Proceedings of the 1997 IEEE
International Conference on Robotics and Automation
Albuquerque, New Mexico - April 1997

Trajectories for Optimal Temporal Integration in Active Vision Systems

James J. Clark[1] and Lei Wang[2]

[1]Centre for Intelligent Machines, McGill University
[2]Centre for Imaging Science, Washington University

Abstract— **We describe a general technique for specifying trajectories of controllable imaging parameters in an active vision system so that temporal integration processes are optimized. The technique assumes that a Kalman filter is used to perform the temporal integration of measurements and is based on determining, at each point in time, the set of imaging parameter values that minimizes the trace of the state estimate error covariance matrix. We present the application of this technique to the active vision task of extracting the location and orientation of a plane from shadows cast on it with a position controlled light source.**

I. Introduction

Active vision systems are characterized by control over various aspects of the imaging process. A number of active vision techniques have been proposed that use such control to improve the process of obtaining 3D information about object surfaces. For example, Whaite and Ferrie [15] introduced a method for determining the location of a sensor that minimizes some measure of uncertainty about the shape of an object. They applied this to the task of building up object surface descriptions from range-finder data. Shmuel and Werman [11] presented a Kalman Filter based temporal integration scheme for depth from stereo where the cameras are positioned so as to minimize an uncertainty measure. The Kalman filter allows information from a sequence of sensor readings to be integrated, resulting in a reduction in the effect of noise on the information derived from the sensor data. Each of these techniques shares a common thread in that they have control over some aspect of the imaging process, and use this control in a way so as to optimize the extraction of information from the sensor data.

In this paper we present a technique for determining the trajectory of an arbitrary control variable in an active sensing system so as to minimize the error in a result derived from the sensor data. This technique, based on the use of the extended Kalman filter, has a general applicability to active sensing systems, and has the advantage over the techniques used in the work cited above in that the control law is expressible in closed form, and its implementation is straightforward. To demonstrate its use we apply it to the task of obtaining surface information using the shadows cast by a position controlled point light source.

II. Optimal Control of Imaging Parameters

Although the assumptions used in the derivation of the Kalman filter are not usually satisfied in vision problems, the Kalman filter, it has been successfully used in many vision algorithms. For example, the various vision tasks described in [1], [?], [4], [6], [8], [9], [12] use the Kalman filter or its many extensions. Likewise, we will use the extended Kalman filter as the basic temporal integration process in our technique.

The extended Kalman filter equations are [5]:

$$\hat{x}_k = \hat{x}_{k-1} + K_{k-1}[y_k - h_k(\hat{x}_{k-1})] \tag{1}$$

where $\hat{x}_k$ is the estimate of the state vector and y_k is the vector of measurement values at time k. The function $h(\hat{x})$ is the measurement function which relates the noise free measurement vector to the state vector (i.e. $y = h(x) + \eta$, where η is a vector of Gaussian random variables with mean zero and covariance matrix R). measurement. K is the Kalman gain,

$$K_k = P_{k-1}H_k^T(\hat{x}_{k-1})[H_k(\hat{x}_{k-1})P_{k-1}H_k^T(\hat{x}_{k-1}) + R_k]^{-1} \tag{2}$$

The state estimate error covariance is updated is using

$$P_k = [I - K_k H_k(\hat{x}_{k-1})]P_{k-1} \tag{3}$$

In the above equations $H_k(\hat{x}_{k-1})$ is the linear term of the taylor's series expansion of the measurement function $h(\vec{x})$ evaluated at $\vec{x} = \hat{x}_{k-1}$,

$$H_k(\hat{x}_{k-1}) = \left.\frac{\partial h_k(\vec{x})}{\partial \vec{x}}\right|_{\vec{x}=\hat{x}_{k-1}} \tag{4}$$

The state estimate error covariance matrix P provides a measure for the amount of uncertainty in the estimate of the state variables. At each step of the Kalman filter estimation process the information provided by the new measurement reduces the error in the state estimates. In all active vision systems there are certain controllable parameters which affect the measurements that are made. We can specify the values of these parameters so as to maximize the reduction in the state estimate error covariance due to the incorporation of the measurement.

In order to implement such a procedure we require a well-posed optimization criterion. In the work of

0-7803-3612-7-4/97 $5.00 © 1997 IEEE

Whaite and Ferrie [15], the determinant of the state estimate error covariance matrix was used as an objective function for optimization purposes. The trace of the state estimate error covariance matrix, suitably weighted, is also a reasonable quantity to be optimized. If the state estimate errors are uncorrelated, then the trace of P is the sum of the state estimate error variances. The diagonal elements of the covariance matrix must be weighted to account for difference in units between the state estimate error variances. This weighting can be implemented by pre- and post- multiplying the covariance matrix by a constant diagonal matrix before computing the trace. In terms of the notation introduced earlier, the optimization problem associated with this choice of objective function can be expressed as follows (we drop the time step subscript from H and K for clarity):

$$\vec{u} = \arg\min_{\vec{u}} tr[WP_kW] = \arg\min_{\vec{u}} tr[W(I-KH)P_{k-1}W] \tag{5}$$

where $\vec{u}$ is the vector of parameters under our control, and W is a suitable diagonal weighting matrix. This problem is equivalent to

$$\vec{u} = \arg\max_{\vec{u}} \epsilon(\vec{u}) \tag{6}$$

where

$$\epsilon(\vec{u}) = tr[WK(\vec{u})H(\vec{u})P_{k-1}W] \tag{7}$$

Noting that $K = P_{k-1}H^T[HP_{k-1}H^T + R]^{-1}$ it can be seen that we can express ϵ in the simple form:

$$\epsilon = tr[A^TBA] \tag{8}$$

by taking

$$A = HP_{k-1}W \tag{9}$$

and

$$B = [HP_{k-1}H^T + R]^{-1} \tag{10}$$

The matrix B is symmetric. The matrices A and B typically depend on the value of the control vector $\vec{u} = (u_1, u_2, ..., u_n)$ to be used in obtaining the measurement. This dependency arises through the dependency of the measurement matrix H and the measurement noise covariance matrix R on the control variables. The prior covariance matrix P_{k-1} does not depend on the new control variables, only on their previous values.

The quantity, $\epsilon(\vec{u})$ that is to be maximized is, in general, a nonlinear function of $\vec{u}$. As such, solving for the value of $\vec{u}$ that minimizes ϵ may be difficult. The situation is exacerbated by the fact that, in practice, there will be constraints on the values of $\vec{u}$ that can be used. For example, a camera may be mounted on the end of a robot arm which has a limited range of motion. This means that solution of a constrained nonlinear optimization problem is required. We can, however, use the constraints on the values of the control parameters to our advantage in solving the optimization problem. In many active vision tasks that involve temporal integration, measurements are taken after small, incremental, changes of the imaging parameters (see, for example, [9]). In such a case we can replace the global constraints (those due to physical limits on the actuators that implement the imaging parameter changes) with artificially imposed limits that constrain the space of possible control parameter values much more. In what follows we will make the assumption that the range of possible control parameter values is restricted to a spherical region centred on the current value of the control parameter vector, and that the radius of this region is sufficiently small so as to allow the assumption that the gradient of ϵ with respect to the components of $\vec{u}$ is constant. We further assume that this constant gradient is non-zero. These assumptions imply that the extrema will lie on the boundary of the constraint sphere and that the extremal values are equal to

$$\vec{u}^* = \vec{u}_0 \pm r\frac{\nabla_{\vec{u}}\epsilon(\vec{u})}{|\nabla_{\vec{u}}\epsilon(\vec{u})|} \tag{11}$$

where r is the radius of the constraint sphere, $\vec{u}_0$ is the current control vector, and

$$\nabla_{\vec{u}} = \left(\frac{\partial}{\partial u_1}, \frac{\partial}{\partial u_2}, ..., \frac{\partial}{\partial u_n}\right) \tag{12}$$

In general the limits on the changes of the different control parameters will not all be the same, as is required by the sphericity assumption on the shape of the constraint region. In fact, the different control variables may not even have the same units. We handle this by normalizing the control parameters by dividing each of them by their respective maximum allowed changes. Thus the constraint region becomes the unit sphere.

It may appear that such an incremental approach results in a gradient descent on the objective function $\epsilon(\vec{u})$. The objective function depends, however, on the current state vector estimate, $\hat{x}_k$. Thus, the form of the objective function is constantly changing, and so the "landscape" which is being "descended" is constantly changing its topography. In such a situation the objective function will never increase, but the system is not guaranteed to converge to a minimum. This aspect of incremental active vision control systems was noted by Whaite and Ferrie [15]. They observed that the change in the objective function due to the information gained at a sensor location is of a

form that decreases the value of the objective function at the current location relative to other locations, thereby forcing the sensor away from its current location. This type of repulsion from the current location prevents the system from being trapped in local minima.

We now derive an expression for the gradient of $\epsilon(\vec{u})$ with respect to the components of the control parameter vector $\vec{u}$. It suffices to find a general form for the gradient components, so we will determine the partial derivative of $\epsilon(\vec{u})$ with respect to a specific component u_i. In the derivation of the expression for the gradient we will use the following matrix differentiation identities:

Let $X = X(t)$ and $Y = Y(t)$ be matrix functions of a scalar t. Then

$$\frac{\partial}{\partial t} tr[X^T Y X] = tr\left[\frac{\partial}{\partial t}(X^T Y X)\right] \tag{13}$$

Let $X(t) = Y^{-1}(t)$ be a symmetric matrix function of a scalar t. Then

$$\frac{\partial X}{\partial t} = -X\left(\frac{\partial Y}{\partial t}\right)X \tag{14}$$

From equation (8) we have that:

$$\frac{\partial \epsilon}{\partial u_i} = \frac{\partial}{\partial u_i} tr[A^T B A] \tag{15}$$

Using the first of the matrix differentiation identities this can be written as:

$$\frac{\partial \epsilon}{\partial u_i} = tr\left[\frac{\partial A^T}{\partial u_i}(BA) + A^T\frac{\partial B}{\partial u_i}A + A^T B\frac{\partial A}{\partial u_i}\right] \tag{16}$$

From the definition of the matrix A we have that:

$$\frac{\partial A}{\partial u_i} = \frac{\partial H}{\partial u_i} P_{k-1} W \tag{17}$$

(since P_{k-1} and W do not depend on the new value of the control vector $\vec{u}$). Using the second identity we can write:

$$\frac{\partial B}{\partial u_i} = -B\left(\frac{\partial H}{\partial u_i}P_{k-1}H^T + HP_{k-1}\frac{\partial H^T}{\partial u_i} + \frac{\partial R}{\partial u_i}\right)B \tag{18}$$

Expanding equation (16), using the linearity of the trace operator, we get

$$\begin{aligned}&\frac{\partial \epsilon}{\partial u_i} = \\ &tr\left[WP_{k-1}\frac{\partial H^T}{\partial u_i}BA\right] - tr\left[A^T B\frac{\partial H}{\partial u_i}P_{k-1}H^T BA\right] - \\ &tr\left[A^T BHP_{k-1}\frac{\partial H^T}{\partial u_i}BA\right] + \\ &tr\left[A^T B\frac{\partial H}{\partial u_i}P_{k-1}W\right] - \\ &tr\left[A^T B\frac{\partial R}{\partial u_i}BA\right]\end{aligned} \tag{19}$$

Since the trace of a matrix is equal to the trace of its transpose, the above equation can be written as

$$\begin{aligned}&\frac{\partial \epsilon}{\partial u_i} = 2tr\left[A^T B\frac{\partial H}{\partial u_i}P_{k-1}W - A^T B\frac{\partial H}{\partial u_i}P_{k-1}H^T BA\right] - \\ &tr\left[A^T B\frac{\partial R}{\partial u_i}BA\right]\end{aligned} \tag{20}$$

Using $A^T B = WK$, we get

$$\begin{aligned}&\frac{\partial \epsilon}{\partial u_i} = 2tr\left[WK\frac{\partial H}{\partial u_i}P_{k-1}(I - H^T K^T)W\right] - \\ &tr\left[WK\frac{\partial R}{\partial u_i}K^T W\right]\end{aligned} \tag{21}$$

Note that $P_{k-1}(I - H^T K^T) = (I - KH)P_{k-1} = P_k$, the update to the covariance matrix. Thus we can write:

$$\frac{\partial \epsilon}{\partial u_i} = tr\left[WK\left(2\frac{\partial H}{\partial u_i}P_k - \frac{\partial R}{\partial u_i}K^T\right)W\right] \tag{22}$$

Usually R depends on u only indirectly, through the influence of u on the measurement function h. So we use

$$\frac{\partial R}{\partial u_i} = \frac{\partial R}{\partial h}\frac{\partial h}{\partial u_i} \tag{23}$$

in the above equation. The quantity $\frac{\partial R}{\partial h}$ is obtained from a model for how the measurement noise varies with the measured value.

III. Application to Shape From Controlled Shadow Motion

Control of imaging parameters in active vision systems is not limited to control over the position of the image sensor. One can control other parameters such as the position of the illumination source. For example, Clark [2] has described a technique for obtaining the shape and absolute position of objects from shading information that rely on the control over the position of a light source. Likewise, Wang and Clark [13] describe a technique for obtaining the shape of an object onto which a shadow is cast which uses the controlled motion of a light source to provide the required data. In this section we show how the trajectory optimization procedure given above can be used in this application.

There have been proposed many techniques that use shadow motion to obtain information about surfaces (e.g. [7], [10]). The technique described here, developed by Wang and Clark [13], are different than the previously cited approaches in that it uses controlled motion of a nearby light source, and provides absolute surface depth information as well as surface normal information. The geometry of the Wang-Clark method is shown in figure 1. For simplicity we show the 2-D case, but the analysis given below is for the more

general 3D case. We assume that the light source is a point source, and that the background object (onto which the shadow is cast) is planar and that the foreground object (which casts the shadow) is a quadric surface. As the algorithm determines shape locally, the assumptions on the background and occluding object shape need only be satisfied locally. The background plane can be defined by a vector, $\vec{n}$, normal to its surface, and by the vector from the origin to any point on the plane, $\vec{P}$, via the following equation:

$$\vec{n}_p^T \vec{P} + 1 = 0 \tag{24}$$

The three components of $\vec{n}$ define the plane. The quadric surface can be defined with the following equation:

$$\vec{Q}^T M_q \vec{Q} + \vec{t}_q^T \vec{Q} = 1, \tag{25}$$

respectively. where M_q is a positive definite symmetric matrix that describes the curvature of the object's surface, and $\vec{Q}$ is any point on the surface. The six independent parameters of M and the three components of $\vec{t}$ define the quadric surface. In total, there are twelve parameters which define the shape of the background plane and the shadowing object. In addition to these shape parameters, we also must determine the location in space of the shadow boundary on the background plane, $\vec{r}_p$, and on the shadowing object, $\vec{r}_q$. These will change as the light source moves. The shape parameters will not change as the light source moves.

With the measurements $\vec{i}_p$ and $\vec{i}_q$, and the known light source position $\vec{r}_l$, we seek to find the least number of such equations from which we can solve for the unknown shape and shadow boundary position parameters: $\vec{n}_p$, $\vec{t}_q$, M_q, $\vec{r}_p$, and $\vec{r}_q$. Each light source position gives rise to a set of constraint equations that describe the geometric relations among points on the surfaces and in the image plane. Assuming perspective projection and a pinhole camera model, the image formation process gives rise to the following *image formation equations*:

$$\vec{i}_p = \frac{f}{r_{pz}} \vec{r}_p; \quad \vec{i}_q = \frac{f}{r_{qz}} \vec{r}_q \tag{26}$$

Noticing that, for each light source position, the vectors $\vec{r}_l$, $\vec{r}_p$, and $\vec{r}_q$, are coplanar, we have the *illumination constraint equation*:

$$(\vec{r}_p - \vec{r}_l) \times (\vec{r}_q - \vec{r}_l) = 0, \tag{27}$$

Only two of the three equations provided by this constraint are independent, however. To see this, consider the case where both $\vec{r}_l$ and $\vec{r}_q$ are known. These then define a line on which $\vec{r}_p$ must fall, but leave free where on the line $\vec{r}_p$ falls.

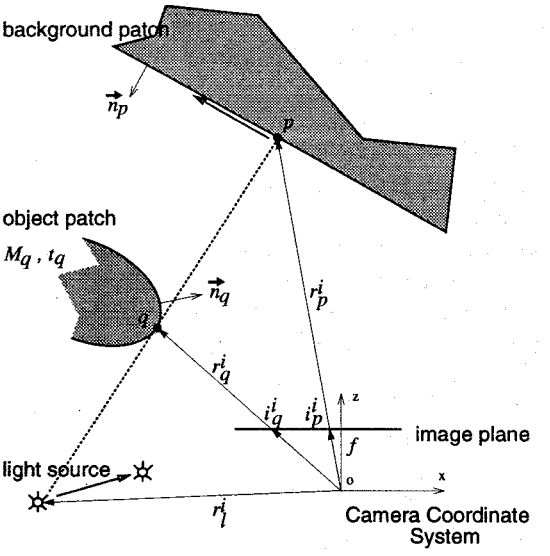

Fig. 1. Geometry for the active shape from shadows technique.

Finally, the light rays for which the cast shadow boundaries are formed on the background are tangential to the quadric surface patch, therefore we have the following *tangent equations*:

$$(\vec{r}_q - \vec{r}_l)^T \vec{n}_q = 0, \tag{28}$$

where $\vec{n}_q = 2 M_q \vec{r}_q + \vec{t}_q$

Counting the equations provided by the above geometric constraints, we see that there are $7 * i$ independent equations, where i is the number of different light source positions we make measurements at. There are a total of $12 + 6 * i$ unknowns (the 12 shape parameters and 6 shadow position parameters for each light source position). Thus we need at least 12 light source positions to solve for the unknowns. This results in a system of 84 quadratic equations in 84 unknowns! It can be shown (see [14] for details) that the unknowns can be decoupled so that we need only solve a system of 12 quadratic equations for the 12 shape parameters, from which the position parameters can be straightforwardly computed. Even so, the solution of a system of 12 quadratic equations cannot be obtained in closed form and numerical solutions will be plagued by sensitivity to noise and multiple solutions.

If, however, we make the assumption that the object that is casting the shadow contains a "sharp" edge at the shadow boundary then the problem simplifies greatly. At a sharp edge the self-shadow boundary does not move as the light source moves. In this case

the shadowing object can be defined by only the location of the points along its sharp edge, which will coincide with the location of the shadow boundary $\vec{r}_q$. Thus the number of unknown shape parameters has been reduced to those of the background plane only. Furthermore, in the sharp corner case, the location of the self-shadow boundary $\vec{r}_q$ does not vary with the light source position, further reducing the number of unknowns to be determined.

Under these conditions we get an equation, for each light source position, which is linear in four of the unknowns $\vec{n}_p$ and $1/r_{qz}$:

$$\begin{pmatrix} \frac{i_{px}}{f} & \frac{i_{py}}{f} & 1 & \frac{\gamma_x}{\omega_x} \end{pmatrix} \begin{pmatrix} n_{px} \\ n_{py} \\ n_{pz} \\ 1/r_{qz} \end{pmatrix} = -\frac{\xi_x}{\omega_x}. \quad (29)$$

where

$$\begin{aligned} \omega_x &= f\, r_{lx} - i_{qx}\, r_{lz} \\ \gamma_x &= f\, r_{lx} - i_{px}\, r_{lz} \\ \xi_x &= i_{px} - i_{qx}. \end{aligned} \quad (30)$$

A linear system of four such equations, obtained from four different positions of the light source,

$$\begin{pmatrix} \frac{i_{px}^{(1)}}{f} & \frac{i_{py}^{(1)}}{f} & 1 & \frac{\gamma_x^{(1)}}{\omega_x^{(1)}} \\ \frac{i_{px}^{(2)}}{f} & \frac{i_{py}^{(2)}}{f} & 1 & \frac{\gamma_x^{(2)}}{\omega_x^{(2)}} \\ \frac{i_{px}^{(3)}}{f} & \frac{i_{py}^{(3)}}{f} & 1 & \frac{\gamma_x^{(3)}}{\omega_x^{(3)}} \\ \frac{i_{px}^{(4)}}{f} & \frac{i_{py}^{(4)}}{f} & 1 & \frac{\gamma_x^{(4)}}{\omega_x^{(4)}} \end{pmatrix} \begin{pmatrix} n_{px} \\ n_{py} \\ n_{pz} \\ 1/r_{qz} \end{pmatrix} = - \begin{pmatrix} \frac{\xi_x^{(1)}}{\omega_x^{(1)}} \\ \frac{\xi_x^{(2)}}{\omega_x^{(2)}} \\ \frac{\xi_x^{(3)}}{\omega_x^{(3)}} \\ \frac{\xi_x^{(4)}}{\omega_x^{(4)}} \end{pmatrix} \quad (31)$$

will give a unique solution for the shape variable $\vec{n}_p$, and depth of the object edge point, r_{qz}, as long as the coefficient array is invertible. From these values, the remaining parameters can be recovered as follows:

$$\vec{r}_q = \frac{r_{qz}}{f}\vec{i}_q; \quad \vec{r}_p = \frac{-1}{\vec{i}_p{}^T \vec{n}_p}\vec{i}_p \quad (32)$$

Using these equations, we can solve, in closed form, for the desired quantities from image measurements obtained at four different light source positions. The results of doing this straightforward solution are unsuitable in practice due to noise in the localization of the images of the shadow boundaries and to uncertainties in the location of the light source. To alleviate this sensitivity to noise somewhat, we can apply an extended Kalman filter to provide estimates of the shape and position parameters, integrating information obtained over many light source positions. This indeed does provide acceptable results in real-world applications, as detailed in [14]. We can, however, do even better if we have control over the light source position. In this approach we make small displacements of the light source, as prescribed by equation (11). After each displacement we make a measurement of the shadow boundary locations, and update the Kalman filter. To implement this technique we need to determine expressions for H, the linearized measurement equation, and $\frac{\partial H}{\partial u}$, the derivatives of H with respect to the control parameters. Due to their complexity we do not include these expressions here. They can be found in [14].

In the remainder of the paper we will present the results of an implementation of this approach in a real robotic system. In this experiment the light source was a 250W quartz light with a fiber-optic light guide output. The light guide terminated in a 2mm aperture, over which a diffusing layer was placed. The position of the light guide terminus was controlled via a six degree-of-freedom robotic manipulator, the American Cimflex Merlin. The camera was mounted on a tripod which was fixed relative to the base of the robot. The camera viewed a scene consisting of a sharp edged object that cast shadows onto a flat background wall. The coordinate transformations between the camera image plane and the initial light source position were determined via a calibration process. Details of the calibration process can be found in [14].

In figures 2, 3, and 4 are shown results of two separate experiments. In the first, an iterated extended Kalman filter was used to integrate information where the light source was moved in a pre-set helical trajectory. In the second experiment the Kalman filter was used to integrate information where the light source was moved according to the optimal control specified by equation (11).

Figures (2) and (3) shows the convergence of the state variable estimates. Note that using the optimal control results in the estimates converging much more rapidly than when using the preset trajectory (by about 10 steps versus 40 steps). Note also the shape of the optimal trajectory.

Figure (4) shows a depiction of the surface patches reconstructed by the active shape-from-shadowing process in the case of the optimally controlled trajectories. Shown are the locations of the recovered 3D locations of the shadow boundaries on the shadowing and background objects. The recovered surface normal of the background is not depicted.

IV. Acknowledgements

This research was performed at the Harvard Robotics lab and was supported in part by the National Science Foundation through grants IRI-92-23676 and CDR-85-00108.

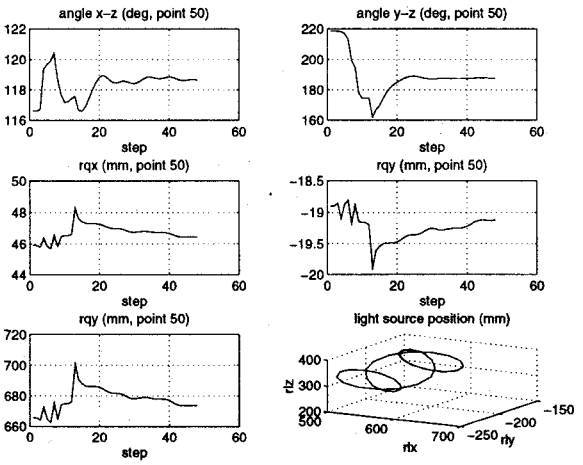

Fig. 2. The convergence of the state variable estimates for the pre-determined light source trajectory.

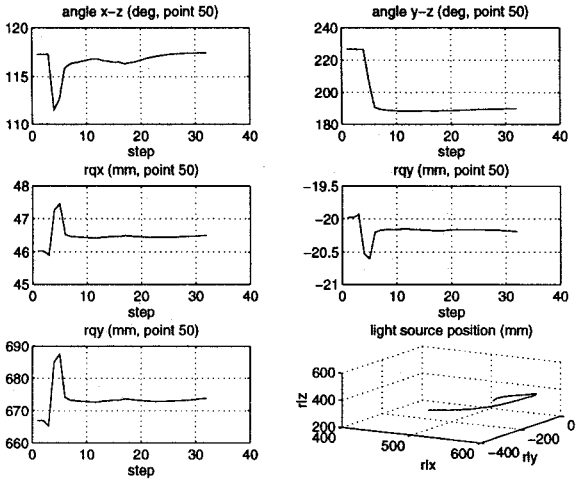

Fig. 3. The convergence of the state variable estimates for the optimally controlled light source trajectory.

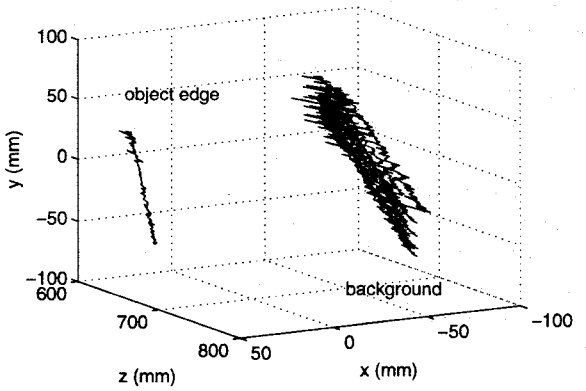

Fig. 4. The reconstruction of the shadowing object edge and the background plane at the 50th time step for the optimally controlled light source trajectory.

References

[1] Chaumette, F. and Boukir, S., "Structure from motion using an active vision paradigm", *Proc. 11th Int. Conf. Pattern Recognition*, The Hague, Netherlands, 1991

[2] Clark, J.J., "Active Photometric Stereo", in *Proceedings of the 1992 IEEE Computer Vision and Pattern Recognition Conference*, Champaign, IL, June 1992, pp 29-35

[3] Durrant-White, H.F., **Integration, Coordination, and Control of Multi-Sensor Robot Systems'**, Kluwer Academic Publishers, 1988

[4] Faugeras, O., Ayache, N., and Faverjon, B., "Building visual maps by combining noisy stereo measurements", Proceedings of the 1986 *IEEE International Conference on Robotics and Automation*, pp 1433-1438

[5] Gelb, A. (ed.), **Applied Optimal Estimation,** MIT Press, Cambridge, MA, 1974

[6] Harris, C., "Tracking with rigid models" in **Active Vision,** Blake, A. and Yuille, A., eds., MIT Press, Cambridge, MA 1992

[7] Kender, J. and Smith, E., "Shape from darkness: Deriving surface information from dynamic shadows", *Proceedings of AAAI*, 1987, pp 539-546

[8] Martinez, J.M. and Montano, L., "A camera motion strategy to localize uncertain 3D lines", *Proc. 1993 IEEE Int. Conf. on Systems, Man and Cybernetics*, Le Touquet France, pp 517-522, 1993

[9] Matthies, L.H., Kanade, T., and Szeliski, R., "Kalman filter-based algorithms for estimating depth from image sequences", in *International Journal of Computer Vision*, Vol. 3, pp 209-236, 1989

[10] Raviv, D., Pao, Y. and Loparo, K., "Reconstruction of three-dimensional surfaces from two-dimensional binary images", *IEEE Transactions on Robotics and AUtomation*, Vol. 5, No. 5, 1989, pp 701-710

[11] Shmuel, A. and Werman, M., "Active vision: 3D from an image sequence", Proceedings of the 10th International Conference on Pattern Recognition, June 1990, pp 48-54

[12] Terzopoulos, D. and Szeliski, R., "Tracking with Kalman snakes", in **Active Vision**, Blake, A. and Yuille, A., eds., MIT Press, Cambridge, MA 1992

[13] Wang, L. and Clark, J.J., "Depth from active shadow motion", SPIE Conference on Intelligent Robots and Computer Vision XII, Boston, MA, September 1993, pp. 2-13

[14] Wang, Lei, **3D Structure from Active Shadowing,** Ph.D. Thesis, Division of Applied Sciences, Harvard University, 1995

[15] Whaite, P, and Ferrie, F.P., "Autonomous exploration: driven by uncertainty", *Proceedings of 1994 IEEE Conference on Computer Vision and Pattern Recognition*, Seattle WA, June 1994, pp 339-346