A Geometric Model for the Lumen of Blood Vessels in X-ray Angiography

[1]C.R.Hardingham, [1]D.J.Hawkes and [2]A.C.F.Colchester

[1]Radiological Sciences and [2]Neurology, U.M.D.S., Guy's Hospital, London Bridge, London SE1 9RT, U.K.

We have developed a model for representing the image of segments of blood vessels in x-ray angiograms, and a software package for generating the model from digitized angiograms. The model is based on projections of linearly tapering cylinders and is suitable for the analysis of vessel lumen position, size and shape. This is of interest for research into the disease atherosclerosis. The system has been validated using a set of post mortem samples. The accuracy of representation of the lumen radius is between +0% and +10% (95% confidence limits) and the accuracy of simple shape information is 81%.

Atherosclerosis is a common arterial disease. It is obliterative and its principal effect is due to the encroachment of a lesion on the lumen, with consequent effect on blood flow through the artery. It is a focal process and many foci may be present (Davies, 1986). X-ray angiography is a method of imaging the lumen of an arterial system by introducing an iodine-containing contrast medium into the blood stream before taking an x-ray. Traditionally, the cross sectional area of the residual lumen is measured and compared with the normal area, to give an estimate of the severity of a lesion. In recent years there have been several publications indicating that the shape of the lesion is an important indicator of the severity and prognosis of disease (Ambrose et al, 1986, Brown, 1981, Levin et al, 1982). Furthermore, there has been research into the 3-dimensional reconstruction of vascular trees from two or more projections (Hawkes et al, 1987). There is therefore a need for a model to represent the position, size and shape of the image of an artery, describing the features required for further analysis.

In the work preceding this project (Hawkes et al, 1988), a technique of linear calibration and background subtraction was developed, which produces a grey level image proportional to the projection of contrast, if noise and blurring caused by the finite MTF of the system is ignored. This linear response of the image to the projection of the object of interest permitted the development of the image analysis system described below. We have developed a model of the vessel image as a set of projections of generalized cylinders of elliptical cross section, and have implemented a package of computer programs, written in C and Prolog, to generate the model for digitized angiograms. The system has so far been used primarily in the analysis of lumen shape in atherosclerosis (Hawkes et al, 1989), but is also of general applicability where projected images of cylindrical or near cylindrical objects are of interest.

Although there is much research activity into quantitative angiography, there have been few publications regarding lumen shape reconstruction. Reiber et al. (1981) described research into a system for the three-dimensional reconstruction of lumen shape from two orthogonal views. There are two problems with this approach. First, two views are not always enough to reconstruct the lumen shape unambiguously. Secondly, it is not always possible to obtain two good orthogonal views of a vessel segment. The approach described here is to try and generate a useful representation of the lumen image, as seen from a single view. This can only provide limited information: as with angiographic assessment by eye, the appearance of the lesion is sensitive to the viewing angle. However, the aim is to provide information which is both useful on its own, and can be combined with the information from another view, if available.

When the shape of a blood vessel's lumen is investigated, the usual assumption is that a normal vessel, in a segment between branch points, has a lumen which is close to being cylindrical. A certain amount of taper is expected, the artery narrowing distally, but this is small compared with the effect of significant arterial disease. Deviation in shape from that of a slowly tapering cylinder is usually taken as a sign of abnormality. The projection of a cylinder is semi elliptical in cross section. A semi ellipse is easily represented and manipulated computationally. The model chosen uses projections of tapering cylinders (PTC's), constrained to taper linearly. These can be used to identify two types of feature in the image. First, they can represent grey level profiles in the image which could be projections of parts of cylindrical objects. Second, they can represent parts of "ribbon like" objects, i.e. strips with near parallel edges and height consistently higher than the background. Both these features are of interest, the first in looking for possibly normal artery, and the second in identifying artery which may or may not be diseased. In order to interpret the PTC's, it is necessary to have information about the goodness of fit between the image and the PTC.

The aims of this work were to find a method of representing the image of a non-branching, nearly straight segment of artery, using projections of tapering cylinders. The model must represent the artery image accurately and efficiently, identifying in particular the projections of parts of cylindrical objects.

The Model

The image may be thought of as a surface in three dimensions, the two dimensions of the digitized array and the grey level dimension. The convention used for describing this volume is to use x, y and z coordinates, where x and y are in

139

the plane of the image. and z is the grey level dimension. The y direction is along the axis of the vessel image, the x direction is normal to the vessel axis. The units of x and y are pixels in the digitized angiogram, z is in terms of "grey level units".

An artery slice is defined as a short length of artery whose image, the transverse density profile (TDP), is defined as the plot of grey level intensity across the x-ray image in a direction perpendicular to the image of the blood vessel centre line or axis. A TDP is therefore the image of an artery slice. Since the thickness of a slice is one pixel width in the image, and this is the limit of resolution of the digitized angiogram, the TDP is treated as the image of a cylinder of constant, unknown cross section along the thickness of the slice.

The Segment and Slice Models

The model of a segment of artery is referred to as a segment model. The segment model consists of a set of projections of tapering cylinders (PTC's). There are two types of PTC, the slice fit and the taper fit. A slice fit is a projection of a cylinder of length one slice and of constant cross section. It is a match to a single TDP. A taper fit is a projection of a cylinder which is more than one slice long. Its cross section is semi elliptical, varying linearly along the length of the taper fit. A slice model consists of one or more slice fits for a given slice. A segment model consists of the union of the slice models for the slices making up the segment, plus zero or more taper fits to parts of the segment. A slice model is never empty and therefore a segment model always contains at least one PTC matching every slice.

Parameters for the Model

The model is made up of a set of slice and taper fits, which are both PTC's. Two types of parameters are used to describe a PTC: Cylinder parameters and Fit parameters. The cylinder parameters specify the tapering cylinder. The fit parameters describe how well the cylinder fits the image.

The Cylinder Parameters

Each PTC is semi elliptical in cross section at all points along its length. A constrained semi ellipse of this type can be specified using three parameters: the centre position, the spatial radius and the height. A slice fit can be described using these three parameters, plus the position of the TDP along the vessel image axis. For a taper fit, the cylinder parameters are the positions along the vessel image axis of the ends of the PTC, and two triplets defining the semi elliptical cross sections at each end of the PTC. The cross section is defined to vary linearly along the PTC and may be calculated at any point by linear interpolation from the values at the ends.

The Fit Parameters

There are four fit parameters, called Enclosure, Fullness, Fit Points and Skew. Enclosure is the proportion of the TDP enclosed by the slice fit. Fullness is the proportion of the slice fit which is enclosed by the TDP. The fit points parameter is the proportion of TDP points lying close to the slice fit. The skew of a slice fit is the difference between the

TDP area enclosed by the right and left hand halves of the PTC, divided by the total TDP area.
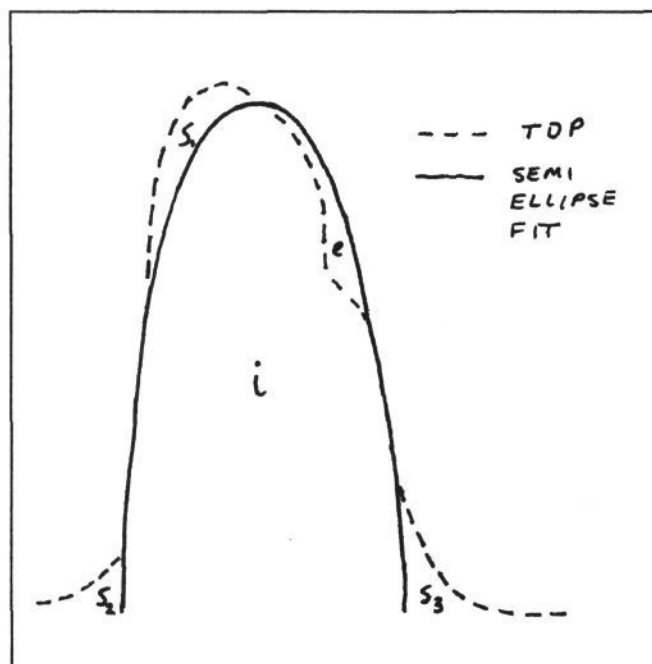

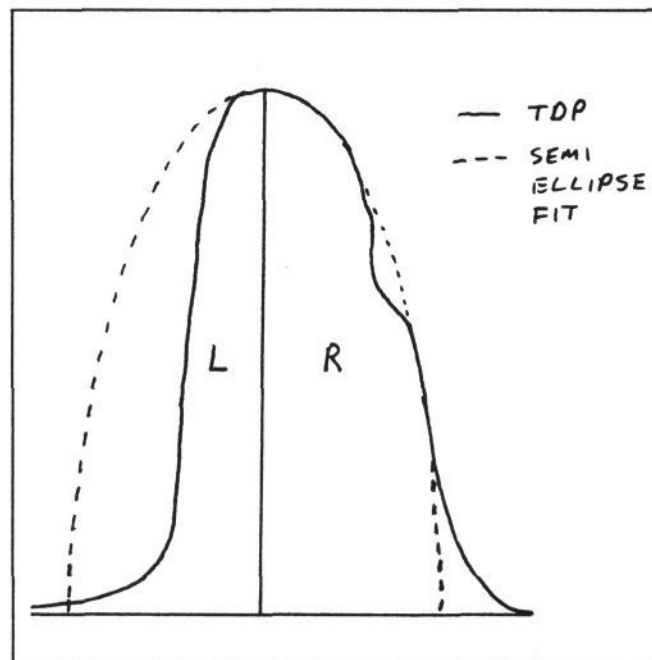Figure 1. The Enclosure and Fullness Parameters. Enclosure = i / (i + s1 + s2 + s3), Fullness = i / (i + e).

Figure 2. The Skew Parameter. Skew = (R - L) / (R + L)

The Goodness of Fit Mark

In order to make decisions about PTC's, several of the algorithms described below generate a goodness of fit mark for each fit, based on a weighted mean on the enclosure, fullness and fit points parameters. Strictly speaking the goodness of fit mark is not part of the model. However it is stored with the fits for convenience and is an important factor in how or whether a PTC is used.

Generating the Model

Generating the model is in two main steps. The first step is to split the segment image into slices (TDP's) and generate the slice model for each slice. This consists of a set of slice fits. The second step is to combine the slice fits into taper fits.

The software which generates the slice model is written in the C programming language. The program which generates the taper fits is written in Prolog.

Generating Slice Fits From the TDP

There are two main reasons for wanting to fit a semi ellipse to a TDP. One is to find the edge positions and height, in order to locate the vessel and estimate its rough size. The other is to identify regions of part-elliptical shape in the TDP, which might be evidence of the original boundary in a vessel which has been affected by a lesion. The two different algorithms used to find semi elliptical fits to a TDP are intended to identify these two different types of information. The first, the enclose algorithm, works entirely on the basis of area enclosed by the TDP. It therefore seeks out the edges and height. The second, the Hough algorithm, looks for ellipses which lie close to many points of the TDP, and so identifies part-elliptical regions in the TDP.
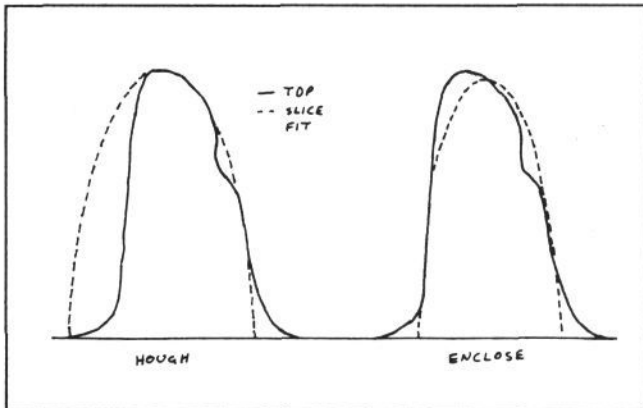

Figure 3. Slice fits Generated by the Hough and Enclose Algorithms. The Hough algorithm works by fitting to points on the TDP boundary, whereas the Enclose algorithm matches the area under the TDP. Although the ellipse shown here for the Enclose algorithm matches a significant length of the TDP boundary, the number of points it lies close to is small, since the gradient is steep.

Generating the slice fits for a TDP has the following steps (the algorithms are described in more detail below): The Hough algorithm is run for the first time, using a large parameter space. Then the Enclose algorithm is run. The Hough algorithm is run for a second time, using a reduced parameter space for improved accuracy. Finally the Prune algorithm is run. This reduces the number of semi ellipse fits in the model by eliminating semi ellipses which are very similar to higher marked semi ellipses.

At each of the above stages, if a new semi ellipse is found which may be put into the slice model, it is marked and put into an internal list of semi ellipse fits. It is on this list that the prune algorithm operates.

The Enclose Algorithm

The Enclose algorithm was developed to find semi-ellipses which match the whole area of the TDP, lying as close as possible to it, and enclosing it, possibly allowing a small spill, or area where TDP exceeds candidate semi-ellipse, to account for noise. There are two main problems to be overcome in finding such semi-ellipses. First, there are in general infinitely many semi-ellipses which enclose a given TDP, but touch it in two or more places. Second, if the generated semi-ellipse completely encloses the TDP, it almost certainly overestimates the area of the vessel cross-section, because of image noise.

The enclose algorithm is iterative and generates a sequence of candidate semi-ellipses, converging on a limit semi-ellipse which encloses the TDP completely. The semi-ellipses alternately enclose and clip (i.e. just fail to enclose) the TDP in all of the three regions. The semi-ellipses are divided into left, middle and right regions. The amount by which the TDP exceeds the semi ellipse in each of the three regions is used to control the adjustment made to the semi ellipse parameters at each iteration. At each stage, the program halves the steps and thresholds according to which the semi-ellipse parameters are adjusted. The sequence starts inside the TDP, and ends outside it, moving in smaller and smaller steps as it approaches the TDP.

The candidate semi-ellipses are marked according to goodness of fit and the best are saved. This approach overcomes the above difficulties in that: The limit semi-ellipse encloses the TDP, but is also constrained to touch it on the left, in the middle and on the right. The curve-length of the middle region is equal to the sum of the left and right regions (which are symmetrical). This generates a limit semi-ellipse which favours height and width equally. Thus various candidate semi-ellipses are generated which clip the TDP, but only by small amounts. These small amounts of spill allow the effects of noise to be minimized, as the marking system favours semi ellipses which lie close to many points.

Noise at the edges is dealt with by a "squelch" method, similar to squelch in FM radio, where a signal is considered to be of zero amplitude if below a given threshold. Spill, or area where the TDP exceeds the semi ellipse, at a point x is calculated by subtracting the candidate semi ellipse z value from the image z value at x. This is modified by treating the image z value as zero if it is smaller than a certain threshold. The threshold is a proportion p of the semi ellipse height z. The whole enclose algorithm is executed several times with different values of p (e.g. 0.01, 0.03, 0.1, 0.3).

The Hough Transform

The Hough transform (Ballard & Brown 1982, pp123-131) can be used to search for many types of shape and is not specific to ellipses. A general explanation of the Hough transform is beyond the scope of this report, but a brief description follows, using this application as the example.

A transform is defined between "image space" (here the TDP) and "parameter space" for the matched object (a semi-ellipse). Parameter space is N-dimensional, where there are N parameters defining the matched object (here the parameters are the centre position, spatial radius and height). A point in parameter space defines an instance of

the matched object. First the parameter space is quantized and an array is allocated to represent it between set limits for each dimension. The array is set to zero. Next the transform is used to map the image onto the parameter space array. In this case, we traverse all point pair combinations in the TDP. For any pair of points there is in general an infinite number of semi-ellipses which pass through the two points, therefore a line in hough space is incremented. Finally the parameter space array is interrogated for local maxima. These represent instances of the matched object which appear most in the image.

The Prune Algorithm

In general, the ellipse fitting algorithms find many semi ellipses. Some of these are very similar, indeed there may be duplication. The prune algorithm is intended to remove redundant semi ellipse fits, i.e. fits which are very close to another, better fit. It works as follows: The list of fits is searched to find the highest marked fit. This is moved to the start of the list. The other fits in the list are checked for closeness to the highest marked fit. Any which are very close are removed from the list and discarded. The above process is repeated for the list of fits beginning after the highest marked fit which was just found, and so on until the whole list has been processed.

Deriving the Fit Parameters and Mark

As each slice fit is generated, its fit parameters and mark are calculated and stored with it. The fit parameters are enclosure, fullness, fit points and skew. The mark is a weighted average of enclosure, fullness and fit points, calculated as 0.3*enclosure + 0.3*fullness + 0.4*fit points. This mark is used during the generation of the slice model, including the prune procedure at the end.

Generating the Segment Model

The segment model consists of the slice fits contained in the slice models of all the slices making up the segment, plus taper fits. The taper fits are generated from the slice fits, in two phases. First sequences of close slice fits are found, then taper fits are generated from the sequences of close slice fits, using a modified line fitting algorithm described below. In both phases, the algorithms use only the cylinder parameters and the mark of each slice fit.

Finding Sequences of Close Slice Fits

The algorithm which finds sequences of close slice fits does so using a restricted priority-first search. The aim is to find all pairs of end points which can be connected by sequences of close slice fits, with the best sequence for each pair, and ignoring sub-sequences. The cost function used is a called the semi ellipse closeness factor. This is a measure of closeness for two semi ellipses, having value 1 for two identical semi ellipses and 0 for any two semi ellipses with no area in common. It is calculated from the two sets of semi ellipse parameters and is an approximation of the proportion of common area enclosed by the two semi ellipses. Two slice fits are considered "close" if this factor exceeds 0.85.

The algorithm works in three phases, the forward, reverse and merge scans. In the forward scan, sequences are built up in the direction of increasing slice number. In all cases, the sequence is extended using the closest slice fit to the current end-of-sequence, unless the closeness factor is less than 0.85. Sequences are propagated from all slice fits which do not already form part of a sequence (these have already been propagated). In the backward scan, sequences generated by the forward scan are propagated in the direction of decreasing slice number, as far as possible, using the closest slice fit if it is close enough. In the merge scan, overlapping sequences are combined if possible, i.e. if there is a pair of close fits in the overlapping region, and if this creates a longer sequence. At the same time as the merge scan, subsequences and lower marked sequences between the same end points (slice fits) are removed.

Making Taper Fits

In order to find segments which are projections of tapering cylinders, a modified least squares line fitting algorithm is applied. This produces a poly line fit (PLF) to each sequence of close ellipses. The poly line fit algorithm works as follows, to find a PLF between slices L and R: If L and R are close (currently less than 8 slices apart), one straight line fit is found by the method of least squares. Straight line fits by least squares are found between L and R, L and M, and M and R, where M is the mid point between L and R. If the sum of squares of the errors for the fit between L and R is less than double the sum of squares of the errors for the two straight lines between L and M and M and R, the single straight line fit between L and R is used. Otherwise, more than one straight line is used. An initial break point is chosen by considering all possible points between L and R (except within 2 of L or R), and using the one which minimizes the sum of squares of errors. If this point is T, then the whole algorithm is run for the segments between L and T, and T and R. If the result of each of these PLF's is a single line, the choice of T as break point is confirmed. If either or both results in more than one straight line, then the whole algorithm is rerun for a length including T, between the rightmost break point between L and T, and the leftmost break point between T and R.

The above poly line fit algorithm is run for the three semi ellipse parameters . If any of these fits results in a break point, the segment is cut at that point. The mark of a taper fit is derived by averaging the marks for the derived slice fits along it.

Validation

The data used to develop and verify the methods described were already in existence at the start of the project, having been prepared for a previous piece of research. (Hawkes et al. 1988). For each of seven artery segments obtained at post mortem, the data available were: two digitized image files (256 *256 pixels, 1 byte per pixel), digitized at 10 pixels per mm, giving orthogonal views of the segment, and photographs, including mm scales, of cross sections through the artery segment at at least four points.

The software package was used to generate an instantiated model for each of the segment images. These were then investigated for accuracy and content.

The photographic data does not contain information about the vessel centre position and therefore this could not be validated. However, accurate determination of spatial radius implies accurate measurement of the centre line. The spatial radius of both the best taper fit and the best slice fit were compared with the values from the photographs. The Pearson correlation coefficient was calculated for each set of paired values, and the least squares regression line with 95% confidence interval at the mean.

The goodness of fit mark of the best slice fit was used to estimate the degree of circularity of the lumen. The mark was compared with a categorization between "rough" and "smooth", obtained by agreement between two observers, using the cross section photographs.

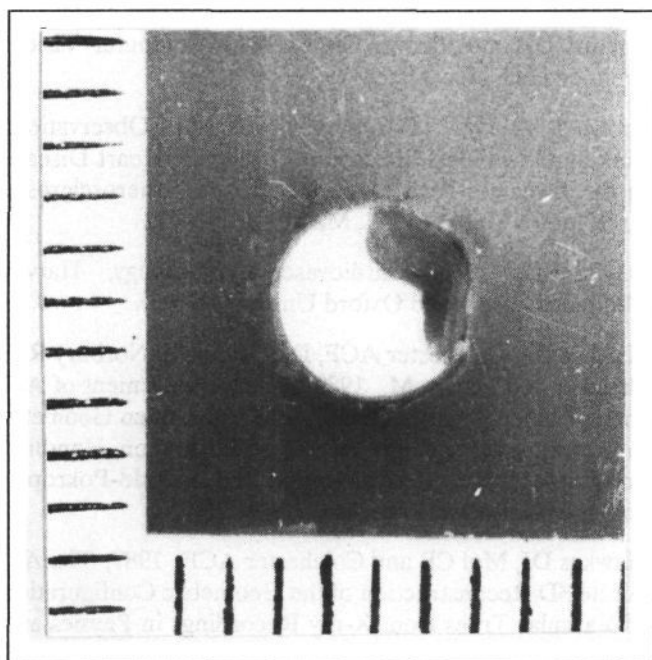The information content of the slice and segment models was also assessed.

Figure 4. Sample Photograph of an Artery Cross Section

Results

The correlation between spatial radius in the best taper fit and in the photographs yielded a correlation coefficient of 0.96 and a 90% confidence interval at the mean of width 0.14 mm (mean 1.4mm). This corresponded to between +0% and +10% of the value measured from the photographs. (See Figure 5).

The correlation between spatial radius in the best slice fit and in the photographs yielded a correlation coefficient of 0.94 and a 90% confidence interval at the mean of 0.15 mm.

The correlation between the goodness of fit mark and the Rough/Smooth categorization gave 81% accuracy for using the mark to predict the Rough/Smooth category, if the best threshold was used.

Investigation of the information content of the model showed that the average number of slice fits in a slice model was 10.8 (range 1 to 26). The mean length of a taper fit was

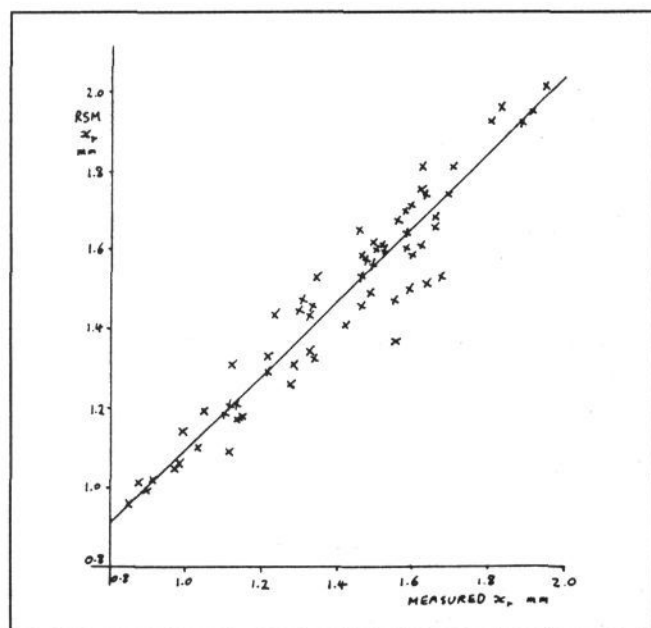11.9 pixels, and each point along the artery had on average 4.8 taper fits passing through it.

Figure 5. Scatter Plot of the Best Fit Spatial Radius (RSM xr) against Measured Spatial Radius.

Discussion

The accuracy validation showed the taper fit values to be slightly (probably not significantly) more accurate than the best slice fit results, which confirmed the validity of the algorithms generating the taper fits from the slice fits. The 95% confidence limit of + 0% to + 10% for the spatial radius estimate was very satisfactory, bearing in mind that the errors in measuring the photographs and the geometric magnification factor for the images will have contributed to this.

The predictive accuracy of 81% for the Rough/Smooth categorization indicates that the model provides a sound basis for the investigation of shape. The fit parameters described provide useful shape information and we have also used them for the identification of concentric and eccentric lesions (Hawkes et al, 1989). A possible method for further shape investigation is to analyze the difference between the image and the PTC's of the model.

The information content results show the taper fits to be roughly ten times more compact than the slice fits: while a taper fit has twice as many parameters as a slice fit, there are half as many passing through each point, but each taper fit describes on average ten slices. In spite of being so much more compact, the taper fits describe the spatial radius slightly more accurately than the slice fits. Therefore they are a considerable improvement as a basis for subsequent processing. The average length of a taper fit of 1.19 mm is comparable to the average artery radius, 1.4 mm. This is the order of length one might expect for a cylindrical feature, in the type of diseased artery under investigation.

Limitations of the Results and Method

The tests were carried out on high quality angiograms of straight artery segments. It remains to be seen whether the

143

technique works on noisy data, or lower resolution images. However, the arteries were fairly small (average 1.4 mm radius), and so the system could be expected to hold at least for the larger arteries in vivo.

The method as described has only been tested for artery segments parallel to the imaging plane. It is thought that the method is applicable to oblique artery segments, provided that the obliquity is less than 30 degrees, but this remains to be tested.

A further limitation is that the technique does not apply to junctions of arteries. Atheroma is commonly found at junctions and this is therefore a significant limitation.

A more immediate problem, if the program is to be applied to other angiographic data, is the number of arbitrary thresholds in the program. Before the system can be applied generally, robust ways must be found of setting these thresholds.

Comments on the Computational Techniques

Throughout the system, non deterministic techniques are used in that the modules which identify possible features of interest are independent of the modules which assess the identified features, and of each other. This approach is heuristic in that the program cannot be guaranteed to extract all projected cylindrical features. However there is the advantage that new feature extraction algorithms can be added extremely easily (if only one could think of them!). This approach, while speeding up development time, is not particularly expensive in terms of execution speed, and in addition has a transparent parallelism which suggests that it would transfer well to multiprocessing architectures.

The algorithms used to generate the slice model, e.g. the Enclose and Hough algorithms, are entirely procedural and were easily and efficiently programmed in the "C" language. The algorithms used to generate the taper fits were written in Prolog. We originally decided to do this because it was felt that there would be a fair amount of rule based or declarative content in the algorithms, and because we wish to use a rule based system to interpret the model, once instantiated. The parts written in Prolog turned out to have a mostly procedural rather than declarative nature, which created problems. However these were partly counterbalanced by Prolog's excellent list processing and pattern matching features, as well as the automatic data structure allocation and easy recursion. In spite of slow execution speeds and frustrations with the amount of memory space needed by the run time system, we feel that Prolog is very suitable for interpreting the data once that the feature extraction process has taken place.

Conclusions

This technique of model matching in the analysis of lumen images in x-ray angiograms is very promising. The model using projections of tapering cylinders was shown to be an adequate and efficient representation of the image of normal and abnormal straight artery segments, and there is good reason to hope that it is extendable to curved, oblique segments. The techniques used for instantiating the model were successful in detecting areas of projected cylindrical shape, as well as the more traditional requirement of edge and height detection.

References

Ambrose JA, Winters SL, Arora RR, Eng A, Riccio A, Gorlin R and Fuster V, 1986, Angiographic Evolution of Coronary Artery Morphology in Unstable Angina, J Am Coll Cardiol 7:472-8

Ballard DH and Brown CM, 1982, Computer Vision, Prentice Hall, New Jersey

Brown BG, 1981, Coronary Vasospasm. Observations Linking the Clinical Spectrum of Ischemic Heart Disease to the Dynamic Pathology of Coronary Atherosclerosis, Arch Intern Med-Vol 141, May 1981: 716-721

Davies MJ, 1986, Cardiovascular Pathology, Harvey Miller Publishers and Oxford University Press

Hawkes DJ, Colchester ACF, De Belder M, Norbury RN, Camm J and Davies M, 1988, The Measurement of Absolute Lumen Cross Sectional Area and Lumen Geometry in Quantitative Angiography in The Formation, Handling and Evaluation of Medical Images (ed A Todd-Pokropek and MA Viergever) in press

Hawkes DJ, Mol CB and Colchester ACF, 1987, The Accurate 3D Reconstruction of the Geometric Configuration of Vascular Trees from X-ray Recordings in Physics and Engineering of Medical Imaging (ed R Guzzardi) pp 250-256 (Martinus Nijhoff, Holland)

Hawkes DJ, Hardingham CR and Colchester ACF, 1989, Matching of Geometric Models in the Analysis of Lumen Shape in Digital X-ray Angiography, Presented at Information Processing in Medical Imaging, 11th Conference, San Francisco June 1989 and submitted for publication.

Levin DC and Fallon JT, 1982, Significance of the Angiographic Morphology of Localized Coronary Stenoses: Histopathologic Correlations, Circulation 66, No. 2, 316-320

Reiber JHC, Troost GJ, Gerbrands JJ, Booman F, den Boer A, SerruysPW, 1981, Densitometric Assessment Severity of Coronary Obstruction from Monoplane Views; Three-Dimensional Reconstruction Arterial Segment from Two Orthogonal Views, Computers in Cardiology 1981, 333-336.