

Controllability and Observability: Tools for Kalman Filter Design

B Southall^{‡†}, B F Buxton[†] and J A Marchant[‡]

[‡]Silsoe Research Insitute, Wrest Park, Silsoe,
Bedfordshire MK45 4HS, UK

[†]Department of Computer Science, University College London,
London WC1E 6BT, UK

B.Southall@cs.ucl.ac.uk

Abstract

Kalman's optimum linear filter has proved to be immensely popular in the field of computer vision. A less often quoted contribution of Kalman's to the control theory literature is that of the concepts of *controllability* and *observability* which may be used to analyse the state transition and observation equations and give insights into the filter's viability. This paper aims to highlight the usefulness of these two ideas during the design stage of the filter and, as well as presenting the standard solutions for linear systems, uses a practical vision application (that of tracking plants for an autonomous crop protection vehicle) to illustrate a useful special case where these methods may be applied to a non-linear system. The application of tests for controllability and observability to the practical non-linear system give not only confirmation that the filter will be able to produce stable estimates, but also gives a lower bound on the number of features required from each image for it to do so.

1 Introduction

The Kalman filter [1] has proved to be an immensely popular tool in computer vision, particularly in the field of tracking, e.g. [3, 5, 2], and also in parameter estimation [13] and stereo matching [7]. By contrast, another two issues raised by Kalman, *controllability* and *observability*, are seldom seen in the machine vision literature, with the exception of [12], who has shown their value as a design aid for linear time-invariant filters.

This paper gives the definitions of the two concepts, along with the standard techniques for testing whether linear systems are controllable and observable, and the implications of uncontrollability or unobservability are discussed. In general, there is no equivalent test for controllability and observability for non-linear systems, but a useful exception for machine vision algorithms is illustrated here with reference to a real application; that of controlling an autonomous crop protection vehicle. The class of non-linear problems where the exception applies are those where several observations are compared to a single prior state prediction, for example when a group of features extracted from an

image are evaluated with respect to a prediction of the group's position in the image. The tests not only provide confirmation that, if it is initialised correctly, the extended Kalman filter will (given good numeric conditioning) yield stable state estimates, but they also give a lower bound on the number of observed features required from each image of a sequence.

2 Controllability and observability in linear systems

The definitions and results given here for controllability and observability of linear systems may be found in any standard textbook, e.g. [4, 6]. In the explanations given below, the discrete time invariant linear systems are given by

$$\mathbf{x}(k+1) = \mathbf{A}\mathbf{x}(k) + \mathbf{B}\mathbf{u}(k) \quad (1)$$

$$\mathbf{z}(k) = \mathbf{H}\mathbf{x}(k) \quad (2)$$

where $\mathbf{x}(k)$ is the system state at time step k , $\mathbf{u}(k)$ the control input, $\mathbf{z}(k)$ is the observation vector, \mathbf{A} is the state evolution matrix, \mathbf{B} the control gain matrix, and \mathbf{H} the observation matrix. In the time variant case, \mathbf{A} , \mathbf{B} and \mathbf{H} are matrix functions of time. The tests given below are also applicable to continuous time systems.

2.1 Controllability

A system is said to be *controllable* if every state vector $\mathbf{x}(k)$ can be transformed to a desired state in finite time by the application of unconstrained control inputs $\mathbf{u}(k)$. Evidently then, an *uncontrollable* system is one where some elements of the state vector $\mathbf{x}(k)$ cannot be affected by the control input.

The test for controllability of a time-invariant system is given by:

A system with state vector \mathbf{x} of dimension n , is controllable if the *controllability matrix*

$$\mathbf{C} = [\mathbf{B}, \mathbf{A}\mathbf{B}, \dots, \mathbf{A}^i\mathbf{B}, \dots, \mathbf{A}^{n-1}\mathbf{B}] \quad (3)$$

has column rank n (i.e. n linearly independent columns).

The proof of this statement is based on successive substitutions of (1) to find a solution for $\mathbf{x}(N)$ in terms of an original state $\mathbf{x}(0)$ and a series of control inputs $\mathbf{u}(0) \dots \mathbf{u}(N-1)$. For a time-varying system, where the *state transition matrix* $\Phi(k_0, k_1)$ describes the cumulative change of $A(k)$ in the interval $k_0 + 1 \leq k_1$, the controllability is defined by the rank of the matrix $\Phi(k_0, k+1)\mathbf{A}(k)$. Proofs for both cases may be found in [4].

2.2 Observability

A system is said to be *observable* at a time step k_0 if for a state $\mathbf{x}(k_0)$ at that time, there is a finite $k_1 > k_0$ such that knowledge of the outputs \mathbf{z} from k_0 to k_1 are sufficient to determine state k_0 . Following this, an *unobservable* system is one where the values of some elements in the state vector at time k_0 may not be determined from examination of the system output regardless of the number of observations taken. That observability is

specified over an interval highlights that whilst a single observation of the system at time k may not be enough to obtain the complete state, additional observations may allow the full state information to be accumulated. Evidently, for time invariant systems, the time k_0 is unimportant.

For linear time invariant systems, the test for observability is given by:

A system with state vector \mathbf{x} of dimension n is observable if the *observability matrix*

$$\mathbf{O} = \begin{bmatrix} \mathbf{H} \\ \mathbf{HA} \\ \vdots \\ \mathbf{HA}^i \\ \vdots \\ \mathbf{HA}^{n-1} \end{bmatrix} \quad (4)$$

has row rank n (i.e. n linearly independent rows)

The proof of this test [4] uses (1) and (2) to determine the value of $\mathbf{z}(k)$ for $0 \leq k \leq n - 1$ in terms of $\mathbf{x}(0)$ and the known control inputs $\mathbf{u}(k)$ during that time period. It is then straightforward to show that $\mathbf{x}(0)$ can be completely determined (and is therefore observable) if the matrix \mathbf{O} has full row rank. A test for observability between time steps k_0 and k_1 for linear time varying systems may be found in [4] (and relates to the rank of the matrix $\mathbf{H}(k)\Phi(k, k_0)$).

3 Implications for Kalman filtering

The above tests for controllability may be applied to stochastic systems which provide the Kalman filter prediction and observation equations:

$$\mathbf{x}(k + 1) = \mathbf{Ax}(k) + \mathbf{Bu}(k) + \mathbf{Vv}(k) \quad (5)$$

$$\mathbf{z}(k) = \mathbf{Hx}(k) + \mathbf{w}(k). \quad (6)$$

In (5) the matrix \mathbf{V} describes the gains applied to the vector of independent, identically distributed unit variance Gaussian noise sources \mathbf{v} which stipulate how noise affects the state evolution. The noise terms may simply be treated as another control input, and the controllability test described above may be applied, using matrix \mathbf{V} in place of \mathbf{B} ; the result indicates the ability of the noise to affect the state. The additive zero-mean Gaussian noise term \mathbf{w} in (6) clearly does not affect the measure of observability, which is solely dependent on \mathbf{H} and \mathbf{A} .

From (5) and (6) the Kalman filter prediction equations may be written; these will be used to illustrate the implications of controllability and observability in Kalman filter performance.

$$\hat{\mathbf{x}}(k + 1|k) = \mathbf{A}\hat{\mathbf{x}}(k|k) + \mathbf{Bu}(k) \quad (7)$$

$$\mathbf{P}(k + 1|k) = \mathbf{AP}(k|k)\mathbf{A}^T + \mathbf{Q}. \quad (8)$$

$$\hat{\mathbf{z}}(k + 1|k) = \mathbf{H}\hat{\mathbf{x}}(k + 1|k). \quad (9)$$

The term $\hat{\mathbf{x}}$ is the filter state estimate, and \mathbf{P} the filter estimate of the state covariance. \mathbf{Q} is the covariance matrix of the system noise, i.e. $\mathbf{V}^T\mathbf{V}$. The notation $k|k$ indicates the quantity's value at time step k , given k observations, and $k+1|k$ indicates the value of the variable predicted for time step $k+1$ given k observations. The quantity $\hat{\mathbf{z}}$ is the predicted observation, i.e. the position of the predicted state in observation space.

The action of the Kalman filter is to improve the prediction $\hat{\mathbf{x}}(k+1|k)$ by blending the predicted observation $\hat{\mathbf{z}}(k+1|k)$ with the real observation $\mathbf{z}(k+1)$ in such a way that the total mean square error on the estimate (the trace of the matrix \mathbf{P}) is minimised.

3.1 Controllability

If the pair \mathbf{A} and \mathbf{B} form an uncontrollable system, this has no implications for the Kalman filter other than that the control inputs \mathbf{u} will not affect every element of the state estimate $\hat{\mathbf{x}}$ during the prediction step. This will not affect tracking performance unless, of course, this reveals a flaw in the modelling of the system dynamics. However, if the system formed by \mathbf{A} and \mathbf{V} is uncontrollable (and there is no coupling between the states via the matrix \mathbf{B}), this means that the Gaussian noise sources in \mathbf{v} do not affect all of elements of the state, i.e. they are uncorrupted by the noise. The diagonal elements of \mathbf{P} corresponding to these “uncorruptible” states will be driven to zero by the Kalman filter (which minimises the trace of \mathbf{P}), and once this has happened, the estimates of these states are fixed; no further observations will alter their values.

Although such an approach may at first seem a little curious, there is sometimes good reason to allow the variance of an estimate to collapse to zero. In [8] the mean of a process is estimated; clearly a mean is a single, fixed quantity which does not vary with time. If the estimation process allows accurate evaluation of the mean it is quite correct to use the process covariance to constrain the value of the estimated mean by allowing the estimate variance to reduce to zero (the issue of controllability of this system is discussed in depth in the later work [12]). In [11], where the covariance matrix is used to provide partial constraints on aligning patches of range data, it is noted that in certain degenerate cases there are problems with their algorithm, although this is not named as a controllability problem.

Inspection of (8), which governs the evolution of the state covariance \mathbf{P} , will confirm that certain diagonal elements cannot increase if the \mathbf{A}, \mathbf{V} pair are uncontrollable. The first term in the equation, $\mathbf{A}\mathbf{P}(k|k)\mathbf{A}^T$ describes the process of noise transfer between states during state evolution – by definition, in an uncontrollable system noise cannot transfer into any uncontrollable state via this mechanism, therefore it cannot be responsible for any increase in the variance of uncontrollable states. The second term in (8) is the covariance \mathbf{Q} ($= \mathbf{V}^T\mathbf{V}$) which represents the direct noise input; evidently in an uncontrollable system this contribution to \mathbf{P} must be zero for any uncontrollable states (this follows from the form of \mathbf{V} , which *must* have zero elements in the row corresponding to the uncontrollable state; otherwise noise is being directly input to that state at each iteration). If the diagonal elements of \mathbf{P} corresponding to uncontrollable states cannot increase, the Kalman filter will drive these elements to zero, after which point new observations will not alter the estimate of any uncontrollable states.

3.2 Observability

Whilst uncontrollable systems are sometimes desirable for Kalman filtering, a Kalman filter built around a system with unobservable states will simply not work. By definition, an unobservable state is one about which no information may be obtained through the observation equations; in the absence of information, the filter estimate for that state will not converge on a meaningful solution.

4 Non-linear systems – a practical application

Non-linear systems are those where the state transition and/or observations of a system are functions of the state of the system:

$$\mathbf{x}(k+1) = \mathbf{a}(\mathbf{x}(k), \mathbf{v}(k), k) \quad (10)$$

$$\mathbf{z}(k) = \mathbf{h}(\mathbf{x}(k), \mathbf{w}(k), k) \quad (11)$$

where \mathbf{a} is the state evolution function and \mathbf{h} the observation function. Both functions here are stochastic, with random noises \mathbf{v} and \mathbf{w} in the state and observation equations respectively. There is no rule for determining the controllability and observability of a general non-linear system; so no prior judgements may be made about the extended Kalman filter's viability. Despite this lack of a general test of controllability/observability for non-linear systems, it is possible to show that the test for observability may be adapted for a certain class of non-linear systems which can be particularly useful for vision applications.

4.1 A practical example

The practical application used to illustrate the techniques discussed above is that of steering an autonomous agricultural vehicle across a field of crop (such as cauliflower or sugar beet), whilst simultaneously picking out the individual plant positions. The extended Kalman filter described below is based upon that in [9], but has been altered to address some of the problems highlighted in that work. The system model is illustrated and the controllability and observability of the system is discussed.

4.2 The system model

Figure 1 shows a schematic view of a patch of crop, with the plants being represented by black circles. There are two sets of axes in the figure, (x_w, y_w) and (x_c, y_c, z_c) which represent the world and camera co-ordinate systems respectively, with the world z_w axis projecting out of the page, and camera axis z_c projecting into the page. It can be seen that the world y axis is coincident with the middle plant row. The model of the crop consists of the two measurements; r , the spacing between the rows and l the space between plants in the row.

Three parameters (t_x , Y and Ψ), specify the position of the model relative to the camera as shown in the diagram. The measurement Y is the offset in world co-ordinates from the world origin of the plant in the central row at the bottom of the image. The offset of the camera y_c axis from the world origin is given by t_x . It can be seen then that the t_x and

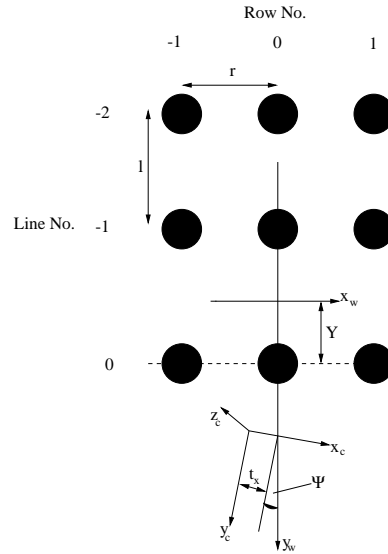


Figure 1: The crop planting model

Ψ parameters may be used to provide navigation information in terms of a heading angle and offset of the camera (and hence the vehicle) relative to the rows, and the parameter Y can yield the position of individual plants via (12) and (13) below.

$$x_w = nr \quad (12)$$

$$y_w = ml + Y \quad (13)$$

The quantities $n \in \{-1, 0, 1\}$ and $m \in \{0, -1, -2, \dots, -(m_{max} - 1)\}$ index into the $3 \times m_{max}$ grid formed by the planting pattern. It should be noted that the plant centres are assumed to be on the ground plane ($z_w = 0$). It is stressed that the model describes the grid on which individual plants should lie, rather than the actual location of the individual plants.

In [9], the quantities estimated by the Kalman filter are t_x , Y and Ψ , whilst the model parameters r and l are assumed fixed (the imperfect plant positioning being somewhat inappropriately modelled as observation noise). To rectify this situation, and allow for the fact that r and l are imperfect, the Kalman filter also estimates the mean values \bar{r} and \bar{l} of r and l respectively for the plants currently in view. The implications of this will be discussed further below.

4.3 System dynamics and controllability

With the state vector $\mathbf{x} = [t_x, Y, \Psi, \bar{r}, \bar{l}]^T$, the system dynamics are described by

$$\mathbf{x}(k+1) = \mathbf{I}_5 \mathbf{x}(k) + \mathbf{I}_5 \begin{bmatrix} U_{t_x} \\ U_Y \\ U_\Psi \\ 0 \\ 0 \end{bmatrix} + \mathbf{V} \mathbf{v}(k) \quad (14)$$

where \mathbf{I}_5 is the 5×5 identity matrix and the subscripted U values control inputs relating the vehicle's motion between image acquisitions. It should be noted that, except for noise, the mean plant spacings \bar{r} and \bar{l} are assumed constant. The noise gain matrix is diagonal, such that

$$\mathbf{V}^T(k)\mathbf{V}(k) = \mathbf{Q}(k) = \begin{bmatrix} \sigma_{t_x}^2(k) & 0 & 0 & 0 & 0 \\ 0 & \sigma_Y^2(k) & 0 & 0 & 0 \\ 0 & 0 & \sigma_\Psi^2(k) & 0 & 0 \\ 0 & 0 & 0 & \sigma_{\bar{r}}^2(k) & 0 \\ 0 & 0 & 0 & 0 & \sigma_{\bar{l}}^2(k) \end{bmatrix} \quad (15)$$

The σ^2 terms relate to the mean-square error introduced at each prediction step. For t_x , Y and Ψ they are provided by the vehicle control Kalman filter, but for \bar{r} and \bar{l} the σ^2 terms are non-zero *only* when a new row of plants come into view. The reason for this is that plant noise \mathbf{Q} should be added only when new information is available, and for the estimated mean values of the plant spacing parameters, new information is introduced whenever a new row of plants appears in the image. By inspection of $\mathbf{Q}(k)$ with $\sigma_{\bar{r}}^2$ and $\sigma_{\bar{l}}^2$ zero (i.e. when no new plants have come into view) it can be seen that the \mathbf{AV} pair ($\mathbf{A} = \mathbf{I}_5$ and $\mathbf{V} = \sqrt{\mathbf{Q}(k)} = \text{diag}(\sigma_{t_x}, \sigma_Y, \sigma_\Psi, \sigma_{\bar{r}}, \sigma_{\bar{l}})$) is uncontrollable, so the estimated covariance on the estimates of both \bar{r} and \bar{l} will start to converge to zero. This is quite acceptable, however, because during the period where no new plants come into view, the filter is refining its estimate of the plant spacing using several views of the same patch of crop. When the new plants are seen, process noise *is* added which reflects the fact that the mean spacing parameters may well be different for the patch of crop containing these new plants.

Figure 2 shows the estimate of \bar{r} over a set of 20 images for the uncontrollable system described above (figure 2, right) and a fully controllable system where noise is added to \bar{r} with each new image (figure 2, left). As can be seen, the estimate from the controllable system is not as smooth as that from the uncontrollable one, and the variance of the estimate does not shrink over the sequence. Both of these behaviours are to be expected when new plant noise is injected for each image. The uncontrollable system has a smoother estimate, and it can be seen that the variance of this estimate converges as the same set of plants is seen over several images. The increase in estimate variance at image numbers 4 and 12 correspond to the appearance of new plants in the image. Similar behaviour has been observed for the estimate of \bar{l} .

4.4 Observability in a non-linear system

The system described above is observed as plant centres in a perspective image [9]. This perspective projection leads (by using (12) and (13) with the imaging model of Tsai [10]) to the following non-linear observation system:

$$\mathbf{z}(k, m, n) = \mathbf{h}[k, \mathbf{x}(k), m, n, \mathbf{w}(k)] = \begin{bmatrix} x_f(\mathbf{x}(k), m, n) \\ y_f(\mathbf{x}(k), m, n) \end{bmatrix} + \mathbf{w}(k) \quad (16)$$

where (x_f, y_f) is the image position in pixels of a plant centre on the ground plane:

$$x_f(x, Y, \Psi, \bar{r}, \bar{l}, m, n) = \frac{f}{dx} \frac{n\bar{r} + \Psi(m\bar{l} + Y) + t_x}{n\bar{r}\Psi \sin \phi - (m\bar{l} + Y) \sin \phi + t_z} + C_x \quad (17)$$




Figure 2: A comparison of state estimates and variances from a controllable and uncontrollable system. The middle trace shows the state estimate, the upper and lower traces showing the estimate ± 2 standard deviations. Left: controllable system output. Right: uncontrollable system output.

$$y_f(x, Y, \Psi, \bar{r}, \bar{l}, m, n) = \frac{f}{dy} \frac{(m\bar{l} + Y - \Psi n\bar{r}) \cos \phi}{n\bar{r}\Psi \sin \phi - (m\bar{l} + Y) \sin \phi + t_z} + C_y \quad (18)$$

C_x and C_y are the co-ordinates of the camera's optic centre on the CCD in pixels, f the focal length of the camera, and dx and dy the side-lengths of the pixels.

The observability criteria expressed in (4) does not apply to the system described in (16). However, it can be shown by substituting the Taylor series for $\mathbf{h}[k, \mathbf{x}(k), \mathbf{w}(k)]$ into (16) and truncating second order and higher terms, that a test similar to that of (4) may be derived, with the Jacobian of $\mathbf{h}[k, \mathbf{x}(k), m, n]$ evaluated at $\mathbf{x} = \hat{\mathbf{x}}(k+1|k)$, replacing the observation matrix \mathbf{H} . In this example, \mathbf{h}_x (after taking out common factors on each row, which is acceptable for rank analysis) is given by :

$$\mathbf{h}_x = \begin{bmatrix} 1 & \Psi & m\bar{l} + Y - n\bar{r} & n(1 - Y) & m\Psi \\ 0 & 1 & -n\bar{r} & -n\Psi & m \end{bmatrix} \quad (19)$$

On inspection of (14), it can be seen that the matrix \mathbf{A} of (7) is simply the identity matrix. Therefore, the rank of the observability matrix (4) is the same as the rank of \mathbf{h}_x in (19).

For a system to be observable, its row rank must be equal to the dimension of the state vector \mathbf{x} , in this case 5. As can be seen in (19), \mathbf{h}_x has only two rows (although they are independent), so at first glance it would appear that the system is unobservable, and the extended Kalman filter cannot work. However, in this application, where there is one state prediction (prior) for a set of many features, it is possible to incorporate all the measurements in parallel, as in [9], where the features observed in each image are incorporated simultaneously in batches. As part of this batch update process a stacked observation matrix is formed by stacking the matrix of (19) associated with each individual feature extracted from the image into a single matrix, \mathbf{h}_{xstack} , with $2s$ rows for observations of s features, where each feature has its own unique pair of values m, n .

The observability of the system may now be assessed by performing a rank analysis on the stacked matrix \mathbf{h}_{xstack} (in cases where \mathbf{A} is not the identity matrix, \mathbf{A} and \mathbf{h}_{xstack} should be substituted into (4)). Inspection of the second row of (19) shows that regardless of the values of the state variables ($t_x, Y, \Psi, \bar{r}, \bar{l}$), i.e. whatever the point of linearisation of (16), the stacked matrix \mathbf{h}_{xstack} will have one independent row for each different

value of m found in the feature set. For the first row of (19) to be unique for different features (values of m and n) places some constraints on the state variables; however, a little thought shows that sufficient conditions are that $\bar{r} \neq 0$ and $\bar{l} \neq 0$, which stipulates that the crop must lie on a two dimensional grid (i.e. that the model of figure 1 is valid). In short, provided a sufficient number of observations (see below) are available, \mathbf{h}_{xstack} will be of full row rank and the linearised system will be observable.

To generate the required minimum of 5 independent rows, only 3 features are required (which will give an \mathbf{h}_{xstack} matrix with 6 rows), providing that all three do not share the same value of m or n , or in other words, providing that the three features do not all lie within a single row or column of the grid structure. This result confirms that provided by common sense; one feature can partially locate the grid (giving t_x and Y), a second feature determines the orientation Ψ and grid parameter \bar{r} (or \bar{l}), and a third measurement, provided that it does not lie on the same row (or column) will provide the second grid parameter \bar{l} (or \bar{r}).

5 Conclusions

The control theoretic concepts of controllability and observability have been introduced and their implications for Kalman filter design underlined. Despite the lack of a general test for the controllability and observability of non-linear systems, it has been shown that for a non-linear problem where several observations are evaluated with respect to a single prior (here the task of determining ego-motion from perspective views) it is possible to test for the observability of the system, and that this test provides a lower bound on the number of observations required for the system to function.

Acknowledgements

This work was funded by the BBSRC. Thanks to Tony Hague for useful discussions on many aspects of the Kalman filter.

References

- [1] Y Bar-Shalom and T Fortmann. *Tracking and Data Association*. Academic Press, New York, 1988.
- [2] A Blake, R Curwen, and A Zisserman. A framework for spatio-temporal control in the tracking of visual contours. *International Journal of Computer Vision*, 11(2):127–145, 1993.
- [3] R Deriche and O Faugeras. Tracking line segments. *Image and Vision Computing*, 8(4):261 – 270, November 1990.
- [4] M Gopal. *Modern Control System Theory*. Wiley Eastern Ltd., 1984.
- [5] C Harris. Tracking with rigid models. In A Blake and A Yuille, editors, *Active Vision*, chapter 4. MIT Press, 1992.
- [6] A H Jazwinski. *Stochastic Process and filtering theory*. Academic Press, 1970.
- [7] S B Pollard, T P Pridmore, J Porrill, J E W Mayhew, and J P Frisby. Geometrical modelling from multiple stereo views. *International Journal of Robotics Research*, 8(4):132–138, 1989.

- [8] D Reynard, A Wildenberg, A Blake, and J A Marchant. Learning dynamics of complex motions from image sequences. In B Buxton and R Cipolla, editors, *Computer Vision – ECCV ‘96*, Lecture Notes in Computer Science. Springer, April 1996.
- [9] B Southall, J A Marchant, T Hague, and B F Buxton. Model based tracking for navigation and segmentation. In H Burkhardt and B Neumann, editors, *Proceedings 5th European Conference on Computer Vision*, volume 1, pages 797–811, June 1998.
- [10] R Y Tsai. An efficient and accurate camera calibration technique for 3D machine vision. In *Proc. IEEE Computer Soc. Conf. Comp Vis Patt Recog.*, Miami Beach, June 1986.
- [11] M Waite, M Orr, R Fisher, and J Hallam. Statistical partial constraints for 3D model matching and pose estimation problems. In *Proc. British Machine Vision Conference*, pages 105–114, 1993.
- [12] A P Wildenberg. *Learning and Initialisation for Visual Tracking*. PhD thesis, University of Oxford Dept. Engineering Science, Oxford, UK, 1997.
- [13] Z Y Zhang. Parameter estimation techniques: a tutorial with respect to conic fitting. *Image and Vision Computing*, 15(1):59–76, 1997.