Image Feature Extraction Using a Method Derived from the Hough Transform with Extended Kalman Filtering

Sergio A. Velastin and Chengping Xu

Digital Imaging Research Centre, Kingston University,
Kingston upon Thames, KT1 2EE, United Kingdom
sergio.velastin@kingston.ac.uk

Abstract. The conventional implementation of the Hough Transform is inadequate in many cases due to its integrative effects of the discrete spaces. The design of an algorithm to extract optimal parameters of curves passing through image points requires a measure of statistical fitness. A strategy for image feature extraction called Tracking Hough Transform (THT) is presented that combines Extended Kalman Filtering with a Hough voting scheme that incorporates a formal noise model. The minimum mean-squares filtering process leads to high accuracy. Computing cost for real-time applications is addressed by introducing a converging sampling scheme. Extensive performance tests show that the algorithm can achieve faster speed, lower storage requirement and higher accuracy than the Standard Hough Transform.

Keywords: Hough Transform, Parametric curve detection, line detection, Kalman Filtering.

1 Introduction

The Standard Hough transform (SHT) [1] provides a technique in image processing for extracting the parameters of a straight line from its feature points (or edgels) and that involves applying a co-ordinate transformation to the image, such that all the feature points belonging to a contour of a given type in the image space map into a single location in the transformed space. Although this technique has been widely studied [2], conflicts in accuracy, computing cost and memory requirement are still serious. The accuracy achieved depends on the resolution of the parameter space or the number of accumulator cells, but with a corresponding increase in computation cost. A class of solutions has been proposed [2] which employs non uniform or multiple resolutions techniques, based on the observation that it is only necessary to have high accumulator resolution where a high density of votes accumulates. Alternatively, higher accuracy can also be achieved by a curve fitting post-process [3] or interpolation [4]. However, the conventional least square distance method of fitting a line to a set of feature points is unreliable when feature points due to noise and to other edges are present [5-6]. A class of Hough-like techniques appeared in the literature [2] to reduce the computing cost through the redundancy of the range of the parameter or image space , such as the Randomised Hough Transform (RHT) [7] and the subimage processing strategy [8]. The RHT only uses feature points that have a

D. Mery and L. Rueda (Eds.): PSIVT 2007, LNCS 4872, pp. 191–204, 2007.

high probability of forming lines. This can also avoid the use of a conventional quantisation scheme, which greatly influences the accuracy and detection capabilities of the algorithm, as well as the computational and storage requirements. However, in practice, due both to noise in the coordinates of the feature points and to the quantisation of the parameter space, the sampled, quantized line segments do not in general intersect precisely at a common point in the parameter space. Thus, variations of the RHT [9,10] have been proposed such as the curve fitting post-process after the accumulation. In the subimage strategy, as only partial features or subsets of the whole image are used, its performance depends on the quality of the image, the type and the fraction of the image feature points that is selected [2]. Further work has been carried out to derive stopping rules for selection of image feature points [11]. Behrens *et al* [12] report a method to Kalman track a set of features (ellipses) from a resulting from a prior Hough transform stage. Accuracy is determined by the accuracy of the preliminary stage whereas in our case Kalman filtering is used as an integral part of the voting process. Hills *et al* [13] and French *et al* [14] report how objects can be tracked in a sequence of frames using groups of features detected using a Hough Transform, in this case tracking is used to denote correspondence between frames rather than a trajectory in a single image.

In this paper a Hough transform method (based on previous results [15][16], but explained in more detail here) called the Tracking Hough Transform, is presented combining the Extended Kalman filter technique with a converging sampling scheme, which achieves faster speed, lower memory requirement and higher accuracy than the SHT.

2 Standard Hough Transform

A straight line in 2D space can be represented by the equation,

$$\mathbf{f}(\mathbf{Z}_k, \mathbf{a}_i) = x_k \cos\theta_i + y_k \sin\theta_i - \rho_i = 0 \tag{1}$$

where $\mathbf{Z}_k = [x_k \quad y_k]^T \in \Re^2, (k = 1,\ldots,M), \mathbf{a}_i = [\theta_i \quad \rho_i] \in \Re^2$, $(i = 1,\ldots,N)$, $\mathbf{f}$, is a function from $\Re^2 \times \Re^2$ into $\Re^1$, ρ_i is the length of the normal vector from the image origin to the line, θ_i is the orientation of the normal vector and (x_k, y_k) are the coordinates of the image feature points. The SHT uses the idea that the constraint equation above can be viewed as a mutual constraint between image feature points and parameter points. Therefore, it can be interpreted as defining a many-to-one coordinate transformation from the space of image points to the space of possible parameter values and a one-to-many mapping from an image feature point to a set of possible parameter values. The intersection points in the parameter space represent the parameters of the possible lines in the image space. Thus, the objective is to find a subset of significant $\mathbf{a}_i$ from the superset of *all possible* $\mathbf{a}$ in the parameter space, represented by a discrete "accumulator" array whose elements contain a count of the number of image feature points associated with a straight line and whose size depends on the required resolution for the parameter space. The procedure is performed in two stages: an *incrementation* or *voting* stage followed by an exhaustive *search* for maximum counts in the accumulator array. The accumulating function can be expressed formally as

$$H(\mathbf{a}_i)=\sum_{k=1}^{M} I[\mathbf{f}(\mathbf{Z}_k,\mathbf{a}_i)]=\sum_{k=1}^{M} I[\mathbf{f}_i] \quad (i=1,\ldots,N) \tag{2}$$

where M is the total number of image feature points. The *indication* function $I[\mathbf{f}_i]$ is defined by

$$I[\mathbf{f}_i]=\begin{cases}1 & \mathbf{f}_i=0\\ 0 & \mathbf{f}_i\neq 0\end{cases}. \tag{3}$$

The search stage is a peak detection process on the $\mathbf{a}_i$ stored in the accumulator array, where the number of significant peaks is selected using global thresholding, local peak enhancement, etc. In this way, the SHT converts a difficult global detection problem in the image space into a more easily solved local peak detection problem in the parameter space.

3 Tracking by Extended Kalman Filtering

3.1 Models

The Kalman Filter (KF) is a minimum mean squares filtering technique based on a state-space formulation, whose recursive nature makes it appropriate for use in systems without large data storage capabilities. In this paper, the EKF is used for tracking a sequence of positions (coordinates of feature points) in the image space. The special characteristic in this case is that the trajectory of the feature points followed is restricted to be a straight line (or any particular known curve) and the final results from the "tracking" will be the parameters of the trajectory.

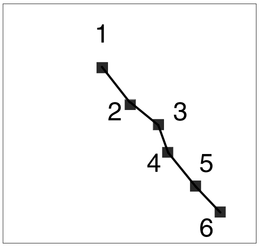

Fig. 1. Viewing line detection as a position tracking process

The KF is based on three probabilistic models: the system or state model, the measurement model and the prior model. The system model considers the process as the result of passing white noise through some system with linear or non-linear dynamics. For straight line detection by the Hough transform, the probabilistic models can be considered as follows:

- The measurement model, which relates the measurement vector $\mathbf{Z}_k$ to the state vector $\mathbf{a}$, is non-linear. Therefore, the Extended Kalman Filter (EKF) is used.
- During a sequence of independent detection of a particular straight line, the state vector representing the line remains constant. Thus, the system model can be thought of as being a deterministic random process that satisfies the differential equation $\dot{\mathbf{a}} = 0$.

The prior model which describes the knowledge about the initial system state $\hat{\mathbf{a}}_0$ is obtained by sampling or digitising the parameter space for the HT accumulation.

3.2 Basic Formulation

Considering a static case only (single image), let the feature points $\mathbf{Z}_k$ associated with an accumulator cell vector $\mathbf{a}_i$ satisfy the general non-linear relationship,

$$\mathbf{f}(\mathbf{Z}_k, \mathbf{a}_i) = 0 \tag{4}$$

where $\mathbf{Z}_k \in \Re^m$ and $\mathbf{a}_i \in \Re^n$ and f is a function from $\Re^m \times \Re^n$ into $\Re^p$. It is assumed that $\mathbf{Z}_k$ and $\mathbf{a}_i$ are independent zero-mean stochastic processes for which only estimate values of $\hat{\mathbf{Z}}_k$ and $\hat{\mathbf{a}}_i$ are available i.e.

$$E\left[\mathbf{Z}_k - \hat{\mathbf{Z}}_k\right] = 0 \tag{5}$$

$$E\left[\left(\mathbf{Z}_k - \hat{\mathbf{Z}}_k\right)\left(\mathbf{Z}_k - \hat{\mathbf{Z}}_k\right)^T\right] = \mathbf{R}_k \tag{6}$$

$$E[\mathbf{a}_i - \hat{\mathbf{a}}_i] = 0 \tag{7}$$

$$E\left[(\mathbf{a}_i - \hat{\mathbf{a}}_i)(\mathbf{a}_i - \hat{\mathbf{a}}_i)^T\right] = \mathbf{P}_i \tag{8}$$

Here, $\mathbf{R}_k$ is the *measurement* covariance matrix (directly related to image space resolution) and $\mathbf{P}_i$ is the *model* error covariance matrix (directly related to parameter space resolution and deviations from an ideal model). Using a first-order Taylor's expansion around $\left(\hat{\mathbf{Z}}_k, \hat{\mathbf{a}}_i\right)$,

$$\mathbf{f}(\mathbf{Z}_k, \mathbf{a}_i) = 0 \approx \mathbf{f}\left(\hat{\mathbf{Z}}_k, \hat{\mathbf{a}}_i\right) + \frac{\partial \mathbf{f}}{\partial \mathbf{Z}}\left(\mathbf{Z}_k - \hat{\mathbf{Z}}_k\right) + \frac{\partial \mathbf{f}}{\partial \mathbf{a}}(\mathbf{a}_i - \hat{\mathbf{a}}_i) \tag{9}$$

which can rewritten as the linear measurement equation:

$$\mathbf{Y} = \mathbf{M}\mathbf{a}_i + \mathbf{U} \tag{10}$$

where

$$\mathbf{Y} = -\mathbf{f}\left(\hat{\mathbf{Z}}_k, \hat{\mathbf{a}}_i\right) + \frac{\partial \mathbf{f}}{\partial \mathbf{a}} \hat{\mathbf{a}}_i \text{, } (p \times 1) \text{modified measurement on the process } \mathbf{a}_i \quad (11)$$

$$\mathbf{M} = \frac{\partial \mathbf{f}}{\partial \mathbf{a}} \text{ , } (p \times n) \text{measurement-state transformation matrix} \quad (12)$$

$$\mathbf{U} = \frac{\partial \mathbf{f}}{\partial \mathbf{Z}}\left(\mathbf{Z}_k - \hat{\mathbf{Z}}_k\right) \text{, } (p \times 1) \text{ modified measurement noise} \quad (13)$$

It can then be shown that $\mathbf{U}$ is a random variable with zero-mean uncorrelated noise, i.e.

$$E[\mathbf{U}] = 0 \quad (14)$$

$$E\left[\mathbf{U}\mathbf{U}^T\right] = \mathbf{W}_k = \frac{\partial \mathbf{f}}{\partial \mathbf{Z}} \mathbf{R}_k \frac{\partial \mathbf{f}}{\partial \mathbf{Z}}^T \quad (15)$$

Therefore, the linear Kalman Filter equations can now be applied directly leading to the recursive EKF algorithm:

$$\mathbf{K}_k = \mathbf{P}_k \mathbf{M}\left(\mathbf{M}\mathbf{P}_k \mathbf{M}^T + \mathbf{W}_k\right)^{-1} \text{ , \quad Kalman "gain"} \quad (16)$$

$$\hat{\mathbf{a}}_{ik} = \hat{\mathbf{a}}_{ik} + \mathbf{K}_k\left(\mathbf{Y}_k - \mathbf{M}\hat{\mathbf{a}}_{ik}\right) \text{, \quad update state} \quad (17)$$

$$\mathbf{P}_k = \left(\mathbf{I} - \mathbf{K}_k \mathbf{M}\right)\mathbf{P}_k \text{, \quad update state covariance matrix} \quad (18)$$

where k is the iteration number. It can be seen that the previously estimated parameter $\hat{\mathbf{a}}_{ik}$ is corrected by an amount proportional to the current error $\left(\mathbf{Y}_k - \mathbf{M}\hat{\mathbf{a}}_{ik}\right)$ called the *innovation*. The proportionality factor (Kalman gain) $\mathbf{K}_k$ minimises the mean-square estimation error [17] (i.e. the terms along the major diagonal of the ***P*** matrix that represent the estimation error variances for the elements of the state vector being estimated).

3.3 Voting Function of the THT

The clustering criterion used to reject outliers when dealing with multiple lines is the Mahalanobis distance (MD) test, defined as

$$\mathbf{d}\left(\hat{\mathbf{Z}}_k, \hat{\mathbf{a}}_i\right) = \mathbf{d}_{ik} = \mathbf{f}\left(\hat{\mathbf{Z}}_k, \hat{\mathbf{a}}_i\right)^T \mathbf{D}_{ik}^{-1} \mathbf{f}\left(\hat{\mathbf{Z}}_k, \hat{\mathbf{a}}_i\right) < \varepsilon \quad (19)$$

where ε is a suitable threshold (normally selected from a χ^2 distribution table), and

$$\mathbf{D}_{ik} = E\left[\mathbf{f}\left(\hat{\mathbf{Z}}_k, \hat{\mathbf{a}}_i\right)\mathbf{f}\left(\hat{\mathbf{Z}}_k, \hat{\mathbf{a}}_i\right)^T\right] = \frac{\partial \mathbf{f}}{\partial \mathbf{Z}} \mathbf{R}_k \frac{\partial \mathbf{f}}{\partial \mathbf{Z}}^T + \frac{\partial \mathbf{f}}{\partial \mathbf{a}} \mathbf{P}_i \frac{\partial \mathbf{f}}{\partial \mathbf{a}}^T \quad (20)$$

Thus, the voting stage of the THT takes place through the computation of the MD accumulating function

$$M(\hat{\mathbf{a}}_i)=\sum_{k=1}^{M} I\left[\mathbf{d}\left(\hat{\mathbf{Z}}_k,\hat{\mathbf{a}}_i\right)\right]=\sum_{k=1}^{M} I[\mathbf{d}_{ik}]\ \left(i=1,\cdots,N\right) \tag{21}$$

$$I[\mathbf{d}_{ik}]=\begin{cases}1 & \mathbf{d}_{ik}\le\varepsilon\\ 0 & \mathbf{d}_{ik}>\varepsilon\end{cases} \tag{22}$$

During the EKF, a feature point is rejected if it does not satisfy the MD test. Otherwise, the point is processed to update the tracked trajectory. The refined value for $\mathbf{a}_i$ and the updated $\mathbf{P}_i$ (eqs. (17) and (18)) are fedback to the MD test to be used by the next image point. At the same time, the vote in the accumulating cell $\mathbf{a}_i$ is incremented by one unit. All the points used by the EKF are then removed if the vote is larger than a threshold. Because of the central role played by the EKF, we have called this approach the “Tracking Hough Transform” (THT).

4 Converging Tracking Strategy

4.1 Converging Sampling of Parameter Space

The SHT is a one-to-many transform where the whole parameter space is sampled, i.e. each image point is mapped to a curve in the Hough space (intersections of such curves indicate the presence of a significant feature). This exhaustive sampling wastes a great deal of computing time, especially for low line-density images, as the voting stage of the SHT usually dominates the execution time. A converging feature tracking strategy is proposed by combining the EKF with a converging sampling strategy. This maps a set of image feature points into a single location in the parameter space and achieves high accuracy. In other words, the scheme combines voting and feature refinement in a *single* stage. The converging sampling mechanism is based on the fact that a pair of feature points defines a straight line and hence a single value $\mathbf{a}_i$ (c.f. Randomised Hough Transform [9]). During the sampling process, two feature points (x_1, y_1) and (x_2, y_2) are selected to obtain initial parameter values (θ,ρ)

$$\theta=\arctan\left(\frac{x_1-x_2}{y_2-y_1}\right) \tag{23}$$

$$\rho=x_i\cos\theta+y_i\sin\theta \qquad\qquad i=1 \text{ or } i=2 \tag{24}$$

Obviously, these two selected image feature points must be different. Then, the EKF process is started from these initial parameter values. The method reduces storage requirement and accumulating time significantly. An overview of the THT algorithm is shown in diagrammatic form in Figure 2.

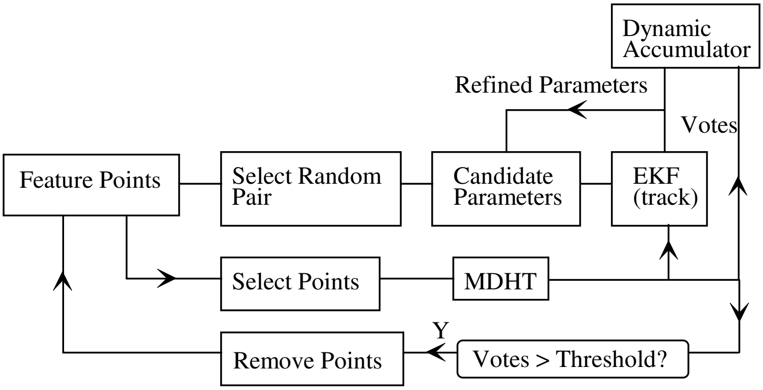

Fig. 2. THT (overview)

4.2 Criteria

Random Selection:
As there is no prior knowledge about the features in an image, any pair of image points is equally likely to form part of a significant feature. Therefore, these points are selected at random.

Stopping Criteria:
Selection and tracking is carried out until the candidate image points are *exhausted*. Two situations are considered to satisfy the *exhausted* condition (Exhausted Iteration Control):

(1) There are no feature points left in the image.
(2) All the image points have been at least selected once.

Using these criteria, even if there no feature points are removed, every feature point can be selected once. However, some feature points might have a chance to be selected again, when there are some feature points removed from the image space after a few selection and tracking processes. Thus, the computing cost of the selection and tracking process can be further reduced if the following criterion is applied (Minimum Iteration Control): "No repetitive selection of any image feature points will be allowed even if they can never be removed from the image", e.g. they are noise.

5 Properties of the THT

5.1 Computing Complexity

One of the effects of removing feature points from the image space is that the subsequent THT process can be regarded as performing on a "subimage" of the original image, but without losing any useful feature points. This is an advantage over Kiryati's subimage strategy. With less feature points in the subimage, the computing cost is also reduced. Combined with the converging sampling process, the THT can reduce the computing complexity even further by *continuously* reducing the number of feature points in the image, subimage, sub-subimage, …, and so on.

5.2 Memory Requirement

In the THT, once the tracking process on a candidate line is finished, the accumulated peak value is compared with a threshold to decide if a line has been detected. This can be regarded as "on-line" peak detection. Thus, there is no need in the THT to use a multi-dimensional accumulator array to register the parameter values for subsequent peak detection process or for further accumulation when the counts are not large enough, such as in the RHT. Instead, only a one-dimensional array is needed to register the counts. In this way, the THT reduces memory requirement significantly from a multi-dimensional accumulator array into a one-dimensional array. Furthermore, converging sampling reduces the range of the parameter space. By using EKF tracking and feature point removal, the THT reduces the range of the parameter space and hence memory requirements even further. The memory required is only proportional to the number of lines (curves) in the image.

For example, suppose that the resolution of the parameter space is high enough for detection, e.g. $(\Delta\rho = 1, \Delta\theta = 180/I_{size})$ where $I_{size} = I_x = I_y$. Typical accumulator sizes required for the SHT are shown in Table 1. These are typically much larger than the expected number of lines in the image. As the dimensionality of the accumulator array in the SHT is proportional to the dimensionality of the curves to be detected, the situation is aggravated for curves other than straight lines (circles, ellipses and so on).

Table 1. Accumulator size required by the SHT

Image Size $(I_x \times I_y)$	64×64	128×128	256×256
Accumulator Size	4K $(= 2^6 \cdot 2^6)$	16K $(= 2^7 \cdot 2^7)$	64K $(= 2^8 \cdot 2^8)$

5.3 Ease of Peak Detection

In the SHT, the peak detection is carried out after the accumulation. As the SHT is a one to-many mapping from the image space to the Hough space, votes from a feature point spread among several accumulator cells. When there are other edges present in the image, the peak detection process sometimes becomes difficult. In the THT, on the other hand, the selection and tracking is a detection-rejection process. When a parameter is obtained, the THT tracks all the feature points on the trajectory. If the number of feature points tracked is high enough, these feature points are removed from the image space; otherwise a new selection and tracking process starts. The process that removes feature points from the image space totally cancels the chances of the contributions to the later detected parameter cell from these feature points. This has an effect similar to the back-projection strategy for peak detection proposed elsewhere [18,19].

5.4 Accuracy

The tracking of feature points by the EKF uses a stochastic model and avoids the need for a quantised parameter space. Therefore, it provides higher accuracy compared to conventional HTs or other algorithms such as the RHT or the subimage strategy.

5.5 Connectivity

The connectivity problem can be directly addressed in the THT during tracking of a line from the starting to the terminating positions, as the positions of the feature points are tracked continuously and recorded in a dynamic array. Thus, after the THT the end-points of lines are automatically obtained [20]. Linking and merging techniques reported in the literature, such as in [21], can be further used to locate the end-points.

6 Tests

The performance of the algorithm presented here has been studied using the HT Test Framework (HTTF) developed by Hare and Sandler [22]. The HTTF generates a large number of images with randomly distributed geometric features (e.g. position and length of straight lines) for gathering statistical data (parameter accuracy, detection and false alarm rates) on the behaviour of a given HT algorithm. This avoids the problem of choosing a representative data set to compare different algorithms which, unless carefully selected, might result in little more than "anecdotal" evidence. In the HTTF, characteristics such as detection rates, false alarm rates, average errors and relative computing cost are used to characterise detection capability, location/accuracy and speed. From a signal detection theory point of view, the performance of a HT algorithm can be considered as a composite of detection capabilities and location accuracy [23]. Thus, alternative HT algorithms can be applied to each random image set and compared in terms of these characteristics. The

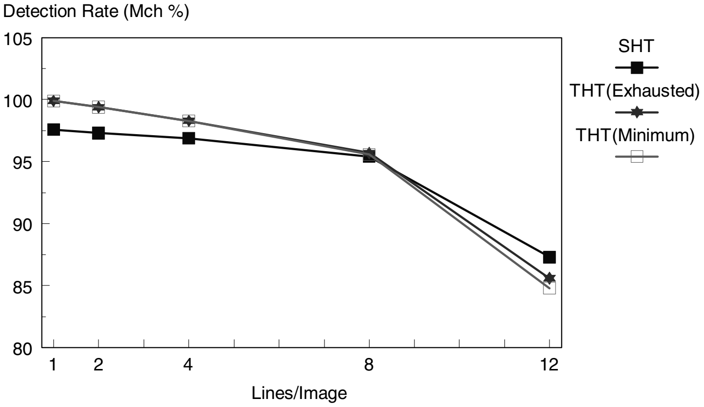

Fig. 3. Detection rates (%) as a function of line density

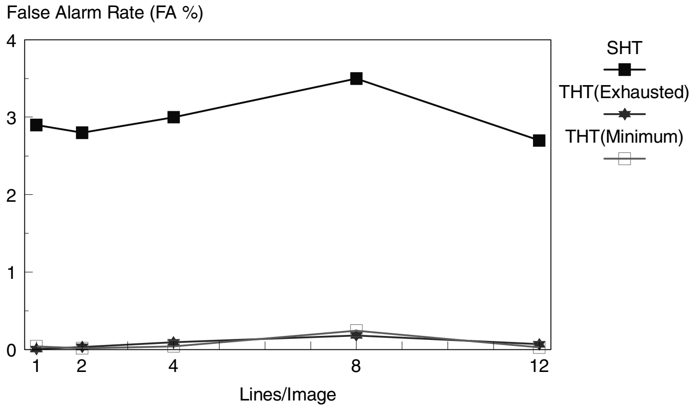

Fig. 4. False alarm rates (%) as a function of line density

SHT is used here as a reference as it is the algorithm with best performance reported in [22]. In these tests $\mathbf{R}_k$ (equation 6) is set to a 2x2 matrix where the diagonal elements are equal to 0.25 (corresponding to a pixel standard deviation σ of 0.5) and the non-diagonal elements are set to zero (i.e. the x and y image components are uncorrelated). Following usual practice the **P** matrix is set initially to have all elements set to a high value (10,000), implying there is no prior knowledge of the process.

Figures 3 to 5 show detection rates, false alarm rates, and relative computing cost for the SHT and the THT (using exhausted and minimum iteration control) respectively, for a randomly generated set of 12000 images (image size 128×128, parameter space resolution is $(\Delta\theta = \pi/128, \Delta\rho = 1)$, threshold set to 30 and 40 votes for low and high line densities respectively [22]). The Relative Computing Cost in Figure 5 is defined as the ratio of CPU time between the SHT and the THT.

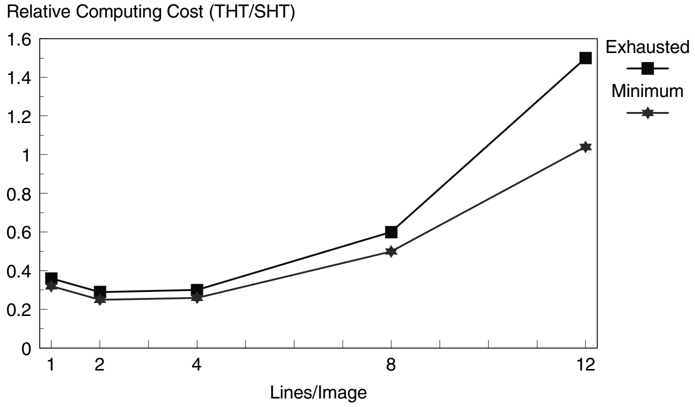

Fig. 5. Relative Computing Cost (THT/SHT)

Table 2. Average parameter errors

Lines/Image	THT (Exhausted)	THT (Minimum)	SHT
1	(0.057, 0.141°)	(0.057, 0.142°)	(0.246, 0.429°)
2	(0.063, 0.152°)	(0.062, 0.152°)	(0.254, 0.444°)
4	(0.067, 0.173°)	(0.071, 0.171°)	(0.256, 0.450°)
8	(0.093, 0.221°)	(0.092, 0.221°)	(0.262, 0.462°)
12	(0.087, 0.196°)	(0.084, 0.189°)	(0.258, 0.448°)

It can be seen that the computing costs in the THT (using either exhausted or minimum iteration control) is less than that of the SHT for detection in low density images. For the high density case, the THT with the minimum iteration control criterion still costs less. It also achieves similar detection performance to the THT that uses exhausted iteration control criterion. It should be noted (Figure 3, Figure 4 and Table 2) that the computing cost saved by the minimum iteration control does not sacrifice performance.

Table 2 shows average parameter errors, between generated and measured features, obtained by the SHT and the THT over the complete test (12000 images). Since the THT provides sub-parameter space accuracy, for direct comparison votes are assigned to discrete cells using tolerance bands of $\Delta\rho$ and $\Delta\theta$.

Detection performance or capabilities can be characterised by the interaction between probability of detection and the probability of false alarm [24]. For a given detection scheme, these two probabilities will depend on the value of the threshold is used. Broadly speaking, higher thresholds improve detection but, at the same time, increase false alarms. For practical applications, therefore, a compromise has to be reached to achieve acceptable levels of detection and false alarms. To visualise the performance of a HT algorithm in this context, we use "Performance Characteristic

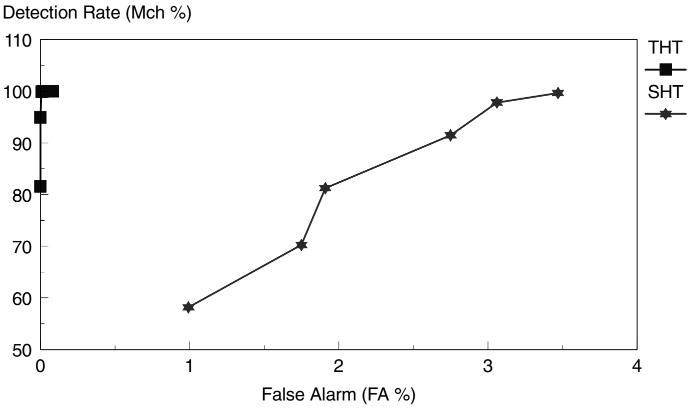

Fig. 6. PCCs of the THT and the SHT (1 line/image)

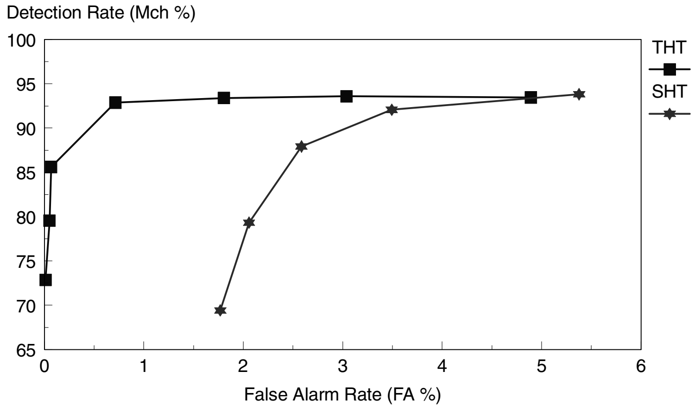

Fig. 7. PCCs of the THT and the SHT (12 lines/image)

Curves" (PCC). The PCC is a plot of the probability of detection versus the probability of false alarm for different thresholds and shows the maximum detection rate that can be achieved for a given maximum false alarm rate. Since the performances of the THT with minimum or exhausted iteration control criterion are similar except in terms of computing cost, only the THT with exhausted iteration control is used here.

In the low density case (1 line/image) the superiority of the THT over the SHT is clear. High detection rates can be achieved for low false alarm rates (e.g. 99.92% detection for 0.008% false alarm). Thus the PCC of the THT is close to the ideal case, essentially on the *y*-axis (Figure 6). In the higher density case (12 lines/image) the PCC of the THT also illustrate the superiority of this algorithm over the SHT (Figure 7).

7 Conclusions

The THT algorithm presented here achieves faster speed, lower memory requirement, and higher accuracy than the SHT. The minimum iteration control strategy achieves even faster speed without sacrificing performance. This has been demonstrated by extensive statistical performance tests. The method is based on a converging sampling scheme which avoids sampling the whole image space, thus saving significant storage and reducing iteration times. Unlike usual post-processing strategies combined with the HT algorithms, high accuracy is obtained by a single-stage combination of voting with feature refinement based on Extended Kalman filtering. A Mahalanobis Distance test is used to reject outliers, so that points which are far from the candidate line do not contribute to voting or refinement. This addresses one of the common practical weaknesses of least-squares methods, as the MD test is dynamically updated by the refinement process. The incorporation of a noise model for image quantisation deals with the so called "errors in the variables" problem. A formal noise model for parameter space quantisation allows control of Hough space "coarseness" (e.g. when establishing a compromise between accumulator array size and accuracy) while maintaining accuracy. Computing cost will increase with line density, in a way similar to the SHT, but the effect can be minimised by dividing the image into several subimages for parallel processing. This algorithm has also been extended to deal with features of higher dimensionality (e.g. circles, ellipses, etc.).

References

1. Duda, R.O., Hart, P.E.: Use of the Hough transform to detect lines and curves in pictures. Communications of the Association of Computing Machinery 15, 11–15 (1972)
2. Leavers, V.F.: Which Hough Transform? Computer Vision, Graphics and Image Processing: Image Understanding 58, 250–264 (1993)
3. Liang, P.: A new and efficient transform for curve detection. Journal of Robotic Systems 8, 841–847 (1991)
4. Niblack, W., Petkovic, D.: On improving the accuracy of the Hough transform. Machine Vision and Applications 3, 87–106 (1990)
5. Weiss, I.: Line fitting in a noisy image. IEEE Transactions on Pattern Analysis and Machine Intelligence 11, 325–329 (1989)
6. Kiryati, N., Bruckstein, A.M.: Antialiasing the Hough Transform. Computer Vision. Graphics and Image Processing: Graphical Models Image Processing 53, 213–222 (1991)
7. Xu, L., Oja, E., Kultanen, P.: A new curve detection method: Randomized Hough Transform (RHT). Pattern Recognition Letters 11, 328–331 (1990)
8. Kiryati, N., Eldar, Y., Bruckstein, A.M.: A probabilistic Hough transform. Pattern Recognition 24, 303–316 (1991)
9. Xu, L., Oja, E.: Randomized Hough transform (RHT): Basic mechanisms, algorithms, and computational complexities. Computer Vision, Graphics and Image Processing: Image Understanding 57, 131–154 (1993)
10. Kalviainen, H., Hirvonen, P., Xu, L., Oja, E.: Comparisons of probabilistic and non-probabilistic Hough transforms. In: Eklundh, J.-O. (ed.) ECCV 1994. LNCS, vol. 801, pp. 351–360. Springer, Heidelberg (1994)
11. Shaked, D., Yaron, O., Kiryati, N.: Deriving stopping rules for the probabilistic Hough transform by Sequential Analysis. Computer Vision and Image Understanding 63, 512–526 (1996)
12. Behrens, T., Rohr, K., Stiehl, H.S.: Using an extended Hough transform combined with a Kalman filter to segment tubular structures in 3D medical images. In: Proceedings of the Vision, Modelling, and Visualization Conference 2001, pp. 491–498
13. Hills, M., Pridmore, T., Mills, S.: Object tracking through a Hough space. In: VIE 2003. Visual Information Engineering Conference, pp. 53–56. IEE, Guildford (2003)
14. French, A., Mills, S., Pridmore, T.: Condensation tracking through a Hough space. In: ICPR 2004. 17th International Conference on Pattern Recognition, vol. 4, pp. 195–198.
15. Xu, C., Velastin, S.A.: A comparison between the standard Hough transform and the Mahalanobis distance Hough transform. In: Eklundh, J.-O. (ed.). LNCS, vol. 800, pp. 95–100 (1994)
16. Xu, C., Velastin, S.A.: The Mahalanobis Distance Hough Transform with Extended Kalman Filter Refinement. IEEE International Symposium on Circuits & Systems 3, 5–8 (1994)
17. Brown, R.G., Hwang, P.Y.C.: Introduction to Random Signal Analysis and Kalman Filtering, 2nd edn. John Wiley and Sons Inc, Chichester (1992)

18. Gerig, G.: Linking image-space and accumulator-space: A new approach for object-recognition. In: Proceedings of IEEE 1st International Conference on Computer Vision, pp. 112–115 (1987)
19. Dambra, C., Serpico, S.B., Vernazza, G.: A new technique for peak detection in the Hough-transform parameter space. In: Proceedings of Signal Processing V: Theories and Applications, pp. 705–708 (1990)
20. Xu, C.: The Mahalanobis Distance Hough Transform with Kalman Filter Refinement. PhD thesis, King's College London, University of London (1995)
21. Weiss, R., Boldt, M.: Geometric grouping applied to straight lines. In: ICPR 1986. IEEE International Conference on Pattern Recognition, pp. 489–495.
22. Hare, A.R., Sandler, M.B.: General test framework for straight-line detection by Hough transforms. In: ISCAS 1993. IEEE International Symposium on Circuits and Systems, pp. 239–242.
23. Hunt, D.J., Nolte, L.W.: Performance of the Hough transform and its relationship to statistical signal detection theory. Computer Vision, Graphics and Image Processing 43, 221–238 (1988)
24. Peterson, W.W., Birdsall, T.G., Fox, W.C.: The theory of signal detectability. IRE Transactions on Information Theory 4, 171–211 (1954)