to a specific grasp-affordance that needs to be respected.

C. Comparison with baselines

We first consider baseline methods that do not use active perception, in the sense of using an information gain objective, as they are intuitive and can show in which cases active perception is necessary. We then consider a state-of-the-art method in the active grasping literature. Namely, we compare against (i) a naive approach (Naive) in which we navigate the robot towards the approximate target location and activate grasp execution if, within a 0.85 m distance

TABLE III: Comparison with baselines

Approach	SR (%) \uparrow	AR (%) \downarrow	GFR (%) \downarrow	d_{total} (m) \downarrow	v_{total} \downarrow
Simple scenes					
Naive	95.2	1.2	3.6	1.36\pm0.28	5.93\pm3.77
Random	93.2	5.0	1.8	4.38 \pm 1.98	15.24 \pm 6.47
Breyer et al. [2]	92.0	8.0	0.0	3.71 \pm 1.78	12.56 \pm 6.12
ActPerMoMa (ours)	95.4	1.4	3.2	3.59 \pm 1.69	12.67 \pm 5.39
Complex Scenes					
Naive	86.8	2.6	10.6	1.55 \pm 0.45	8.72\pm8.89
Random	90.4	5.8	3.2	3.97 \pm 1.56	13.98 \pm 5.39
Breyer et al. [2]	90.0	6.0	4.0	3.32\pm1.47	12.32 \pm 6.13
ActPerMoMa (ours)	92.6	2.2	3.2	4.57 \pm 2.51	16.20 \pm 8.65
Complex Scenes (Hard Grasps)					
Naive	43.8	49.2	7.0	6.06 \pm 3.81	27.25 \pm 13.95
Random	24.6	68.4	7.0	3.82\pm2.23	16.29\pm11.07
Breyer et al. [2]	47.2	48.8	4.0	5.11 \pm 3.83	24.03 \pm 14.88
ActPerMoMa (ours)	61.8	29.6	8.6	7.19 \pm 4.17	24.81 \pm 13.24

from the object, a high-quality grasp has been detected; (ii) a random approach (Random) in which, at each time step, we randomly select a feasible base goal (i.e. no collision found with the currently perceived scene, also considering the reachability of the object from the base goal) around the approximate target object location. If a grasp has been detected, we execute the grasp. If the grasp is not successful, we resample a feasible goal with a smaller radius (0.75 m); and (iii) the method by Breyer et al. [2] adapted for mobile manipulation, in which we compute the IG per view (and not accumulated over paths) sampled on a hemisphere of radius 1m around the approximate object location. In this case, we always move to the viewpoint with the highest IG. If no grasp is found, we resample views with a smaller radius, and if we are within reach of the object and a grasp has been found, we execute it. In case of grasp failure, we move to the NBV.

Table III presents the comparative results both for simple and complex scenes. Interestingly, both the Naive and Random approaches outperform the active grasping approach of Breyer et al. [2]. Thus, heuristic approaches like the Naive approach can work in simple scenes, alleviating the need for path planning. The landscape changes when looking at the results for the complex scenes. ActPerMoMa still has the highest performance, while the naive approach now performs worse than Breyer et al. [2]. The random method performs well by moving close to the target and continuously sampling different base goals close to it. We posit that, as with ActPerMoMa-IG-only, already sampling base goals that are collision-free according to the currently perceived scene and within the reachability radius of the robot is a strong inductive bias for planning successful robot placements for robotic grasping. Nevertheless, as with the previous experiments in sec. III-B, the hard-grasp case shows the significant benefit of ActPerMoMa w.r.t baselines. Notably, we highlight the large benefit in finding grasps (at least $\sim 20\%$ lower abort rate) compared to baselines.

D. Real robot demonstration

To evaluate the applicability of ActPerMoMa we conduct real-world experiments in a cluttered room, as shown in Fig. 1. We ran 10 experiments, re-arranging the setup every time and selecting a different target object. The target object is placed on the table while the rest of the scene objects act as obstructions. The grasp quality threshold is increased to 0.9, higher than in simulation, to account for

the sim-to-real gap for the grasping network. We achieved real-world performance with an 80 % success rate, 20 % aborts, and no grasp failures, with an average distance traveled of 3.09 ± 1.06 m, and 13.50 ± 13.86 number of views. Additional details and demonstrations are provided in the supplementary video and the project website: <https://sites.google.com/view/actpermoma>.

E. Limitations

An issue of methods that use reactive planning (such as ours) is that, depending on the resolution of the sampled base goals and the sampling frequency, the robot can get stuck in deadlocks trying to switch between base goals leading to oscillating motions. Although we introduce a penalty for this behavior, namely the momentum, some amount of deadlocks due to changes in direction can still exist. Another limitation is the limited volumetric information in the target area in very occluded scenes, making the IG computation difficult. This can prominently be observed in Breyer et al. [2], as very occluded objects that get partially discovered from different views often prompt a ‘zigzag’ path. We improve this behavior by not just considering the best NBV to decide where to go but instead using the whole spatial distribution provided by our sampled base goals. A possible mitigation for this in future work is to train a reinforcement learning agent on these POMDP problems, leveraging a combination of our active and some random exploration.

While the real-world demonstration shows promise toward autonomous perceptive MoMa, some practical challenges remain. Most prominently, the camera frame rate and RGBD integration can significantly limit the algorithm’s control frequency (ActPerMoMa frequency is about 10 Hz, but can drop to 5 Hz in the real world). Secondly, continuous stable grasp detection while moving can be challenging, especially with imperfect localization.

IV. CONCLUSION

In this paper, we delved into the intricate challenges of active perception for mobile manipulation in unknown and cluttered environments. We introduced a novel formulation that generates robot paths toward objects of interest without any prior scene knowledge, drawing upon principles of active perception and mobile manipulation coordination. Our approach seamlessly combines exploration — for maximizing visual information gain—and exploitation of task-specific parameters such as grasp executability. Using a receding horizon control strategy, we ensure the robot’s motion can adapt dynamically to new data. Our experiments with the dual-arm TIAGo++ mobile manipulation robot have further validated the feasibility and efficiency of our proposed method in cluttered environments. While our results have shown how active perception can provide performance gains for efficient mobile grasping, the vast potential of active perceptive mobile manipulation remains uncharted territory. Looking ahead, we aim to explore deep learning techniques to predict scene information gain to use in robot learning for mobile manipulation to tackle even more challenging tasks.

REFERENCES

- [1] R. Bajcsy, Y. Aloimonos, and J. K. Tsotsos, "Revisiting active perception," *Autonomous Robots*, vol. 42, pp. 177–196, 2018.
- [2] M. Breyer, L. Ott, R. Siegwart, and J. J. Chung, "Closed-loop next-best-view planning for target-driven grasping," 2022. [Online]. Available: <https://arxiv.org/abs/2207.10543>
- [3] C. Chi, S. Feng, Y. Du, Z. Xu, E. Cousineau, B. Burchfiel, and S. Song, "Diffusion policy: Visuomotor policy learning via action diffusion," *arXiv preprint arXiv:2303.04137*, 2023.
- [4] A. Zeng, P. Florence, J. Tompson, S. Welker, J. Chien, M. Attarian, T. Armstrong, I. Krasin, D. Duong, V. Sindhwani *et al.*, "Transporter networks: Rearranging the visual world for robotic manipulation," in *Conference on Robot Learning*. PMLR, 2021, pp. 726–747.
- [5] C. Huang, O. Mees, A. Zeng, and W. Burgard, "Visual language maps for robot navigation," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 10 608–10 615.
- [6] H. Wang, W. Wang, W. Liang, S. C. Hoi, J. Shen, and L. V. Gool, "Active perception for visual-language navigation," *International Journal of Computer Vision*, vol. 131, no. 3, pp. 607–625, 2023.
- [7] D. Shah, B. Osiński, S. Levine *et al.*, "Lm-nav: Robotic navigation with large pre-trained models of language, vision, and action," in *Conference on Robot Learning*. PMLR, 2023, pp. 492–504.
- [8] N. Yokoyama, A. W. Clegg, E. Undersander, S. Ha, D. Batra, and A. Rai, "Adaptive skill coordination for robotic mobile manipulation," *arXiv preprint arXiv:2304.00410*, 2023.
- [9] M. Ahn, A. Brohan, N. Brown, Y. Chebotar, O. Cortes, B. David, C. Finn, C. Fu, K. Gopalakrishnan, K. Hausman *et al.*, "Do as i can, not as i say: Grounding language in robotic affordances," *arXiv preprint arXiv:2204.01691*, 2022.
- [10] M. Mittal, D. Hoeller, F. Farshidian, M. Hutter, and A. Garg, "Articulated object interaction in unknown scenes with whole-body mobile manipulation," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 1647–1654.
- [11] F. Xia, C. Li, R. Martín-Martín, O. Litany, A. Toshev, and S. Savarese, "ReImogen: Integrating motion generation in reinforcement learning for mobile manipulation," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 4583–4590.
- [12] J. Pankert and M. Hutter, "Perceptive model predictive control for continuous mobile manipulation," *IEEE Robotics and Automation Letters*, vol. 5, no. 4, pp. 6177–6184, 2020.
- [13] J. Aloimonos, I. Weiss, and A. Bandyopadhyay, "Active vision," *International journal of computer vision*, vol. 1, pp. 333–356, 1988.
- [14] Y. Aloimonos, *Active perception*. Psychology Press, 2013.
- [15] J. I. Vasquez-Gomez, L. E. Sucar, R. Murrieta-Cid, and E. Lopez-Damian, "Volumetric next-best-view planning for 3d object reconstruction with positioning error," *International Journal of Advanced Robotic Systems*, vol. 11, no. 10, p. 159, 2014.
- [16] C. Potthast and G. S. Sukhatme, "A probabilistic framework for next best view estimation in a cluttered environment," *Journal of Visual Communication and Image Representation*, vol. 25, no. 1, pp. 148–164, 2014.
- [17] T. Zaenker, C. Lehnert, C. McCool, and M. Bennewitz, "Combining local and global viewpoint planning for fruit coverage," in *2021 European Conference on Mobile Robots (ECMR)*. IEEE, 2021, pp. 1–7.
- [18] L. Schmid, M. Pantic, R. Khanna, L. Ott, R. Siegwart, and J. Nieto, "An efficient sampling-based method for online informative path planning in unknown environments," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 1500–1507, 2020.
- [19] S. Isler, R. Sabzevari, J. A. Delmerico, and D. Scaramuzza, "An information gain formulation for active volumetric 3d reconstruction," *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 3477–3484, 2016.
- [20] A. Bircher, M. Kamel, K. Alexis, H. Oleynikova, and R. Siegwart, "Receding horizon" next-best-view" planner for 3d exploration," in *2016 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2016, pp. 1462–1468.
- [21] D. Watkins-Valls, P. K. Allen, H. Maia, M. Seshadri, J. Sanabria, N. Waytowich, and J. Varley, "Mobile manipulation leveraging multiple views," 2021. [Online]. Available: <https://arxiv.org/abs/2110.00717>
- [22] M. Naazare, F. G. Rosas, and D. Schulz, "Online next-best-view planner for 3d-exploration and inspection with a mobile manipulator robot," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 3779–3786, apr 2022.
- [23] L. Bartolomei, L. Teixeira, and M. Chli, "Semantic-aware active perception for uavs using deep reinforcement learning," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 3101–3108.
- [24] J. Delmerico, S. Isler, R. Sabzevari, and D. Scaramuzza, "A comparison of volumetric information gain metrics for active 3d object reconstruction," *Autonomous Robots*, vol. 42, no. 2, 2018.
- [25] T. Zaenker, J. Rückin, R. Menon, M. Popović, and M. Bennewitz, "Graph-based view motion planning for fruit detection," *arXiv preprint arXiv:2303.03048*, 2023.
- [26] A. Bircher, M. Kamel, K. Alexis, H. Oleynikova, and R. Siegwart, "Receding horizon" next-best-view" planner for 3d exploration," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*, 2016, pp. 1462–1468.
- [27] D. Morrison, P. Corke, and J. Leitner, "Multi-view picking: Next-best-view reaching for improved grasping in clutter," in *2019 International Conference on Robotics and Automation (ICRA)*, 2019, pp. 8762–8768.
- [28] M. Breyer, J. J. Chung, L. Ott, R. Siegwart, and J. I. Nieto, "Volumetric grasping network: Real-time 6 DOF grasp detection in clutter," in *CoRL*, ser. Proceedings of Machine Learning Research, vol. 155. PMLR, 2020, pp. 1602–1611.
- [29] Z. Jiang, Y. Zhu, M. Svetlik, K. Fang, and Y. Zhu, "Synergies between affordance and geometry: 6-dof grasp detection via implicit representations," in *Robotics: Science and Systems*, 2021.
- [30] S. Jauhuri, J. Peters, and G. Chalvatzaki, "Robot learning of mobile manipulation with reachability behavior priors," *IEEE Robotics and Automation Letters*, vol. 7, no. 3, pp. 8399–8406, 2022.
- [31] T. Birr, C. Pohl, and T. Asfour, "Oriented surface reachability maps for robot placement," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 3357–3363.
- [32] F. Zacharias, C. Borst, and G. Hirzinger, "Capturing robot workspace structure: representing robot capabilities," in *2007 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2007, pp. 3229–3236.
- [33] N. Vahrenkamp, T. Asfour, and R. Dillmann, "Robot placement based on reachability inversion," in *ICRA*, 2013.
- [34] A. Makhmal and A. K. Goins, "Reuleaux: Robot base placement by reachability analysis," in *IRC*. IEEE Computer Society, 2018, pp. 137–142.
- [35] B. Calli, A. Singh, A. Walsman, S. Srinivasa, P. Abbeel, and A. M. Dollar, "The ycb object and model set: Towards common benchmarks for manipulation research," in *2015 international conference on advanced robotics (ICAR)*. IEEE, 2015, pp. 510–517.