A measurement model for tracking hand-object state during dexterous manipulation

Craig Corcoran and Robert Platt Jr.

Abstract— It is frequently accepted in the manipulation literature that tactile sensing is needed to improve the precision of robot manipulation. However, there is no consensus on how this may be achieved. This paper applies particle filtering to the problem of localizing the pose and shape of an object that the robot touches. We are motivated by the situation where the robot has enclosed its fingers around an object but has not yet grasped it. This might be the case just prior to grasping or when the robot is holding on to something fixtured elsewhere in the environment. In order to solve this problem, we propose a new model for position measurements of points on the robot manipulator that tactile sensing indicates are touching the object. We also propose a model for points on the manipulator that tactile measurements indicate are *not* touching the object. Finally, we characterize the approach in simulation and use it to localize an object that Robonaut 2 holds in its hand.

I. Introduction

One of the fundamental barriers to autonomous robot manipulation in unstructured environments is perception. Estimating the combined state of the manipulator and the objects that the robot touches has proven to be difficult. Attempts to estimate hand-object configuration visually are hampered by occlusions. Instead, force and tactile sensing is a natural way to track the combined state of the manipulator and the objects acted upon during manipulation. Although this type of measurement is information-poor relative to camera images, it does not suffer from occlusions and it has the potential to enable more precise position and force estimates than is possible using only visual information. This paper focuses on the problem of tracking the pose and shape of an object that a robot holds between compliant fingers. We are motivated by the situation where the robot has enclosed its fingers around an object but has not yet grasped it. This might be the case just prior to grasping or when the robot is holding on to something fixtured elsewhere in the environment. Our objective is to localize the object more precisely using force and position sensing in order to assist subsequent interactions with the object. The result should be more accurate than the visual estimate alone and it should account for displacements caused by the manipulator itself.

Particle filtering is a statistical approach to robust nonlinear state estimation that is well suited to the problem of tracking object configuration based on a series of force and position measurements [1], [2], [3]. In [1], the authors applied Markov localization to the problem of localizing the planar pose (3 DOFs) of an inelastic fixtured part based on tactile measurements. The likelihood of a particular position was calculated by numerically integrating a Gaussian error model over the contour of the object. The resulting likelihood model was stored in a look-up table. Similarly, Petrovskaya *et. al.* also localize an inelastic object by making repeated contact with a single end-effector. In this work, localization occurred in the space of spatial object poses (6 DOFs) [2]. The high-dimensional likelihood space caused Petrovskaya *et. al.* to propose a variant of particle filter annealing that iteratively increases measurement model entropy while decreasing the search space. In related work, Chhatpar and Branicky apply particle filtering to the problem of localizing the pose of a peg with respect to a hole [3]. In contrast to the above, their work samples measurements from across the state space on-line rather than creating an analytical model for the measurement distribution.

This paper expands on the work described above by proposing a new measurement model that can be used to track hand-object configuration during manipulation. We are specifically interested in the case where the pose of the robot hand is known but the object configuration is unknown and possibly moving in the hand. A key step in particle filtering is the weighting phase where each particle is weighted depending upon how likely the observed measurement would be if the system were in the state hypothesized by that particle. Our measurement model integrates the likelihood of a contact position measurement over the space of all possible true contact positions on the surface of the object. This is different than that approach taken in [2] where the likelihood of an observed contact is evaluated with respect to the maximum likelihood point on the surface of the object. This paper also proposes using negative information obtained from contacts that are known *not* to be touching the object. As in the above, we integrate the likelihood of the negative contact measurement over the volume outside the object. Also, we briefly describe a dynamic approach to particle filter annealing that enables us to localize in five dimensions using the particle filter. Finally, after demonstrating the advantage of the new measurement model in simulation, we apply the approach to a practical object localization problem where we estimate the pose of an object that is captured by the Robonaut 2 hand. The Robonaut 2 experiments demonstrate that the approach is a viable solution to an important practical problem in humanoid robotics.

Craig Corcoran is an undergraduate in the Mechanical Engineering department at Rice University. (ccor@rice.edu)

Robert Platt Jr. is a Research Scientist in the Computer Science and Aritificial Intelligence Laboratory at MIT. This research was performed while he was at NASA Johnson Space Center. (rplatt@csail.mit.edu)

II. BAYESIAN FILTERING

The goal of Bayesian filtering is to track the state of a stochastic system as it changes. It is assumed that state, x, is a stochastic Markov function of time. At every timestep, the measurements, z, depend only on the current state. Starting with a prior distribution over state, $P(x_0)$, Bayesian filtering recursively updates a posterior distribution, $P(x_t|z_{1:t})$, where x_t is the state at time t and $z_{1:t} = \{z_1, \ldots, z_t\}$ is the set of measurements between time 1 and time t. The update to the posterior (also called the "belief state") is accomplished in two steps. First, the prediction step updates the distribution by applying a system model:

$$P(x_t|z_{1:t-1}) = \int P(x_t|x_{t-1})P(x_{t-1}|z_{1:t-1})dx_{t-1} \quad (1)$$

The above uses the Markov assumption that $P(x_t|x_{t-1}, z_{1:t-1}) = P(x_t|x_{t-1})$. In the second step, the posterior distribution is updated in proportion to the likelihood of having generated the observed measurements, z_t:

$$P(x_t|z_{1:t}) = \eta P(z_t|x_t)P(x_t|z_{1:t-1}), \quad (2)$$

where

$$\eta = \frac{1}{P(z_t|z_{1:t-1})}$$

is a normalizing constant.

Equations 1 and 2 constitute an optimal solution to the problem of tracking state in a Markov system. However, they ignore the question of how the posterior distribution is represented. Two popular solutions to this problem are the Kalman filter and the particle filter. The Kalman filter is optimal, but makes strict (linear system, Gaussian noise) assumptions regarding the system and measurement models. The particle filter does not make these assumptions, but relies on Monte Carlo methods that depend on an adequate sampling the posterior distribution. This paper uses the sample importance resampling (SIR) version of the particle filter [4] to track hand-object state.

III. LIKELIHOOD OF CONTACT POSITIONS

In the context of mobile robot localization, range measurements are functions of the relative configuration of the robot in the environment. Since relative robot configuration is generally the variable of interest, the likelihood of the measurements can be used to infer robot configuration. However, in manipulation, contact positions are *not* functions of object configuration alone. These measurements also depend on manipulator configuration. Rather than requiring a model of how the manipulator and the object interact, we evaluate the likelihood of contact positions by integrating over all possible manipulator configurations.

A. General case

Let x describe the object configuration (*i.e.* shape and pose). Let R be a set of contact positions on the robot manipulator equipped with force sensors that measure whether a point $r \in R$ is touching the object or not. Let $p = (p_1 \ldots p_P) \subseteq R$ be the portion of contacts that are touching.

Let $q = (q_1 \ldots q_Q) \in R - p$ be the contacts that are not touching. Assume that $\hat{p}$ and $\hat{q}$ are noisy measurements of p and q, but that there is no uncertainty regarding membership in p and q. This corresponds to an assumption that perfect contact force sensors determine whether a contact is touching or not but that the position of the contacts cannot be measured accurately. In our hardware experiments (see Section VI), the uncertainty in the position measurements was caused by modeling inaccuracies in the manipulator geometry and kinematics.

Assume that the contact position measurements, $\hat{p}$ and $\hat{q}$, are independent and identically distributed given object configuration, x. Then the likelihood of the measurement can be written as a product:

$$P(\hat{p}, \hat{q}|x) = \prod_{i=1}^{P} P(\hat{p}_i|x) \prod_{j=1}^{Q} P(\hat{q}_j|x). \quad (3)$$

If an accurate model of the manipulator-object interaction were available, then the likelihood of a given position measurement could be evaluated in terms of its proximity to an expected position measurement: $P(\hat{p}_i|model(x, u))$, where $model(x, u)$ denotes the expected contact position given an object configuration x and manipulator control parameters, u. However, since the ultimate position of manipulator contacts on an object is a complex function of the second-order impedances of the manipulator and object, creating such a model can be prohibitively difficult. Instead, we propose a simpler (but less informative) measurement model created by integrating over all possible contact positions as a function of object pose:

$$P(\hat{p}_i|x) = \int P(\hat{p}_i|p_i)P(p_i|x)dp_i.$$

In principle, p_i depends on *both* x and u, and we should integrate over u:

$$P(p_i|x) = \int P(p_i|u, x)P(u|x)du.$$

However, in the absence of a model, assume that $P(p_i|x)$ is uniformly distributed over all possible contact positions on the surface of the object:

$$P(\hat{p}_i|x) = \int_{p \in \delta S(x)} P(\hat{p}_i|p)dp, \quad (4)$$

where $\delta S(x)$ is the set of points on the surface of the object. Similarly, since the contact points, q, do not touch the object, assume that they are uniformly distributed over the set of possible contact positions outside of the object (but within a gross region about the object):

$$P(\hat{q}_i|x) = \int_{p \in \bar{S}(x)} P(\hat{q}_i|q)dq, \quad (5)$$

where $\bar{S}(x)$ is the finite set of points outside of but within a gross region of the object.

B. Positive contact on a polyhedron

In general, Equations 4 and 5 have no closed form solution and must be evaluated numerically or approximated by a simple surface. This section explores approximations to Equation 4 for polyhedrons. Let F be the set of faces that comprise the polyhedron. Then Equation 4 becomes:

$$P(\hat{p}_i|x) = \sum_{f \in F} P(\hat{p}_i|f), \qquad (6)$$

with

$$P(\hat{p}_i|f) = \int_{p \in f(x)} P(\hat{p}_i|p)dp, \qquad (7)$$

where $f(x)$ is the set of positions on face f when the object is in configuration x. Suppose the contact position measurement noise is Gaussian:

$$\begin{aligned} P(\hat{p}_i|p_i) &= \mathcal{N}(\hat{p}_i|p_i, \Sigma) \\ &= \mathcal{N}(p_i|\hat{p}_i, \Sigma), \end{aligned} \qquad (8)$$

where $\mathcal{N}(\cdot|\mu, \Sigma)$ denotes the normal distribution about μ with a covariance matrix, Σ. For each facet, define an orthonormal basis described by the rotation matrix, $R_f = (t_x, t_y, n)$, where n is a basis vector normal to f and t_x and t_y span the plane containing f. Let

$$p_t = \begin{pmatrix} t_x^T \\ t_y^T \end{pmatrix} p$$

be the projection of p onto the plane and let

$$p_n = n^T p$$

be the projection onto the normal. As a result, we can write:

$$\begin{aligned} P(\hat{p}_i|p) &= P(p|\hat{p}_i) \\ &= P(p_t|p_n, \hat{p}_i)P(p_n|\hat{p}_i). \end{aligned} \qquad (9)$$

$P(p_t|p_n, \hat{p}_i)$ and $P(p_n|\hat{p}_i)$ can be evaluated using standard Gaussian manipulation techniques [5]. The Gaussian distribution in Equation 8 can be treated as a joint distribution over p_t and p_n with a covariance matrix:

$$\Sigma = \begin{pmatrix} \Sigma_{tt} & \Sigma_{tn} \\ \Sigma_{nt} & \Sigma_{nn} \end{pmatrix}.$$

Then

$$P(p_n|\hat{p}) = \mathcal{N}(p_n|\hat{p}_n, \Sigma_{nn}), \qquad (10)$$

and

$$P(p_t|p_n, \hat{p}_i) = \mathcal{N}(p_t|\hat{p}_{t|n}, \Sigma_{t|n}), \qquad (11)$$

where

$$\Sigma_{t|n} = \Sigma_{tt} - \Sigma_{tn}\Sigma_{nn}^{-1}\Sigma_{nt},$$

and

$$\hat{p}_{t|n} = \hat{p}_t - \Sigma_{tn}\Sigma_{nn}^{-1}(p_n - \hat{p}_n).$$

Substituting Equations 10 and 11 into Equation 7, we have:

$$\begin{aligned} P(\hat{p}_i|f) &= \int_{(p_t, p_n) \in f(x)} \mathcal{N}(p_t|\hat{p}_{t|n}, \Sigma_{t|n})\mathcal{N}(p_n|\hat{p}_n, \Sigma_{nn}) \\ &= \mathcal{N}(f_n|\hat{p}_n, \Sigma_{nn}) \int_{p_t \in f_t(x)} \mathcal{N}(p_t|\hat{p}_{t|n}, \Sigma_{t|n}), \quad (12) \end{aligned}$$

(a)　　　　　　　　　(b)

Fig. 1. The dotted line in (a) shows a hypothetical path of a contact measurement in the neighborhood of a planar rectangle. (b) illustrates the likelihood of the measurement along this path. The dip in the likelihood function occurs as the path turns the corner.

where $f_t(x)$ is the set of tangent coordinates of the points in face f and $f_n = p_n \in f_n(x)$ is the constant position of points in face f measured along the normal vector, n.

If the facet f is indefinite, then the integral in Equation 12 goes to one and

$$P(\hat{p}_i|f) = \mathcal{N}(p_n|\hat{p}_n, \Sigma_{nn}). \qquad (13)$$

If the facet, f, is bounded by a rectangle,

$$p_t \in R^2 : p_t^x \in [x_h, x_l] \wedge p_t^y \in [y_h, y_l],$$

and $\Sigma_{t|n}$ is isotropic, then Equation 12 becomes:

$$P(\hat{p}_i|f) = \frac{1}{4}\mathcal{N}(p_n|\hat{p}_n, \Sigma_{nn})Y_f X_f, \qquad (14)$$

where

$$Y_f = \left[erf\left(\frac{y_h - \hat{p}_y}{\sqrt{2}\sigma_y}\right) - erf\left(\frac{y_l - \hat{p}_y}{\sqrt{2}\sigma_y}\right) \right],$$

and

$$X_f = \left[erf\left(\frac{x_h - \hat{p}_x}{\sqrt{2}\sigma_x}\right) - erf\left(\frac{x_l - \hat{p}_x}{\sqrt{2}\sigma_x}\right) \right],$$

and erf denotes the error function and σ_x and σ_y are the singular values of Σ associated with eigenvectors directed along the rectangle axes, t_x and t_y. Note that in order to apply this technique on multiple differently oriented rectangular faces, the requirement for $\Sigma_{t|n}$ to be isotropic essentially requires Σ to be isotropic. Also, note that $P(\hat{p}|f)$ tends toward Equation 13 as the facet becomes larger relative to σ_x and σ_y.

Figure 1 illustrates the behavior of the likelihood function in the neighborhood of a planar rectangle. The dotted line in Figure 1(a) illustrates a hypothetical path of a contact measurement, $\hat{p}$, through the space around the rectangle. Figure 1(b) illustrates the likelihood function for the path. The dip in Figure 1(b) occurs as the path turns the corner and shows that the likelihood of the measurement decreases in the neighborhood of the corner.

C. Negative contact on a polyhedron

Whereas it is possible to give a good closed-form approximation of the likelihood of contacts touching a polyhedral object, there is no similar closed form expression for the likelihood of negative contacts (contacts that do not touch such the object). Our analysis in Section III-B was possible because of the constraint that the Gaussian was integrated over rectangular surfaces. However, it is not possible to use this method to evaluate the integral over the space outside of the object unless the object itself is rectangular. Rather than considering only rectangular objects, we propose approximating the likelihood function by integrating over an appropriate half plane.

We are interested in integrating over the half plane (in Cartesian 3-space) outside the object that contains the largest part of the probability mass, $\mathcal{N}(q|\hat{q}, \Sigma)$. Consider the set of half planes bounded by planes that contain the faces of the polyhedron. If $\hat{q}$ is inside the object, then the largest part of the probability mass is contained in the half plane associated with the closest face. If $\hat{q}$ is outside the object, then this is the half plane that contains $\hat{q}$ and is bounded by the plane furthest from $\hat{q}$. Let $c(f)$ be the plane that contains that face $f \in F$. Let $c^+(f)$ be the half plane associated with $c(f)$ that does *not* contain the object. Let $h(\hat{q}) = \{f | \hat{q} \in c^+(f)\}$ be the set of faces that bound half planes containing $\hat{q}$. Let $d(f, \hat{q})$ be the distance from $\hat{q}$ to $c(f)$. We integrate $\mathcal{N}(q|\hat{q}, \Sigma)$ over the half plane outside the object associated with the following face:

$$f^*(\hat{q}) = \begin{cases} \arg\max_{f \in h(\hat{q})} d(f, \hat{q}) & \text{if } \hat{q} \text{ outside object.} \\ \arg\min_{f \in F} d(f, \hat{q}) & \text{if } \hat{q} \text{ inside object.} \end{cases}$$
(15)

Now, we expand Equation 5 by integrating over the positive half plane for face $f^*(\hat{q})$:

$$P(\hat{q}|x) = \int_{(q_t, q_n) \in f^*(\hat{q})} \mathcal{N}(q_t|\hat{q}_{t|n}, \Sigma_{t|n}) \mathcal{N}(q_n|\hat{q}_n, \Sigma_{nn}).$$

Since we are integrating over the entire half plane, the tangent integral goes to one and we have:

$$\begin{aligned} P(\hat{q}_i|x) &= \int_{q_n^*}^{\infty} \mathcal{N}(q_n|\hat{q}_n, \Sigma_{nn}) dq_n \\ &= \frac{1}{2}\left[1 - erf\left(\frac{q_n^* - \hat{q}_n}{\sqrt{2}\sigma_n}\right)\right], \end{aligned}$$
(16)

where q_n^* is the normal coordinate of face $f^*(\hat{q})$ and σ_n is the square root of the variance in the normal direction (again, we have assumed an isotropic covariance matrix, Σ).

One interesting point about using negative contact information is that *all* hand surfaces are known to be outside the object – not just those hand surfaces equipped with force sensors that indicate they are not touching. If a large number of appropriate negative contact surfaces are used, then in principle, good object localization is possible just using negative information and without using contact force sensors at all. Essentially, these negative contacts roughly encode the geometry of the hand or manipulator and extrapolate the object configuration based on the available free space.

Combining positive and negative contact information is a unified way of combining the information about where the manipulator touches the object and the available free space.

IV. MAINTAINING PARTICLE DIVERSITY USING DYNAMIC ANNEALING

Until this point, the proposed measurement model can be applied equally well in the context of a particle filter, Kalman filter, monte carlo maximum likelihood estimate, or a different form of inference. Exactly which method should be used depends on the exact nature of the localization problem. For problems where object configuration is known to be fixed, a filtering solution should be discarded in favor of a maximum likelihood or maximum a priori estimate. The scaling series approach in [2] performs inference in a six dimensional space. However, since we are interested primarily in tracking the unknown motions of an object captured by the robot hand, our focus is on a filtering solution.

Although the particle filter has had success in three-dimensional tracking problems, it is not clear that it is suitable for localization problems in five, six, or higher dimensions (as in the present object localization scenario). The key problem is ensuring the particle set maintains a diversity suitable for the level of "confidence" present in the system. When system state is uncertain, a higher entropy particle set is appropriate while a lower entropy sampling improves the accuracy of a highly-confident track. Deutscher *et. al.* address this problem in the context of human motion tracking (a very high dimensional state space) by proposing an annealing technique where the entropy of the measurement distribution is gradually decreased over time by taking the distribution to an increasing power [6]. While this approach assists initial localization, it does not help if the track gets "lost" because it is impossible to increase the expected entropy of the sample set. While this is irrelevant to the problem of locating a static object, it is important when the object is moving in an unknown way in the robot hand.

We address this problem with a dynamic annealing approach that adjusts measurement model entropy as a function of the normalized likelihood of the most recent measurements. Large measurement likelihoods indicate that the particle set is distributed in a likely region of space and it is possible to decrease measurement model entropy. Small measurement likelihoods indicate that the particles are not focused in likely regions of space and a higher entropy distribution is needed in order to "find" the peaks. We control the entropy of the distribution by varying the eigenvalues of the measurement model covariance matrix, Σ, in Equations 16 and 14 between a minimum, σ_{min}, and maximum, σ_{max}. This happens in inverse exponential proportion to the measurement likelihood of Equation 3. Assuming an isotropic covariance matrix, Σ, let σ^2 denote the variance of Σ. Then set σ according to:

$$\sigma = (\sigma_{max} - \sigma_{min})exp\left(-\frac{P(\hat{p}, \hat{q}|x) - S_{min}}{S_{max} - S_{min}}\right) + \sigma_{min},$$
(17)

Fig. 2. Illustration of simulation setup. Three fingers (the lines) tracked the rotating motion of a rectangle (in the plane) by applying small simulated inward forces.


Fig. 3. Comparison of the maximum likelihood model (dotted line) for positive contacts with the likelihood model proposed in this paper (solid line). Results are averaged over ten runs for each likelihood model.

where S_{min} and S_{max} are convenient minimum and maximum values for the likelihood, $P(\hat{p}, \hat{q}|x)$.

V. SIMULATIONS

The first experiment compared the positive contact measurement model proposed in this paper to the maximum likelihood model used in [2] and [3]. Let

$$p^* = \arg \max_{p \in \delta S(x)} \mathcal{N}(p|\hat{p}, \Sigma)$$

be the most likely point on the object surface given the position measurement, $\hat{p}$. Then, under the maximum likelihood model, the likelihood of this position measurement is:

$$P(\hat{p}|x) = \mathcal{N}(\hat{p}|p^*, \Sigma).$$

Figure 2 illustrates the experimental setup. A three finger manipulator touches a moving rectangle two inches wide and one inch high. The fingers apply small inward forces such that the contacts always touch the object but do not impede its motion. The objective is to localize and track the rectangle using measurements of the three fingertip positions. The rectangle rotates between $-\frac{\pi}{7}$ and $\frac{2\pi}{5}$ radians about an interior point in 86 time steps while the "palm" position remains fixed. Localization occurred over the three-dimensional space of planar object poses. The height and width of the rectangle were assumed to be known.

Simulation results are illustrated in Figure 3. The figure shows localization error (measured as an L2 norm in state space) averaged over ten trials for identical runs using the maximum likelihood model and the proposed measurement model. The results show that measurably better performance is obtained by the proposed model. Beyond that, a couple of features are apparent. First, localization convergence takes 30 or 40 time steps. This is surprising since, in principle, a minimum of only three measurements are needed to localize the planar rectangle. However, note that three of the *right* measurements are needed. For example, when the manipulator is in the configuration shown in Figure 2, it is impossible to accurately localize displacement along the long axis of the box. Therefore, complete localization only


Fig. 4. Robonaut 2.

occurs after one of the contacts moves over the a corner of the rectangle. The other reason that localization takes several steps is that the particle filter update occurs only once at each time step. Although it is possible to execute multiple filter updates on each time step, notice that position measurements are plentiful in this scenario and there is no need to conserve this information. It is a better use of computational resources to track using the latest data. The second feature that is apparent in Figure 3 is that the localization error actually begins to increase after time step 50. It turns out that this is an artifact of measurement aliasing after timestep 50. During this period, the rectangle was in a configuration relative to the contacts similar to that shown in Figure 2. As a result, it was impossible for the system to localize position error along the long axis of the rectangle and error slowly began to integrate. This would continue until a contact again turned a corner of the rectangle.

Fig. 5. Robonaut 2 hand. The black dots indicate the contacts used in the experiments.


Fig. 6. Localization error for position (cyan), radius (green), and orientation (blue) as a function of time (tenths of a second). The top plot shows error; the bottom plot shows standard deviation of the particle cloud. The results are averaged over the nine different tube orientations shown in Table I.

VI. EXPERIMENTS WITH HARDWARE

Experiments were performed using Robonaut 2, shown in Figure 4. All hardware experiments used the positive and negative contact measurement likelihood models described in Sections III-B and III-C. Although our goal is to track object pose in dynamic scenarios (for example, when the object moves in the hand), our experiments consider the converse problem: localizing a fixed object by moving the hand over the object surface. This is slightly easier than the original problem because the robot measures hand velocities relative to the object. If the object were moving in the hand, this information might not be available. Our experiments evaluate the efficacy of localization when the Robonaut 2 hand makes the same sequence of compliant moves for nine different relative tube configurations.

A. Setup

In all hardware experiments, Robonaut 2 interacted with a piece of rubber tubing fixtured to the ground approximately 1.5 inches in diameter. Robonaut 2 is equipped with actively compliant fingers [7] that allow the stiffness of the finger joints to be controlled programmatically. In addition, Robonaut 2 has torque-controlled arm joints that similarly allow the stiffness of the palm Cartesian position to be specified programmatically. This arm and hand compliance enabled Robonaut 2 to compliantly move along the surface of the tube. After wrapping its fingers around the object, the manipulator reference configuration was adjusted according to a fixed pattern that pulled the manipulator approximately along the length of the tubing in a twisting motion. Because of the manipulator compliance, the resulting motion of the robot hand was a function of the cylinder pose and radius.

The pose and radius of the rubber tubing was estimated using a particle filter operating in a six-dimensional state space comprised of 5 pose DOFs (no axial orientation) and one dimension describing radius. Pose localization occurred in five dimensions rather than four (the pose configuration space of an infinite cylinder is only four dimensional) for programming convenience. Since the measurement models pro-

posed in this paper operate on polyhedra only, the shape of the tubing was approximated by an infinite length prismatic hexagon for the purposes of localization. Figure 5 illustrates the contact points used during localization. The six contact positions on the index finger, middle finger, and thumb are equipped with embedded force sensors that measure contact. These positions were included in $\hat{p}$ when the corresponding force sensor registered an above-threshold force and in $\hat{q}$ otherwise. The contact position on the palm did not have a force sensor and was always assumed to be out of contact.

B. Experiment

The robot interacted with the tubing in the nine different relative configurations shown in Table I. In each localization trial, the tube was fixtured in a different orientation and the robot executed the same sequence of compliant motions. The filter used a set of 1000 particles. Figure 6 shows localization error averaged over nine trials in each of the different relative configurations. Position measurements were made every 0.1 seconds. The results show that particle variance has converged after 20 iterations of the filter (two seconds of data). Localization converges to a position and radius error approximately one tenth of an inch. Orientation error converges to approximately 8 degrees. Orientation error seems large because the kinematics of the hand-tube system are poorly configured to measure orientation. When the robot grasps the tube, a change of 8 degrees in tube orientation results in little movement of the contacts. (One might imagine placing two hands on a tube to measure its orientation more precisely.)

The underlying cause of localization error was a result of modeling errors measuring contact location. First, rather than calculating the exact contact location on the (complex) surface geometry of the finger, our experiments simplified localization by assuming contact locations at the center of

Label	A	B	C	D	E	F	G	H	I
Angle	0°	8°	16°	8°	16°	8°	16°	8°	16°
Axis	NA	z	z	z	z	x	x	x	x

TABLE I

THE NINE DIFFERENT TUBE ORIENTATIONS USED IN EXPERIMENT 1. THE x AND z AXES ARE AN ORTHOGONAL BASIS PERPENDICULAR TO THE CYLINDER AXIS WHEN IT IS THE ORIENTATION IS AT ZERO.


Fig. 7. Comparison of localization using only positive contact measurements (dashed line) and localization using positive and negative contact measurements (solid line). The results show orientation error averaged over ten runs in tube orientation F. The top plot shows error; the bottom plot shows standard deviation of the particle cloud.

the finger on the corresponding phalanges. If all finger contact surfaces were spherical, then this assumption would be accurate for a corresponding extruded object. However, for the actual Robonaut 2 hand, this approximation is clearly a source of error. More significant, however, were errors caused by incorrect kinematic modeling or incorrect joint calibration. Given the large number of joints in the hand and the particularly complex kinematics of the thumb, maintaining a very accurate kinematic model in the context of period recalibration of the finger joint angle sensors proved to be difficult.

Figure 7 compares the accuracy of localization using only positive contact measurement and using both positive and negative contact information. The cylinder was fixtured in configuration F (see Table I). Performance of localization using only positive contact information improves when negative information is incorporated.

VII. CONCLUSIONS

This paper considers the problem of hand-object state estimation during mechanical interactions between the robot hand and the object. We are particularly interested in localizing a partly or incompletely grasped part that is moving in an unknown way. This capability could be used to identify or characterize objects that the robot touches. Or, it could be used during grasping to confirm that the robot is holding the part correctly and to provide information about how to adjust the grasp. Or, it could enable the robot to achieve a desired hand-object relative pose in the context of a task or assembly. The paper makes three main contributions. First,

we provide a new model of the likelihood of contact position measurements and demonstrate a measurable improvement in localization accuracy. Second, we propose modeling negative contact information to improve localization. Finally, we demonstrate that the methods can be used to localize an object touched by a humanoid robot hand.

REFERENCES

[1] K. Gadeyne and H. Bruyninckx, "Markov techniques for object localization with force-controlled robots," in *10th Int'l Conf. on Advanced Robotics*, 2001.
[2] A. Petrovskaya, O. Khatib, S. Thrun, and A. Ng, "Bayesian estimation for autonomous object manipulation based on tactile sensors," in *IEEE Int'l Conf. on Robotics and Automation*, 2006, pp. 707–714.
[3] S. Chhatpar and M. Branicky, "Particle filtering for localization in robotic assemblies with position uncertainty," in *IEEE Int'l Conf. on Intelligent Robots and Systems*, 2005, pp. 3610– 3617.
[4] S. Arulampalam, S. Maskell, G. N., and T. Clapp, "A tutorial on particle filters for on-line non-linear/non-gaussian bayesian tracking," *IEEE Transactions on Signal Processing*, vol. 50, pp. 174–188, 2001.
[5] C. Bishop, *Pattern nrecognition and machine learning*. Springer, 2006.
[6] J. Deutscher, A. Blake, and I. Reid, "Articulated body motion capture by annealed particle filtering," in *IEEE Int'l Conf. on Computer Vision and Pattern Recognition*, vol. 2, 2000, pp. 126–133.
[7] M. Abdallah, C. Wampler, R. Platt, and B. Hargrave, "Applied joint-space torque and stiffness control of tendon-driven manipulators," in *ASME 2010 International Design Engineering Technical Conferences and Mechanisms and Robotics Conference (Submitted)*, 2010.