

Using Bayesian filtering to localize flexible materials during manipulation

Robert Platt Jr., Frank Permenter, Joel Pfeiffer

Abstract—Localizing and manipulating features such as buttons, snaps, or grommets embedded in fabrics and other flexible materials is a difficult robotics problem. Approaches that rely too much on sensing and localization that occurs before touching the material are likely to fail because the flexible material can move when the robot actually makes contact. This paper experimentally explores the possibility of using proprioceptive and load-based tactile information to localize features embedded in flexible materials during robot manipulation. In our experiments, Robonaut 2, a robot with human-like hands and arms, uses particle filtering to localize features based on proprioceptive and tactile measurements. Our main contribution is to propose a method of interacting with flexible materials that reduces the state space of the interaction by forcing the material to comply in repeatable ways. Measurements are matched to a “haptic map”, created during a training phase, that describes expected measurements as a low-dimensional function of state. We evaluate localization performance when using proprioceptive information alone and when tactile data is also available. The two types of measurements are shown to contain complementary information. We find that the tactile measurement model is critical to localization performance and propose a series of models that offer increasingly better accuracy. Finally, the paper explores this localization approach in the context of two flexible materials insertion tasks that are relevant to manufacturing applications.

I. INTRODUCTION

Flexible materials manipulation is an important class of problems. Many “general assembly” tasks in automobile factories that are currently performed by humans involve installing cables, carpets, and flexible plastics. Similarly, manufacturing clothing, shoes, and other soft goods is labor-intensive because robots are unable to manipulate flexible materials reliably. Aside from its practical value, studying flexible materials manipulation is interesting for its own reasons because many existing approaches cannot easily be applied to the problem. It is admittedly possible to manipulate flexible material without estimating the state of the interaction once manipulation has begun (for example, see the towel folding work in [1]). However, if there is no mechanism for tracking state during manipulation, then there is no possibility of reacting to unforeseen events. Given that the system is already interacting with the object, it is natural to attempt to use a sense of touch to track state.

This paper applies ideas used in mobile robot localization to manipulation. There is a strong analogy: whereas the goal of



Fig. 1. Robonaut 2 hand localizing a bump in a piece of flexible plastic.

mobile robot localization is to track the position of the robot in the environment, the goal of manipulation localization is to track the position of the object held in the hand. Also, the kind of information available from range sensors or landmark bearing estimates is of a similar complexity to that which is available from touch sensors. Our basic approach is to interact with a known object during a controlled training phase whereby a map is created that describes how the material “feels.” Then, during localization, the touch measurements are matched to the map using Bayesian filtering. Many approaches to flexible materials state estimation utilize high-dimensional models of the space of possible material deformations (for example [2], [3]). Instead, a key insight of this paper is that it is frequently possible to manipulate a flexible material in such a way that it always deforms in a certain way. As a result, it is possible to reduce the dimensionality of the model by assuming that this deformation always takes place. Our work applies this idea to the problem of localizing “haptic features” such as buttons, grommets, or snaps in flexible materials through touch.

The details of this approach are explored experimentally using Robonaut 2 [4] for three features embedded in flexible materials: a bump in flexible plastic, a snap in fabric, and a grommet in fabric (see Figure 4). Two types of touch information are considered: proprioceptive measurements (the configuration of a compliant hand during manipulation) and tactile measurements using load-based sensors. We experimentally characterize the localization accuracy using proprioceptive information alone and demonstrate that an improvement is possible by also incorporating tactile information. We evaluate the dimensionality in the tactile data that contains information relevant to localization and show that the information con-

Robert Platt Jr. is with the Computer Science and Artificial Intelligence Laboratory at MIT. This work was performed while he was at NASA Johnson Space Center. rplatt@csail.mit.edu

Frank Permenter is with Oceanering Space Systems. frank.n.permenter@nasa.gov

Joel Pfeiffer is with the Computer Science Department at Purdue University. jpfeiffer@purdue.edu

tained in the tactile data is qualitatively different from that in the proprioceptive data. Finally, we demonstrate an additional improvement in performance that results from modeling the tactile data as a mixture of Gaussians. Bringing the pieces together, we are able to demonstrate an expected localization accuracy of less than 0.2 inches using a combination of proprioceptive information and load-based tactile information. The practical advantages of the approach are illustrated in the context of two insertion tasks (see Figures 13 and 14). This paper is an expanded and more complete review of this work relative to [5].

A. Related Work

This paper is one of the first to consider the problem of tactile state estimation while manipulating a flexible material. Nevertheless, there is a large body of relevant prior work. The problem of localizing inflexible objects using tactile information has received considerable attention from a number of different intellectual directions. An early approach considers the problem of localizing an object with unknown object shape parameters by fitting contact position and surface normal measurements to a model [6], [7]. Noting that object shape is known in many practical situations, Jia and Erdmann propose an application of observability theory that estimates the contact position and pose of a known object when single point contact is made [8]. Okamura and Cutkosky take a geometric approach to localizing surface features on inflexible objects using haptic exploration [9].

Recently, there has been an effort to apply Bayesian filtering to the problem of localizing inelastic objects through touch interactions. Chhatpar and Branicky apply particle filtering to the problem of localizing the pose of a peg with respect to a hole [10]. Rather than using an analytical measurement model, they create a model during a training phase where the robot slides the peg over the hole in a series of “sweeps” [11], [12]. This approach to interacting with a material (especially the notion of “sweeping” over a material) is related to the fabric interaction procedure described in this paper (Section II-B). In [13], Petrovskaya *et al.* 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) using a particle filter and a maximum likelihood measurement model. Gadeyne and Bruyninckx take a similar approach where Markov localization is applied to the problem of localizing the planar pose (3 DOFs) of an inelastic fixtured part based on tactile measurements [14]. In this work, the measurement model incorporated a numerical integration step. Corcoran and Platt found an analytic solution to the above integration for polyhedral objects and use it to realize spatial object localization using contact position information [15].

Much flexible material manipulation literature focuses on knot tying, and surgical suturing in particular. Remde *et al.* perform a comprehensive analysis of the contact states and feasible transitions that can occur for a deformable linear object (a rope or cable) [16]. As pointed out in [17], it is not strictly necessary to model the material compliance in order to plan knots [18], [19]. However, planning techniques that take the




Fig. 2. Robonaut 2 hand. Note the three tactile sensor caps on each finger (one cap on each phalange).

flexible dynamics into account have more broad applications. One way of incorporating better material models into the planning process is to calculate low-energy states for the material given end-point configurations and plan accordingly [2], [20], [21]. Wakamatsu and Hirai consider the more general problem of manipulation planning for arbitrary flexible objects [22]. However, this work assumes linear strain dynamics. Tian and Jia propose a non-parametric extension of the above linear model [17]. Their work also considers the grasping problem where the ramifications of object deformation on grasp point selection is explicitly considered.

Another related body of work is concerned with flexible materials modeling. This is important in computer graphics as well as robotics applications. A standard approach models the deformable object using a set of small masses that interact with each other through springs or other potential function elements [23], [24], [3], [25]. For example, Burion *et al.* find mass-spring parameters that generate model deformations that best fit a series of mechanical tests performed on the object using a particle filter [3]. Morris and Salisbury find parameters for a potential function-based model that are damped and generate object geometries closest to what is observed [25].

II. SYSTEM AND SETUP

This section introduces the finger tactile sensors and finger torque control and then describes the interaction scenario.

A. Tactile sensors and finger torque control

The tactile sensors used in this work are composed of strain gauges mounted in the load path between the contact surfaces of the Robonaut 2 (R2) finger and the finger structure through which contact loads are reacted to the ground [26]. Figure 2 shows the basic structure of the hand. Notice that each finger has three “contact caps” on it – one cap on each phalange. Each of these caps is mounted to a spring element instrumented with strain gauges. Strain gauges are small patches of silicone or metal that measure mechanical strain and are affixed to surfaces on the load path. When a




Fig. 4. The three features embedded in flexible materials used in the experiments.




Fig. 3. Internals of the tactile load cell used in the experiments.

load is applied to an elastic material (aluminum or steel, for example), the load causes elastic deformations in the material that can be measured using strain gauges. The principle of operation is that when the R2 hand touches something (for example, refer to Figure 1), it is these caps that actually make contact with the environment. When this occurs, the sensors in the spring element measure the load. Figure 3 illustrates the spring element itself. Notice that it has a roughly cylindrical shape that facilitates mounting on the human-sized R2 finger. The spring element is grounded to the robot finger at the edges of the cylinder and attached to the contact shell by a center plate with two screw holes. Each tactile sensor produces a total of eight signals. No two different loads applied to the sensor produce the same measurements. In order to minimize the effects of uncontrolled variables such as temperature or mechanical shifts in the sensor itself, the vector of signals produced by a single sensor is normalized on every timestep.

Since the R2 hand is extrinsically actuated (it is driven by motors located in the forearm), it is necessary to actuate the tendons in order to realize joint torques.

where \dot{x} is the vector of tendon velocities, $\dot{\theta}$ is the internal tendon velocity, q is the vector a finger joint positions, \dot{q} is the vector of joint velocities, and P is full rank and non-diagonal in general.

Following [27], our control law calculates a desired tendon position, x_d , that decouples joint velocities:

$$x_d = x - k_d \dot{x} + P^T K_p (\tau_d - P f),$$

where x describes tendon positions, \dot{x} describes tendon velocities, f describes tendon tensions, P describes the linear relationship between tendon velocities and joint velocities, and K_p and k_d are the PD parameters of the torque controller. This control law moves the tendons so as to maintain the desired torque, τ_d . If a joint stiffness is desired rather than a contact torque, the desired torque is a function of joint

position: $\tau_d = K(q_d - q)$. Finger joint positions are measured using hall sensors on the output of each joint. The arm joint positions are measured using accurate optical absolute position sensors. All the joint position sensors are calibrated relatively accurately. Hand position estimates relative to the base frame are accurate to within 0.25 inches. Fingertip position estimates relative to the palm are accurate to within hundredths of an inch.

B. Interaction scenario

One key idea of this work is to interact with the flexible material such that it deforms in repeatable ways. As a result, it is unnecessary to model all possible deformations of the material. We only need to model the particular interaction scenario illustrated in Figures 1 and 5(a). In this scenario, the flexible material loosely hangs from a test rig such that it swings freely with respect to the robot in different directions. The robot grasps the material between its thumb and forefingers (index and middle fingers). The forefingers apply a constant light squeezing force against the thumb which is held fixed. Then, the robot pulls its hand away from the fixed point in the direction of the arrow illustrated in Figure 5(a). We will refer to a single pull as a “swipe.” Each swipe covers a distance of typically two or three inches at a speed of approximately 1.3 inches per second. During each swipe, the thumb is commanded to hold a constant position with large stiffnesses in its four joints. In the index and middle fingers, the adduction/abduction and the medial/distal joints are commanded to hold fixed positions with large stiffnesses. The proximal finger joints apply a constant closing torque such that each finger pushes with approximately 0.75 Newtons against the thumb in the direction of closing (see Figure 5(a)).

As the hand pulls, the material is squeezed between the thumb and fingers so that it complies with the hand in a particular and repeatable way. As the fingers move over the material, proprioceptive and tactile sensor measurements respond to the mechanical stiffness characteristics of the material. Haptic features such as buttons or grommets have mechanical properties different from that of the surrounding material. As a result, we expect to be able to localize these features based on sensor measurements. The evaluations in this paper are performed for the three features illustrated in Figure 4. The bump in Figure 4(a) is used to fixture the flexible plastic in the context of a factory assembly task. The snap in Figure 4(b) and the grommet in Figure 4(c) are embedded in

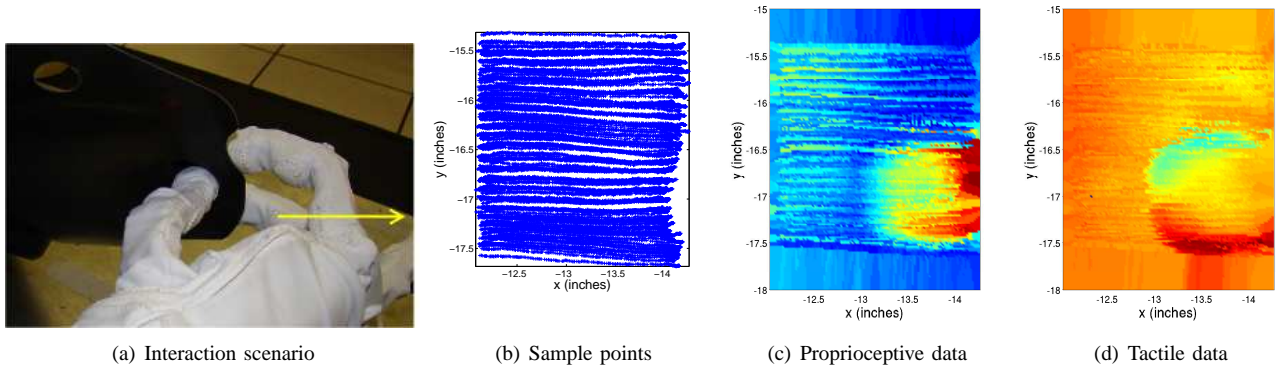


Fig. 5. Illustration of the training phase. (a) illustrates the robot hand performing a “swipe” from left to right. (b) shows the state locations of the roughly 25000 state-measurement sample pairs collected during training over a 2.25×2.25 inch patch on the surface of the material. (c) illustrates a single proprioceptive signal (distance between the middle fingertip and the thumb, color variation: 0.01 to 0.97 inches) over the state space. (d) illustrates a single tactile sensor signal (from the middle fingertip force sensor, color variation: -0.457 to +0.351 volts) over the state space.

a simulated thermal blanket that is an important part of many extra-vehicle NASA tasks.

III. LOCALIZATION

When the robot interacts with a haptic feature such as a button or grommet, it “feels” a characteristic signal that enables it to localize the feature. We consider two types of sensor information: *proprioceptive* information that measures finger displacements and force sensor information that directly senses the magnitude and direction of loads applied to the finger.

A. Training Phase

During training, a haptic “map” is created that associates each point in state space with a measurement. Since the material is squeezed between the thumb and forefingers, we know already that it is touching the thumb and that it is locally tangent to the finger and thumb surfaces at the point of contact. The remaining dimensions of uncertainty describe where contact occurs on the surface of the material. We parametrize the surface of the material by a local two-dimensional coordinate frame. State is defined to be the location of the thumb tip in this coordinate frame. During training, a corpus of data is collected that pairs state with proprioceptive and force sensor measurements in the context of the swipe interaction described earlier. In order to obtain accurate measurements of state, the material is held in a jig so that it is roughly immobile with respect to the base frame. Then, data is collected by systematically performing a series of swipes so that the entire region of interest has been “scanned.” In this paper, each “scan” consists of 23 swipes. A series of scans are performed during training. Then, while the material remains immobilized in the jig, one or more additional “test” scans are performed. The fact that the material location is exactly the same during training and testing enables us to evaluate localization accuracy by comparing localization estimates with the known position of the hand and the material during the test scans.

Figures 5(b), 5(c), and 5(d) illustrate data collected during a training session. The data corpus represented consists of approximately 25000 state-measurement pairs sampled in the locations indicated in Figure 5(b) over approximately a 2.25×2.25 inch patch in the neighborhood of the plastic bump. The data was collected by performing 6 “scans” of the entire region. Each of the 23 swipes in a scan are approximately 0.1 inches apart. Each swipe consists of approximately 182 data points collected approximately 0.011 inches apart. In principle, one would expect this procedure to generate samples in a series of parallel lines 0.1 inches apart. However, stiction, Coriolis, and inertial effects in the robot arm joints as well as forces generated by interaction with the material cause the variation evident in Figure 5(b). Figures 5(c) and 5(d) illustrate an example of a proprioceptive signal and a force sensor signal sampled from a nearest neighbor function on a regularly spaced grid defined over the 2.25×2.25 inch patch where each grid cell is a 0.02×0.02 inch square. Each point in the grid takes the measurement value of the nearest sample in the corpus. Figure 5(c) shows the distance between the thumb and middle finger. The measurement values range between 0.01 inches (blue) and 0.97 inches (red). Figure 5(d) shows the response of one of the tactile sensor signals in the middle finger. Although the physical quantity being measured is strain, we only report the voltage response of the sensor because the sensor is uncalibrated. Voltage is related to strain through an unknown (but constant) linear parameter. The measurement values range between -0.457 volts (blue) and 0.351 volts (red). As one might expect, the two types of measurements are aligned. The same forces that cause the thumb and middle finger to separate as they travel over the bump are also recorded by the force sensor. Notice that the proprioceptive data (Figure 5(c)) has the largest response when the middle finger is on top of the bump while the tactile data (Figure 5(d)) has the greatest response on the edges of the bump.

Figures 5(c) and (d) are characterized by variations in measurements that form horizontal lines. Comparison with Figure 5(b) indicates that these lines are associated with the geometry of the scan process during training. If two swipes that are performed nearby to each other at different times

have slightly different measurement responses, then this is manifested by a line. There are two main sources for this variation: sensor error and shifts in the flexible material during training. Sensor error has two effects. First, sensor error in the finger tension sensors causes the finger torque controller to produce slightly different torques, thereby squeezing the material slightly more or less tightly and causing variation in the fingertip load cell measurements. Second, error in the fingertip sensors themselves directly contributes to the variation. This paper models both of the above sources of sensor error as independent and identically distributed (iid) Gaussian noise.

The other main source of variation in the training data is shifts in the position of the flexible material during training. Our training procedure is to fixture the material such that the position of the thumb in the base frame is roughly proportional to state (the position of the thumb in the coordinate frame of the material). If the material is perfectly fixtured with respect to the jig (which is itself fixtured with respect to the ground), then the system should make the same measurements in the same state on average. However, we have observed some degree of uncontrolled shifting in the material during training. These shifts appear to be stochastic in some regions of state space and relatively deterministic in others. For example, when a finger swipes near the edge of a feature, it will stochastically either remain on top of the feature or it will slide off (this effect can be observed on the top edge of the bump in Figure 5(c) where there are a few relatively pronounced lines). Whether the finger slides off or not is stochastic. However, this particular effect only occurs on the edges of the features – in the middle of a bump or in a featureless region of state space, state measurements are likely to be less noisy. This paper handles the possibility of state estimation errors in the training set by modeling the likelihood of a measurement in terms of a neighborhood of states in the training set surrounding the query state. In Section III-C and III-D, we model this likelihood with a Gaussian fit to the measurements from the training set neighborhood. In Section IV, we model the likelihood as a mixture of Gaussians fit to measurements from the neighborhood.

B. Bayesian filtering

One way to localize the unobserved state of the flexible material in the grasp is to use Bayesian filtering. Bayesian filtering is especially appropriate for flexible materials state estimation because it handles noisy observations and process dynamics well. The goal of Bayesian filtering is to track the unobserved state of a stochastic system as it changes. It is assumed that state, x , is Markov. At every time step, 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, $P(x_t|z_{2:t}, u_{1:t-1})$, where x_t is the state at time t and $z_{2:t} = \{z_2, \dots, z_t\}$ is the set of measurements between time 2 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:

$$P(x_t|z_{2:t-1}, u_{1:t-1}) = \int P(x_t|x_{t-1}, u_{t-1})P(x_{t-1}|z_{2:t-1}, u_{1:t-2})dx_{t-1}. \quad (1)$$

In the second step, the posterior distribution is updated in proportion to the likelihood of having generated the observed measurements, z_t :

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

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. Another alternative, the particle filter, does not make these restrictive assumptions. However, it can fail when the particle sample set does not estimate the posterior distribution with sufficient accuracy.

The observation dynamics in the flexible materials domain of this paper that are modeled during the training phase are highly non-linear. (We assume observation noise to be Gaussian but the observation model itself to be non-linear.) As a result, the Kalman filter (or extended Kalman filter) is inappropriate. The experiments in this paper were all performed using the standard sample importance resampling (SIR) version of the particle filter [28] using a 75-particle sample set. At each time step in the SIR particle filter, the process update (Equation 1) is implemented by sampling from the posterior distribution over states conditioned on action. We assume a Gaussian motion model:

$$P(x_{t+1}|u_t) = N(x; f(x_t, u_t), Q), \quad (3)$$

where $x_{t+1} = f(x_t, u_t)$ denotes the nominal process dynamics and Q is the covariance of the process noise (Q is set to $Q = \text{diag}(0.0004)$ in all experiments presented in this paper). The measurement update (Equation 2) is implemented by weighting each of the particles proportional to the measurement likelihood. In order to prevent the sample set from collapsing at one of the modes of the posterior distribution, 13 percent of the particles are chosen uniformly randomly at each time step.

C. Proprioceptive measurements

Bayesian filtering can be used to perform localization using proprioceptive information alone. We encode proprioceptive information in terms of the pairwise distances between the three fingers. Recall that during interaction with the material, only the proximal flexion joints in the index and middle fingers are under torque control. The rest of the joints in the hand are commanded to hold fixed positions with a high stiffness. As a result, there are no more than two dimensions of finger position variation. These two dimensions are represented to the system in terms of the three pairwise distances. Although




Fig. 6. Relative finger positions as a function of palm position. Color denotes the magnitude of each pairwise distance with red indicating a large distance and dark blue indicating a small distance. (a) shows the distance between the tips of the index and middle fingers (color variation: 0.01 – 0.97 inches); (b) shows the same for the index finger and thumb (color variation: 0.04 – 0.93 inches); (c) shows the same for the middle finger and thumb (color variation: 0.09 – 0.96 inches). (d) illustrates average localization performance using only pairwise distance measurements.

this is a redundant representation, the extra data helps average out the sensor and state estimation error in the training set described in Section III-A.

During the measurement update, the particle filter weights each particle by the likelihood of the measurements. The likelihood of a proprioceptive measurement, z_d , given that the system is in state x is modeled by a locally-weighted Gaussian distribution defined with respect to the k states nearest (Euclidean distance) x ,

$$P(z_d|x) = \mathcal{N}(z_d; \hat{z}_d(x), \Sigma_d(x)),$$

where $\mathcal{N}(x; \mu, \Sigma)$ denotes the Gaussian pdf over x with mean, μ , and covariance, Σ . The mean is

$$\hat{z}_d(x) = \frac{1}{k} \sum_{x_i \in N_k(x)} z_d(x_i), \quad (4)$$

where $z_d(x)$ denotes the distance measurement associated with state x in the training set, and $N_k(x) = \{x_1, \dots, x_k\}$ denotes the set of k states nearest (Euclidean distance) to x . The covariance is

$$\Sigma_d(x) = \frac{1}{k} \sum_{x_i \in N_k(x)} (z_d(x_i) - \hat{z}_d) (z_d(x_i) - \hat{z}_d)^T. \quad (5)$$

Notice that we are not fitting a measurement function with constant measurement noise. At a query point, our model estimates both the mean and covariance parameters of the Gaussian based on a local neighborhood of data points in the training set. This model incorporates state uncertainty in the training set. In regions of state space where the average gradient of the measurement function with respect to state is large, Equation 5 calculates a large covariance. In contrast, the locally weighted sample covariance in a region where all neighboring states have the same expected measurement should be similar to the underlying measurement noise. Notice that this approach lumps together measurement covariance caused by material slippage during training with measurement covariance variation intrinsic to the flexible material itself.

Figure 6(a) through (c) shows the neighborhood means for the three pairwise distances as a function of state for the plastic




Fig. 7. Comparison of average localization performance for the three flexible materials shown in Figure 4 when only proprioceptive information is used. Average performance for the flexible bump is the blue solid line, for the snap is the green dashed line, and for the grommet is the black dotted line. Results are aligned with feature location.

bump training set. As in Figure 5(c) and (d), each of these images is sampled over a grid covering a 2.25×2.25 inch area with each grid cell 0.02 inches on a side. The color denotes the magnitude of the pairwise distance averaged over a local neighborhood of 30 nearest neighbors (Equation 4). Figure 6(a) through (c) can be understood intuitively. Either the index finger or the middle finger travel over the bump. When a finger crosses the bump, the bump pushes it away from the thumb. At rest, the middle finger is raised slightly above the index finger. When the middle finger crosses the bump, it moves away from both the index finger and the thumb. When the index finger crosses the bump, it moves away from the thumb and towards the middle finger.

The localization performance of this model using the three pairwise distance measurements for a plastic bump dataset is illustrated in Figure 6(d). This experiment was performed using a test scan collected just following collection of the training scans. Recall that the material was fixtured in the same jig and in the same way during both training and testing. Since material position can be assumed to be constant in both the training and testing data, measurements of hand position dur-

ing testing correspond give us a “ground-truth” measurement of hand-material position relative to the training data. In the experiments, we compare state (the location of the hand relative to the material) estimated using proprioceptive information with the ground-truth measurement of state calculated based on hand position measurements. Figure 6(d) shows localization error averaged over 20 test swipes in an additional test scan of the material. As in training, the test swipes comprising the scan are approximately 0.1 inches apart over approximately a 2.25 square inch area. Approximately 182 measurements are made during each swipe with each measurement an average of 0.011 inches away from its neighbors. A single swipe takes approximately two seconds. The particle filter is updated once per measurement. Error is equal to the L2 distance between weighted average particle location (the mean of the sampled distribution) and the ground truth state. Figure 6(d) shows a fast initial drop in localization error that is caused by the system immediately realizing that it is *not* on top of the bump. After this, localization error begins to fall again between -12.5 and -13 . This is exactly the point where the thumb-index distance begins to change significantly in Figure 6(b). Localization error reaches its minimum between -13.5 and -14 inches. Since the three pairwise distances also reach their maxima in this region, we know that error is minimized when one finger is completely on top of the bump. Average localization error briefly reaches a minimum near 0.25 inches. However, since this low error estimate does not persist, it may be difficult to assure that the particle filter converges with a low error estimate.

Figure 7 shows a comparison with average localization performance for the snap (dashed green line) and the grommet (dotted black line). Training data was collected for these two other features similarly to how the plastic bump data was collected as described in Section III-A. The data are aligned with the center of the feature at zero. Localization error for all three features becomes smallest just before reaching the center of the feature. This suggests that the most relevant measurements are made as the fingers are just beginning to move over the feature. Notice that as the fingers move past the center of the feature, localization error for the bump and snap gets worse while error on the snap remains roughly constant. This suggests that the proprioceptive measurements made after reaching the feature center are less informative for the bump and grommet but continue to be informative for the snap. When the measurements are not informative, notice that our Gaussian noise assumption (Equation 3) causes a gradual increase in the entropy of the distribution, leading to an increase in the expected error. But why are the measurements less informative for the bump and the grommet but not for the snap? Since the grommet is relatively narrow compared with the snap and bump, the fingers quickly leave the surface of the grommet and measurement informativeness drops. For the bump, once the fingers are on top of it, the proprioceptive measurements are equally consistent with any other location on top of the bump. Therefore, there is some flexibility for motion error to integrate once the fingers reach the top of the bump. In contrast to the grommet and the bump, the snap is both large and haptically informative over its entire extent.




Fig. 9. Comparison of average localization performance using proprioceptive measurements alone (the dotted blue line) and average localization performance when both proprioceptive and tactile measurements are used (the black line).

Measurements continue to be informative for the entire time while the fingers are touching the snap.

D. Tactile measurements

The fact that it is possible to achieve localization accuracy of approximately 0.25 inches briefly using only proprioceptive information suggests that it should be possible to do very well if tactile data is incorporated as well. The fingertip tactile sensors provide more descriptive information – although our load-based sensors are still limited because they only provide force and torque information rather than direct information regarding the contours or texture of the surface. As was the case for the proprioceptive measurements, the tactile measurements are also subject to sensor noise that will be assumed to be Gaussian. In addition, the tactile data measurement model must also take into account the state estimate noise caused by shifts in the flexible material. As a result, we model the tactile data as a single Gaussian defined over locally-weighted sample moments:

$$P(z_t|x) = \mathcal{N}(z_t; \hat{z}_t(x), \Sigma_t(x)).$$

The mean is,

$$\hat{z}_t(x) = \frac{1}{k} \sum_{x_i \in N_k(x)} z_t(x_i), \quad (6)$$

where $z_t(x)$ is a function that evaluates to the vector of tactile signals for state x in the training set and $N_k(x)$ is the set of $k = 30$ nearest states. The covariance over the local region is:

$$\Sigma_t(x) = \frac{1}{k} \sum_{x_i \in N_k(x)} (z_t(x) - \hat{z}_t(x)) (z_t(x) - \hat{z}_t(x))^T. \quad (7)$$

Assuming that the proprioceptive and tactile data are conditionally independent given state, the joint likelihood is the product:

$$P(z|x) = P(z_d|x)P(z_t|x). \quad (8)$$

The tactile data can be visualized using a singular value decomposition. We perform the analysis for a grid with 0.02

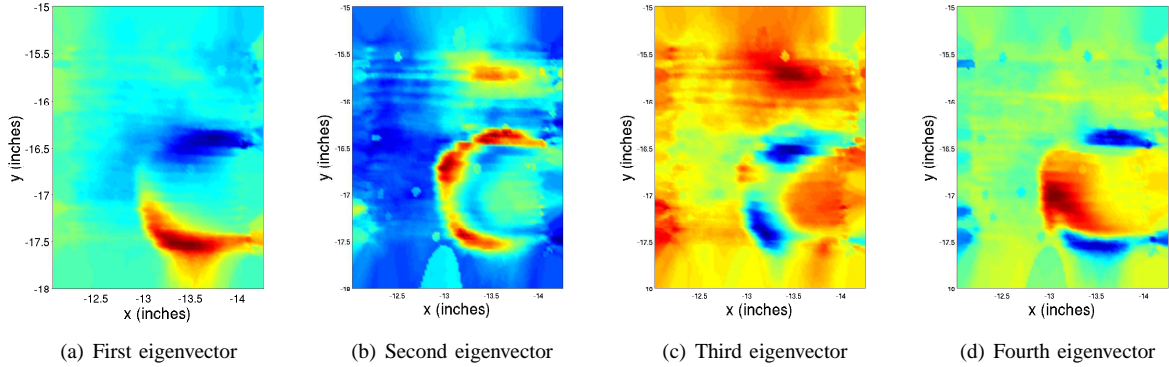


Fig. 8. First four eigenvectors of the tactile data.

square inch cells over a 2.25×2.25 square inch patch (the same patch illustrated in Figures 5 and 6). Let $\mathbf{x} = (x_1, \dots, x_n)^T$ be the vector of $n = 24802$ cells. Let $\hat{z}_t^i(x)$ be i^{th} element of $\hat{z}_t(x)$. Let $\hat{z}_t^i(\mathbf{x}) = (\hat{z}_t^i(x_1), \dots, \hat{z}_t^i(x_n))^T$. For m measurements, the dimensionality of the information contained in the smoothed measurements is the rank of:

$$\Gamma = (\hat{z}_t^1(\mathbf{x}), \dots, \hat{z}_t^m(\mathbf{x})).$$

For the flexible bump training data, the middle fingertip sensor produced seven dimensions of tactile data. The singular values of Γ for this 7×24802 matrix are 1.9361, 1.2055, 1.0716, 0.7418, 0.2446, 0.1883, and 0.0664. The first four eigenvectors are illustrated in Figure 8. A couple of points bear mentioning. First, in contrast to the proprioceptive information (Figure 6(a) through (c)), most of the sensor response occurs on the edges of the bump. Furthermore, the first four eigenvectors respond differently to different parts of the edge of the bump. Using only the first four eigenvectors, it should be possible to do a good job localizing where along the edge of the bump contact with the finger occurs. The plot shows localization error

Figure 9 compares the performance of Bayesian localization using a combination of proprioceptive and tactile data (the solid line) with the performance using just the proprioceptive data (the blue dotted line – same as in Figure 6(d)). The particle filter parameters as well as the flexible plastic bump training and test data sets are the same as those used in Section III-C. Error (L2 norm) is measured with respect to the relative hand-material state during testing (recall that the material continued to be fixtured in exactly the same way during training and testing). As before, these results are averaged over 20 test swipes comprising an additional test scan. The first thing to notice about Figure 9 is that incorporating the tactile data definitely improves localization accuracy – especially between -13 and -13.5 inches. This is consistent with what may be observed by comparing Figures 6 and 8: the tactile data has a larger response earlier than the proprioceptive data. When only proprioceptive information is used, the fingertips must actually be displaced by the feature before localization is possible. The tactile information allows localization to occur while the forces that cause the fingertip displacements are acting. The other notable feature of Figure 9 is that localization performance is actually worse between -13.95 and -14.25




Fig. 10. Measurements associated with the 30 nearest states for a trajectory through state space.

inches. This is counter-intuitive because in the Bayes optimal setting, additional data should only improve the estimate. This suggests that below -13.95 inches, the tactile data likelihood model is inaccurate and causes localization errors. The next section shows that a more accurate tactile measurement model can reduce the impact of this effect.

IV. GAUSSIAN MIXTURE MEASUREMENT MODEL

Until this point, we have modeled state uncertainty in the training set by fitting a single Gaussian to the measurements associated with a neighborhood of training set states about the query point. However, Figure 10 illustrates that this uncertainty is not always Gaussian. Figure 10 shows measurements from one tactile signal in the middle fingertip associated with the $k = 30$ nearest states in the training set for a particular trajectory in state space. In this trajectory, the middle finger skirts the edge of the bump. For states less than -13.2 , there is little variance among the signals of the 30 neighbors. However, during the portion of the trajectory where the finger interacts with the bump, there is a clear bimodal distribution over signals within the neighborhood. Sometimes the finger slips off of the bump and produces the lower trajectory in Figure 10. Sometimes the finger remains on the bump and produces the upper trajectory. Clearly a single Gaussian distribution is a poor fit for this data. Given state uncertainty in the training

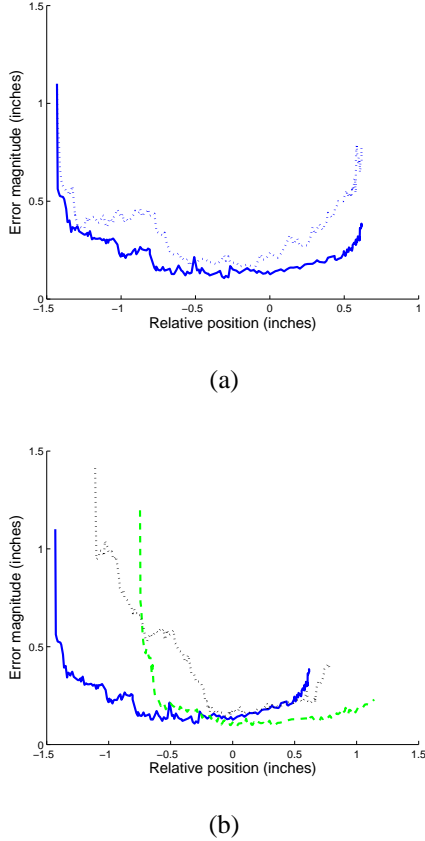


Fig. 11. Performance of the Gaussian mixture measurement model averaged over a test scan consisting of 20 swipes. The solid line in (a) illustrates localization error on the plastic bump for the mixture of Gaussians measurement model. The dotted line shows average localization error for the single Gaussian model (repeated from Figure 9). (b) compares the average performance for the plastic bump (the blue solid line) with the average performance for the snap (the green dashed line) and the grommet (the black dotted line) using the mixture of Gaussians measurement model. The centers of the three features are aligned with zero on the horizontal axis.

set, we need a measurement model that associates some states with a multimodal measurement distribution.

A number of techniques can be used to fit a model to a multimodal distribution. A variant of EM could be used to fit a mixture of Gaussians [29]. Alternatively, Gaussian process regression might be used to fit a non-parametric model [29], [30]. However, this paper leaves these more sophisticated models to future work. Currently, we take a lazy-learning approach that models the multimodal distribution as a Gaussian mixture defined directly over the training data. In particular, we model the likelihood of a tactile measurement vector, z_t , as:

$$P(z_t|x) = \eta \sum_{x_i \in \mathcal{D}_x} \phi(x_i) \mathcal{N}(z_t; z_t(x_i), \Sigma_t), \quad (9)$$

where \mathcal{D}_x is the set of all states in the training data set, $z_t(x)$ is the tactile measurement in the training set corresponding to state x , Σ_t is a user-defined spherical variance, and η is a normalizing constant. $\phi(x_i)$ is a radial basis function that penalizes the contributions from elements of the data set with

associated states that are far from the query state:

$$\phi(x_i) = \mathcal{N}(x_i|x, \Sigma_x),$$

where Σ_x is another user-defined parameter. In the subsequent experiments, we have set $\Sigma_x = \text{diag}(0.0075)$ and $\Sigma_t = \text{diag}(0.001)$.

The results of incorporating this model into Bayesian localization are illustrated in Figure 11. As before, these results are averaged over 20 test swipes. This version of localization is identical with that used in Section III-D except that the likelihood of tactile measurements, $P(z_t|x)$ in Equation 8, is modeled as the Gaussian mixture. Figure 11(a) compares the performance of the Gaussian mixture version of localization (the solid line) with the performance of the single Gaussian model (the dotted blue line) from Figure 9. Just as localization performance was improved by incorporating tactile information in addition to proprioceptive information, performance is again improved by adopting the Gaussian mixture model over a single Gaussian model. Correct localization occurs earlier than it did with the single Gaussian model and there is less integration of error once the fingertips move off the edge of the bump.

Figure 11(b) compares average localization performance for the flexible plastic bump with the average performance for the snap and the grommet. Comparing with Figure 7, the mixture of Gaussians model improves localization performance for all three features. However, notice that Figure 11(b) indicates that the mixture of Gaussians is capable of localizing the plastic bump and the snap before the fingers actually touch the feature (the center of the feature is at the origin of the coordinate frame). This suggests that the model is overfit to the training data. The early localization is a result of information content in the “featureless” region of the flexible material prior to contacting the feature. Looking at Figure 8, notice that there is subtle information content prior to touching the bump (otherwise, we would expect the non-bump measurements to be perfectly uniform). This subtle information does not exist in the proprioceptive information alone (Figure 6). From the perspective of contact mechanics, we hypothesize that the stiffness and surface properties of the flexible plastic have slight variations over the “featureless” region as a function of the distance of the contact point to the edge of the plastic, the position of the contact with respect to the bump, or differences in the surface properties of the plastic. Although we have found this pre-feature model to be repeatable with respect to data collected on different days, we expect that over longer time horizons, this pre-feature tactile variation is not repeatable.

A. Modeling off-feature states as a single state

One way to address the long time horizon overfitting problem is to divide state space in the training set into an *on-feature* region and an *off-feature* region that are defined manually. For all states in the off-feature region, the measurement likelihood is modeled by a single likelihood function that models data taken from the entire region. This prevents the filter from differentiating between off-feature states. Essentially, we are

lumping all off-feature state hypotheses into a single null hypothesis with a single likelihood model.

Consider the case of two tactile sensors (for example, the index and middle finger tips) with positions a and b and corresponding measurement vectors z_t^a and z_t^b such that $z_t = \{z_t^a, z_t^b\}$. Whereas in earlier sections, the measurement likelihood was conditioned on the palm position, now marginalize over the two sensor positions:

$$P(z_t^a, z_t^b | x) = \sum_{a,b} P(z_t^a | a) P(z_t^b | b) P(a, b | x). \quad (10)$$

Define functions, $A(x)$ and $B(x)$, that evaluate to the position of sensors a and b , respectively, when the palm is at x . Approximate $P(a, b | x)$ to be 1 when $a \in A(x)$ and $b \in B(x)$ and zero otherwise. Then, Equation 10 becomes:

$$P(z_t | x) = \sum_{(a,b) \in A(x) \times B(x)} P(z_t^a | a) P(z_t^b | b). \quad (11)$$

If a is in the on-feature region, then we estimate $P(z_t^a | a)$ as before using Equation 9. Otherwise, we estimate:


$$P(z_t^a | a) = \mathcal{N}(z_t^a | \hat{z}_{off}, \Sigma_{off}), \quad (12)$$

where \hat{z}_{off} and Σ_{off} are the sample mean and covariance taken over all points in the off-feature region.


Figure 12 illustrates the results of aggregating off-feature states. These results were obtained using the same plastic bump dataset that was used to produce the results in Figure 11. The solid line in Figure 12(a) shows the error for the on-feature/off-feature approach averaged over a test scan comprised of 20 swipes. The dashed line shows error for our previous approach reproduced from Figure 11. As expected, this new model does not localize the feature before the fingers come into contact with it. Figure 12(b) shows variance in the particle set averaged over the 20 test swipes. The new model has a high variance that persists until the fingers come into contact with the feature at approximately 0.5 inches prior to the bump center (the bump has approximately a one inch outer diameter). From a practical perspective, the decrease in variance when the fingers contact the feature is useful for signaling that the localization system has reached the on-feature region and probably has a good state estimate. Essentially, this on-feature/off-feature approach transforms the continuous state estimation problem into a hybrid estimation problem where the hypothesis space consists of the space of on-feature states and the binary possibility that the system is in an off-feature state. The likelihood of the binary off-feature hypothesis is the marginal likelihood of all particles in the off feature region.

B. Applications

The main motivation for using touch sensing to localize haptic features is that it can improve the robustness of manipulation tasks involving soft materials. This subsection illustrates this advantage in the context of two tasks: a plastic manipulation task and a grommet insertion task. The objective of the plastic manipulation task is to locate a bump in the flexible plastic using touch sensing and move the tip of the



(a)



(b)

Fig. 12. Performance of the on-feature/off-feature approach (solid lines) compared with the undifferentiated mixture of Gaussians approach (dotted lines). The solid lines show the performance of the featureless averaging method. The dashed lines show the Gaussian mixture performance for comparison. (a) shows average localization error. (b) shows average localization variance. The origin on the horizontal axis denotes the center of the feature.

thumb inside the recess of the bump. This is part of a larger factory assembly task. The objective of the grommet insertion task is to localize a grommet using touch sensing and insert the grommet onto a fastener. This is part of a NASA task.

First, consider the thumb-in-bump task (illustrated in Figure 13). Before localizing the bump, it was assumed that the bump position was known to within a square region two inches on a side. Given this approximate location, the robot reached to the nominal bump position and compliantly closed its fingers around the plastic using the interaction procedure described in Section II-B. Then the robot performed a swipe. During the swipe, the bump was localized using the single-Gaussian model of the proprioceptive information, the mixture of Gaussians model of the tactile information, and the separate modeling of the featureless regions (all the techniques proposed in this section). If, at any point during filtering, the variance (the trace of the covariance matrix) of the filter particles fell below a given threshold, then filtering stopped and the thumb was inserted into the bump. Otherwise, an additional swipe was performed. The insertion itself was performed using a hand-coded procedure, parametrized by the maximum likelihood bump location, that changed all finger joints to stiffness mode, moved the thumb into the bump,




Fig. 13. Illustration of the thumb insertion task. The objective is to insert the thumb into the recessed plastic bump. Frames (a) through (d) illustrate the swipe. Frames (e) and (f) illustrate the thumb insertion.




Fig. 14. Illustration of the grommet insertion task. Frames (a) through (e) illustrate the swipe. Frame (f) illustrates the insertion.

and simultaneously gripped the plastic from the other side using the fingers. The diameter of the interior of the bump was approximately 0.85 inches. The diameter of the thumb tip was approximately 0.65 inches. In order to successfully insert the thumb into the bump, the localization error could be no greater than approximately 0.35 inches. Any greater error would cause the thumb to “pop” out of the bump during the insertion.

While we do not have quantitative statistics on the success and failure rate of this insertion task, it was qualitatively very successful. We believe that most failures were associated with sensor calibration problems. As a result, we developed a short calibration procedure that was performed before running localization experiments or demonstrations. This procedure automatically relaxes all finger tendons, resets the tension affine offsets, and recalculates tendon gains after re-tensioning [27]. Out of the more than 100 attempts, the thumb insertion task succeeded approximately 95 percent of the time. Of those attempts that did not succeed, almost all failures were attributed to a failure to run the calibration procedure prior to the test or a failure in another part of the system.

We also applied our localization technique to a grommet insertion task. The objective was to localize a grommet embedded in fabric that was placed in the robot hand in an unknown position, grasp the grommet, and insert the grommet onto a fastener. (This was actually a *quarter turn* fastener that must be turned after insertion to lock the fabric in place. However, in this paper we ignore the turning part and just perform the insertion.) The grommet was placed in the robot hand in an unknown (but constrained to the region of states from which the swipe would cause the fingers to pass over the bump) configuration (Figure 14(a)). Then, the hand compliantly closed

around the fabric and performed a swipe (Figure 14(b-d)). As in the bump insertion experiment, the localization techniques proposed in this section were applied. Filtering was stopped when particle variance dropped below a threshold. At this point, the robot gripped the fabric tightly (Figure 14(e)) and moved to an insertion location (Figure 14(f)) calculated using the maximum likelihood grommet position and the fastener location that is assumed to be known (we assume that the fastener is fixtured to a large object that can be localized using other methods.) The insertion was performed under Cartesian stiffness control with a stiffness center located at the grip point. This task was much more difficult than the thumb insertion task because the required tolerances were very small. In order to successfully insert the grommet, localization error could be no greater than 0.2 inches. Since this is very close to the expected localization error for the grommet (see Figure 11(b)), even a small errors in force sensor calibration caused this task to fail. Compared with the thumb-in-bump insertion, we executed this task relatively few times (only approximately 20 times). The task was likely to succeed when executed directly after taking a training data set. However, our system was subject to sufficient drift in the sensors that we could not execute successfully on a different day without taking a new training set.

V. DISCUSSION

This paper has examined methods of using proprioceptive and tactile measurements to estimate the position of a feature (such as a button, snap, or grommet) embedded in a flexible material such as thin plastic or fabric. We have characterized the relative utility of the two types of measurements with respect to localization performance and shown that they

contain different kinds of information. We have demonstrated that using both types of information rather than just proprioceptive information results in a sizable gain in localization performance. Given the state estimation errors inherent in our training mechanism, we have found the tactile measurement model to be multimodal and proposed a mixture of Gaussians model that results in an additional improvement in localization performance. Finally, we have explored two applications of our approach that are relevant to manufacturing and space applications: a flexible plastic manipulation application (Figure 13) and a grommet insertion application (Figure 14). Although the study in this paper of localization during manipulation has been experimental, the conclusions can be expected to generalize beyond the particular hardware platform used. Although Robonaut 2 is an extremely sophisticated platform, only two hardware capabilities are needed in order to apply the conclusions from this paper: finger compliance and tactile sensing. These two features can be realized with less expensive hardware.

The idea of using Bayesian filtering to localize materials held in a robot hand is attractive. However, the approach proposed in this paper has significant limitations that should be understood. First, from an implementation perspective, we have found that the approach is very sensitive to force sensor calibration errors. Although this sensor error was not a problem for the thumb-in-bump insertion task because of the relatively large tolerances, it became more of a problem for the tight-tolerance grommet insertion task. This highlights the continuing need for more robust and accurate force and tactile sensing. Perhaps a more significant limitation is that localization based on a single haptic map does not generalize well to different interaction scenarios or to different haptic features. Given a particular haptic map, our method can be expected to localize the feature used to create the map. However, localization using the same map will not perform as well when used to localize differently shaped haptic features or even when the same haptic feature is presented in a different orientation. For example, the haptic map created for the plastic bump will not generalize to plastic bumps of different sizes. In addition, the haptic map used to localize the grommet in one orientation cannot be used to localize the same grommet when it is presented in different orientations. Note that this sensitivity to orientation only applies to non-symmetric haptic features. For example, the plastic bump training set can be used to localize the plastic bump regardless of bump orientation because the bump is symmetric (of course, localization is always relative to the orientation of the hand itself during interaction). It should also be noted that our approach inherits a degree of robustness relative to errors in the haptic map from the general robustness properties of the particle filter. Small differences in size or orientation between the haptic feature during training and testing should not be expected to affect localization accuracy significantly.

In the future, these limitations might be addressed in several ways. One approach might be to train the system on a class of features rather than on a single feature. For example, one might train the system to recognize plastic bumps of different sizes or grommets presented in different orientations. However,

since our current procedure requires 45 minutes of training time to recognize a single feature, it would clearly become infeasible to train the system on large classes of objects in a reasonable period of time. An alternative approach might be to take a compositional approach where the system is trained to recognize *parts* of a feature rather than an entire monolithic feature. For example, the system might be trained to recognize a library of curves with different curvatures and orientations. Features would be described in terms of located collections of curves. While this approach would extend the representational capabilities of this approach, the challenge would be to identify the relevant atomic shape primitives.

REFERENCES

- [1] J. Maitin-Shepard, M. Cusumano-Towner, J. Lei, and P. Abbeel, "Cloth grasp point detection based on multiple-view geometric cues with application to robotic towel folding," in *IEEE Int'l Conf. on Robotics and Automation*, 2010.
- [2] F. Lamiroux and L. Kavraki, "Planning paths for elastic objects under manipulation constraints," *International Journal of Robotics Research*, vol. 20, no. 3, pp. 188–208, 2001.
- [3] S. Burion, F. Conti, A. Petrovskaya, C. Baur, and O. Khatib, "Identifying physical properties of deformable objects by using particle filters," in *IEEE Int'l Conference on Robotics and Automation*, 2008.
- [4] M. Diftler, J. Mehling, M. Abdallah, N. Radford, L. Bridgwater, A. Sanders, S. Askew, D. Linn, J. Yamokoski, F. Permenter, B. Hargrave, R. Platt, R. Savely, and R. Ambrose, "Robonaut 2 the first humanoid robot in space," in *IEEE Int'l Conference on Robotics and Automation*, 2011.
- [5] R. Platt, F. Permenter, and J. Pfeiffer, "Inferring hand-object configuration directly from tactile data," in *Electronically published proceeding of the Mobile Manipulation Workshop, ICRA*, 2010.
- [6] P. Allen and P. Michelman, "Acquisition and interpretation of 3-d sensor data from touch," *IEEE Transactions on Robotics and Automation*, vol. 6, no. 4, pp. 397–404, 1990.
- [7] P. Allen, A. Miller, P. Oh, and B. Leibowitz, "Integration of vision, force, and tactile sensing for grasping," *Int'l Journal of Intelligent Mechatronics*, vol. 4, no. 1, pp. 129–149, 1999.
- [8] Y. Jia and M. Erdmann, "Pose and motion from contact," *International Journal of Robotics Research*, vol. 18, no. 5, pp. 466–490, 1999.
- [9] A. Okamura and M. Cutkosky, "Feature detection for haptic exploration with robotic fingers," *International Journal of Robotics Research*, vol. 20, no. 12, pp. 925–938, 2001.
- [10] S. Chhatpar and M. Branicky, "Localization in robotic assemblies with position uncertainty," in *IEEE Int'l Conf. on Intelligent Robots and Systems*, 2003, pp. 2534–2540.
- [11] S. Chhatpar, "Localization for robotic assemblies with position uncertainty," Ph.D. dissertation, Case Western Reserve University, 2006.
- [12] S. Chhatpar and M. Branicky, "Localization for robotic assemblies with position uncertainty," in *Video proceedings of the IEEE Int'l Conf. on Robotics and Automation*, 2006.
- [13] 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.
- [14] K. Gadeyne and H. Bruyninckx, "Markov techniques for object localization with force-controlled robots," in *10th Int'l Conf. on Advanced Robotics*, 2001.
- [15] C. Corcoran and R. Platt, "A measurement model for tracking hand-object state during dexterous manipulation," in *IEEE Int'l Conf. on Robotics and Automation*, 2010.
- [16] A. Remde, H. Herrich, and H. Worn, "Manipulating deformable linear objects – contact state transitions and transition conditions," in *IEEE Int'l Conf. on Intelligent Robots and Systems*, 1999.
- [17] J. Tian and Y. Jia, "Modeling deformations of general parametric shells grasped by a robot hand," *IEEE Transactions on Robotics (to appear)*, 2010.
- [18] M. Saha and P. Ito, "Motion planning for robotic manipulation of deformable linear objects," in *IEEE Int'l Conf. on Robotics and Automation*, 2006.
- [19] H. Wakamatsu, A. Tsumaya, E. Arai, and S. Hirai, "Manipulation planning for knotting/un-knotting and tightly tying of deformable linear objects," in *IEEE Int'l Conf. on Robotics and Automation*, 2005.

- [20] O. Bayazit, J. Lien, and N. Amato, "Probabilistic roadmap motion planning for deformable objects," in *IEEE Int'l Conf. on Robotics and Automation*, 2002.
- [21] M. Moll and L. Kavraki, "Path planning for deformable linear objects," *IEEE Transactions on Robotics*, vol. 22, no. 4, pp. 625–636, 2006.
- [22] H. Wakamatsu and S. Hirai, "Static modeling of linear object deformation based on differential geometry," *International Journal of Robotics Research*, vol. 23, pp. 293–311, 2004.
- [23] S. Malassiotis and M. Srinivasan, "Tracking textured deformable objects using a finite-element mesh," *IEEE Trans. on circuits and systems for video technology*, vol. 8, no. 6, 1998.
- [24] M. Teschner, B. Heidelberger, M. Müller, and M. Gross, "A versatile and robust model for geometrically complex deformable solids," in *IEEE Computer Graphics International*, 2004.
- [25] D. Morris and K. Salisbury, "Automatic preparation, calibration, and simulation of deformable objects," *Computer methods in biomechanics and biomedical engineering*, vol. 11, no. 3, 2008.
- [26] R. Platt, C. Ihrke, L. Bridgwater, M. Linn, M. Diftler, M. Abdallah, S. Askeew, and F. Permenter, "A miniature load cell suitable for mounting on the phalanges of human-sized robot fingers," in *IEEE Int'l Conference on Robotics and Automation*, 2011.
- [27] M. Abdallah, R. Platt, C. Wampler, and B. Hargrave, "Applied joint-space torque and stiffness control of tendon-driven fingers," in *IEEE Int'l Conf. on Humanoid Robots*, 2010.
- [28] 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.
- [29] C. Bishop, *Pattern recognition and machine learning*. Springer, 2006.
- [30] J. Ko and D. Fox, "GP-BayesFilters: Bayesian filtering using gaussian process prediction and observation models," *Autonomous Robots*, vol. 27, no. 1, 2009.



Robert Platt Jr. Robert Platt is a Research Scientist in the Computer Science and Artificial Intelligence Laboratory at the Massachusetts Institute of Technology. Between 2005 and 2009, he was a researcher at NASA Johnson Space Center working on the Robonaut 2 project. He received a Ph.D. in Computer Science in 2006 from the University of Massachusetts, Amherst. He is interested in robot planning and control in noisy and uncertain environments with a particular focus on manipulation and assembly.



Frank Permenter Frank Permenter received the BSEE (Summa Cum Laude) from the University of Houston and the MSEE from Stanford. From 2004–2007, he was a member of the Guidance, Navigation and Control team for the International Space Station, specializing in attitude control flight software. Since 2007, he has been a design engineer on NASA's Robonaut 2, the first humanoid robot to fly in space.



Joseph Pfeiffer Joseph Pfeiffer is a Graduate Student in the Network Learning and Discovery Laboratory in the Purdue Computer Science Department, working under the Purdue University Frederick N. Andrews fellowship. He is also a cooperative education student at NASA Johnson Space Center in the Robotic Systems Technology Branch. His research interests are in the broad area of statistical machine learning.