State Estimation and Tracking of Deforming Planar Elastic Rods

Andy Borum[1], Dennis Matthews[2], and Timothy Bretl[1]

***Abstract*— In this paper, we address the problem of estimating the shape of a planar elastic rod (e.g., a thin flexible strip of metal) using images of the rod. This is done by treating configurations of the elastic rod as solutions of a geometric optimal control problem. The necessary conditions for optimality provide coordinates over which to perform inference, and the sufficient conditions provide the gradient of the shape of the rod with respect to these coordinates. This optimal control formulation allows for configurations of the rod to be represented as points in a finite-dimensional space without having to discretize the shape of the rod. We consider the estimation problem with and without fiducial markers attached to the rod. Results from both simulations and hardware experiments demonstrate the ability of our approach to track the shape of a deforming elastic rod.**

I. INTRODUCTION

Consider a thin, flexible elastic rod of fixed length that is confined to deform in a plane. One end of the rod is held fixed while the other is held by a robotic gripper that can translate and rotate in the plane of the rod. When a camera is positioned so that its line of sight is orthogonal to the plane of the rod, we show how to estimate the shape of the elastic rod as it is being deformed by the robot using images captured by the camera. The magenta curve in Figure 1 was obtained using our approximation method.

To demonstrate why this problem seems difficult, consider an elastic rod that begins in the stable equilibrium configuration shown in Figure 2(a) and is then deformed by moving the robotic gripper downward. At a certain point during this deformation, shown in Figure 2(b), the rod experiences an instability. At this point, a pitchfork bifurcation occurs, and two stable equilibrium shapes emerge, while the original stable configuration becomes unstable. The three resulting equilibrium shapes are shown in Figure 2(c)-(e). Knowing only the trajectory of the robotic gripper does not provide enough information to accurately track the rod's configuration. By using a sequence of camera images, our approach is able to track the shape of the rod no matter which branch of the bifurcation the rod follows.

The formulation of our method relies on a description of elastic rods in equilibrium as local solutions to a geometric optimal control problem [1], [2]. Pontryagin's Maximum Principle [3] provides necessary conditions for optimality, and the initial values of the costates that aries from these necessary conditions provide coordinates that describe every equilibrium configuration the rod can take. These coordinates can be interpreted physically as the forces and the torque acting on one end of the rod. Using these coordinates, it can be shown that the set of all configurations is a smooth three-dimensional manifold. This manifold is exactly the space over which to search for configurations that match the observed shape. A sufficient condition for optimality provides the gradient of the configuration with respect to these coordinates. Thus, by using this geometric optimal control formulation, the state estimation problem can be cast as a nonlinear optimization problem.

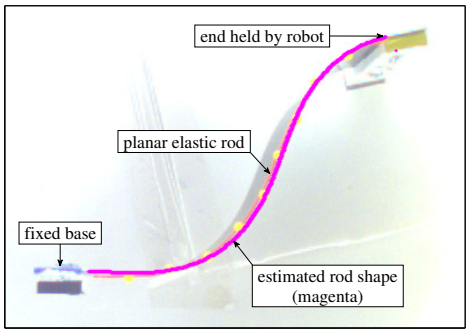

Fig. 1: An elastic rod held by a robotic gripper and an approximation of the rod's shape.

This optimal control formulation offers advantages over previously proposed methods of tracking deformable objects. Many previous approaches use finite-dimensional approximations of contours, such as cubic splines [4], so that they may perform estimation in a finite-dimensional space rather than the infinite-dimensional space of continuous curves. In [5], the finite element method is used to track a deforming surface, while the boundary element method is used in [6]. In [7], a discrete elastic rod model, developed in [8], is used to estimate configurations of elastic rods. These approaches rely on a discretization of the deformable object, and while this discretization reduces the problem to a finite-dimensional space, the dimension of the space depends on the resolution of the discretization. Other approaches work with infinite-dimensional curves directly, by representing contours as the locus of zeros of a function, but performing estimation in these spaces produces many challenges [9].

The optimal control formulation used in this paper produces a three-dimensional description of configurations of an elastic rod. However, this finite-dimensional representation is based upon a continuous description of the rod. Therefore,

*This work was supported by the NSF under Grant No. IIS-1320519. The work of A. Borum was supported by the NSF-GRFP under Grant No. DGE-1144245.

[1]Andy Borum and Timothy Bretl are with the Department of Aerospace Engineering and [2]Dennis Matthews is with the Department of Electrical and Computer Engineering, University of Illinois at Urbana-Champaign, Urbana, IL, 61801, USA, {borum2,matthws1,tbretl}@illinois.edu

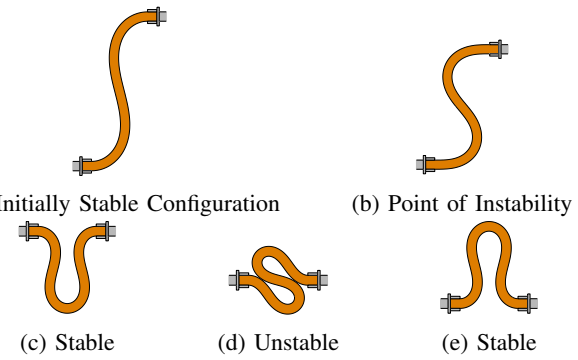

(a) Initially Stable Configuration (b) Point of Instability

(c) Stable (d) Unstable (e) Stable

Fig. 2: An elastic rod passing through a bifurcation.

the dimension of the space in which we perform estimation remains constant when the accuracy of the numerical method used to solve for the shape of the rod increases. Thus we are able to work with a continuous description of the rod while using tools from finite-dimensional estimation. Also, the estimate provided by our method gives physical information about the forces and torques acting on the rod, which is not provided by the curve fitting methods described above.

One application of our estimation and tracking procedure is a vision-based feedback controller for manipulation of elastic rods. The coordinates derived from the necessary conditions for optimality can be used to plan a path of the robotic gripper that takes the elastic rod from an initial configuration to some goal configuration while remaining in static equilibrium [1]. While executing this path, uncertainties in the stiffness or length of the rod may lead to instabilities. If an instability were to occur, our tracking and estimation procedure could determine the new configuration of the rod, and then a path from the current configuration to the goal configuration could be planned. While the previously proposed approaches to tracking can be used to estimate the shape of the rod, they do not provide the coordinates needed to perform manipulation planning.

A variety of industrial processes that involve robotic handling of deformable objects could benefit from this work. Examples include fixturing of sheet metal [10], cutting and layup of composites [11], installation of a wire harness [12], assembly of flexible circuit boards [13], and cable routing [14]. Medical procedures and equipment that involve manipulation of flexible objects include automated knot tying and surgical suturing [15], [16], and manipulation of flexible needles [17]. Other possible applications include elastic rod based image segmentation [18] and reconstruction [19].

The remainder of this paper proceeds as follows: Section II reviews the formulation of the planar elastic rod as an optimal control problem. In Section III, we formulate the state estimation problem and show how to approximate configurations of an observed rod. Section IV presents state estimation and tracking results from simulation experiments, while Section V presents similar results from hardware experiments. In Section VI, we close by discussing the limitations of this approach to state estimation.

II. NECESSARY AND SUFFICIENT CONDITIONS FOR STABILITY OF PLANAR ELASTIC RODS

In this section, we introduce a model for planar inextensible elastic rods in equilibrium based on a geometric optimal control problem. We then review the necessary and sufficient conditions derived in [2] for a configuration of the rod to be in static equilibrium.

A configuration of a planar elastic rod of unit length is described by the two functions $q : [0,1] \to SE(2)$ and $u : [0,1] \to \mathbb{R}$. $q(t)$ describes the position and orientation of points along the rod in $SE(2)$, and $u(t)$ is the rate of change of the rotation component of $q(t)$, i.e., $u(t)$ is the curvature of the rod. Since the rod is inextensible, the relationship between $q(t)$ and $u(t)$ can be expressed as

$$\dot{q} = q(X_1 + uX_3)$$

where

$$X_1 = \begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \quad X_2 = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{bmatrix} \quad X_3 = \begin{bmatrix} 0 & -1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}$$

is a basis for $\mathfrak{se}(2)$, the Lie algebra of $SE(2)$. Let $\{P_1, P_2, P_3\}$ be the corresponding dual basis of $\mathfrak{se}(2)^*$. We assume that the base of the rod is held fixed at the origin, so that $q(0) = e$, where e is the identity element in $SE(2)$. The other end of the rod is held by a robotic gripper at some arbitrary $q(1) \in SE(2)$. Denote the space of all possible $q(1)$ by $\mathcal{B} = SE(2)$.

A configuration of the rod is a stable equilibrium configuration if it is a local minimizer of elastic potential energy. Assuming the rod behaves in a linear-elastic manner, the elastic energy is proportional to the integral of the curvature squared. Therefore, for a fixed $b \in \mathcal{B}$, the rod will be in static equilibrium if it is a local optimum of

$$\begin{aligned} \underset{q,u}{\text{minimize}} \quad & \frac{1}{2}\int_0^1 u^2\, dt \\ \text{subject to} \quad & \dot{q} = q(X_1 + uX_3) \\ & q(0) = e \quad q(1) = b. \end{aligned} \tag{1}$$

We now state the necessary conditions derived in [2] for a configuration (q,u) to be a local optimum of (1).

Theorem 1: (Necessary Conditions) A configuration (q,u) is a normal extremal of (1) if and only if there exists $\mu : [0,1] \to \mathfrak{se}^*(2)$ that satisfies

$$\dot{\mu}_1 = \mu_2 u \qquad \dot{\mu}_2 = -\mu_1 u \qquad \dot{\mu}_3 = -\mu_2 \tag{2}$$

$$\dot{q} = q(X_1 + uX_3) \qquad u = \mu_3 \tag{3}$$

with $q(0) = e$ and $\mu(0) = \sum_{i=1}^3 a_i P_i$ for some $a \in \mathcal{A}$, where $\mathcal{A} = \{a \in \mathbb{R}^3 : (a_2, a_3) \neq (0,0)\}$.

Proof: See the proof of Theorem 5 in [2]. ■

Theorem 6 in [2] states that the set of all (q,u) which satisfy (2)-(3) for some $a \in \mathcal{A}$ is a smooth 3-manifold parameterized by a single global coordinate chart, and coordinates for this global chart are $a \in \mathcal{A}$. Thus if we want to estimate the configuration of some observed rod, $\mathcal{A}$ is exactly the space in which we should search. These coordinates have

a physical interpretation, as described in [2]. The vector $a = (a_1, a_2, a_3)$ can be viewed as the vector of reactions at the fixed base of the rod, i.e., a_1 and a_2 are the forces acting on the rod at $t = 0$ and a_3 is the torque acting on the rod at $t = 0$. The following theorem provides sufficient conditions for (q, u) to be a local solution of (1).

Theorem 2: (Sufficient Conditions) Let (q, u) and μ be the functions obtained by solving (2)-(3) for a particular choice of $a \in \mathcal{A}$. Define the matrices

$$\mathbf{F} = \begin{bmatrix} 0 & \mu_3 & \mu_2 \\ -\mu_3 & 0 & -\mu_1 \\ 0 & -1 & 0 \end{bmatrix} \quad \mathbf{G} = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad \mathbf{H} = \begin{bmatrix} 0 & \mu_3 & 0 \\ -\mu_3 & 0 & 1 \\ 0 & 0 & 0 \end{bmatrix}$$

Solve the (linear, time-varying) matrix differential equations

$$\dot{\mathbf{M}} = \mathbf{FM} \qquad \dot{\mathbf{J}} = \mathbf{GM} + \mathbf{HJ} \tag{4}$$

with initial conditions $\mathbf{M}(0) = I$ and $\mathbf{J}(0) = 0$. Then (q, u) is a local optimum of equation (1) for $q(1) = b$ if and only if $\det(\mathbf{J}(t)) \neq 0$ for all $t \in (0, 1]$.

Proof: See the proof of Theorem 7 in [2]. ■

While the matrix $\mathbf{J}(t)$ can be used to determine the stability of a configuration of the rod, it also contains information about the relationship between small changes in a and small changes in q. In particular, the j^{th} column of the matrix $\mathbf{J}(t)$ is the coordinate representation of $\frac{\partial q(t)}{\partial a_j}$ with respect to the basis $\{X_1, X_2, X_3\}$. As we will see in the next section, these gradients of q with respect to a will be useful when trying to determine the state of the rod from observations.

III. FORMULATION OF THE STATE ESTIMATION PROBLEM

We saw in the previous section that any configuration of a planar elastic rod corresponds to a point $a \in \mathcal{A}$. We will now formulate the problem of determining the value of $a \in \mathcal{A}$ that corresponds to an observed configuration of an elastic rod.

Consider a planar elastic rod which has a known length of L (not necessarily unit length). The base of the rod is fixed, while the other end is held by a robotic gripper. Assume that the robot is able to provide the position and orientation of the end of the rod relative to the frame situated at the base of the rod. Denote this position by $(\bar{X}, \bar{Y})$ and this orientation by $\bar{\theta}$, and let $\bar{x} = \bar{X}/L$ and $\bar{y} = \bar{Y}/L$. Also assume that we are observing the rod with a camera whose line of sight is orthogonal to the plane in which the rod is deforming and is placed at a fixed distance from this plane. Let B be a linear transformation that takes points in the image plane of the camera to points in the frame situated at the base of the rod and then scales these points by L. After applying this transformation to the rod observed by the camera, the resulting rod will have unit length.

This rod is in some configuration (q, u) which, from the previous section, we know corresponds to some $a \in \mathcal{A}$. Thus determining the configuration (q, u) of the rod (which is an infinite-dimensional curve) is equivalent to determining the finite-dimensional value of a, and we denote this dependence by $(q(t, a), u(t, a))$. We will consider the problem of estimating the value of a in two cases of observations.

A. Estimation with fiducial markers

In the first case, we assume that fiducial markers are placed at known locations along the rod. Assume that n markers are located at distinct positions along the rod, given by $0 < L_1 < L_2 < ... < L_n < L$, and let $l_j = L_j/L$. Let $(\tilde{X}_j, \tilde{Y}_j)$ be the location of the j^{th} marker observed in the image plane, and let $(\tilde{x}_j, \tilde{y}_j) = B((\tilde{X}_j, \tilde{Y}_j))$, i.e. $(\tilde{x}_j, \tilde{y}_j)$ is the position of the j^{th} marker in the plane of the rod scaled by the length of the rod. For any $a \in \mathcal{A}$, let $x(t, a)$ and $y(t, a)$ be the position coordinates and let $\theta(t, a)$ be the orientation of point t along the configuration $(q(t, a), u(t, a))$. Define $F(a)$ to be

$$F(a) = \begin{bmatrix} x(l_1, a) - \tilde{x}_1 \\ y(l_1, a) - \tilde{y}_1 \\ \vdots \\ x(l_n, a) - \tilde{x}_n \\ y(l_n, a) - \tilde{y}_n \\ x(1, a) - \bar{x} \\ y(1, a) - \bar{y} \\ \theta(1, a) - \bar{\theta} \end{bmatrix} \tag{5}$$

If the configuration corresponding to a matches the observed configuration exactly, then we will have $F(a) = 0$. We now have a system of $2n + 3$ nonlinear equations. If we let $\hat{\mathbf{J}}(a)$ denote the Jacobian of $F(a)$, we can use the Newton-Raphson method to recursively approximate a by

$$\Delta a^{(i)} = \left(\hat{\mathbf{J}}\left(a^{(i)}\right)^T \hat{\mathbf{J}}\left(a^{(i)}\right) \right)^{-1} \hat{\mathbf{J}}\left(a^{(i)}\right)^T F\left(a^{(i)}\right) \tag{6}$$

$$a^{(i+1)} = a^{(i)} - \Delta a^{(i)} \tag{7}$$

All that remains is to find $\hat{\mathbf{J}}(a)$. From the discussion after Theorem 2, we see that the k^{th} column of $\mathbf{J}(t)$ is

$$\begin{bmatrix} \frac{\partial x(t,a)}{\partial a_k} & \frac{\partial y(t,a)}{\partial a_k} & \frac{\partial \theta(t,a)}{\partial a_k} \end{bmatrix}^T \tag{8}$$

Denoting the dependence of $\mathbf{J}(t)$ on a explicitly and denoting the k^{th} row of $\mathbf{J}(t, a)$ by $\mathbf{J}_k(t, a)$, we have

$$\hat{\mathbf{J}}(a) = \begin{bmatrix} \mathbf{J}_1(l_1, a) \\ \mathbf{J}_2(l_1, a) \\ \vdots \\ \mathbf{J}_1(l_n, a) \\ \mathbf{J}_2(l_n, a) \\ \mathbf{J}_1(1, a) \\ \mathbf{J}_2(1, a) \\ \mathbf{J}_3(1, a) \end{bmatrix} \tag{9}$$

We can now use (6)-(7) to recursively approximate a.

B. Estimation without fiducial markers

In the previous case, we assumed that fiducial markers can be placed along the rod at known lengths and observed by the camera. In practice, however, this may not be possible. Thus we will need to estimate a using a continuous image of the rod.

We begin by assuming that the image of the rod obtained from the camera can be segmented so that the rod is an

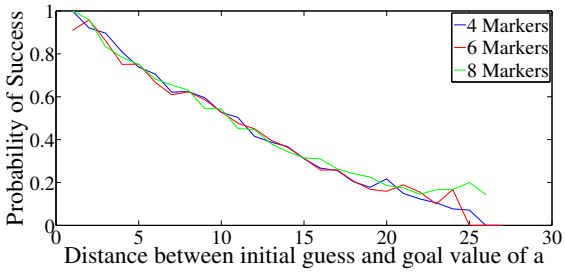

Fig. 3: Probability of successfully estimating $a \in \mathcal{A}$ using fiducial markers.

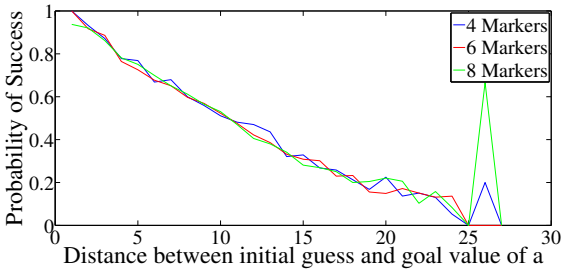

Fig. 4: Probability of successfully estimating $a \in \mathcal{A}$ without using fiducial markers.

isolated object in the image. The observed rod will then consist of a collection of points in the image plane of the camera. Denote this collection by P and let p be the result of applying the transformation B to the points in P. We now pick a positive integer n, which will be the number of points we will try to match along the observed rod. Define $l_j = j/(n+1)$ for $j = 1, 2, ..., n$, and let

$$F(a) = \begin{bmatrix} x(l_1, a) - \tilde{x}_1(a) \\ y(l_1, a) - \tilde{y}_1(a) \\ \vdots \\ x(l_n, a) - \tilde{x}_n(a) \\ y(l_n, a) - \tilde{y}_n(a) \\ x(1, a) - \bar{x} \\ y(1, a) - \bar{y} \\ \theta(1, a) - \bar{\theta} \end{bmatrix} \tag{10}$$

where $\tilde{x}_j(a)$ and $\tilde{y}_j(a)$ are now defined as follows:

$$(\tilde{x}_j, \tilde{y}_j)(a) = \arg\min_{(\tilde{x}, \tilde{y}) \in p} \left\| \begin{pmatrix} x(l_j, a) \\ y(l_j, a) \end{pmatrix} - \begin{pmatrix} \tilde{x} \\ \tilde{y} \end{pmatrix} \right\| \tag{11}$$

The point $(\tilde{x}_j(a), \tilde{y}_j(a))$ is the closest point on the observed rod to $(x(l_j, a), y(l_j, a))$. We are still using the position and orientation of the end of the rod provided by the robot.

When using fiducial markers, the points $(\tilde{x}_j, \tilde{y}_j)$ are constant for $j = 1, 2, ..., n$, so the Jacobian of $F(a)$ is independent of $(\tilde{x}_j, \tilde{y}_j)$, as seen in (9). Without fiducial markers, the points $(\tilde{x}_j(a), \tilde{y}_j(a))$ depend on the current guess of a according to (11). However, since p is a finite collection of points, $(\tilde{x}_j(a), \tilde{y}_j(a))$ changes discretely as a varies. If we assume that small changes in a do not cause $(\tilde{x}_j(a), \tilde{y}_j(a))$ to change, then the Jacobian matrix of (10) is independent of $(\tilde{x}_j(a), \tilde{y}_j(a))$. Therefore, the Jacobain matrix of (10) is given by (9), and the Jacobian is the same with or without fiducial markers.

IV. SIMULATION RESULTS

A. *State estimation from a random initial guess*

The method of approximating a described in the previous section was implemented using data from simulated elastic rods. 60,000 values of a were randomly selected from the set $\bar{\mathcal{A}} = \{a \in \mathcal{A} : |a_1| \leq 10, |a_2| \leq 10, |a_3| \leq 3\}$. These values were used to generate simulated rods by integrating (2)-(3). The values of $\tilde{x}_1, ..., \tilde{x}_n, \tilde{y}_1, ..., \tilde{y}_n, \bar{x}, \bar{y}$, and $\bar{\theta}$ from these rods were used as input data in the approximation algorithm. Another 60,000 values of a were randomly selected from $\bar{\mathcal{A}}$ and used as initial guesses in the approximation algorithm. Of the 60,000 pairs of goal a values and initial guesses, 30,000 were used for estimation with fiducial markers, and 30,000 were used for estimation without fiducial markers. In each case, 10,000 were used with 4 markers along the rod, 10,000 with 6 markers, and 10,000 with 8 markers.

To improve the performance of the iteration procedure given in (6)-(7), a backtracking line search method was used to adjust the magnitude of the step $\Delta a^{(i)}$ in (7) [20]. In this backtracking line search procedure, the parameter α, which governs the acceptable decrease in the error during each step, was chosen to be 0.1, and the parameter β, which governs the decrease in step size when the acceptable decrease in the error is not met, was chosen to be 0.5. For each pair of start and goal a values, we iterated (6)-(7) (along with the backtracking procedure) until either 50 iterations of (6)-(7) were performed, the norm of (5) (or (10) in the case without fiducial markers) was less than 10^{-6}, or the step size in the backtracking line search was less than 10^{-10}.

Figures 3 and 4 show the probability of success of these simulations as a function of the distance between the goal value of a and the initial guess of a with and without fiducial markers, respectively. A simulation was deemed successful if the Euclidean distance between the estimated a value and the goal a value was less than some specified tolerance. For 48.01% of the 60,000 simulations performed, the Euclidean distance between the estimated and goal a values was less than 1. For the other 51.99%, the distance was greater than 10. Therefore, in Figures 3 and 4, the tolerance for success was chosen to be 1.

The results exhibit a clear monotonic trend, with the probability of success decreasing as the initial guess moves further away from the goal value (the spike near 25 in Figure 4 is due to a small sample of points at this particular distance between starting and goal a values). Both cases exhibited nearly identical behavior, and there was no clear distinction between having 4, 6, or 8 markers in either case. Thus, for the elastic rods considered, having fiducial markers attached to the rod did not provide any advantage in terms of the probability of success, nor did having more than 4 markers along the rod.

Having markers attached to the rod did, however, decrease

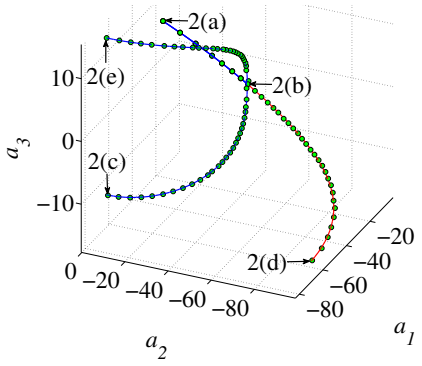

Fig. 5: Bifurcation in $\mathcal{A}$ corresponding to Figure 2.

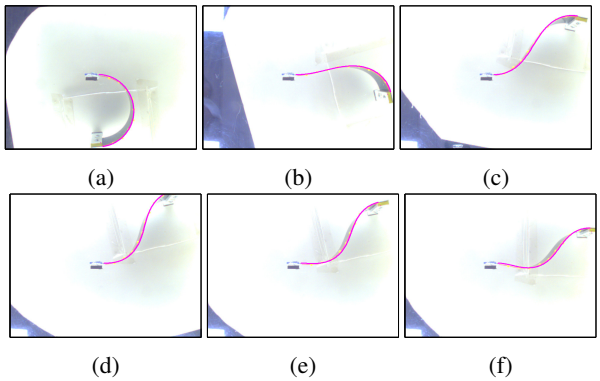

Fig. 6: Tracking the shape of a deformed elastic rod. Magenta curves are produced using our state estimation procedure.

the computation time of the estimation procedure. The average computation times for successful simulations was 0.1545 sec and 0.3470 sec for simulations with and without markers, respectively. When the simulation was not successful, the average computation times were 2.2320 sec and 1.9972 sec for simulations with and without markers, respectively. The computations were performed using Matlab on a MacBook Pro laptop with a 2.53 GHz Intel Core 2 Duo processor.

B. Tracking through a bifurcation

We next used our algorithm to track the bifurcation shown in Figure 2. To do this, we assumed no fiducial markers were attached to the rod and chose $n = 4$. The solid lines in Figure 5 represent the trajectories in $\mathcal{A}$ of the three deformations in Figure 2, with blue lines corresponding to stable equilibrium shapes and red lines corresponding to unstable shapes. The labels in Figure 5 show the a values that correspond to each configuration in Figure 2. At the bifurcation point (labeled 2(b)), the previously stable configuration becomes unstable and two new stable configurations are created. The green circles in Figure 5 represent the estimated values of a along these three trajectories.

To obtain these estimates, the simulation was initialized at the configuration in Figure 2(a) with the exact value of a given as the initial guess. Then the observed rod was updated to the next configuration (corresponding to moving the gripper downward) and the estimate was allowed to converge. The estimate of each configuration was used as the initial guess for the next configuration. This was done for the three branches shown in Figure 5. When passing through the bifurcation point, the tracking algorithm had no trouble estimating the correct branch of the bifurcation to follow. If the only information used was the end position and orientation of the rod, the tracking procedure followed the unstable (red) branch no matter which branch the actual rod followed. In a physical experiment, the rod would never follow this unstable branch.

V. EXPERIMENTAL RESULTS

Hardware experiments were conducted using an industrial Adept robot to deform an elastic rod. One end of the rod was held by the robot, and the other end was fixed to a stationary table. Markers were placed along the rod and tracked using a camera. The intrinsic parameters of the camera were calculated using OpenCV's cvCalibrate2 function for a set of 70 images of a checkerboard pattern in different poses. Using the location of the robot's end effector in various positions, the extrinsic camera parameters were calculated using the function cvFindExtrinsicCameraParams2. The experimental setup is shown in Figure 1.

A. Tracking through a continuous motion

We first tracked the rod through a continuous motion in $\mathcal{A}$, i.e., a path that contained no unstable configurations. Figure 6 shows six configurations sampled along the trajectory. The magenta curve in each image is the approximation of the shape obtained using our tracking procedure. The path in $\mathcal{A}$ corresponding to the trajectory in Figure 6 is shown in Figure 7 by the solid blue curve. The green circles in Figure 7 represent the estimated shapes. Both Figures 6 and 7 show that we were able to accurately track the shape of the rod as it was deformed by the robot.

B. Tracking through an instability

We next attempted to track the rod along a trajectory that contained an unstable configuration. Figures 8(a) and (b) show the rod in the instants shortly before and after the instability occurred. We found that using the value of a corresponding to the estimate in Figure 8(a) as the initial guess for the rod shown in Figure 8(b) resulted in the estimate converging to a local minimum of the error and not to the correct value of a. Two such local minima are shown in Figures 8(c) and (d). When this occurred, we provided a new randomly selected initial guess from $\bar{\mathcal{A}}$ and restarted the estimation procedure. We repeated this process 800 times for the instability shown in Figure 8, and an average of 3.6 new initial guesses needed to be supplied before the procedure converged to the correct value of a.

As can be seen in Figures 8(c) and (d), the local minima do not pass through the fiducial markers on the rod. Therefore, one can determine whether the algorithm has converged to the correct value of a or a local minima of $\|F(a)\|$ simply by comparing the value of $\|F(a)\|$ to some specified tolerance. Thus, the process of determining when to randomly pick a new guess of a can be done automatically.

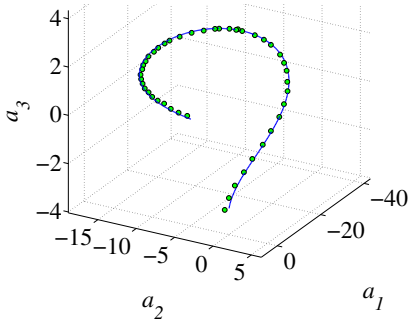

Fig. 7: Tracking a continuous path in $\mathcal{A}$. Green circles are estimated points along the trajectory.

VI. CONCLUSIONS

We have presented a procedure for shape estimation and tracking of an elastic rod being deformed by a robot based on an optimal control formulation of the rod. The advantages of our approach over other methods are: (1) we work with a continuous description of the rod and do not rely on a discretization, (2) even though we model the rod as a continuous curve, we are able to represent configurations of the rod as points on a three-dimensional smooth manifold, and (3) the dimension of the space in which we are performing estimation does not grow as the accuracy of the numerical method used to solve for the shape of the rod increases. Results from simulation and hardware experiments show that our estimation procedure is able to accurately estimate and track the shape of a deforming elastic rod.

One particular area in which our procedure could be improved is the process of picking initial guesses of a after an instability. Selecting new guesses of a at random, as was done in Section V(B), is not an efficient strategy, since it was reported in Section IV(A) that unsuccessful estimations require roughly 2 seconds of computational time. One way to decrease the computation time would be to place a force/torque sensor at the base of the rod. Since the value of a can be interpreted as the vector of forces and torques acting on the rod, the measurements from the sensor could be used as an initial guess of a. A further extension would be to use the force and torque measurements along with visual tracking as an estimator for a.

Other areas for future work include extending this method to three dimensional rods. An optimal control formulation for three dimensional elastic rods is given in [1]. While fiducial markers could be tracked using a motion capture system, the problem of tracking from a single camera image seems difficult due to the possibility of occlusions. Another possible extension is estimation of material properties. Just as the the gradient of the rod's configuration (q, u) with respect to a is given by $\mathbf{J}(t)$, one can also compute the gradient of (q, u) with respect to the material parameters of the elastic rod. Observations could then be used to estimate the rod's bending and torsional stiffnesses.

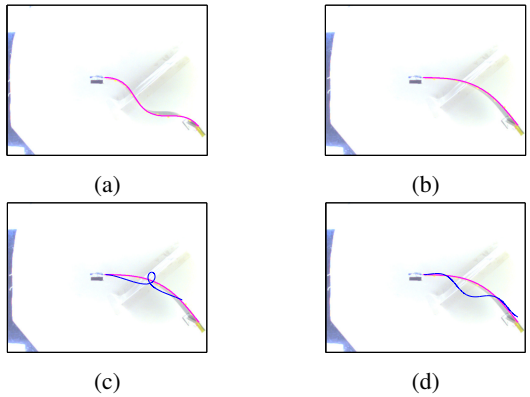

Fig. 8: Tracking the shape of an elastic rod before (a) and after (b) an instability. The estimation procedure sometimes converges to a local minima of the error, as shown by the blue curves in (c) and (d).

REFERENCES

[1] T. Bretl and Z. McCarthy, Quasi-static manipulation of a Kirchhoff elastic rod based on a geometric analysis of equilibrium Configurations, *IJRR*, 33(1):48-68, 2014.

[2] D. Matthews and T. Bretl, Experiments in quasi-static manipulation of a planar elastic rod, In *IEEE/RSJ IROS*, pp. 5420-5427, 2012.

[3] L. S. Pontryagin, V. G. Boltyanskii, R. V. Gamkrelidze, and E. F. Mishchenko, The mathematical theory of optimal processes, New York: Wiley, 1962.

[4] G. Medioni and Y. Yasumoto, Corner detection and curve representation using cubic *B*-splines, *Comput. Vision Graph.*, 39(3):267-278, 1987.

[5] S. Wuhrer, J. Lang, and C. Shu, Tracking complete deformable objects with finite elements, In *IEEE 3DIMPVT*, pp. 1-8, 2012.

[6] M. Greminger and B. Nelson, A deformable object tracking algorithm based on the boundary element method that is robust to occlusions and spurious edges, *Int. J. Comput. Vision*, 78:29-45, 2008.

[7] S. Javdani, S. Tandon, J. Tang, J. F. O'Brien, and P. Abbell, Modeling and perception of deformable one-dimensional objects, In *IEEE ICRA*, pp. 1-8, 2011.

[8] M. Bergou, M. Wardetzky, S. Robinson, B. Audoly and E. Grinspun, Discrete elastic rods, *ACM Trans. Graph.*, 27(3): 112, 2008.

[9] J. Jackson, A. Yezzi, and S. Soatto, Tracking deformable moving objects under severe occlusions, In *IEEE CDC*, Volume 3, pp. 2990-2995, 2004.

[10] W. Nguyen and J. Mills, Multi-robot control for flexible fixtureless assembly of flexible sheet metal auto body parts, In *IEEE ICRA*, Volume 3, pp. 2340-2345, 1996.

[11] A. Angerer, C. Ehinger, A. Hoffman, W. Reif, G. Reinhart, and G. Strasser, Automated cutting and handling of carbon fiber fabrics in aerospace industries, In *IEEE CASE*, pp. 861-866, 2010.

[12] X. Jiang, K. M. Koo, K. Kikuchi, A. Konno, and M. Uchiyama, Robotized assembly of a wire harness in a car production line, *Adv. Robotics*, 25(3-4):473-489, 2011.

[13] Y. Asano, H. Wakamatsu, E. Morinaga, E. Arai, and S. Hirai, Deformation path planning for manipulation of flexible circuit boards, In *IEEE/RSJ IROS*, pp. 5386-5391, 2010.

[14] H. Inoue and H. Inaba, Hand-eye coordination in rope handling, In *ISRR*, pp. 163-174, 1985.

[15] M. Saha and P. Isto, Manipulation planning for deformable linear objects, *IEEE Trans. Robot.*, 23(6):1141-1150, 2007.

[16] H. Wakamatsu, E. Arai, and S. Hirai, Knotting/unknotting manipulation of deformable linear objects, *IJRR*, 24(4):371-395, 2006.

[17] N. Chentanez, R. Alterovitz, D. Ritchie, J. Cho, K. Hauser, K. Goldberg, J. R. Shewchuk, and J. F. O'Brian, Interactive simulation of surgical needle insertion and steering, *ACM Trans. Graph.*, 28(88):1-10, 2009.

[18] M. Krueger, P. Delmas, and G. Gimel'farb, Efficient image segmentation using weighted pseudo-elastica, *Pattern Recogn. Lett.*, 34(8):833-845, 2013.

[19] T. Chan, S. Kang, and J. Shen, Euler's elastica and curvature-based image inpainting, *SIAM J. Appl. Math.*, 63(2):564-592, 2002.

[20] S. Boyd and L. Vandenberghe, Convex Optimization, Cambridge University Press, 2004.