Reconstruction of Blood Vessel Networks from X-Ray Projections and a Vascular Catalogue

Peter Hall[1], Milton Ngan, and Peter Andreae
Department of Computer Science
PO Box 600
Victoria University of Wellington
Wellington
New Zealand
Email: {peter | milton | pondy}@comp.vuw.ac.nz

Abstract

Reconstruction of blood vessel networks from their x–ray projections is a challenging problem. This is because the correspondence problem must be solved for vessels which appear to twist, turn and overlap. Moreover, vessel networks vary in both branching structure and vessel shape. The extent of these variations is not catalogued in the clinical literature, or elsewhere. We have built a working system that accepts a few x–rays – separated by an angle of about ninety degrees – and reconstructs a three dimensional model of the vessel network. This task is impossible unless a priori information is used; how this information is represented is widely regarded as a key issue. Our representation makes a contribution by building and using a catalogue of anatomy that explicitly accounts for the wide variations in branching structure and shape. It is extensible in the sense that new information can be added to it at any time, and it is task independent in the sense that it can be used in many applications. We demonstrate its use in the problem of reconstructing vasculature from angiograms. Our reconstruction algorithm seeks to explain angiograms in terms of the vascular model. It can reconstruct vessel structures, even when vessels appear highly tangled, are missing, or extra vessels are present. The ability to recover complicated structure is the contribution made by our reconstruction method.

Keywords

Blood vessel networks, reconstruction, representation, x–rays

1 Introduction

Our task is to build a three–dimensional (3D) model of blood vessels in the brain (cerebral vasculature). An angiogram is an image of vasculature; biplane x–ray angiograms are perspective projections separated by an angle of about 90°. For brevity we refer to biplane x–ray angiograms as angiograms. The clinical motivation (see [6]) is to locate any arterio-venous malformations (AVM). We have previously presented an algorithm for automatic segmentation and 3D reconstruction of AVM's [4]. Here we focus on a method for reconstructing the structural elements of vasculature in 3D.

Angiograms separated by an angle of about 4° may be reconstructed using stereopsis. This is because well known techniques exist for solving the correspondence problem. We work with biplane angiograms for two reasons: (1) clinicians routinely acquire them; and (2) reconstructions from them are far more spatially accurate [2]. We face an under-determined problem that is solvable only with *a priori* information.

---

1. Author now at University of Wales Cardiff, Cardiff CF2 3XF, UK. peter@cs.cf.ac.uk

Systems are made individual by the way they represent the *a priori* information. It is widely held that a suitable representation is a key issue in solving the correspondence problem. This is because vasculature exhibit wide inter–individual variances of both branching structure (topology) and individual vessel shape (geometry). Moreover, the extent of these variations are not recorded in the clinical literature or elsewhere.

Previous approaches to the problem have used expert systems. Stansfield [13] provided three levels of rules that enabled coronary arteries (around the heart) to be segmented and identified. The low level rules relate to image processing methods for segementing pieces of vessels as trapeziums. Medium level rules combine these into strands. High level rules recognise strands as particular blood vessels. This type of scheme has been used by others [11, 14]. Delaere *et al.* [1] uses slightly different rules and, in addition, a Boolean function that predicts valid structures as seen in angiograms. The Boolean function comprises statements of the form 'this vessel has this clinical label', and evaluates to TRUE when a valid set of labels is used. We note that all of these systems are characterised by a fixed rule base that is used to describe the appearance of vasculature in angiograms acquired from specific points of view; the rules relate to angiograms and not to vasculature. New rules are needed for each vascular sub–system and for each point of view. Garreau *et al.*[3] provide an exception by modelling vasculature in three–dimensions. They do not appear to make the geometry explicit, although possible variances in topology are accounted for. However, the representation is fixed – so the system cannot "learn" from new instances of vasculature. None of these systems seem able to cope with the complexity of arteries in the brain, the cerebral vasculature, where inter–individual differences are most manifest and the vessels are smaller and more "tortuous" in their paths. All the systems described so far seem able to support only the task of reconstruction from angiograms; they are task–dependent.

Our representation is a fully three–dimensional representation of a collection of 3D models in which both topology and geometry are explicitly represented. It is a Vascular Catalogue (VC) in which each entry is an anatomical model of an individual vascular system. The content of the VC is incrementally updatable, so the VC can "learn" new instances. This is an important property since the full extent of vascular variations are not recorded in the literature. Inter–individual variances in both topology and geometry are made explicit and are maintained throughout the learning process. The VC uses manually reconstructed instances to initially learn the common vasculatures. As the VC learns more of the common structures, the need for the manually reconstructed instances becomes less. Because our VC contains an anatomical model it exhibits a high degree of task–independence. We have used it for simulating x–ray angiograms [8] and combining information from images and text [7]. The VC representation is such that it can be used to reconstruct other types of vasculature, for example coronary arteries.

We describe the VC and how to use it for reconstruction from x–ray angiograms. With regard to that task it offers two advantages: 1. it does not require different rule bases for different sub–systems; 2. it can reconstruct using angiograms from any point of view. The reconstruction process is a search for the model that best explains all input angiograms simultaneously, before recovering 3D information. The VC provides the *a priori* information needed to solve the correspondence problem. We do not consider methods for processing angiograms to extract vessels from their background, nor do we consider reconstruction of geometry once correspondence is solved, both issues are addressed in the literature. We focus on the difficult problem of solving correspondence.

2 The Vascular Catalogue

The VC is constructed from instances, each a valid anatomical model of a vascular system drawn from a library, L. It separates topological information from geometrical information. It stores the information in a compact form that allows easy access and generalisation. The VC contains a labelled graph, $G^+ = (N^+, A^+)$, that we call the *proto–graph*; this is built from L. Individual vasculature are also represented as a graph, $G_i = (N_i, A_i)$. Each instance, G_i, represents the topology of a particular vasculature. In all the graphs, nodes correspond to places where vessels divide (*furcations*) and arcs correspond to *vessel fragments* that connect furcations. Vessels as recognised by clinicians [12] are paths, or possibly small subgraphs, in a graph. The manner in which an individual graph corresponds to a real vasculature is shown in Figure 1.

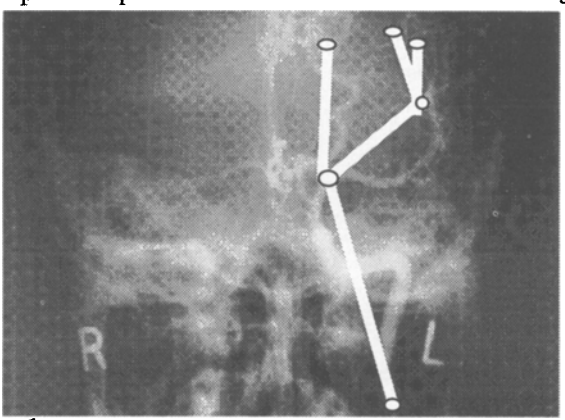

Figure 1. A real angiogram with part of a graph model drawn over it.

The labels of the nodes N^+ and the arcs A^+ contain information relating to topology and geometry, and their variances. These are generalised versions of the labels in each of the N_i and A_i, which contain only geometric information specifying the shape of a furcation or vessel fragment. The set N^+ is formed by merging elements in each of the N_i. Similarly, A^+ is formed by merging elements in each of the A_i. For example, vessel fragments with a sufficiently close geometrical similarity all map into a single element of A^+. Hence the merging operation results in a compact form for the proto–graph since many nodes (or arcs) may map to one. Unfortunately topological information may be destroyed by the merger, how this is avoided is explained next.

2.1 Topology

Topological information about the instance graphs can be maintained at a node (or arc) by a list of the graphs in L that contain the node (or arc). We can pull these labels off the nodes and arcs and use them as terms in a Boolean valued function that we call the *discrimination proposition*. This discriminates subgraphs of G^+ that belong to L and those that do not. The discrimination proposition is most easily represented by a matrix in which each column contains the label that came from a node or arc, and each row is an individual graph in L. A VC contains a proto–graph and a discrimination proposition. An example VC is depicted in Figure 2.

Rows of the matrix specify those topologies that are in L. We can add new instances to the proto–graph and simultaneously maintain the discrimination proposition (see [5] for

details). This means that topological information in the VC can be incrementally updated. Updating the VC requires looking for the maximal common–subgraph (MCS) between the new instance and the current proto–graph. Searching for the MCS is an NP–complete problem in general. We are fortunate in having nodes and arcs that are richly labelled with geometric information, which makes our problem tractable. When searching for the MCS we ignore the discrimination proposition and match against the whole proto–graph. Once a match is found we check the common–subgraph, using the discrimination proposition, to decide whether the new instance is already in the library. If it is not we are able to decide which of its parts come from which instances in the library, and which of its parts are new. These operations are independent of the source of the new instance and is, therefore, independent of reconstruction from x–ray angiograms. However, we can decide which parts (of a new instance) came from which individuals (in the library). This is useful for reconstruction tasks because it provides information needed to generalise from examples.

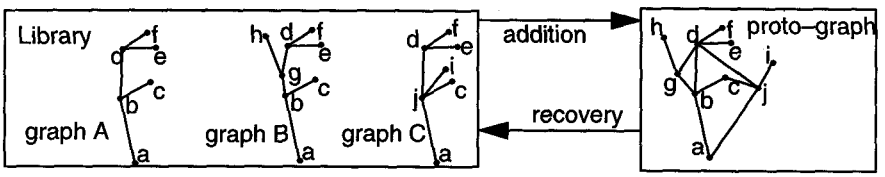

discrimination proposition: columns are labels on proto–graph.

	a	b	c	d	e	f	g	h	i	j	ab	aj	bc	bd	bg	cj	de	df	dg	dj	gh	ij
graph A	1	1	1	1	1	1	0	0	0	0	1	0	1	1	0	0	1	1	0	0	0	0
graph B	1	1	1	1	1	1	1	1	0	0	1	0	1	0	1	0	1	1	1	0	1	0
graph C	1	0	1	1	1	1	0	0	1	1	0	1	0	0	0	1	1	1	0	1	0	1

Figure 2. Example vascular catalogue with discrimination proposition.

2.2 Geometry

Geometric information is held in labels and specifies the 3D shape of the vasculature. The shape of each furcation is represented by a list of the direction of the outgoing vessel fragments relative to the incoming parent. Vessel fragments dominate the geometry of a vasculature and so we concentrate upon their description. Each vessel fragments a generalised cone whose medial–axis is a piecewise cubic curve and whose cross–section is circular. The vessel fragments in a set $V = \{ v_i : i = 1...n \}$ each correspond to a particular real vessel and all map to the same element of A^+. Each vessel fragment, v_i, is represented as a sequence of point, tangent, radius triplets; $v_i = < (\mathbf{p}_{ij}, \mathbf{t}_{ij}, r_{ij}) : j = 0...M >$, M is constant over all elements of V. The points, $\mathbf{p}_{ij}$, are separated by a distance D/M, along the medial axis; D is the length of the medial axis. Thus we obtain M+1 triplets, with the *jth* at position jD/M. To describe variances in geometry, we assume: 1. the sequence of triplets in each v_i is consistently ordered, and 2. each component in the point, tangent, and radius triplet is an independent random variable. Triplets of equal index, j, correspond by definition and are generalised by computing mean and normal distribution of each component. In this way we produce a generalised vessel–fragment for the set V. Before computing any statistical measure, we translate each vessel fragment so that its first point lies at the origin and its principal axes align to a standard frame.

The labels in elements of A^+ are generalised vessel–fragments. The representation is incrementally updatable. An important property of the generalised vessel fragment is that it includes vessel instances not yet seen: as in the topological case, we have generalised from examples. The representation can be projected onto a viewplane, which is a useful property for reconstruction, which is discussed next.

3 Reconstruction

The reconstruction module accepts angiograms as input and produces as output a reconstruction that accounts for the input angiograms. The role of the VC within the reconstruction process is to provide a model that is used to determine correspondence between angiograms. For example, if a furcation in one angiogram and a furcation in another angiogram have each been identified as corresponding to the same furcation in the proto-graph, then we can deduce that they correspond to each other. Correspondences such as this are a prerequisite to recovery of a vascular model in 3D. The description below concentrates on the central issue of how the VC is used to find such correspondences.

During reconstruction it is helpful to represent each angiogram as a graph that we call an *α–graph*. The α–graphs are obtained from the angiograms via segmentation. Angiograms are usually acquired from more than one point of view and we write $\alpha_i = (X_i, S_i)$ for the *ith* angiogram. In such a graph, each node in X_i corresponds either to a furcation in the vasculature or to a *crossing*. Crossings are artefacts of projection created by the apparent intersection of vessels. The arcs in the graph (elements of S_i) represent fractions of a vessel segment (*slivers*) that connect furcations and crossings as seen in an angiogram. This construction is shown in Figure 3. The set of correspondences between the α–graphs and the proto-graph generate a subgraph in G^+. This subgraph can be tested to determine whether it belongs to the vascular library, L, an operation discussed previously.

Reconstruction is now a search for the maximal, consistent set of mappings from nodes in the VC to nodes and arcs in each α–graph, and from arcs in the VC to nodes and arcs in each α–graph. A node in the VC will usually correspond to a node in the α–graph. Under certain circumstances, such as the node being hidden behind a vessel fragment, the node may correspond to a localised region of a sliver. An arc in the VC may map to a sequence of nodes and arcs (*sliver path*) in an α-graph because vessel fragments (arcs in the VC) may appear to cross in angiograms. In addition, more than one vessel fragment may account for a single sliver. This may occur in situations where the model has more pieces than the vasculature being reconstructed. There are a large number of possible ways correspondences can occur. We have found it is useful to record these correspondences in a *match space*, as explained next.

3.1 Match Space

The set of consistent, maximal mappings from G^+ to the set of n α–graphs, $\{\alpha_i\}$, exists in a match space with dimension $|\{\alpha_i\}|+1$. One axis of this match space comprises all elements in the proto–graph, $M^+ = N^+ \cup A^+ \cup \{\emptyset\}$; $\emptyset$ is the null element. The other axes each comprise elements in a particular α–graph, $M_i = X_i \cup S_i \cup \{\emptyset\}$. The order of these elements is of no consequence. Points within the volume of match space represent assertions of the form $CR(h, h_1, \ldots, h_n)$ where $h \in M^+$ and each $h_i \in M_i$. We assign these points a value in the range [0,1] to represent the degree to which the correspondence has been validated. The value -1 is used to assert nothing is known

about the correspondence. The null element is required on each axis so that we can explicitly represent those cases where an element has no correspondence. Of most interest are the planes of match space that are generated by pairs of axes – we must find correspondences on these planes. In particular, the correspondences in the planes defined by pairs of distinct axes M_i, M_j are of interest because these are needed to recover geometry in 3D by back–projection. The plane defined by M^+ and one of the M_i records correspondences between the proto–graph and the i*th* α-graph. Pairs of such planes are used to deduce the correspondences on the M_i, M_j plane.

To make these remarks a little more concrete, consider the example shown in Figure 3, which is a mock up of the experiments reported later in this paper. A model (the proto–graph without the vessel fragment BF), has been projected onto a biplane angiogram and α–graphs formed for each. We have supposed that the correspondences between the proto–graph and each of the α–graphs have been found – typically by matching the projected proto–graph (including the vessel–fragment BF). This projection is not shown on the diagram which is intended to be an illustration of a biplane angiogram. We have also supposed that B projects to somewhere in sliver 12 in α–graph$_1$ and to somewhere in sliver bb in α–graph$_2$; this helped fix the correspondences. Three planes of match space can be seen. Two planes record correspondences between the proto–graph and an α-graph. The third plane records correspondences that were deduced between the α-graphs by using the other two planes.

The example shows simple correspondences, such as CR(A, 1), on the plane defined by the proto–graph and α-graph$_1$. A situation in which a node of the proto–graph may correspond to a localised region on a sliver is also shown in the example: node C is hidden in α–graph$_2$, which we record as CR(C, de). Multiple matches can be recorded; CR(AB, 12), CR(B, 12), and CR(BD, 12) in the proto–graph / α–graph$_1$ plane shows how many elements from the proto–graph account for a single sliver. In the same plane we see CR(CD, 23_2), CR(CD, 3), and CR(CD, 35), showing that vessel fragment CD accounts for many elements in the α-graph. For brevity we shall indulge in abuse of notation and write CR(AB–B–BD, 12) or CR(CD, 23_2–3–35) for multiple correspondences. This notation has the advantage of clearly denoting paths in the proto–graph (which may be vessels as recognised by clinicians) and sliver paths. Once the set of correspondences has been constructed from the proto–graph to each α-graph, we can deduce how elements of the α-graphs correspond. Simple cases exist: CR(A, a) and CR(A, 1) lead us to conclude CR(a,1), for example. A more complex example is CR(CD, cd_1–d–de) which pairs with CR(CD, 23_2–3–35). We write these corresponding sliver paths as CR(cd_1–d–de, 23_2–3–35). In the plane defined by the α–graphs this is recorded in each point of the Cartesian product $\{23_2, 3, 35\}\times\{cd_1, d, de\}$. In Figure 3 such Cartesian products are highlighted by connecting the elements; without further geometric constraints this is the best we can do. We can deduce CR(cd_2–d–de, 23_1–3–34), which has elements in common with the previous example – this is a result of the physical overlap. The most complex example arises from the set of correspondences CR(AB, ab–bb–bb), CR(B, bb), CR(BB, bb–b–bc), and CR(AB-B-BD, 12); which can be used to deduce CR(12, ab-b-bb-b-bc). Match space contains the correspondences we desire and also allows us to build structures within it that limit the search and record its passage. The example demonstrates the role of the VC in solving the correspondence problem. However, determining the correspondences between the model and an α–graph has not been discussed so far. We address that issue next.

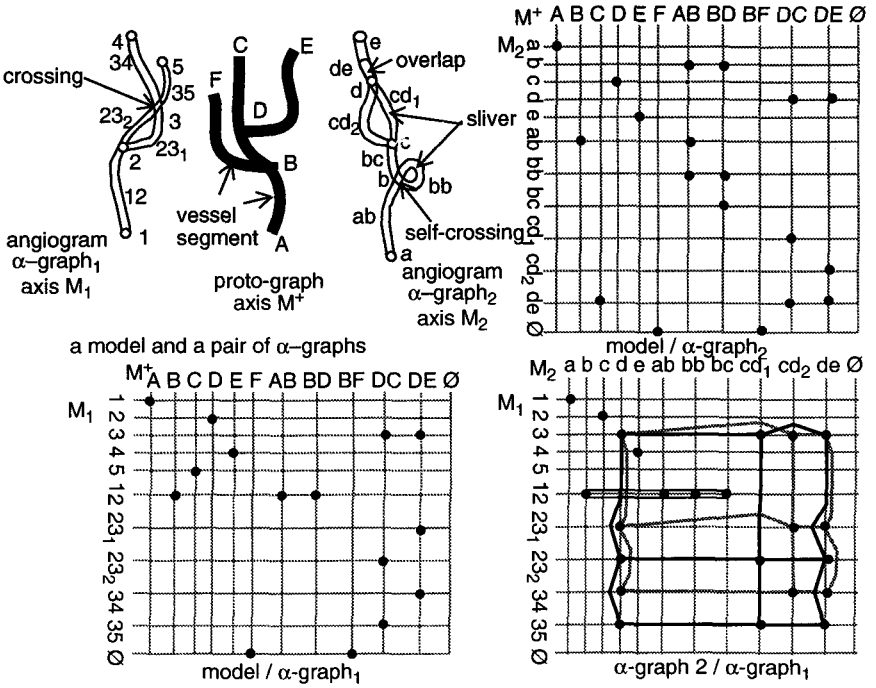

Figure 3. A model, its α-graphs, and three planes in match space

3.2 Search Algorithm

To find correspondences between the proto–graph and an α–graph we traverse the proto-graph in breadth-first order, accounting for the largest most important vessels first. By choosing such a strategy we are taking advantage of our knowledge of vascular structures. The search is initiated by manual selection of a node–arc–node sequence in the proto–graph. We assume the model is properly aligned with the angiograms and proceed by searching for correspondences in the α–graphs as follows:

1. Traverse the proto–graph, pulling off node pairs and connecting arc. This provides us with a connected sequence of elements from M^+, and sliver paths we suppose must exhibit similar connectivity. The first furcation in this sequence is usually assumed to match nodes in each of the α–graphs. In the proto–graph and each α–graph we call this the current node. We discuss the case where we relax this assumption below.
2. Hypothesise that the vessel fragment, h, exists as a sliver path in each of the α–graphs. We test this by searching in each of the α–graphs; the search is initiated with a set of mutually independent slivers. Each such sliver is a potential start for the sliver path and must be connected to the current node. The sliver path is constructed by building paths from each starting sliver, and pruning when some geometric tolerance is exceeded; beam search is used. Geometric matching is explained below.
3. The result of step 2 is a sequence CR(h, h_{ij} ...h_{ik}) for the ith α-graph. If such a

exists for every α-graph, the correspondences are recorded in match space. If more than one such sequence exists for all α-graphs then only best is retained; we use a best–first search for the subgraph that explains the angiograms. The current node is then moved to the end of the vessel fragment and traversal of the proto–graph continues.

As mentioned, we may relax our assumption that search continues from the current node. This is so we can "jump" sections of vasculature that are missing or which are otherwise unexplained by the model. If at some point in the search all geometrical matches are poor, then the model cannot satisfactorily explain the angiograms at that point. We proceed by supposing a correspondence between particular candidates in the VC and elements of the α–graphs. This is a useful property for it tends to maximise both the fraction of the model used in an explanation and the amount of angiogram that is explained. Moreover, the supposition can be used as a basis to update the VC.

3.3 Geometric Matching

Geometrical matching treats generalised vessel–fragments as straight lines by using the chord between the first and last point in the triplet sequence. We chose this geometrical match because: 1. at this stage of system development, we are more concerned with recovering the structure of vasculature than we are with recovering specific vessel geometry; and 2. clinical texts indicate that where a vessel ends up (the territory it feeds) is more important than specific geometry. The geometrical match is defined in two parts. 1. Matching a sliver against a projected vessel–fragment. This is used in stage 2 of the search algorithm explained above. 2. Accumulate of such matches. This is in line with stage 3 of our search. An individual sliver is matched to a vessel–fragment using

$$m = \begin{cases} \left(\frac{|\mathbf{S}|}{|\mathbf{F}|}\right)\cos\theta + \left(1 - \frac{|\mathbf{F}_1 - \mathbf{S}_1|}{r}\right) & \text{if } |\mathbf{S}| > |\mathbf{F}| \\ \left(2 - \frac{|\mathbf{S}|}{|\mathbf{F}|}\right)\cos\theta + \left(1 - \frac{|\mathbf{F}_1 - \mathbf{S}_1|}{r}\right) & \text{otherwise} \end{cases}$$

in which $\mathbf{F}_1$ is a point on the projected vessel–fragment; $\mathbf{F}_2$ is the end–point of the vessel–fragment; $\mathbf{F} = \mathbf{F}_2 - \mathbf{F}_1$; $\mathbf{S}_1$ is the starting point of the sliver, and $\mathbf{S}$ is the sliver chord. $\cos\theta$ is the dot product of unit vectors $\mathbf{S}/|\mathbf{S}|$ and $\mathbf{F}/|\mathbf{F}|$. r is a positive constant which measures vessel width. Initially, $\mathbf{F}_1$ lies at the start of the vessel–fragment and moves over it as the sliver and chord are matched. The measure is designed to penalise large slivers should they turn away from the expected direction, and any sliver that is too distant from the vessel–fragment. The positive constant, r, penalises distance. We form a complete measure of how N such slivers, $\mathbf{S}_i$, match a vessel–fragment by a weighted sum of these matches. Despite the crudeness of this measure, our initial results have been highly satisfactory and vindicate our approach. In particular, we have been able to recover complicated structures, even when vessels appear to be highly tangled, when vessels are missing, or when extra vessels are present.

4 Results

We have tested our method with biplane angiograms produced by our simulator [8]. This uses the same vascular catalogue as the reconstruction. Our test method was as follows:

1. Extract a model from the VC and modify it by: (a) by adding new pieces; (b) removing old pieces.

2. Use the simulator to produce both x–ray angiograms and an equivalent α–graph for each. In this we assume that the problem of segmenting slivers and crossings has been solved; an assumption justified by reference to the literature that performs a similar task [9, 10].
3. Reconstruct a vascular model using the α–graphs and the VC. Geometric form can be "borrowed" from the VC rather than recovered by back projection. This is because, here, we are more concerned with recapturing structure. We are developing mechanisms to recover geometry from the correspondences found in the match space.
4. Compare the reconstructed vascular model to the original from stage 1

Figure 4 shows an x–ray simulation of complete cerebral vasculature; this is our test VC. The α-graph representation and the reconstructed model are also shown in the same figure. This example shows we can reconstruct complicated structures. Notice that the reconstruction process has been able to "unravel" highly tangled blood vessels, including vessels that appear to overlap along their entire length from some points of view. Our comparison of the reconstructed model to the original showed the structures to be isomorphic. We can reconstruct sub–vasculature too; the image in Figure 5 shows a simulated x–ray and reconstruction of the left carotid system, which is a sub–vasculature of that in Figure 4. We found a perfect structural match between the reconstructed model and the original. This experiment was repeated with the exception that pieces of vasculature were removed. The simulated x–ray of the original, and the reconstructed model, are shown in Figure 6. The experiment was repeated again, this time with pieces added. The simulated x–ray of the original, and the reconstructed model, are again shown in Figure 6. Each time we found a perfect structural match between the original and the reconstructed model. The only reason the extra vessel does not appear in the reconstructed model is that, currently, we "borrow" reconstructed geometry from the model – and are unable to do so for pieces not in the VC.

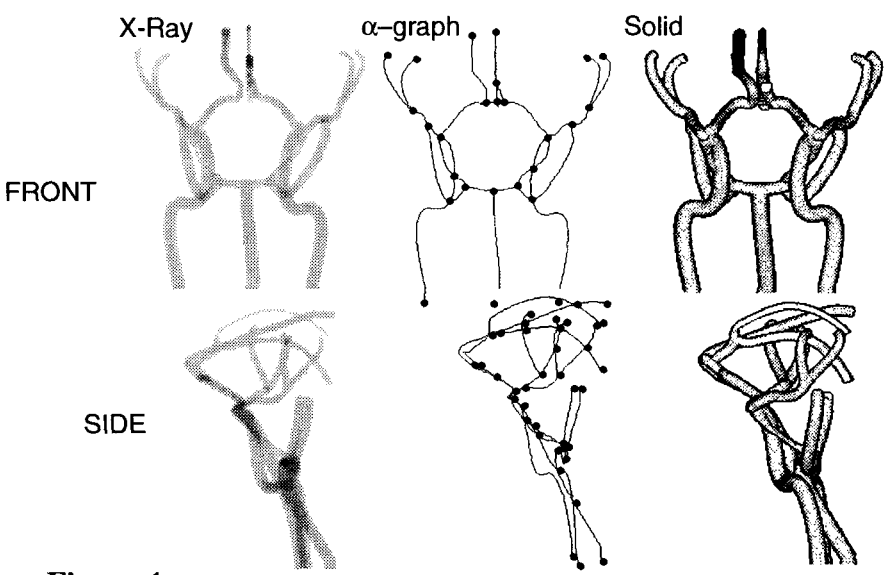

Figure 4. A complete VC as x–ray, α–graph, and reconstructed model.

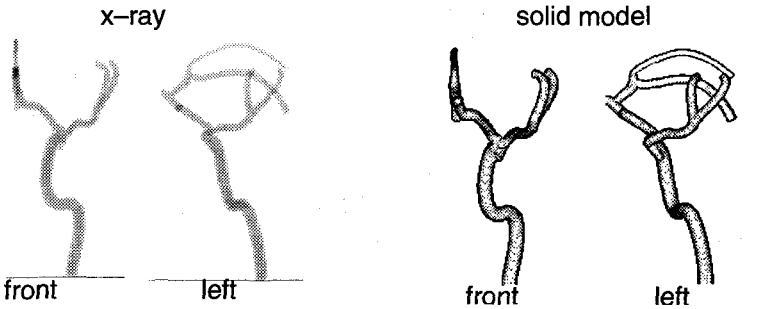

Figure 5. X-ray and reconstructed model of left carotid sub–vasculature.

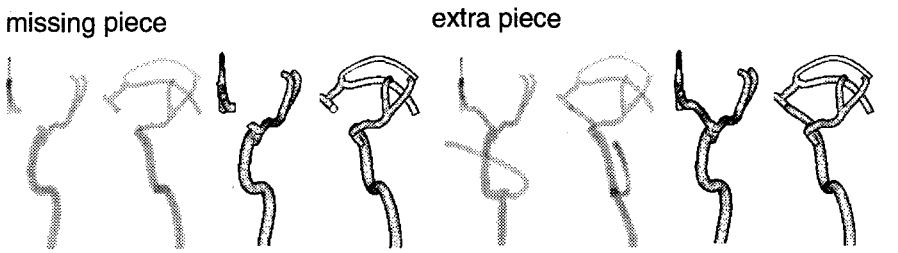

Figure 6. Reconstruction with a piece missing and with an extra piece.

5 Conclusion

We have introduced a Vascular Catalogue to model a collection of vasculature. It offers a number of improvements over previous systems:

- Structural and geometrical variations are made explicit.
- It is task–independent.
- The information it contains can be incrementally updated.
- It generalises from examples to create possibly unseen topologies and geometries.

In addition, it is feasible that our VC could be used in many contexts other than reconstruction from x–rays. For example, clinicians use blood vessels as a map of regions of the brain. We suggest that our VC provides a suitable ground structure from which a catalogue of brain anatomies could be constructed.

Because our VC describes anatomy it can be used to reconstruct angiograms from any point of view and there is no restriction on their number. We have shown its use in recovering vascular structure from simulated biplane angiograms. Our reconstruction method is able to cope with:

- Complex structures that appear to overlap.
- Complete substructures.
- Substructures with missing or additional parts.

Our geometric match is crude but surprisingly effective. This is an important result for it indicates specific geometry may not be as important as structural features. Future work is aimed at: (1) predicting structural features; (2) developing a more sophisticated geometrical match; and (3) segmenting the angiograms rather than assuming it.

6 References

[1] D. Delaere, C. Smets, P. Suetens, and G. Marchal, "A knowledge-based system for the angiographic reconstruction of blood vessels from two angiographic projections," In *Proc. of the North Sea Conference on Biomedical Engineering,* Nov. 1990.

[2] D. Delaere. (personal communication).

[3] M. Garreau, J.L. Coatrieux, R. Collerec, and C. Chardenon, "A knowledge based approach for reconstruction and labelling of vascular networks from biplane angiographic projections," *IEEE Trans Med Imag*, 10(2), 1991, pp 122-131.

[4] P.M. Hall, J. Brady, A.H. Watt, L. Walton, and U. Bergvall, "Segmenting and reconstruction vascular lesions from biplane angiogram projections," In *Proc. Digital Image Computing Techniques and Applications*, Sydney, Australia, 1993, pp 802-809.

[5] P.M. Hall, and J.J. McGregor, "A graph based model of a collection of physical vasculature." *Proc. Digital Image Computing Techniques and Applications,* 1993, Sydney, Australia, pp 414-421.

[6] P.M. Hall, R. Feltham, and T. Fitzjohn, "Automated analysis of x-ray angiograms," In *Proc. New Zealand Computer Society Conference*, Wellington, New Zealand, Aug. 1995, pp 253–260.

[7] P.M. Hall, and P. Mc Kevitt, "Automatic Interpretation of Vasculature," *AI Review*, 10(3-4), 1996, (forthcoming).

[8] P.M. Hall, "Simulating angiography," *Mathematical Modelling and Scientific Computing*, 6, 1996, (forthcoming).

[9] K.R. Hoffmann, D. Kunio, H.P. Chan, L. Fencil, H. Fujita, and A. Muraki, "Computer reproduction of the vasculature using an automated tracking method," In *Proc. SPIE 767 Medical Imaging*, 1987, pp 449–453.

[10] T.V. Nguyen, and J. Sklansky, "Reconstructing the 3–D medial axes of coronary arteries in single–view cineangiograms," IEEE Trans. Med. Imag. 13(1), March 1994, pp 61–73

[11] S.T. Rake, and L.D.R. Smith, "The interpretation of x-ray angiograms using a blackboard control architecture," *Proc. Computer Assisted Radiology* 1987, pp 681-685.

[12] G. Salamon, and Y.P. Huang, "*A radiological anatomy of the brain*," Springer-Verlag, Berlin. 1976.

[13] S.A. Stansfield, "ANGY: A rule base expert system for automatic segmentation of coronary vessels from digital subtracted angiograms," *IEEE Transactions on Pattern Analysis and Machine Intelligence*. 8(2), 1986, pp 188-199.

[14] P. Suetens, J. van Cleynenbreugel, F. Fierens, C. Smets, A. Oosterlink, and G. Wilms, "An expert system for blood vessel segmentation on subtraction angiogram," *SPIE 767 Medical Imaging*, 1987, pp 454-459.