

# Value of Assistance for Grasping

Mohammad Masarwy, Yuval Goshen, David Dovrat and Sarah Keren

**Abstract**—In many realistic settings, a robot is tasked with grasping an object without knowing its exact pose and needs to rely instead on a probabilistic estimation of the pose in order to decide how to attempt the grasp. We support settings in which it is possible to provide the robot with an observation of the object before a grasp is attempted. However, since this possibility is limited, there is a need to predict the expected benefit of possible sensing actions before choosing which one to perform. For this, we offer new ways to compute *Value of Assistance (VOA)* for grasping to predict the expected effect a sensing action will have on the probability of a successful grasp. Our empirical evaluation shows how our suggested VOA measure can be used to identify helpful sensing actions.

## I. INTRODUCTION

Task-driven agents often need to decide how to act based on partial and noisy state estimations, which may greatly compromise performance. We consider settings in which an agent is tasked with grasping an object based on a probabilistic estimation of its pose. Before attempting the grasp, it may be possible to provide the agent with an additional observation. We offer *value of assistance (VOA)* for grasping as a principled framework for assessing the expected benefit of performing a sensing action and suggest ways to compute it for two different sensors.

VOA is relevant to settings in which sensing and communication may be costly or limited and there is a need to select a sensing action to perform. Notably, we do not offer a new approach to searching for the best action. Instead, we offer a measure that can guide any off-the-shelf heuristic search in finding the best sensing action to apply. It can also guide an anytime search in which it is unknown a-priori how much time there is for searching for the best intervention before the actor reattempts the grasp. In both cases, there is a need to iteratively select the next sensing action to perform and therefore a need to efficiently estimate the expected benefit of each action.

To demonstrate, consider the simplified automated manufacturing setting depicted in Figure 1. The acting agent marked as the *actor*, is a robotic arm with a parallel-jaw gripper that is tasked with grasping an object (here, an adversarial object from the Dex-Net Dataset<sup>1</sup>). After a successful grasp, the object drops unexpectedly. Since the actor does not have the means to sense the state of the dropped object on its own, it can attempt to grasp the object based only on its estimation of the current position of the object. Alternatively, it can attempt the grasp after receiving an observation from another agent that is equipped with a sensor (here, an OnRobot 2.5D Vision System). The question




Fig. 1: Collaborative grasping example.

we pose is *what is the best position from which to acquire a sensor reading to maximize the actor’s probability of successfully grasping the object*. Naturally, the same question arises in single-agent settings in which the actor needs to decide whether to perform a costly sensing action by itself.

Beyond this illustrative example, grasping is an essential task in a wide range of robotic applications, including industrial automation, household robotics, agriculture, and more [1], [2]. Accordingly, research on effective grasping has resulted in many solution approaches that can be generally divided into two main categories [3], [2]. In analytical approaches, a representation of the physical and dynamical models of the agent and the object are used when choosing a configuration from which to attempt a grasp [4], [5], [6], [7]. In contrast, data-driven approaches rank labeled samples to come up with grasping policies. The ranking is usually based on a heuristic or on experiences collected from simulated or real robots [8], [9], [10], [11], [12].

We assume the actor is associated with a procedure for choosing a grasp given its *belief* which represents its knowledge about the position of the object. To support the decision of which observation would be most beneficial, we formulate VOA for grasping and offer ways to compute it. This involves accounting for how the actor’s estimation will change based on the acquired observation and assessing how this change will affect the actor’s decision of how to attempt the grasp.

As a measure for assessing the informative value of an observation, VOA is closely related to the well-established notions of *value of information (VOI)* and *information gain (IG)* [13], [14], [15], [16], [17], which are widely used across multiple AI frameworks to assess the impact information will have on agents’ decisions and expected utility. We adapt these ideas to robotic settings. While the notion of VOA was previously used for assessing the effect localization information would have on a navigating robot’s expected cost [18], we adapt it to a grasping task.

Our work lies within the well established fields of *active perception* and *sensor planning* [19], [20], [21], [22], [23],

<sup>1</sup>The authors are from the Taub Faculty of Computer Science, Technion - Israel Institute of Technology, Israel




Fig. 2: Example stable poses.




Fig. 3: Example grasp configurations from which the actor can attempt to grasp the object - each configuration is associated with a score, i.e., probability of success.

[24], [25], [26], [27] which refer to the integration of sensing and decision-making processes within a robotic system. These involve actively acquiring and utilizing information from the environment and selecting viewpoints or trajectories likely to reveal relevant information or reduce uncertainty [28], [29]. While these include work on active perception in manipulation tasks, they mostly focus on assessing the effect various perspectives will have on the ability to correctly locate and classify objects [19], [30]. We adopt a task-driven approach and offer a general formulation of VOA for grasping and use it for estimating the effect an observation will have on the probability of accomplishing a grasp.

Since our focus is on settings in which information acquisition actions are limited and may be performed by another agent, our work is also highly related to *decision-theoretic communication*, where agents communicate over a limited-bandwidth channel and messages are chosen to maximize the utility or effectiveness of the communication [31], [26], [25], [32], [33]. We offer measures that account for the manipulation and sensing capabilities of robotic agents when assessing the value of communicating an observation.

Our key contributions include the formulation of *Value of Assistance (VOA)* for grasping and its instantiation for grasping settings with a robotic arm equipped with a gripper and another agent equipped with either a lidar or a depth camera. We empirically demonstrate in both simulated and real-world robotic settings how VOA predicts the effect an observation will have on performance and how it can be used to identify the best among a set of possible assistive actions.

## II. PRELIMINARIES

To support a grasping task, where the object is assumed to be in a stable static pose, we use a function that assigns

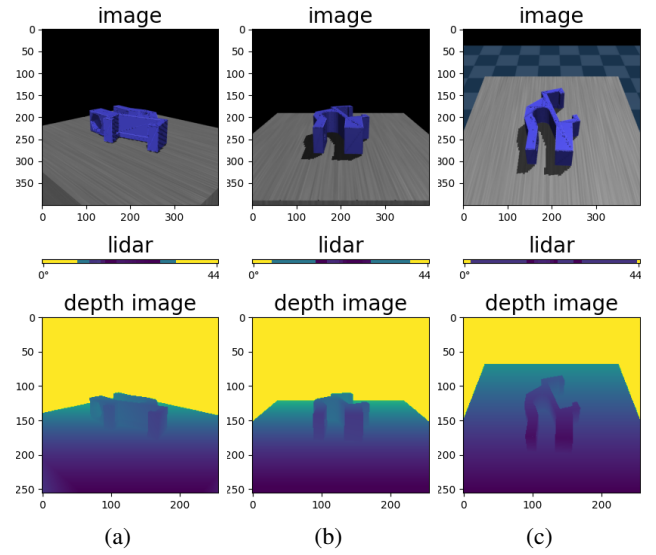


Fig. 4: Example sensor configurations and corresponding observations for a given stable pose. Each column represents the RGB image [top] lidar reading [middle] and depth image [bottom] for a sensor configuration-object pose pair.




Fig. 5: Grasps for HOLDER: best grasp  $g^*$ , initial chosen grasp  $g_i$  and grasp chosen after the intervention  $g_x$ .

a score to a grasping configuration - object pose pair.

*Definition 1 (Grasp Score):* Given a set of object poses  $\mathcal{P}$  and a set of grasp configurations  $\mathcal{G}$ , a *grasp score function*  $\gamma : \mathcal{G} \times \mathcal{P} \mapsto [0, 1]$  specifies the probability that an actor applying grasp configuration  $g \in \mathcal{G}$  will successfully grasp an object at pose  $p \in \mathcal{P}$ , i.e.  $\gamma(p, g) = P(s|p, g)$ , where  $s$  is the event of a successful grasp.

The grasp score function may be evaluated analytically, by considering diverse factors such as contact area, closure force, object shape, and friction coefficient [34], [2] or empirically, by using data-driven approaches where a deep learning model is trained to predict the quality of grasps based on depth images of the objects [35].

The actor's choice of a grasp configuration relies on a *pose belief* describing the perceived likelihood of each object pose within the set of possible stable poses  $\mathcal{P}$ . We let  $\mathcal{B}$  denote the set of pose beliefs and define a pose belief as follows.

*Definition 2 (Pose Belief):* A *pose belief*  $\beta : \mathcal{P} \mapsto [0, 1]$  is a probability distribution over  $\mathcal{P}$ .

The pose belief is affected by different factors, including

the model of the object and its dynamics and the collected sensory information. We formulate the initial belief after the object is dropped (Figure 1) using a joint probability model that captures the prior probability of stable poses, a von Mises PDF for the angle [36], and a multivariate normal distribution for the position on the plane (see the online appendix for the complete formulation<sup>1</sup>).

*Definition 3 (Expected Grasp Score):* For grasp configuration  $g$  and pose belief  $\beta$ , the *expected grasp score*  $\bar{\gamma}(g, \beta) = \mathbb{E}_{p \sim \beta}[\gamma(g, p)]$  is the weighted aggregated grasp score over the set of possible poses. A *maximal grasp* of  $\beta$ , denoted  $g^{max}(\beta)$ , maximizes the expected grasp score, i.e.,  $g^{max}(\beta) = \arg \max_{g \in \mathcal{G}} \bar{\gamma}(g, \beta)$ .

The actor receives an *observation* which corresponds to a reading from a specific sensor configuration  $x \in \mathcal{X}$  and object pose  $p \in \mathcal{P}$ . Our formulation of the *observation space*  $O$  is general and represents the set of readings that can be made by the sensor that is available in the considered setting. In our evaluations, we used a planar lidar sensor for which a reading is an array of non-negative distances per angle  $\mathbb{R}^{360}$  and a depth camera which emits a 2D array  $\mathbb{R}^{w \times h}$ , where  $w \times h$  are the image dimensions.

As is common in the literature, we consider obtaining an observation as a stochastic process [37], [38].

*Definition 4 (Sensor Function):* Given object pose  $p \in \mathcal{P}$  and sensor configuration  $x \in \mathcal{X}$ , sensor function  $\mathcal{O} : \mathcal{P} \times \mathcal{X} \mapsto O$  is a random function, such that if  $o = \mathcal{O}(p, x)$  then  $P(o|p, x)$  provides the conditional probability of obtaining the observation  $o$  when the sensor configuration is  $x$  and the object pose is  $p$ .

Notably, an agent may not be aware of the actual distribution of sensor functions or observations, and may instead only have a *predicted sensor function*  $\hat{\mathcal{O}}$  and a *predicted observation probability*  $\hat{P}$ , based on distributions which may be incorrect or inaccurate. For example, a predicted sensor function based on a deterministic model would assign a cell in a lidar reading based on the predicted distance between the lidar and the object surface at a specific angle. A stochastic sensor model would sample from a Gaussian distribution with this value as the mean and the specified error margins as the standard deviation. If  $\hat{P}$  is unknown, a similarity score,  $\omega : O \times O \mapsto [0, 1]$ , may be used to compare the predicted and received observations and to compose a valid distribution function for  $\hat{P}$ :

$$\hat{P}(o|p, x) = \frac{\omega(\hat{\mathcal{O}}(p, x), o)}{\int_{p' \in \mathcal{P}} \omega(\hat{\mathcal{O}}(p', x), o) dp'} \quad (1)$$

The literature is rich of various definitions for  $\omega$ , which may vary between applications and sensor types. Our appendix describes several approaches including using MSE for assessing the similarity between lidar sensor readings and an SSIM-based measure [39] for depth images.

When receiving an observation  $o$ , the actor updates its belief using its *belief update function* which defines the effect an observation has on the pose belief.

*Definition 5 (Belief Update):* A belief update function  $\tau : \mathcal{B} \times O \times \mathcal{X} \mapsto \mathcal{B}$  maps belief  $\beta \in \mathcal{B}$ , observation  $o \in O$  and sensor configuration  $x \in \mathcal{X}$  to an updated belief  $\beta^{o,x}$ .

There are many approaches for belief update [37], [16], [22], [40]. We use a Bayesian filter such that for any observation  $o \in O$  taken from sensor configuration  $x \in \mathcal{X}$ , the updated pose belief  $\beta^{o,x}(p)$  for pose  $p \in \mathcal{P}$  is given as

$$\beta^{o,x}(p) = \frac{\hat{P}(o|p, x) \beta(p)}{\int_{p' \in \mathcal{P}} \hat{P}(o|p', x) \beta(p') dp'} \quad (2)$$

where  $\beta(p)$  is the estimated probability that  $p$  is the object pose prior to considering the new observation  $o$ .

### III. VALUE OF ASSISTANCE (VOA) FOR GRASPING

We offer a way to assess *Value of Assistance (VOA) for grasping* as the expected benefit an observation collected from a sensor configuration will have on the probability of a successful grasp. For this, we seek a way to estimate beforehand the effect an expected observation will have on the actor's belief and thus on its choice of configuration from which to attempt the grasp.

Whereas the object's exact pose is unknown, its shape is given and it is assumed to be in a stable pose  $p$  (Figure 2). The actor chooses a feasible *grasp configuration*  $g \in \mathcal{G}$  from which to attempt a grasp (Figure 3) based on the grasp score function  $\gamma$  (Definition 1) and its *pose belief*  $\beta$  (Definition 2). Before attempting the grasp, it is possible to acquire an observation from a *sensor configuration*  $x \in \mathcal{X}$  but there is a need to choose among a set of configurations, each offering a different point of view of the object (Figure 4). Each observation may have a different effect on the pose belief and on the decision of which grasp to attempt. Figure 5 depicts the relationship between the grasp chosen by the agent based on its initial belief and the objective of using a sensing action to induce the selection of a pose with a higher probability of success.

Our perspective is that of the agent that needs to decide which sensing action to perform. This can be the actor itself or another agent in the environment. If the observation is provided by another agent, its belief may be different. We denote the beliefs of the actor and helping agent as  $\beta_a$  and  $\beta_h$ , respectively.

A key element in VOA computation is the expected difference between the utility of the actor with and without the intervention.

*Definition 6 (Value of Assistance (VOA) for Grasping):* Given the actor's belief  $\beta_a \in \mathcal{B}$ , the belief of the helping agent  $\beta_h \in \mathcal{B}$ , the predicted observation probability  $\hat{P}$ , sensor configuration  $x \in \mathcal{X}$ , and the actor's belief update function  $\tau_a$ ,

$$U_x^{VOA}(\beta_h, \beta_a) \stackrel{\text{def}}{=} \quad (3)$$

$$\mathbb{E}_{p \sim \beta_h} \left[ \mathbb{E}_{o \sim \hat{P}(o|p, x)} [\gamma(g^{max}(\beta_a^{o,x}), p)] - \gamma(g^{max}(\beta_a), p) \right].$$

In the definition above, the expected agent utility is based on the grasp score  $\gamma$  of the configuration  $g^{max}$  chosen by the actor based on its beliefs  $\beta_a^{o,x}$  and  $\beta_a$  with and without the intervention, respectively. We note that the pose  $p$  is sampled

<sup>1</sup>Online appendix: <https://github.com/CLAIR-LAB-TECHNION/VOAGrasp>

from  $\beta_h$ , since VOA is computed from the perspective of the helping agent and based on its belief.

Next, we consider a special case in which there is only one belief  $\beta$ , either because a single agent is performing both the sensing and grasping or because the two agents share a common belief. In addition, for simplicity, we assume the predicted observation  $\hat{o} = \hat{O}(p, x)$  is generated using a deterministic sensor function. This allows us to use a simplified VOA formulation which considers the difference between the expected grasp score of the actor’s choice with and without collecting an observation from  $x$ .

$$U_x^{VOA}(\beta) \stackrel{\text{def}}{=} \mathbb{E}_{p \sim \beta} [\gamma(g^{max}(\beta^{\hat{o}, x}, p)) - \gamma(g^{max}(\beta), p)] \quad (4)$$




Fig. 6: Processed observations (using masking) for a given sensor configuration for different poses of HOLDER.

Figure 6 demonstrates the role of VOA. Here, there are 6 possible stable poses, with a uniform prior probability, and the predicated observation from a given sensor configuration for each pose. This sensor configuration has a high VOA value since it clearly distinguishes between poses P1-P4 for which a single grasp has a high probability of success and P5-P6 for which a different grasp will have a high probability.

Our proposed VOA measure can be used within any algorithm that searches for an optimal intervention to perform and in which a heuristic is used to assess the effect of an intervention. While developing such an algorithm is out of scope for this work, we provide an example of a state space search that uses VOA to support the decision of which sensing action to perform in our online appendix.

#### IV. EMPIRICAL EVALUATION

The objective of our evaluation is to examine the ability of our VOA measures to predict the effect sensing actions will have on the probability of a successful grasp. After presenting the empirical setup both in simulation and at the lab, we measure the accuracy of our predicted sensor functions and show how these are used to compute VOA.

##### A. Experimental Setting

We performed our evaluation in a two-agent robotic setting. In our lab setup, depicted in Figure 1, the actor is a UR5e robotic arm<sup>2</sup> with an OnRobot 2FG7 parallel jaw gripper<sup>3</sup>. We examined two different sensors: a LDS-01 lidar<sup>4</sup> that could be moved on the x-y plane and a 2.5D Onrobot vision system<sup>5</sup> mounted on an adjacent UR5e arm. For simulation, we used a MuJoCo<sup>6</sup> environment (depicted in Figure 3) [41]. We simulated a lidar sensor using the MuJoCo depth camera, taking only one row of the camera’s readings. The simulated gripper was a Robotiq 2F-85 parallel jaw<sup>7</sup>.




Fig. 7: Four grasp configurations for FLASK at the lab




Fig. 8: Evaluation objects

We used five objects for the simulation and lab experiments (Figure 8). Object meshes are based on the Dex-Net dataset.

For each object, we sampled a set  $\tilde{\mathcal{P}} \subseteq \mathcal{P}$  of six stable poses and considered four possible grasps (see Figure 7). We empirically evaluated the grasp score by recording the grasp success rate for the set of pose-grasp pairs for each object in both simulated and lab settings. The full details and results can be found in the appendix (complete results as well as our implementation, are in our online appendix).

a) *Evaluating the Predicted Sensor Function  $\hat{O}$* : For both sensors, synthetic images were rendered using the pyrender library [42] which involved projecting the 3D model of an object (transformed into a specific pose), onto a 2D plane using the sensor’s intrinsic and extrinsic parameters. The rendered images were computed by considering for each pixel the closest intersection between the simulated ray and the object meshes for the relevant field of view (FoV). For the lidar,  $\hat{O}$  included only one row, while for the depth

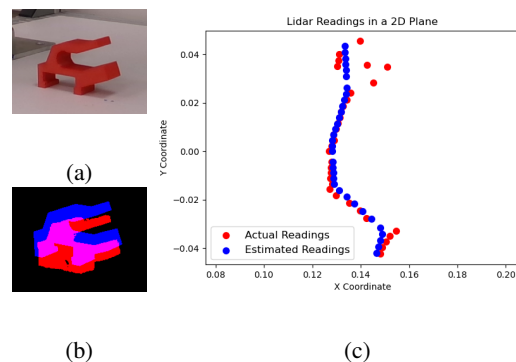


Fig. 9: Comparing the predicted (blue) and the lab-recorded image (red) for HOLDER. (a) actual scene (b) depth camera (c) 2D representation of the lidar reading.

camera, we considered the complete rendered image.

For each object pose  $p \in \tilde{\mathcal{P}}$ , we recorded the actual observation  $o$  and the predicted observation  $\hat{o}$ . For the lidar, we used four sensor configurations corresponding to the four cardinal directions. Noteably, while in the lab the lidar was able to capture all objects, in simulation MOUSE and MARKER, could not be captured.

For the depth camera, the set  $\tilde{\mathcal{X}}$  of sensor configurations was generated by randomly sampling 10K robot configurations  $q \in (-\pi, \pi]^6$  and translating them into camera poses using forward kinematics i.e.  $x = FK(q)$ . Each sensor configuration  $x \in \tilde{\mathcal{X}}$  was filtered using a heuristic:

$$H(x) = \left(1 - \frac{D(x)}{D_{max}}\right) + V(x) \quad (5)$$

where  $D(x)$  is the Euclidean distance between the camera’s center and the point of interest (PoI), representing the mean landing position after the object is dropped,  $D_{max}$  is a maximum acceptable distance used for normalization, and  $V(x)$  is a visibility score, which assesses how centered the PoI is within the camera’s FoV and is computed as the distance between the projection of the PoI onto the image plane and the center of the image divided by  $R_{ref}$ , a reference radius within the image plane that represents the boundary of acceptability. The set of six sensor configurations we examined was selected from the filtered set.

To measure the accuracy of  $\hat{\mathcal{O}}$  for the lidar, for each sensor configuration-object pose pair, we recorded the mean error of the difference between the measured and predicted readings. Evaluation of  $\hat{\mathcal{O}}$  for the depth camera, compared the predicted and actual images using the Intersection Over Union (IoU) measure: we pre-processed the RGB images to extract the object masks and computed IoU between the actual and synthetic images (see Figure 9).

**Results:** Table I presents results per sensor for HOLDER (results for all objects are in our online appendix). For each sensor configuration  $I - IV$  of the lidar, the table shows the average, minimal and maximal error (Avg. Err., Min. Err. and Max. Err., respectively) in mm over object poses  $\tilde{\mathcal{P}}$ . Similarly, for the depth camera, we computed the average, maximal, and minimal IoU values for configurations  $I - VI$ .

Results for the lidar show that the prediction errors (in mm) are negligible given the dimensions of the objects examined. In contrast, for the depth camera, errors (in  $[0, 1]$ ) are more substantial with a maximal average of 0.6. At the same time, results show varying performance across different configurations, depending on the object and its pose. For example, configuration  $I$  gives high accuracy for some objects, while configuration  $IV$  excels for others.

Inconsistencies between the observations are the result of several factors including the noise of the sensor itself, inconsistent scaling of the meshes with regard to the real objects, mismatches between objects and the meshes used for estimation, and inaccuracies in the placement of the objects in the lab (while the observation prediction is based on perfect object positioning). In addition, as the distance between the sensor and the object increases, the object

Lidar [mm]				Depth Camera [ratio]			
$x$	Avg. Err.	Min. Err.	Max. Err.	$x$	Avg. IoU	Max IoU	Min. IoU
$I$	1.6	0.7	3.2	$I$	0.6	0.8	0.4
$II$	3.9	0.7	4.8	$II$	0.5	0.7	0.4
$III$	0.8	0.5	1.0	$III$	0.6	0.7	0.5
$IV$	3.0	0.6	4.7	$IV$	0.6	0.7	0.6
-	-	-	-	$V$	0.4	0.6	0.3
-	-	-	-	$VI$	0.3	0.5	0.1

TABLE I: Sensor prediction for the lab setting of HOLDER. For the lidar, lower values are better while for the depth camera higher values are better.

		$\delta$	$\hat{q}$	$\mathcal{A}$
<i>HOLDER</i>	$\tau_1$	<b>0.19</b>	0.23	0.23
	$\tau_2$	0.15	<b>0.23</b>	<b>0.29</b>
	$\tau_3$	0.2	<b>0.26</b>	0.28
	$\tau_4$	0.03	0.15	0.13
	$\tau_5$	0.0	0.0	0.0
	$\tau_6$	0.0	0.0	0.0
<i>EXPO</i>	$\tau_1$	<b>0.29</b>	0.31	<b>0.29</b>
	$\tau_2$	<b>0.29</b>	0.31	<b>0.29</b>
	$\tau_3$	<b>0.29</b>	0.31	<b>0.29</b>
	$\tau_4$	0.04	<b>0.37</b>	<b>0.29</b>
	$\tau_5$	0.00	0.00	0.00
	$\tau_6$	0.00	0.29	<b>0.29</b>
<i>MOUSE</i>	$\tau_1$	0.00	0.06	0.04
	$\tau_2$	0.00	0.05	0.04
	$\tau_3$	0.00	0.07	0.00
	$\tau_4$	<b>0.04</b>	<b>0.44</b>	<b>0.53</b>
	$\tau_5$	0.00	0.00	0.00
	$\tau_6$	0.00	0.00	0.00
<i>CUP</i>	$\tau_1$	0.30	0.33	0.30
	$\tau_2$	0.30	0.33	<b>0.41</b>
	$\tau_3$	<b>0.37</b>	<b>0.63</b>	0.37
	$\tau_4$	0.24	0.38	0.27
	$\tau_5$	0.00	0.00	0.00
	$\tau_6$	0.00	0.00	0.00
<i>FLASK</i>	$\tau_1$	<b>0.25</b>	<b>0.25</b>	<b>0.25</b>
	$\tau_2$	<b>0.25</b>	<b>0.25</b>	<b>0.25</b>
	$\tau_3$	<b>0.25</b>	<b>0.25</b>	<b>0.25</b>
	$\tau_4$	0.02	<b>0.25</b>	<b>0.25</b>
	$\tau_5$	0.00	0.00	0.00
	$\tau_6$	0.00	0.00	0.00
AVG	$\tau_1$	0.19	0.23	0.23
	$\tau_2$	0.19	0.23	0.29
	$\tau_3$	<b>0.20</b>	<b>0.26</b>	0.28
	$\tau_4$	0.08	0.22	<b>0.34</b>
	$\tau_5$	0.00	0.00	0.00
	$\tau_6$	0.00	0.00	0.00

TABLE II: Results per belief update function (best results per criteria are highlighted)

occupies less of the sensor’s FoV and the readings include fewer and less informative data points. We also note that bright light reflections on shiny surfaces distort object shapes as perceived by the depth camera.

Figure 9 demonstrates inconsistencies between the predicted and actual observation of HOLDER at the lab. Here, this is due to the misplacement of the object. As we show next, the estimated observations are still useful for VOA computation despite these inconsistencies.

### B. Assessing VOA

In our experiments, both agents share the pose belief. We compare 3 different belief update functions to compute VOA. Each update function uses a different similarity metric  $\tau_i$  to compare the predicted observations that would be collected

for the different object poses. For the lidar,  $\tau_1$  deterministically considers two observations  $o_i, o_j$  as equivalent if for all angles, the values are within a margin of 8 mm.  $\tau_2$  uses the similarity metric  $\omega(o_i, o_j) = e^{-\|o_i - o_j\|}$  to update the belief based on Equation 1.  $\tau_3$  uses a multidimensional Gaussian where one observation defines the expectation vector, the covariance matrix is the identity matrix, and the similarity is the value of the PDF at the point corresponding to the second observation. For the depth camera,  $\tau_4$  is based on the structure element of SSIM [39],  $\tau_5$  is based on IoU between the two observations, and  $\tau_6$  employs the cv2 library [43] for contour matching such that similarity between two masks is computed by detecting their primary contours and comparing their shapes through a shape-matching algorithm.

The computation of VOA uses the similarity metric to compare the predicted observations of each examined pair of poses. In principle, lower similarity values indicate a greater ability to distinguish between the poses, and thus an ability to select a better grasp. However, as we demonstrate in Figure 6, even if the distance between the predicted images for different stable poses is small, VOA may still be high if these poses share a grasp configuration with a high score.

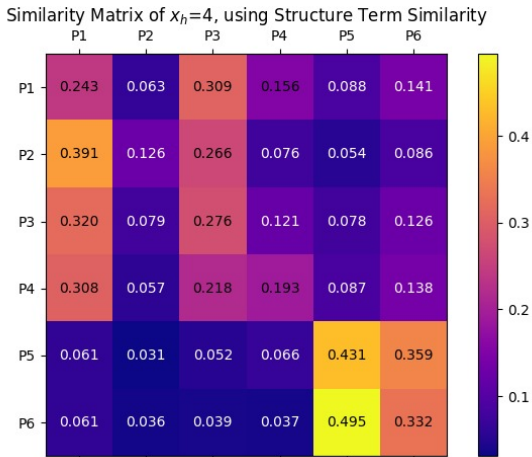


Fig. 10: Similarity score between actual observation (rows) for each pose and predicted observations (columns) for  $\tau_4$  (see the poses in Figure 6). The matrix shows a clear distinction between the standing positions P5 and P6 and the laying positions P1-P4, but the distinction within these two groups is ambiguous.

**Results:** For each experiment we consider 3 grasps as shown in Figure 5: an optimal grasp  $g^*$ , the initially chosen grasp  $g_i$ , and the grasp  $g_x$  chosen after collecting an observation from sensor configuration  $x$ , and compute the following:

- $\delta(x) = \mathbb{E}_{p \sim \beta} [\gamma(g_x, p) - \gamma(g_i, p)]$  is the weighted difference between the score of  $g_x$  and  $g_i$ .
- $\hat{\delta} = \delta(\arg \max_{x \in \tilde{\mathcal{X}}} U_x^{VOA}(\beta))$  is the  $\delta$  value of the sensor configuration with the highest VOA.
- $\delta^* = \mathbb{E}_{p \sim \beta} [\gamma(g^*, p) - \gamma(g_i, p)]$  is the weighted difference between the grasp score of  $g^*$  and  $g_i$ .
- $q(x) = \frac{\delta(x)}{\delta^*}$  is the ratio between  $\delta(x)$  and  $\delta^*$ .

- $\hat{q} = q(\arg \max_{x \in \tilde{\mathcal{X}}} U_x^{VOA}(\beta))$  is the ratio for the sensor configuration with the highest VOA in  $\tilde{\mathcal{X}}$ .
- $\mathcal{A} = \hat{q} - E_{x \in \tilde{\mathcal{X}}} [q(x)]$  is the advantage of choosing the maximal VOA configuration defined as the difference between  $\hat{q}$  and the weighted average over all sensor configurations  $\tilde{\mathcal{X}}$ . This represents the difference between choosing a sensor configuration using VOA and choosing randomly.

Table II presents for each object and for each belief update function the average values of  $\hat{\delta}$ ,  $\hat{q}$ , and  $\mathcal{A}$ . Results show that selecting the sensor configuration with the highest VOA value yields better outcomes compared to choosing a random sensor configuration for all but  $\tau_5$  and  $\tau_6$ , and in all cases is not worse. In other words, using VOA for decision-making when selecting a sensor configuration from which to provide an observation leads to better grasp choices. The smallest benefit is for MOUSE for which the initial grasp is optimal for all poses except one for which the optimal grasp has only a slight advantage. This subtlety is captured only by  $\tau_4$ .

Notably, belief update functions  $\tau_5$  and  $\tau_6$  of the depth camera which rely on IoU and contour matching, respectively, did not perform well on average for any of the objects. We associate this with the fact that the examined objects are small relative to their distance from the sensor, which is something these update functions are sensitive to. In contrast,  $\tau_4$  of the depth camera, yielded the best overall performance.

Figure 10 presents for  $\tau_4$  the similarity scores between the actual (rows) and predicted (columns) observations. Although it is hard to differentiate between object poses that occupy the same area, there are grasps that can succeed with high probability for all poses within the two distinguishable groups, as depicted in Figure 6.

In contrast, for  $\tau_5$  and  $\tau_6$  the groups of poses that have similar successful grasps are not sufficiently distinguishable, thus the distances metrics are not effective for the considered setting, which is captured by the VOA values. The reason for these results is the high sensitivity of  $\tau_5$  to noise in the actual image, where even slight translation can dramatically impact prediction quality. A similar trend was observed for  $\tau_6$  where noise dramatically distorts contours. See the online appendix for the complete analysis of the results achieved for all measures.

## V. CONCLUSION

We introduced *Value of Assistance* (VOA) for grasping and showed in both simulation and real-world experiments how our VOA measures predict the effect an observation will have on performance and how it can be used to support the decision of which observation to perform in order to enhance the probability of success in the context of grasping tasks.

Future work will include optimization considerations of the agent performing the sensing and for integrating VOA in long-term and complex tasks. Another extension will consider multi-agent settings in which VOA can be used not only for choosing which assistive action to perform but also for choosing which agent to assist.

## NOTES

- <sup>1</sup> <https://berkeleyautomation.github.io/dex-net>  
<sup>2</sup> <https://www.universal-robots.com/products/ur5-robot/>  
<sup>3</sup> <https://onrobot.com/en/products/2fg7>  
<sup>4</sup> [www.robotis.us/360-laser-distance-sensor-lds-01-lidar](http://www.robotis.us/360-laser-distance-sensor-lds-01-lidar)  
<sup>5</sup> <https://onrobot.com/en/products/onrobot-eyes>  
<sup>6</sup> <https://mujoco.org/>  
<sup>7</sup> <https://robotiq.com/products/2f85-140-adaptive-robot-gripper>

## REFERENCES

- [1] O. Kroemer, S. Niekum, and G. Konidaris, "A review of robot learning for manipulation: challenges, representations, and algorithms," *J. Mach. Learn. Res.*, vol. 22, no. 1, jan 2021.
- [2] A. Bicchi and V. Kumar, "Robotic grasping and contact: A review," in *Proceedings of the 2000 IEEE International Conference on Robotics and Automation, (ICRA)*. IEEE, 2000.
- [3] A. Sahbani, S. El-Khoury, and P. Bidaud, "An overview of 3d object grasp synthesis algorithms," *Robotics Auton. Syst.*, 2012.
- [4] A. M. Hands, "Kinematic and force analysis of," *Journal of Mechanisms, Transmissions, and Automation in Design*, 1983.
- [5] D. Prattichizzo and J. C. Trinkle, "Grasping," in *Springer Handbook of Robotics*, B. Siciliano and O. Khatib, Eds. Springer, 2008.
- [6] F. T. Pokorny and D. Kragic, "Classical grasp quality evaluation: New algorithms and theory," in *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*. IEEE, Year, p. Page Range.
- [7] D. Berenson, S. Srinivasa, and J. Kuffner, "Task space regions: A framework for pose-constrained manipulation planning," *The International Journal of Robotics Research*, 2011.
- [8] J. Bohg, A. Morales, T. Asfour, and D. Kragic, "Data-driven grasp synthesis—a survey," *IEEE Transactions on Robotics*, vol. 30, no. 2, pp. 289–309, 2014.
- [9] R. Balasubramanian, L. Xu, P. D. Brook, J. R. Smith, and Y. Matsuoka, "Physical human interactive guidance: Identifying grasping principles from human-planned grasps," *IEEE Transactions on Robotics*, vol. 28, no. 4, pp. 899–910, 2012.
- [10] L. Pinto and A. Gupta, "Supersizing self-supervision: Learning to grasp from 50k tries and 700 robot hours," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2016.
- [11] S. Nahavandi, R. Alizadehsani, D. Nahavandi, C. P. Lim, K. Kelly, and F. Bello, "Machine learning meets advanced robotic manipulation," *Information Fusion*, vol. 105, p. 102221, 2024. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1566253523005377>
- [12] M. Khansari, D. Kappler, J. Luo, J. Bingham, and M. Kalakrishnan, "Action image representation: Learning scalable deep grasping policies with zero real world data," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, 2020, pp. 3597–3603.
- [13] R. A. Howard, "Information value theory," *IEEE Transactions on Systems Science and Cybernetics*, vol. 2, no. 1, pp. 22–26, 1966.
- [14] S. Russell and E. Wefald, "Principles of metareasoning," *Artificial Intelligence*, vol. 49, no. 1, pp. 361–395, 1991.
- [15] S. Zilberstein and V. Lesser, "Intelligent information gathering using decision models," Computer Science Department, University of Massachusetts Amherst, Tech. Rep. 96-35, 1996.
- [16] C. Stachniss, G. Grisetti, and W. Burgard, "Information gain-based exploration using rao-blackwellized particle filters," in *Robotics: Science and systems*, vol. 2, 2005, pp. 65–72.
- [17] C. Cai and S. Ferrari, "Information-driven sensor path planning by approximate cell decomposition," *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, vol. 39, no. 3, pp. 672–689, 2009.
- [18] A. Amuzig, D. Dovrat, and S. Keren, "Value of assistance for mobile agents," 2023.
- [19] J. Shin, S. Chang, J. Weaver, J. C. Isaacs, B. Fu, and S. Ferrari, "Informative multiview planning for underwater sensors," *IEEE Journal of Oceanic Engineering*, vol. 47, no. 3, pp. 780–798, 2022.
- [20] É. Pairet, J. D. Hernández, M. Lahijanian, and M. Carreras, "Uncertainty-based online mapping and motion planning for marine robotics guidance," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 2367–2374.
- [21] V. Chandrasekhar, W. K. Seah, Y. S. Choo, and H. V. Ee, "Localization in underwater sensor networks: survey and challenges," in *Proceedings of the ACM international workshop on Underwater networks*, 2006, pp. 33–40.
- [22] V. Indelman, L. Carlone, and F. Dellaert, "Planning in the continuous domain: A generalized belief space approach for autonomous navigation in unknown environments," *The International Journal of Robotics Research*, vol. 34, no. 7, pp. 849–882, 2015.
- [23] V. Indelman, S. Williams, M. Kaess, and F. Dellaert, "Information fusion in navigation systems via factor graph based incremental smoothing," *Robotics and Autonomous Systems*, vol. 61, no. 8, pp. 721–738, 2013. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S092188901300081X>
- [24] Y. Tian, Y. Chang, F. H. Arias, C. Nieto-Granda, J. P. How, and L. Carlone, "Kimera-multi: Robust, distributed, dense metric-semantic slam for multi-robot systems," *IEEE Transactions on Robotics*, vol. 38, no. 4, 2022.
- [25] S. Koenig and Y. Liu, "Sensor planning with non-linear utility functions," in *Recent Advances in AI Planning*, S. Biundo and M. Fox, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2000, pp. 265–277.
- [26] R. Becker, A. Carlin, V. Lesser, and S. Zilberstein, "Analyzing myopic approaches for multi-agent communication," *Computational Intelligence*, 2009.
- [27] R. Mirsky, W. Macke, A. Wang, H. Yedidsion, and P. Stone, "A penny for your thoughts: The value of communication in ad hoc teamwork," *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence*, pp. 254–260, July 2020.
- [28] A. T. Taylor, T. A. Berrueta, and T. D. Murphey, "Active learning in robotics: A review of control principles," *Mechatronics*, vol. 77, p. 102576, 2021. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0957415821000659>
- [29] R. Bajcsy, Y. Aloimonos, and J. K. Tsotsos, "Revisiting active perception," *Autonomous Robots*, vol. 42, pp. 177–196, 2018.
- [30] Q. V. Le, A. Saxena, and A. Y. Ng, "Active perception: Interactive manipulation for improving object detection," *Stanford University Journal*, 2008.
- [31] A. Fern, S. Natarajan, K. Judah, and P. Tadepalli, "A decision-theoretic model of assistance," *Journal of Artificial Intelligence Research*, vol. 50, pp. 71–104, 2014.
- [32] C. V. Goldman and S. Zilberstein, "Optimizing information exchange in cooperative multi-agent systems," in *Proceedings of the Second International Joint Conference on Autonomous Agents and Multiagent Systems*, ser. AAMAS '03. New York, NY, USA: Association for Computing Machinery, 2003, p. 137–144.
- [33] R. J. Marcotte, X. Wang, D. Mehta, and E. Olson, "Optimizing multi-robot communication under bandwidth constraints," *Autonomous Robots*, vol. 44, no. 1, pp. 43–55, Jan. 2020.
- [34] R. M. Murray, S. Sastry, and Z. Li, "A mathematical introduction to robotic manipulation," 1994.
- [35] J. Mahler, J. Liang, S. Niyaz, M. Laskey, R. Doan, X. Liu, J. A. Ojea, and K. Goldberg, "Dex-net 2.0: Deep learning to plan robust grasps with synthetic point clouds and analytic grasp metrics," 2017.
- [36] K. V. Mardia and P. E. Jupp, *Directional statistics*. John Wiley & Sons, 2009.
- [37] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics*, ser. Intelligent Robotics and Autonomous Agents series. MIT Press, 2005. [Online]. Available: <https://dl.acm.org/doi/10.5555/1121596>
- [38] M. Lauri, D. Hsu, and J. Pajarinen, "Partially observable markov decision processes in robotics: A survey," *IEEE Transactions on Robotics*, vol. 39, no. 1, pp. 21–40, 2022.
- [39] Z. Wang, A. Bovik, H. Sheikh, and E. Simoncelli, "Image quality assessment: From error visibility to structural similarity," *Image Processing, IEEE Transactions on*, vol. 13, pp. 600 – 612, 05 2004.
- [40] H. Kurniawati, "Partially observable markov decision processes and robotics," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 5, pp. 253–277, 2022.
- [41] P. Daniel, "Mujoco.rl.ur5," <https://github.com/PaulDanielML/MuJoCo.RL.UR5>, 2022.
- [42] mmatl, "Pyrender," 2019, python 3D rendering toolkit. [Online]. Available: <https://github.com/mmatl/pyrender>
- [43] G. Bradski, "The OpenCV Library," *Dr. Dobb's Journal of Software Tools*, 2000.