Mitigating Uncertainty by Learning to Grasp Under Blindness

Philipp Zech and Justus Piater
Department of Computer Science, University of Innsbruck, Austria

I. Introduction

Over the last two decades we have seen a tremendous change in robotics from simple workhorses to complex, inter- and reactive autonomous agents. However, this move from a *closed* to an *open* world comes at the cost of both uncertainty in the robot's perception of and its belief about the world. Further, and probably even more serious, this uncertainty obstructs learning severely due to missing or incomplete information. Hence, for robotics to advance and evolve further effectively mitigating uncertainty is paramount.

To succeed in an open world, we believe that robust grasping is a key factor due to its importance in shaping the world to make it maneuverable. Unfortunately, one of the main drawbacks with current research in grasping is the prevalent reliability on 2D or 3D object models [2], [7]. Obviously, in an open, and thus, noisy world suitable object models are seldom available, thence rendering futile grasping approaches that rely on them.

Motivated by the remarkable object interaction and manipulation capabilities of blind people, we propose a novel approach to grasp learning inspired by blindness that (i) does not rely on an object model and (ii) effectively mitigates uncertainty by only relying on salient and robust object features. Blind people essentially learn to grasp novel objects by (i) being manually guided by a teacher and (ii) fostering their capabilities by playing around with the object given their prior. After a sufficient amount of training this finally results in a cognitive representation of the spatial relationship between and object's pose and the hand's pose for grasp planning.

In robotics, such a learning scheme can be imitated by trial and error learning of grasp affordance densities, bootstrapped by grasps learned from Programming by Demonstration (PbD). Grasp affordance densities are a popular tool for modeling the spatial relationship of successful grasps by relating gripper poses to object poses using probability distributions. The additional biasing effect that results from bootstrapping the learning phase accelerates the learner and guides its search for successful grasps towards interesting regions on an object's surface. Finally, given an object's pose ξ, a successful grasp is readily available by maximizing the object's grasp affordance density π for some grasp g,

$$g^* = \arg\max_{g \in \mathcal{G}} \pi_\xi(g),$$

where $\mathcal{G}$ denotes the set of all feasible grasps with $\mathcal{G} \subset \mathbb{R}^3 \times SO(3)$ and π_ξ the object's grasp affordance density aligned to pose ξ.

II. Contribution

The goal of our work is to study the design of efficient active grasp learning methods that are robust under uncertainty. This poses two main challenges:

- Reducing the necessary amount of object geometric information to a minimum so as to rely only on essential, but at the same time highly resilient information.
- The construction of efficient anytime learning methods that by design treat the learning problem at hand as a black box.

The latter statement may sound confusing at first. However, observe that we aim at learning grasp affordance densities which essentially can be treated as learning of black box functions.

A. Reducing Object Geometric Information

Uncertainty generally grows if an agent lacks necessary information. One common strategy to mitigate such a lack is to learn missing links from available data. However, such an approach may not scale well and is also prone to overfitting. In our work, we take a different approach. Contrary to relying on precise object geometric information, we instead study two different ideas of *blind* grasp learning, bootstrapped with as little necessary information as possible. Obviously, the less information is required the more robust learning methods and the yielded knowledge become.

Our first learning method (Section II-B.1) is bootstrapped by a few demonstrated grasps g by their 6D poses, i.e., $g \in \mathcal{G}$ and a scalar quality indicator based on the Grasp Wrench Space (GWS) [5][1]. From these we then learn new grasps around

[1]We want to mention that we are well aware that the GWS is not and ideal grasp quality metric. However, we chose to use it as (i) we are primarily interested in learning of feasible grasps that allow to characterize an object's grasp affordance density, and (ii) at this point do not consider the notion of a task, which requires an alternate grasp quality metric [1].

the demonstrated object by active sampling. For our second learning method (Section II-B.2) we reduce this bootstrapping information even further to just a few randomly chosen points in $\mathbb{R}^3$ distributed on the object. Again, using active sampling we then learn good grasps at those points.

Clearly, there is a crucial difference between our two proposed learning schemes. The former performs global optimization, whereas the latter only searches for a locally optimal solution. Nevertheless, as discussed in the next section both our learning methods allow for global characterization of an object's grasp affordance density by requiring only very little prior information on the problem at hand. In both cases, we learn this prior information by guiding the robot's end effector either into a valid grasping pose or a good grasping point on the object, respectively. Especially in case of demonstrating a valid grasping pose, the use of PbD is motivated by Balasubramanian et al.'s recent work, where they showed that empirical grasp learning grounded upon PbD can achieve results superior to planner based, analytic methods [1].

B. Efficient Anytime Learning of Grasp Affordances

The algorithms we employ in our study for designing our learning methods are based on the famous Monte Carlo formulation, originally designed for approximating highly complex integrals (or probability densities) by repeated sampling from regions of interest. We study both the applicability of adaptive Metropolis Hastings (Section II-B.1) and Monte Carlo Tree Search (Section II-B.2).

1) Adaptive Metropolis-Hastings: We first study the applicability of Kernel-adaptive Metropolis Hastings (KAMH) [8], an adaptive Metropolis Hastings sampler, for global characterization of an object's grasp affordance density. The general scheme of KAMH is visualized in Figure 1a. The underlying idea of KAMH is to map the input space $\mathcal{X}$ into a Reproducing Kernel Hilbert Space (RKHS) $\mathcal{H}$ to render it Gaussian. Subsequently, new samples f are proposed in terms of sampling Gaussian processes, as an RKHS is a Hilbert space of functions. These samples f are then mapped back using a single gradient descent step along a cost function [8] and accepted or rejected given the Metropolis Hastings acceptance criterion. The intuition behind mapping the input space $\mathcal{X}$ into an RKHS is to make it easy to sample from. This idea makes KAMH a promising candidate for learning grasp affordance densities, as these generally are highly non-linear and thus difficult to sample from in their original (input) space.

The second intriguing feature of KAMH is its ability to learn a covariance operator $C_{\mathbf{z}}$ of its target. This is achieved by repeated subsampling of the constructed Markov chain during the burn-in duration of KAMH. Observe that the chain is initialized with $\mathcal{X}$. To prevent the algorithm from diverging, the input space $\mathcal{X}$ is thus required to provide exploratory properties of the target. This is exactly where the demonstrated grasps come to aid. We define the input space $\mathcal{X}$ to be a subset of $\mathbb{R}^3 \times SO(3)$ that contains those demonstrated grasps, as they provide necessary exploratory properties of the grasp affordance density sought, i.e., $\mathcal{X} \subset \mathbb{R}^3 \times SO(3)$ and $\mathcal{X} \cap \mathcal{G} \neq \emptyset$. The reason why we can populate $\mathcal{X}$ by only a few demonstrated grasps and otherwise random points is KAMH's learning capability during its burn-in phase which capitalizes on the repeated subsampling of the Markov chain. This essentially trains the covariance operator $C_{\mathbf{z}}$ to express how far changes in the gripper pose affect grasp success in case of a specific object.

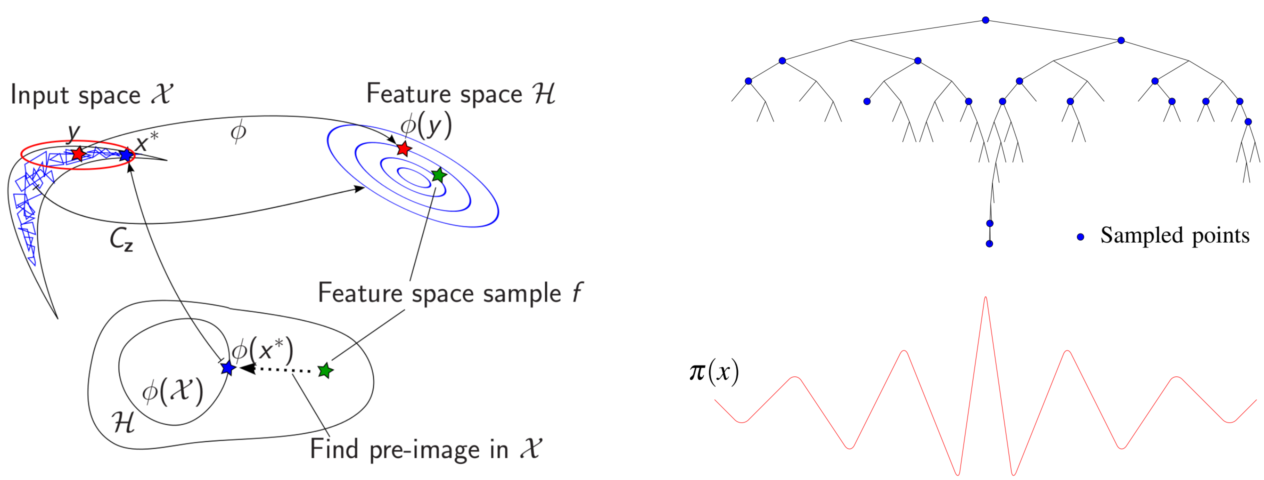

(a) Kernel-adaptive Metropolis Hastings [8] (b) Bayesian Multiscale Optimistic Optimization [9]

Fig. 1: Visualization of the Monte Carlo variants applied in this paper.

2) Monte Carlo Tree Search: The second algorithm whose applicability we study for learning an object's grasp affordance density is Bayesian Multiscale Optimistic Optimization (BaMSOO) [9], a variant of optimistic optimization [6] grounded upon Monte Carlo Tree Search (MCTS) [4]. The general idea of this class of optimization algorithms, as depicted in Figure 1b, is to create a search tree by evenly splitting the search space $\mathcal{X}$ along the centers of each of its dimensions, evaluate the target

$\pi(x)$ at the centers of the resulting subsets (the leaves) and then – in an optimistic manner – dig further into those subsets that may contain the global optimizer x^*. In case of BaMSOO, this optimistic choice is mathematically formulated by computing, at each iteration, upper confidence bounds for each new subset that reflect the estimated optimizer x' contained in a subset. Consequently, BaMSOO only looks further into those subsets whose estimated optimizer x' is better than the current optimizer x^*. This intuitive and effective mechanism both drastically reduces the search space and speeds up learning the optimizer x^* of $\pi(x)$.

In the face of grasp learning where we define the search space $\mathcal{X} \subset \mathbb{R}^3 \times SO(3)$, this mechanism of tree construction however is not applicable, since $SO(3)$ lacks a distinct notion of origin relative to which the search space could be split. Not to mention further that these are also different metric spaces. Thus, we relax our problem to only optimize a gripper's orientation in $SO(3)$ at a demonstrated point p on the object. We do so by mapping $SO(3)$ onto the unit quaternion manifold $\mathcal{M}_{\mathbb{H}}$ which represents the unit hypersphere $\mathcal{S}^3$ embedded in $\mathbb{R}^4$. The key advantage of this mapping is $\mathcal{S}^3$'s property of being parametrizable by three angles that live inside specific intervals along which we can safely split[2] to search for an optimal orientation of the gripper. As for KAMH, the scalar quality indicator of a grasp, that is, the value on which BaMSOO optimizes, is based on the GWS.

Observe that in both cases we can stop the algorithms at any time, guaranteeing to return the best grasp g found so far. Further, also both algorithms assure that the longer they run, the better the found grasps g become.

III. Preliminary Results

We performed a series of preliminary experiments to evaluate whether the application of the suggested algorithms (Section II-B) is feasible under the assumption that we reduce available information on the problem at hand to a minimum (Section II-A) and points into a promising direction for future research. Figure 2 shows the results of these experiments.

All in all we evaluated both algorithms on 6 objects, where KAMH was run for the objects from Figure 2a and BaMSOO for the objects from Figure 2b. We chose default parameters for both algorithms as suggested in the original works [8], [9]. The number of iterations for KAMH was set to 200+2000, where the former denotes the burn-in duration, and for MCTS to 200. Observe that KAMH requires this burn-in duration to learn an informative covariance operator $C_{\mathbf{z}}$. After the burn-in, $C_{\mathbf{z}}$ remains unchanged to assure convergence. We ran both algorithms once for each of the corresponding objects.

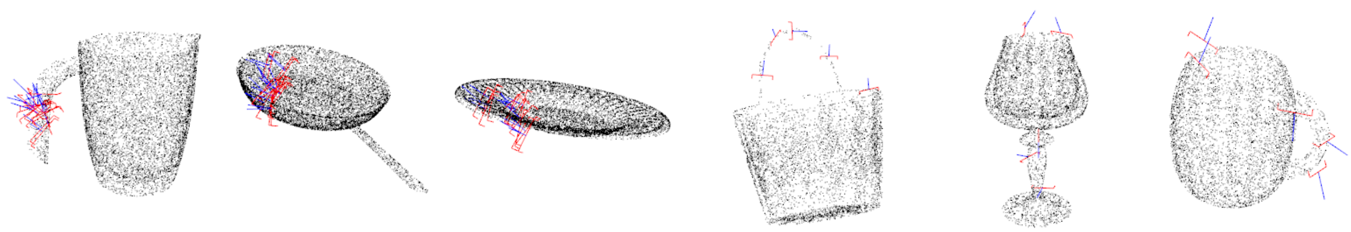

(a) Kernel-adaptive Metropolis Hastings

(b) Monte Carlo Tree Search

Fig. 2: Preliminary results for both our learning models (best viewed in color).

It is obvious from Figure 2 that the learning schemes sketched in this study are well feasible. In both cases, the reformulation of the search space in terms of learning valid gripper poses and orientations, respectively, seems to be promising for the application of Monte Carlo based optimization methods. However, there is also some clear discrepancy between the grasps learned by both algorithms. As stated earlier (Section II-B.1), KAMH aims at globally characterizing an object's grasp affordance density. It is thus obvious that KAMH yields a large number of grasps, densely spread over the object, whose quality increases over time. The high density of grasps results from the circumstance that the step size of Metropolis Hastings samplers generally is very small. Under the bottom line, this means that at the cost of longer run times and accepting gradually better samples we can learn the complete grasp affordance density for a specific object by only requiring a few user-demonstrated grasps. We argue that learned grasps are then both (i) easy to recall and apply and (ii) robust under uncertainty as they only rely on an object's pose. Figure 2a shows, for each of the objects under study, a few successfully learned grasps for KAMH. Observe that they are dense at the depicted regions.

Contrary to KAMH, the results achieved for BaMSOO look quite sparse (Figure 2b). However, this is a result of our reformulation of the search space $\mathcal{X}$ which is the space of 3D rotations, i.e., $SO(3)$. Thus, the algorithm cannot move around the object but instead is forced to remain at the demonstrated points. Luckily, this locality comes at the benefit that the grasps learned using BaMSOO are of substantially higher quality as early grasps found using KAMH. This is evident from Figure 2b where it can be seen that specific grasps found for each of the objects appear very natural, e.g., the grasps learned at the neck of the goblet.

[2] As a side note, observe that this kind of tree construction also yields an even uniform discretization of $SO(3)$.

It is also noticeable that for our study we did not tune any of the algorithms. As stated earlier, we used default parameters as mentioned in the original works. This further underpins the fitness of the selected algorithms, as they, without any fine-tuning, already yield promising results.

To summarize, in this study we investigated the feasibility of grasp learning using only very little information under the assumption of uncertainty. Contrary to most existing work that relies on 2D, 2.5D or 3D object models [2], [7] we have shown that grasp learning bootstrapped only by a few user demonstrated grasps or points and no further object geometric information is feasible. We have found that the application of Monte Carlo methods seems to be promising for such a task. This essentially capitalizes on the circumstance that these methods are designed to find solutions to complex problems by approximating them while learning from the past. We thus argue that our learning schemes are both robust and efficient.

IV. Conclusions and Future Work

In this study we have presented a promising, new direction for grasp learning in robotics motivated by the remarkable learning and interaction capabilities of people with blindness. Our results show that both our grasp affordance learning methods are feasible given only very little information on the problem at hand. This shows empirically that for grasp learning it is possible to mitigate uncertainty by reducing necessary information to only very salient and robust object features, e.g., a pose. However there still remain open issues to be addressed in future work if our ideas shall grow into mature grasp learning methods.

First of all, the current implementation of KAMH lacks a valid geometrical interpretation of the search space. However, we intend to look into dual quaternions for representing gripper poses as they provide a unified view on 6D poses like $SE(4)$, yet are easier to handle compared to elements of $SO(4)$. Further, in its current form, KAMH does not address the multimodlity of a density, which however is paramount in the event of grasp affordance densities. Clearly, such a density has multiple modes in terms of multiple ideal regions to grasp the object. Second, in order to globally characterize an object's grasp affordance density by BaMSOO, we essentially have to run an exhaustive search all over the object's surface. To speed this up we also want to study the applicability of dual quaternions, as this would allow us to not only optimize over the orientation anymore but instead the full 6D pose of the gripper. Clearly, this would come at the cost of a sparse grasp affordance density as BaMSOO would only look into specific subregions on the object, yet with the benefit that the resulting grasps are of high quality. Finally, another aspect to study in our learning methods is an alternate quality metric to the current one based on the GWS. For this, we will look into formalizations of the notion of a task for computing a grasp quality measure [3].

Acknowledgments

The research leading to these results has received funding from the European Community's Seventh Framework Programme FP7/2007–2013 (Specific Programme Cooperation, Theme 3, Information and Communication Technologies) under grant agreement no. 610532, SQUIRREL.

References

[1] 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*, 28(4):899–910, 2012.

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

[3] Ch. Borst, M. Fischer, and G. Hirzinger. Grasp Planning: How to Choose a Suitable Task Wrench Space. In *Robotics and Automation, 2004. Proceedings. ICRA'04. 2004 IEEE International Conference on*, volume 1, pages 319–325. IEEE, 2004.

[4] C. B. Browne, E. Powley, D. Whitehouse, S. M. Lucas, P.I. Cowling, P. Rohlfshagen, S. Tavener, D. Perez, S. Samothrakis, and S. Colton. A Survey of Monte-Carlo Tree Search Methods. *IEEE Transactions on Computational Intelligence and AI in Games*, 4(1):1–43, 2012.

[5] A. T. Miller and P. K. Allen. Examples of 3D Grasp Quality Computations. In *IEEE International Conference on Robotics and Automation, 1999. Proceedings, ICRA'99.*, volume 2, pages 1240–1246. IEEE, 1999.

[6] R. Munos. *From Bandits to Monte-Carlo Tree Search: The Optimistic Principle Applied to Optimization and Planning*. now publishers, Boston-Delft, 2014.

[7] A. Sahbani, S. El-Khoury, and P. Bidaud. An Overview of 3D Object Grasp Synthesis Algorithms. *Robotics and Autonomous Systems*, 60(3):326–336, 2012.

[8] D. Sejdinovic, H. Strathmann, M. L. Garcia, C. Andrieu, and A. Gretton. Kernel Adaptive Metropolis-Hastings. In Eric P. Xing and Tony Jebara, editors, *Proceedings of the 31st International Conference on Machine Learning*, volume 32, pages 1665–1673, 2014.

[9] Z. Wang, B. Shakibi, L. Jin, and N. de Freitas. Bayesian Multi-Scale Optimistic Optimization. In *Proceedings of the 17th International Conference on Artificial Intelligence and Statistics*, pages 1005–1004, 2014.