

Automated measurement of retinal vascular tortuosity*

William E. Hart
Applied & Numerical Math Dept
Sandia National Laboratories

Michael Goldbaum
Ophthalmology Department
University of California, San Diego

Brad Côté
bcote@eagle.ais.net

Paul Kube
Dept Elec & Comp Engineering
University of California, San Diego

Mark R. Nelson
Data Vector

Abstract

Automatic measurement of blood vessel tortuosity is a useful capability for automatic ophthalmological diagnostic tools. We describe a suite of automated tortuosity measures for blood vessel segments extracted from RGB retinal images. The tortuosity measures were evaluated in two classification tasks: (1) classifying the tortuosity of blood vessel segments and (2) classifying the tortuosity of blood vessel networks. These tortuosity measures were able to achieve a classification rate of 91% for the first problem and 95% on the second problem, which confirms that they capture much of the ophthalmologists' notion of tortuosity.

1 Introduction


Normal retinal blood vessels are straight or gently curved. In some diseases, the blood vessels become tortuous, i.e., they become dilated and take on a serpentine path. The dilation is caused by radial stretching of the blood vessel, and the serpentine path occurs because of longitudinal stretching. The tortuosity may be focal, occurring only in a small region of retinal blood vessels, or it may involve the entire retinal vascular tree. Figure 1 shows images with tortuous and non-tortuous blood vessels.

Many disease classes produce tortuosity, including high blood flow, angiogenesis, and blood vessel congestion. Information about disease severity or change of disease with time may be inferred by measuring the tortuosity of the blood vessel network. Consequently, there is a benefit in measuring tortuosity in a consistent, repeatable fashion.

Given the increased availability of digitized fun-



(a)



(b)

Figure 1: Images with (a) tortuous and (b) non-tortuous blood vessel segments.

us photographs, automated tortuosity measurements are now feasible. Kaupp et al. [8] have reported unpublished results of an automated tortuosity measurement that uses a Fourier analysis of the perpendicular along the blood vessel. Smedby et al. [9] describe five tortuosity measures used to measure tortuosity in femoral arteries. Included are several measures of the integral curvature along the blood vessel, the number of inflection points of the vessel and the fraction of the vessel that has high curvature. Their experiments

*Supported by NIH Grant #RO1LM05759

examine properties of these measures like reproducibility and scalability. Capowski, Kylstra and Freedmen [1] describe a measure of blood vessel tortuosity based on spatial frequencies. Zhou et al. [11] have also described a method for distinguishing tortuous and nontortuous blood vessels in angiograms. In a preliminary abstract, we have proposed a tortuosity measure based on the integral curvature along a blood vessel [7].

In this paper we describe tortuosity measures that are used to measure the tortuosity of retinal blood vessels as well as the retinal blood vessel network. To assess the relative utility of these measures, they were used to classify blood vessel segments and blood vessel networks. The segments used in these classification experiments were extracted manually and automatically, and we discuss how the tortuosity measures can be influenced by properties of the method used to extract the vessel segments. The classification rate was as high as 91.5% for blood vessel segments and 95% for blood vessel networks. While no single tortuosity measure was clearly superior, the experiments recommend a total squared curvature measure.

2 Tortuosity Measures

2.1 Abstract Properties

This section discusses properties of tortuosity measures that are motivated by the ophthalmologist’s notion of tortuosity. These properties are defined for parametrized differentiable curves $C = (x(t), y(t))$, with t in an interval $[t_0, t_1]$. We model retinal blood vessels with such curves. A tortuosity measure τ takes a curve C as its argument and returns a real number.

Invariance to translation and rotation. While ophthalmologists’ judgements of vessel tortuosity do seem to incorporate the relative curvature of other vessels on the fundus, we believe their judgement is largely independent of the location or orientation of the vessels. Our measures assume that vessel tortuosity is not dependent on the location or orientation of the vessel.

Response to scaling. Ophthalmologists do not seem to have unequivocal intuitions about whether the tortuosity of a vessel depends on its scale (i.e. its absolute size). However, it is clear that if scale does affect the tortuosity then it does so in a multiplicative manner. That is, if curve $C = (x(t), y(t))$ has tortuosity $\tau(C)$, curve $C' = (\gamma x(t), \gamma y(t))$ has tortuosity $\beta(\gamma) \cdot \tau(C)$ for some function β . If tortuosity is invariant to scal-

ing, $\beta(\gamma) \equiv 1$; and if tortuosity is inversely related to scale, $\beta(\gamma) < 1$ when $\gamma > 1$.

Vessel Compositionality. It appears to be the intuition of ophthalmologists that a vessel that is composed of two segments of different tortuosities would have a degree of tortuosity between the tortuosities of its constituent segments. This assumes that the two segments are smoothly connected, but the “order” in which segments occur in the vessel is not important. A method of computing the tortuosity of a whole vessel that is consistent with this intuition is to weight the tortuosity of each constituent segment by the fraction of arc length which that segment contributes to the vessel. We call this method *weighted additivity*.

It also appears that if a vessel segment with a given tortuosity were extended with a segment of the same tortuosity, the resulting vessel would have the same tortuosity as either of its segments. That is, a vessel doesn’t increase (or decrease) its tortuosity by having the same serpentine pattern extended at greater length. We say that a measure has the *property of chord-colinear compositionality* if a vessel C is segmented such that each segment has the same tortuosity and the chords of the segments are colinear, then the tortuosity of the vessel is the same as the constituent segments.

Network Compositionality Another type of compositionality concerns the way in which the tortuosity measures for vessel segments are combined to determine a tortuosity measure for an entire vessel network. We have yet to develop a method for extracting a complete vessel network, so we calculate the tortuosity of a blood vessel network using the weighted additivity of all of the blood vessel segments in the image.

2.2 Definitions

Using our previous definition of a curve segment C , we define a suite of tortuosity measures. We begin by defining the components used to construct these measures. The *arc length* of C is $s(C) = \int_{t_0}^{t_1} \sqrt{x'(t)^2 + y'(t)^2} dt$. The *chord length* of C is $\text{chord}(C) = \sqrt{(x(t_0) - x(t_1))^2 + (y(t_0) - y(t_1))^2}$. The *curvature* of C at t is

$$\kappa(t) = \frac{x'(t)y''(t) - x''(t)y'(t)}{[y'(t)^2 + x'(t)^2]^{3/2}}.$$

The *total curvature* of a curve segment is $\text{tc}(C) = \int_{t_0}^{t_1} |\kappa(t)| dt$, and the *total squared curvature* of C is $\text{tsc}(C) = \int_{t_0}^{t_1} \kappa(t)^2 dt$.

Table 1 defines the measures that we examine in our experiments and describes their response

to scale. These measures have zero measure for straight vessel segments and increasing positive measure for segments as they become tortuous. Except for τ_2 and τ_3 , all of these measures have the property of chord-colinear compositionality. Only τ_4 and τ_5 have the property of compositionality.

Tortuosity Measure		Response to Scale
τ_1	$s(C)/\text{chord}(C) - 1$	1
τ_2	$tc(C)$	$1/\gamma$
τ_3	$tsc(C)$	$1/\gamma^2$
τ_4	$tc(C)/s(C)$	$1/\gamma^2$
τ_5	$tsc(C)/s(C)$	$1/\gamma^3$
τ_6	$tc(C)/\text{chord}(C)$	$1/\gamma^2$
τ_7	$tsc(C)/\text{chord}(C)$	$1/\gamma^3$

Table 1: Summary of the tortuosity measures and their response to scale.

The measure τ_1 simply measures the tortuosity of the segment by examining how long the curve is relative to its chord length. This measure is the same as the *distance factor* tortuosity measure described by Smedby et al. [9]. Measures τ_2 and τ_3 directly calculate the curvature of the curve. The τ_3 measure differs from τ_2 in that it places a greater emphasis on the parts of the curve that have high curvature, and deemphasizes the parts of the curve that have low curvature. Since the curvature is greater for small vessels, τ_3 will emphasize the tortuosity of smaller vessels more than τ_2 . The τ_2 measure is the same as the total curvature measure described by Smedby et al. [9]. The remaining measures are “length-normalized.” The measures τ_4 and τ_5 average the total curvature measures by the arclength, while τ_6 and τ_7 average by the chord length.

2.3 Tortuosity Calculation

The abstract definitions for the tortuosity measures were modified to work with skeletonized blood vessels that represent the center line of blood vessels with sequences of pixel locations. Following a suggestion of Flynn and Jain [6], we smoothed the pixel representation of a blood vessel segment, using a low-pass filter, before making our tortuosity calculations. This eliminated undesirable noise that is due to the discrete nature of the pixel representation. For example, a blood vessel at a 45 degree angle in a image can have a pixel representation that is a zigzag line of pixels along the blood vessel. The smoothing operation

also eliminated noise that was introduced when independent segments were linked together to form a longer segment. Preliminary experiments indicated that two applications of this smoothing method were sufficient to reduce this noise.

3 Methods

We examined the utility of our tortuosity measures by using them as features in two different classification problems. In the first problem, we classified blood vessel segments as tortuous or non-tortuous. In the second problem, we classified the tortuosity of the blood vessel network.

3.1 Data

Blood vessel segments were extracted from a set of 20 retinal images in two different ways: automatically and manually. To extract blood vessel segments, we applied the blood vessel filter described in Chaudhuri et al. [3] to the green plane of an RGB image. The green plane was selected because it typically exhibits the greatest contrast. The filter was applied at 12 orientations over 180° . The final response map was computed by taking the maximum response of the 12 filters at each location. We thresholded and thinned the response map of the blood vessel filter to produce an image containing binary edge segments. The edge segments were primarily blood vessels, but also included some edges of large objects like the optic nerve.

To create the set of *automatically* extracted blood vessel segments, the edge segments were classified as blood vessels or non-blood vessels using the linear classifier described in Côté et al. [2]. There were 981 automatically extracted blood vessel segments, of which 252 were tortuous and 729 were non-tortuous.

The unclassified edge segments were also used to extract blood vessel segments *manually*. The edge segments were manually identified as blood vessels and linked together to form the final blood vessel segments. Care was taken to link vessel segments only if the smoothness of the link reflected the curvature of the underlying blood vessel. There were 284 blood vessels extracted manually, of which 133 were tortuous and 151 were non-tortuous.

After extraction, the tortuosity measures were calculated for each segment, and the segments were labeled by one of the authors (MG), a retinal specialist who has experience with diseases causing tortuosity in retinal blood vessels. These two

sets of extracted blood vessels were used for the first classification problem.

For the second classification problem, the extracted blood vessels were used to calculate the network tortuosity measures for each image. The vessel network tortuosity of the 20 retinal images were labeled; there were 10 tortuous and 10 non-tortuous images.

3.2 Classification

We used a logit model [4] to classify the data for both classification problems. For problems with two classes, a logit model computes a weighted sum of the input features passed into a logistic function. The output of the logit model is between zero or one. To perform classification, the output is thresholded to zero or one, depending on whether the output is greater or less than 0.5.

Two measures were used to compare the performance of classifiers using the different tortuosity measures: the classification rate and the integrated relative operating characteristic (ROC). The classification rate is simply the proportion of test samples that are correctly classified. The ROC measures the proportion of test samples that are correctly classified as positive instances (true-positives) as a function of the proportion of negative test samples that are classified as positive instances (false-positives) [10]. The integrated ROC measures how sensitive a classifier is to the choice of the classification threshold; values closer to 1.0 are better.

The performance on the samples used to train a classifier is typically an optimistic estimate of the classifier’s performance on a new set of data [5]. To estimate the true error rate, we used cross-validation to partition the data into two subsets that are used to train and test the data. For both classification problems a variety of partitions of the data were selected. A different classifier was trained and tested on each partition, and the mean classification rate and integrated ROC on the testing subsets were used to evaluate the expected performance of a classifier.

4 Results

In both classification problems, we performed experiments that examine the classification rate using each tortuosity measure by itself. We also performed classification using all of the tortuosity measures together. Table 2 summarizes the experimental results for the two classification prob-

lems. The experiment name indicates whether the data were manually extracted or automatically extracted, and indicates which tortuosity measures were used in the experiment. For most of the classifiers, the relative integrated ROC values mirrored the relative classification rates. An interesting exception is τ_1 , which had a relatively high classification rate for automatically extracted segments but a relatively low integrated ROC value.

5 Discussion

The experimental results demonstrate that the proposed tortuosity measures can be used to classify the tortuosity of blood vessel segments and blood vessel networks. In particular, these results show that the tortuosity measures can be effectively used with automatically extracted blood vessel segments, which is important for the development of retinal image analysis tools.

There are several trends in our results. First, the classifiers tend to have better performance on the automatically extracted data set. This can be attributed to the short, simple blood vessel segments in this data set. The shorter lengths of these segments avoid difficulties when an entire vessel segment is labeled tortuous, but subsegments of the vessel are not tortuous, which makes it difficult to classify.

It is surprising that the classification performance increased for τ_4 and τ_5 when classifying the blood vessel network. These measures have the compositionality property, so the blood vessel networks should have roughly the same tortuosity measure for both the manually and automatically extracted data sets. An analysis of the data indicated that the performance difference was due to a difference in the extraction process on one of the images. The automatically extracted segments included some highly curved, erroneous blood vessel edges that were not included in the edges for the manually extracted segments.

We have observed several different ways in which the extraction process influenced the performance in the classification tasks. First, the segmentation method can affect the tortuosity of the extracted vessels. For example, our segmentation algorithm breaks up vessels at points of very high curvature and tends to smooth out very thin blood vessels. The length of the extracted segments can also affect the utility of the tortuosity measure for classification. The shorter segments in the automatically extracted data sets is the prin-

Measure	Segments				Networks			
	Classification Rate		Integrated ROC		Classification Rate		Integrated ROC	
	Manual	Auto	Manual	Auto	Manual	Auto	Manual	Auto
τ_1	79.5	91.3	0.897	0.955	65	90	0.61	0.91
τ_2	79.3	82.7	0.875	0.915	70	85	0.79	0.85
τ_3	82.9	89.5	0.927	0.960	90	90	0.86	0.86
τ_4	80.4	88.5	0.905	0.951	80	95	0.83	0.99
τ_5	81.2	87.7	0.906	0.941	85	90	0.84	0.83
τ_6	82.1	89.1	0.921	0.956	80	90	0.81	0.91
τ_7	82.0	88.1	0.914	0.944	85	90	0.88	0.91
ALL	85.6	91.5	0.935	0.970	95	90	0.95	0.86

Table 2: Results for the problem of classifying blood vessel segments and blood vessel segments. The mean of the cross-validation classification rate and integrated ROC.

principle reason for the better classification results for these data sets. Finally, extracting blood vessel segments automatically can effect the tortuosity measures of blood vessel networks since they can include misclassified segments. This problem did not seriously impact our results because the blood vessel classifier used to automatically extract the blood vessel segments is quite accurate [2].

Finally, we note that although our classification results do not strongly support the use of one measure over the others, the τ_3 measure can be recommended for a tortuosity measure. Both the τ_3 and τ_4 measures seem closest to the ophthalmologist’s notion of tortuosity, though the τ_3 measure performed slightly better on the classification tasks. Classification with all of the tortuosity measures was better than τ_3 , but not in a statistically significant manner.

References

- [1] J. J. Capowski, J. A. Kylstra, and S. F. Freedman. A numeric index based on spatial-frequency for the tortuosity of retinal-vessels and its application to plus disease in retinopathy of prematurity. *Retina*, 15(6):490–500, 1995.
- [2] B. Côté, W. E. Hart, M. Goldbaum, P. Kube, and M. R. Nelson. Classification of blood vessels in images of the ocular fundus. Technical Report CS94-350, Univ Calif, San Diego, 1994.
- [3] S. Chaudhuri, S. Chatterjee, N. Katz, M. Nelson, and M. Goldbaum. Detection of blood vessels in retinal images using two dimensional blood vessel filters. *IEEE Trans Medical Imaging*, 8(3), Sept. 1989.
- [4] J. Cramer. *The Logit Model: An Introduction for Economists*. Edward Arnold, 1991.
- [5] R. O. Duda and P. E. Hart. *Pattern Classification and Scene Analysis*. John Wiley and Sons, 1973.
- [6] P. J. Flynn and A. K. Jain. On reliable curvature estimation. In *IEEE Proceedings on Computer Vision and Pattern Recognition*, pages 110–116, 1989.
- [7] N. P. Katz, M. H. Goldbaum, S. Chaudhuri, and M. R. Nelson. Automated measurements of blood vessels in digitized images of the ocular fundus. In *Invest Ophthalmol Vis Sci*, volume 31, 1990. #1185.
- [8] A. Kaupp, H. Toonen, S. Wolf, K. Schulte, R. Effert, D. Meyer-Ebrecht, and M. Reim. Automatic evaluation of retinal vessel width and tortuosity in digital fluorescein angiograms. In *Invest Ophthalmol Vis Sci*, volume 32, 1991. #952.
- [9] Ö. Smedby, N. Högman, S. Nilsson, U. Eriksson, A. G. Olsson, and G. Walldius. Two-dimensional tortuosity of the superficial femoral artery in early atherosclerosis. *J Vascular Research*, 30:181–191, 1993.
- [10] J. A. Swets. Measuring the accuracy of diagnostic systems. *Science*, 240:1285–1293, 1988.
- [11] L. A. Zhou, M. S. Rzeszotarski, S. L. J., and J. M. Chokreff. The detection and quantification of retinopathy using digital angiograms. *IEEE Trans Med Imaging*, 13(4):619–626, 1994.