

A Model-based Approach to the Reconstruction of Three Dimensional Arteries from Biplane Angiograms

S. T. Rake¹
IBM UK Scientific Centre,
Winchester

K. D. Baker
Department of Computer Science,
University of Reading

G. D. Sullivan
Department of Computer Science,
University of Reading

This paper describes the automatic production of three dimensional computer models of arteries, where the originating data is a pair of biplane X-ray angiograms. By developing the process of reconstruction using a hierarchy of generic models, it is possible to build a limited number of plausible reconstructions. There are indications that further research will be able to identify the unique reconstruction that matches the original structure. The process avoids the extensive human operator intervention and simplifying assumptions used by other researchers.

Introduction

This paper describes the automatic reconstruction of three dimensional arteries from biplane X-ray angiograms. Angiography is the study of the vascular system by means of radio-opaque material. The X-ray images of the vessels are captured, after the injection of the material, on two image planes which are at widely diverging angles (approximately 90°). The task is to use these two images to automatically build a computer model of the original artery structure.

Three dimensional reconstruction of arteries is important for both qualitative and quantitative reasons. It is well known that visualising three dimensional structures from two dimensional data can be difficult due to misleading depth cues. If the three dimensional structure can be presented so that it can be viewed from any angle, then an accurate perception of its three dimensional shape can be quickly acquired, and ambiguities will be resolved. Once a three dimensional reconstruction has been achieved a number of quantitative measurements are possible.

The processes used here to build individual arterial structures in three dimensions are well understood and have been described by Mol(9). By identifying corresponding points in each view for all the points of interest, a three dimensional reconstruction can be achieved (Figure 1).

This research shows that developing processes under the *model based* paradigm (Brady(2)) accurate reconstruction can be achieved. In this task a geometric or parameterised model is inappropriate and descriptive or *generic* models are used. The development of the three dimensional pipe structures takes place in two stages. The first stage is to identify parallel edged structures (ribbons) in each view and has been described elsewhere (12). The second stage is to use

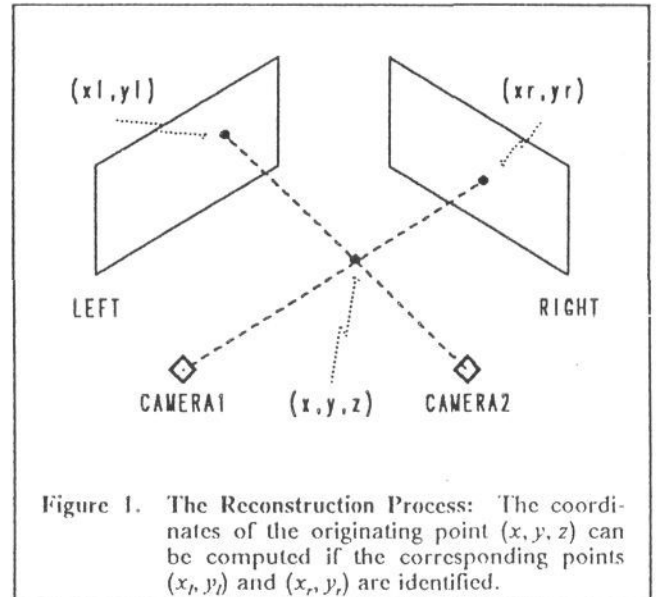


Figure 1. The Reconstruction Process: The coordinates of the originating point (x, y, z) can be computed if the corresponding points (x_l, y_l) and (x_r, y_r) are identified.

knowledge from the models to generate an interpretation of the ribbon features. The result is a three dimensional reconstruction that corresponds to the features in the images.

Unlike other research in which there is extensive operator interaction (1,11) or in which there are significant simplifying assumptions (4,13), the processes described here require a human operator to identify **one** pair of corresponding features and there are no simplifying assumptions.

A Hierarchy of Models

Marr(8) proposes a hierarchical model of three dimensional structures with a change of scale at each level in the hierarchy. This research uses a hierarchy of models, each level in the hierarchy containing a model which is appropriate to a particular stage of the developing solution. A pipe model describing arteries in terms of three dimensional generalised cylinders, a ribbon model describing a two dimensional projection of those cylinders and an edge model all appear in the hierarchy.

Two levels of three dimensional model are needed by this task. At the lower level there are arteries or pipes in three dimensions which are modelled by generalised cylinders. At the higher level there is a set of connected pipes. The purpose of the lower level of the three dimensional model employed in this task is to achieve the

¹ This research has been conducted while the author has been registered for a PhD with the Department of Computer Science at the University of Reading.

placement of pipes in space, and higher level three dimensional model is used to achieve connectivity of those pipes. The model needed in this task is a *generic* one, in which the objects to be identified are described in very general terms, avoiding geometric and/or parameterised descriptions. In fact the model is more *qualitative* than *quantitative*. As a result, the model will apply to any structure that can be described as connected sets of generalised cylinders.

Each pipe, in an uninterrupted or obscured projection into the image, will produce a feature which will be bounded by two edges which are close. The centre line of the feature will be the axis of symmetry between the edges and it will be the projection of the axis of the generalised cylinder. The arteries are essentially thin, with a smoothly changing cross section, producing a silhouette which has approximately parallel edges. These approximately parallel edged features are called *ribbons*. This suggests two further levels of model, a ribbon model and an edge model.

The hierarchy of models consists of

Connected Pipe Model
Pipe Model
Ribbon Model
Edge Model

The edge and ribbon models are used to develop the descriptions of ribbons found in the images. The details of these models are described elsewhere (12).

Building Pipe Hypotheses

Mol(9) and others(1,11) have demonstrated the use of the 'auxiliary line' (also called the 'epi-polar line') to aid the identification of corresponding points and features in the two images. After the point of interest has been identified in the first view, it is possible to compute the location of an 'auxiliary line' in the other view on which the corresponding point should lie. The operator uses the line to help identify and plot the corresponding point. When both points have been identified the location of the feature in three dimensions is calculated (Figure 2).

If enough data points are accumulated a complete three dimensional reconstruction can be built.

The principle problem is the identification of corresponding points and features. If one is prepared to take on the extra computing load involved, the correspondence problem can be avoided by looking for *all possible* correspondences between ribbons in the left and right views and developing a set of structures in three dimensions which represent the total possible set of correspondences. Each possible correspondence allows a *pipe hypothesis* to be constructed. The complete set of pipe hypotheses will be a *Hypothesis Set* and will contain the three dimensional structure with the correct correspondences (see Figure 3).

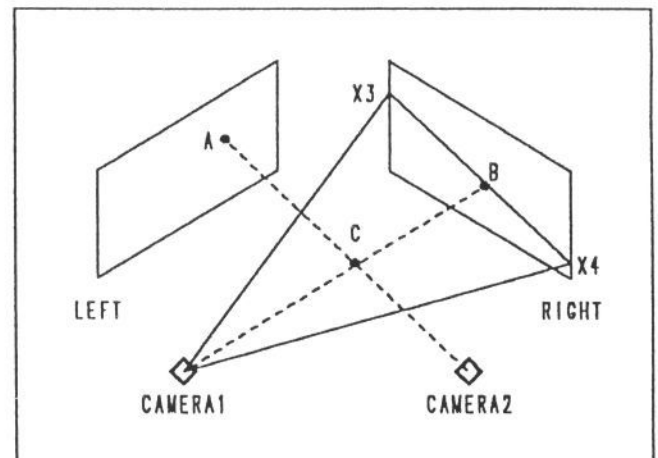


Figure 2. Using the auxiliary line: Point A is chosen by the operator. The location of the auxiliary line X3-X4 is then computed and displayed. The system invites the operator to plot the corresponding point B. The location of C in three dimensional space is then computed.




Figure 3. The Hypothesis Set: Ribbons L_1 , L_2 and L_3 are matched with R_1 , R_2 and R_3 creating the Hypothesis Set containing Pipes 1 to 7.

The nature of the problem is now changed from one of identifying correspondences between two dimensional features, to one of extracting a set of pipes from the Hypothesis Set which represent the original three dimensional structure.

The auxiliary line is used both to locate ribbons which may be in correspondence (to develop a pipe hypothesis) and to locate the corresponding points within those ribbons. If the auxiliary lines being originated in one ribbon encounter a ribbon in the other view, then a pipe hypothesis is constructed from these ribbons. An exhaustive search of all ribbon pairs results in the Hypothesis Set.

Recovering Connectivity

Having avoided the difficulties of identifying corresponding features, the problem now is to *generate*, from the Hypothesis Set, an *interpretation* (a subset of

the Hypothesis Set) which constitutes a plausible reconstruction of the original three dimensional structure. However, the Hypothesis Set does not contain adequate data for the generation process. There are two problems. One of the constraints to be used during generation is connectivity between pipes, but, due to errors in the locations of the ends of pipes in the Hypothesis Set, connectivity between the pipes is lost and has to be recovered. This is achieved by transforming the Hypothesis Set into a graph(3), in which connectivity of the ends of the pipes has been recovered, and where the format is ideal for the generation process. Nodes in the graph represent locations in three dimensional space which correspond to the ends of pipes while arcs represent the pipes. Connectivity between pipes is revealed by arcs which are incident with a particular node.

The graph can be built by identifying ends of pipes which have the same location in space and creating a node to represent that location. However, due to various types of accumulated error, the ends of the pipes are not located with such precision. Making 'close' (in three dimensions) nodes coincident allows the connectivity to be recovered. The definition of 'close' is controlled by a parameter which has been determined by experiment.

Generating Interpretations

The problem is to extract interpretations I from the graph G . For I to be acceptable, there must be a match between I and some prior expectation of the acceptability of I . One approach(5) is to 'generate and test'. Interpretations are generated from the available data, and then tested to verify their validity. The key to success is to powerfully constrain the generation step, thus reducing the search space and limiting the computation required by the testing process. In this particular task there are a number of constraints which can be applied to reduce the number of generated interpretations.

Three Dimensional Structure. The three dimensional structure must be retained.

An interpretation must be developed so that the structural relationships revealed in the graph G are retained.

Connectivity. The connectivity of the structure must be retained.

No pipe is removed from the graph if it results in the disconnection of parts of the graph.

The Ambiguity Table

The third type of constraint can be deduced from the Hypothesis Set and is a result of the exhaustive search that is used to develop the Hypothesis Set. Suppose that a ribbon L_1 in the LEFT view, and a ribbon R_2 in the RIGHT view, lead to a pipe hypothesis P_2 . Suppose also that, later in the process, L_1 and R_3 lead to a hypothesis P_3 . These two hypotheses are *potentially ambiguous*. The ambiguity suggests that if P_2 is believed to be part of the true three dimensional structure and should be retained in the developing interpretation, P_3 should be discarded. Notice that P_3 can only be dis-

carded if it does not disconnect the interpretation. The complete set of ambiguous pairs among the pipes in the Hypothesis Set is called the *Ambiguity Table*.

If an ambiguous pipe is included in an interpretation, its pair in the Ambiguity Table can be removed from the interpretation, but only if the graph is not disconnected. The ambiguity between the pipes is then resolved. If both pipes in an ambiguous pair are included in an interpretation, the ambiguity is considered to be resolved. An interpretation is complete if all ambiguities have been resolved.

It is possible to cast the application of the constraints to the Graph in terms of the Constraint Satisfaction Problem(7). This research has implemented a simple recursive, backtracking search algorithm very similar to the backtracking algorithm described by Nadel(10). The significant difference to that algorithm is the way in which choices are enumerated at each branch point. In classical CSP the choices are made sequentially from the domain of values available to the variable. In this task the choice is made from the developing interpretation which is carried forward to each branch point. The current state of the graph (i.e. the developing interpretation) has the information necessary to allow the search process to produce the candidate branches for exploration at the next lower level of the search tree. When the selection of a branch in the search tree results in all the ambiguities in the Ambiguity Table being resolved, an interpretation (the current state of the graph) is generated.

To start the generation process, the operator is required to identify one pipe as the product of a correct correspondence. This is the root of the search.

Testing Interpretations

The second stage of the 'generate and test' process is to analyse the generated interpretations for some form of acceptability. (5) assumes that there is a model against which the generated interpretation can be tested. However, in this task it is not obvious how to test for acceptability since the constraints derived from the connected pipe model have been exhausted during the generation process. That is, the underlying physical reality of the original graph has also been retained, connectivity has been retained, and all the ambiguities have been eliminated.

However, visual inspection shows that, in many of the interpretations, all is not well. Each interpretation can be displayed on a device that can represent the three dimensional structure as a set of vectors which can be rotated by a user. If the interpretation is rotated so that its projection is into a plane which is parallel to the original image planes, it can be observed, in some cases, that the interpretation is producing a projection which does not match the ribbons in that view. The projection shows that either a pipe is shorter than it should be (the relevant ribbon has an *incomplete end*), or that a pipe is missing completely (the projection has a *missing ribbon*), or that a pipe appears to have a long straight section within its length (the relevant ribbon has an *incomplete middle*). These conditions are all caused by pipes, created from incorrect correspond-

ences, being forced into an interpretation by the exhaustive search during the generation step.

Tests have been developed which identify each of these conditions and any interpretation which passes the tests is regarded as 'correct'.

Results

The processes have been verified against biplane images generated from four different CSG models. The use of such models has a number of advantages. The shape of the original model is known and the validity of a reconstruction can be verified and measured if required. The CSG model can be accurately rotated to produce orthogonal (90°) images, and can be accurately projected orthographically into those images. This will eliminate errors that occur in genuine angiograms due to camera geometry and image distortion, and will allow accurate measurement of the performance of the processes discussed here.

The assessment of the results is essentially qualitative rather than quantitative, with the verification of plausible reconstructions being performed by eye. If the reconstructions pass a qualitative test then a more accurate, quantitative measurement becomes possible and appropriate. This quantitative measurement has not yet been performed.

The reconstruction of one model is shown. Figure 6 and Figure 7 are the biplane images generated from the CSG model, labelled LEFT and RIGHT respectively. Figure 8 is a plot of the ribbon centre lines that have been identified in the LEFT image. Notice that the network of ribbons is not complete. This does not prevent the completion of a correct three dimensional structure. Figure 9 is a plot of the ribbons found in the RIGHT image. Figure 10 is the Hypothesis Set which has been rotated so that the full set of pipes in the set can be seen. Pipe disconnections can be seen. Figure 11 is the graph, showing that connectivity has been recovered. In this example 8 interpretations were generated and 4 tested as correct. Figure 12 shows the interpretation tested as correct which matches the original model. It has been rotated approximately 45° about the vertical axis. Notice that the connectivity of the pipes has been achieved. Figure 13 shows one of the interpretations that is tested to have an *incomplete end*. The interpretation has been rotated and projected into the RIGHT image. The pipe that has been identified as having an *incomplete end* is marked with a small circle.

Performance and Complexity

The processes described above have been implemented in LISP/VM and run on an IBM 3081 which delivers about 6 MIPS. The volumes of generated data are shown in Figure 4. The process times for each stage are shown in Figure 5.

Ex	Data Quantities			
	Pipes in Hyp. Set	Pipes in Graph	Inter'tions Generated	Inter'tions Tested OK
1	12	11	8	1
2	20	17	8	1
3	32	27	7	3
4	53	39	90	40

Figure 4. The Data Quantities

Ex	Process Times (seconds)			
	Hyp. Set	Graph	Generate	Test
1	8	1	1	1
2	125	2	1	1
3	280	3	134	3
4	103	5	275	22

Figure 5. The Process Times

Analysis shows that the performance of the construction of the Hypothesis Set by exhaustive search is $O(n^2k^2)$, where n is the number of ribbons and k is the number of description points (points defining the linear piece-wise locus of the ribbon centre line) in each ribbon. Building the graph is $O(IP^2)$ where I is the number of iterations of the algorithm to merge the ends of the pipes which are close, and P is the number of pipes in the Hypothesis Set. As can be seen, this process is not computationally expensive. Testing the interpretations is $O(T)$ where T is the number of generated interpretations. Determining the complexity of the process that generates interpretations is currently under further study.

Further Research is Required

There are indications from other examples that, as the complexity of the ribbons increases, the number of generated interpretations increases. For instance, a helix is a simple unbranched structure which can project, in both views, into looping ribbons. The biplane images produce 90 interpretations of which 40 are tested as correct and visual inspection shows that one of the correct interpretations is a helix. A method must be identified which will automatically identify this interpretation.

Lowe(6) has shown that many combinations of features in an image can only occur when there is a coincidence between two or more features and the camera position. Although such coincidences are possible, they are less likely than a combination of features that does not depend upon the coincidence. In this task, if all interpretations with such coincidences are rejected (currently by visual inspection) only the interpretation that matches the original structure is left in the set of correct interpretations. Figure 13 shows an example of a Line of Sight coincidence, which will cause the interpretation to be rejected. Visual inspection of the interpretations for all examples shows that this technique will identify the interpretation which matches the original object in all cases. Further research is neces-

sary to design and implement an algorithm which will identify such coincidences.

Conclusions

In this task the major problem is to solve the correspondence problem. Other researchers have either identified the corresponding points by hand, or have made significant simplifying assumptions about the nature of the original structure and its projections into the biplane views. The correspondence problem has been avoided by building the Hypothesis Set, a set of three dimensional reconstructions (pipes), which contain all possible correspondences. The identification of a plausible reconstruction then involves identifying a suitable subset of pipes in the Hypothesis Set. This has been achieved by transforming the Hypothesis Set into a graph and performing an exhaustive search for all possible reconstructions, applying the constraints available from the pipe models, and from an understanding of the way in which the Hypothesis Set has been constructed. The generated reconstructions (interpretations) are then tested against the original projections to identify those which are acceptable. This Generate and Test process has shown that a limited number of acceptable interpretations are produced.

There have been a number of notable successes. Only one operator action is required. The number of plausible reconstructions is limited. A reconstruction that matches the original object has been achieved in all examples. The processes have demonstrated a reconstruction technique that produces sufficiently encouraging results that further research becomes worth while.

References

1. Barba J, Fenster P and Suardiaz M. 3D Arterial Trace Reconstruction from Biplane Multi-valued Projections. SPIE Vol. 974, Applications of Digital Image Processing XI, 1988.
2. Brady M. Computational Approaches to Image Understanding. Association for Computing Machinery Computing Surveys, Vol 14, No 1, March 1982.
3. Carre B. Graphs and Networks. Clarendon Press, 1979.
4. Gerbrands J J, Reiber J H C, Schols B, Langhout G and Kooijman C J. Structural Analysis of the Coronary Artery Tree. Symposium on Medical Imaging and Image Interpretation. Berlin. October, 1982.
5. Grimson W and Lozano-Perez T. Model-based Recognition and Localisation from Sparse Range and Tactile Data. A.I. Memo 738, Massachusetts Institute of Technology, 1983.
6. Lowe D G. Perceptual Organisation and Visual Recognition. The Kluwer International Series in

Engineering and Computer Science, Kluwer Academic Publishers, 1985

7. Mackworth A. Consistency in Networks of Relations. Artificial Intelligence 8, 1977, pp 77-98.
8. Marr D. Vision. W H Freeman and Co., 1982.
9. Mol C R. 3D Reconstruction from Biplane X-Ray Angiograms. IBM UK Scientific Centre Report 124. IBM Scientific Centre, Winchester. 1984.
10. Nadel B A. Constraint Satisfaction Algorithms. Computational Intelligence, January 1989.
11. Parker D L, Pope D L, White S, Tarbox L R and Marshall H W. Three dimensional Reconstruction of Vascular Beds. In Information Processing in Medical Imaging. Ed S. L. Bacherach. Martinus Nijhoff Publishers, 1986.
12. Rake S T. Finding Curved Ribbons. IBM UK Scientific Centre Report 200, Winchester, England. November 1988.
13. Sun Y and Bergerson E. Automated 3D Reconstruction of Tree-like Structures from Two Orthogonal Views. Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing, Vol 2, pp 1296-1299. New York. April 1988.




Figure 6. The original image, LEFT view: A CSG model is used to generate the image




Figure 7. The original image, RIGHT view: The same CSG model is rotated by 90° and used to generate the image




Figure 7. Ribbon centre lines, LEFT view:




Figure 10. The Graph: The connectivity of the pipes have been recovered.




Figure 8. Ribbon Centre lines, RIGHT view:




Figure 11. The correct interpretation: The interpretation has been projected into the plane of the LEFT image with a rotation of approximately 45° about the vertical axis.




Figure 9. The Hypothesis Set: The 3D structure has been rotated to reveal all the pipes




Figure 12. An incorrect interpretation: This interpretation has been projected into the plane of the RIGHT image. The incomplete end is marked The Line of Sight coincidence can be seen