Probabilistic Approach to Sensor-based Grasping

Janne Laaksonen and Ville Kyrki

Abstract— **In this paper, we present a probabilistic framework for grasping. In the framework, we consider grasp and object attributes, on-line sensor information and the stability of a grasp, through probabilistic models. We describe how sensor-based grasp planning can be formulated in a probabilistic framework and how information about object attributes can be updated simultaneously using on-line sensor information gained during grasping. The feasibility and advantages of the framework are demonstrated in a 2D simulation environment, with simulated tactile sensors used to update object information. In the demonstration, particle filters are used to model the evolving probability distributions.**

I. INTRODUCTION

Current grasp planning approaches are usually based on an assumption of perfect knowledge of target objects. While geometric models are good approximations of the objects in the real world, the models are not exactly accurate, especially when speaking of household items. Thus, a difference between the expected and the realized grasp arises from these approximations, although in many cases the difference is small enough to achieve a stable grasp. However, this discrepancy is usually left unused.

On the other hand, methods using sensor information to grasp using corrective motions or reacting to the tactile sensor information have been proposed. Contrary to grasp planners, accurate object models are not usually available in this type of grasping. However, it has been shown that using for example tactile sensors it is possible to estimate stability of the grasp [1], the pose of the object [2], [3], or even the identity of the grasped object [4].

In this paper, we present a probabilistic framework, which unifies the ideas behind grasp planning and the possibilities of sensor-based grasping. The framework considers the required variables and models for grasping as probability distributions and allows thus the representation of the current belief probabilistically, that is, the uncertainty in the knowledge can be represented. The framework allows interplay between grasp planning and corrective motions, in situations where object attributes, such as pose, are not precisely known, by utilizing sensor information gained during grasping. Such a situation can arise for example when visual sensing is used to initially estimate the target objects. We demonstrate the framework with a simple 2D example. In the demonstration, we use particle filters, a MCMC (Monte Carlo Markov Chain) method, to estimate the evolving probability distributions. A Bayesian approach is used, that is, instead of using the maximum likelihood or maximum a posteriori solution, the result is obtained by marginalizing over the current knowledge. The main advantage of the proposed approach is that grasp planning can be performed with uncertainties in environment as well as measurements, for example, if the target object's pose is uncertain.

Section II collects the related work about grasp planning and other related fields and Section III describes the probabilistic framework. In Section IV, a practical implementation based on the probabilistic framework is presented where we simulate a gripper grasping a rectangle in 2D. We conclude with discussion and possibilities of the probabilistic framework, in addition to our focus of future work in Section V.

II. RELATED WORK

Our approach to find good grasps is closely related to the field of grasp planning. In grasp planning the goal is to find as good as possible grasp on a given object. The goodness of the grasp is usually measured with a grasp quality measure [5]. However, compared to our method, most current grasp planning methods do not account for the uncertainty present in the object or in the object's pose information. Also most of the grasp planning methods require a known geometric model of the object.

To simplify the grasp planning, many methods employ some form of object decomposition. The goal of the decomposition is to reduce the amount of feasible grasps without trying every grasp on an object. In [6], the object is decomposed to minimum volume bounding boxes, in an effort to understand the underlying shape of the object. The primitive shape is then used to reduce the search space for stable grasps. Instead of boxes, superquadrics are used in [7]. In addition to the construction of the superquadric decomposition, heuristic is used to define the trial grasps based on the superquadric form of the object limiting the space of grasps significantly.

The Columbia Grasp Database [8] takes a different approach to most grasp planners and computes best grasps for a set of hundreds of objects. The grasp planning problem is then transformed to a problem of matching a new object with an object found in the precomputed database of grasps. The work has also been extended to consider partial data [9].

If the object is not known, i.e. a geometric model is not available, the grasp planning methods can still be used if the model of the object can be constructed. The model construction can either be done by vision or tactile exploration.

The research leading to these results has received funding from the European Community's Seventh Framework Programme under grant agreement n° 215821.

J. Laaksonen and V. Kyrki are with Department of Information Technology, Lappeenranta University of Technology, P.O. Box 20, 53851 Lappeenranta, Finland, jalaakso@lut.fi, kyrki@lut.fi

However, the geometric model in this case is usually a mesh or a point cloud, and contains no information about the inherent uncertainty related to the perception. Approaches such as [6] can be applied here as well but the results can be worse than in the cases where the full geometric model is known and the decomposition may fail in cases where large volumes are missing from the perceived object. It is also possible to use more simple approaches to grasp objects instead of grasp planners described above when geometric model is not known. In [10] unknown objects are grasped by first segmenting an user indicated object and then grasping the object using a parallel jaw gripper. The grasp orientation and location is based on simple rules applied to the segmented point cloud.

Another approach for finding grasps is object affordance modeling. While object affordance is a broader subject, the affordances can also be thought in the sense of grasp stability. In some of the grasp related studies, grasp affordances consider the overall stability of the grasp [11], [12] or, for example, the grasp affordance in specific tasks [13].

Learning to find good grasps is another view on the problem. [11] utilizes learning on a real robot to learn the grasp affordance of an object. The learning process reduces a vision bootstrapped distribution of grasps to a smaller set of grasps containing only good grasps. Reinforcement learning [14] can also be applied, so that a sequence of grasps can be learned which will lead to a stable grasp of an object.

Our approach to grasping is more related to the methods found in [15] and [2]. The aim of [15] is to reduce the uncertainty of a object's pose to enable grasping the object. In [2], the shape of the object is also uncertain in addition to the pose. In both of the studies, the method is presented with a parallel jaw gripper grasping a 2D-object. However, these methods do not utilize sensor information gained during grasping. Also in [16], the authors propose a decision-theoretic controller which minimizes the uncertainty of the object pose using arm trajectories to enable task specific grasps on objects. Tactile sensors were used to detect contacts between the hand and the objects. A new algorithm, Guaranteed Recursive Adaptive Bounding (GRAB), for inference was developed in [17]. The algorithm was also tested in a manipulation environment where the algorithm made accurate inference of object's pose in both simulation and real environments. However, only the problem of object localization was studied.

This paper will present a probabilistic approach for reasoning about the grasp stability and the object attributes, so that grasp planning can be utilized even if object's pose is uncertain. As can be seen from our survey of recent grasp planners and other grasping methods, similar grasping frameworks have not been yet published to our best knowledge.

III. GRASPING IN PROBABILISTIC FRAMEWORK

The probabilistic framework is now presented in a general form. We model sensor-based grasping using the following variables: S denotes the stability of a grasp as a binary value, G the grasp attributes (e.g. the pose of the end-effector), O the object attributes (e.g. the pose of the target object) and T represents on-line measurements, for example, tactile information. The variables have characteristics: G, the grasp attributes, can be controlled, T can be measured for each grasp attempt, while O is uncertain, that is, we assume we only have an uncertain initial estimate of the object attributes.

In our framework, traditional grasp planning algorithms try to maximize the stability, S, by controlling the grasp attributes, G, with perfect knowledge of the object attributes O,

$$\max_G \ P(S|G,O) \ . \tag{1}$$

In our model, O is not precisely known but instead represented as a probability distribution. Moreover, we do not assume that the available model (1) is precise, that is, the stability of a grasp given information about the object and grasp need not be binary, but instead the model itself can exhibit uncertainty, for example, due to simplifications made in a simulator to compute a grasp quality metric.

It has been shown that grasp stability can be estimated using tactile information [1]. Thus, we can build a probabilistic model for the stability given the other variables, $P(S|G,O,T)$. That model can be used to assess the stability of a single grasp attempt, as shown in [1]. Moreover, for stability detection with uncertain object knowledge, we can marginalize over the uncertain object attributes, such that the probability of a stable grasp given the grasp attributes and tactile measurements is given by

$$P(S|G,T) = \int P(S|G,O,T)P(O|G,T)\,\mathrm{d}O \ . \tag{2}$$

If the grasp attributes are also uncertain, we can marginalize over them in a similar fashion to find $P(S|T)$. This is also the model for grasp stability for the case where no information about the object or grasp is used for stability recognition.

In order to perform grasp planning, we again need to marginalize over the distribution of object attributes. That is we need to find the mode of $P(S|G)$. The marginalization can be written

$$P(S|G) = \int P(S|G,O)P(O)\,\mathrm{d}O \ . \tag{3}$$

Note that since the tactile information for a future grasp attempt is not available, we can use the tactile information only from the previous grasp attempts to update the posterior distribution for the object attributes $P(O)$. That is, we can use the model $P(O|G,T)$ to update the posterior of object attributes. Thus, after some tactile information has been collected, for grasp planning we find the maximum

$$\max_G P(S|G) \approx \max_G \int P(S|G,O)P(O|G_{t-1},T_{t-1})\,\mathrm{d}O \ . \tag{4}$$

Equation (4) shows that the stability S can be maximized by finding the best grasp G, when G_{t-1} and T_{t-1} are known (subscripts denoting that these are from the previous attempt). To build a working system based on the Equation (4), two models are needed:

- Model for $P(O|G,T)$, describing relation between tactile information and grasp and object attributes.
- Model for $P(S|G,O)$, stability as a function of grasp and object attributes

Unfortunately, these models are not trivial to build and depend on the object and the manipulator used to grasp the object. Still, there are existing models for both cases, e.g. see [3] for a model for $P(O|G,T)$ and [18], [1] for a model for $P(S|G,O,T)$. It should also be noted that $P(O|G,T)$ can be obtained from a prediction model of sensor measurements $P(T|G,O)$ using the Bayes formula. One approach to generate the models is to simulate the object and the manipulator to produce the required tactile information and stability models. We have used this approach to demonstrate the framework in action in Section IV.

Our framework does not place constraints on the actual models, and the attributes G, O, T can be freely chosen. For example, G and O can include the poses of the manipulator and the object. The benefit of the presented probabilistic framework is that throughout the grasping process the uncertainty of the actions arising from equation (4) is known. Also, measurement errors can be accounted for during both grasp planning as well as on-line grasp stability detection.

IV. DEMONSTRATION

We will demonstrate the validity of the proposed method using a simulated environment. The purpose of the demonstration is purely to demonstrate the theoretical possibilities of the basic idea. The simulation environment is depicted in Figure 1. The environment consists of a parallel jaw gripper with finger width $l_{\texttt{finger}}$ and a rectangular object with side lengths of 6 and 2. Note that in the demonstration we do not consider the cases where the gripper fingers would touch the shorter side or the corners of the rectangle. The angle of the gripper in degrees is denoted by θ, which is zero when the gripper is perpendicular to the long side of the object. The gripper center is denoted with (x, y), which is relative to the object center (x_0, y_0). In the demonstration, the object is static. When grasping, we can close the two fingers of the gripper independently of each other and we will use the distance to contacts from gripper center d_1 and d_2 as the measurements, representing the tactile information T. We assume that the fingers have the capability to detect when they come into contact with the object and that we can stop the fingers at those instants. We will use a 3-tuple (x, y, θ) to denote the gripper variables, which are in relation to the object center (x_0, y_0, θ_0). (x, y, θ) represent G, the grasp attributes, while (x_0, y_0, θ_0) represents the O, object attributes.

Note that because of the symmetry of the setup, there is ambiguity about the orientation of the gripper if sensor measurements are used to estimate the orientation when the fingers are in contact with the top and bottom sides of the object. Also due to the symmetry, we can not reduce uncertainty along the x-axis.

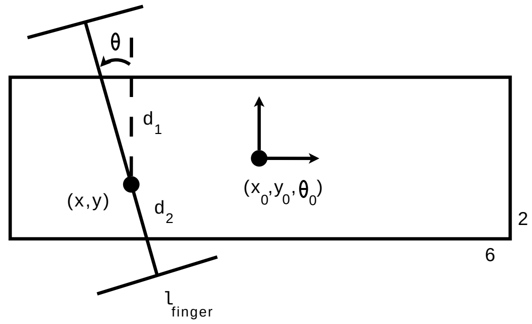

Fig. 1. Simulation environment.

1) Implementation: Our general approach is based on the sequence of actions shown in Figure 2. We assume that some type of initial estimate (with associated uncertainty) of the object pose is obtained in phase 1, e.g., from vision. Using the estimate, we can plan for a grasp with the uncertainty from the initial estimate, phase 2. Then a grasp is performed, phase 3, giving measurement data (we assume tactile and joint configuration data is available). Using the measurement data, we can make a decision of the grasp stability, phase 4. If the grasp is stable, the object can be manipulated, if not, we can plan for a new grasp, phase 5, with the new information from the attempted grasp. This loop can then be further iterated until grasp stability conditions are satisfied.

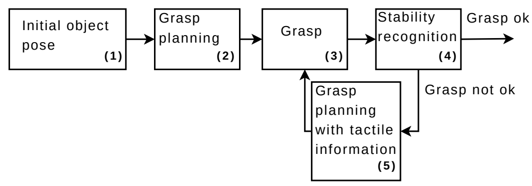

Fig. 2. Sequence of actions.

The theoretical framework described in Section III is implemented with particle filters to make the computation of evolving probability distributions tractable. The particle filter method is an MCMC method which models probability distributions with a cloud of particles. More information on particle filtering, especially applied to robotics can be found in [19]. Particle filters have been used in manipulation, for example in [3], to estimate object pose using tactile sensors. We use two separate particle processes to estimate the two different posterior densities, $P(O|G,T)$ and $P(S|G,O)$, introduced in Section III. Likelihoods, which are shown in Figure 3, were chosen by hand for the purposes of this example. Figure 3(a) shows how likely a measurement d is a correct measurement in relation to the true measurement d^*, this model is used for both d_1 and d_2, and represents additive Gaussian noise in the measurement. Figure 3(b) presents how likely a grasp is stable relative to the belief of object pose O. The stability model defines that a grasp is stable when the grasp is performed close enough to the center of the object. Furthermore, for the purpose of the example, we defined that the stability model factorizes into independent factors for different coordinate axes, i.e., $P(S|G,O) = P(S|x,x_0)P(S|y,y_0)P(S|\theta,\theta_0)$.

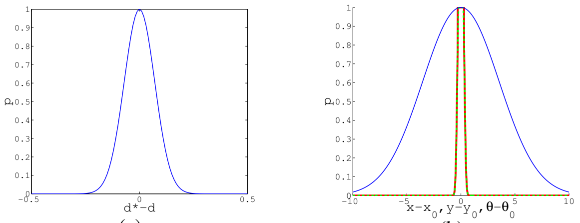

Fig. 3. Likelihoods for: (a) Measurement model, $p(d|d^*)$; (b) Grasp stability, $P(S|G,O)$, for (x, y, θ), x in red, y in dashed green, θ in blue.

Algorithms 1 and 2 describe our method of finding stable grasps. The algorithms also contain the variables and distributions that we have used in the particle filter processes. Algorithm 1 requires the initial estimates of the uncertainty, given in $\sigma_{\mathbf{init}}$, for each of the variables (x, y, θ). The particle set O_1 in Algorithm 1 represents the probability distribution of the object, $P(O|G,T)$, while particle set G_1 in Algorithm 2 represents the relative or corrective motion to the actual grasp, and by applying the relative motion to each of the particles in O_1, we can find the probability of a stable grasp $P(S|G,O)$. In Algorithm 2 the maximum of distribution, $\max_G \int P(S|G,O)$, is searched for and the corresponding relative motion is then applied.

Relating the algorithms to Figure 2, Algorithm 2 takes care of the grasp planning, that is, phases 2 and 5. Algorithm 1 handles phase 3, grasping the object and updating the belief of object pose. In line 13 of Algorithm 1, the grasp stability probability is computed and corresponds to phase 4 of the action sequence.

2) Results: Figure 4 shows a single example run of Algorithm 1. The example was run with the true object pose $(x_0,\ y_0,\ \theta_0)$ set to $(0,\ -0.3,\ -15)$. Figure 4(a) shows the initial distribution of O_1, which was initialized around zero with $\sigma_{\mathbf{init}} = [0.3\ 0.3\ 6]$. Particle locations are shown in green, indicating the possible object location and $\frac{1}{4}$ of the particles are plotted with blue line, indicating the orientation, θ_0, of the object. As the initial distribution is zero mean, during the first iteration the grasp planning stage will produce a near zero relative motion as there are no measurements yet. In Figure 4(b), the first grasp attempt has been made, and the distribution of object pose changes to account for the measurements, d_1 and d_2. The figure also shows the symmetry of the problem and two modes arising from this symmetry, one for $\theta_0 = 15$, other for $\theta_0 = -15$. This grasp does not satisfy the threshold of 0.5 for the grasp stability probability. Maximizing $P(S|G,O)$ yields solution $(0.07,\ -0.32,\ -14.3)$ for the grasp G. Figure 4(c) shows the posterior distribution of O_1 after the information from the second grasp is included. This grasp is determined as stable as the probability of a stable grasp is greater than the probability of an unstable grasp. The mean of the final

Algorithm 1 find_stable_grasp($\sigma_{\mathbf{init}}$)

1: Generate initial particle set, O_1 according to $\mathcal{N}(0, {\sigma_{\mathbf{init}}}^2)$
2: $q \leftarrow 1$
3: **while** $q = 1$ **do**
4: $(x,\ y,\ \theta) \leftarrow$ **find_best_relative_motion**(O_1,σ_{init})
5: Apply motion $(x,\ y,\ \theta)$ to gripper
6: Grasp object
7: **while** O_1 is not converged **do**
8: For each particle, simulate the finger lengths, d_1 and d_2
9: Weigh particles O_1, $w_1 \propto p(d|d^*)$, i.e. estimate $P(O|G,T)$
10: Do importance filtering according to w_1
11: Use $\mathcal{N}(0, {\sigma_{\mathbf{1}}}^2)$ as proposal distribution with $\sigma_{\mathbf{1}} \leftarrow$ [0.02 0.02 2]
12: **end while**
13: Approximate $P(S|G,O)$ by $\sum_i P(S|G,O_{1_i})$
14: **if** $\sum_i P(S|G,O_{1_i}) > 0.5$ **then**
15: $q \leftarrow 0$
16: **end if**
17: **end while**

Algorithm 2 find_best_relative_motion(O_1,σ_{init})

1: Generate particle set, G_1 according to $\mathcal{N}(0, \mathbf{5}{\sigma_{\mathbf{O_1}}}^2)$
2: **while** G_1 is not converged **do**
3: Weigh particles G_1, $w_2 \propto P(S|G,O)$
4: $(x_{max}, y_{max}, \theta_{max}) \leftarrow \max_G P(S|G,O)$
5: Do importance filtering according to w_2
6: Use $\mathcal{N}(0, {\sigma_{\mathbf{2}}}^2)$ as proposal distribution with $\sigma_{\mathbf{2}} \leftarrow$ 0.2 $\sigma_{\mathbf{init}}$
7: **end while**
8: **return** $(x_{max}, y_{max}, \theta_{max})$

posterior distribution O_1 was $(-0.24,\ -0.32,\ -14.28)$ compared to the set pose which was $(0,\ -0.3,\ -15)$. As can be seen, the method was able to find a corrective motion for the gripper and produce a stable grasp and after the two grasps, the particle cloud converged to near optimal values for the object pose, except for the x-variable for which the uncertainty can not be reduced.

However, as the method is probabilistic, for the second grasp attempt maximizing the probability $P(S|G,O)$ can produce a motion that is opposite to the correct one in θ. Nevertheless after the mistaken grasp, the wrong mode will be eliminated, and during the next iteration of Algorithm 1, the correct motion will be found. The particle clouds are not shown for this case as they are essentially similar to the first case. Overall, in the case of this example, the system will thus usually make two or three grasps before finding a stable grasp, depending on the first corrective motion.

One of the benefits of the probabilistic approach is that we always know the uncertainty behind the actions. Figure 4(d)

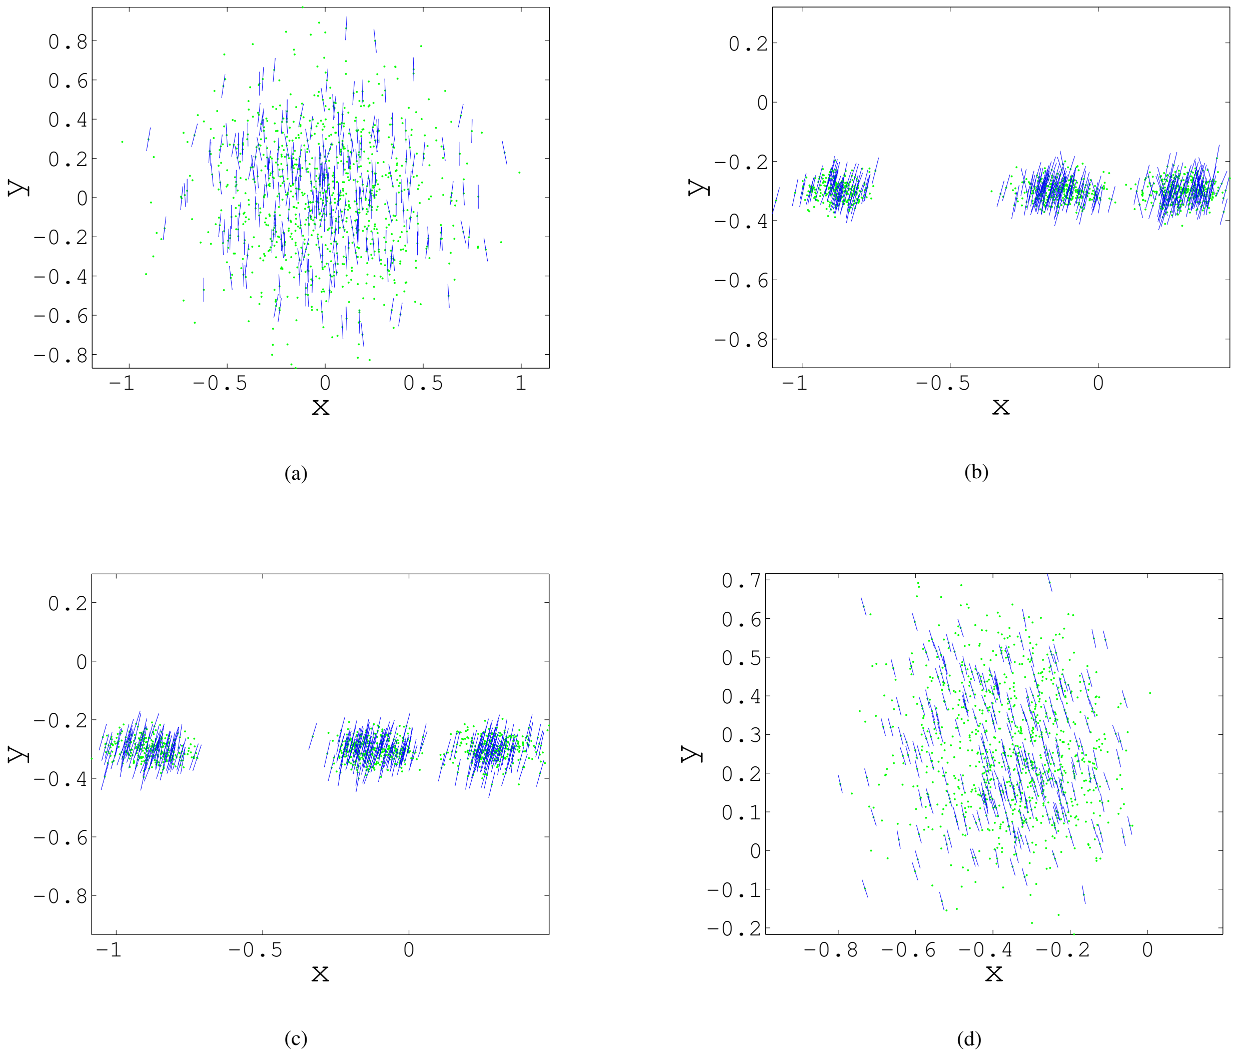

Fig. 4. Sequential distributions of particles modeling $P(O|G,T)$, (a)-(c), for a single run of Algorithm 1: (a) Initial distribution ; (b) Distribution after the first grasp; (c) Distribution after second grasp, for which $P(S|G,O) = 0.503$; (d) An example of distribution $P(S|G,O)$

shows an example of probability distribution of $P(S|G,O)$, the particle locations and orientations now show the possible relative motion of the gripper. There are numerous possibilities for using this information. For example, the uncertainty of the corrective motion can be observed and by marginalizing over O, the uncertainty is directly taken into account in the stability estimation and grasp planning. Also, if non-contact measurements would be preferred, a low probability of a stable grasp in the planning could be used to trigger additional visual measurements to improve the object pose estimate.

V. CONCLUSIONS AND FUTURE WORK

We have presented a novel framework for grasping, which operates in a probabilistic setting. The framework allows grasp planning, measurements, and corrective motions to interact, leading to a system where we can estimate uncertain object attributes, such as pose, and improve grasp stability simultaneously. We also presented a demonstration of our framework utilizing particle filters. The demonstration showed that the approach is able to find a stable grasp and simultaneously update the pose estimate of the object. However, what we demonstrated is only a small part of the possibilities presented by the probabilistic framework. In addition to pose, we could estimate other attributes, for example, the size, the shape or the identity of an object. The models used in the demonstration can also be extended, for example, the model for stability, $P(S|G,O)$, can be extended to include tactile information, $P(S|G,O,T)$, which has been done in previous work [18], [1].

To extend the practical use of the probabilistic framework for sensor-based grasping, ongoing work is focused on applying the framework presented here to more complex manipulators, objects and statistical models. Our plans include working both with simulation for the repeatability and real hardware. Finally, it should be mentioned that the framework could be applied in any manipulation scenario where a success function, attempted actions, sensor models and environment variables can be defined. In the case where a single step action is not sufficient, the approach can be understood in the context of partially observable Markov decision processes (POMDPs), which provide an approach for the solution, although the POMDP approach might well be computationally intractable.

REFERENCES

[1] Y. Bekiroglu, J. Laaksonen, J. A. Jorgensen, and V. Kyrki, "Learning grasp stability based on haptic data," in *Robotics: Science and Systems (RSS 2010) Workshop on Representations for Object Grasping and Manipulation in Single and Dual Arm Tasks*, 2010.

[2] V. Christopoulos and P. Schrater, "Handling shape and contact location uncertainty in grasping two-dimensional planar objects," in *Intelligent Robots and Systems, 2007. IROS 2007. IEEE/RSJ International Conference on*, Nov. 2007, pp. 1557 –1563.

[3] C. Corcoran and R. Platt, "A measurement model for tracking hand-object state during dexterous manipulation," in *Robotics and Automation (ICRA), 2010 IEEE International Conference on*, May 2010, pp. 4302 –4308.

[4] S. Chitta, M. Piccoli, and J. Sturm, "Tactile object class and internal state recognition for mobile manipulation," in *Proceedings of the IEEE International Conference on Robotics and Automation*, 2010, pp. 2342–2348.

[5] C. Ferrari and J. Canny, "Planning optimal grasps," in *Robotics and Automation, 1992. Proceedings., 1992 IEEE International Conference on*, May 1992, pp. 2290 –2295 vol.3.

[6] K. Huebner, S. Ruthotto, and D. Kragic, "Minimum volume bounding box decomposition for shape approximation in robot grasping," in *Robotics and Automation, 2008. ICRA 2008. IEEE International Conference on*, May 2008, pp. 1628 –1633.

[7] C. Goldfeder, P. K. Allen, C. Lackner, and R. Pelossof, "Grasp Planning Via Decomposition Trees," in *IEEE International Conference on Robotics and Automation*, 2007, pp. 4679–4684.

[8] C. Goldfeder, M. Ciocarlie, H. Dang, and P. K. Allen, "The Columbia Grasp Database," in *IEEE International Conference on Robotics and Automation*, 2009.

[9] C. Goldfeder, M. Ciocarlie, J. Peretzman, H. Dang, and P. K. Allen, "Data-Driven Grasping with Partial Sensor Data," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2009.

[10] A. Jain and C. C. Kemp, "El-e: an assistive mobile manipulator that autonomously fetches objects from flat surfaces," *Auton. Robots*, vol. 28, pp. 45–64, January 2010. [Online]. Available: http://dx.doi.org/10.1007/s10514-009-9148-5

[11] R. Detry, D. Kraft, A. Buch, N. Kruger, and J. Piater, "Refining grasp affordance models by experience," in *Robotics and Automation (ICRA), 2010 IEEE International Conference on*, May 2010, pp. 2287 –2293.

[12] C. Barck-Holst, M. Ralph, F. Holmar, and D. Kragic, "Learning grasping affordance using probabilistic and ontological approaches," in *Advanced Robotics, 2009. ICAR 2009. International Conference on*, June 2009, pp. 1 –6.

[13] D. Song, K. Huebner, V. Kyrki, and D. Kragic, "Learning task constraints for robot grasping using graphical models," in *Intelligent Robots and Systems (IROS), 2010 IEEE/RSJ International Conference on*, Oct. 2010, pp. 1579 –1585.

[14] R. Platt, "Learning grasp strategies composed of contact relative motions," in *Humanoid Robots, 2007 7th IEEE-RAS International Conference on*, Dec. 2007, pp. 49 –56.

[15] K. Goldberg and M. Mason, "Bayesian grasping," in *Robotics and Automation, 1990. Proceedings., 1990 IEEE International Conference on*, May 1990, pp. 1264 –1269 vol.2.

[16] K. Hsiao, L. Kaelbling, and T. Lozano-Perez, "Task-driven tactile exploration," in *Proceedings of Robotics: Science and Systems*, Zaragoza, Spain, June 2010.

[17] A. Petrovskaya, S. Thrun, D. Koller, and O. Khatib, "Guaranteed inference for global state estimation in human environments," in *RSS 2010 Mobile Manipulation Workshop*, 2010.

[18] Y. Bekiroglu, J. Laaksonen, J. A. Jorgensen, V. Kyrki, and D. Kragic, "Assessing grasp stability based on learning and haptic data," Accepted to IEEE Transactions on Robotics.

[19] S. Thrun, W. Burgard, and D. Fox, *Probabilistic robotics*. MIT Press, 2005.