The detection and tracking of pipe joints in noisy images

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

Centre for Digital Image Measurement and Analysis, School of Engineering,
City University, Northampton Square, London. EC1V 0HB. U. K.
Email: X.Pan@city.ac.uk

ABSTRACT

The remote and automatic inspection of the inside of pipes and tunnels is an important industrial application area. The main characteristics of the environment found in commonly used pipes such as sewers are: limitations on the camera spatial position; a large variety of surface features; a wide range of surface reflectivity due to the orientation of parts of the pipe, e.g. the joints; and many disturbances to the environment due for example to: mist; water spray; or hanging debris. The objective of this research is defect detection and classification; however, a first stage is the construction of a model of the pipe structure by pipe joint tracking. This paper describes work to exploit the knowledge of the environment to: build a model of defects, reflectivity characteristics and pipe characteristics; develop appropriate methods for grouping the pipe joint features within each image from edge information; fit a pipe joint model (a circle, or connected arcs) to the grouped features; and to track these features in sequential images. Each stage in these processes has been analysed to optimise the performance in terms of reliability and speed of operation. The methods that have been developed are described and results of robust pipe joint tracking over a large sequence of images are presented. The paper also presents results of experiments of applying several common edge detectors to images which have been corrupted by JPEG encoding and spatial sub-sampling. The subsequent robustness of a Hough based method for the detection of circular images features is also reported.

Keywords: pipe joints, tracking, Hough transform, edge detection, JPEG.

1. INTRODUCTION

There has been considerable research into the inspection of industrial objects using computer vision techniques[1,2]. An increasing area of interest occurs when either the camera or the objects are moving. In this case, two different procedures may be involved in the inspection process: (i) tracking to determine the position of the camera with respect to the subject of interest and (ii) feature detection. In this paper, the inspection of pipe joints is discussed. The current method of pipe inspection is by manually controlling a remote TV camera and classifying defects from images displayed on a monitor. This is both time consuming and costly. The inspection is also stored on video tape for subsequent archiving and analysis. These video tapes are used in this research to formulate strategies and develop software for analysis which may, in the future, be used in the field. The ultimate objective of the research is to provide an objective measurement of pipe defects which matches or exceeds the performance of a human operator. What makes this research particularly challenging is the high level of noise encountered. The origins of this noise are many. For example: camera instability; poor illumination; occlusion; gross distortions in the pipe; the build up of extraneous matter on the walls and joints of the pipe; and the environment that the images are acquired in. As a first step in the inspection process the pipe joints, which are readily visible, are identified. The subsequent extraction of information about deformation, the build up of extraneous matter, or a large number of other features will be overlaid on the basic pipe model that will be constructed from the pipe joints. Alternative methods of inspection by more direct means have been suggested[3] and may be used in conjunction with the proposed method in the future. One aspect of the overall project is to investigate efficient methods of minimising the amount of data typically stored during a sewer inspection. Many thousands of video tapes are used to store the images, and JPEG image compression has been employed to significantly reduce the storage requirement for digital images.

Real time defect detection depends on reliable tracking of features from a sequence of images. In this paper, edges are used as the main features since they often represent physical discontinuities. With well defined object features, the edge can be used directly for detection and tracking. However, because of the large amount of noise in the images, the detected edges are often short broken segments which makes extraction difficult. In this application, there are no reliable lines or corners which can be used so tracking cannot be based on such localised features. In this paper the characteristics of the pipe images, the process of feature extraction, and the tracking of pipe joints are considered. Results of testing the algorithms on a range of images are presented.

2. THE CHARACTERISTICS OF THE PIPE IMAGES

The pipe images were acquired using a remotely controlled camera and video tape storage. The illumination is provided by a number of small lights surrounding the lens which project forwards. This has the tendency to accentuate some features, while reducing the visibility of others. Furthermore, the images contain a considerable range of additional features which complicates the process of feature detection and tracking. The reliable extraction of features in the images is facilitated by using *a priori* knowledge of the camera/pipe configuration to model the expected characteristics of the images so that algorithms can be correctly designed. Six example pipe images are shown in Figures 1 - 6.

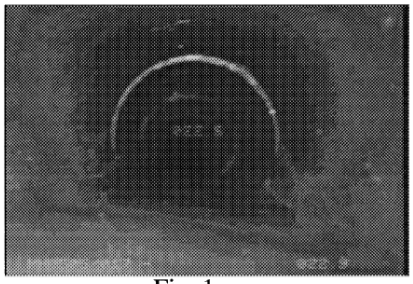
Fig. 1

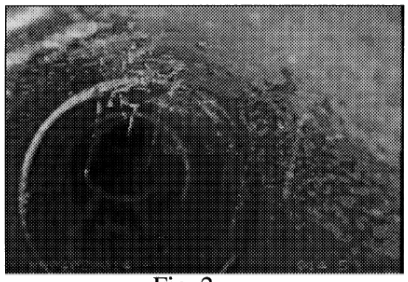
Fig. 2

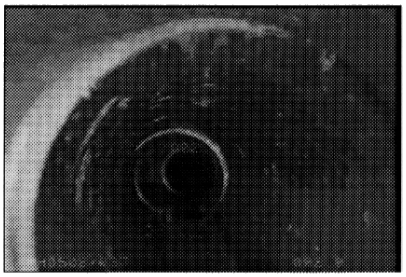
Fig. 3

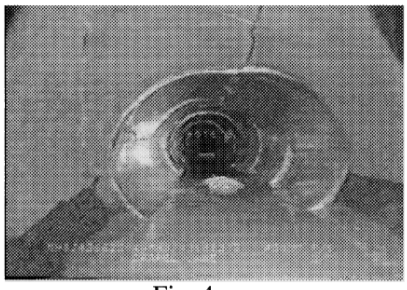
Fig. 4

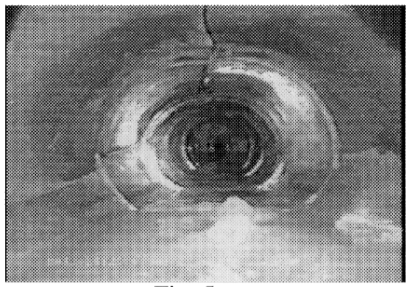
Fig. 5

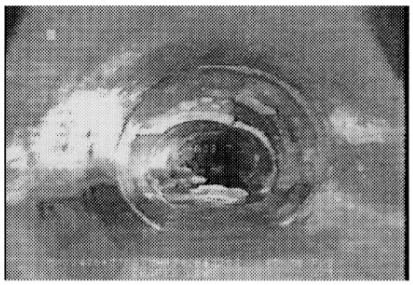
Fig. 6

Figures 1-6. Examples of pipe images.

2.1 The reflectivity model.

The structures which are the subject of this research are constructed from short pipe sections, typically of 200 - 300 mm. in diameter and of approximately 1 meter in length. The camera views the pipe joints obliquely and because the pipes do not join exactly, the ends of the pipe present a face that is almost perpendicular to the camera lens axis. The lighting that is generally used is projected along the length of the pipe. Hence: (i) the pipe joint will appear very bright in comparison to the sides of the pipe because of the preferential angle presented to the light source (Figures 1, 2, & 3.); (ii) the light reflected from the pipe side wall will gradually diminish as the distance between the camera / light source is increased (assuming the characteristics of the pipe surface are approximately constant); (iii) variations in the camera / light source orientation will cause further complications to the model (Fig. 3); (iv) the geometry of the pipe will vary considerably, especially when defects are present (Figures 4, 5, & 6.); (v) successively darker areas will generally be found inside the pipe joints until the joints cannot be distinguished or disappear from view. A knowledge of these characteristics will be required for the production of efficient and robust algorithms.

2.2 The camera and pipe model.

The pipe joint visibility is limited by the lighting used by the monitoring system and generally a maximum of three to four pipes joints can be detected by the sensor. The spatial resolution of the sensor is not usually a problem, as the third or fourth joints are still reasonably large size. Hence, the following assumptions can be made :

(i) the pipe joints are generally circular and have the same radius;

(ii) the normals of the plane of the joints will either be parallel or will only gradually change in direction;

(iii) the intersection angle between the optical axis of the camera and the pipe joint plane normal will be small, e.g. less than 20 degrees; and

(iv) at least one joint should be visible in each image.

These assumptions suggest that the pipe joint will usually approximate to a circle in the images. The camera will generally move in a direction close to the axis of the pipe and will also always be inside the pipe; hence, an image of a joint which is further away will always be within the boundary of the previous pipe joint image or occluded by it. Since the shape of joints is assumed to be circular, tracking the joints is then a problem of tracking the movements of circles.

2.4 The defect model.

The typical sewer can be categorised into eight main classes[4]. The most serious defects, such as gross deformations of the pipe and full, or partial, collapse, usually result in large changes in the shape of the original pipe which will be particularly obvious at the joints, as shown in Figures 4-6. Hence, the pipe model used will be only approximately correct in these cases. The classification of these defects is beyond the scope of this paper, but will form the subject of future investigations. The edges detected using Canny detector in images shown in Figures 1-6 are shown in Figures 7-12. These edges will be used in the subsequent tracking of pipe joints.

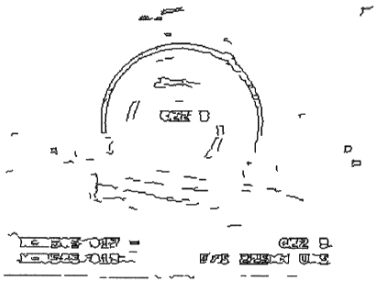

Fig. 7

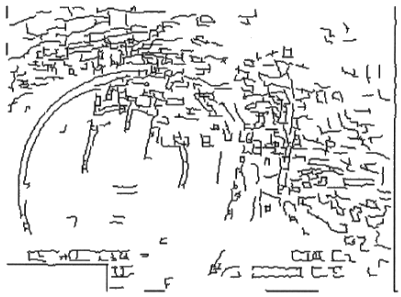

Fig. 8

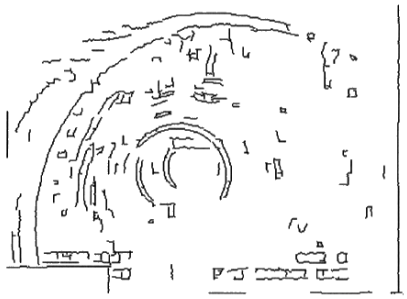

Fig. 9

Fig. 10

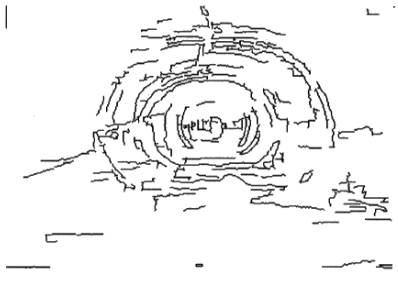

Fig. 11

Fig. 12

Figures 7 - 12 Edges detected from the images in Figures 1 - 6.

3. THE ANALYSIS OF SUITABLE EDGE DETECTION METHODS

3.1 Introduction.

In contrast to many vision tasks in man-made environments, in the images used here, curved lines dominate. Under these circumstances conic curves are reasonable approximations to pipe joints in the edge data. Circles occur often in real world scenes so the the extraction of conic curves from edges has been investigated by many researchers and various methods have been suggested, such as conventional and recursive least squares fitting[5-8] and the Hough transform[9]. In this application the circle detection method has to deal with considerable levels of noise, as shown Figures 1 - 6. As a result, the edges which may result from a particular view of a pipe joint may be broken into small segments, and many extraneous edges may survive the thresholding. For a small segment, many circles of varying radii and locations may fit the edge data, all of which may match the segments reasonably well. A major problem in work of this kind is fitting the correct curve to the segments which originate from the feature of interest. The Hough transform is particularly robust in the presence of image noise for these types of feature extraction tasks, and has been selected for use in this research. However, the high levels of noise that must be dealt with requires careful consideration in implementation of the algorithm.

The first step in the process of considering the Hough Transform is the analysis of the factors that degrade the performance. JPEG (Joint Photographic Experts Group) compressed images have been used in this project because sequences of pipe images can be taken from video tape and efficiently stored for subsequent analysis. JPEG is a commonly used form of lossy compression that uses several techniques to compress images with a compression ratio that is dependent on the level of detail in the image. One of the prime benefits of the method is its implementation in numerous hardware boards such as the Parallax Xvideo[10]. This hardware has been used to grab long near-real-time image sequences from a video tape recorder. However, before using the compressed images it was considered necessary to analyse the effect of the compression on the lossy images.

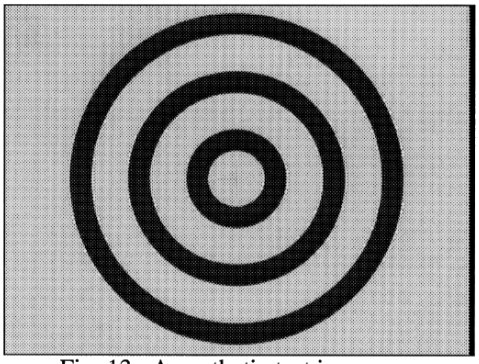

Fig. 13. A synthetic test image.

3.2 An analysis of the effect of using JPEG compression.

The detection of circular features is the prime objective in this work, so a test image was constructed to simulate various conditions that it was expected to encounter in the real world images (Fig. 13). The test image was made using a drawing package and printed out at 600 dpi. The three black circles were arranged to be concentric with a cross placed in the centre. The images were collected using the Parallax frame grabber with two focus settings, three different ambient light levels, and four compression ratios from zero (no compression) to a Q factor of 300 (maximum compression). An analysis of the compression ratios resulting from the twenty colour sewer images (size: 768x576) revealed that a range of compression from 22:1 of the original image size for a Q factor of 30 to 120:1 for a Q factor of 300 were obtained. Figure 14 shows a sectional profile through a uncompressed and compressed image (Q factor 300) illustrating the blockiness resulting from the JPEG compression. Under these circumstances it would be expected that presence of such noise would affect the edge detection

process. Because accurate edge direction information is required in the Hough Transform it was considered necessary to conduct tests to analyse the effect.

Fig. 14 The profiles of cross-section of circles in the test image in Fig. 13.
(a) Before compression. (b) After compression and decompression (Q factor 300).

3.3 Experiments to analyse the effect of JPEG compression on edge direction.

3.3.1 Analysis on test images.

The test images discussed earlier were used to assess the effect of JPEG compression on the computed edge direction. The centre of the circles was known and so this point was used to measure the edge direction error (in degrees) for all of the six rings from the edge detection process. In the majority of cases no significant differences were observed at the differing radii so the average of the standard deviation calculated for each of the rings was used to assess the compression effects. Two edge detectors were selected from the filters analysed by Lyvers[11] and by Davies[12]: the Canny[13] and the Sobel[14]. These detectors were chosen as they were expected to give good results without being too computationally expensive. It is always difficult to compare filters as they use differing thresholds and the results are often viewed qualitatively; however, in this case each filter was tuned to give the minimum R.m.s error in edge direction, which is used as a reasonable judge of performance. In addition, each of the detectors was also tested on sub-sampled images of 2:1, 4:1, and 8:1. The overall results for these tests are now summarised.

Sobel: For the Sobel edge detector under all conditions of spatial sampling, compression, and image quality, there was a reduction in edge direction error with larger filter size (5x5, 7x7, 9x9, & 11x11 were used). The effect of sub-sampling the image gave better results for 2:1 and 4:1 than for the 1:1 and 8:1 images. Blurred edges gave worse results compared to the sharp edges except in the case of sub-sampled edges when the blurred edges were sharpened by the sub-sampling process.

Canny: For the Canny edge detector in 1:1 images there was an improvement in edge detection accuracy with an increase in σ (a measure of the size of the filter). However, this was not carried through into the sub-sampled images where the performance worsened for a larger σ. The optimum σ value was found to be 2.0 for most images. The sub-sampled images provided better angle direction for the 2:1 and 4:1 sampled images than for the 1:1 and 8:1, but the 8:1 images were significantly worse.

Further tests were conducted to analyse the effect of compression on the two algorithms. These test revealed an increase in error for both filters as the compression was increased from zero to the maximum. It was noted that for the smaller images a large filter used on a highly compressed image would yield results similar to a small filter used on an uncompressed image.

Typical direction errors for the Canny and the Sobel detectors are illustrated in Figure 15.

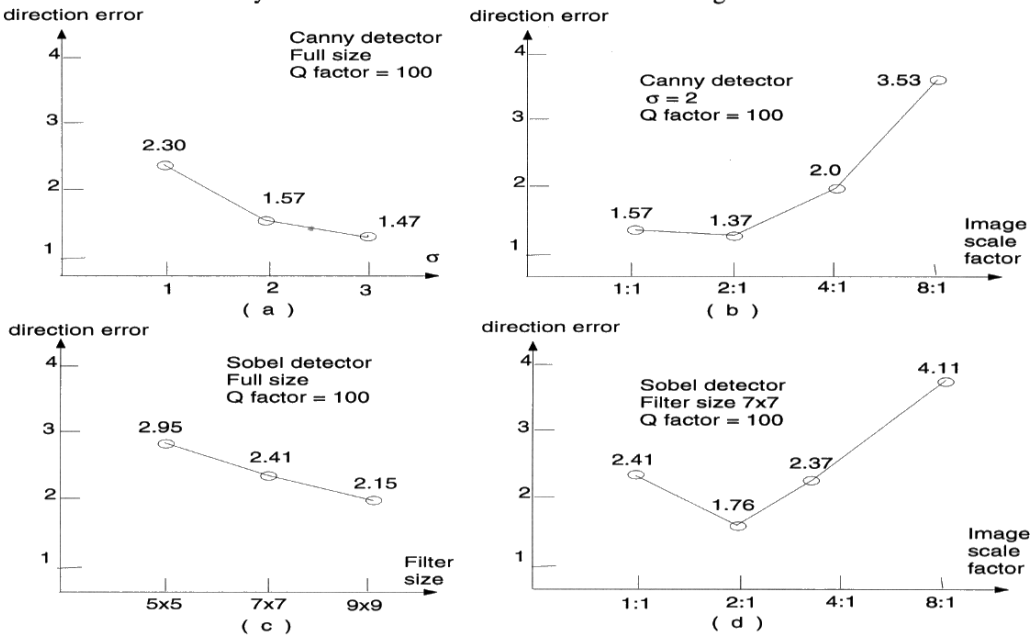

Fig. 15: The direction error deviation of the Canny and Sobel edge detectors in degrees.
(a) Canny detector with different values of σ. (b) Canny detector with σ = 2 on images of different sizes.
(c) Sobel detector with different sizes of the filter. (d) Sobel detector with size = 7x7 on images of different sizes.

3.3.2 Analysis on real images.

Test were conducted on the real images shown in Figure 16. To find the centre of a particular pipe joint a circle was manually drawn on the screen until it was overlayed on the pipe joint. The centre of the circle was then used to compute the angular errors of the edges. An error in the location of the centre of two or three pixels was considered unlikely to significantly change the result of the computation of the standard deviation.

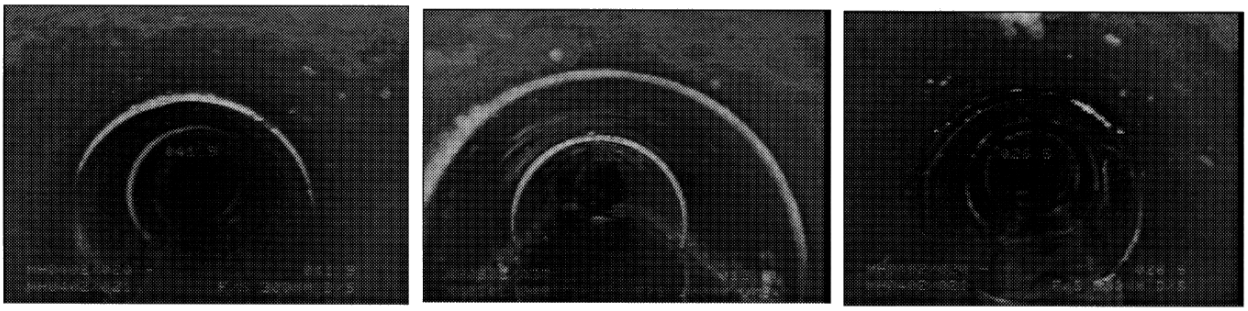

Fig. 16 The pipe images used in the experiment.

Sobel: The Sobel operator used in the previous tests was used to assess three real images. These tests revealed that there was a significant degradation of the edge direction quality compared to the synthetic images by almost an order of magnitude. However, in most cases the error in edge direction was reduced with smaller images to levels, in the case of the 8:1 sub-sampled images comparable with the synthetic images. Increasing the filter size had a positive effect for most images. There was only a marginal difference between compressed and uncompressed images and sub-sampling caused a reasonable improvement.

Canny: In general the results for the Canny algorithm were similar to those for Sobel when compared to the simulated images. The Canny filter gave better results for full size images, and worse results for the 8:1 scaled images. A σ of 2 in these trials gave the best results. Subsequent tests showed that the optimum σ on the sub-sampled image (4:1) for the Canny filter is 1.5.

3.3.3 Timing conclusions.

The level of computation required for the Canny and Sobel edge detector depends on the parameters used and the image size. On a Sun Sparc Classic the speed of the Canny was 1.5 second with $\sigma = 1$ and an image size of 192x144 and 47 seconds with $\sigma = 3$ and image size of 768x576. The speed of the Sobel was 1 second with a 5x5 mask size and image size of 192x144, and 34 second with 11x11 mask size and image size of 768x576, (Recursive implementation of the Canny algorithm[15] would improve the speed of the filter, whilst minimizing the loss of image at the boundary).

3.3.4 Conclusions.

In all of the tests on real images an increased error in edge direction was noted. This occurs even for the best pipe joints because of three effects: (i) the edges are not circular but have debris attached to the joint giving a genuine variation in edge direction; (ii) the edges do not have uniform reflectivity which will cause the edge detector to report various angles for the edges; and (iii) the real images will suffer from noise in the original video collection as it is collected at more than a hundred metres from the tape recorder and other degrading effects caused by the tape play back and frame grabbing. The level of noise found in the pipe joint edges resulted in edge direction errors that are significantly larger than those resulting from the JPEG compression. The effect of sub-sampling has the desirable effect of reducing the noise in the pipe joint so that the edge direction error is reduced for smaller images. For pipe joint tracking the 4:1 size appears to be adequate in most cases, a multi-scalar approach being possible if this proves not to be the case over a series of images.

3.4 A method for improving the reported edge direction.

To avoid the edge direction errors reported in Section 3.3, initial trials were conducted using the measured distance of the radius to a postulated centre in Hough space but such a method is highly computationally expensive. Hence, the edge direction information was chosen as the method most likely to give realistic results. As a consequence it has been taken as axiomatic that if the edge direction is improved for a given pipe joint there will be a corresponding improvement in the Hough transform performance. Image space sub-sampling and the use of smoothing filters were considered as methods of improving the edge direction information passed to the Hough transform, but such methods are likely to have a limited effect, and in the case of the filtering likely to be inefficient. An alternative was sought to improve the edge direction without adding undue computations. As a result a method was developed for smoothing in the linked edge space which has beneficial effects not only on the edge direction information, but also in de-cluttering the edges. The smoothing of edge direction information from the detectors (Canny and Sobel) was achieved for a given point "c" along a linked edge list by computing the average of the edge directions of the neighbours, i.e. $s = (a + b + c + d + e)/5$ (Fig. 17). In this way, local perturbation are

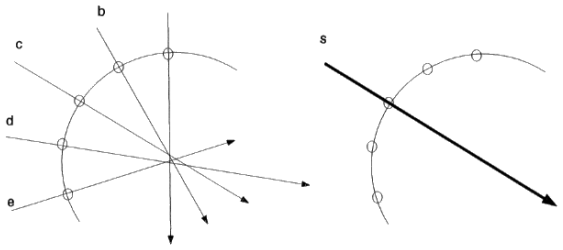

Fig. 17 The averaged direction of an edge point.

smoothed, and what is more significant, for points which lie on a circle, no error is introduced in the direction due to the

averaging. The results of tests using the angle filter showed that a filter length of seven used on sub-sampled images resulted in an reduction of 30% in the edge direction errors.

4. HOUGH TRANSFORM ON REAL IMAGES

4.1 Introduction.

As discussed earlier, the pipe joints can be approximated by circles. The detection of pipe joints using the Hough Transform (HT)[16] is achieved by looking for peaks in the Hough space using the parameters of the centre (a, b) and the radius **r** of the circles. The circular HT accumulates all the detected edge pixels and uses them to vote into a (a, b, r) Hough space defined by:

$$(x - a)^2 + (y - b)^2 = r^2 \quad (1)$$

Each edge pixel at (x, y) will vote to a cone in the 3D Hough space according to Eq. (1). To save computational cost, the standard approach is for an edge pixel to vote only along the direction of its maximum gradient and hence vote along a line for accumulating into the Hough space[17]. After the voting by all the edge pixels, a real circle in the image will then correspond to a peak in the Hough space at the location (a, b, r). The circle detection then depends upon the detection and location of peaks corresponding to genuine circles. The errors in the detected edge directions will cause errors in the location and the values of the peaks. An adaptation to make the HT less computationally expensive is to perform it in two steps: (i) vote into a 2D Hough space (a, b) and detect the location of the peaks as the centre of the circles, (ii) use these centres to determine the radius of the circles[18,19].

The main characteristic of the HT is that it is robust in the presence of noise since the peaks in the Hough space can only correspond to edge pixels meeting the criteria of the parameters used. The main problem with the HT is the large storage and computational requirements. Suppose the image size is $N_x * N_y$ and assume the centres of the circles are within the image. Then for a bin size of unity, the largest size of 3D Hough space is then $N_x * N_y * R_{max}$. (R_{max} is the largest radius of circle). In sewer images, when the camera moves towards a pipe joint, its projection expands until it moves out of the periphery of the camera's vision and R_{max} can therefore be larger than N_x or N_y. Another problem of the HT in this application is that the centres of the circles corresponding to different pipe joints may be very close to each other, and hence their corresponding peaks in the Hough space will interfere with each other. (This characteristics makes the two step HT less useful, although it is efficient.)

4.2 The adaptation the Hough Transform for this application.

As shown in Figures 1 - 6 and 7 - 12, there are at most three pipe joints visible in a sewer image and because of the width of the joint there may be two co-centred circles for each joint. The pipe joints, and their corresponding circles, have different radii. Hence, in the 3D Hough space (a, b, r), they should correspond to peaks with different **r**. The problem is that because the image noise and edge direction errors, the peaks in the Hough space are poorly defined, resulting in errors in the location and shape of the peaks. Furthermore, in this application the edge direction information is noisy giving indistinct peaks when **r** is divided into pixel units. The use of the HT with (a, b) parameters causes peaks which interfere with each other in the Hough space and makes it difficult to find the corresponding edges that originated the peaks to identify the radius **r**. A compromise was identified that would have some of the benefits of both methods. First of all, the **r** parameter was subdivided by a ratio of B, where B depends on the size of the images and R_{max}. In the experiment on a 192x144 sewer image, R_{max} is limited to 100 and B = 2. Hence the whole Hough space is $N_x * N_y * R_{max}/B$. Secondly, to avoid the problems of not knowing **r** and not being able to find the peaks, it was decided to use a coarse quantization of the **r** ordinate (6 pixels) when searching for peaks. The search for a peak at a radius **r** is conducted in three neighbouring sections in the Hough space corresponding to (r-1, r, r+1). The resulting peaks give both **a**, **b** and **r**, with a lower precision. However, a disadvantage of using a coarse **r** space was that it was difficult to distinguish between similar peaks in neighbouring **r** space, since they may correspond to the same edge pixels. Hence it would be necessary to back project to find a more accurate **r**[20]. The advantage

of this method is that after tracking has begun it is only necessary to look for peaks in a region corresponding to the tracked **r** in the Hough space. Searching for peaks in each **r** section in the Hough space is still a problem, as the pipe joints frequently appear only as arcs and have correspondingly weak peaks in the Hough space. On the other hand, there is only one circle at a particular radius **R**, and hence only one peak in a plane **r** = **R** in the Hough space. To find these weak peaks, the contribution of the Hough space surrounding a peak was accumulated in each plane **r** = **R** ($R_{min} < \mathbf{R} < R_{max}$), in the following manner:

(i) find the peak at (a_0. b_0) with the largest magnitude.

(ii) find the distribution of the peak along the **a** and **b** directions in a small area in the Hough space by use of

$$a' = (\Sigma a_i\, h_i) / (\Sigma h_i) \qquad a'' = (\Sigma a_i^2\, h_i) / (\Sigma h_i)$$
$$b' = (\Sigma b_i\, h_i) / (\Sigma h_i) \qquad b'' = (\Sigma b_i^2\, h_i) / (\Sigma h_i) \qquad (2)$$

where h_i = the number of votes at a particular **a**, **b** location and a′, a″, b′ and b″ are first and second moments.

(iii) based on Eq. (2), determine an area that is centred at a′ and b′ and with the size of a″/2*b″/2.

(iv) calculate the mean and variance of all the heights in the area using:

$$M_R = \Sigma h_i / n \qquad \sigma^2_R = (\Sigma h_i^2/n) - M_R^2 \qquad (3)$$

where n is the number of samples in the area.

(v) use M_R as the value of the peak in the Hough plane at **r** = **R** if it satisfies:

(a) $M_R > 2\sigma_R$.

(b) $M_{R-1} < M_R > M_{R+1}$

Otherwise, there is not enough evidence to support the existence of a pipe joint.

Having computed these values the next stage is to select the peaks that correspond to the best results. to achieve this, a simple form of back projection is used. A give edge chooses those among the previous determined peaks in different R plane in the Hough space if the **R** equals the distance between the pixel and the peak. The pixel will vote for the one that has the largest value of M_R among those chosen. After all the edge pixels have voted, only those peaks that maintain their values are considered to be correct Others are discarded because they correspond to the same edge pixels of the true peaks.

4.3 Conclusion.

A method has been developed where the locations and radii of pipe joints can be reliably identified using the noisy edges found in this environment. The use of a quantized **r** space with accumulation within a small range and the addition of a search in the region surrounding the selected **r** has proved reliable and reasonably efficient. This method can be further improved by using the (**a**, **b**, **r**) found in the previous images to localise the search for the new (**a**, **b**, **r**).

5. RESULTS OF TRACKING OF PIPES JOINTS

After edge detection and the Hough Transform, for each image, a number of circles are detected. Some are projections of pipe joints and others are false, corresponding to noisy edges and interference from other pipe joints. The circles corresponding to pipe joints will expand gradually and consistently. They should appear in each subsequent image with a slightly increased radius and a slightly changed location. The false circles are likely to change more dramatically and therefore methods can be used to differentiate them from genuine joints.

Because the bottom of the pipe is not smooth and the speed of the camera is not constant, the precise prediction of the location and orientation of the camera between frames is not easy. At this stage of the investigation tracking is based on the output the Hough Transform only. Each input circle from a new image is checked against all the tracked circles from

previous images. For each pair of tracked circles: (a_1, b_1, r_1) and the new circle: (a_2, b_2, r_2), three parameters are examined,

$d_r = r_2 / r_1$

$d_h = h_2 / h_1$ (h_n represents the the number of edge pixels supporting the circle.)

$d_{ab} = (d_a^2 + d_b^2)^{1/2}$ where $d_a = a_2 - a_1$ and $d_b = b_2 - b_1$.

If, $d_{rmin} < d_r < d_{rmax}$, $d_h < d_{hmax}$, $d_{ab} < d_{abmax}$ the two circles are matched and the tracking counter for the circle is increased by one. The values of these thresholds used in the experiments are $d_{rmin} = 0.9$, $d_{rmax} = 1.2$, $d_{hmax} = 0.6$, $d_{abmax} = \sqrt{2}*10$. If a circle is not matched by any currently tracked circle, it is regarded as a new circle. If a tracked circle is not matched by any new circle, the counter for the circle is decreased by one. If a circle is tracked for more than two sequential frames, it is deemed acceptable as a pipe joint. If a tracked circle is missed for one or more frame, it is not accepted but it will still be tracked for three frames before it is discounted.

For a test sequence of seventy-five frames of typical pipe images, after JPEG compression with Q factor = 100 and sub-sampled at 4:1 into the size of 192x144, the tracked circles were measured against the known pipe joints in these images, where the centres and radius had been measured by hand. A circle was regarded as well matched if the error in the centre and radius (d_{ab}, δ_r), was less than five pixels, where d_{ab} is defined as earlier and δ_r is the difference in radius. The match was regarded as rough if d_{ab} or δ_r was bigger than five and all values ware less than ten. If any value of d_{ab} or δ_r was larger than ten it was considered that a pipe joint did not exist. Five pipe joints appear in the test sequence. The results of tracking are shown in Table 1.

	Visible in frame	Tracked in frame	Well matched dab, δr < 5	Roughly matched dab, δr < 10	Not tracked dab, δr ≥ 10
Pipe joint 1	1 - 6	1 - 3	0	3	3
Pipe joint 2	6 - 26	6 - 20, 21 - 24	15	4	2
Pipe joint 3	21 - 39	22 - 38	12	4	3
Pipe joint 4	34 - 57	34 - 35, 38 - 56	12	9	3
Pipe joint 5	53 - 75	53 - 54, 57 - 58, 61 - 75	17	2	4

Table 1 Tracking of pipe joints over 75 image frames.

In the table, column 2 illustrates the frames in which the pipe joints were significantly visible to a human observer, column 3 shows the frames in which successful trackings took place, columns 4 - 5 illustrate the confidence levels obtained and column 6 shows the numbers of the frames where the pipe joints have not been tracked. The only tracking of wrong circles happens at the frame 3 - 6 at the start of the image sequence during the initialization stage.

The initial appearance of a pipe joint when it starts to be visible is weak and of small size. As the camera moves along, it becomes larger until it goes out of the view of the camera. Tracking and detection are difficult when the circle is either very small, or when it is very large and less likely to be circular. The best matching period occurs between the start and end of its visibility. The results of these early tests show that it is possible to detect and track the pipe joints in these image sequences. Furthermore, information concerning the reliability of the tracking is also collected and used to ensure the maximum overallreliability. Examples of the tracking in the 22nd and 53rd frame are shown in Figure 18.

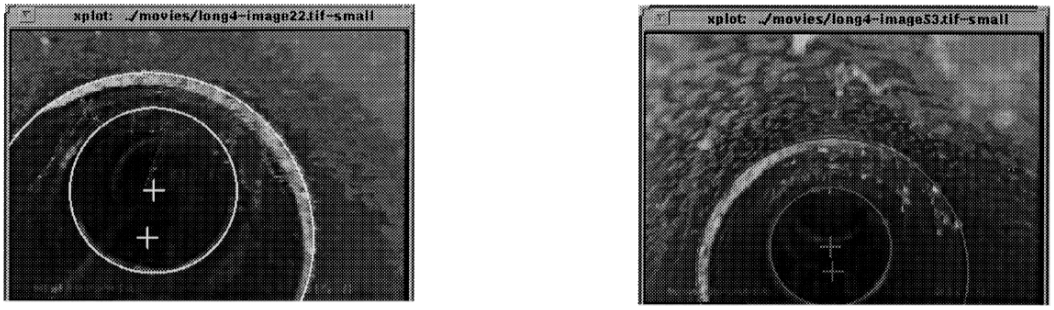

Fig. 18 The 22nd and 53r frame of the test sequence.
(The tracked pipe joints are presented as circles and their centers as crosses.)

6. CONCLUSIONS

In this paper the detection and tracking of pipe joints has been considered. A model of the pipe reflectivity, the camera, the pipe itself, and pipe defects has been described which allows the pipe joints to be described by circles. The Hough Transform was selected as the method most likely to achieve reliable results in the noisy images used in these trials. However, due to the practical requirements of this application, a number of modifications and tests were thought necessary. First, the effect of JPEG compression was considered using ideal images and the performance of several filters noted under various conditions. Second, real images were used to compare the results. As a result of this investigation, a method of minimising the edge direction error was developed, hence optimising the information available to the Hough Transform. Third, a sequence of images was used to perform tracking of pipe joints. It was found that for the images that make the up the majority of pipe inspections, reasonable reliability could be achieved.

Future work will continue in two directions. The first is to improve the tracking procedure and investigate more complex tracking methods and use feedback of tracking results in the HT and edge detection to constrain the search space and improve the tracking of more distant pipe joints. The second is to construct the pipe model and determine the camera location and orientation with respect to this, and thus facilitate the detection and classification of defects.

7. ACKNOWLEDGEMENTS

The project is carried out under the grant number: GR/441898 from EPSRC, U.K.

8. REFERENCES

1. F. Cohen, Z. Fan and S. Attali, "Automated inspection of textile frabics using textural models", *IEEE Trans. on* PAMI. Vol . 13 no. 8, pp. 803-808, August 1991.
2. N. Bryson, R.N. Dixon, J. J. Hunter and C. J. Taylor, "Contextual classification of cracks", *The Proceedings of British Machine Vision Conference 1993*, Surrey, U.K, pp. 409-418, September. 1993.
3. T A. Clarke and N. Lindsey, A triangulation based cross sectional profile, SPIE Conf. on Close range photogrametry and Machine Vision, Vol. 1395, pp. 940- 947, Sept. 3-7, 1990.
4. Water Research Centre and WAA Sewers and Water Mains Committee, "Manual of Sewer Condition Classification," WAC Engineering, England. 1988.
5. T. Ellis, A. Abbood and B. Brillault, "Ellipse Detection and Matching with uncertainty," *The Proceedings of British Machine Vision Conference 1990*, Glasgow, U. K, pp. 136-143, September, 1990.
6. U. Landau, "Estimation of a circular arc centre and its radius", Computer Vision, Graphics and Image Processing. Vol 38, pp 317-326, 1987.

7. P.L. Rosin and G.A.W. West, "Segmentation of edges into lines and arcs,", Image and Vision Computing, Vol 7, No. 2, pp 109-114, May 1989.
8. S. Pollard and J. Porrill, "Robust recovery of 3D ellipse data," *The Proceedings of British Machine Vision Conference 1992*, Oxford, U..K, pp. 38-48. September, 1992.
9. H.K. Yuen, J. Illingworth and J. Kittler, "Ellipse detection using the Hough transform," *The proceedings of 4th Alvey conference, Manchester,* pp. 265 - 271, U. K 1988.
10. "Xvideo software developer's guide," Parallax Graphics, Inc. 2500 Condensa St., Santa Clara, CA 95051, 1991.
11. E.P. Lyvers and R.O. Mitchell, "Precision edge contrast and orientation estimation," *IEEE Trans. on PAMI* , Vol 10, No. 6, pp. 927 - 937. November 1988.
12. 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
13. J.F. Canny, "A computational approach to edge detection," *IEEE Trans on PAMI*, Vol 8, pp. 679-698, 1986.
14 . P.E. Danielsson and O. Seger, "Generalized and separable sobel operators," *Machine vision acquiring and interpreting the 3D scene, Freeman (ed)*, Academic Press, New York, pp. 347-379, 1989.
15. R. Deriche, "Using Canny's criteria to derive a recursively implemented optimal edge detector,", *International Journal of Computer Vision*, Vol. 1 No.2, May 1987.
16. J. Illingworth and J. Kittler, "The adaptive Hough transform,", *IEEE trans. on PAMI*, Vol PAMI-9, No. 5, pp 690-698, September 1987.
17. E.R. Davies, "A new frame work for analysing the properties of the generalized Hough transform", *Pattern Recognition Letters*, pp. 1-7, June 1987. .
18. 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.
19. P. Kierkegaard, "A method for detection of circular arcs based on the Hough Transform," *Machine Vision and Applications*, vol. 5, pp. 249-263, 1992.
20. G. Grig and F. Klein, "Fast contour identification through efficient Hough Transform and simplified interpretation strategy", 8th international Joint Conference on Pattern Recognition, Paris, France pp. 498-500 (1986).