

An Adaptive Minimal Path Generation Technique for Vessel Tracking in CTA/CE-MRA Volume Images

Brian B. Avants, James P. Williams

Imaging & Visualization Department - Siemens Corporate Research
755 College Road East, Princeton, NJ 08540
avants, jwilliams@scr.siemens.com

Abstract. We present an efficient method for the segmentation and axis extraction of vessels and other curvilinear structures in volumetric medical images. The image is treated as a graph from which the user selects seed points to be connected via 1-dimensional paths. A variant of Dijkstra's algorithm both grows the segmenting surface from initial seeds and connects them with a minimal path computation. The technique is local and does not require examination or pre-processing of the entire volume. The surface propagation is controlled by iterative computation of border probabilities. As expanding regions meet, the statistics collected during propagation are passed to an active minimal-path generation module which links the associating points through the vessel tree. We provide a probabilistic basis for the volume search and path-finding speed functions and then apply the algorithm to phantom and real data sets. This work focuses on the contrast-enhanced magnetic resonance angiography (CE-MRA) and computed tomography angiography (CTA) domains, although the framework is adaptable for other purposes.

1 Introduction

The domain for this investigation is the volumetric images produced by contrast-enhanced magnetic resonance angiography (CE-MRA) and computed tomography angiography (CTA.) In the CE-MRA imaging protocol, a contrast agent, usually based on the rare-earth element Gadolinium (Gd) (a highly paramagnetic substance), is injected into the bloodstream. In such images, blood vessels and organs perfused with the contrast agent appear substantially brighter than surrounding tissues. In CTA, a contrast agent is injected which increases the radio-opacity of the blood making the vessels appear dense.

The goal of the majority of CTA/CE-MRA examinations is diagnosis and qualitative or quantitative assessment of pathology in the circulatory system. The most common pathologies are aneurysm and stenosis caused by arterial plaques.

The modern clinical workflow for the reading of these images increasingly involves interactive 3D visualization methods, such as volume rendering, for quickly pinpointing the location of the pathology. Once the location of the pathology is determined, quantitative measurements can be made on the original 2D slice data or, more commonly, on 2D multi planar reformat (MPR) images produced at user-selected positions and orientations in the volume.

In the quantification of stenosis, it is desirable to produce a cross-sectional area/radius profile of a vessel so that one can compare pathological regions to patent (healthy) regions of the same vessel. In order to accurately measure cross-sectional area or radius a definition of the cross-sectional plane is required.

Segmentation research in this area has focused in recent years with the application of curve propagation methods (such as snake and active contours), region competition and mathematical morphology to analysis of the vessel tree [4,8]. Sethian [5,6] provides a fast marching method for vessel tree segmentation that requires pre-processing of the image. Kimia [7] and others have used balloons. Morphological methods, combined with the eikonal approximation, have also proven fruitful [3]. However, our approach uses segmentation simply as a learning procedure that provides both a metric and bounds for an explicit vessel centerline search. Vessel axes encode the vessel tree more sparsely. Our method makes the extraction of such abstractions sufficiently fast to make axis creation and manipulation part of a real-time, interactive clinical workflow.

The foundation given here provides a clinical framework that is related to the above work but novel in its design and in its efficiency. No pre-processing of the images is required. The interface is semi-automated in which the image reader places seed points within the vessel tree. A max-probability surface propagation algorithm then makes a dynamic segmentation of the vessels containing the seeds. Parameters are also set for the subsequent minimal path calculation used to approximate the vessel centerlines. This piece-wise linear approximation can then be used as a visual handle to the vessel region, as a navigation tool for automatic placement of MPR's orthogonal to the vessel and also as a starting point for further segmentation or automated quantification algorithms.

2 Methods: Segmentation Module

2.1 Graph Based Approach

In our application, we view the image volume as a k -regular graph (where ' k ' denotes the number of edges at each node/voxel) through which we would like to compute a small set of the minimal connectivity. The user contributes the location of nodes to be connected through the image volume. The algorithm then consists of two parts: a volume search and a path search. The first part searches for the subset of the full image graph that is most likely to contain the paths between the user's seed points. The associated volumes are also given a statistical characterization. The edge cost function for the second phase is determined from the statistical distribution of the voxels explored during the first phase. The second phase computes the minimal connectivity of the selected vessel tree points according to parameters set by the partial segmentation results.

The constraint of the minimal path computation to only those regions that contain the objects to be connected is an important part of the efficiency and correctness of the algorithm if it is to operate on volumetric data. If unnecessary (non-object) regions of the image are searched, not only does the execution time rapidly increase

(as the algorithm is $O(N \log N)$), but the chance that the path will mistakenly cross a object boundary increases. It is imperative, then, to restrain the search to only those regions in which it is probable to have a true path.

2.2 Surface Expansion

Our approach to the segmentation aspect of this problem is motivated by Sethian's fast marching method. He observes in [6] that the fast marching method, which is a level set surface propagation tool, is fundamentally the same as Dijkstra's algorithm in a regular graph. The former, however, computes velocities (inversely proportional to edge costs) incrementally according to the local solution to the eikonal partial differential equation (pde). The latter is defined, usually, over a simple network with fixed edge costs.

Sethian describes his method as producing a gradient field for the arrival time of the front such that:

$$|\nabla T| = \frac{1}{F} \quad (1)$$

In his application of the fast marching method on medical images [2], the velocity function is given by:

$$F = \frac{1 - \varepsilon \kappa}{(1 + \nabla G * I)} \quad (2)$$

The first term is an expansion term; the second term slows expansion where curvature, κ , is large. ∇G is the gradient of the image pre-processed with a gaussian filter. Time is zero at all initialized source points and arrival time is calculated such that $T > 0$ at all successive iterations. The proof of the algorithm's correctness is given in [5].

We incorporate a similar approach to Sethian's, but do not focus entirely on segmentation. The surface propagation is simultaneously used to learn an appropriate data based cost function for calculating the minimal path. Also, in our case, the velocity model is probabilistic and non-viscous thereby allowing velocity to be determined by data at a particular voxel rather than the voxel's neighborhood. This probabilistic model effectively learns the intensity distribution of the segmented region and is subsequently used to generate a data-based closed form for the edge cost function.

We redefine the speed function such that we are examining the probability of the local data being included in a known set. Then, the speed function for each node is determined by maximizing the probability of advancement in each possible direction of new exploration. We sample the data locally according to Dijkstra's method (analogous to the eikonal model of front advancement, if connectivity is k-regular) and then define speed to be:

$$F_{ijk} = V * P(\text{data included in the known set}) \quad (3)$$

If V is 1 (expansive), then F_{ijk} can be written:

$$F_{ijk} = 1 - P(\text{data not included in the known set}) \quad (4)$$

and we see the parallel between this and Sethian's purely geometric $F = 1 - \epsilon\kappa$. The change in time is then given by:

$$\frac{dT}{ds} = \frac{1}{\max(F_{ijk}(P))} \quad (5)$$

Initially the front will expand most rapidly in the most probable directions. After long times, however, lower probability regions may be explored. This fits with the intuitive notion that probable events occur often and improbable ones only at large intervals.

If the probability calculation varies only according to data at each voxel, and the exploration will execute in $O(N \log N)$. If probability changes according to the delivery from the min-heap, the entire boundary should be recomputed (as it must be with pure level-set methods.) This guarantees the maximum probability exploration, given the elapse of time, but reduces efficiency with a worst-case bound of N^2 .

Instead, we use a bounded-error approach where recomputation is performed only when potential maximum error exceeds a given threshold. This generates an error-bounded approximation of the maximum probability exploration with relatively little loss in efficiency.

2.3 Edge Cost / Velocity Model

The probability model given for the data determines our speed function. In vessel path finding for CTA/CE-MRA we are interested in quickly connecting vessel-like regions without crossing neighboring non-vessel regions. Stated another way, we are interested in discrimination between energy levels, vessel membership being analogous to a particular energy level. Therefore, we want rapid propagation through regions of similar energy and slow movement elsewhere.

2.3.1 Partitioning Distributions

We formulate the vessel search in terms of a thermodynamic partitioning problem. Sections of the image are allowed two possibilities, vessel and non-vessel, where the existence of one state excludes the other. Each image region, whether it is a single voxel or collection of voxels, is then treated as a canonical ensemble. The comparison of two regions is performed with the partition function [1] given by:

$$Q = \sum_{all-states} \exp (E_i / kT) \quad (6)$$

Where each state is described by its energy E and all states share the constants kT .

Equilibrium requires that, given a system, we have a time independent minimum of information to describe it. That information for equilibrium systems (and our image) is given by N , the number of particles (proportional to the number of voxels), volume V (the sum of voxel volumes) and temperature T , which we can set. We assume that, from the macroscopic perspective, we can also measure the system's average energy.

2.3.2 The Fermi-Dirac Distribution

These assumptions lead directly to the Fermi-Dirac probability function for two states [1]:

$$P_{FD} = \frac{1}{1 + \exp\left(\frac{E_1 - E_0}{kT}\right)} \quad (7)$$

P_{FD} is a step function with slope determined by the size of kT . The inflection point occurs where E_1 (the ensemble temperature) equals E_0 , beyond which probability approaches one.

Therefore, the kT term determines the parsimony in the allocation of probability. Increasing kT makes the slope gentle, supplying slowly changing probabilities to a broad range of energies about E_1 . As the temperature approaches absolute zero, the function approaches zero slope. Higher temperatures result in steep vertical slopes.

2.3.3 Implementation


In our framework, given fixed temperature, we can examine the change of an equilibrium system surrounded by a narrow border of ensembles that may or not contain the vessel. We choose a low temperature (generally, such that kT is within an order of magnitude of unity) to make propagation through high probability energy regions fast, and low probability regions slow.

Initialization occurs when the user selects seeds. Each seed is considered a vessel containing system for which we have N , V , T and the energy,

$$E_1 = \frac{1}{N} \sum_{ensemble\ members}^N E_i(I) \quad (8)$$

which is the standard averaging procedure, where $N=1$ and $E(I)$ is the intensity. These points are treated as ensembles as stated above. For each source, surrounding regions are explored and compared point-by-point to the expanding ensemble $< E >$ with which they are in contact. The border point with the highest probability of being vessel tissue will have the smallest arrival time. That voxel is added to its source

ensemble such that the latter becomes a system of size $N+1$. $\langle E \rangle$ is then re-measured and the process repeats.



Figs. 1a (left),b (right). Slices from a segmented 3-D carotid CE-MRA data set illustrate the region explored and segmented by phase 1 of the algorithm. The explored region is highlighted by a white boundary.

2.4 Approximation Methods

We could simplify our calculation by weighting the source points with large N values and giving all border systems $N = 1$ such that adding a new voxel won't appreciably change E_1 , the mean of the ensemble values. However, for calculating the average and standard deviation of the vessel intensity (needed for the path finding parameters), we would like more exact discrimination. We therefore may be forced to recompute border probabilities when we re-measure E_1 . This increases the order by a factor of N . However, we can use the monotonicity of the p.d.f. as well as the relative stability of $\langle E \rangle$ over time to limit the frequency of re-computation. The derivative of P_{FD} is:

$$\frac{d}{dE} P_{FD} = \frac{-m}{(m + \exp(E_0 - E_1))(m + \exp(E_0 - E_1))} \quad (9)$$

where $m = 1/kT$. We first note that the derivative of our p.d.f. approaches the delta function if temperature decreases. As temperature increases, the derivative spreads and its magnitude decreases. Generally speaking, the error is only appreciable along

the region of the p.d.f where there is a significant slope. For the low temperature fermi function, this region is narrow.

We can illustrate the nature of the measurement error as follows

$$P_{FD}(E_{+error}) = \frac{1}{1 + \exp(E_1(1 + \varepsilon) - E_0)} \quad (10)$$

We see that our error can be thought of as occurring on the measurement of E_1 . We can, then, limit our error concern to only those values of E_1 where the error might be appreciable. We expect to encounter three cases:

- 1) $E_0 > (E_1 + |\Delta E|)$, error rapidly diminishing as E_0 increases
- 2) $(E_1 - |\Delta E|) < E_0 < (E_1 + |\Delta E|)$, error significant
- 3) $E_0 < (E_1 - |\Delta E|)$, error rapidly diminishing as E_0 decreases

We store, for each E_0 probability computed, the value of E_1 against which it was compared. Then, at the time of delivery, we compute

$$Error \approx P_{FD}(E_0 | E_{1old}) - P_{FD}(E_0 | E_{1new}) \quad (11)$$

If this value is greater than χ , the error tolerance, we re-compute the probability and re-insert the voxel. This process is repeated for all delivered nodes with error greater than the tolerance. Generally, we apply a 5% rule to case 1 and 2 errors and 10 % to case 3 errors that could benefit from updating. Fortunately, not only is the range that we need to recompute small, but the value $\langle E \rangle$ stabilizes over time. In practice this means that the vast majority of border recomputations occur early in the expansion process when the boundaries contain relatively few nodes.

3 Methods: Path Finding Module

3.1 Path Finding Speed Function

The speed function for the path finding derives naturally from our former statements. Here, we do not have to choose between two states. The only state we are concerned with is the vessel state. Due to our surface exploration, then, we know the degeneracy (histogram) of energy levels in the vessel. So, we use that knowledge along with the canonical distribution to compute our speed function:

$$F_{ijk} = L * W_j * P_{canonical}(E_j) = L * W_j * \exp(-E_j / kT) \quad (12)$$

Where W_j is the degeneracy of energy E_j and L is an euclidean term that depends on the distance in patient space. The advantage that we have here is that we know all energies locally and can normalize the computation by the 'known' natural temperature of the distribution. It is proportional to the standard deviation of the intensities in the region, which we measured in our surface propagation algorithm. This guarantees that we have a provably minimal path that will move through the vessel region with a very high probability.

Even if there is error in the segmentation, the use of the canonical distribution for the speed function makes it highly unlikely the actual minimal path will deviate into




Fig. 2 (left), Fig. 3 (center), and Fig. 4 (left) Volume-rendered MR angiography data. Renal arteries have been tracked by the selection of two seeds, one near each kidney. The segmented search region and extracted path are shown.

the erroneously segmented parts of the volume. This is because the canonical distribution peaks around $\langle E \rangle$ due to the large degeneracy of that energy level. [1]

3.2 Active Backtracking and System Equilibration

Although we instantiate a min-heap for each source, we compute the minimum over all heaps for the expansion of exploration. In this way, we have a virtual single min-heap. The movement of all surfaces, then, is governed by the interaction of the set of heaps. This, naturally, leads to surfaces in more homogeneous regions of the image propagating more rapidly and thus initiating meetings with other surfaces first.

Surface meetings occur when different surfaces share osculating border points. This is the condition for producing a path. We set the parameters for the path finding by measuring the degeneracy of the energy levels within each surface and computing the (approximate) standard deviation of the energies. An instance of Dijkstra's algorithm with the above speed function is then initiated from the osculating points backward into each surface thus producing a minimal path as it goes. The stopping condition for this search is the discovery of either an original source or a minimal path computed at a previous meeting.

After the paths are computed, the two osculating regions then share statistical information and merge into a single heap. Then, the algorithm continues as before, but with one less surface. The final termination condition is that all surfaces have met and all paths connecting them have been computed.

The data structures containing the vessel boundaries, paths and sources are maintained. This allows subsequent sources to be added to the image volume.

The algorithm will then begin a new set of surface searches and meetings will be handled in the same way as before, except that if the old surface is met, active backtracking is not symmetric. Instead, the new surface is allowed to find paths both backward into its volume and forward towards the paths computed in prior executions of the algorithm.



Fig. 5 (left) Volume-rendered CTA of the coronary arteries and 3 seed points. **Fig 6 (center),** Segmented exploration region **Fig 7 (right)** Axes obtained from the search.

4 Performance Evaluation

4.1 Narrow Vessels – Coronary CTA, Renal MRA

The segmentation of the narrow vessels is very rapid and the path is smooth and remains within the boundaries of the volume-rendered vasculature. Computation times for the renal paths are about 1-2 seconds on a 733MHz Pentium III machine. These times vary with the spatial resolution of the images.

4.2 Large Vessels – Carotid/Upper Aorta CE-MRA

Larger vessels with a larger enhanced region require the algorithm to do more exploration. There is, then, greater possibility for deviation of the minimal path relative to the narrow vessels. For aorta tracking in 256x256x64 volumes, typical times for tracking and segmentation of the aorta from above the kidneys to the iliac branch takes 2-3 seconds. Re-centering the path with a distance transform takes an additional 2-3 seconds on a 73MHz Pentium III machine.

4.3 Segmentation Results

[Fig. 1a,b] Shows original slice data from a CE-MRA carotid/upper aorta study. This data was seeded with four points, two at the visible extrema of the subclavian arteries and two in vessel extrema near the top of the head, one on each side. The images depict the region explored by the phase 1 expansion in light gray with a white border.

5 Summary & Future Work

We have demonstrated a physics-based framework for the extraction of vessel paths from CE-MRA and CTA volumes. The two-phase computation is adaptive and rapid. It is not dependent on arbitrary thresholds or similar parameters. With minimal optimization, we feel that this algorithm is capable of real-time operation. In current tests the algorithm designed for vessels is being successfully used for segmenting and tracking the large intestine CT colonography studies and for delineating bone in peripheral CT studies.

References

- [1] Andrews, Frank C., *Equilibrium Statistical Mechanics* (2nd edition), John Wiley & Sons, Inc. New York, NY, 1975.
- [2] Malladi R., Sethian J.A., "A Real-Time Algorithm for Medical Shape Recovery," *Proceedings of International Conference on Computer Vision*, pp. 304-310, Mumbai, India, January 1998.
- [3] Masutani Y., Schiemann Th., Hohne K. "Vascular Shape Segmentation and Structure Extraction Using a Shape-Based Region-Growing Model", *MICCAI '98*, pp. 1242-1249, Cambridge, MA, USA, October 1998.
- [4] McInerney T. et al. "Medical Image Segmentation Using Topologically Adaptable Surfaces," LNCS Vol. 1205, *Proceedings of Computer Vision Virtual Reality and Robotics in Medicine '97*, pp 23-32, 1997.
- [5] Sethian J. "A Fast Marching Level Set Method for Monotonically Advancing Fronts," *Proceedings of the National Academy of Sciences*, Vol. 93, pp. 1591-1595, 1996.
- [6] Sethian J. *Level Set Methods and Fast Marching Methods: Evolving Interfaces in Computational Geometry, Fluid Mechanics, Computer Vision, and Materials Science*, Cambridge University Press, 2nd Edition, 1999.
- [7] Tek H., Kimia B. "Volumetric Segmentation of Medical Images by Three-Dimensional Bubbles," *CVIU*, 64(2):246-258, 1997.
- [8] Zhu S., Yuille A., "Region Competition: Unifying Snakes, Region Growing, and Bayes/MDL for Multiband Image Segmentation" *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 18, No. 9, September 1996.