Action Selection for Robotic Manipulation of Deformable Planar Objects

Saúl Cuén-Rochín, Juan Andrade-Cetto and Carme Torras

***Abstract*— This paper deals with the manipulation of planar deformable objects for typical service robot applications. Specifically, we present a system that straightens pieces of cloth from any arbitrary initial wrinkle condition using a robotic manipulator. The cloth is modeled with a Finite Element Method, and its state is estimated with a physical-based implicit integration scheme that computes particle velocities as a function of internal and external forces acting on the object. The state of the object is tracked with a stochastic observer, in which measurements come from a stereo vision system. Manipulation actions are chosen maximizing an a-optimal information measure.**

To our knowledge, this is the first time that a stochastic state estimator has been derived for an implicit integration model of a deformable planar object, bridging the gap between computer simulation and vision-based tracking of the state of deformable planar objects for manipulation.

I. INTRODUCTION

Service robots, and in particular those aimed at helping humans in daily tasks, are gaining a lot of attention. The variability of tasks and environments in which they are to operate pose new research problems not tackled within industrial robotics. The non-repetitive manipulation of deformable objects is one such problem, since these objects are plentiful in homes and assistive environments.

While a lot of work has been devoted to grasping, motion planning and manipulation of rigid objects [1], [2], similar research for deformable objects is just starting. The European Project PACO-PLUS [3] is addressing the grasping and manipulation of both types of objects within a kitchen environment. The long-term goal is to plan and execute manipulation tasks with the fingered hands of the ARMAR robot, but as a first step this paper deals with action selection for cloth straightening with just one finger.

The existing work on manipulation of deformable objects deals mainly with linear objects [4], [5], [6], such as ropes threads and wires. These works usually rely on a Finite Element Method (FEM) to model the objects, and make use of knot topology to plan motions. Modelling deformable planar objects –those of interest to us– in the same way may be computationally costly, and the alternative of using a Boundary Element Method (BEM) has been proposed [7], where BEM differs from FEM in that only the contour of the object needs to be meshed. However, BEM-based simulation does not provide enough detail on cloth state for our purposes, so we don't adopt this method in the present work. Other works focus on grasping skills for cloth manipulation [8], and on iterative learning of the force required to lift a deformable object [9]. A compilation of systems for the industrial manipulation of deformable objects is discussed in [10], going from sewing systems to fish manipulation processes.

For action planning, a physical-based simulation that accurately predicts the outcome of actions is crucially needed. In the Computer Graphics field, there are two approaches to cloth simulation that use FEM to describe the particles positions and velocities as a mesh of primitives such as triangles and rectangles. One is the implicit integration scheme [11], which at the expense of a high computational cost, remains stable despite taking long time steps. The other is the explicit integration scheme [12], with lower computational burden, but restrained to take short time steps to assure stable solutions. The reader is referred to [13] for a detailed discussion of cloth simulation approaches, and to [14] for a wider compilation of physical models of deformable objects.

In the current work we have adopted the implicit integration scheme, given its better accuracy in estimating the state of the object over extended periods of time. In this approach, a large sparse linear system is solved through a preconditioned conjugate gradient (CG) iterative method. The preconditions of the CG method permit imposing external constraints on the velocities of some particles, which comes handy when we need to restrain the motion of the particles fixed by the finger. State estimation techniques are used to track the state of the cloth. In particular, an Extended Kalman filter on the implicit integration scheme has been implemented. The selection of the best next action to straighten the cloth is then tackled using tools from information theory. This approach to action selection has previously been pursued by our group in the context of active vision for Simultaneous Localization and Mapping (SLAM) [15], and wire-based robot pose tracking [16]. Here, we have adapted it to cope with the uncertainty inherent to the manipulation of deformable objects.

The paper is structured as follows. Section II details the model used to predict the deformation of cloth under the presence of three types of forces: i) internal forces such as stretch, shear and bend, ii) the forces exerted by our manipulation strategy, and iii) external forces such as gravity and the collision with other objects. In Section III the action selection strategy is described. The strategy has the dual objective of straightening the cloth while at the same time maintaining good estimation of its state. Section IV presents both simulated and experimental results, and Section V contains some concluding remarks. A short video

The authors are with the Institut de Robòtica i Informàtica Industrial, CSIC-UPC. Llorens Artigas 4-6, Barcelona 08028, Spain. scuen,cetto,torras@iri.upc.edu.

sequence accompanies the paper validating the approach both in simulation and in a real robotic work-cell.

II. PHYSICAL MODEL OF CLOTH DEFORMATION

To model its deformation, a piece of cloth may be modeled as a triangular mesh of particles. Each vertex of the triangular mesh has coordinates $\mathbf{p}_i$, and moves with velocity $\mathbf{v}_i$. The deformation of the entire mesh after a time step h is given by the difference equation

$$\begin{bmatrix} \mathbf{p}(t+h) \\ \mathbf{v}(t+h) \end{bmatrix} = \begin{bmatrix} \mathbf{p}(t) + \mathbf{v}(t)h \\ \mathbf{v}(t) + \Delta\mathbf{v} \end{bmatrix} \tag{1}$$

in which the vectors $\mathbf{p}$ and $\mathbf{v}$ are vertical concatenations of all the vertex locations and velocities, respectively.

The change in velocity for the interconnected particles, $\Delta\mathbf{v}$, follows the backward Euler method for implicit time integration given in [11]. The method, in contrast to a more simple forward Euler integration technique, finds an output state whose time derivative is consistent with the initial state. The method is used to simulate the effect of internal and external forces applied to the cloth. These forces acting on each particle are defined in terms of precondition functions that permit to impose constraints on the velocities, effectively allowing us to model the effect of motion commands on those particles fixed by the manipulator gripper:

$$\Delta\mathbf{v} = \frac{h\mathbf{W}\left(\mathbf{f}_0 + h\frac{\partial\mathbf{f}}{\partial\mathbf{p}}\mathbf{v}_0\right) + \mathbf{u} + \delta\mathbf{u}}{\left(\mathbf{I} - h\mathbf{W}\frac{\partial\mathbf{f}}{\partial\mathbf{v}} - h^2\mathbf{W}\frac{\partial\mathbf{f}}{\partial\mathbf{p}}\right)} \tag{2}$$

The solution for $\Delta\mathbf{v}$ depends on the initial particle velocities $\mathbf{v}_0 \in \mathbb{R}^{3n}$, initial particle vector forces $\mathbf{f}_0 \in \mathbb{R}^{3n}$ (that may include external forces such as gravity, wind, etc), the internal forces $\mathbf{f} \in \mathbb{R}^{3n}$ and the particle positions and velocities. Particle internal forces are modeled as the sum of resistance and damping effects on specific stretch, shear and bend conditions

$$\mathbf{f} = -k\frac{\partial\mathbf{C}(\mathbf{p})}{\partial\mathbf{p}}\mathbf{C}(\mathbf{p}) - d\frac{\partial\mathbf{C}(\mathbf{p})}{\partial\mathbf{p}}\dot{\mathbf{C}}(\mathbf{p})$$

where

$$\dot{\mathbf{C}}(\mathbf{p}) = \frac{\partial\mathbf{C}(\mathbf{p})}{\partial\mathbf{p}}\dot{\mathbf{p}}$$

and $\mathbf{v} = \dot{\mathbf{p}}$.

A. Internal Forces

In Baraff's formulation [11], $\mathbf{C}(\mathbf{p})$ is a condition vector which we want to be zero. Its associated energy $E = \frac{k}{2}\mathbf{C}(\mathbf{p})^\top\mathbf{C}(\mathbf{p})$ is used to derive simple stretch, shear and bend conditions. So, for example if $\mathbf{w}_u(\mathbf{p}_i)$ and $\mathbf{w}_v(\mathbf{p}_i)$ are the vectors indicating the stretch or compression of a particular triangle in the mesh, with unit length when the material is unstretched, and a is the area of the triangle in uv coordinates, the condition

$$\mathbf{C}(\mathbf{p}_i) = a\begin{pmatrix} \|\mathbf{w}_u(\mathbf{p}_i)\| - 1 \\ \|\mathbf{w}_v(\mathbf{p}_i)\| - 1 \end{pmatrix}$$

TABLE I

SET OF POSSIBLE ACTIONS TO STRAIGHTEN A PIECE OF CLOTH.

Action	
`drag-upright`	$v_x, v_y > 0$
`drag-upleft`	$v_x < 0, v_y > 0$
`drag-downright`	$v_x > 0, v_y < 0$
`drag-downleft`	$v_x, v_y < 0$

can be used to model stretch energy. Similarly, by the small angle approximation, shear can be measured as the inner product between $\mathbf{w}_u(\mathbf{p}_i)$ and $\mathbf{w}_v(\mathbf{p}_i)$

$$\mathbf{C}(\mathbf{p}_i) = a\mathbf{w}_u(\mathbf{p}_i)^\top\mathbf{w}_v(\mathbf{p}_i)\,.$$

Finally, if we let $\mathbf{n}_i$ and $\mathbf{n}_j$ denote the unit normals of two adjacent triangles, and let $\mathbf{e}$ be a common vector parallel to the common edge, the angle between the two faces defined by the relations $\sin\theta = (\mathbf{n}_i \times \mathbf{n}_j)^\top\mathbf{e}$ and $\cos\theta = \mathbf{n}_i^\top\mathbf{n}_j$, the condition that counters bending along that edge is

$$\mathbf{C}(\mathbf{p}_{ij}) = \theta\,.$$

B. Motion commands

Our cloth dynamics model, the time varying partial differential Equation (2), differs from the original equation in [11] in that we have included a set of external induced velocities $\mathbf{u}$, representing the actions exerted by our manipulator. To account for the effects of linearization and unmodelled artifacts we also add to each external action on the cloth a stochastic term $\delta\mathbf{u}$ with zero mean white Gaussian distribution with covariance $\mathbf{Q}$.

Input commands belong to a limited set of actions depending on the task to be performed. Table I shows, for example, a set of possible actions for the straightening of a cloth on the table. Each such action is intended to drag a corner in the cloth at a constant speed and for a short period of time, $\mathbf{u}_i = (v_x, v_y, 0)^\top$.

C. Managing External Collisions

Collision of the deformable planar object with external rigid objects is handled with the help of a particle constrainer matrix $\mathbf{W}$, whose block diagonal elements are defined as $\mathbf{W}_i = \frac{1}{m_i}\mathbf{S}_i$, where m_i is the mass of the *i*-th particle, and $\mathbf{S}_i$ is a 3×3 matrix used to constraint the three degrees of freedom affecting the particle mass at any given location. Our approach does not handle yet internal cloth collisions. This is a topic of further research.

III. ACTION SELECTION

A. Predicting the Outcome of Actions

To estimate the state of the deformable planar object after an action is executed, particle positions and velocities are considerd as a Gaussian random vector, with an initial covariance $\mathbf{P}_{0|0}$. A Kalman filter is then used to track the state of the deformable planar object. For every possible action, the state mean can be obtained from Eqs. (1) and

(2) with $\delta \mathbf{u} = \mathbf{0}$, and an estimated change in covariance can be computed with the linearized expression

$$\mathbf{P}_{t+h|t} = \mathbf{A}\mathbf{P}_{t|t}\mathbf{A}^\top + \mathbf{B}\mathbf{Q}\mathbf{B}^\top$$

where the Jacobian $\mathbf{A}$ takes the form

$$\mathbf{A} = \begin{bmatrix} \mathbf{I} & h\mathbf{I} \\ \frac{\partial \mathbf{v}}{\partial \mathbf{p}} & \mathbf{I} \end{bmatrix}$$

in which the partial derivative of a particle velocity with respect to its position is

$$\frac{\partial \mathbf{v}}{\partial \mathbf{p}} = (AB - CD)(B^\top B)^{-1}$$

$$A = h^2 \mathbf{W} \frac{\partial^2 \mathbf{f}}{\partial \mathbf{p} \partial \mathbf{p}} \mathbf{v}_0$$

$$B = \mathbf{I} - h\mathbf{W}\frac{\partial \mathbf{f}}{\partial \mathbf{v}} - h^2 \mathbf{W}\frac{\partial \mathbf{f}}{\partial \mathbf{p}}$$

$$C = h\mathbf{W}(\mathbf{f}_0 + h\frac{\partial \mathbf{f}}{\partial \mathbf{p}}\mathbf{v}_0) + \mathbf{u} + \delta\mathbf{u}$$

$$D = -h\mathbf{W}\frac{\partial^2 \mathbf{f}}{\partial \mathbf{v} \partial \mathbf{p}} - h^2 \mathbf{W}\frac{\partial^2 \mathbf{f}}{\partial \mathbf{p} \partial \mathbf{p}}$$

and the Jacobian $\mathbf{B}$ is

$$\mathbf{B} = \begin{bmatrix} \mathbf{0} \\ (\mathbf{I} - h\mathbf{W}\frac{\partial \mathbf{f}}{\partial \mathbf{v}} - h^2 \mathbf{W}\frac{\partial \mathbf{f}}{\partial \mathbf{p}})^{-1} \end{bmatrix} .$$

The state of the object can then be revised from the observation of some points as measured by our stereo vision system. Assuming that the error from our sensor $\delta \mathbf{z}_i$ is also zero mean Gaussian with covariance $\mathbf{R}$, each particle observed, whereas it is a corner or not, contributes to revise the state estimate with

$$\begin{bmatrix} \mathbf{p}_{t+h|t+h} \\ \mathbf{v}_{t+h|t+h} \end{bmatrix} = \begin{bmatrix} \mathbf{p}_{t+h|t} \\ \mathbf{v}_{t+h|t} \end{bmatrix} + \mathbf{K}(\mathbf{z}_i - \mathbf{p}_i) .$$

If using sequential innovation for each particle, the measurement Jacobian is a row block of zeros, only with a selective 3×3 identity matrix at the i-th block cell, and the Kalman gain becomes the $6n \times 3$ matrix

$$\mathbf{K} = \mathbf{P}_{t+h|t,i}(\mathbf{P}_{t+h|t,ii} + \mathbf{R})^{-1}$$

where $\mathbf{P}_{t+h|t,ii}$ is the position covariance for the *i*-th particle, and $\mathbf{P}_{t+h|t,i}$ represents the *i*-th column block of the full state covariance matrix. The update of the state covariance becomes

$$\mathbf{P}_{t+h|t+h} = (\mathbf{I} - [\ \mathbf{0}_{6n \times 3i-1} \quad \mathbf{K} \quad \mathbf{0}_{6n \times 3(2n-i)-2}\ ])\mathbf{P}_{t+h|t} .$$

TABLE II

PARAMETER VALUES FOR IMPLEMENTATION.

Parameter	Symbol	Value
Stretch resistance	k_{st}	5000
Shear resistance	k_{sr}	500
Bend resistance	k_b	0.00001
Stretch damping	d_{st}	1000
Shear damping	d_{sr}	100
Bend damping	d_b	2×10^{-6}
Gravity		9.81

B. Action Selection

A strategy is developed to straighten the cloth by choosing from a limited set of possible actions, the one that maximises the information gain for our state estimate. The set of actions under inspection are: one possible motion command from Table I for each corner of the object. The commands evaluated are those that drive the cloth corners away from the center. In essence, the strategy is aimed at choosing, from four possible choices, which corner is to be dragged next, based on the current estimate that we have about its location.

A classic approach would be to chose the action that maximizes the relative entropy between prior and posterior covariance estimates [15], [16], [17], that for our multivariate Gaussian case reduces to computing the expression

$$I = \frac{1}{2}(\log|\mathbf{P}_{t+h|t}| - \log|\mathbf{P}_{t+h|t+h}|) .$$

A D-optimality measure of information gain however may become unreliable when one or more of the state space directions is constrained, since it is computed from the product of the eigenvalues of $\mathbf{P}$. The conditions that constraint the motion of particles include contact with an obstacle, or the mere dragging action. In such cases, there is absolute information about some components of the location (and/or velocity) of such particle, with the consequence of semi-defínteness on the estimation covariance. For this reason, we use an A-optimality meause of information instead [18] to minimize the squared error of the model, which uses the sum of the eigenvalues (instead of their product), given by the trace of the covariance matrix.

$$I = \text{tr}\,(\mathbf{P}_{t+h|t}) - \text{tr}\,(\mathbf{P}_{t+h|t+h}) \tag{3}$$

IV. IMPLEMENTATION AND RESULTS

In our experiment setup, our workcell is composed of a robotic manipulator *Stäubli RX-60* with a *FTC-Schunk* force sensor attached to the end-effector, and a *Bumblebee* stereo camera. The force sensor is used to ensure that a sufficiently large perpendicular force is maintained while dragging a piece of cloth against a table.

To measure the state of the cloth at any given instance, a set of feature points must be observed. One possiblity is to select scale invariant salient features on the object and match them against a previously trained dataset [19]. The experiments reported here are less complex in terms of

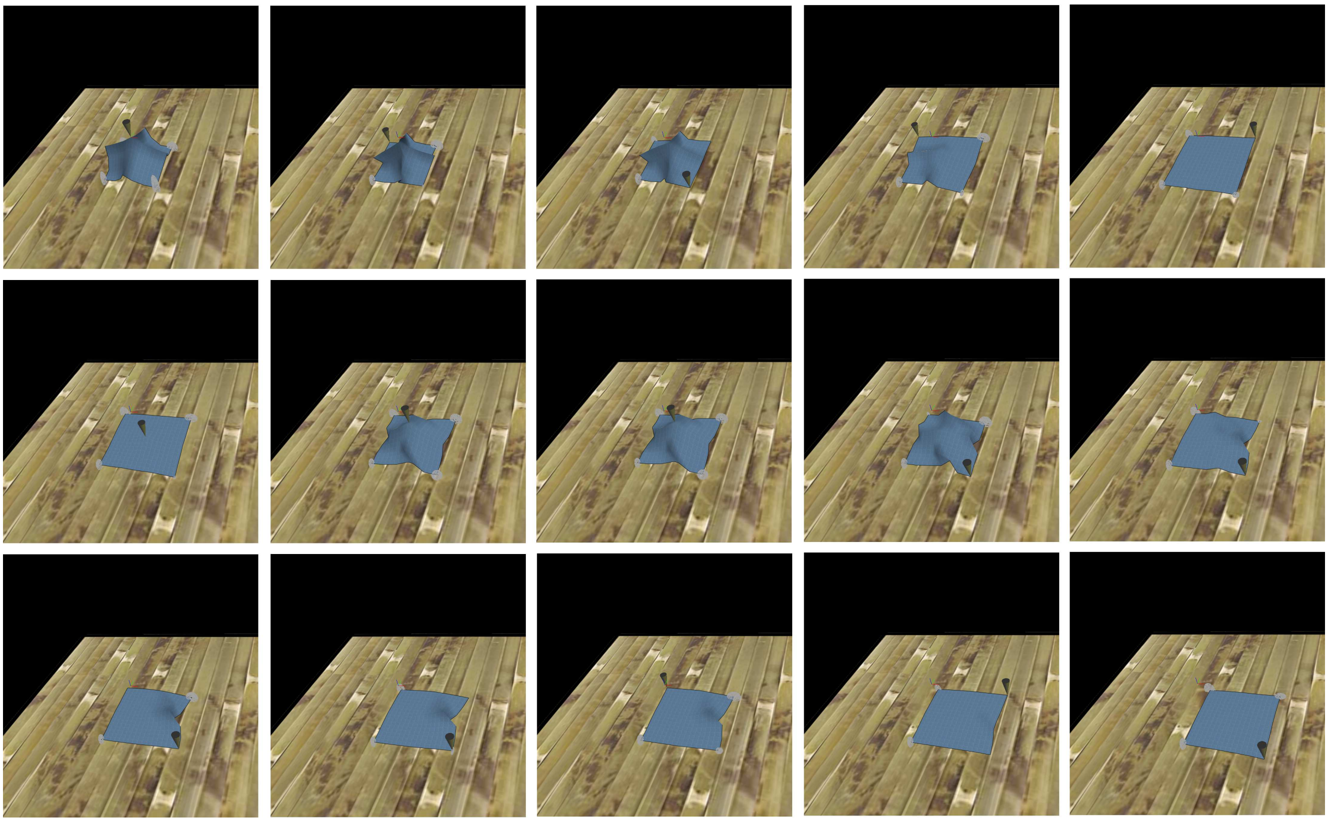

Fig. 1. Computer simulation of action selection of planar deformable obejcts. Time goes from left to right, then from top to bottom. The hyperellipsoids on the corners indicate surfaces of equal probability for the corner location estimates.

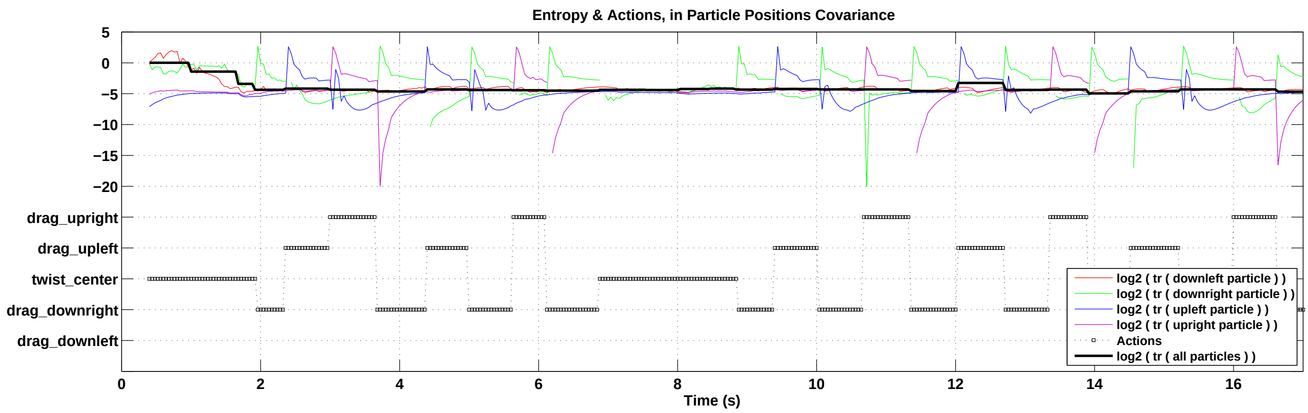

Fig. 2. Evolution of the trace of the covariance at each particle for the same simulation, as well as a history of the chosen motion commands

the computer vision tools used. For the porpouses of our straightening task, we are content with tracking the four corners of the cloth piece binarizing the image of the object, and detecting the corners by selecting the discontinuities of the 1-D signal along the object contour using multiresolution and non-maximum supression [20]. Once the corner points are located, their location with respect to the camera is computed with stereo triangulation. Given that the camera location is calibrated with respect to the robot workcell, the measurement of the corner points can be given in world coordinates.

To constraint the velocity of the particles during manipulation, a velocity profile must be generated for each possible action command. The set of possible actions is restricted to the actions in Table I. For the simulations however, an extra motion `twist` has been implemented, which drags two center particles to emulate a wrinkling effect to restart the simulation.

The cloth parameters used are given in Table II. Fig. 1 shows a simulated manipuation sequence. Each frame in the sequence represents the current state of the deformable object. The ellipsoids drawn at the corners represent surfaces of equal probability at one-standard deviation, and are used to indicate the value of the estimation covariance at that point. To easy their visibility, these ellipsoids have been magnified by a factor of ten. The sequence shows two instances of the evolution of the cloth straightening task. The cone pointing to each of the corners emulates the manipulator end-effector. Figure 2 contains the evolution of the trace of the covariance at each particle for the same simulation, as well as a history of the chosen motion commands.

Finally, Fig. 3 presents a sequence with a real straightening task on our robotic workcell. These images, as well as the video accompanying this paper illustrate the feasibility of the presented approach for the information-oriented action selection for the manipulation of planar deformable objects for simple household applications.

V. CONCLUSIONS AND FUTURE WORK

If the adaptive robot manipulation of rigid objects is already a challenging research topic, the manipulation of deformable objects poses additional difficulties. An important one is the representation of the state of such objects. For the rigid ones, once a CAD model of the object is available, its state at a given time is uniquely determined by the six parameters of its pose. Contrarily, flexible objects require models that accommodate their possible deformations as a result of their manipulation or other external causes. In this work we have coupled a stochastic state estimator with a physical-based implicit integration model of a deformable planar object (a cloth). This model has then been used to predict the effect of manipulation actions on the cloth, a critical feature for planning sequences of such actions. Here, as a first step to test state estimation, only the best next action to achieve a given goal is determined. The particular goal pursued has been a weighted combination of two objectives, namely, straightening the cloth while at the same time maintaining a good estimation of its state. Action selection relies on a maximization information criterion. The obtained results, both in simulation and in a real robotic work-cell, have been satisfactory. In sum, we are proposing a framework for goal-driven manipulation of deformable planar objects.

Envisaged future work will be along three lines. First, we would like to incorporate detection of self-collisions into the FEM model of cloth [21]. Second, we aim to come up with a characterization of qualitative-different states of deformable planar objects (e.g., foldings), in a similar way as knots describe states of deformable linear objects. And third, we like to go beyond single action selection to develop planning strategies for manipulating pieces of cloth, more specifically, for unfolding and then folding them in prescribed ways.

VI. ACKNOWLEDGMENTS

This work is supported by the EU PACO-PLUS project FP6-2004-IST-4-27657, by the Spanish Ministry of Education and Science through a Ramn y Cajal Fellowship to JAC, and by the Catalan Research Commission, through a FI scholarship to SCR, and the Consolidated Robotics Group.

REFERENCES

[1] R. M. Murray, Z. Li, and S. S. Sastry, *A Mathematical Introduction to Robotic Manipulation*. CRC Press, 1994.

[2] M. T. Mason, *Mechanics of Robotic Manipulation*. MIT Press, 2001.

[3] "PACO-PLUS: Perception, action & cognition through learning of object-action complexes," 2006. [Online]. Available: http://www.paco-plus.org/

[4] J. Brown, J. C. Latombe, and K. Montgomery, "Real-time knot-tying simulation," *Visual Comput.*, vol. 20, no. 2-3, 2004.

[5] M. Saha and P. Isto, "Motion planning for robotic manipulation of deformable linear objects," in *Proc. IEEE Int. Conf. Robot. Automat.*, Orlando, May 2006, pp. 2478–2484.

[6] M. Moll and L. E. Kavraki, "Path planning for deformable linear objects," *IEEE Trans. Robot.*, vol. 22, no. 4, pp. 625–636, 2006.

[7] M. Greminger, Y. Sun, and B. Nelson, "Boundary element deformable object tracking with equilibrium constraints," in *Proc. IEEE Int. Conf. Robot. Automat.*, New Orleans, Apr. 2004, pp. 3896–3901.

[8] K. Salleh, H. Seki, Y. Kamiya, and M. Hikizu, "Tracing manipulation of deformable objects using robot grippers with roller fingertips," in *Proc. SICE-ICASE Joint Conf.*, Busan, Oct. 2006, pp. 5882–5887.

[9] A. M. Howard and G. A. Bekey, "Intelligent learning for deformable object manipulation," *Auton. Robot.*, vol. 9, no. 1, pp. 51–58, 2000.

[10] D. Henrich and H. Wörn, Eds., *Robot Manipulation of Deformable Objects*, ser. Advanced Manufacturing. Springer, 2000.

[11] D. Baraff and A. Witkin, "Large steps in cloth simulation," in *Computer Graphics. Proc. ACM SIGGRAPH Conf.* Orlando: ACM Press, Jul. 1998, pp. 43–54.

[12] X. Provot, "Deformation constraints in a mass-spring model to describe rigid cloth behavior," in *Proc. Graphics Interface*. Quebec: Canadian Human-Computer Communications Society, 1995, pp. 147–154.

[13] H. Ng and R. Grimsdale, "Computer graphics techniques for modeling cloth," *IEEE Comput. Graphics Applicat.*, vol. 16, no. 5, pp. 28–41, Sep. 1996.

[14] S. Gibson and B. Mirtich, "A survey of deformable modeling in computer graphics," MERL, Cambridge, Tech. Rep. TR-97-19, Nov. 1997.

[15] T. Vidal-Calleja, A. Davison, J. Andrade-Cetto, and D. Murray, "Active control for single camera SLAM," in *Proc. IEEE Int. Conf. Robot. Automat.*, Orlando, May 2006, pp. 1930–1936.

[16] J. Andrade-Cetto and F. Thomas, "Wire-based tracking using mutual information," in *Advances in Robot Kinematics*. Ljubljana: Springer-Verlag, Jun. 2006, pp. 3–14.

[17] D. J. C. MacKay, "Information based objective functions for active data selection," *Neural Comput.*, vol. 4, no. 4, pp. 589–603, 1992.

[18] R. Sim and N. Roy, "Global A-optimal robot exploration in SLAM," in *Proc. IEEE Int. Conf. Robot. Automat.*, Barcelona, Apr. 2005, pp. 673–678.

[19] D. Pritchard and W. Heidrich, "Cloth motion capture," in *Proc. Eurographics*, vol. 22, no. 3. Granada: Blackwell Publishing, Sep. 2003, pp. 263–271.

[20] R. Deriche, "Fast algorithms for low-level vision," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 12, no. 1, pp. 78–88, Jan. 1990.

[21] M. Teschner, S. Kimmerle, B. Heidelberger, G. Zachmann, L. Raghupathi, A. Fuhrmann, M.-P. Cani, F. Faure, N. Magnenat-Thalmann, W. Strasser, and P. Volino, "Collision detection for deformable objects," *Computer Graphics Forum*, vol. 24, no. 1, pp. 61–81, Mar. 2005.

[22] D. Pritchard, "Freecloth 0.7.1. a free, open-source cloth simulation tool," 2003. [Online]. Available: http://sourceforge.net/projects/freecloth

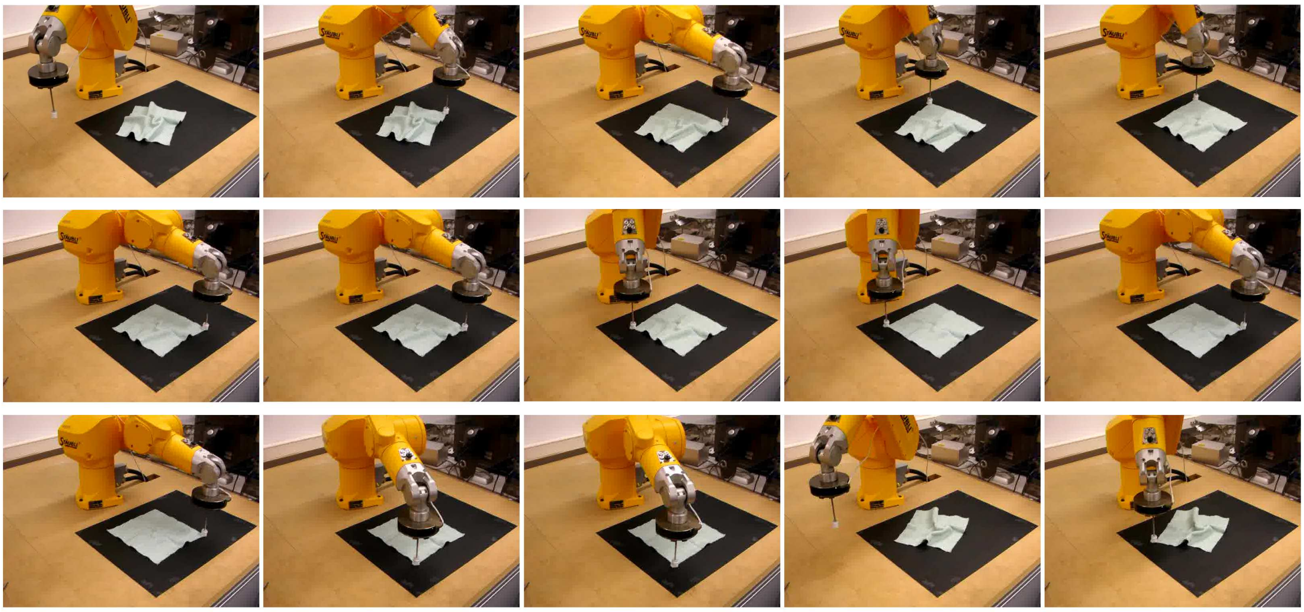

Fig. 3. A view of the workcell setup. Time goes from left to right, then from top to bottom