Data-Driven Statistical Modeling of a Cube Regrasp

Robert Paolini and Matthew T. Mason
The Robotics Institute, Carnegie Mellon University
rpaolini@cmu.edu, matt.mason@cs.cmu.edu

***Abstract*— Regrasping is the process of adjusting the position and orientation of an object in one's hand. The study of robotic regrasping has generally been limited to use of theoretical analytical models, and cases with little uncertainty. Theoretical analytical models and simulations have so far proven unable to capture the complexity of the real world. Empirical statistical models are more promising, but collecting good data is difficult. In this paper, we collect data from over 3000 robot regrasps, and use this data to learn two probability functions: 1) The probability that the object is still in the robot's hand after a regrasp action; and 2) Given an initial pose, action, and the object is still grasped, the probability distribution of the object pose after the regrasp. Both of these functions are learned using kernel density estimation with a similarity metric over object pose. We show that our data-driven models achieve comparable accuracy to a geometric model and an off the shelf simulator in classification and prediction tasks, while also enabling us to predict probability distributions.**

I. INTRODUCTION

Humans are experts at reorienting objects in their hands. They use this skill to adjust the grip of a pencil to write with it, or to adjust the grip of a key from its teeth to its head to unlock a door. By contrast, once a robot has picked up an object, it generally maintains the same grasp as long as the object is in contact with the hand. If a robot does adjust an object grip, it is generally a predetermined operation with deterministic results, and only works for a constrained set of initial and desired final grasp poses. In contrast, humans can adapt to different objects, with arbitrary initial and desired final poses. One possible explanation of this discrepancy is that humans have better models of how the object pose changes as a function of their actions. In this paper, we show how robots can build better models of regrasp actions.

By a regrasp action, we mean any sequence of movements that results in a change of the object pose with respect to the hand. Often, the final pose of an object is critical to a task. If the initial pose of an object is arbitrary, then the robot must use models to determine what regrasps to use to move the object to the desired pose. Physics-based models can only take us so far. Modeling multiple contacts along with impact, friction, and uncertainty in object size, mass, finger shape, dirt, etc can make it difficult to compute models a priori. In this paper, we encapsulate the noise and uncertainty in the model as a probability distribution. A simple Gaussian is not a good model of the probability distribution. Manipulation is the iconic non-Gaussian problem. Flipping a coin, or the difference in object pose based on whether or not contact occurs, cannot be represented as a unimodal Gaussian

This work was supported by National Science Foundation [NSF-IIS-0916557] and Army Research Laboratory [W911NF-10-2-0016]. This work does not necessarily reflect the position or the policy of the U.S. Government or ARL. No official endorsement should be inferred. Some equipment was provided by ABB.

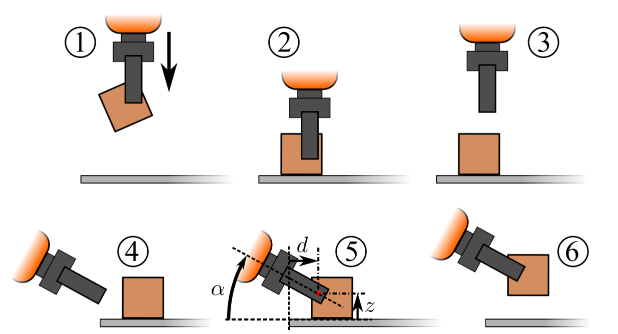

Fig. 1. Place and pick regrasp action studied in this paper. Initially, the robot is holding a block between its fingers. Then, it moves downwards to a specific pose above the platform, where the block may move to conform to the new contact. It then opens its fingers and repositions itself at a certain position and orientation with respect to the edge of the platform. It then closes its fingers and moves upwards, completing the regrasp. There are several scenarios where the robot will not be holding the block in step 6. In step 1, the object could slip out of the hand. In step 2, the object-hand system could collide with the table. In step 5, the robot could miss the block and fail to pick it up.

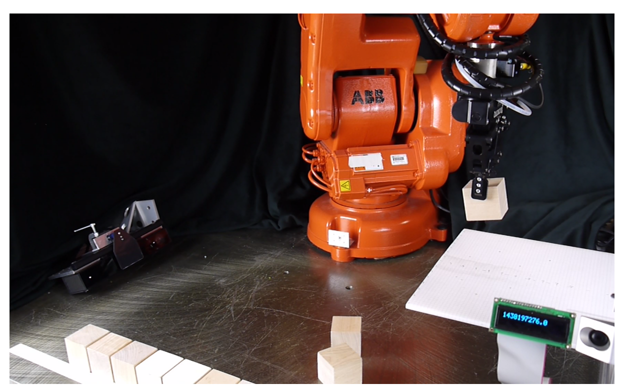

Fig. 2. Place and pick regrasp data collection setup. We use an industrial arm with a parallel jaw gripper and place an pick objects. A three dimensional vision system based on point clouds is used to record the object position before and after a regrasp. In the event of a failure, the robot either picks up the dropped block from on top of the platform or retrieves a new block from a stack of fresh blocks. With this setup, the robot performed over 3000 regrasp experiments. We use this data to fit models for predicting the probability of not dropping the object, and estimating the final pose of the block after a regrasp.

distribution. In this paper, we look at other ways to handle the modeling challenges presented by regrasps.

To accurately model manipulation actions, a lot of real world data is necessary. For this paper, we performed over 3000 robot regrasp experiments by reducing human intervention.

We study a place and pick regrasp of a cube (Figure 1) as a first step to understand the challenges involved in statistical modeling of regrasp actions. We collect real world data to learn two models: 1) Given an initial object pose and regrasp action, how likely is it that the object remains grasped? and 2) Given an initial pose and action, where do we expect the object to end up?. We show that our learned model, despite not having any prior knowledge of the task, achieves comparable accuracy to simulation and a geometric model.

The rest of the paper is outlined as follows. In Section II we look at prior work in this area. In Section III we outline our method for modeling regrasps. In Section IV we explain the data collection process and experiments performed. In Section V we compare our model with an off-the-shelf simulator and geometric model, and in Section VI we summarize our work and discuss future directions.

II. PRIOR WORK

Regrasping has been studied for a long time, starting with Paul [1], Tournassoud et al. [2], Fearing [3], and Brock [4]. Early regrasping work assumed a known world model with deterministic actions. Most regrasping work falls under 3 categories: pick and place [2], [5], [6], closed-loop dynamic regrasping [7][8][9], or what is generally referred to as dexterous manipulation or finger gaiting [3], [10][11][12][13][14][15]. Chavan Dafle et al. [16] present work on "extrinsic dexterity", which uses gravity, inertia, and external contacts to vary the pose of the object within the hand.

Uncertainty during manipulation has been represented using two approaches: "possibilistic" and probabilistic. The "possibilistic" approach [17][18] maintains a set of possible object poses, and the robot makes motions that reduce the size of the set. Brost [19] uses pushing, squeezing and offset grasping with a parallel jaw gripper to reliably grasp objects with high position uncertainty. Dogar and Srinivasa [20] explicitly propagate object uncertainty regions to plan robust grasp plans. Probabilistic approaches [21][22][23][24] maintain a probability distribution of object poses in order to plan the best action. Bayesian estimation [25][26] and particle filters [27][28] are the most common ways to deal with the non-Gaussian, multi-modal probability distributions inherent in manipulation tasks.

To model manipulation actions, researchers often use simulation[29], [30], imitation learning[31], or build models with collected robot data [32][33]. In this paper, we expand on prior work by using real data to model uncertain manipulation actions.

The most similar work to ours is probably the work of Kopicki et. al. [34]. They use regression to learn the resulting motions of real robotic push actions. They also fit multi-modal probability distributions to their data and show improvement over regression. Our work focuses on learning both the probability of maintaining a grasp after a regrasp and the resulting probability distributions of robotic regrasp actions, along with paying closer attention on how to collect a large amount of robot manipulation data.

III. METHOD

A. Task Description

The regrasp action we will learn is shown in Figure 1. The robot moves down vertically to a fixed height above a platform, releases the object, and then grasps it again at a specified position in the workspace. Note that in step 2, the object pivots and slides in the fingertips when it comes into contact with the platform, which we expect to be difficult for physics based models to capture. Our regrasp action $\boldsymbol{a}$ is parameterized by 3 continuous variables, d, z, and α, which correspond to the pose of the hand frame with respect to the edge of the platform.

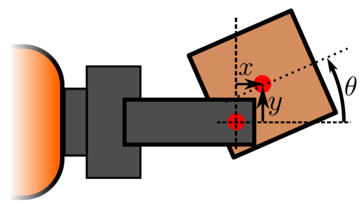

Fig. 3. State space we will be using in this paper. While the world is 6-dimensional, because we are grasping a cube with a parallel jaw gripper, we can reduce the state space to 3 dimensions. Note that as the cube is symmetric, we only allow θ to be between $-\pi/4$ and $\pi/4$.

Our state space is shown in Figure 3. Our state $\boldsymbol{s}$ is represented by 3 continuous variables, x, y, and θ, corresponding to the relative pose of the cube with respect to the hand. Note that this is a planar state space; we will not consider out of plane rotations or grasps. Any grasps of this kind will be considered "not grasped" for the purposes of this paper.

In order for the robot to successfully model this regrasp action, we must learn two probability functions:

1) The probability that the object is still in the robot's hand after a regrasp action, $P(\text{grasped}|\boldsymbol{s}, \boldsymbol{a})$
2) Given an initial state, action parameters, and that the object is still grasped, what the distribution of the final state is, $P(\boldsymbol{s}'|\boldsymbol{s}, \boldsymbol{a}, \text{grasped})$

Note that if the object is not still grasped after a regrasp action, it's final state $\boldsymbol{s}'$ does not exist, since $\boldsymbol{s}'$ is the in-hand pose of the object. Learning these two probability distributions enables us to solve problems in the future such as: 1) what action maximizes the chance of maintaining a grasp? or 2) what action maximizes the chance of the center of the object being at most 1cm away from the center of the fingers? In this paper, we focus solely on learning the above distributions from data, and leave planning with these models as future work.

B. Predicting the Probability of Maintaining a Grasp

To estimate the probability of retaining the object after a regrasp action, we will use kernel density estimation with Bayes discriminant rule [35], [36]. The basic idea is to estimate the probability of retaining and not retaining the object using kernel density estimation, and then for a query point, determine which of the two probabilities are greater. That is, we would like to calculate:

$$P(\pi_i|\boldsymbol{s},\boldsymbol{a}) = \frac{p_i P(\boldsymbol{s},\boldsymbol{a}|\pi_i)}{\sum_j^g p_j P(\boldsymbol{s},\boldsymbol{a},\pi_j)}$$

where $p_i = P(\pi_i)$ is the prior probability of a randomly selected observation being in class π_i, g is the total number of classes, and $P(\boldsymbol{s},\boldsymbol{a}|\pi_i)$ is the conditional probability density of an observation given that it is in class π_i. In our case, we have two classes, grasped and not grasped, so we will learn two probability densities using kernel density estimation:

$$P(\boldsymbol{s},\boldsymbol{a}|\pi_i) = \frac{1}{N_i}\sum_j^{N_i} K^s_{h_1}(\boldsymbol{s},\boldsymbol{s}_j)K^a_{h_2}(\boldsymbol{a},\boldsymbol{a}_j)$$

where N_i is the number of training observations belonging to class π_i. Note that we set $p_i = N_i/\sum_j^g N_j$ as our prior probabilities.

We define our kernel functions by first expressing distances in state and action space, and then use a Gaussian kernel over these distance functions:

$$D_a(\boldsymbol{a}_1,\boldsymbol{a}_2) = \left\| \begin{bmatrix} d_1 \\ z_1 \\ \rho\alpha_1 \end{bmatrix} - \begin{bmatrix} d_2 \\ z_2 \\ \rho\alpha_2 \end{bmatrix} \right\|^2$$

$$D_s(\boldsymbol{s}_1,\boldsymbol{s}_2) = \left\| \begin{bmatrix} x_1 \\ y_1 \end{bmatrix} - \begin{bmatrix} x_1 \\ y_1 \end{bmatrix} \right\|^2 + 2\rho^2(1-\cos(\theta_1-\theta_2))$$

$$K^s_{h_1}(\boldsymbol{s},\boldsymbol{s}_j) = \frac{1}{\eta_s(\rho,h_1)}\exp\left(-\frac{1}{2h_1^2}D_s(\boldsymbol{s},\boldsymbol{s}_j)\right)$$

$$K^a_{h_2}(\boldsymbol{a},\boldsymbol{a}_j) = \frac{1}{\eta_a(\rho,h_2)}\exp\left(-\frac{1}{2h_2^2}D_a(\boldsymbol{a},\boldsymbol{a}_j)\right)$$

Where ρ is the radius of gyration for the object, which allows us to properly trade off distance and angle, while η_s and η_a normalize the kernels so they represent probability distributions. Note that for the state distance function, we use a cosine function to handle angle wrap around for object pose (i.e. $-\pi = \pi$). Both of the distance functions D_a and D_s represent squared distance in action and state space. The units for d, z, x, and y are mm, while α and θ are in radians. Thus, our distance functions have units of mm^2, and our bandwidths h_1 and h_2 have units of mm.

We choose the values of bandwidths h_1 and h_2 that minimize the cross-validated negative log likelihood of the observed data:

$$NLL(h_1,h_2) = -\frac{1}{N_i}\sum_j^{N_i}\hat{P}_j(\boldsymbol{s}_j,\boldsymbol{a}_j|\pi_i)$$

$$\hat{P}_j(\boldsymbol{s},\boldsymbol{a}|\pi_i) = \sum_{k\neq j}^{N_i} K^s_{h_1}(\boldsymbol{s},\boldsymbol{s}_k)K^a_{h_2}(\boldsymbol{a},\boldsymbol{a}_k)$$

C. Predicting the Final Object Pose

To predict the resulting probability distribution of the cube after a regrasp action, we will use kernel *conditional* density estimation. We formulate our conditional density estimate using our kernels from above and roughly following Hall, Racine and Li[37]:

$$P(\boldsymbol{s}'|\boldsymbol{s},\boldsymbol{a},\mathrm{gra}) = \frac{P(\boldsymbol{s}',\boldsymbol{s},\boldsymbol{a}|\mathrm{grasped})}{P(\boldsymbol{s},\boldsymbol{a}|\mathrm{grasped})}$$

$$P(\boldsymbol{s}',\boldsymbol{s},\boldsymbol{a}|\mathrm{gra}) = \frac{1}{m}\sum_i^m K^s_{h_3}(\boldsymbol{s}',\boldsymbol{s}'_i)K^s_{h_4}(\boldsymbol{s},\boldsymbol{s}_i)K^a_{h_5}(\boldsymbol{a},\boldsymbol{a}_i)$$

$$P(\boldsymbol{s},\boldsymbol{a}|\mathrm{gra}) = \frac{1}{m}\sum_i^m K^s_{h_4}(\boldsymbol{s},\boldsymbol{s}_i)K^a_{h_5}(\boldsymbol{a},\boldsymbol{a}_i)$$

We will choose values for h_3, h_4 and h_5 that minimize the integrated squared error, again using cross validation (see [37] for more details):

$$ISE(h_3,h_4,h_5) = \int\left(\hat{P}_{ssa} - P_{ssa}\right)^2 P_{sa}\,d\boldsymbol{s}\,d\boldsymbol{a}d\boldsymbol{s}'$$

where

$$P_{ssa} = P(\boldsymbol{s}'|\boldsymbol{s},\boldsymbol{a},\mathrm{grasped}) \quad \text{and} \quad P_{sa} = P(\boldsymbol{s},\boldsymbol{a}|\mathrm{grasped})$$

Note that for both $P(\mathrm{grasped}|\boldsymbol{s},\boldsymbol{a})$ and $P(\boldsymbol{s}'|\boldsymbol{s},\boldsymbol{a},\mathrm{grasped})$, we could have chosen more complex kernels or used different learning algorithms. However, in this paper we select rather simple models to understand the viability of a data-driven framework for modeling regrasps. In our future work, we plan to evaluate different non-parametric methods for estimating these probability functions.

IV. DATA COLLECTION

Our data collection setup is shown in Figure 2. For our experiments, we use an ABB IRB 140 industrial robot arm and a Robotiq C-85 2-fingered gripper that place and pick an object from a metal platform. We use a 50mm wooden cube as our object. Initially, the block is resting on the platform and the robot locates it and picks it up. The vision system records the initial state $\boldsymbol{s}$. Then, the robot places the cube and picks it up again using an action $\boldsymbol{a}$. The $\boldsymbol{a}$ parameters $[d, z, \alpha]$ are sampled uniformly at random and cover the entire range of actions we wish to model for this regrasp. The vision system first checks whether or not the cube is in the robot's hand and then records the final state $\boldsymbol{s}'$. If the object is grasped, it repeats the process with a new action $\boldsymbol{a}$. If the object is not grasped, it enters a recovery procedure and then runs a new regrasp experiment after that. In this way, we collect a series of $D = (\boldsymbol{s},\boldsymbol{a},\mathrm{grasped},\boldsymbol{s}')$ data points. In this paper, we collected 3304 data points. The robot successfully maintained its grasp of the object after a regrasp 2642 times, and failed 662 times.

The vision system consists of four Microsoft Kinect v2 sensors arranged so that we have multiple views of the object both on the platform and in the robot's hand. Depth point clouds are fused together and after an initialization, Iterative Closest Point is used to find the closest match between our object model and the point cloud. In practice, we were able to achieve accuracies on the order of 5mm.

The recovery procedure is split into 2 parts. If the object is not in the robot's hand, it is either resting on the platform, or has fallen off the platform. If it is resting on the platform, we command the robot to pick up the object and continue with the next experiment. If the object has fallen off of the platform, we consider the object lost, and grasp a new block from a queue of identical blocks resting on the table.

V. VALIDATION

We now compare our learned model with an off the shelf simulator and a rudimentary geometric model on our data set D. We randomly select a hold out test set of 1000 data points, which we will use to compare all 3 methods. We describe the geometric and simulation models below.

A. Geometric Model

The challenging part of this regrasp to model is what happens during step 2 of Figure 1, as there are many possible contact modes including: no contact followed by an impact and settling, sliding / pivoting in finger tips, sliding / rotating against the platform. For simplicity, we'll say that during step 2, once an object corner contacts the platform, the object rotates about the contact point until it is lying flat on the platform. Given an initial object pose $s = [x, y, \theta]$, if c is the distance from the center of the hand to the edge of the platform when placing and w is the width of the block, we can calculate the distance from the edge of the platform to the center of the block q as:

$$q = \begin{cases} \theta \geq 0, & c + y + \frac{w}{2}(\sin(\theta) - \cos(\theta) + 1) \\ \theta < 0, & c + y + \frac{w}{2}(\sin(\theta) + \cos(\theta) - 1) \end{cases}$$

Now, given an action $a = [d, z, \alpha]$, if g is the maximum horizontal distance away from the center of the block that the robot can still grasp the object without missing it, then we will successfully grasp the block if $|d - q| \leq g$, and that final pose will be:

$$\boldsymbol{s}' = \begin{bmatrix} x' \\ y' \\ \theta' \end{bmatrix} = \begin{bmatrix} (q-d)\cos(\alpha) + (z - w/2)\sin(\alpha) \\ (q-d)\sin(\alpha) - (z - w/2)\cos(\alpha) \\ \gamma \end{bmatrix}$$

$$\text{with} \quad \gamma = \begin{cases} 0 \leq \alpha \leq \pi/4, & \alpha \\ \pi/4 \leq \alpha \leq \pi/2, & (\alpha - \pi/2) \end{cases}$$

Creating probability distributions from this kinematic model is difficult, as we do not know the distribution of errors on our parameters. Note that even if we did, even for this rudimentary model, the probability distributions would be multimodal and non-Gaussian.

B. VREP with ODE as a Simulation Model

Figure 4 shows our simulation framework with VREP [38]. We have put in an ABB IRB 140 robot with a Robotiq C-85 2-fingered gripper just as in our real experimental setup. The platform is placed in the same location, and we use a 50mm cube with the same density and frictional properties as our real wooden cube. We can now place the object into the simulated robot's hand at a given initial state $\boldsymbol{s}$, ask the robot to perform the regrasp action $\boldsymbol{a}$, and then we can observe whether the object was grasped, and if so, what the final state $\boldsymbol{s}'$ was. Note that getting the simulator to work was a challenge in and of itself. Even the well-tuned ODE in VREP still cannot handle parallel grasping well, and once the block is also made to slide against the table and in the hand, it is difficult to get stable results.

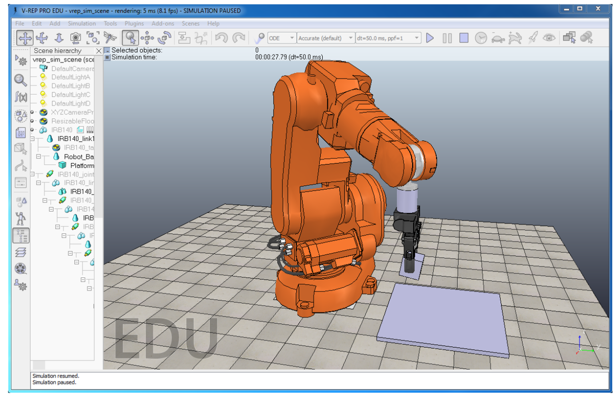

Fig. 4. Simulation environment used in the paper. Using VREP, we have put in the same industrial robot arm and gripper we are used in our real experiments, along with placing the platform in the same relative location. We place the block in the simulated robot's hand in the same initial pose as our real trials, and record whether or not the regrasp succeeds in simulation, and if it does, what the final pose of the block is.

In our opinion, the two most difficult phenomena to model in simulation / with physics is 1) how the contact patch between the parallel jaws and the cube changes as the hand slightly loosens its grip on the object, and 2) what happens to the cube at the onset of contact with the platform.

Again, creating probability distributions using a simulator is difficult, as the simulation is deterministic. We could vary initial parameters slightly and observe results, however it is unclear how much to vary parameters by in order to get plausible results.

Classification Accuracy			
	Geometric	Simulation	Data-Driven
In Hand	74.8 %	72.8 %	76.2 %
Platform	89.3 %	-	90.7 %

TABLE I

PREDICTING WHETHER THE OBJECT IS STILL GRASPED FOR DIFFERENT MODELS

C. Validation Results: Predicting if the Object is Still Grasped

Table I show our results for predicting if the object is still grasped after a regrasp action. We compared the classification accuracy of the three models for two separate conditions. First, we look at the condition where we are given the pose of

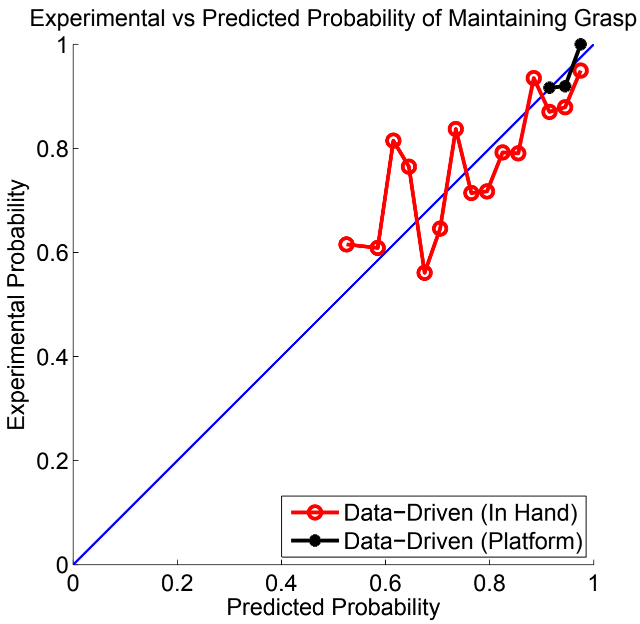

Fig. 5. Comparison between the experimental and predicted probability of maintaining the object after a regrasp

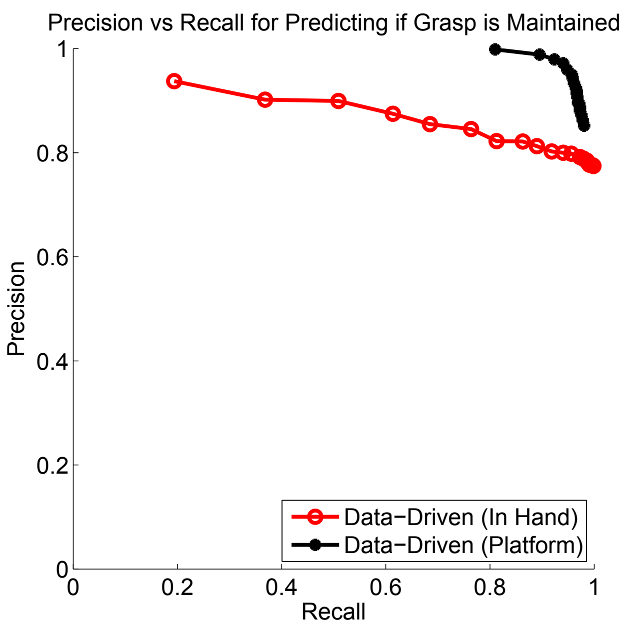

Fig. 6. Precision-recall curve for predicting whether the robot is still holding the object after a regrasp

the object in the robot's hand and the parameterized regrasp action it is to perform. Second, we look at the condition where we know the pose of the object on the platform, and predict the probability of whether or not the robot will be able to successfully pick it up. All three models perform comparably, even though our data-driven model is given no prior information about the task.

In Figure 5, we have binned our predicted probability, and then looked at the percentage of those points where the object was still grasped, and plotted the results. If our predictions are good, the mean of the true grasp probability should follow the straight line. Our predicted probabilities for the platform condition match better than the in-hand condition, which is expected. If we can predict the probability of maintaining the object after a regrasp, this means we can adjust the decision boundary to achieve different precision and recall values. This is plotted in Figure 6. Note that the platform case gives us a much better precision-recall curve, and that these precision-recall curves are not easily achievable without a data-driven model.

Mean Pose Estimation Accuracy (mm)			
	Geometric	Simulation	Data-Driven
In Hand	11.7	10.8	13.0
Platform	5.7	-	6.3

TABLE II
MEAN POSE ESTIMATION ACCURACY AMONG DIFFERENT MODELS

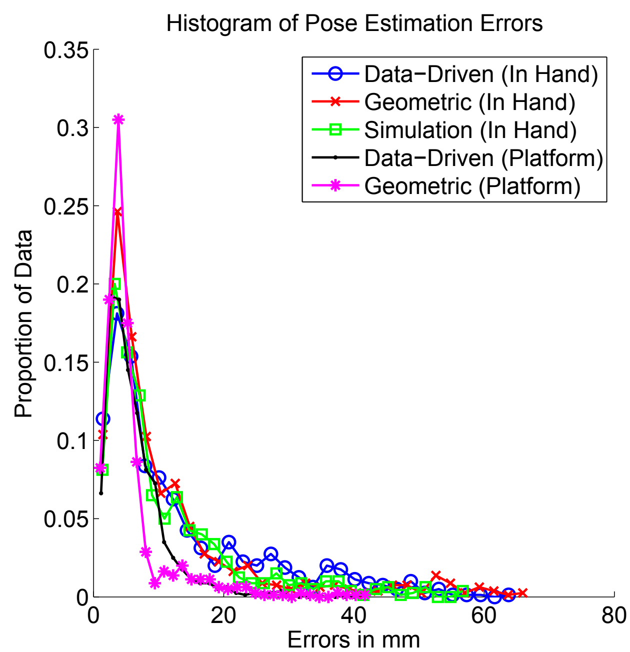

Fig. 7. Histogram of Prediction Errors. Note that most of the error is less than 10mm, and the data-driven approach achieves comparable accuracy to the other approaches.

D. Validation Results: Final Pose Estimation

To evaluate the predictive power of our pose estimation models, we looked at the mean pose estimation accuracy. We used the square-root of our distance function $D_s(s_1, s_2)$ as a measure of accuracy. Note that if the distribution is multi-modal, this measures does not reward capturing that multi-modality. However, since we do not have the true underlying distribution, we use the mean pose estimation accuracy as a

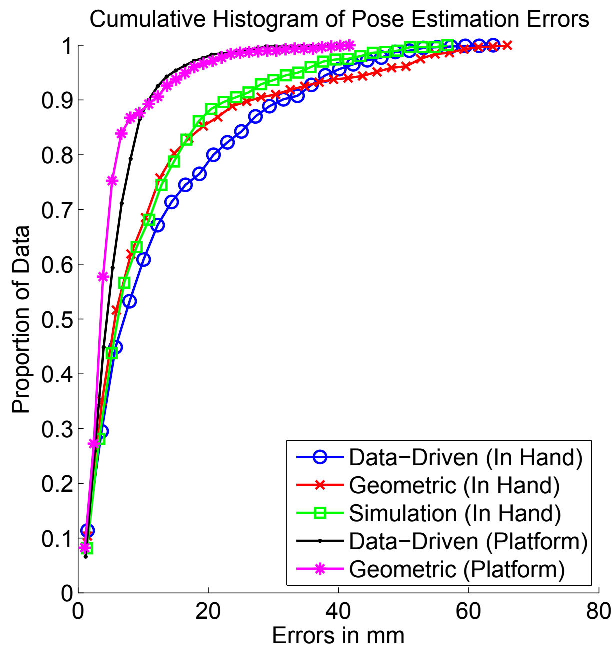

Fig. 8. Cumulative Histogram of Prediction Errors. Over 80% of the data has an error of less than 10mm for the platform case. The data-driven method achieves comparable accuracy to the other approaches.

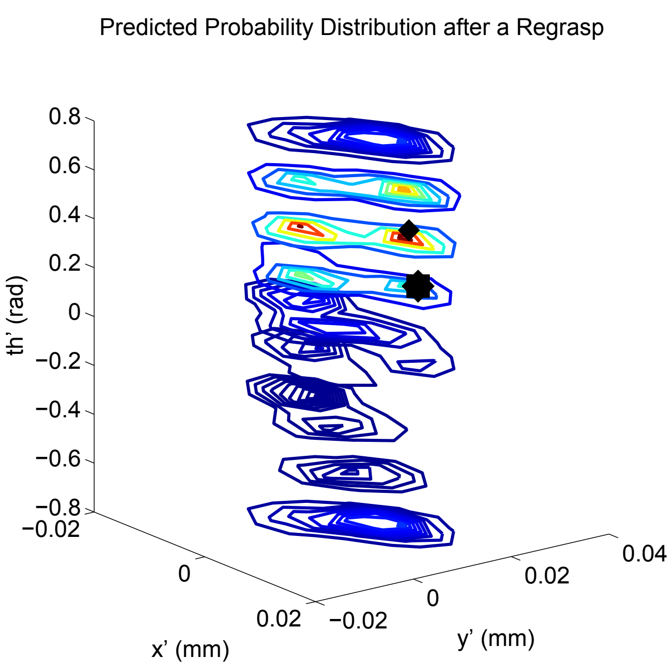

Fig. 9. Probability Distribution of a Predicted Final Pose. Predicted pose is the small diamond and the true pose is the larger star. Note the multi-modal nature of the distribution.

baseline. Our results are shown in Table II, Figure 7 and Figure 8. Again, our data-driven model achieves comparable accuracy with no prior information.

With our data-driven model, we can also calculate the entire resulting probability distribution in pose space, which is shown in Figure 9. Note the multi-modal nature of the distribution.

VI. CONCLUSIONS

In this paper, we introduced a way to model robotic regrasping using a large amount of real data. First, we briefly discussed how we collected the real robot manipulation data needed for our models. We then showed how to predict the probability of maintaining the grasp of an object given an initial position and robot regrasp action using this data. In addition, we showed how to estimate the probability distribution of where the object will end up in the robot's hand given an initial pose and a robotic regrasp action. We compared our models with a simulator and a rudimentary physics model and showed that our data-driven models have comparable performance even with no prior knowledge of the task.

In the future, we are interested in extending these models to other objects, regrasp actions, and hands. We are especially interested in extending our models to SE(3) space to handle three dimensional rigid body transformations. We are also interested in exploring other non-parametric models in an attempt to achieve higher fidelity.

REFERENCES

[1] R. P. Paul, "Modeling, trajectory calculation and servoing of a computer controlled arm." Stanford Artificial Intelligence Laboratory, Technical Report AIM– 177, 1972.

[2] P. Tournassoud, T. Lozano-Pérez, and E. Mazer, "Regrasping," in *IEEE International Conference on Robotics and Automation*, 1987, pp. 1924–1928.

[3] R. S. Fearing, "Simplified grasping and manipulation with dextrous robot hands," *Robotics and Automation, IEEE Journal of*, vol. 2, no. 4, pp. 188–195, 1986.

[4] D. L. Brock, "Enhancing the dexterity of a robot hand using controlled slip," in *Robotics and Automation, 1988. Proceedings., 1988 IEEE International Conference on*. IEEE, 1988, pp. 249–251.

[5] S. Stoeter, S. Voss, N. P. Papanikolopoulos, and H. Mosemann, "Planning of regrasp operations," in *Robotics and Automation, 1999. Proceedings. 1999 IEEE International Conference on*, vol. 1. IEEE, 1999, pp. 245–250.

[6] K. Harada, T. Foissotte, T. Tsuji, K. Nagata, N. Yamanobe, A. Nakamura, and Y. Kawai, "Pick and place planning for dual-arm manipulators," in *Robotics and Automation (ICRA), 2012 IEEE International Conference on*. IEEE, 2012, pp. 2281–2286.

[7] A. Cole, P. Hsu, and S. S. Sastry, "Dynamic control of sliding by robot hands for regrasping," *Robotics and Automation, IEEE Transactions on*, vol. 8, no. 1, pp. 42–52, 1992.

[8] T. Schlegl, M. Buss, T. Omata, and G. Schmidt, "Fast dextrous re-grasping with optimal contact forces and contact sensor-based impedance control," in *Robotics and Automation, 2001. Proceedings 2001 ICRA. IEEE International Conference on*, vol. 1, 2001, pp. 103–108 vol.1.

[9] N. Furukawa, A. Namiki, S. Taku, and M. Ishikawa, "Dynamic regrasping using a high-speed multifingered hand and a high-speed vision system," in *Robotics and Automation, 2006. ICRA 2006. Proceedings 2006 IEEE International Conference on*. IEEE, 2006, pp. 181–187.

[10] J. K. Salisbury Jr., "Kinematic and Force Analysis of Articulated Hands," PhD Dissertation, Stanford University, 1982.

[11] J. Hong, G. Lafferriere, B. Mishra, and X. Tan, "Fine manipulation with multifinger hands," in *Robotics and Automation, 1990. Proceedings., 1990 IEEE International Conference on*. IEEE, 1990, pp. 1568–1573.

[12] T. Schlegl and M. Buss, "Hybrid closed-loop control of robotic hand regrasping," in *Robotics and Automation, 1998. Proceedings. 1998 IEEE International Conference on*, vol. 4. IEEE, 1998, pp. 3026–3031.

[13] M. Farooqi, T. Tanaka, Y. Ikezawa, T. Omata, and K. Nagata, "Sensor based control for the execution of regrasping primitives on a multifingered robot hand," in *Robotics and Automation, 1999. Proceedings. 1999 IEEE International Conference on*, vol. 4. IEEE, 1999, pp. 3217–3223.

[14] M. Cherif and K. K. Gupta, "Planning quasi-static fingertip manipulations for reconfiguring objects," *Robotics and Automation, IEEE Transactions on*, vol. 15, no. 5, pp. 837–848, 1999.

[15] D. Rus, "In-hand dexterous manipulation of piecewise-smooth 3-d objects," *The International Journal of Robotics Research*, vol. 18, no. 4, pp. 355–381, 1999.

[16] N. Chavan-Dafle, A. Rodriguez, R. Paolini, B. Tang, S. Srinivasa, M. Erdmann , M. T. Mason , I. Lundberg, H. Staab, and T. Fuhlbrigge, "Extrinsic dexterity: In-hand manipulation with external forces," in *IEEE International Conference on Robotics and Automation (ICRA)*, May 2014.

[17] T. Lozano-Perez, M. T. Mason, and R. H. Taylor, "Automatic synthesis of fine-motion strategies for robots," *The International Journal of Robotics Research*, vol. 3, no. 1, pp. 3–24, 1984.

[18] M. Erdmann, "Using backprojections for fine motion planning with uncertainty," *The International Journal of Robotics Research*, vol. 5, no. 1, pp. 19–45, 1986.

[19] R. C. Brost, "Automatic grasp planning in the presence of uncertainty," *The International Journal of Robotics Research*, vol. 7, no. 1, pp. 3–17, 1988.

[20] M. Dogar and S. Srinivasa, "A framework for push-grasping in clutter," *Robotics: Science and Systems VII*, 2011.

[21] K. Y. Goldberg and M. T. Mason, "Bayesian grasping," in *Robotics and Automation, 1990. Proceedings., 1990 IEEE International Conference on*. IEEE, 1990, pp. 1264–1269.

[22] R. C. Brost and A. D. Christiansen, "Probabilistic analysis of manipulation tasks: A conceptual framework," *The International journal of robotics research*, vol. 15, no. 1, pp. 1–23, 1996.

[23] Y. Bekiroglu, J. Laaksonen, J. A. Jørgensen, V. Kyrki, and D. Kragic, "Assessing grasp stability based on learning and haptic data," *Robotics, IEEE Transactions on*, vol. 27, no. 3, pp. 616–629, 2011.

[24] M. Li, Y. Bekiroglu, D. Kragic, and A. Billard, "Learning of grasp adaptation through experience and tactile sensing," in *Intelligent Robots and Systems (IROS 2014), 2014 IEEE/RSJ International Conference on*. Ieee, 2014, pp. 3339–3346.

[25] A. Petrovskaya, O. Khatib, S. Thrun, and A. Y. Ng, "Bayesian estimation for autonomous object manipulation based on tactile sensors," in *Robotics and Automation, 2006. ICRA 2006. Proceedings 2006 IEEE International Conference on*. IEEE, 2006, pp. 707–714.

[26] K. Hsiao, M. Ciocarlie, and P. Brook, "Bayesian grasp planning," in *ICRA 2011 Workshop on Mobile Manipulation: Integrating Perception and Manipulation*, 2011.

[27] L. E. Zhang and J. C. Trinkle, "The application of particle filtering to grasping acquisition with visual occlusion and tactile sensing," in *Robotics and Automation (ICRA), 2012 IEEE International Conference on*. IEEE, 2012, pp. 3805–3812.

[28] M. C. Koval, N. S. Pollard, and S. S. Srinivasa, "Pose estimation for planar contact manipulation with manifold particle filters," *The International Journal of Robotics Research*, p. 0278364915571007, 2015.

[29] C. Goldfeder and P. K. Allen, "Data-driven grasping," *Autonomous Robots*, vol. 31, no. 1, pp. 1–20, 2011.

[30] D. Kappler, J. Bohg, and S. Schaal, "Leveraging big data for grasp planning," in *Robotics and Automation (ICRA), 2015 IEEE International Conference on*. IEEE, 2015, pp. 4304–4311.

[31] F. Stulp, E. Theodorou, J. Buchli, and S. Schaal, "Learning to grasp under uncertainty," in *Robotics and Automation (ICRA), 2011 IEEE International Conference on*, May 2011, pp. 5703–5708.

[32] A. Morales, E. Chinellato, A. H. Fagg, and A. P. Del Pobil, "Using experience for assessing grasp reliability," *International Journal of Humanoid Robotics*, vol. 1, no. 04, pp. 671–691, 2004.

[33] R. Paolini, A. Rodriguez, S. Srinivasa, and M. T. Mason, "A data-driven statistical framework for post-grasp manipulation," *International Journal of Robotics Research (IJRR)*, vol. 33, no. 4, pp. 600–615, April 2014.

[34] M. Kopicki, S. Zurek, R. Stolkin, T. Mörwald, and J. Wyatt, "Learning to predict how rigid objects behave under simple manipulation," in *Robotics and Automation (ICRA), 2011 IEEE International Conference on*. IEEE, 2011, pp. 5722–5729.

[35] M. Rosenblatt, "Remarks on some nonparametric estimates of a density function," *Ann. Math. Statist.*, vol. 27, no. 3, pp. 832–837, September 1956.

[36] J. S. Marron *et al.*, "Optimal rates of convergence to bayes risk in nonparametric discrimination," *The Annals of Statistics*, vol. 11, no. 4, pp. 1142–1155, 1983.

[37] P. Hall, J. Racine, and Q. Li, "Cross-validation and the estimation of conditional probability densities," *Journal of the American Statistical Association*, vol. 99, no. 468, 2004.

[38] M. Freese, S. Singh, F. Ozaki, and N. Matsuhira, "Virtual robot experimentation platform v-rep: a versatile 3d robot simulator," in *Simulation, modeling, and programming for autonomous robots*. Springer, 2010, pp. 51–62.