Robust tracking of circular features

X. Pan, T. J. Ellis & T. A. Clarke

Centre for Digital Image Measurement and Analysis,
Department of Electrical, Electronic & Information
Engineering, City University, Northampton Square,
London EC1V 0HB, U.K. (t.j.ellis@city.ac.uk)

Abstract

This paper considers the problem of internally inspecting underground pipes, such as sewers, in order to identify and locate major structural defects or damage to the pipes, and to identify potential defects which may lead to failure. It considers three aspects of the pipe inspection task; namely, the extraction of geometric primitives from the pipe image sequences, the tracking of these primitives over time, and the detection of (potential) structural collapse of the pipes. Knowledge of the environment (i.e. the reflectivity and geometry characteristics of the pipes) is exploited to develop appropriate methods for extracting the pipe joint features from edge information and to fit a pipe joint model (a circle, or set of connected arc segments) to the grouped features; and to track these features in sequential images. The paper describes a novel implementation of a Hough-based circle detection algorithm and compares its operation with a least-squares curve fitting algorithm for the detection of circular image features. Results of robust tracking of these pipe joints over a large sequence of image frames (some 1500 frames) are presented. Finally, these circle data are employed to aid the detection of partial pipe collapse.

Introduction

There are many thousands of miles of underground pipes in the UK alone, for which there is a requirement for regular surveying and inspection. Such surveying is intended to detect existing and potential failure of the pipe, normally associated with either substantial cracking or collapse of the pipe structure and a range of other defect types [1]. These surveys are performed by pulling or driving a sledge-mounted camera through the pipe. This project is concerned with considering the means of augmenting the existing video and manually recorded data by the off-line analysis of recorded video sequences.

The analysis of image frames from these VHS recorded image sequences is far from trivial. The images are significantly corrupted by high levels of noise stemming from a variety of sources: the quality of the VHS recording, the poor illumination, and the extraneous matter on the walls of the pipe. In addition, because long data sequences are needed to test for the robustness of the algorithms, JPEG encoding of the sequences has

BMVC 1995 doi:10.5244/C.9.55

been used to reduce the storage requirements. This lossy encoding introduces further noise into the image data.

For the present work the analysis is limited to the inspection of clay and concrete pipes. These have a well defined geometry (when undamaged) which can be used to simplify the determination of the camera location within the pipe, and also to aid the location of defects. The pipes (which range in diameter from 150-450 mm) are composed of long straight sections (approximately 1m).

2. Characteristics of the pipe imaging

A dominant and characteristic feature of all of the pipes used in this work arises from the strong reflection from the pipe joints. The projection of the joints are approximately circular (though usually incomplete circles) depending on the condition of the pipes and the level of water in the pipe. Hence, determining the shape of pipe joints can be an important first step for defect detection. Edges detrected for a typical joint should have a common radius and centre. For a defective or broken joint, where the segments cannot be fitted to a single circle, they can be interpreted as a small cluster (2-4) of similar radii circles, but with displaced centres. One of the problems discussed in this paper is how the distorted pipes can be identified.

The following observations can be made regarding the imaging environment: the light reflected back from the pipe joint diminishes for features further from the camera; a maximum of 4 joints will be visible in each image, depending on the level of illumination, the reflectivity of the walls, the orientation of the camera, and the rate of change of direction of the pipe; the pipe joints further from the camera are constrained to be enclosed within and approximately concentric with the joints nearer to the camera; because the sledge is not tightly constrained to point down the axis of the pipe, the intersection angle between the optical axis of the camera and the pipe joint cross-section normal can vary by up to 20°.

The main problems of the inspection come from two sources: the wide variation in the appearance of pipe joints and the generally poor visibility within the pipelines. These contribute to the image noise and hence to the reliability of feature extraction. Although the pipe joints are prominent features in the pipes, they are not ideal image features to detect, appearing thick and irregular in the image. For a given joint, the width of different parts of the joint can vary by up to 50% under normal conditions, and up to 400% if the joint is broken. The reflected intensity of different parts of a pipe joint may also vary substantially, with some parts reflecting strongly whilst other sections are less visible. In addition, the build up of extraneous matter on the pipe walls can obscure the joint or generate erroneous edges.

3. Circle Detection

A variety of circle detection algorithms exist in the literature [3-8]. In this research, two popular methods have been investigated to detect arcs from the edge detected image

data: the Hough transform (HT) and least-squares circle fitting. The current work compares the 'standard' HT with a computationally more efficient method which restricts the voting region in the Hough accumulator space by first grouping edgels into connected arc segments (voting in groups), and then determining the (restricted) range of parameter values (radius, x,y centres) that may be associated with each group. These extensions to the HT algorithm are discussed in [9] in more detail.

Alternatively, given a group of edgels, circle fitting can be performed using a least-squares method. The problem is to choose which pixels should be included in the calculation to generate an accurate result, since outliers can significantly effect the least-squares estimate. Using an hypothesise-and-test strategy, given a few features, an initial estimate of the circle location and size can be made. The hypothesis is then tested by locating other features detected in close proximity to the circle boundary. To reduce the combinatorics of the search, each group generates a separate circle hypothesis, and circles with similar parameters are merged to give the most consistent result (see [9]).

4. Grouping edgels into arcs using gradient directions

Following the application of a standard edge detector (e.g. Canny), the pipe joints are detected as arcs in sewer images. A joint may generate one or several arcs, depending on the thickness of the joint. It is known that neighbouring edgels on an arc will have a similar gradient Hence, if the edgels are grouped by their gradient direction, it is more likely that the edgels in a group will belong to an arc.

A threshold of direction differences is used to decide which edges should be included in the group. Ideally, it should allow the tracking to continue along an arc and disallow other edges. The change of direction along an arc depends on the corresponding radius. The length (L) of an arc on a circle is $L = \theta R$, where R is the circle radius and θ is the angle subtended by the arc from the circle centre (in radians). Given a minimum arc length of 1, the minimum angle θ_0 correpsonding to the smallest circle which can be detected, has a radius of 10, then θ_0 is approximately 6°. A threshold of 10° has been used in this work, which corresponds to a smallest detected radius of R=12, and only neighbouring pixels with direction differences less than the threshold will be grouped. In the experiments, this threshold prevented the edgels from different arcs from being included in the current group. Because of noise in the edge directions, the thresholding is likely to break an arc into different groups, though this does not have very serious effect on the final result.

5. Hough Transform using pixel groups

Having grouped edgels in the image into groups, the next stage is to determine valid circle fits. If the grouped edgels originate from a circular arc in the image, then the edgel data can be used to restrict the voting in the Hough accumulator space. In the following analysis, the possible error in the predicted circle parameters (centre and radius) which are estimated over the group, are used to define the region in Hough space where the voting can occur. Each pixel group is assumed to be an arc (or line) and is represented

by three points: two endpoints and the centre point. The possible circles that pass through an arc with a given error can be determined by the following.

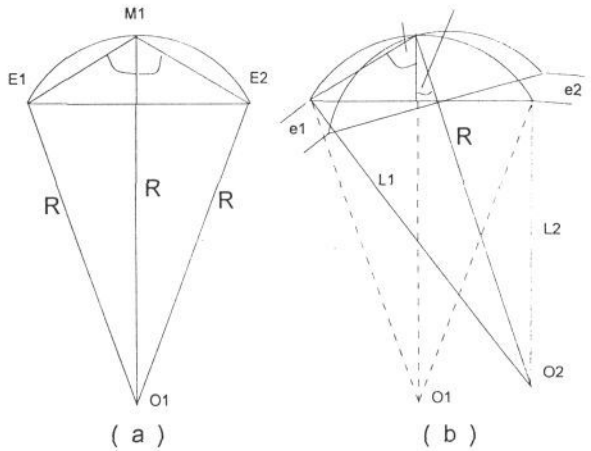

Figure 2. The parameters associated with the circle error estimation.

The largest difference between the original and the rotated arc is at the two endpoints. Considering E_1, we have (from [9]):

$$R \geq (D^2 - e_i^2)/2(D\,Cos(\theta + \phi) + e_1) = k_1 \qquad \text{Eq. 1}$$

And at endpoint E_2

$$R \leq (D^2 - e_2^2)/2(D\,Cos(\theta - \phi) - e_2) = k_2 \qquad \text{Eq. 2}$$

Hence, along the line O_2-M_1 for an error e ($e = e_1 = e_2$), the circles that pass the point M_1 should have a radius that satisfies Eq. 1 and Eq. 2. In the other words, the centre O_2 should locate on the line within a distance $k_1 < R < k_2$. It is important to note that the larger the angle ϕ, the closer the parameters k_1 and k_2. When ϕ is too large, neither k_1 or k_2 may exist (if R is to remain positive) when:

$$-D\,Cos(\theta + \phi) > e \text{ or } D\,Cos(\theta - \phi) < e$$

Additionally, when $e = 0$ and $\theta = 0$, then there is no error, and the circle radius R can be calculated directly from: $k_1 = k_2$ and $R = D/2\cos(\phi)$

For small edgel groups, the arc is small, e.g. $D < e$ and $R < 0$, and θ may range from 0 to 360°. In this case, the pixel group is discarded. An example of the voting mechanism based on Eq. 1 and 2 is shown in Fig. 3. Since the swing angle θ cannot be very large for a given error e, an arc only votes for a small strip in the (x,y) space. Hence, the Hough space is reduced. Further reductions in the size of the accumulator array can be achieved by reducing the resolution of the circle radius, and limiting the largest circle to the linear

dimensions of the image (i.e. not across the diagonal), since the joint is starting to disappear from view.

In practice, since only a small number of circles are expected in the image, it is more efficient to represent the radius dimension with a linked list. When an (x,y) location receives a vote with a radius value of R, the entry in the list at (x, y) is checked. If there is a record with this value of R, its vote is increased by one, otherwise a new record is set up. A maximum limit of N_r is set for each (x, y) location and if there are already N_r records at location (x, y), the vote is discarded. The problem with this arrangement is that the information from later votes may be lost. However, if pixel groups are arranged according to the number of edgels in the group prior to Hough voting, and the largest groups vote first, then the later votes will come from small pixel groups and hence have a reduced effect.

The voting procedure is as follows. For each pixel grouping, the parameters D, θ and direction O_1-M_1 are known. The voting is along the direction ϕ plus a swing angle (θ). For each ($\phi+\theta$), there is a possible radius range of k_1 - k_2 (as given by Eq. 1 and Eq. 2) with error e. Hence, the voting region is fan-shaped and bounded by (k_1, k_2, $\phi+\theta$, $\phi-\theta$). For each location (x, y), the **r** records are searched. If, within a threshold, there is a similar record, a weighted vote is added. Otherwise the new record is initialised with the weighted vote.

The weighting is in the form:

$$V = N - \theta - \text{error}$$

where N is the number of edgels in a group and θ is the swing angle. Hence, the voting gives more weight to possible circles in the middle of the area in the accumulator array bounded by (k_1, k_2, $\phi+\theta$, $\phi-\theta$) than for groups near the edges, and is in proportion to the number of edgels in the group. The pixels groups from a real image and their votes are shown in figure 3. The strips in figure 3b) indicate the voting 'paths' of the groups, and the two best circles are plotted into the image in 3b).

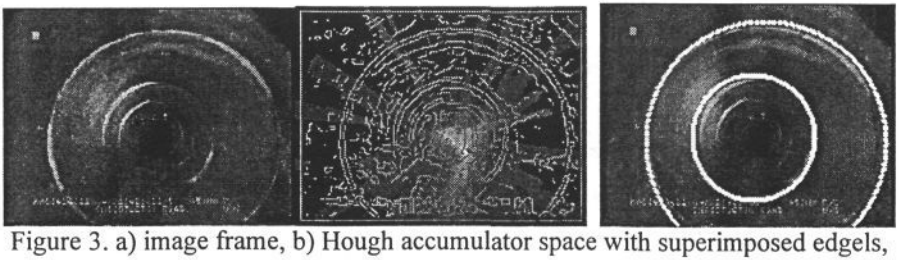

Figure 3. a) image frame, b) Hough accumulator space with superimposed edgels, c) image with superimposed circles.

The peak in the Hough space is determined by a back voting procedure [9]. Each group searches the same locations (x, y, [r]) as with the initial vote accumulation, and finds the location with the largest vote. All the other locations will be then decreased by a

weighted vote, V. The peaks will be at the intersection of voting strips of the different pixel groups and each peak will correspond to at least one pixel group.

6. Detecting arcs using the least squares method

The main advantage of using the HT for arc detection is that it is a global detection method which finds un-connected arc segments which originate from the same circular feature. Its main disadvantage is that it is computationally expensive and does not work reliably with only small arc sections.

The least-squares (LSE) method can estimate the circle parameters for an arc segment from a group of pixels or a combination of pixel groups. It is assumed that an arc may be composed of different pixel groups, and the problem is how to choose the best combinations of groups, and to exclude outliers. Ideally, all the possible combinations should be examined, which is computationally expensive. Fortunately, unlike ellipse detection [2] where the LSE method tends to under-estimate the size of the ellipse, circle detection is less prone to such errors. Using this characteristics, arcs can be detected from N pixel groups by a recursive method.

The least-squares formulation is described in [9], and uses two two thresholds: e_1 is the initial acceptance threshold and e_2 the group membership threshold and $e_1 > e_2$. In experiments, e_1=15 and e_2 = 10 have been found to give satisfactory results. The circles extracted from real images are shown in the next section.

7. Experiments with the HT, modified HT and LSE method.

Following results reported in an earlier paper [8], which investigated the effects of the JPEG encoding on spatial resolution, the current work adopts a reduced resolution image (192 x 144), sub-sampled from the original recording sampled at 768 x 576 pixels. Naturally, this substantially reduces the computation time, but has only a negligible effect on the detectability of the circles.

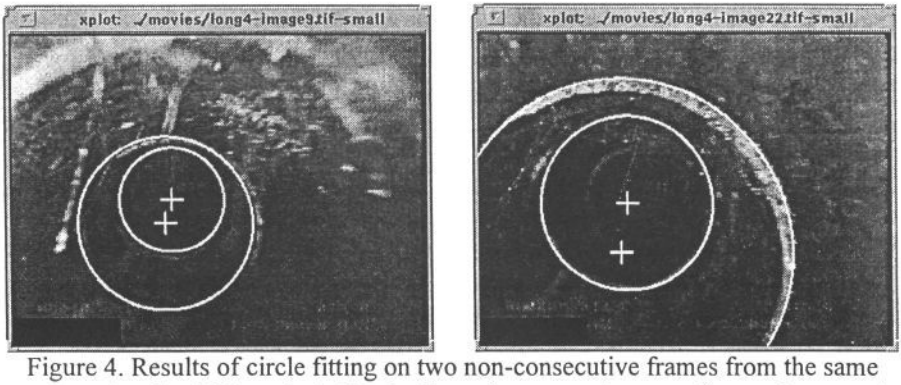

Figure 4. Results of circle fitting on two non-consecutive frames from the same sequence using the LSE method. Circle circumference and centres (crosses) are marked.

Three methods, the standard HT, the HT based on the pixel grouping and LSE method were applied to three image sequences composed of 596, 507 and 395 image frames. Circles are detected independently in each image frame and no information is passed between frames. The results are then compared against a set of manually measured data. Since the number of circles detected in each image may be different, only the two best matched circles from each method are used in the comparison.

Two examples of the circle detection algorithm are shown in figure 4, detected using the LSE method. The accuracy and speed of computation of the three detection algorithms are compared in figure 5, where the plots indicate the measurement error between the manual measurements and the detection algorithms, measured for more than 800 circles. (Note: The abscissa of the graphs 5a)-5c) show the percentage error in the radius of the estimated circles - 100% indicates a factor-of-two difference between manual and detected circles). As can be seen, the grouped HT has few errors greater than 50%. Figure 5d) shows the relative computational speed of the three methods, indicating that the LSE is fastest, followed by the grouped HT and then the standard HT.

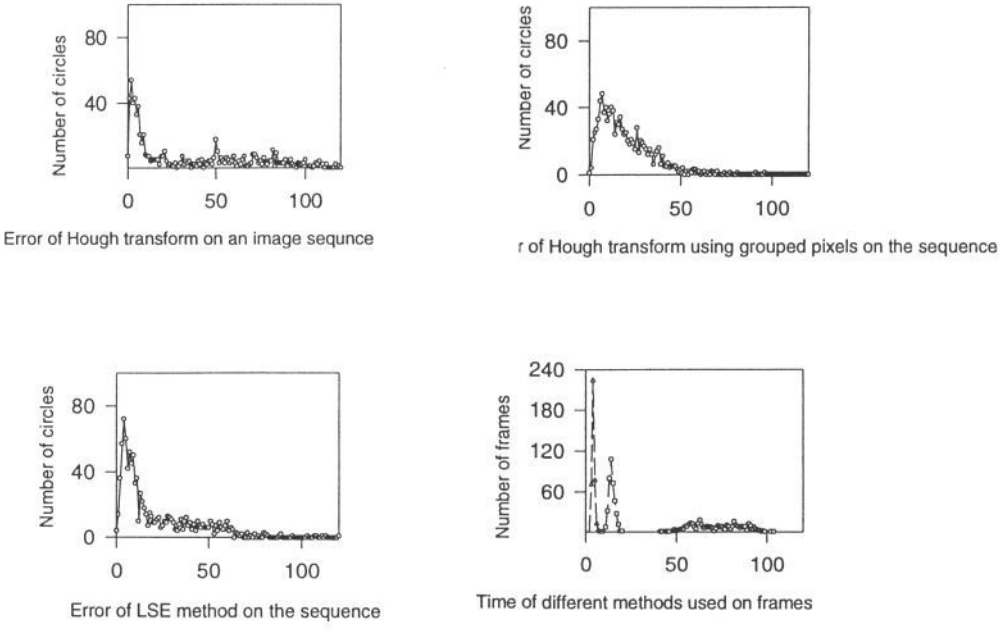

Figure 5. a)-c) Illustrates the error in the size estimate in comparing the manually extracted circles with those extracted by the three methods - the HT using grouped edgels, the standard HT, and LSE circle-fitting. (d) shows the comparative processing times for the three methods (key: triangle - LSE, square - grouped HT, circle - standard HT).

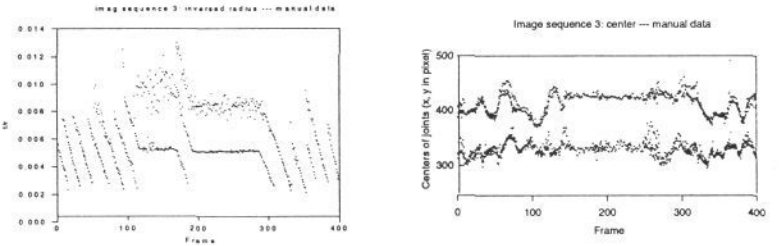

Figure 6. Illustrates the results of extracting circle data manually from one sequence.

Figure 6 presents the results of manually measuring the visually detectable circles in an image sequence (usually 2 or 3 per image) using a software tool which allows the user to point to three points in the image and fit a circle. These points can then be edited to ensure a good fit. The left graph shows the radii (1/r, in pixels) of measured circles over 400 image frames. Overall, the camera was moving with a roughly uniform speed, though it was stationary between frames 150-240, as can be seen from the horizontal trace, where the circles size remains static. The right graph shows the variation in the position of the circle centres over the same frames. As can seen, although this varies by 50-60 pixels over the sequence, the variation is not significantly correlated with the radius.

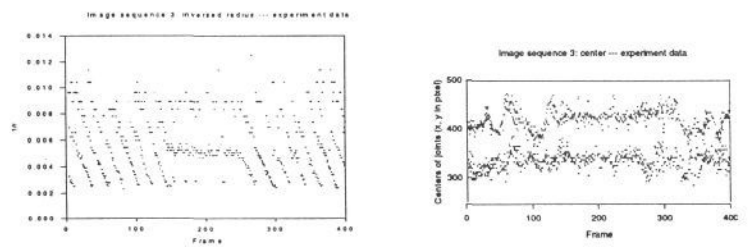

Figure 7. Automatically extracted data corresponding to the manually extracted data in fig. 6.

Figure 7 shows similar results generated by the LSE circle detection algorithm. For this data, circles generated by pipe joints were crudely selected by identifying only concentric circles in the measurement data. As can be seen, the results are strongly comparable. The seemingly lower density of points, in comparison with figure 6, is largely attributable to the coarse re-sampling of the image (from 768 x 576 to 192 x 144) applied to the images before circle detection. The results on the manual data (figure 5) were performed at the original image resolution. The total number of circles detected over the two data sets were 1140 and 816 from the manual and automatically detected circles respectively. The straight lines that could be fitted to the data in figure 6a) and 7a) indicates that the sledge is being pulled at a uniform velocity, except where it comes to rest in the middle of the sequence.

A simple algorithm has been applied to the circle data in order to track the pipe joints over a long frame sequence. The algorithm matches circle data over two consecutive frames using the centre and radius parameters. These parameters are allowed to vary by only 20% of the circle radius in order to ensure that tracking is maintained. The results of the tracking over one of the image sequences (1140 circles in 507 frames) is summarised in table 1. This shows the percentage of correctly tracked joints between the manually extracted circle data, and the LSE circles shown in figure 6 and 7. The labels relate to the sequences that have been extracted from the manually measured circles. Even though the tracking is based on rather simple criteria, the results indicate that it will be quite feasible to obtain a track% substantially closer to 100% by improving the predictive capability of the algorithm, based on the results shown in figure 8.

label	1	5	6	9	13	16	17	22	29	32	33	35	36	38	47
manual	36	47	56	44	45	24	121	51	22	39	44	46	59	64	23
auto	35	40	50	43	40	15	111	41	20	35	38	36	44	42	17
track%	97	85	89	98	89	63	92	80	91	90	86	78	75	66	74

Table 1. Summary of the result of tracking automatically labelled circles in one of the image sequences, in comparison with the manually labelled data.

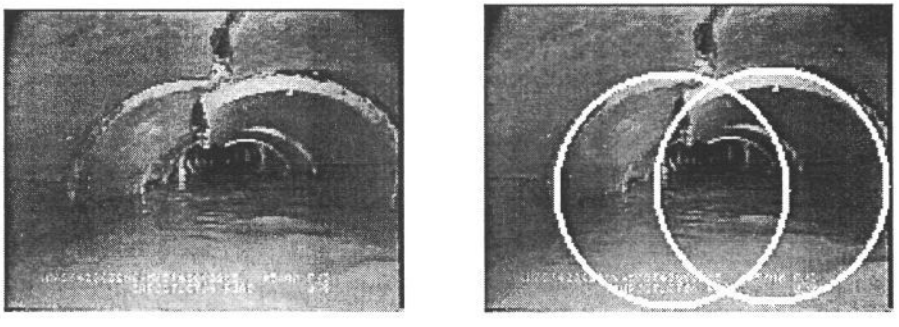

Figure 9. Shows the detection of pipe cracking and collapse. The overlapping circles of equal size indicate the presence of this type of defect.

Finally, figure 9 shows the preliminary results of the pipe collapse detection. The concrete pipes retain their circular shape when cracked, and the capability of the LSE algorithm to reliably detect the pipe fragments leads to a practical method for detecting these types of pipe defect.

6. Conclusions

In this paper, the detection of circular segments from noisy image has been discussed. Two new methods are introduced. A Hough Transform algorithm which employs voting with pixel grouping and a least squares method, also based on pixel grouping. The main advantage of pixel grouping is that the computation is considerably simplified, and hence faster.

The methods have been extensively tested on a number of long image sequences, and the performance of the methods evaluated. Although the modified HT and LSE give roughly comparable errors in determining the circle parameters over the sequence data, the LSE scores by being computationally more efficient, and also more robust in detecting the collapsing pipes.

To further improve the tracking performance, a predictive (Kalman) filter will be used to track the joint sections. The next stage in the defect detection will combine crack detection with the cues generated from the collapsed pipes in order to validate the defective pipe section. Other defect detection schemes, which will cover a fuller range of the common defect types [1] are also in development.

Acknowledgements

This research was supported by the Engineering and Physical Science Research Council (EPSRC) as part of the Information Technology in Engineering (ITE) initiative under grant number GR/H41898.

References

1. Water Research Centre and WAA Sewers and Water Mains Committee, "Manual of Sewer Condition Classification," WAC Engineering, England. 1988.

2. T. Ellis, A. Abbood and B. Brillault, "Ellipse Detection and Matching with uncertainty," BMVC 1990,, Glasgow, U. K, September, 1990. pp 136-143

3. U. Landau, "Estimation of a circular arc centre and its radius", Computer Vision, Graphics and Image Processing. Vol 38, pp 317-326, 1987.

4. E. R. Davies, "Circularity - a new principle underlying the design of accurate edge orientation operators," Image and Vision Computing, Vol 2, No. 3, pp 134 - 142. August 1984

5. J. Illingworth and J. Kittler, "The adaptive Hough transform,", IEEE trans. on PAMI, Vol PAMI-9, No. 5, September 1987. pp 690-698.

6. H. K.Yuen, J. Princen, J. Illingworth and J. Kittler, "Comparative study of Hough Transform methods for circle finding", Image and Vision Computing, vol. 8, no. 1, pp 71-77. February 1990.

7. G. Gerig and F. Klein, "Fast contour identification through efficient Hough Transform and simplified interpretation strategy", 8th international Joint Conference on Pattern Recognition, Paris, France (1986), pp 498-500.

8. Pan, X, Clarke, TA, Ellis, TJ, "The detection and tracking of pipe joints in noisy images", In: Videometrics III, SPIE vol. 2350, Boston 1994, pp 136-147.

9. Pan, X, Ellis, TJ, Clarke, TA, "Pipe Joint Tracking", Tech Report MVG95/1, Electrical, Electronic & Inf. Eng., City Univ., London EC1V OHB, 1995.