A Data-Driven Statistical Framework for Post-Grasp Manipulation

Robert Paolini, Alberto Rodriguez, Siddhartha S. Srinivasa, and Matthew T. Mason
The Robotics Institute, Carnegie Mellon University
{rpaolini, albertor}@cmu.edu, {siddh, matt.mason}@cs.cmu.edu

***Abstract*— Grasping an object is usually only an intermediate goal for a robotic manipulator. To finish the task, the robot needs to know where the object is in its hand and what action to execute. This paper presents a general statistical framework to address these problems. Given a novel object, the robot learns a statistical model of grasp state conditioned on sensor values. The robot also builds a statistical model of the requirements of the task in terms of grasp state accuracy. Both of these models are constructed by offline experiments. The online process then grasps objects and chooses actions to maximize likelihood of success. This paper describes the framework in detail, and demonstrates its effectiveness experimentally in placing, dropping, and insertion tasks. To construct statistical models, the robot performed over 8000 grasp trials, and over 1000 trials each of placing, dropping and insertion.**

I. Introduction

Knowledge of the grasp state is often critical to any subsequent manipulation task. Intuitively, harder tasks demand a more accurate estimation of the state of a grasp than simpler ones. For example, balancing a cylinder on a table requires more accuracy than dropping it into a hole. More generally, consider a manipulator, an object to manipulate, a task, and a set of actions designed to accomplish the task. In this paper we build a data-driven framework to automate the process of deciding whether the task is solvable with the available hardware and set of actions, and find the action most likely to succeed.

The statistical framework proposed in this paper is best suited to model the execution of tasks that require grasping an object prior to execution, i.e., post-grasp manipulation tasks. We address the problem by separating it into two independent steps. First, estimate the state of the grasp with in-hand sensors, and second, model the accuracy requirements that the particular task imposes on our state estimation. This separation yields the benefit that we can use the same model of state estimation for different tasks, and the same model of task requirements for different manipulators. Using this framework, each sensor reading generates a probability distribution in task action space, enabling us to find not only the optimal action, but to understand just how likely that action is to succeed.

Figure 1 illustrates the process for placing an object. Sensors in the hand provide information of the grasp state. First, we estimate the probability distribution of the pose of the object in the hand. Second, we predict the probability of success of each available action. Both of these are computed based on data-driven models. Finally, we choose the action most likely to succeed.

We test the framework with three different manipulation tasks: placing an object, dropping it into a hole, and inserting it. In this extended abstract, we show the results for the first task. The experimental setup in Figure 1 consists of a simple gripper [1], [2] mounted on a robotic arm that iteratively grasps an object from a bin, estimates the distribution of the pose of the object, computes the probability of success for all available actions, chooses the optimal one, and executes it.

II. Statistical Framework

Our goal in this paper is to find the action a from a set of available actions $\mathcal{A}$ that, given sensor inputs $z \in Z$, maximizes the expected performance of succeeding at a task.

Two possible options are to model the performance of an action directly as a function of sensor observations, or to project the sensor inputs z to a more compact representation of state, x. However, in this work, we choose to encapsulate uncertainty by representing the system by its state belief $P(x|z)$ rather than just by its most likely value x. By maintaining the distribution of all possible poses of the object, we can later make a more informed prediction on the probability of success of a given action.

The dimension of the space of belief distributions $\mathrm{Bel}(X)$ is too large to model the probability of success of an action $P(a|z)$ directly from the belief $P(x|z)$. We simplify this problem by marginalizing the probability of success of an action $P(a|z)$ with respect to the true state of the system x:

$$P(a|z) = \int_X P(a|z,x) \cdot P(x|z)\mathrm{d}x = \int_X P(a|x) \cdot P(x|z)\mathrm{d}x \tag{1}$$

where in the last step, we make the assumption that the state representation x is informative enough such that the output of an action is conditionally independent of z, given the true state x.

A. Learning Sensing Capabilities

In this section, we model the posterior distribution $P(x|z)$. Learning $P(x|z)$ directly is usually expensive in terms of data required, since the complexity of the model depends on the dimension of z. To simplify the process we use Bayes rule to flip the conditioning and assume a Gaussian likelihood $P(z|x) \sim \mathcal{N}(z; \mu(x), \sigma(x)^2)$. The likelihood of the system is the distribution of sensor readings given the true state of the system, which is usually unimodal.

$$P(x|z) = \frac{P(z|x)P(x)}{P(z)} \simeq \mathcal{N}(z; \mu(x), \sigma(x)^2) \cdot \frac{P(x)}{P(z)} \tag{2}$$

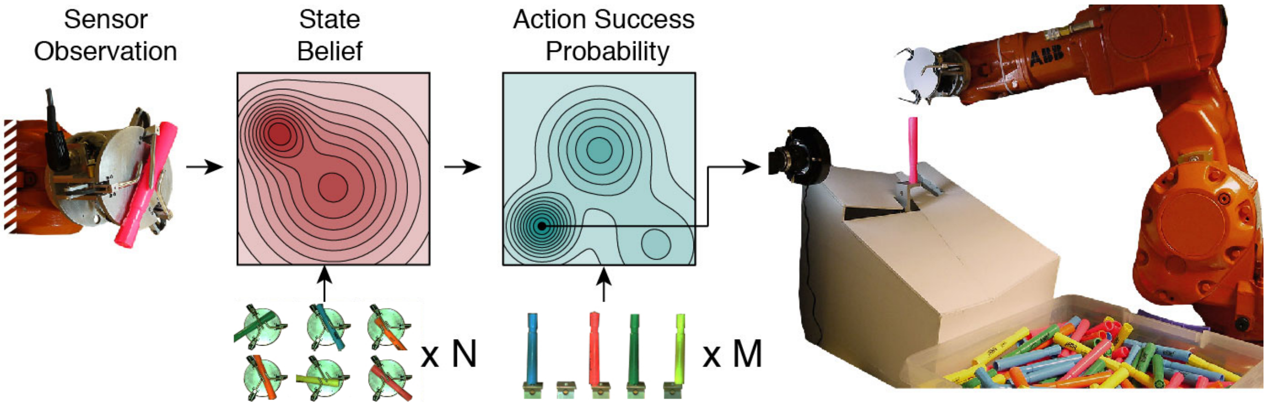

Fig. 1. Procedure to choose the optimal action to accomplish a manipulation task. First, we learn the belief of the state of the system from sensor readings. Based on that belief, we then estimate the probability of success of available actions and choose on the best action to take. Both the state estimation and task requirements are learned using real data.

where $\frac{P(x)}{P(z)}$ is the normalized state prior, and both μ and σ are functions of the true state of the system x. We normalize $P(x|z)$ a posteriori, rather than explicitly computing $P(z)$. In our implementation, we use Kernel Density Estimation to model the prior distribution $P(x)$, and Gaussian processes [3] to regress the functions $\mu(x)$ and $\sigma(x)$.

Figure 2 shows data from 2000 grasps in the $r - \theta$ plane, the estimated prior distribution $P(x)$, and the posterior distribution $P(x|z)$ corresponding to one example grasp.

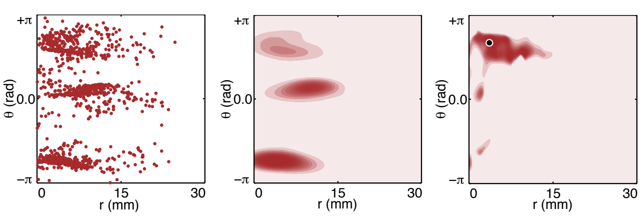

Fig. 2. [Left] State from 2000 graps collected to model the *sensing capabilities* of the manipulator. [Center] Prior state distribution $P(x)$. [Right] $P(x|z)$ for an example grasp, which shows us the most likely pose of the object, and where else it could be.

B. Learning Task Requirements

We now model the probability of success of an action, $P(a|x)$. This will tell us how accurate our estimation of grasp state must be for an action to succeed at a task. While not required for our framework, we choose to state parameterize the set of actions. For example, for the task of placing a cylindrical object, if assumed to be at pose p, we design an action a_p that turns the cylinder so it is upright with respect to the ground, and then set it down. Note that a_p is an action parameterized by a chosen state p.

In general, the success of an action depends both on the action a_p itself and the true state of the system x. Since we assume state parametrized actions, we assume that the probability of success only depends on the difference $(x-p)$. When placing a cylinder whose estimated axis is 1 degree off from its true state, we are more likely to succeed than if we try to place an object several degrees off. We model the outcome of an action a_p as a Bernoulli random variable of parameter ϕ_{a_p}, so that: $P(a_p = 1|x) = \phi_{a_p}(x) = \phi(x-p)$

The use of state parameterized actions allows us to randomly sample the space of mismatches $(x-p)$ by choosing to execute the action a_p with $p = x + \epsilon$, where ϵ is a uniformly distributed error in the space of system states instead of the optimal one a_x. In our implementation, we use a Gaussian process to regress the Bernoulli parameter ϕ with the outcome of 1000 placements. Figure 3 shows the requirements $\phi(x-p)$ for placing a highlighter marker with a simple hand. As $|x-p|$ increases, our likelihood of task success decreases, which is as expected.

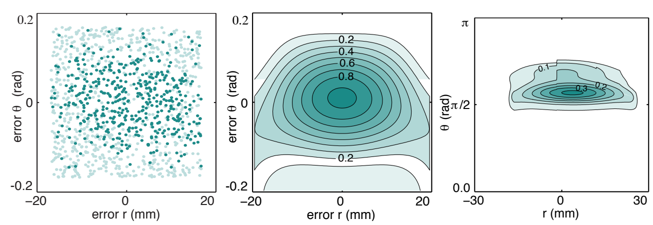

Fig. 3. [Left] Data from 1000 purturbed placing experiments. Darker points are successes. [Center] $P(a_p = 1|x)$, requirements for the task of placing a marker subject to state perturbation. [Right] $P(a_p|z)\ \forall\ p$, distribution of predicted placing accuracy for an example grasp.

C. Matching Task Requirements with Sensing Capabilities

Here we combine the models of $P(x|z)$ and $P(a|x)$ to estimate the probability of success of an action a_p. For that, we extend (1) as:

$$\begin{aligned} P(a_p = 1|z) &= \int_X P(a_p = 1|x)P(x|z)\mathrm{d}x \\ &= \int_X \phi(x-p)\mathcal{N}(z; \mu(x), \sigma(x)^2)\frac{P(x)}{P(z)}\mathrm{d}x \\ &\simeq \int_{\mathcal{E}} \phi(\epsilon)\mathcal{N}(z; \mu(p+\epsilon), \sigma(p+\epsilon)^2)\frac{P(p+\epsilon)}{P(z)}\mathrm{d}\epsilon \end{aligned} \quad (3)$$

where we apply the change of variables $\epsilon = x - p$. Depending on the value of $\max_{a_p} P(a|x)$ we can decide either to execute the task with the optimal action, or to abort the execution. Figure 3 shows the predicted task success distribution for an example grasp.

III. Experimental Validation and Discussion

We performed an additional 500 placing trials with placing parameters as outlined in Section II. Figure 4 compares experimental results to model predictions. We bin grasps by their predicted placing probability and compute the experimental success rate of those grasps. If we look at all of the grasps that were predicted to succeed around 40% of the time, the average experimental success rate for those grasps should also be around 40%. The experimental probability tends to follow the predicted probability, indicating the validity of this framework. We can significantly improve our system performance by making an informed decision based on the predicted probabilitiy of success of a given action.

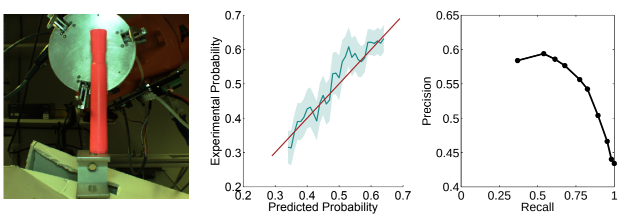

Fig. 4. [Left] shows the robot succeeding at placing a marker. [Center] shows that the experimental success rates of grasps binned by their predicted probability support that predicted probability value. [Right] shows how we can increase our success rate by rejecting low placing probabilities.

In this paper, we have laid out a general statistical framework to model the sensing capabilities of a hand and task success requirements. In addition, we show how to combine them together in a principled way to yield a mapping between sensors and the best action for a task, as well as the expected success rate. All of this was done in a data-driven way, requiring over 8000 real grasps.

References

[1] A. Rodriguez, M. T. Mason, and S. S. Srinivasa, "Manipulation Capabilities with Simple Hands," *International Symposium on Experimental Robotics (ISER)*, 2010.

[2] M. T. Mason, A. Rodriguez, S. S. Srinivasa, and A. S. Vazquez, "Autonomous Manipulation with a General-Purpose Simple Hand," *The International Journal of Robotics Research*, vol. 31, no. 5, pp. 688–703, 2012.

[3] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*. Press, MIT, 2006.