Grasp selection for simple manipulation tasks

Ana C. Huamán Quispe

***Abstract*—Manipulation tasks are sequential in nature. Depending on the specific task constraints, more than one solution usually exist. Even for simple objects, several candidate grasps can be considered, each of which can potentially be used to plan the required arm motions. Grasp selection approaches that take into account the constraints at each task step are critical, since they allow to both identify grasps that are more likely to produce feasible arm motions and analogously, to discard grasps that might not be executable due to future task constraints. In this abstract, we present a manipulation metric for grasp selection based on a combined arm-and-grasp measure evaluated at each task step. We analyze 3 simple tasks (pick-up, pick-and-place and pouring tasks) and show the advantages of using our metric: (1) Shorter end-effector displacements and (2) Higher planning success rates. We present quantitative results in simulation and validate our approach's practicality with experimental results in our physical robot platform.**

I. Introduction

Given a manipulation task and a target object, many possible candidate grasps can be used to accomplish the task. Finding a suitable grasp among the infinite set of candidates is a challenging problem that has been addressed frequently in robotics, resulting in an abundance of approaches [2]. Interestingly, the vast majority of these methods have two aspects in common: (1) The metrics used for grasp selection focus on the hand-centric aspect of a manipulation task, such as grasp robustness measured with either analytical [3] or heuristic measures [1]. (2) Manipulation is implicitly seen as a single-step task, in which the main goal is to reach an object without further regard to what will be done with it once it is grasped.

In general, even the simplest of manipulation tasks, such as pick-and-place, have 2 or more steps. And while grasp robustness is perhaps the most important aspect of the manipulation task, it is not the only factor to consider. For a grasp to be executable in a task, feasible arm motions are needed. We argue that a metric that considers both the grasp robustness and arm kinematics (dependent on the task constraints) is a more useful way to select grasps that will in turn produce arm motions that can be planned without the need of testing several grasps, are fast to calculate, and produce short end-effector displacements. In this abstract we present a simple manipulation metric that combines arm and grasp measures and we show its direct applicability to select grasps in 3 standard, simple manipulation tasks.

II. Arm + Grasp Metric

A. Arm Metric (m_a)

When humans perform simple reaching actions, they select a grasp such that their arm is comfortable at the end of the reaching movement. This inherently simple phenomenon,

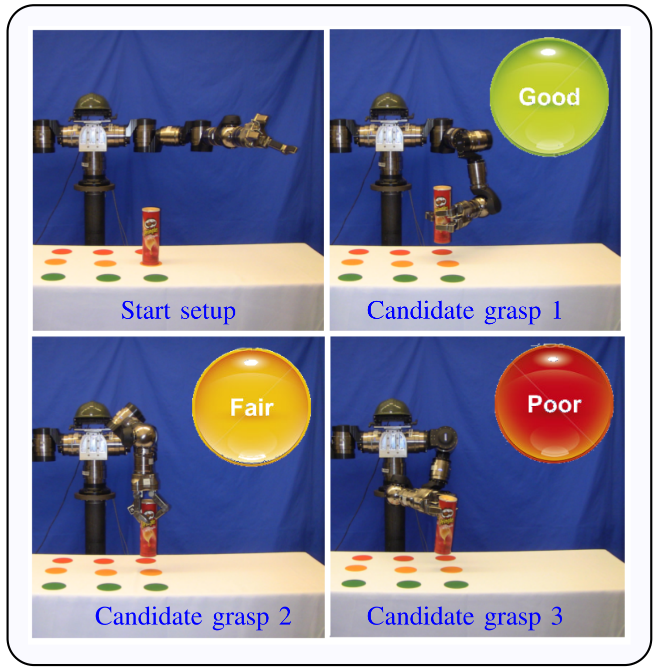

Fig. 1: Grasp selection application based on our arm-grasp metric for pick-up tasks

known as the *end comfort effect*, has been observed in adult humans as well as in primates [6].

Our proposed arm-centric metric intends to capture the comfort factor for a given grasp. Formally, for a given grasp $\mathbf{g}$ applied on an object located at ${}^{w}T_{o}$ we define our arm metric as the number of collision-free inverse kinematic solutions that allow the hand to execute $\mathbf{g}$.

$$m_a(\mathbf{g}) = |\mathcal{Q}| \text{ such that } \forall \mathbf{q_i} \in \mathcal{Q} \begin{cases} \mathbf{q_i} & \text{is collision-free} \\ FK(\mathbf{q_i}).(\mathbf{g}.{}^{h}T_o) = {}^{w}T_o \end{cases}$$

For our specific setup, the redundant robot arm presents a standard S-R-S configuration for which a pseudo-analytic solution is available [7] given as an input an end-effector's goal pose and a parameter $\phi \in [-\pi, \pi]$ which determines the elbow pose. In the equation above, the initial set of inverse kinematic solutions are calculated by discretizing ϕ and evaluating which of them are collision-free ($\mathbf{q_i} \in \mathcal{Q}$).

B. Grasp Metric (m_g)

The arm-centric metric presented above only considers the arm comfort. Consider the scenario in Figure 2, where 3 candidate grasps are depicted for a cylindrical object (a grasp here being parameterized by two elements: (1) The relative rigid transform of the end-effector frame with respect to the object frame, and (2) The finger's initial joint configurations). Let us assume that these grasps have similar m_a values, hence they are all deemed equally desirable. From human experience, we can all agree that the second grasp is the most likely to be stable since the hand is closer to the center of mass of the object being held. We incorporate this heuristic on the proposed grasp metric.

Our second metric attempts to favor grasps that hold the object near its center of gravity. We propose to quantify this heuristic as the distance between the object's center of mass and the hand's approach direction vector. We select this metric because it is easy to calculate, as it is just the distance between a line and a point. This metric is similar to the existing metric B_1[5], which measures the distance between the center of the contact polygon and the center of mass of the object. We prefer our metric over B_1 mainly because our system does not provide finger contact information.

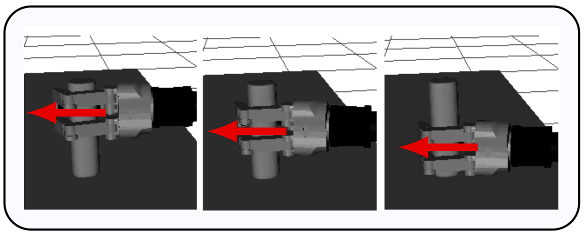

Fig. 2: Examples of similar grasps with a different distance from the hand approach direction (red arrows) and the object center of mass

C. Arm-Grasp Metric (m_{ag})

Now that we have both metrics, we must combine them. A direct way to do this could be using a weighted sum of both. However, both metrics have different units (m_a is adimensional and m_g has length units), hence adding them does not have a real meaning. Instead, we propose to calculate m_{ag} using 2 consecutive steps, each of them using one of the metrics for partial ordering: In the first step, m_a is used to obtain a partial ordering of the grasps, whereas in the second step, m_g is used to reorder within each partial subgroup. This can be explained in simple terms as:

1) Calculate the mean μ_a and the standard deviation σ_a of the arm metric (m_a) over all the candidate grasps ($\mathcal{G}$).
2) Divide the grasps in 4 groups, similarly as [5]:
   a) Very good m_a quality: $m_a(\mathbf{g}_i) > \mu_a + \sigma_a$
   b) Good m_a quality: $\mu < m_a(\mathbf{g}_i) < \mu_a + \sigma_a$
   c) Fair m_a quality: $\mu_a - \sigma_a < m_a(\mathbf{g}_i) < \mu_a$
   d) Bad m_a quality: $m_a(\mathbf{g}_i) < \mu_a - \sigma_a$
3) Within each of the 4 groups, reorder the grasps according to their grasp metric m_g.
4) The final ordered set of grasps will contain 4 m_a-based ordered sets (very good, good, fair and bad), inside each of which grasps are ordered according to m_g.

III. Application to Simple Manipulation Tasks

We apply our proposed metric for grasp selection in the 3 simple manipulation tasks which are briefly described in the following subsections.

A. Pick-up tasks

The most basic of the 3 manipulation tasks considered. In this work this is modeled as a single-step task in which the proposed metric m_{ag} is measured considering the grasps being executed at the object start pose. Table I shows simulation results comparing the success rates and hand displacement corresponding to arm motions planned by selecting grasps according to m_{ag} (best and worst). We observe that by using our metric, our system is able to select grasps that demand substantially less planning time than in the worst case. Additionally, the success rate of the arm planning is considerably higher for the grasps deemed best. Some snapshots showing the physical robot using the metric to perform simple pick-up tasks is shown in Figure 3.

TABLE I: Simulation results of 100 randomized scenarios per each object

Object	Success		Plan time(s)	
	Best	Worst	Best	Worst
1. Pringles	84%	52%	0.88	4.62
2. Cheezit	91%	31%	1.31	1.33
3. Coffee	100%	20%	0.88	1.70
4. White cup	84%	54%	0.88	5.46
5. Ball	100%	12%	0.77	10.37
6. Plushie	100%	3%	0.83	0.84
7. Raisins	99%	47%	0.83	6.64
8. White cleaner	99%	35%	0.93	1.84
9. Yellow cone	100%	46%	0.83	2.65

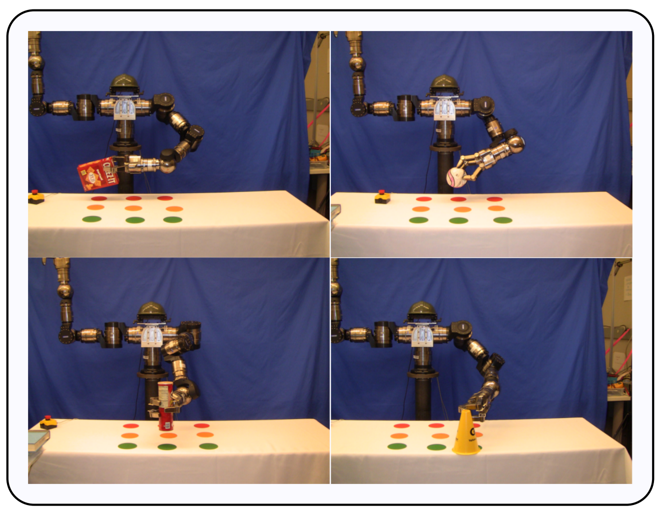

Fig. 3: Pick-up experiments using m_{ag} for grasp selection

B. Pick-and-place tasks

In this case, the task evaluated has two steps (reach and transport), hence it is not clear in which step it should be evaluated. In order to determine how to best use the metric in this case, we evaluate it under 3 modalities: (1) At the start pose of the object, (2) At the goal pose, and (3) Using an average of the arm metric both at the start and pose. From experience, we have observed that by considering mode (1), the reach arm path tends to be shorter. Analogously, by using mode (2), the transport path is shorter. The average measuring attempts to consider both the start and goal to achieve an intermediate comfort during both reach and transport phases. We performed randomized simulation experiments and the results are shown in Table II. We observe that the mode (3) - average - produces the highest success results. Figure 4 shows sample executions of pick-and-place tasks putting objects inside a box using our average metric for grasp selection.

TABLE II: Simulation results of randomized on-table pick-and-place scenarios (250) per each object

Object	Success		
	Goal	Start	Avg
Master Chef	171/188	150/188	182/188
	79/188	146/188	164/188
Green plushie	108/110	89/110	109/110
	63/110	97/110	106/110
Pringles	234/250	232/250	243/250
	149/250	165/250	180/250
Soft Scrub	235/250	244/250	247/250
	175/250	170/250	200/250
Sun maid	195/250	220/250	230/250
	104/250	194/250	225/250
Yellow cone	229/250	151/250	215/250
	87/250	182/250	199/250

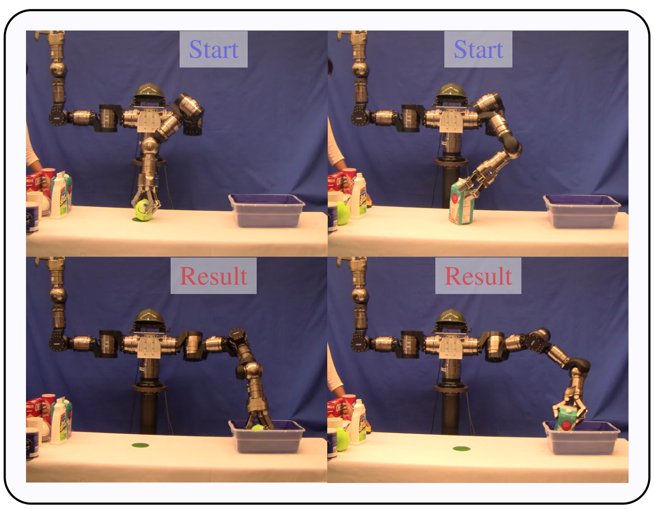

Fig. 4: Pick-and-place experiments using a grasp selected by evaluating the average of m_{ag} measured at the start and goal object poses.

C. Pouring tasks

For the final task considered - pouring task - we model it with 3 canonical steps: Reach, transport and tilt. Since tilt is mainly a rotation in place, we evaluate the m_{ag} in a similar manner to the pick-and-place case, only considering the reach and transport phases. Surprisingly, for this case we found (Table III) that the best results in terms of success rate and end-effector displacement were produced by evaluating the metric only at the start pose, without further regard to the goal pose. We hypothetize that this is due to the fact that for pouring tasks, the goal pose for the *arm* is mostly invariant, depending only on the pose of the container. Given this, the goal configuration does not play a factor in the grasp selection. Some snapshots showing the robot performing the task using m_{ag} to select the grasp executed are shown in Figure 5.

TABLE III: Simulation results of randomized scenarios. Container object: Red cup

Object	Success			Hand Disp.(m)		
	Goal	Start	Avg	Goal	Start	Avg
Pringles	221/250	250/250	240/250	2.20	2.14	2.18
	187/250	184/250	221/250	2.37	2.44	2.37
White cup	230/250	250/250	237/250	2.18	2.10	2.13
	204/250	205/250	218/250	2.24	2.30	2.29
Soft Scrub	234/250	250/250	243/250	2.22	2.10	2.16
	200/250	200/250	228/250	2.43	2.42	2.42

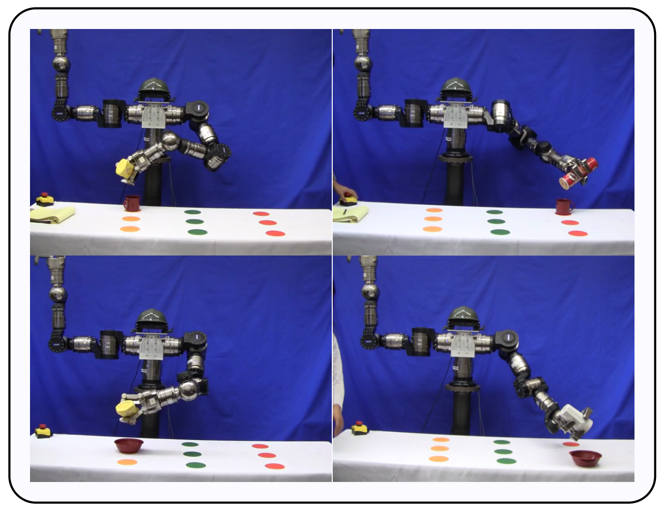

Fig. 5: Pouring experiments using m_{ag} for grasp selection evaluated at the object start pose

IV. Conclusions

In this abstract we presented a grasp selection approach consisting on measuring our proposed manipulation metric (m_{ag}) at each differentiated step in the manipulation task and select the grasp that presents the highest value for a measure based on this metric. For one-step tasks, such as pick-up, m_{ag} was directly used as the selection measurement, whereas for two-step tasks, such as pick-and-place (pouring), we observed that the m_{ag} measure calculated at the initial step (as an average of the start and goal step) were more efficient, based on our simulated experiments.

As future work, more challenging scenarios can be faced when the task at hand has n steps (with $n > 2$), as there exist many more possible ways in which the metric m_{ag} measured at each step $i \in [1, n]$ can potentially be used. There are a few studies in psychology that strongly suggest that each step in an action sequence influences the choice of grasp in humans [4], although the exact influence of each step (or how this influence takes place) has not been yet defined.

REFERENCES

[1] R. Balasubramanian, L. Xu, P. Brook, J. Smith, and Y. Matsuoka. Physical human interactive guidance: Identifying grasping principles from human-planned grasps. *IEEE Transactions on Robotics*, 28(4):899–910, 2012.

[2] J. Bohg, A. Morales, T. Asfour, and D. Kragic. Data-driven grasp synthesisa survey. *IEEE Transactions on Robotics*, 30(2):289–309, 2014.

[3] C. Ferrari and J. Canny. Planning optimal grasps. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 2290–2295. IEEE, 1992.

[4] Patrick Haggard. Planning of action sequences. *Acta Psychologica*, 99(2):201 – 215, 1998. ISSN 0001-6918. doi: http://dx.doi.org/10.1016/S0001-6918(98)00011-0. URL http://www.sciencedirect.com/science/article/pii/S0001691898000110.

[5] B. León, C. Rubert, J. Sancho-Bru, and A. Morales. Characterization of grasp quality measures for evaluating robotic hands prehension. In *International Conference on Robotics and Automation (ICRA)*, pages 3688–3693. IEEE, 2014.

[6] D. A. Rosenbaum, C. M. van Heugten, and G. E. Caldwell. From cognition to biomechanics and back: The end-state comfort effect and the middle-is-faster effect. *Acta psychologica*, 94(1):59–85, 1996.

[7] Masayuki Shimizu, Hiromu Kakuya, Woo-Keun Yoon, Kosei Kitagaki, and Kazuhiro Kosuge. Analytical inverse kinematic computation for 7-DOF redundant manipulators with joint limits and its application to redundancy resolution. *Transactions on Robotics*, 24(5):1131–1142, 2008.